

Software Model Checking

RANJIT JHALA

University of California, San Diego

and

RUPAK MAJUMDAR

University of California, Los Angeles

We survey recent progress in software model checking.

Categories and Subject Descriptors: D.2.4 [**Software Engineering**]: Software/Program Verification; F.3.1 [**Logics and Meanings of Programs**]: Specifying and Verifying and Reasoning about Programs

General Terms: Languages; Verification; Reliability

Additional Key Words and Phrases: Software model checking, enumerative and symbolic model checking, abstraction, counterexample-guided refinement, safety, liveness

ACM Reference Format:

Jhala, R. and Majumdar, R. 2009. Software model checking. *ACM Comput. Surv.* 41, 4, Article 21 (October 2009), 54 pages. DOI = 10.1145/1592434.1592438 <http://doi.acm.org/10.1145/1592434.1592438>

Software model checking is the algorithmic analysis of programs to prove properties of their executions. It traces its roots to logic and theorem proving, both to provide the conceptual framework in which to formalize the fundamental questions and to provide algorithmic procedures for the analysis of logical questions. The undecidability theorem [Turing 1936] ruled out the possibility of a sound and complete algorithmic solution for any sufficiently powerful programming model, and even under restrictions (such as finite state spaces), the correctness problem remained computationally intractable. However, just because a problem is hard does not mean it never appears in practice. Also, just because the *general* problem is undecidable does not imply that specific *instances* of the problem will also be hard.

As the complexity of software systems grew, so did the need for some reasoning mechanism about correct behavior. (While we focus here on *analyzing* the behavior of

This work was sponsored in part by the National Science Foundation grants CCF-0546170, CCF-0702743, and CNS-0720881.

Authors' addresses: R. Jhala, University of California, San Diego, Department of Computer Science and Engineering, 9500 Gilman Drive #0404, La Jolla, CA 92093-0404; email: rjhala@ucsd.edu; R. Majumdar, UCLA Comp Sci, Box 951596, 4532H Boelter, Los Angeles, CA 90095-1596; email: rupak@cs.ucla.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

©2009 ACM 0360-0300/2009/10-ART21 \$10.00

DOI 10.1145/1592434.1592438 <http://doi.acm.org/10.1145/1592434.1592438>

a program relative to given correctness specifications, the development of *specification mechanisms* happened in parallel, and merits a different survey.)

Initially, the focus of program verification research was on manual reasoning, and the development of axiomatic semantics and logics for reasoning about programs provided a means to treat programs as logical objects [Floyd 1967; Hoare 1969; Dijkstra 1976; Apt and Olderog 1991]. As the size of software systems grew, the burden of providing entire manual proofs became too cumbersome, and brought into questions whether long and laborious proofs of correctness could themselves be trusted [de Millo et al. 1979]. This marked a trend toward automating the more mundane parts, leaving the human to provide guidance to an automatic tool (for example, through loop invariants and function pre- and post-conditions [Dijkstra 1976]). This trend has continued since: the goal of software model checking research is to expand the scope of automated techniques for program reasoning, both in the scale of programs handled and in the richness of properties that can be checked, reducing the burden on the expert human programmer.

More recently, software model checking has been influenced by three parallel but somewhat distinct developments. First, development of program logics and associated decision procedures [Nelson 1981; Nelson and Oppen 1980; Shostak 1984] provided a framework and basic algorithmic tools to reason about infinite state spaces. Second, automatic *model checking* techniques [Clarke and Emerson 1981; Queille and Sifakis 1981; Vardi and Wolper 1994] for temporal logics [Pnueli 1977; Emerson 1990] provided basic algorithmic tools for state-space exploration. Third, compiler analysis, formalized by *abstract interpretation*, provided connections between the logical world of infinite state spaces and the algorithmic world of finite representations. Throughout the 1980s and 1990s, the three communities developed with only occasional interactions. However, in the last decade, there has been a convergence in the research directions and modern software model checkers are a culmination of ideas that combine and perhaps supersede each area alone. In particular, the term “software model checker” is probably a misnomer, since modern tools simultaneously perform analyses traditionally classified as theorem proving, or model checking, or dataflow analysis. We retain the term solely to reflect historical development.

In this survey, we trace some of these ideas that have combined to produce tools with more and more automation and precision for the analysis of software systems.

1. PRELIMINARY DEFINITIONS

1.1. Simple Programs

We assume a transition-relation representation of programs, following the style of Lamport [1983] and Manna and Pnueli [1992]. Over the course of this chapter, we define several classes of programs, starting with a simple model, and adding more features as we go along. To begin with, we consider *simple* programs which are defined over a set of integer variables. In the following sections, we augment this simple model with pointers and procedure calls.

A *simple imperative program* $P = (X, L, \ell_0, T)$ consists of a set X of typed variables, a set L of control locations, an initial location $\ell_0 \in L$, and a set T of transitions. Each transition $\tau \in T$ is a tuple (ℓ, ρ, ℓ') , where $\ell, \ell' \in L$ are control locations, and ρ is a constraint over free variables from $X \cup X'$. The variables from X denote values at control location ℓ , and the variables from X' denote the values of the variables from X at control location ℓ' . The sets of locations and transitions naturally define a directed labeled graph, called the *control-flow graph* (CFG) of the program. We denote by *imp* the class of all simple programs.

An imperative program with assignments and conditional statements can be translated into a simple program as follows. The control flow of the program is captured by the graph structure. An assignment $x := e$ (where e is an expression over X) is translated to the relation:

$$x' = e \wedge \bigwedge_{y \in X \setminus \{x\}} y' = y$$

and a conditional p (a Boolean expression over X) is translated to

$$p \wedge \bigwedge_{x \in X} x' = x$$

We shall exploit this translation to provide our examples in a C-like syntax for better readability.

Similarly, there is a simple encoding of simple programs executing concurrently with an interleaved semantics into a single simple program, for example, by taking the union of the variables, the cross product of the locations of the threads, and transitions $(\langle l_1, l'_1 \rangle, \rho, \langle l_2, l'_2 \rangle)$ where either (l_1, ρ, l_2) is a transition of the first thread and $l'_2 = l'_1$, or (l'_1, ρ, l'_2) is a transition of the second thread and $l_2 = l_1$.

Thus, we shall use simple programs in our exposition of model checking algorithms in the following sections. Of course, particular model checkers may have more structured representations of programs that can be exploited by the model checker. For example, input formats can include primitive synchronization operations (e.g., locks, semaphores, or atomic sections), language-level support for channels that are used for message passing, etc. In each case, such features can be compiled down to the “simple” model.

A *state* of the program P is a valuation of the variables from X . The set of all states is denoted $v.X$. We shall represent sets of states using constraints. For a constraint ρ over $X \cup X'$ and a valuation $(s, s') \in v.X \times v.X'$, we write $(s, s') \models \rho$ if the valuation satisfies the constraint ρ .

Sometimes we shall consider simple programs with an explicitly provided *initial state* that sets all variables in X to specific values in their domains. If an initial state is not given explicitly, we assume that the program can start executing from an arbitrary state. In this latter case, any state is initial.

A *finite computation* of the program P is a finite sequence $\langle \ell_0, s_0 \rangle, \langle \ell_1, s_1 \rangle, \dots, \langle \ell_k, s_k \rangle \in (L \times v.X)^*$, where ℓ_0 is the initial location, s_0 is an initial state, and for each $i \in \{0, \dots, k-1\}$, there is a transition $(\ell_i, \rho, \ell_{i+1}) \in T$ such that $(s_i, s_{i+1}) \models \rho$. Likewise, an *infinite computation* of the program P is an infinite sequence $\langle \ell_0, s_0 \rangle, \langle \ell_1, s_1 \rangle, \dots, \langle \ell_k, s_k \rangle \dots \in (L \times v.X)^\omega$, where ℓ_0 is the initial location, s_0 is an initial state, and for each $i \geq 0$, there is a transition $(\ell_i, \rho, \ell_{i+1}) \in T$ such that $(s_i, s_{i+1}) \models \rho$. A computation is either a finite computation or an infinite computation. A state s is *reachable* at location ℓ if $\langle \ell, s \rangle$ appears in some computation. A location ℓ is reachable if there exists some state s such that $\langle \ell, s \rangle$ appears in some computation. A *path* of the program P is a sequence $\pi = (\ell_0, \rho_0, \ell_1), (\ell_1, \rho_1, \ell_2), \dots, (\ell_{k-1}, \rho_{k-1}, \ell_k)$ of transitions, where ℓ_0 is the initial location.

We define two useful operations on states. For a state s and a constraint ρ over $X \cup X'$, we define the set of *successor states* $Post(s, \rho) = \{s' \mid (s, s') \models \rho\}$. Similarly for a state s' and constraint ρ , we define the set of *predecessor states* $Pre(s', \rho) = \{s \mid (s, s') \models \rho\}$. The operations $Post$ and Pre are extended to sets of states in the obvious way: $Post(S, \rho) = \bigcup_{s \in S} Post(s, \rho)$ and $Pre(S, \rho) = \bigcup_{s \in S} Pre(s, \rho)$. The $Post$ and Pre

Thread 1	Thread 2
$L0 : \text{lock} = 0 \wedge \text{lock}' = 1 \wedge x' = x$	$L0' : \text{lock} = 0 \wedge \text{lock}' = 1 \wedge x' = x$
$L1 : x' = 1 \wedge \text{lock}' = \text{lock}$	$L1' : x' = 2 \wedge \text{lock}' = \text{lock}$
$L2 : \text{lock} = 1 \wedge \text{lock}' = 0 \wedge x' = x$	$L2' : \text{lock} = 1 \wedge \text{lock}' = 0 \wedge x' = x$
$L3 :$	$L3' :$

Fig. 1. Simple program.

operations are also called the *strongest postcondition* and *pre-image* operations respectively [Dijkstra 1976]. An operator related to the pre-image is the *weakest liberal precondition* operation WP [Dijkstra 1976] defined as $\text{WP}(s', \rho) = \{s \mid \forall t. (s, t) \models \rho \Rightarrow t = s'\}$. The WP and Pre operators coincide for deterministic systems.

Example. Figure 1 shows a simple program consisting of two threads. Each thread has four locations $\{L0, L1, L2, L3\}$ and $\{L0', L1', L2', L3'\}$, respectively. There are two global variables x and lock , let us assume x can take values in the set $\{0, 1, 2\}$ and lock is Boolean. The transitions for thread 1 are given by

$$\begin{aligned}
 & (L0, \text{lock} = 0 \wedge \text{lock}' = 1 \wedge x' = x, L1) \\
 & (L1, \text{lock}' = \text{lock} \wedge x' = 1, L2) \\
 & (L2, \text{lock} = 1 \wedge \text{lock}' = 0 \wedge x' = x, L3)
 \end{aligned}$$

and similarly for thread 2. For readability, we write programs in an imperative syntax as shown in Figure 1. The initial location of the program is $\langle L0, L0' \rangle$. Let us additionally assume that the initial state is $\text{lock} = 0$ and $x = 0$. The set of reachable states are given by:

$$\begin{aligned}
 & \langle L0, L0', \text{lock} = 0, x = 0 \rangle \\
 & \langle L1, L0', \text{lock} = 1, x = 0 \rangle \quad \langle L0, L1', \text{lock} = 1, x = 0 \rangle \\
 & \langle L2, L0', \text{lock} = 1, x = 1 \rangle \quad \langle L0, L2', \text{lock} = 1, x = 2 \rangle \\
 & \langle L3, L0', \text{lock} = 0, x = 1 \rangle \quad \langle L0, L3', \text{lock} = 0, x = 2 \rangle \\
 & \langle L3, L1', \text{lock} = 1, x = 1 \rangle \quad \langle L1, L3', \text{lock} = 1, x = 2 \rangle \\
 & \langle L3, L2', \text{lock} = 1, x = 2 \rangle \quad \langle L2, L3', \text{lock} = 1, x = 1 \rangle \\
 & \langle L3, L3', \text{lock} = 0, x = 2 \rangle \quad \langle L3, L3', \text{lock} = 0, x = 1 \rangle
 \end{aligned}$$

Notice that the location $\langle L1, L1' \rangle$ is not reachable.

1.2. Properties

The main goal of software model checking is to prove *properties* of program computations. Examples of *properties* are simple assertions, that state that a predicate on program variables holds whenever the computation reaches a particular control location (e.g., “the variable x is positive whenever control reaches ℓ ”), or global invariants, that state that certain predicates hold on every reachable state (e.g., “each array access is within bounds”), or termination properties (e.g., “the program terminates on all inputs”). Broadly, properties are classified as *safety* and *liveness*. Informally, safety properties stipulate that “bad things” do not happen during program computations, and liveness properties stipulate that “good things” do eventually happen. This intuition was formalized by [Alpern and Schneider 1987] as follows.

Mathematically, a property is a set of infinite sequence from $(L \times v.X)^\omega$. We say an infinite sequence σ *satisfies* a property Π if $\sigma \in \Pi$. A *safety property* $\Pi \subseteq (L \times v.X)^\omega$ is a set of infinite computations such that for every infinite computation $\sigma \in$

Π , for every finite prefix σ' of σ , there exists $\beta \in (L \times v.X)^\omega$ such that $\sigma' \cdot \beta \in \Pi$, where \cdot is string concatenation. Taking the contrapositive, a computation σ does not satisfy a safety property Π if there is a finite prefix σ' of σ such that no extension of σ' satisfies Π . That is, every violation of a safety property has a finite witness. A *liveness property* Π , on the other hand, is such that every partial computation can be extended to a computation satisfying Π . That is, Π is a liveness property for every finite computation $\alpha \in (L \times v.X)^*$, there is an infinite computation $\beta \in (L \times v.X)^\omega$ such that $\alpha \cdot \beta$ satisfies Π .

The *verification problem* takes as input a program P and a property Π , and returns “safe” if every computation of P is in Π , and returns “unsafe” otherwise. In the former case, we say P *satisfies* Π . If Π is a safety property (respectively, a liveness property), we refer to the *safety* (respectively, *liveness*) verification problem.

For most of the survey, we focus on the problem of checking if a program P satisfies a safety property Π . We consider the verification of liveness properties in Section 8.

In the following, we formulate the safety verification problem as a check for reachability of a particular location. Let P be a simple program, and let $\mathcal{E} \in L$ be a special *error location*. We say the program is *safe* with respect to the error location \mathcal{E} if the location \mathcal{E} is not reachable. An *error trace* is a computation ending in the location \mathcal{E} . Clearly, reachability of location \mathcal{E} is a safety property, and it is known that checking any safety property (expressed, e.g., in a temporal logic) can be reduced to the above reachability question. Thus, we consider safety verification problems specified in the following form: the input is a program P and an error location $\mathcal{E} \in L$, and the output is “safe” if P is safe with respect to \mathcal{E} , and “unsafe” if \mathcal{E} is reachable.

An alternate and common way to specify safety properties for programs is through *assertions*. The programmer explicitly puts in a predicate p over program variables (called the assertion) at a program location, with the intent that for every execution of the program that reaches the location, the program state satisfies the assertion. Our simple formulation subsumes *assertion checking* in a simple way: each assertion is replaced by a conditional on the predicate p , the program continues execution if the predicate p holds, and enters an error location otherwise. Conversely, reachability of an error location can be encoded as an assertion violation, by putting an assertion *false* at the desired location. Thus, reachability of error location and assertions are equivalent ways to specify safety properties.

In the following, we use the following terminology. An algorithm for the safety verification problem is *sound* if for every program P and error location \mathcal{E} of P , if the algorithm returns “safe” then P is safe with respect to \mathcal{E} . It is *complete* if for every program P and error location \mathcal{E} of P , if P is safe with respect to \mathcal{E} , then the algorithm returns “safe”.

The undecidability of the halting problem implies that there is no general sound *and* complete algorithm for the verification problem. In practice, algorithmic tools maintain soundness, but compromise on completeness. Interestingly, there are two ways this can be done. One way is to explore a part of the reachable state space of the program, hoping to find a computation that violates the property. In this case, the model checker is geared towards *falsification*: if it finds a violation, then the program does not satisfy the property, but if it does not, no conclusion about correctness can be drawn (either the program satisfies the property, or the unexplored part of the state space has a computation that is not in the property). Another way is to explore a *superset* of program computations. In this case, the model checker is geared towards *verification*: if it finds the property is satisfied, then the program does satisfy the property, however, if it finds a violation, no conclusion can be drawn (either the original program has a computation not in the property, or the violation is due to adding extraneous computations in the analysis).

1.3. Organization

The rest of the survey is organized as follows. We first describe two main ways of representing state: enumerative (in which individual states are represented, Sections 2) and symbolic (in which sets of states are represented using constraints, Section 3). We then describe abstraction techniques (Section 4), which reduce the state space at the expense of precision, and automatic techniques to make abstract analyses more precise (Section 5). While we describe each facet in isolation, in practice, several notions can be combined within the same tool: the program state can be represented partly in enumerated form, and partly symbolically, and combined with (several different) abstractions.

The next few sections describe extensions to the basic approach: dealing with (potentially recursive) functions (Section 6), dealing with program memories with dynamic allocation and heap abstractions (Section 7), and extending the techniques to reason about liveness properties (Section 8). We then make the connection between model checking techniques and related dynamic (testing) and static (type systems) techniques in software quality (Section 9). We conclude with the limitations of current tools and some future challenges.

2. CONCRETE ENUMERATIVE MODEL CHECKING

Algorithms for concrete enumerative model checking essentially traverse the graph of program states and transitions using various graph search techniques. The term *concrete* indicates that the techniques represent program states exactly. The term *enumerative* indicates that these methods manipulate individual states of the program, as opposed to *symbolic* techniques (see Section 3) which manipulate sets of states. Concrete enumerative model checking grew out of testing and simulation techniques in the late 1970's, most notably from techniques for testing network protocols [Sunshine 1978; Sunshine et al. 1982], as a comprehensive methodology to ensure correctness. In independent lines of work, Clarke and Emerson [1981] and Queille and Sifakis [1981] generalized the techniques to temporal logic specifications. Since then, the method has been applied successfully in analyzing many software domains, most notably asynchronous, message-passing protocols [Holzmann 1997] and cache coherence protocols [Dill 1996].

2.1. Stateful Search

Let imp_{fin} be the class of simple programs in which each variable ranges over a finite domain. In this case, the safety verification problem can be solved by explicitly constructing the (finite) set of reachable states and checking that \mathcal{E} is not reachable.

Figure 2 shows Algorithm EnumerativeReachability, a procedure that computes the set of reachable states of a program in imp_{fin} by performing graph search. The algorithm maintains a set *reach* of *reachable states* and a set *worklist* of *frontier states* that are found to be reachable but whose successors may not have been explored. Initially, the set *reach* is empty, and the frontier *worklist* contains all the *initial states*. The main *reachability loop* of the algorithm explores the states of the frontier one at a time. If the state has not been visited before, the successors of the state are added to the frontier, otherwise, the successors are not added. The process is repeated until all reachable states have been explored, which happens when the frontier *worklist* becomes empty. At this point, *reach* contains exactly the set of reachable states. The loop terminates for all imp_{fin} , in fact, it terminates for all programs for which *reach* is finite.

While we have implemented the check for reachability of \mathcal{E} at the end of the reachability loop, it can be performed whenever a new state is picked from the frontier. Also,

Algorithm: Enumerative Reachability
Input simple program $P = (X, L, T, \ell_0)$, error location $\mathcal{E} \in L$
Output SAFE if P is safe w.r.t. \mathcal{E} , UNSAFE otherwise

```

def EnumerativeReachability( $P, \mathcal{E}$ ):
    reach =  $\emptyset$ 
    worklist =  $\{(\ell_0, s) \mid s \in v.X\}$ 
    while worklist  $\neq \emptyset$  do:
        choose( $\ell, s$ ) from worklist, worklist = worklist  $\setminus \{(\ell, s)\}$ 
        if  $(\ell, s) \notin reach$ :
            reach = reach  $\cup \{(\ell, s)\}$ 
            foreach  $(\ell, \rho, \ell')$  in  $T$  do:
                add  $\{(\ell', s') \mid s' \in Post(s, \rho)\}$  to worklist
        if exists  $(\mathcal{E}, s) \in reach$ :
            return UNSAFE
        else:
            return SAFE

```

Fig. 2. EnumerativeReachability.

Algorithm EnumerativeReachability can be easily modified to produce an error trace in case \mathcal{E} is reachable.

The generic schema of Figure 2 can be instantiated with different data structures for maintaining the set of reachable states and the set of frontier states, and with algorithms for implementing the order in which states are chosen from the frontier set.

For example, maintaining the frontier as a stack (and always choosing the next state by popping the stack) ensures depth-first traversal of the graph, while maintaining the frontier as a queue ensures breadth-first traversal. For efficiency of checking membership, the set *reach* is usually implemented as a hashtable [Holzmann 1997; Dill 1996]. In addition, instead of generating all states and transitions up front, reachability search algorithms usually construct the state space *on-the-fly*, based on currently reached states and the program. This exploits the observation that the reachable state space of the program can be much smaller than the state space.

While we focus on a *forward* algorithm, based on the *Post* operator, a dual *backward* algorithm based on the *Pre* operator is also possible. This algorithm starts at the location \mathcal{E} and searches backward over the set of states that can reach \mathcal{E} . If some initial state can reach \mathcal{E} , the program is unsafe.

Enumerative model checking of finite state concurrent programs has been implemented in several tools, most notably SPIN [Holzmann 1997] and MURPHI [Dill 1996]. Both tools have had significant impact, especially in the protocol verification domain.

The state space of a program can be exponentially larger than the description of the program. This problem, known as *state explosion*, is one of the biggest stumbling blocks to the practical application of model checking. Controlling state explosion has therefore been a major direction of research in software model checking. In the context of enumerative model checking, broadly, the research takes two directions.

First, *reduction-based* techniques compute equivalence relations on the program behaviors, and explore one candidate from each equivalence class. A meta-theorem asserts that the reduced exploration is complete, that is, for any bug in the original system, there is a bug in the reduced one. Primary reduction-based techniques consist of *partial-order reduction* [Valmari 1992; Katz and Peled 1992; Godefroid 1996], *symmetry reduction* [Clarke et al. 1993; Emerson and Sistla 1996; Ip and Dill 1996; Sistla et al. 2000] or minimization based on behavioral equivalences such as simulation or bisimulation [Bouajjani et al. 1990; Loiseaux et al. 1995; Bustan and Grumberg 2003]. Partial order reductions exploit the independence between parallel threads of execution

on unrelated parts of the state. That is, if two transitions τ_1 and τ_2 in parallel threads of execution access independent sets of variables, the final state reached after executing τ_1 and τ_2 in that order is the same as that reached after executing first τ_2 and then τ_1 . An algorithm based on partial order reduction chooses to explore one candidate interleaving among independent transitions rather than all of them. Symmetry reduction determines symmetries in the program, and explores one element from each symmetry class. In general, identifying symmetries in the state space may be difficult, and in practice, the syntax of the programming language is used to identify symmetries. In many examples, such as parameterized protocols, symmetry-based techniques can yield dramatic reductions in the state space [Ip and Dill 1996]. Behavioral equivalences such as similarity and bisimilarity construct a quotient graph that preserves reachability (i.e., there is a path from an initial state to \mathcal{E} in the original graph iff there is a path to \mathcal{E} in the quotient), and then performs reachability analysis on the quotient. This assumes that constructing the quotient is simpler, or more scalable, than computing reachability on the original graph. Some algorithms combine reachability and quotient construction [Lee and Yannakakis 1992].

Second, *compositional* techniques reduce the safety verification problem on the original program to proving properties on subprograms, such that the results of model checking the subprograms can be combined to deduce the safety of the original program. *Assume-guarantee reasoning* is a well-studied form of compositional reasoning [Misra and Chandy 1981; Jones 1983; Stark 1985; Abadi and Lamport 1993, 1995; Henzinger et al. 1998; Alur and Henzinger 1999; Alur et al. 1998]. In this approach, the behavior of a component is summarized as a pair (A, G) of two formulas: an *assumption* on the environment of the component (which restricts the possible inputs presented to the component), and a *guarantee* that the component will satisfy provided the inputs to the component from the environment satisfy the assumption. Let P_1 and P_2 be two components with assumption-guarantee pairs (A_1, G_1) and (A_2, G_2) respectively. To show that the composition of P_1 and P_2 has an assumption-guarantee pair (A, G) , one shows that the assumption A and the guarantee G_2 together imply that P_1 maintains the assumption A_2 , and similarly, the assumption A and the guarantee G_1 together imply that P_2 maintains the assumption A_1 , and finally that G follows from A , G_1 , and G_2 . The reasoning above is circular: assumptions on the behavior of P_1 are used to discharge assumptions on the behavior of P_2 and vice versa. In order to make the reasoning sound, the interpretation of (A, G) must be carefully defined to break this circularity. One common way is to break the circularity by induction on time.

The search heuristic also has a profound influence on the performance of Algorithm EnumerativeReachability on practical problems, where the objective is to find bugs quickly. One direction of research applies *guided* or *directed search* heuristics inspired by search heuristics from the artificial intelligence literature, such as iterative deepening [Korf 1985], best-first search [Russell and Norvig 2003], or A* search [Hart et al. 1968; Russell and Norvig 2003]. These techniques were imported into MURPHI [Yang and Dill 1998] and SPIN [Edelkamp et al. 2004; Lluch-Lafuente 2003], and there have been several extensions (and combinations with orthogonal techniques) since that time [Fraser et al. 2000].

In a different direction, using the observation that model checking is often more useful as a *falsification* or bug-finding aid, one gives up completeness of the search. This is often done by limiting resources available to the model checker (run for a specified amount of time or memory), or by bounding the set of behaviors of the program to be explored (e.g., by bounding the depth of the search, or the number of context-switches). *Bitstate hashing* is a popular technique in which the hash of each reachable state is stored, rather than the state itself. The choice of the range of the hash function is determined

by the available memory. Bitstate hashing is unsound, as two distinct reached states can hash to the same value (a hash *collision*). In order to obtain nicer guarantees on the probability of collision, each state is hashed using several (in practice, two or three) independent hash functions. When the search space is so big that even with bitstate hashing one can only explore a small portion of the state space, it is possible to store the states on disk rather than on main memory [Stern and Dill 1998]. This is a time-for-space tradeoff: while accessing states take much longer on disk, the disk allows storing a much larger state space, and the algorithms have to be designed carefully to ensure disk accesses are fast (e.g., by ensuring disk reads are sequential).

2.2. Systematic Execution Exploration

The execution-based model checking approach, pioneered by VERISOFT [Godefroid 1997], is a special case of enumerative verification. This approach uses the runtime system of a programming language implementation to implement enumerative state space exploration. In the most common approach, implemented in VERISOFT, JAVAPATHFINDER, and several other tools, all the the non-determinism in a concurrent program is factored into two sources: *inputs* from the environment, and *scheduling choices* made by the scheduler. That is, each sequence of user inputs and schedule choices uniquely determines the outcome of an execution, and the set of all behaviors can be explored by analyzing the behavior of the process under all possible inputs and schedules.

In contrast to techniques that exhaustively explore the state space using graph algorithms, systematic execution exploration proceeds by systematically iterating over the space of possible schedules and simply *executing* the process under each schedule. This most appealing benefit of this approach is that it sidesteps the need to be able to formally represent the semantics of the programming language and machine instructions as a transition relation. In essence, the transitions correspond directly to the manner in which the machine state is altered by the execution of instructions between scheduler invocations. Moreover, when a counterexample is found, the model checker can generate a concrete execution demonstrating how the system violates the property. This counterexample is typically detailed enough to be replayed inside a debugger, thereby helping the user pinpoint the cause of error.

Execution-based model checkers are typically used as a “test amplification” mechanism. The user provides a test harness corresponding to a workload under which the program is run. The usual operating system scheduler would execute the workload under a fixed schedule, thereby missing most possible behaviors. However, the execution-based model checker’s scheduler is able to systematically explore the possible executions of the same workload under different schedules, for example, exploring what happens under different interleavings of shared operations like message sends, receives etc., and is able to find various safety errors like assertion failures, deadlocks and divergences, that are only manifested under corner-case schedules.

There are two main technical challenges that must be addressed to make execution-based model checking feasible. These challenges stem from the fact that in this approach the “state” comprises the entire machine state – all the registers, the heap, the stack, state of network resources etc.. First, the size of the machine state makes it infeasible to store the set of visited states. Thus, how can one systematically explore the space of executions *without* storing the visited states? Second, the space of possible schedules grows exponentially in the length of the schedule. However, many of these schedules can lead to states that are either identical, or identical with respect to the properties being analyzed (i.e., differing only in the values of some irrelevant variables like performance counters). Given that one can only run the model checker for a finite amount of time, how can one *bias* the analysis away from exploring redundant executions and

Algorithm: Stateless Search
Input Program P , Depth D .
Output NOERRORSFOUND in schedules of size $\leq D$, UNSAFE otherwise.

```

def StatelessSearch( $P, D$ ):
     $depth = 0$ 
    while  $depth \leq D$  do:
        increment  $depth$ 
        foreach  $schedule$  in Sequences( $depth$ ) do:
            Reset( $sys$ )
            if Execute( $P, schedule$ ) reaches an error :
                return UNSAFE
    return NOERRORSFOUND

```

Fig. 3. StatelessSearch.

towards executions that are more likely to expose bugs? There is the usual time-space-soundness tradeoff here: storing entire states ensure paths are not re-executed but require too much space; storing no states (“stateless search” described below) spends time re-executing similar paths and can diverge if there are loops; and storing hashes of states can miss executions owing to hash collisions.

2.3. Stateless Search

The key innovation in VERISOFT was to introduce the idea of stateless search, wherein the model checker could explore different executions without actually storing the set of visited states. Since all nondeterminism is captured within the schedules, whenever the process needs to make a nondeterministic choice, for example, a choice based on a random number, a user input, the choice of thread to execute or network latency, it queries the scheduler for a number from a finite range. Each execution is characterized by the *schedule*, that is, the sequence of numbers returned by the scheduler to the process. Thus, to iterate over different executions, the model checker needs only to iterate over the space of possible schedules and re-execute the process using the schedule.

While conceptually the schedule does not distinguish user inputs and thread schedule choices, even for a small number of 32-bit user inputs, the number of schedules can become astronomical. VERISOFT (and many similar tools described below) explicitly model the non-determinism arising from scheduling choices, keeping user inputs fixed. Tools based on symbolic execution (described in Section 3) for sequential program focus on the other hand on exploring the space of user inputs, using symbolic representations to avoid the explosion in schedules. Tools such as JAVAPATHFINDER combine both approaches, exploring schedule choices explicitly while tracking user inputs symbolically.

Algorithm StatelessSearch in Figure 3 shows this algorithm. The algorithm takes as input a system and a depth increment, and performs a *bounded-depth depth first search* of the space of executions. The size of an execution is measured by the number of calls made to the scheduler during the execution.

Recall that each execution of size $depth$ is characterized by the sequence of values that the scheduler returns along that execution. Suppose that at each point, the scheduler returns an integer between 0 and k . $Sequences(depth)$ is an iterator over the sequences of $1 \dots k$ of size $depth$. That is, it returns, in order, the sequences,

$$0^{depth}, 0^{depth-1}1, \dots, 0^{depth-1}k, 0^{depth-2}10, \dots, k^{depth}$$

For each sequence or schedule, we analyze the behavior of the system under the schedule by resetting the system to its initial state and executing the system under the schedule. The order in which sequences of a given depth *depth* are generated ensures that algorithm explores executions in a depth-first manner. Once all the executions of size *depth* have been explored, the *depth* is incremented and the process is repeated until the execution returns UNSAFE meaning that an execution along which some safety property has been violated has been found.

The key property of the above algorithm is that it makes very modest demands on the representation of the process being analyzed. We need only to be able to: (1) *reset* it, or restore it to a unique starting state, and, (2) *execute* it, under a given schedule. Notice that to explore executions of larger depths, the algorithm simply re-executes the system from the initial state. This requires that the execution be deterministic under a given schedule. That is, two executions using the same schedule leave the system in identical states. In practice, the most onerous burden upon the user is to ensure that *every* source of non-determinism is replaced with a call to the scheduler.

2.4. Execution-Based Tools

Next, we describe several execution-based model checkers, and the strategies they have adopted to effectively explore the executions in order to find bugs.

Verisoft. As mentioned before, the VERISOFT tool pioneered the idea of execution-based stateless model checking of software. VERISOFT takes as input the composition of several Unix *processes* that communicate by means of message queues, semaphores and shared variables that are visible to the VERISOFT scheduler. The scheduler traps calls made to access the shared resources, and by choosing which process will execute at each trap point, the scheduler is able to exhaustively explore all possible interleavings of the processes' executions. VERISOFT has been used to found several complex errors in concurrent phone-switching software [Chandra et al. 2002].

JavaPathFinder is an execution-based model checker for Java programs that modifies the *Java Virtual Machine* to implement systematic search over different thread schedules [Havelund and Pressburger 2000; Visser et al. 2003]. This language-based approach restricts the classes of software that can be analyzed, but provides many significant advantages. First, the use of the JVM makes it possible to store the visited states, which allows the model checker to use many of the standard reduction-based approaches (e.g., symmetry, partial-order, abstraction) to combating state-explosion. Second, as the visited states can be stored, the model checker can utilize various search-order heuristics without being limited by the requirements of stateless search. Third, one can use techniques like symbolic execution and abstraction to compute inputs that force the system into states that are different from those previously visited thereby obtaining a high level of coverage. JAVAPATHFINDER has been used to successfully find subtle errors in several complex Java components developed inside NASA [Brat et al. 2004; Penix et al. 2005], and is available as a highly extensible open-source tool.

Cmc is an execution based model checker for C programs that explores different executions by controlling schedules at the level of the OS scheduler. CMC stores a hash of each visited state. In order to identify two different states which differ only in irrelevant details like the particular addresses in which heap-allocated structures are located, CMC *canonicalizes* the states before hashing to avoid re-exploring states that are similar to those previously visited. CMC has been used to find errors in implementations of network protocols [Musuvathi and Engler 2004] and file systems [Yang et al. 2004].

MaceMC is an execution-based model checker for distributed systems implemented in MACE, a domain-specific language built on top of C++ [Killian et al. 2007]. MACEMC

uses two techniques to guide state space exploration. First, rather than exploring the interleavings of low-level operations like network sends and receives, MACEMC exploits higher-level constructs in its modeling language MACE to explore only coarse-grained interleavings of event-transitions, each of which can comprise of tens of low-level shared operations. Second, it combines exhaustive search with long random walks executed from the periphery of the exhaustive search. Long walks that end without “something good” happening indicate potential violations of liveness properties.

Chess is an execution based model checker for multithreaded Windows programs [Musuvathi and Qadeer 2007]. Like CMC and VERISOFT, CHESS works by intercepting system calls in order to systematically explore the space of schedules. CHESS employs an innovative search ordering called *iterative context bounding* in which the tool explores executions with at most k context-switches, where k is a parameter that is iteratively increased [Qadeer and Rehof 2005]. The intuition behind this search-ordering is that many bugs are manifested in long executions containing just a few unexpected context-switches. CHESS has been incorporated inside the testing framework for several programs inside Microsoft.

In addition, many of the techniques for execution-based model checking have been incorporated in newer versions of SPIN.

3. CONCRETE SYMBOLIC MODEL CHECKING

While enumerative techniques capture the essence of software model checking as the exploration of program states and transitions, their use in practice is often hampered by severe state space explosion. This led to research on *symbolic* algorithms which manipulate representations of *sets* of states, rather than individual state, and perform state exploration through the symbolic transformation of these representations. For example, the constraint $1 \leq x \leq 10 \wedge 1 \leq y \leq 8$ represents the set of all states over $\{x, y\}$ satisfying the constraint. Thus, the constraint *implicitly* represents the list of 80 states that would be enumerated in enumerative model checking. Symbolic representations of sets of states can be much more succinct than the corresponding enumeration, and can represent infinite sets of states as well. The *symbolic representation* of regions is the crucial component of symbolic algorithms, but in addition to providing an implicit representation of sets of states, the representation must also allow performing certain operations on sets of states *directly* on the representation. For example, given a program with (symbolic) initial region σ_I , the set of reachable states are given by $\bigcup_{i \geq 0} Post^i(\sigma_I)$. This suggests that a symbolic representation of regions must allow at least the *Post* and the \cup operations on symbolic regions, and a way to check inclusion between regions (to check for convergence). Performing these operations by enumerating individual states in the region nullifies the advantage of the symbolic representation.

3.1. The Region Data Structure

Let P be a simple program. For symbolic model checking algorithms, we introduce the abstract data type of *symbolic representations* for states of P . The abstract data type defines a set of symbolic regions symreg , an *extension function* $\llbracket \cdot \rrbracket : \text{symreg} \rightarrow 2^{v.X}$ mapping each symbolic region to a set of states, and the following constants and operations on regions:

- (1) The constant $\perp \in \text{symreg}$ representing the empty set of states, $\llbracket \perp \rrbracket = \emptyset$, and the constant $\top \in \text{symreg}$ representing the set of all states, $\llbracket \top \rrbracket = v.X$.
- (2) The operation $\cup : \text{symreg} \times \text{symreg} \rightarrow \text{symreg}$ that computes the union of two regions, that is, for every $r, r' \in \text{symreg}$ we have $\llbracket r \cup r' \rrbracket = \llbracket r \rrbracket \cup \llbracket r' \rrbracket$.

Algorithm: Symbolic Reachability
Input simple program $P = (X, L, \ell_0, T)$, error location $\mathcal{E} \in L$
Output SAFE if P is safe w.r.t. \mathcal{E} , UNSAFE otherwise

```

def SymbolicReachability( $P, \mathcal{E}$ ):
    reach =  $\lambda \ell. \perp$ 
    worklist =  $(\ell_0, \top)$ 
    while worklist  $\neq \emptyset$  do:
        choose  $(\ell, r) \in$  worklist; worklist = worklist  $\setminus \{(\ell, r)\}$ 
        if  $r \not\subseteq$  reach( $\ell$ ):
            reach = reach[ $\ell \mapsto r \cup$  reach( $\ell$ )]
            foreach  $(\ell, \rho, \ell')$  in  $T$  do:
                worklist = worklist  $\cup \{(\ell', Post(r, \rho))\}$ 
    if reach( $\mathcal{E}$ )  $\neq \perp$ :
        return UNSAFE
    else:
        return SAFE

```

Fig. 4. SymbolicReachability.

- (3) The operation $\cap : \text{symreg} \times \text{symreg} \rightarrow \text{symreg}$ that computes the intersection of two regions, that is, for every $r, r' \in \text{symreg}$ we have $\llbracket r \cap r' \rrbracket = \llbracket r \rrbracket \cap \llbracket r' \rrbracket$.
- (4) The operation $= : \text{symreg} \times \text{symreg} \rightarrow \text{bool}$ such that $r = r'$ returns *true* iff r and r' denote the same set of states, that is, $\llbracket r \rrbracket = \llbracket r' \rrbracket$.
- (5) The operation $\subseteq : \text{symreg} \times \text{symreg} \rightarrow \text{bool}$ such that $r \subseteq r'$ returns *true* iff r denotes a region contained in r' , that is, $\llbracket r \rrbracket \subseteq \llbracket r' \rrbracket$.
- (6) The operation $Post : \text{symreg} \times T \rightarrow \text{symreg}$ that takes a symbolic region r and a constraint ρ and returns a symbolic region denoting the set $Post(\llbracket r \rrbracket, \rho)$.
- (7) The operation $Pre : \text{symreg} \times T \rightarrow \text{symreg}$ that takes a symbolic region r and a constraint ρ and returns a symbolic region denoting the set $Pre(\llbracket r \rrbracket, \rho)$.

In what follows, we shall assume that each operation above is effectively computable for an implementation of `symreg`.

Figure 4 shows Algorithm `SymbolicReachability`, a symbolic implementation of reachability. The basic search procedure is the same as Algorithm `EnumerativeReachability` (Figure 2), but we now manipulate symbolic regions rather than individual states. Note that we maintain the program locations enumeratively, alternately, we can construct a fully symbolic version where the program locations are maintained symbolically as well. The algorithm maintains a map $reach$ from locations $\ell \in L$ to sets of states reachable when the control location is ℓ . The frontier regions are maintained in $worklist$ as before, and the main loop explores regions at the frontier and adds new regions to $reach$. (The notation $reach[\ell \mapsto r]$ denotes a function that maps ℓ to r but agrees with $reach$ on all other locations.)

Notice that the algorithm does not make any assumptions of finiteness: as long as the symbolic regions can represent infinite sets of states, and the symbolic operations are effective, the algorithm can be implemented.

The power of symbolic techniques comes from tremendous advances in the performance of constraint solvers that underlie effective symbolic representations, both for propositional logic (satisfiability solvers [Silva and Sakallah 1996; Moskewicz et al. 2001; Eén and Sorensson 2003] as well as binary decision diagrams [Bryant 1986; Somenzi 1998]) and more recently for combinations of first order theories [Dutertre and de Moura 2006; Bruttomesso et al. 2008; de Moura and Bjørner 2008].

3.2. Example: Propositional Logic

For finite domain programs, Boolean formulas provide a symbolic representation. The encoding is easiest understood when each variable in the program is Boolean (finite domain variables can be encoded using several Boolean variables). Suppose the program has n Boolean variables. Then, a state is a vector of n bits, and each set of states can be represented by its characteristic formula that maps an n bit vector to *true* if it is in the set and to *false* otherwise.

In this representation, the regions \perp and \top are then the formulas *false* and *true*, respectively, Boolean operations on regions are logical operations on formulas, and equality and subset checks reduce to Boolean equivalence and implication checks. Finally, *Post* and *Pre* can be computed by Boolean conjunction followed by existential quantifier elimination (and renaming) as follows. Given a transition relation $\tau(X, X')$ represented as a formula over $2n$ variables (the unprimed and primed variables), and a formula $b(X)$ over the current set of variables, $Post(b(X), \tau(X, X'))$ is given by the formula

$$\text{Rename}(\exists X. b(X) \wedge \tau(X, X'), X', X)$$

where the existential quantifier $\exists X$ existentially quantifies each variable in X , and $\text{Rename}(\phi(X'), X', X)$ renames each primed variable $x' \in X'$ appearing in the formula $\phi(X')$ by the corresponding unprimed variable $x \in X$. Similarly, $Pre(b(X), \tau(X, X'))$ is given by the formula

$$\exists X'. \text{Rename}(b(X), X, X') \wedge \tau(X, X')$$

Symbolic representation can result in very compact representations of states. Compare, for example, a program with n Boolean variables, and the set of all program states. An enumerative representation would explicitly represent 2^n program states, whereas a symbolic representation using Boolean formulas is simply the formula *true*.

Unfortunately, it is often difficult to work directly with Boolean formulas. (Reduced ordered) binary decision diagrams (BDDs) [Bryant 1986] can be used instead as an efficient implementation of Boolean formulas. BDDs are compact and canonical representations of Boolean functions.

For a program with n state bits, regions are represented as BDDs over n variables with some ordering on the variables. As before, the empty region is the BDD for *false*, representing the empty set, and the top region \top is the BDD for *true*, representing all states. Union and intersection are Boolean disjunction and conjunction, respectively, which can be computed directly on BDDs [Bryant 1986]. Equality is Boolean equivalence, but reduces to equality on BDDs since they are canonical representations. Checking containment reduces to equality, using the observation that $r_1 \subseteq r_2$ iff $r_1 \cap r_2 = r_1$. The computation of *Post* and *Pre* require existential quantification and renaming (in addition to Boolean operations), which can again be implemented directly on BDDs.

BDDs are the primary representation in symbolic model checkers such as SMV, and have been instrumental in scaling hardware model checkers to extremely large state spaces [Burch et al. 1992; McMillan 1993]. Each Boolean operation and existential quantification of a single variable can be quadratic in the size of the BDD, and the size of the BDD can be exponential in the number of variables in the worst case. Moreover, the size of the BDD is sensitive to the variable ordering, and many functions do not have a succinct BDD representation [Bryant 1986]. This is the symbolic analogue of the state explosion problem, and has been a major research direction in model checking.

Notice that Algorithm 4 terminates for imp_{fin} where symreg is implemented using BDDs, since each BDD operation is effective, and each iteration of the loop finds at least one new state.

BDD-based model checkers, such as SMV [McMillan 1993], have been extremely successful in hardware model checking. They were also used as back-ends for initial attempts at software model checking [Corbett et al. 2000].

3.3. Example: First-Order Logic with Interpreted Theories

In case program variables range over infinite domains, such as integers, a more expressive symbolic representation is provided by first order logic formulas [Floyd 1967; Hoare 1969]. A symbolic region is represented by a first order formula whose free variables range over the program variables. For a formula φ , we write $\llbracket \varphi \rrbracket$ for the set of states $\{s \mid s \models \varphi\}$. With abuse of notation, we often identify a formula φ with the set of states $\llbracket \varphi \rrbracket$, and omit $\llbracket \cdot \rrbracket$ when clear from the context. The empty region \perp is represented by the logical formula *false*, the region \top by *true*. Union and intersection are implemented by disjunction and conjunction in the logic, and equality and containment by equivalence checking and implication, respectively.

Finally, the *Pre* and *Post* operations can again be implemented using existential quantification, using essentially the same formulations as for the Boolean case above, but where the formulas can now be over the more general logic.

When using the full power of first-order logic, the equivalence and implication checks are not effective. So in practice, one chooses formulas over a *decidable* theory, such as quantifier-free formulas over a combination of the theory of equality with uninterpreted functions and the theory of rationals [Nelson 1981].

While the availability of decision procedures makes the individual operations in Algorithm 4 effective, it could still be that the algorithm runs forever. Consider for example a loop

$$i := 0; \mathbf{while} (i \geq 0) \{i := i + 1;\} \mathcal{E} ;;$$

Clearly, the location \mathcal{E} is not reachable. However, a symbolic representation based on first order formulas could run forever, finding the set of reachable states

$$i = 0 \vee i = 1 \vee \dots \vee i = k$$

at the k th iteration of the while loop, approximating the “fixed point” $i \geq 0$ closer and closer.

3.4. Bounded Model Checking

As in enumerative model checking, one can trade-off soundness for effective bug finding in symbolic model checking. One popular approach, called *bounded model checking* [Biere et al. 1999], unrolls the control flow graph for a fixed number of steps, and checks if the error location can be reached within this number of steps. Precisely, given program P , error location \mathcal{E} , and $k \in \mathbb{N}$, one constructs a constraint which is satisfiable iff the error location \mathcal{E} is reachable within k steps. Satisfiability of this constraint is checked by a constraint solver. The technique is related to *symbolic execution* [King 1976], in which the program is executed on symbolic as opposed to concrete inputs. While BMC techniques search over all program computations using backtracking search within the constraint solver, traditionally, symbolic execution techniques enumerate all program paths, and generate and solve constraints for each enumerated path.

Tools for bounded model checking of software implementations come in two flavors. The first, such as CBMC [Kroening et al. 2003], F-SOFT [Ivancic et al. 2008], Saturn [Xie and Aiken 2005], or Calysto [Babic and Hu 2008] generate constraints in propositional logic and use Boolean satisfiability solvers to discharge the constraints. Scalability of the techniques depend both on the scalability of the underlying SAT solvers as well as carefully tuned heuristics which keep the size of the constraints small. The reduction to propositional satisfiability captures the semantics of fixed-width program datatypes precisely. Thus, one can find subtle bugs arising from mismatches between the algorithm and low-level fixed-width machine semantics, such as arithmetic overflows.

CBMC and Saturn are tools implementing this idea for C programs. Both have been fairly successful in analyzing large pieces of software, including analyzing C models of processors and large parts of the Linux kernel. Saturn improves upon the basic bounded model checking algorithm by computing and memoizing relations between inputs and outputs (“summaries”) for procedures bottom-up in the call graph. This makes bounded model checking scale to large programs.

The second class of tools generates constraints in an appropriate first order theory (in practice, the combination theory of equality with uninterpreted functions, linear arithmetic, arrays, and some domain-specific theories) and use decision procedures for such theories [de Moura et al. 2002; Ganai and Gupta 2006; Armando et al. 2006]. The basic algorithm is identical to SAT-based bounded model checking, but the constraints are interpreted over more expressive theories.

3.5. Invariants and Invariant Synthesis

The set of reachable states R of a program is the smallest set which contains the set of initial states and is closed under the *Post* operator (i.e., $Post(R) \subseteq R$). By definition, a program is safe if the location \mathcal{E} does not appear in the set of reachable states. Instead of computing the set of reachable states, though, one could compute (possibly larger) regions R' which satisfy the two constraints that the set of initial states is contained in R' and R' is closed under *Post*. As long as the location \mathcal{E} does not appear in R' , we could still deduce that the program is safe. Such regions are called *invariants*.

An *invariant* of P at a location $\ell \in L$ is a set of states containing the states reachable at ℓ . An *invariant map* is a function η from L to formulas over program variables from X such that the following conditions hold:

Initiation: For the initial location ℓ_0 , we have $\eta.\ell_0 = true$.

Inductiveness: For each $\ell, \ell' \in L$ such that $(\ell, \rho, \ell') \in T$, the formula $\eta.\ell \wedge \rho$ implies $(\eta.\ell')$. Here, $(\eta.\ell')$ is the formula obtained by substituting variables from X' for the variables from X in $\eta.\ell'$.

Safety: For the error location \mathcal{E} , we have $\eta.\mathcal{E} = false$.

With these properties, it can be proved by induction that at every location $\ell \in L$, we have that $\{s \mid (\ell, s) \text{ is reachable in } P\} \subseteq \llbracket \eta.\ell \rrbracket$, and so the method is sound. In the above example, a possible invariant map associates the invariant $i \geq 0$ with the head of the while loop.

The presence of invariants reduces the problem of iteratively computing the set of reachable states to checking a finite number of obligations, which is possible if the symbolic representation is effective. In fact, relative completeness results [Cook 1978] ensure that relative to deciding the implications, the method is complete for safety verification. Traditional program verification has assumed that the invariant map is provided by the programmer, and several existing tools (ESC-Modula [Leino and Nelson 1998], ESC-Java [Flanagan et al. 2002]) check the conditions for invariants, given an

invariant map, using decision procedures. (For ease of exposition, we have assumed that an invariant map assigns an invariant to each program location, but it is sufficient to define invariants only over a *cutset* of the program, i.e., a set of program locations such that every syntactic cycle in the CFG passes through some location in the cutset. Thus, the tools only require *loop invariants* from the programmer.)

In practice, the requirement that the programmer provide invariants has not found widespread acceptance, so the emphasis has shifted to algorithmic techniques that can synthesize invariant maps automatically, with minimal support from the programmer. Clearly, the set of reachable states at each location constitutes an invariant map. However, there are other invariant maps as well, each of which contains the set of reachable states. This is a crucial observation, and reduces the safety verification problem to the search for appropriate invariants. For example, it enables the use of *abstraction techniques* for invariant synthesis that we look at in the next section.

Several modifications to the inductiveness condition for invariants have been studied. One promising approach is *k-induction* [Sheeran et al. 2000; de Moura and Ruesch 2003], which checks the inductiveness condition not for an edge (ℓ, ℓ') but for paths of length k . That is, instead of performing an induction scheme in which one assumes the inductive invariant at the current state and proves the invariant holds after one more step, *k-induction* assumes the inductive invariant holds for the previous k consecutive steps along a path, and proves that the invariant continues to hold after one more step. Paths of length k in the program are encoded using bounded model checking constraints.

We close this section with a constraint-based algorithm for invariant synthesis. The method starts with a *template invariant map*, that is, a parameterized expression representing the invariant map, and encodes the three conditions (initiation, inductiveness, safety) on the templates to get a set of constraints on the parameters of the template [Giesl and Kapur 2001; Sankaranarayanan et al. 2005]. A solution to these constraints provides an assignment to the parameters in the template and constitutes an invariant map. For templates in linear rational arithmetic, the constraint system can be encoded as a set of *nonlinear* arithmetic constraints on the template parameters [Sankaranarayanan et al. 2005], and decision procedures for real arithmetic can be used to construct invariant maps. Invariant synthesis for templates over more expressive theories, such as combinations of linear arithmetic and equality with uninterpreted functions, can be reduced to nonlinear arithmetic constraints [Beyer et al. 2007a].

While attractive, the technique is limited by two facts. First, the programmer has to guess the “right” template. Second, the scalability is limited by the performance of constraint solvers for nonlinear arithmetic. Thus, so far constraint-based invariant synthesis has so far been applied only to small programs, even though recent techniques combining the approach with other model checking techniques and constraint simplification are promising [Beyer et al. 2007b; Gupta et al. 2009].

4. MODEL CHECKING AND ABSTRACTION

For infinite state programs, symbolic reachability analysis may not terminate, or take an inordinate amount of time or memory to terminate. Abstract model checking trades off precision of the analysis for efficiency. In abstract model checking, reachability analysis is performed on an *abstract domain* which captures some, but not necessarily all, the information about an execution, using an abstract semantics of the program [Cousot and Cousot 1977]. A proper choice of the abstract domain and the abstract semantics ensures that the analysis is sound (i.e., proving the safety property in the abstract semantics implies the safety property in the original, *concrete*, semantics) and efficient.

Algorithm: Abstract Reachability
Input simple program $P = (X, L, \ell_0, T)$, error location $\mathcal{E} \in L$
Input abstract domain $A = (\mathcal{L}, \llbracket \cdot \rrbracket, Pre^\#, Post^\#)$
Output SAFE if P is safe w.r.t. \mathcal{E} , “Don’t know” otherwise

```

def AbstractReachability( $P, \mathcal{E}, A$ ):
    reach# =  $\lambda \ell. \perp$ 
    worklist =  $(\ell_0, \top)$ 
    while worklist  $\neq \emptyset$  do:
        choose  $(\ell, r^\#)$  from worklist; worklist = worklist  $\setminus \{(\ell, r^\#)\}$ 
        if  $r^\# \not\sqsubseteq reach^\#(\ell)$ :
            reach# = reach# $[\ell \mapsto r^\# \sqcup reach^\#(\ell)]$ 
            foreach  $(\ell, \rho, \ell')$  in  $T$  do:
                worklist = worklist  $\cup \{(\ell', Post^\#(r^\#, \rho))\}$ 
    if reach#( $\mathcal{E}$ )  $\neq \perp$ :
        return UNSAFE
    else:
        return SAFE

```

Fig. 5. Abstract Reachability.

Historically, the two fields of model checking and static program analysis have evolved in parallel, with model checking emphasizing precision and static analysis (with applications to program optimization) emphasizing efficiency. It has been long known that theoretically each approach can simulate the other [Steffen 1991; Schmidt 1998; Schmidt and Steffen 1998; Cousot and Cousot 2000]. In what follows, we focus on abstract *reachability* analysis, but the techniques generalize to model checking more expressive temporal properties, for example, those expressed in the μ -calculus [Clarke et al. 1992; Cousot and Cousot 2000].

4.1. Abstract Reachability Analysis

An *abstract domain* $(\mathcal{L}, \llbracket \cdot \rrbracket, Pre^\#, Post^\#)$ for a program P consists of a complete lattice $\mathcal{L} = (L, \sqsubseteq, \perp, \top, \sqcup, \sqcap)$ of *abstract elements*, a *concretization function* $\llbracket \cdot \rrbracket : L \rightarrow 2^{v.S}$ mapping lattice elements to sets of states, and two monotone total functions $Pre^\# : L \times T \rightarrow L$ and $Post^\# : L \times T \rightarrow L$ such that the following conditions hold:

- (1) $\llbracket \perp \rrbracket = \emptyset$ and $\llbracket \top \rrbracket = v.S$;
- (2) for all elements $l, l' \in L$, we have $\llbracket l \sqcup l' \rrbracket \supseteq \llbracket l \rrbracket \cup \llbracket l' \rrbracket$ and $\llbracket l \sqcap l' \rrbracket \supseteq \llbracket l \rrbracket \cap \llbracket l' \rrbracket$; and
- (3) for all $l \in L$ and transition ρ , we have $\llbracket Post^\#(l, \rho) \rrbracket \supseteq Post(\llbracket l \rrbracket, \rho)$ and $\llbracket Pre^\#(l, \rho) \rrbracket \supseteq Pre(\llbracket l \rrbracket, \rho)$.

A lattice element $l \in L$ represents an *abstract view* of a set of program states $\llbracket l \rrbracket$. Note that $\llbracket \cdot \rrbracket$ is not required to be onto: not all sets of program states need to have an abstract representation. A strictly ascending chain is a sequence $l_0 \sqsubset l_1 \sqsubset l_2 \dots$. The *height* of a lattice is the cardinality of the largest strictly ascending chain of elements in L .

Figure 5 shows the abstract reachability algorithm. It is similar to the symbolic reachability algorithm (Algorithm 4), but instead of using symbolic regions, uses an abstract domain. It takes as input a program P , an error location \mathcal{E} , and an abstract domain A , and returns either `SAFE`, signifying P is safe with respect to \mathcal{E} , or `UNSAFE`, signifying that either the program is unsafe or the abstract domain could not prove that the program is safe. From the properties of an abstract domain, the abstractly reachable set $reach^\#$ so computed has the property $reach \subseteq \llbracket reach^\# \rrbracket$, so if the abstract reachability algorithm returns “safe” then the program is indeed safe with respect to the error

```

0: LOCK = 0;
1: do {
    LOCK = 1;
    old = new;
2:   if (*) {
3:     LOCK = 0;
    new++;
    }
4: } while (new != old);
5: if (LOCK==0)
6:   error();
   LOCK = 0;

```

Fig. 6. Program.

location. Unfortunately, the converse is not true in general: the abstract reachability could return “unsafe” if it loses too much precision in the abstraction process, even though the program is safe.

Consider the example program shown in Figure 6. This program is a simplified version of a function from a network device driver [Ball and Rajamani 2002b]. Intuitively, the variable `LOCK` represents a global lock; when the lock has been acquired, the value of `LOCK` is 1, and when the lock is released, the value of `LOCK` is 0. We would like to verify that at the end of the do-loop, the lock is acquired (i.e., `LOCK = 1`). In the code, this assertion is specified by a check that the lock is acquired (on line 5) and a call to `error` if the check fails. If the abstraction consists of the relations `LOCK = 0`, `LOCK = 1`, `old = new`, and `old ≠ new`, the abstract reachability analysis can prove this property. The abstraction of the set of reachable set of states at line 4 is

$$(\text{LOCK} = 1 \wedge \text{new} = \text{old}) \vee (\text{LOCK} = 0 \wedge \text{new} \neq \text{old})$$

which captures the intuition that at line 4, either the lock is acquired and `new` is equal to `old`, or the lock is not acquired and the value of `new` is different from `old` (in fact, `new = old + 1`).

On the other hand, if the program is analyzed using only the predicates `LOCK = 0` and `LOCK = 1`, the abstract reachability analysis does not track the relationship between `new` and `old`, and hence cannot prove the property.

Notice that if the abstract domain has finite height, the abstract reachability algorithm is guaranteed to terminate. Even if the abstract domain has infinite height, the abstract reachability algorithm is still applicable, but usually augmented with *widening* techniques that ensure termination in a finite number of steps, at the cost of reduced precision [Cousot and Halbwachs 1978].

4.2. Example: Polyhedral Domains

In the *polyhedral abstract domain*, the abstract elements are polyhedral sets over an n -dimensional space of program variables ordered by set inclusion. The intersection operation for the domain is implemented simply as polyhedron intersection, and union is implemented as convex hull. The $Pre^\#$ and $Post^\#$ operations are implemented as intersections and projections of polyhedra for transition relations ρ that are linear relations on $X \cup X'$, and approximated for others.

The polyhedral domain has been successfully used to check for array bounds checks [Cousot and Halbwachs 1978], and efficient implementations are available (e.g., Bagnara et al. [2008]). The domain is not finite height, so the abstract reachability algorithm may not terminate after a finite number of steps. To ensure termination, the reachability analysis uses widening, and several widening heuristics have been studied [Cousot and Halbwachs 1978; Bagnara et al. 2005].

Faster, but less expressive, abstract domains that can represent a subclass of polyhedra, such as intervals [Cousot and Cousot 1976] or octagons [Miné 2006] have been used as well.

4.3. Example: Predicate Abstraction

The *predicate abstraction* domain [Agerwala and Misra 1978; Graf and Saïdi 1997; Saïdi and Shankar 1999; Das et al. 1999] is parameterized by a fixed finite set Π of first order formulas with free variables from the program variables, and consists of the lattice of Boolean formulas over Π ordered by implication. Let ψ be a region. The predicate abstraction of ψ with respect to the set Π of predicates is the smallest (in the implication ordering) region $\text{Abs}(\psi, \Pi)$ which contains ψ and is representable as a *Boolean combination* of predicates from Π :

$$\text{Abs}(\psi, \Pi) = \bigwedge \{ \phi \mid \phi \text{ is a Boolean formula over } \Pi \wedge \psi \Rightarrow \phi \}$$

The region $\text{Abs}(\psi, \Pi)$ can be computed by recursively splitting as follows [Das et al. 1999]:

$$\text{Abs}(\psi, \Pi) \doteq \begin{cases} \text{true} & \text{if } \Pi = \emptyset \text{ and } \psi \text{ satisfiable} \\ \text{false} & \text{if } \Pi = \emptyset \text{ and } \psi \text{ unsatisfiable} \\ (p \wedge \text{Abs}(\psi \wedge p, \Pi')) \vee (\neg p \wedge \text{Abs}(\psi \wedge \neg p, \Pi')) & \text{if } \Pi = \{p\} \cup \Pi' \end{cases}$$

The satisfiability checks can be discharged by a decision procedure [Nelson 1981; Dutertre and de Moura 2006; de Moura and Bjørner 2008]. In the worst case, the computation is exponential in the number of predicates, and several heuristics with better performance in practice have been proposed [Saïdi and Shankar 1999; Flanagan and Qadeer 2002].

Many implementations of predicate-based software model checkers (including SLAM and BLAST) implement an over-approximation of the predicate abstraction that can be computed efficiently in order to avoid the exponential cost. *Cartesian* predicate abstraction is one such precision-efficiency tradeoff: it can be computed more efficiently than full predicate abstraction but can be quite imprecise in the worst case. Cartesian abstraction formalizes the idea of ignoring relations between components of tuples, and approximates a set of tuples by the smallest Cartesian product containing the set [Ball et al. 2001]. Formally, the cartesian abstraction of ψ with respect to the set Π of predicates is the smallest (in the implication ordering) region $\text{CartAbs}(\psi, \Pi)$ which contains ψ and is representable as a *conjunction* of predicates from Π . The region $\text{CartAbs}(\psi, \Pi)$ can be computed as:

$$\text{CartAbs}(\psi, \Pi) \doteq \begin{cases} \text{true} & \text{if } \Pi = \emptyset \\ p \wedge \text{CartAbs}(\psi, \Pi') & \text{if } \Pi = \{p\} \cup \Pi' \text{ and } (\psi \wedge \neg p) \text{ unsatisfiable} \end{cases}$$

Cartesian predicate abstraction was implemented for C programs as part of SLAM in a tool called c2bp [Ball et al. 2001], and since then in other software verifiers.

4.4. Example: Control Abstraction

So far, we have used the abstract domain to reduce the space of data values, while keeping each path of the program precise. Thus, our algorithms so far are *path-sensitive*. In an orthogonal direction, we can develop abstract reachability algorithms in which different paths of the program are merged into equivalence classes. For example, in a *flow-sensitive, path-insensitive* analysis, the reachability algorithm will merge the set of abstract states coming into a program location from any of its predecessor locations. In a *flow-insensitive* analysis, the ordering on program transitions is abstracted, and the program is considered as a bag of transitions which can fire in any order. Historically, model checking has always assumed path-sensitivity in the analysis, and dataflow analysis has rarely assumed path-sensitivity (or even flow-sensitivity). A uniform framework for abstract reachability analysis, which gives flow-insensitive, flow-sensitive, and path-sensitive analyses as special cases is provided in Beyer et al. [2007b].

We have primarily considered the verification of sequential programs. We can reduce the verification of nonrecursive concurrent programs (e.g., nonrecursive multithreaded programs), to sequential programs, by taking the product of the control flow graphs of the different threads [Dwyer and Clarke 1994], and then applying the sequential analysis. However, in many cases, combinatorial explosion makes a product construction prohibitively expensive. In such cases, one approach is to perform a control abstraction for the individual threads, before analyzing the product. To do so, one can first use a data abstraction (e.g., predicate abstraction, polyhedra) to compute finite-state abstractions of the individual threads, and next, apply (bi)simulation quotienting [Bouajjani et al. 1990] to obtain a small control skeleton for each thread. While the direct products may be too large to analyze, it can be feasible to analyze the products of the reduced state machines. This approach was suggested by the MAGIC software model checker [Chaki et al. 2004]. One can refine this approach by using compositional reasoning to iteratively compute an abstraction for each thread that is sound *with respect to* the behavior of the other threads, to obtain a *thread-modular* style of reasoning [Flanagan et al. 2002, 2005; Henzinger et al. 2003]. Instead of (bi)simulation quotients, one can use the L^* algorithm from machine learning to compute a small state machine that generates the observable traces for each thread [Pasareanu et al. 2008]. Intuitively, in each of the above cases, the small state machine computed for each thread can be viewed as a *summary* of the behavior of the thread relevant to proving the property of the concurrent program.

4.5. Combined Abstractions

Finally, one can build powerful analyses by combining several different abstractions, each designed for capturing different kind of property of the program. One way to achieve the combination is to analyze the program in two (or more) stages. In the first stage, one can use a particular abstraction (e.g., polyhedra or octagons) to compute invariants, which can be used to strengthen the abstract regions computed in the second stage. This approach, which is implemented in the F-SOFT [Jain et al. 2006] and IMPACT [McMillan 2006] model checkers, can make verification much faster if the first phase can cheaply find invariants that are crucial for the second phase, but which are expensive to compute using a more general abstraction of the second phase. A second approach is to simultaneously combine multiple abstractions using the notion of a *reduced product* [Cousot and Cousot 1979]. The Astree analysis tool Blanchet et al. [2002,

Algorithm: Counterexample Guided Abstraction Refinement
Input simple program P , error location $\mathcal{E} \in L$
Output SAFE if P is safe w.r.t. \mathcal{E} , UNSAFE otherwise

```

def CEGAR( $P, \mathcal{E}$ ):
   $A$  = Initial Abstraction
  while true do:
    match ModelCheck( $P, \mathcal{E}, A$ ) with
    | SAFE:
      return SAFE
    | UNSAFE( $\bar{\rho}$ ):
      match Refine( $P, A, \bar{\rho}$ ) with
      | GENUINE: return UNSAFE
      | SPURIOUS( $A'$ ): Refine  $A$  using  $A'$ 

```

Fig. 7. Counterexample-guided Abstraction Refinement.

2003] is probably the best example of using combinations of abstract domains to enable precise and scalable verification. The software model checker BLAST combines predicate abstraction with arbitrary other abstractions specified via a lattice [Fischer et al. 2005; Beyer et al. 2007b]. This approach allows one to extend model checkers with abstractions targeted at particular domains in a modular manner. Finally, Gulwani and Tiwari [2006] shows a general framework for combining abstract interpretations for different theories, analogous to the manner in which decision procedures for different theories are combined [Nelson and Oppen 1980].

5. ABSTRACTION REFINEMENT

In general, abstract model checking is *sound*, that is, programs proved to be safe by the abstract analysis are actually safe, but *incomplete*, that is, the abstract analysis can return a counterexample even though the program is safe. In case the abstract analysis produces a counterexample, we would like to design techniques that determine whether the counterexample is *genuine*, that is, can be reproduced on the concrete program, or *spurious*, that is, does not correspond to a real computation but arises due to imprecisions in the analysis. In the latter case, we would additionally like to *automatically refine* the abstract domain, that is, construct a new abstract domain that can represent strictly more sets of concrete program states. The intent of the refined domain is to provide a more precise analysis which rules out the current counterexample and possibly others. This iterative strategy was proposed as *localization reduction* in Kurshan [1994] and Alur et al. [1995] and generalized to *counterexample-guided refinement* (CEGAR) in Ball and Rajamani [2000b], Clarke et al. [2000], and Saidi [2000]. Figure 7 formalizes this iterative refinement strategy in procedure CEGAR, which takes as input a program P , and error location \mathcal{E} and an initial, possibly trivial, abstract domain A . The procedure iteratively constructs refinements of the abstract domain A until either it suffices to prove the program safe, or the procedure finds a genuine counterexample.

5.1. Counterexamples and Refinement

The most common form of counterexample-guided refinement in software model checking has the following ingredients. The input to the counterexample analysis algorithm is a path in the control flow graph ending in the error location. The path represents a possible counterexample produced by abstract reachability analysis. The first step

resulting transition constraints to get the following formula:

$$\bigwedge_{i=0}^{n-1} \rho_i(X_i, X_{i+1}) \quad (2)$$

Notice that the trace formula is equivalent to:

- the bounded model checking formula for the unrolled version of the program corresponding to the path,
- the strongest postcondition of *true* with respect to the straight-line program corresponding to the path (when all variables other than those in X_n are existentially quantified), and
- the weakest liberal precondition of *true* with respect to the straight-line program corresponding to the path (when all variables other than those in X_0 are existentially quantified).

Thus, the trace formula is satisfiable iff the path is executable. To avoid constraints of the form $x_{i+1} = x_i$ for each x not modified by an operation, we can convert the path to *static single-assignment (SSA)* form [Flanagan and Saxe 2000; Henzinger et al. 2004], after which the renamed operations directly get translated into constraints. The trace formula can be further optimized by statically slicing out parts of the path that are not relevant to the reachability of the error states [Jhala and Majumdar 2005].

Figure 8 shows the individual constraints of the trace formula on the right side of the trace. The names LOCK_i refer to the different values of the (state) variable LOCK at different points along the trace. We have used the SSA optimization—the constraint corresponding to the incrementing of new stipulates that the incremented value new_4 is one greater than the previous value new_0 .

Syntax-Based Refinement. Suppose the trace formula is unsatisfiable. One simple way to deduce new predicates that are sufficient to rule out the current spurious counterexample is to find an *unsatisfiable core* of atomic predicates appearing in the formula, whose conjunction is inconsistent. There are several ways of finding such a set. First, one can use a greedy algorithm to find a *minimal* set of constraints that is inconsistent. This was the underlying predicate discovery mechanism in SLAM [Ball and Rajamani 2002a]. Second, one can query a proof producing decision procedure [Necula and Lee 2000] to find a proof of unsatisfiability of the constraints, and choose the atomic formulas that appear as the leaves in this proof. After finding the atomic predicates, we can simply drop the subscripts and add the resulting predicates to the tracked set, thereby refining the abstraction. This was originally implemented in BLAST [Henzinger et al. 2002] and subsequently in other tools.

Consider the trace shown in Figure 8. The trace formula, given by the conjunction of the constraints on the right side of the trace, is unsatisfiable, as it contains the conjunction of

$$\text{old}_2 = \text{new}_0, \quad \text{new}_4 = \text{new}_0 + 1, \quad \text{new}_4 = \text{old}_2$$

which is inconsistent. Thus, by dropping the subscripts, we can refine the abstraction with the new predicates

$$\text{old} = \text{new}, \quad \text{new} = \text{new} + 1, \quad \text{new} = \text{old}$$

which, after dropping redundant and inconsistent predicates, leaves just the predicate $\text{new} = \text{old}$. Notice that when this predicate is added to the set of predicates, the resulting

set, namely

$$\{\text{LOCK} = 0, \text{ LOCK} = 1, \text{ new} = \text{old}\}$$

suffices to refute the counterexample. That is, the given path is *not* a counterexample in the abstract model generated from these predicates.

Finally, another syntax-based refinement strategy is to bypass the trace formula construction, and instead, compute the sequence of predecessors

$$\varphi_n = \text{true} \quad \varphi_{i-1} = \text{Pre}(\varphi_i, \rho_i)$$

along the counterexample path, starting at the error location and going all the way back to the initial location. The abstraction can then be refined by adding the atomic predicates appearing in each φ_i . This technique was proposed by Namjoshi and Kurshan [2000]. It is used in F-SOFT in conjunction with several other heuristics such as the use of atomic predicates appearing in the proof of unsatisfiability of the trace formula.

Interpolation-Based Refinement. Though the syntax-based methods suffice to eliminate a particular counterexample, they are limited by the fact that they essentially capture relationships that are *explicit* in the program text, but can miss relationships that are *implicit*. An alternate refinement strategy, suggested in Henzinger et al. [2004], uses *Craig Interpolation* to find predicates that capture the implicit relationships that are required to verify a given safety property.

Let ϕ^- and ϕ^+ be two formulas whose conjunction is unsatisfiable. An *interpolant* ψ for (ϕ^-, ϕ^+) is a formula such that (a) ϕ^- implies ψ , (b) $\psi \wedge \phi^+$ is unsatisfiable, and (c) the free variables in ψ are a subset of the free variables that are common to ϕ^- and ϕ^+ . An interpolant always exists in case (ϕ^-, ϕ^+) are first-order formulas [Craig 1957], and an interpolant can be constructed in time linear in the size of a resolution proof for formulas in the combination theories of equality with uninterpreted functions and linear arithmetic [McMillan 2004].

Now consider an unsatisfiable trace formula (2), and for each $j \in \{0, \dots, n-1\}$, consider the *j-cut*:

$$\left(\bigwedge_{i=0}^j \rho_i(X_i, X_{i+1}), \bigwedge_{i=j+1}^{n-1} \rho_i(X_i, X_{i+1}) \right)$$

Clearly, the conjunction of the two formulas (the trace formula) is unsatisfiable. Also, the common variables between the two formulas of the cut are from X_{j+1} . Intuitively, the first part of the cut defines the set of states that can be reached by executing the prefix of the counterexample trace up to step j , and the second part defines the set of states that can execute the suffix. The common variables between the two parts are variables that are *live* across step j , that is, defined in the prefix and used in the suffix. Thus, an interpolant for the *j-cut* (a) contains states that are reached by executing the prefix, (b) cannot execute the complete suffix, and (c) contain only live variables. Thus, the interpolant serves as an abstraction at step j which is enough to rule out the counterexample. Now suppose we compute interpolants I_j for each *j-cut* which additionally satisfy an *inductiveness condition* $I_j(X) \wedge \rho_j(X, X') \rightarrow I_{j+1}(X')$. Then, adding predicates appearing in the interpolants and performing abstract reachability with these predicates is enough to rule out the counterexample. The counterexample refinement technique in Henzinger et al. [2004] computes the cut for each j and compute interpolants for each *j-cut*. By using the same proof in the construction of interpolants, the procedure additionally ensures the inductiveness condition on interpolants.

The advantage of interpolation as a technique for refinement is that it not only discovers new predicates, but also determines the control locations at which these predicates are useful. Thus, instead of keeping a global set of predicates, one can keep a map from locations to sets of predicates, and perform predicate abstraction with respect to the local set of predicates. In experiments on device drivers (reported in Henzinger et al. [2004]) this locality results in an order of magnitude improvement in the running times of the CEGAR loop.

Consider the trace formula from Figure 8. The 3-cut of the trace formula corresponds to the pair of formulas

$$\begin{aligned} \phi^- &\doteq \text{LOCK}_1 = 0 \wedge \\ &\quad \text{LOCK}_2 = 1 \wedge \text{old}_2 = \text{new}_0 \wedge \\ &\quad \text{true} \\ &\quad \text{LOCK}_4 = 0 \wedge \text{new}_4 = \text{new}_0 + 1 \\ \phi^+ &\doteq \text{new}_4 = \text{old}_2 \wedge \\ &\quad \text{LOCK}_4 = 0 \end{aligned}$$

The variables that are common to ϕ^- and ϕ^+ are new_4 , old_2 and LOCK_4 , which are the SSA renamed versions of new , old and LOCK that are live at the cut label 4. One interpolant for this cut is

$$\psi \doteq \text{new}_4 \neq \text{old}_2$$

Note that this interpolant is over the common variables, is implied by ϕ^- and is inconsistent with ϕ^+ . Indeed this interpolant captures exactly the key relationship that holds at label 4 that is needed to prove safety. The interpolant yields the predicate $\text{new} \neq \text{old}$ at label 4. Similarly, the i -cuts for $0 \leq i \leq 5$, we get the interpolants ψ_i where

$$\begin{array}{ll} \psi_0 \doteq \text{true} & \psi_1 \doteq \text{true} \\ \psi_2 \doteq \text{old}_2 = \text{new}_0 & \psi_3 \doteq \text{old}_2 = \text{new}_0 \\ \psi_4 \doteq \text{old}_2 \neq \text{new}_0 & \psi_5 \doteq \text{false} \end{array}$$

and hence, by using the predicate $\text{old} = \text{new}$ at locations 2, 3 and 4 only, we can prove the program safe.

Relative Completeness. The term *relative completeness* refers to the property that iterative counterexample refinement converges, given that there exist program invariants in a restricted language which are sufficient to prove the property of interest.

Ensuring relative completeness is not trivial. Consider the example shown in Figure 9. To verify that the error is not reachable, we must infer the invariant that $x = y$ and $x \geq 0$. Unlike the program in Figure 6, neither of these facts appears syntactically in the program. Figure 10 shows the (spurious) counterexample for the abstract model generated from the predicates $x = 0$ and $y = 0$. This counterexample corresponds to the unrolling of the loops once, and syntax based refinement strategies return the new predicates $x = 1$ and $y = 1$. In general, the refinement step can diverge, generating

```

1: x = 0;
   y = 0;
2: while (*){
3:   x++;
   y++;
   }
4: while (x>0){
5:   x--;
   y--;
   }
6: if (y != 0) error()
    
```

Fig. 9. Program.

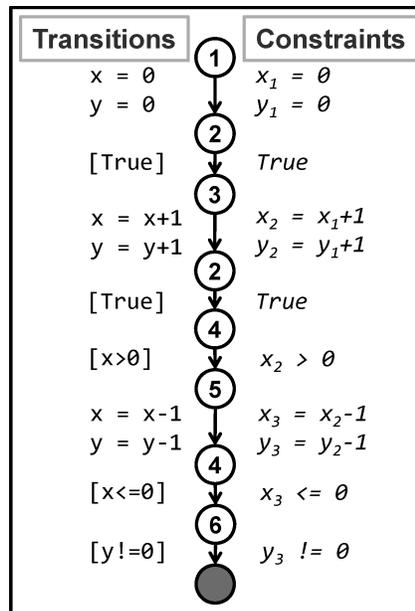


Fig. 10. Abstract Counterexample.

the sequence of predicates

$$\begin{aligned}
 &x = 0, y = 0 \\
 &x = 1, y = 1 \\
 &x = 2, y = 2 \\
 &\vdots \quad \quad \quad \vdots
 \end{aligned}$$

for counterexamples where the loops are unrolled 0, 1, 2, . . . times.

For syntactic techniques [Ball et al. 2002] provides a relative completeness result using a non-deterministic strategy to extract predicates from particular counterexamples. In practice, the strategy is implemented (e.g., in the SLAM model checker) using heuristics specific to the problem domain.

Relative completeness can be ensured for interpolation-based techniques by *restricting* the language from which predicates are drawn, [Jhala and McMillan 2006]. For example, one can stratify the language of predicates \mathcal{L} into $\mathcal{L}_0 \subset \mathcal{L}_1 \dots$ such that $\mathcal{L} = \cup_i \mathcal{L}_i$. Next, we restrict the interpolation to find predicates in \mathcal{L}_i iff there are no predicates in \mathcal{L}_{i-1} that refute the counterexample. In this way, the iterative refinement becomes relatively complete in the sense that if there is some safety proof where the invariants are drawn from \mathcal{L} , then there is some \mathcal{L}_j from which the invariants are drawn, and the restricted iterative refinement will be guaranteed to terminate by finding predicates from \mathcal{L}_j . One way to stratify the language is to let \mathcal{L}_i be the language of predicates where the magnitude of each constant is less than i . In essence, this stratification biases the iterative refinement loop to find predicates involving “small constants”. Jhala and McMillan [2006] shows how to structure decision procedures to force them to produce interpolants from a restricted language. Rybalchenko and Sofronie-Stokkermans [2007] shows a more general technique for computing interpolants that satisfy different kinds of conditions, such as interpolants in which specific variables, even if shared, do not occur.

In the example of Figure 9, the restricted interpolation method directly finds the predicates $x = y$ and $x \geq 0$ which belong in \mathcal{L}_0 , the language of predicates with constants less than 0. This simple restriction goes a long way towards making CEGAR complete.

Syntactic techniques can be shown complete also for certain classes of systems, for example, timed systems and other systems with a finite bisimulation quotient [Henzinger et al. 2002], and broadcast protocols and other *well-structured systems* [Dimitrova and Podelski 2008].

Refining Multiple Paths. Several augmentations to the above refinement scheme have been suggested. First, instead of refining one path at a time, the procedure can be called with a set of abstract counterexamples which are refined together while optimizing the set of inferred predicates [Chaki et al. 2003]. In Beyer et al. [2007b], counterexample analysis is performed on a *path program*, the least syntactically valid subprogram containing the counterexample. A path program represents a (possibly infinite) family of possible counterexamples. The goal of refining path programs (through *path invariants*) is to find suitable program invariants that simultaneously rule out the entire family. Unfortunately, since a path program can contain loops, the simple partitioning and interpolation technique from above is not immediately applicable. In Beyer et al. [2007b], counterexample refinement is performed by inferring path invariants using constraint-based invariant synthesis.

Refining Other Domains. While we have focused on refining predicate abstractions by adding new predicates, the idea of counterexample refinement can be used for other abstract domains. For example, Jhala and McMillan [2005] shows how *symmetric interpolation* can be used to refine (propositional) transition relations in a counterexample guided manner, Gulavani et al. [2008] shows how a polyhedral domain can be refined using counterexample analysis. For control abstractions in thread-modular reasoning, Henzinger et al. [2004] show how environment assumptions for concurrent programs can be automatically constructed and refined.

Refinement techniques have been generalized to domains beyond hardware and software, for example, in the verification of real-time and hybrid systems [Alur et al. 2006; Jha et al. 2007].

5.2. Abstraction-Refinement-Based Model Checkers

Slam. The SLAM model checker [Ball and Rajamani 2002b] was the first implementation of CEGAR for C programs. It introduced *Boolean programs*—imperative programs

where each variable is Boolean—as an intermediate language to represent program abstractions. A tool (called *c2bp*) implemented predicate abstraction for C programs [Ball et al. 2001]. The input to *c2bp* is a C program and a set of predicates, and the output is a Boolean program, where each Boolean variable corresponds to a predicate, and the assignments and conditionals on the variables correspond to the Cartesian predicate abstraction of the C program. A tool called *BEBOP* then implemented a symbolic model checker for (recursive) Boolean programs [Ball and Rajamani 2000a]. Finally, abstraction refinement was performed by *newton*, which implemented a greedy heuristic to infer new predicates from the trace formula. *SLAM* was used successfully within Microsoft for device driver verification [Ball et al. 2006] and developed into a commercial product (Static Driver Verifier, *SDV*).

The *SLAM* project introduced several key ideas in software model checking, including the generalization of predicate abstraction in the presence of pointers and dynamically allocated memory [Ball et al. 2001], modular predicate abstraction [Ball et al. 2005], and BDD-based model checking in the presence of procedure calls [Ball and Rajamani 2000a]. *SLAM* inspired a resurgence of interest in the verification of software implementations and a suite of tools geared to program verification.

Blast. The *BLAST* model checker [Beyer et al. 2007c] implements an optimization of CEGAR called *lazy abstraction*. The core idea of *BLAST* is the observation that the computationally intensive steps of abstraction and refinement can be optimized by a tighter integration which would enable the reuse of work performed in one iteration in subsequent iterations. *Lazy abstraction* [Henzinger et al. 2002] tightly couples abstraction and refinement by constructing the abstract model *on-the-fly*, and locally refining the model *on-demand*. The former eliminates an often wasteful and expensive model-construction phase and instead, performs abstraction on the *reachable* part of the state space. This is achieved by lazily building an *abstract reachability tree* whose nodes are labeled by abstract regions. The regions at different nodes of the tree, and hence, at different parts of the state space, can be over different sets of predicates. To locally refine the search when a counterexample is found, *BLAST* finds the “pivot” node from which the remainder of the counterexample is infeasible and rebuilds the subtree from the pivot node onwards. This ensures that parts of the state-space known to be free of errors, namely different subtrees, are not re-analyzed. Further, this permits the use of different predicates at different program points which drastically reduces the size of the abstract state space that must be analyzed. Upon termination with the outcome “program correct,” the proof is not an abstract model on a global set of predicates, but an abstract model whose predicates change from state to state. Thus, by always maintaining the minimal necessary information to validate or invalidate the property, *lazy abstraction* scales to large systems without sacrificing precision. *BLAST* combines *lazy abstraction* with procedure summaries [Sharir and Pnueli 1981; Reps et al. 1995] and scoped interpolation [Henzinger et al. 2004] to model check recursive programs. Finally, *BLAST* can be extended with arbitrary lattices [Fischer et al. 2005], thus yielding a general mechanism for refining any dataflow analysis with iterative predicate abstraction.

Magic. The *MAGIC* model checker was designed to enable the modular verification of concurrent, message passing C programs. *MAGIC* allows the user to specify arbitrary nondeterministic labeled transition systems (LTS) as specifications, and verifies that the set of traces (over messages and events) generated by a concurrent C program is contained in the language generated by the LTS. *MAGIC* implements a *two-level* abstraction strategy to combat the state explosion that arises from the product of multiple threads. First, an (eager) predicate abstraction is carried out for each individual thread, yielding a finite state machine representing the thread’s behavior. Second, an *action-guided*

abstraction is carried out to minimize each thread’s state machine while preserving the sequences of messages and events generated by the state machine. Finally, the product of the reduced state machines is computed and model checked. The entire procedure is wrapped inside a CEGAR loop, that uses spurious counterexamples to infer new predicates which yield refined (reduced) state machines. Finally, MAGIC also implements several methods to minimize the number of predicates, by finding predicates that simultaneously refute multiple paths [Chaki et al. 2003].

F-Soft. The F-SOFT model checker [Ivancic et al. 2005] combines CEGAR-based predicate abstraction refinement with several other abstract domains that efficiently yield the kinds of invariants needed to check standard runtime errors in C programs (e.g., buffer overflows, null dereferences). For these errors, a CEGAR based predicate abstraction can eventually find the right domain, but a multi-stage framework that eagerly combines numerical domains with CEGAR can be much more efficient. The numerical domains can altogether eliminate some easy checks, or yield invariants that help the subsequent iterative analysis converge more quickly [Jain et al. 2006]. To this end F-SOFT implements symbolic model checking algorithms that combine BDDs and polyhedra [Yang et al. 2006], and techniques that combine widening with iterative refinement [Wang et al. 2007].

Other Tools. The IMPACT model checker [McMillan 2006] implements an algorithm that entirely eliminates all abstract (post) image computations. Instead IMPACT builds an abstract reachability tree by directly using the sequence of interpolants generated from the trace formulas to strengthen the regions labeling the nodes of the tree. This process is repeated using the lazy abstraction paradigm until the program is proved safe, or a counterexample is found. McMillan [2006] shows that by directly using interpolation within the lazy abstraction framework, one can achieve very dramatic improvements in performance.

The ARMC model checker [Podelski and Rybalchenko 2007] implements the CEGAR scheme using a constraint-based logic programming language, resulting in an elegant implementation. The solver can generate interpolants for linear arithmetic constraints in combination with free function symbols [Beyer et al. 2007a]. The full-fledged linear arithmetic constraint solver implemented in the constraint-based logic programming language allows ARMC to handle programs with intensive operations on numerical data, which is needed, for example, for checking real-time bounds for embedded systems [Meyer et al. 2006].

6. PROCEDURAL ABSTRACTION

So far, we have considered programs without procedures. We now extend the class of simple programs `imp` to the class `imp+proc` of programs with potentially recursive procedure calls.

6.1. Programs with Procedures

A *procedural imperative program* P is a tuple (F, f_0) where F is a finite set of procedures, and f_0 is a special *initial* procedure in F from which execution begins. A *procedure* f is a simple imperative program $(\ell^f, X^f, \ell_0^f, T^f)$. Each procedure f has a unique *input parameter* x_0^f in X^f . The states of a procedure f are the elements of $v.X^f$. The set of transitions comprises:

—**Intra-** procedural transitions of the form $(\ell, \text{Intra } \rho, \ell')$,

- Call** transitions of the form $(\ell, \text{Call } x := f(e), \ell')$, and,
- Return** transitions of the form $(\ell, \text{Ret return } e, \ell')$.

For each kind of transition, the source and target locations (i.e., ℓ and ℓ' respectively), belong to the same procedure. For simplicity, we make two assumptions. First, the sets of variables of each procedure are distinct—that is, each procedure has its own *local* variables and there are no global variables. Second, the formal parameters of each procedure are *not modified* by the transitions of the procedure. We write ℓ to abbreviate $\cup_f \ell^f$, X to abbreviate $\cup_f X^f$, T to abbreviate $\cup_f T^f$.

6.2. InterProcedural Reachability

Even if all the variables X range over finite domains, the state space of a procedural program is infinite, as the stack can grow unboundedly due to recursive calls. As a result, we cannot use direct graph reachability (as in Algorithm EnumerativeReachability) to verify safety properties of programs with procedures. The key to analyzing procedural programs lies in designing an algorithm that uses the following two observations that avoid the need for tracking the control stack explicitly at each state. First, the behaviors of the entire program can be reconstituted from the *input-output* behaviors of the individual procedures. Second, even though the number of configurations of the whole program are infinite, each procedure has a *finite* number of input-output behaviors.

To use these observations, we extend the standard reachability procedure with a form of *memoization* thus equipping it to compute input-output *summaries* for each procedure. If all the program variables have a finite domain, then the size of the summaries—which are simply sets of pairs of input parameter values and output expression values—is finite. Consequently, the extended procedure can compute the summaries as a least fixpoint in finite time, thereby yielding an algorithm for model checking programs with procedures. Thus, to compute summaries, we extend the reachability procedure from from Figure 2 with the following data structures.

- (1) An *input* state for a procedure f is a valuation where all the variables except the formal parameter x_0^f are set to 0. The *input of a state* $s \in v.X^f$, written $\text{Init}(s)$, is the input state for f where the formal x_0^f has the value $s(x_0^f)$. As the formal parameters x_0^f of each procedure remain unmodified, each state has encoded within it the input with which the corresponding procedure was called.
- (2) The *callers* of an input state s of f , written $\text{Callers}[s]$, correspond to tuples $(s_c, \text{Call } op_c, \ell'_c)$ such that there is a reachable configuration (ℓ_c, s_c) and a call transition (ℓ_c, op_c, ℓ'_c) that, when executed from state s_c , yields the input state s for f .
- (3) The *exits* of an input state s of f , written $\text{Exits}[s]$, correspond to pairs (s_e, op_e) such that from the input configuration (ℓ_0^f, s) some configuration (ℓ, s_e) is reachable, and there is a return transition $(\ell, \text{return } e, \cdot)$ in f .

Intuitively, the callers and exits allow us to build a memoized version of the enumerative reachability algorithm. Whenever a call transition is reached, the known exits for the corresponding input state are propagated at the caller. Whenever a return transition is reached, the return value is propagated to all the known callers of the corresponding input state.

This intuition is formalized in Figure 11 which shows the interprocedural reachability algorithm that extends Algorithm EnumerativeReachability to handle procedure calls and returns. InterProceduralReachability is a worklist-based algorithm that

Algorithm: Interprocedural Enumerative Reachability
Input simple program $P = (X, L, T, \ell_0, F, \mathbf{f}0)$, error location $\mathcal{E} \in L$
Output SAFE if P is safe w.r.t. \mathcal{E} , UNSAFE otherwise

```

def Propagate( $s_c, s_e, \mathbf{x} := \mathbf{f}(\cdot)$ , return  $\mathbf{e}, \ell'_c$ ):
  add ( $\ell'_c, Post(s_c, (\mathbf{x} \leftarrow \mathbf{e}(s_e)))$ ) to worklist

def PropagateCall( $s_c, \cdot := \mathbf{f}(\cdot)$  as  $op_c, \ell'_c$ ):
   $s' = Post(s_0^f, (\mathbf{x}_0^f \leftarrow \mathbf{e}(s_c)))$ 
  add ( $\ell'_c, s'$ ) to worklist
  add ( $s_c, op_c, \ell'_c$ ) to Callers[Init( $s'$ )]
  foreach ( $s_e, op_e$ ) in Exits[Init( $s'$ )] do:
    Propagate( $s_c, s_e, op_c, op_e, \ell'_c$ )

def PropagateReturn( $s_e, op_e$ ):
  add ( $s_e, op_e$ ) to Exits[Init( $s_e$ )]
  foreach ( $s_c, op_c, \ell'_c$ ) in Callers[Init( $s_e$ )] do:
    Propagate( $s_c, s_e, op_c, op_e, \ell'_c$ )

def InterProceduralReachability( $P, \mathcal{E}$ ):
  reach =  $\emptyset$ 
  worklist =  $\{(\ell_0^f, s) \mid s \in v.X^{f0}\}$ 
  while worklist  $\neq \emptyset$  do:
    pick and remove( $\ell, s$ ) from worklist
    if ( $\ell, s$ )  $\notin$  reach:
      add ( $\ell, s$ ) to reach
      foreach ( $\ell, \text{Intra } \rho, \ell'$ ) in  $T$  do:
        add  $\{(\ell', s') \mid s' \in Post(s, \rho)\}$  to worklist
      foreach ( $\ell, \text{Call } op, \ell'$ ) in  $T$  do:
        PropagateCall( $s, op, \ell'$ )
      foreach ( $\ell, \text{Ret } op, \ell'$ ) in  $T$  do:
        PropagateReturn( $s, op, \ell'$ )
    if exists ( $\mathcal{E}, s$ )  $\in$  reach:
      return UNSAFE
    else:
      return SAFE

```

Fig. 11. Interprocedural Enumerative Reachability.

simultaneously computes the reachable input states for each procedure, the callers of each input state and the exits of each input state. For each procedure, the least fixpoint set of input states corresponds to all the (finitely many) reachable inputs of the procedure, and the respective callers and exits correspond to the possible calling states and return values for that input state.

The algorithm maintains a worklist of configurations, namely pairs of locations and states. At each iteration, a configuration is picked off the worklist. If the configuration has previously been reached, the algorithm moves to the next worklist element. If the configuration is new, its successors are added to the worklist as follows.

- For each enabled intraprocedural transition for the configuration, the successors are computed using *Post*, and added to the worklist, similar to Algorithm EnumerativeReachability.
- For each enabled call transition of the form $\mathbf{x} := \mathbf{f}(\mathbf{e})$ the algorithm: (1) computes the input state for the call, by calling *post* with the transition that assigns the formal \mathbf{x}_0^f the actual parameter \mathbf{e} , (2) adds the calling state to the set of known callers for the input state, and (3) propagates each known *exit* for the input state to the caller, by

adding to the worklist the successors corresponding to assigning the return values at the exit to x .

- For each enabled return transition of the form `return e` the algorithm: (1) computes the input state for the exit state from which the return occurs, (2) adds the exit state to the set of known exits for the input state, and (3) propagates the return value at the exit state to each of the known callers of the input state.

We can prove that this algorithm computes all the reachable configurations by induction on the length of the shortest trace leading to the configuration. We can prove that the algorithm terminates by appealing to the fact that the sets being populated, namely the reach set, callers and exits are all finite.

Graph-Based Algorithms. Algorithm `InterProceduralReachability` is a reformulation of the classical tabulation-based approach to interprocedural dataflow analysis presented in Sharir and Pnueli [1981]. Reps et al. [1995] showed how the technique could be phrased as a special graph-reachability algorithm, leading to an algorithm with running time cubic in the number of reachable configurations. By using bitvectors to represent sets of configurations and bitvector operations to manipulate the sets, one can obtain a subcubic running time [Chaudhuri 2008]. The idea of performing reachability on a graph while ensuring the path defines a word in a context free language has also been used in the context of Datalog query evaluation [Yannakakis 1990], and the approach ultimately uses a dynamic programming algorithm for parsing context free languages.

Symbolic Algorithms. As in the case of (simple) explicit state model checking, the above enumerative algorithms run in time that is polynomial in the number of reachable configurations. However, assuming that each variable can take on at least two values, the number of reachable configurations is *exponential* in the number of local variables of each procedure, and hence, exponential in the size of the program's representation. However, the technique of summarization can be used to lift symbolic algorithms, which work with compact representations of *sets* of states, to the interprocedural setting. For example, to obtain a symbolic version of `InterProceduralReachability`, we need only to:

- (1) view each s as a set of states,
- (2) view the membership in the reach set as inclusion in the reach set,
- (3) view each *Post* operation as a symbolic operation, as in Section 3.

Ball and Rajamani [2000a] describes `BEBOP`, the first BDD-based interprocedural safety model checker that takes as input a *Boolean program*, that is a procedural program where each variable is Boolean valued, and a distinguished error location and determines whether the error location is reachable in the Boolean program. In particular, this article shows how procedure summaries can be symbolically implemented using BDDs and show how to efficiently generate *counterexample* traces from the runs of the checker. Finally, *weighted pushdown systems* [Reps et al. 2005] are a general analysis framework that encompasses interprocedural reachability of Boolean programs as well as a rich class of dataflow analysis problems.

An alternate class of algorithms for the analysis of procedural programs comes from the symbolic analysis of *pushdown processes* [Bouajjani et al. 1994; Walukiewicz 1996]. These algorithms use the language-theoretic observation that the set of stack configurations reachable at any location of a pushdown automaton is a regular language, and use symbolic representations based on automata to represent sets of stacks.

Abstraction. So far, we have considered programs where the variables take on finitely many values. We can use abstraction to extend the finite-state approach to

programs with variables over infinite domains. As for programs in `imp`, there are two approaches. The first is to *eagerly* abstract the program into a finite state procedural program that overapproximates the behaviors of the original program. This method was popularized by the SLAM model checker. The second is to *lazily* carry out the abstraction *during* the model checking. For example, to obtain an abstract version of `InterProceduralReachability`, we need only to:

- (1) view each s as an abstract set of states,
- (2) view the membership in the reach set as inclusion in the abstract reach set,
- (3) view each `Post` operation as an overapproximate $Post^\#$ operation, as in Section 4.

This approach is adopted by BLAST, which uses predicate abstraction to lazily build abstract summaries.

In either case, to ensure soundness, the abstract states constructed when analyzing a procedure f must describe valuations of *only* those variables that are in scope inside procedure f . For example, when predicate abstraction is used eagerly or lazily, to model check procedural programs, we must ensure that all the predicates used to abstract procedure f are *well-scoped*, that is, contain variables that are in-scope at f .

The efficiency of the model checking is greatly enhanced by using abstractions that yield summaries that can be applied at multiple callsites, instead of individual summaries that can only be applied at specific call-sites. For example, it is typically more efficient to use *relational abstractions* [Cousot and Cousot 1979] that can describe the outputs in terms of the inputs, instead of pairs of input-output states, where each pair corresponds to a different callsite. For example, a relational summary that specifies that the output of a procedure is one greater than the input, is more compact and reusable than a tabulational summary that specifies that if the input is 0 (respectively, 1, 2), the output is 1 (respectively, 2, 3).

Ball et al. [2005] shows how relational summaries can be computed by using predicate abstraction over predicates containing *symbolic constants*, which are immutable values representing the input parameters of different functions. Algorithm `InterProceduralReachability` exploits this idea by requiring that the formals remain unmodified. As a result each concrete output state (e.g., $[x_0^f \mapsto 0, \text{ret} \mapsto 1]$) implicitly encodes the input state that generated it, and hence a relational abstraction of the output state (e.g., $\text{ret} = x_0^f + 1$) describes an input-output behavior of the function. Henzinger et al. [2004] shows how Craig Interpolation can be used to address the problem of *automatically discovering* well-scoped predicates with symbolic constants that are relevant to verifying a given safety property.

Top-Down vs. Bottom-Up. The algorithm described above computes the set of reachable inputs for each procedure in a *top-down* manner, that is, using the callsites. Another approach is to aggressively compute the most general behavior for each procedure, by seeding the worklist with all the possible input states for each procedure. This *bottom-up* approach suffers from the disadvantage of computing possibly useless information. On the other hand, it enjoys several engineering advantages. A *procedure call graph* has vertices corresponding to procedures and which has directed edges from f_1 to f_2 if there is a call transition in f_1 to procedure f_2 . Suppose that there is no recursion, and so the procedure call graph is acyclic. In this setting, the bottom-up approach yields two benefits. First, we can analyze each procedure in isolation, processing the leaf procedures of the call-graph first and moving up the graph. This can drastically lower the memory needed to analyze large programs. Second, as procedures can be analyzed in isolation, we can parallelize the analysis as we can concurrently analyze two procedures each of whose callees has been previously analyzed. Bottom-up analyses typically use

relational abstractions, as when analyzing a procedure f , no information is available about the states from which f can be called.

Saturn is a bottom-up interprocedural model checker that exploits the engineering advantages outlined above to scale to the entire Linux kernel [Xie and Aiken 2005]. SATURN computes input-output summaries over a fixed set of predicates as follows. Each input or output configuration is a conjunction of the predicates. SATURN unfolds the transition relation for the body of the procedure (plugging in the summaries for called procedures), and then uses SAT-based Bounded Model Checking to determine which input-output configurations are feasible. The summary computed for the procedure is the *disjunction* of all the satisfiable configurations.

Houdini embodies a dual approach for inferring pre- and post-conditions for procedures [Flanagan et al. 2001]. Intuitively, the input-output behavior of each procedure can be summarized by the pair of the conjunction of all precondition clauses, and, the conjunction of all postcondition clauses. HOUDINI generates pre- and post-conditions from a candidate set of predicates as follows. It begins with summaries corresponding to the conjunction of *all* predicates, and then uses SMT-based Bounded Model Checking (ESC/JAVA [Flanagan et al. 2002]) to iteratively drop the precondition (respectively, postcondition) predicates that do not provably hold at some call (respectively, exit) location. It can be shown that the above converges to a least fixpoint, where each summary's precondition (respectively, postcondition) captures the *strongest* overapproximation of the input states (respectively, exit states) for the procedure, expressible as a conjunction of the candidate predicates.

6.3. Concurrency and Recursion

For concurrent programs where each sequential thread belongs to the class `imp+proc`, the reachability problem is undecidable when synchronizations between threads is taken into account [Ramalingam 2000]. The intuition behind this result is that the executions of procedural programs are isomorphic to context-free languages (CFL) over the alphabet of local (e.g., call) and synchronization (e.g., lock or rendezvous) actions. Thus, a configuration (c_1, c_2) of a concurrent program with two threads can be reached iff the intersection $L_1 \cap L_2$ of two context free languages over the alphabet of synchronization actions is non-empty. Formally, for any two context free languages L_1 and L_2 , one can build a program $P_1 | P_2$ synchronizing over the common alphabet of L_1 and L_2 and a configuration (c_1, c_2) of $P_1 | P_2$ such that the configuration (c_1, c_2) is reachable iff $L_1 \cap L_2$ is non-empty. As the emptiness of CFL intersection is undecidable, it follows that the reachability problem for concurrent procedural programs is undecidable.

Several authors have proposed techniques to over-approximate the set of reachable configurations. These techniques relax the synchronization sequences possible in the original program, for example, by ignoring the ordering of synchronization operations [Bouajjani et al. 2003], or by approximating the context free language by a regular one [Chaki et al. 2006]. Alternatively, one use the notion of *transactions*, that is, sequences of actions of a thread that execute atomically with respect to other threads, to design a sound but incomplete summarization-based model checking algorithm for concurrent procedural programs [Qadeer et al. 2004]. In special cases, such as concurrent threads synchronizing solely using nested locks, the reachability problem is, somewhat surprisingly, decidable [Kahlon and Gupta 2007]. The above paper also gives model checking algorithms for various subclasses of concurrent recursive programs (modeled as interacting pushdown automata) with respect to specifications in fragments of temporal logic.

The reachability problem for concurrent procedural programs becomes decidable if one bounds the number of allowed *context switches*, that is, the number points along

a computation where control passes from one thread to another [Qadeer and Rehof 2005]. Recently, symbolic implementations have been suggested for context bounded reachability [Suwimonteerabuth et al. 2008]. Furthermore, [Lal et al. 2008] shows that the machinery of weighted pushdown systems can be used to generalize the context-bounding approach to a large class of flow-analyses for concurrent programs. Note that by bounding the number of context switches corresponds to an *underapproximation* of the reachable states; in particular, such an analysis can miss unsafe computations that require more than the given number of context switches to reach the error state. However, in practice, context bounding provides a useful underapproximation to reachability. Empirical results suggest that a low value of the context bound suffices to discover subtle bugs [Qadeer and Wu 2004; Musuvathi and Qadeer 2007].

7. HEAP DATA STRUCTURES

So far, we have assumed a simple program model where we have ignored the effect of potentially unbounded data structures on the heap. These, however, represent one of the biggest challenges to scalable and precise software verification. The difficulty of automatically reasoning about data structures stems from the need to reason about relationships between the unbounded number of values comprising the structure. Thus, the verification tool requires some way to generalize relationships over specific values into quantified facts that hold over the structure, and dually, to instantiate quantified facts to obtain relationships over particular values.

The class of `imp+heap` extends the class of `imp` by having a global and unbounded *memory* array, and operations *read* and *write* to read the memory at an index or to write a value (which could be another index) to an index. The reads and writes on the array are governed by McCarthy's axioms:

$$\text{read}(\text{write}(\text{memory}, \text{index}, \text{value}), \text{index}') = \text{read}(\text{memory}, \text{index}')$$

if $\text{index} \neq \text{index}'$ and

$$\text{read}(\text{write}(\text{memory}, \text{index}, \text{value}), \text{index}') = \text{value}$$

if $\text{index} = \text{index}'$. The presence of the heap complicates analysis, since two syntactically distinct expressions e_1 and e_2 could refer to the same memory location, and hence updating the memory by writing to location e_1 require updating information about the contents of the syntactically different location e_2 .

We now briefly discuss some of the approaches that have been proposed. Our coverage of this topic is deliberately brief as the literature on shape analysis merits its own survey, and as the scalable verification of heap-based data structures is perhaps the least understood direction in software model checking.

Alias Analysis determines whether two pointers refer to the same heap cell [Muchnick 1997]. There is a wide variety of alias analyses that span the precision-scalability spectrum [Andersen 1994; Steensgard 1996; Hind 2001; Whaley and Lam 2004; Hardekopf and Lin 2007]. Software model checkers like SLAM, BLAST, and F-SOFT use a precomputed alias analysis to determine whether an assignment $*x := e$ can affect the value of $*y$. The SATURN verifier uses a combination of predicate abstraction, bounded model checking and procedure summarization to compute a precise path- and context-sensitive pointer analysis [Hackett and Aiken 2006] at the same time as the rest of the analysis. However, alias analysis is only useful for reasoning about explicitly named heap cells, but not unbounded data structures. The common assumption in alias analyses is to abstract every memory cell allocated dynamically from the same program

location into one abstract cell. This means, for example, that subsequent algorithms that depend on the alias analysis cannot distinguish between different cells of a data structure (e.g., a linked list) whose cells are allocated from a single program point. This leads to imprecision.

Shape Analysis attempts to characterize collections of heap cells reachable from particular pointers, for example, to determine whether the cells form a list or a tree and so on [Chase et al. 1990; Ghiya and Hendren 1996], by allowing finer distinctions between heap cells. Early shape analyses used dataflow analyses over specialized lattices designed to capture properties of particular structures (e.g., to check whether the heap was organized as a tree or as a cycle). Sagiv et al. [2002] show how three-valued logic could be used as a foundation for a parameterized framework for designing shape analyses. The framework is instantiated by supplying predicates that capture different relationships between cells (e.g., that one cell points to another, one cell is reachable from another), and by supplying the functions that determine how the predicates are updated by particular assignments. The tool TVLA [Lev-Ami and Sagiv 2000] implements these ideas and has been used to verify non-trivial data structure properties. Similar ideas were used to build a model checker capable of verifying concurrent, heap-manipulating programs [Yahav 2001]. The three-valued predicates used in shape analysis can be combined with a numerical domain to verify properties of array manipulating programs [Gopan et al. 2005].

Separation Logic was designed to enable modular reasoning about heap-manipulating programs [Reynolds 2002]. Separation logic extends classical Hoare logic [Hoare 1969] with two operators; *separating conjunction* (written as $*$) and *separating implication* (written as $-*$), which are used to construct assertions over *disjoint* parts of the heap. For example, an assertion of the form $A*B$ says that there is one set of heap cells that satisfy A and a disjoint set of cells that satisfy B . The key advantage of this logic is that it allows one to succinctly specify which parts of the heap are touched by a given piece of code, and allows one to compose properties of sub-heaps to get properties of the entire heap. While the logic was originally designed as a calculus for manually verifying low-level pointer manipulating programs, it has subsequently become the basis for several abstract interpretation based verifiers. To do so, the analysis designer specifies: (1) recursively defined predicates over the separating operators that represent shape properties, and, (2) a mechanism to generalize (i.e., fold) and instantiate (i.e., unfold) such predicates [Distefano et al. 2006; Magill et al. 2007; Yang et al. 2008]. A variant of this approach is to extract the abstract domain from programmer-specified (recursive) checker definitions [Chang and Rival 2008], and shows how the resulting shape domain can be combined with other abstractions to synthesize richer invariants like sortedness.

Reachability Predicates and logics built around them [Nelson 1983] provide another way to reason about heaps. The challenge is to design a *decidable* logic that is expressive enough to capture interesting invariants. Often such logics use a transitive closure operator to express reachability predicates [Immerman et al. 2004]. Among the most expressive of such logics is weak monadic second-order logic (MSOL). MSOL can express transitive closure and is decidable over trees [Doner 1965]. The PALE tool is built on top of MSOL and enables reasoning about tree-like data structures [Møller and Schwartzbach 2001]. Similarly, the HAVOC tool uses an expressive and decidable logic for reasoning about reachability in lists [Lahiri and Qadeer 2008]. Both of these tools give completeness guarantees but require the user to provide loop invariants.

Quantified Loop Invariants are often required to prove properties of unbounded data structures. Techniques for automatically inferring quantified loop invariants are suitable for shape analysis if the underlying logics supports reasoning about reachability

predicates. *Indexed* predicate abstraction [Lahiri and Bryant 2004] and *Boolean heaps* [Podelski and Wies 2005] generalize the predicate abstraction domain such that it enables the inference of universally quantified invariants. In particular, Boolean heaps have been used to infer quantified invariants for linked lists and trees. These methods consider predicates that range over heap cells similar to the predicates used in three-valued shape analysis. Gulwani et al. [2008] show how to combine different abstract domains to obtain universally quantified domains that can capture properties of linked lists. Craig interpolation has also been used to find universally quantified invariants for linked lists [McMillan 2008].

8. LIVENESS AND TERMINATION

Next, we turn from safety properties that specify that nothing bad happens, to *liveness* properties that state, informally, that something good eventually happens.

8.1. Finite State

For finite state programs, and liveness properties specified in a temporal logic such as LTL [Pnueli 1977], there is an automata-theoretic algorithm to check if the program satisfies the temporal logic property [Vardi and Wolper 1986]. Briefly, the algorithm constructs a *Büchi automaton* from the negation of the LTL property, and checks that the intersection of language of program behaviors and the language of the Büchi automaton is empty [Vardi and Wolper 1986; Vardi 1995]. Emptiness of the intersection can be checked by performing a nested depth-first search, looking for accepting cycles in the automaton [Courcoubetis et al. 1992]. This algorithm is implemented in the Spin model checker. A symbolic version of checking Büchi automaton emptiness was given in Emerson and Lei [1986], and is implemented in SMV using BDD operations.

To verify arbitrary LTL properties of procedural programs, we need to track the contents of the control stack. Bouajjani et al. [1994] shows how to precisely model check linear and branching time properties of *pushdown systems* by using automata to symbolically represent sets of stack configurations. Esparza and Schwoon [2001] describes MOPEd, which combines BDD-based symbolic representation for data, that is, program variables, with automata-based representation for stacks, in order to obtain an LTL model checking algorithm for Boolean programs.

8.2. Infinite State

We now move to checking liveness properties for infinite state systems. We focus on *program termination*, a particular liveness property that stipulates that a program has no infinite computations. Formally, P is terminating if every computation $\langle \ell_0, s_0 \rangle \rightarrow \dots \langle \ell_k, s_k \rangle$ reaches some state $\langle \ell_k, s_k \rangle$, which has no successor.

For many systems, termination can be proved only under certain assumptions about the nondeterministic choices made during program execution. The programmer often models certain aspects of the system through non-deterministic choice, with an implicit assumption that such choices are resolved in a “fair” manner. For example, one can model a scheduler as non-deterministically providing a resource to one or other process, with the assumption that both processes are picked infinitely often. Similarly, one can model asynchrony by modeling non-deterministic “stutter” steps, together with the assumption that the process makes progress infinitely often. The standard way to rule out certain undesirable infinite behaviors from the scope of verification is through *fairness conditions* [Francez 1986]. Typically, a fairness condition can be translated to an automaton on infinite words [Vardi 1995]. *Fair termination* is the property that a program terminates on all runs that satisfy the fairness requirements.

Just as safety properties can be reduced to reachability problems, liveness properties can be reduced to checking termination under fairness requirements [Vardi 1991]. The techniques for proving termination generalize to fair termination by taking a product of the program with an automaton on infinite words modeling the fairness condition, and checking for well-foundedness only for final states of the automaton. For this reason, we shall concentrate in the rest of the section on techniques to prove program termination.

Proofs for program termination ultimately rely on well-foundedness and ranking functions. A relation $R \subseteq A \times A$ is *well-founded* if there is no infinite sequence a_0, a_1, \dots such that for each $i \geq 0$ we have $a_i R a_{i+1}$. For example, the usual $<$ relation on the natural numbers is well-founded, but the $<$ relation on the integers is not. Let A be a set and $R \subseteq A \times A$ a well-founded relation on A . A *ranking function* r is a mapping that associates a *rank* from A with each program states, such that for any state s and any successor s' of s , we have $R(r(s), r(s'))$. A program terminates if there is a ranking function from the reachable program states to some set A and well-founded relation R on A [Lehmann et al. 1982]. Intuitively, a program terminates if each reachable state of the program is associated with a *rank*, the rank decreases with every transition, and there is no infinitely decreasing sequence of ranks. It is crucial to restrict attention to the *reachable* states: the transition relation of a terminating program may not by itself be well-founded, for example due to the presence of unreachable non-terminating loops.

In general, checking program termination is undecidable, and so there is no algorithmic technique to compute ranking function. For certain specialized classes of programs ranking functions can be found algorithmically [Colón and Sipma 2001, 2002; Tiwari 2004; Podelski and Rybalchenko 2004a; Bradley et al. 2005], leading to an algorithmic technique to prove termination for these classes.

Furthermore, we cannot compute and reason about the exact set of reachable transitions of an infinite-state program. Modern termination checking tools reduce termination verification to checking the well-foundedness of a *transition invariant*, which is an *overapproximation* of the transitive closure of the transition relation restricted to the reachable states. By explicitly adding program variables to store “previous” values of the program variables, one can reason about transitions in the original program through invariants in the expanded program. Consequently, invariant generation technology developed for safety verification carries over directly to proving termination.

In general, it can be hard to find a suitable single ranking function. In these cases, it is preferable to *compose* a termination argument out of simpler termination arguments. One way to do this is through an idea from Podelski and Rybalchenko [2004b]—the notion of disjunctive well-foundedness. A relation T is *disjunctively well-founded* if it is a finite union $T = T_1 \cup \dots \cup T_k$ of well-founded relations. Every well-founded relation is (trivially) disjunctively well-founded, but not conversely. However, the paper shows that a relation R is well-founded if the transitive closure R^+ of R is contained in some disjunctively well-founded relation T . The relation T can be found as the union of transition invariants computed for parts of the program. In addition, techniques based on transition invariants can be applied to directly reason about fairness [Pnueli et al. 2005]. These observations form the basis of ARMC and TERMINATOR, which prove termination for C programs [Cook et al. 2006].

8.3. Nontermination

In theory, every terminating program has a transition relation whose well-foundedness is witnessed by some rank function. However, in practice, this rank function may not be easy to find, and the inability to find a proof of termination (i.e., a transition invariant

and a well-founded ranking) does not immediately signify that the program has an infinite execution. Thus, for programs for which we cannot prove termination, we require specialized methods that can prove *nontermination* and hence, prove that there is a genuine “counterexample” that refutes the liveness property. Nontermination has received relatively lesser attention in program verification. One way is through finding *recurrence sets* [Gupta et al. 2008; Velroyen and Rümmer 2008]. A set of states R is recurrent if for every state $s \in R$, there is some successor s' of s such that $s' \in R$. It can be shown that a transition relation is not well-founded iff there is some recurrent set, and a program is non-terminating if one can find a recurrent set R which intersects the set of reachable states.

9. MODEL CHECKING AND SOFTWARE QUALITY

In the preceding sections, we have traced the main developments in software model checking. We now indicate some recent interactions and synergies between model checking and allied fields.

9.1. Model Checking and Testing

Testing involves running software on a set of inputs. Testing is the primary technique to ensure software quality in the software industry. Sometimes the term *dynamic analysis* is used for testing, denoting that the program is actually executed and its outputs observed, as opposed to *static analysis* in which a *mathematical model* of the system is analyzed (e.g., in abstraction-based model checking). Like most classifications, the boundaries are somewhat blurred, and some tools mix dynamic and static analysis. For example, Spin allows users to write arbitrary C code for part of the model that is executed during the model checking search.

Systematic Exploration. Each of the execution-based model checkers described in Section 2.4 can be viewed as a testing tool that systematically explores the space of behaviors. For a large software system, the possible number of behaviors is so large that exhaustive testing is unlikely to finish within the software development budget. Thus, the goal of systematic testing is to explore a subset of program behaviors that are “most likely” to uncover problems in the code, and a model checker can be used in “bug-finding” mode in which it searches as many behaviors as allowed within system resources. Indeed, typically these tools work by “amplifying” the effectiveness of a given test suite for the program being verified. They run the program using workloads drawn from the test suite, but they systematically explore the effects of different scheduling choices, failures etc. In this setting, optimizations derived for model checking, such as partial order reduction, are transferrable *mutatis mutandis* into optimizations of the testing process to rule out tests “similar” (in a precise sense) to those already run. Further, state space exploration tools can be configured to provide systematic *underapproximations* to program behaviors by fixing parameters such as input domain, search depth, or the number of context switches in concurrent code, and then checking behaviors within this underapproximation exhaustively.

Test Generation by Symbolic Evaluation. Model checking can be combined with symbolic execution, in order to *generate* test cases. For example, in Beyer et al. [2004] and Xia et al. [2005] a software model checker is used to find program paths satisfying certain coverage conditions (e.g., a path reaching a particular location, or taking a particular branch), and the symbolic constraints generated from the path are solved by a constraint solver to produce test inputs. While the idea of using symbolic execution for test case generation is old [Clarke 1976], the combination with model checking allows

search for test inputs satisfying particular coverage requirements to benefit from the search strategy of the model checker over an abstracted state space.

More recently, combined concrete and symbolic execution (or *concolic execution*) [Godefroid et al. 2005; Sen et al. 2005; Cadar et al. 2006] has been suggested as a strategy to combine symbolic execution and testing. In this technique, first proposed in Godefroid et al. [2005] and independently in Cadar et al. [2006], the program is run on concrete inputs (chosen, e.g., at random) and while it is running, symbolic constraints on the execution path (as a function of symbolic inputs) are generated (e.g., through program instrumentation). By systematically negating symbolic constraints at conditional branches, and solving these constraints using a decision procedure, one generates a set of test inputs exploring every program path. The dynamic execution ensures that (a) parts of the code which do not depend on symbolic inputs are not tracked symbolically (this has been a major performance bottleneck of “pure” symbolic execution), and (b) program semantics hard to capture statically, for example, dynamic memory allocation, can be tracked at run time. Additionally, run time values can be used to simplify symbolic constraints (at the cost of missing certain executions). For example, consider the program

```
x = input();
for (i = 0; i < 1000; i++) a[i] = 0;
```

Pure symbolic execution techniques would symbolically unroll the loop 1000 times, and the resulting constraints would slow down a decision procedure. In contrast, dynamic symbolic execution runs the loop concretely without generating additional symbolic constraints.

In Gulavani et al. [2006] and Beckman et al. [2008], the two ideas of CEGAR-based software model checking and dynamic symbolic execution have been combined to simultaneously search an abstract state space to produce a proof of correctness using predicate abstraction and to search for a failing test case using dynamic symbolic execution on counterexamples returned by the model checker. If dynamic symbolic execution finds the counterexample to be infeasible, usual counterexample analysis techniques can be used to refine the abstraction. The algorithm combines the abilities of CEGAR, to quickly explore abstract state spaces, and of dynamic symbolic execution, to quickly explore particular program paths for feasibility.

9.2. Model Checking and Type Systems

Type systems are perhaps the most pervasive of all software verification techniques. Historically, the goal of types has been to classify program entities with a view towards ensuring that only well-defined operations are carried out at run-time. One can view a simple type system, such as the type system of Java or ML, as a technique for computing very coarse invariants over program variables and expressions. For example, the fact that x has type `int` is essentially the invariant that in every reachable state of the program, an integer value is held in x . Thus, type systems provide a scalable technique for computing coarse-grained invariants. Typically, these invariants have been “flow-insensitive” meaning they are facts that hold at *every* program point, and hence, they cannot be used to verify richer temporal properties. However, several recent approaches relax this restriction, and open the way to applying types to verify richer safety properties.

Typestates. Strom and Yemini [1986] extend types with a finite set of states corresponding to the different stages of the value’s lifetime. For example, a file handle type can be in two states: open if the handle refers to an open file, and, closed if the

handle refers to a closed file. The types of primitive operations like `open()` or `close()` are annotated with pre- and post-conditions describing the appropriate input and output tpestates, and a dataflow (or type-and-effect) analysis is used to determine the tpestates of different objects at different program points, and hence, verify temporal properties like a file should not be read or written after it has been closed. Examples of such analyses include *type qualifiers* [Foster et al. 2002] and *typestate interpretations* [Fahndrich and DeLine 2004], which map from tpestates to predicates over an object's fields, which can be exploited to verify object-oriented programs using inheritance. Each of these algorithms can be viewed as model checking the program over a fixed abstraction generated by the product of the control-flow graph and the tpestates [Chen and Wagner 2002]. Consequently, the algorithms are more efficient than general purpose software model checkers, but less precise as they ignore branches and other path information. The tool ESP [Das et al. 2002] combined tpestates with symbolic execution, and showed that the resulting combination can provide significantly higher precision (compared to path-insensitive tpestate algorithms) but with high scalability (compared to fully path sensitive model checkers).

Dependent Types. Martin-Löf [1984] provide a complementary approach for encoding invariants inside the type system, by refining the types with predicates that describe sets of values. For example, the refined type:

$$\{\nu:\text{int} \mid 0 \leq \nu \wedge \nu < n\} \text{ list}$$

describes a list of integers, where each integer, represented by ν is greater than 0 and less than the value of some program variable n . The NUPRL proof assistant [Constable 1986] introduced the idea of using subset typing with dependent typing in the context of intuitionistic type theory. The PVS proof assistant [Owre et al. 1996] takes this idea further by applying it to the setting of classical higher-order logic. By separating type- and proof- checking, PVS reduces the number of proof obligations that need to be discharged. Further, PVS can also generate code for a fairly large executable fragment, and is one of the first languages whose type system ensures that well-typed programs can never crash (except by running out of resources). The above systems allow very rich invariants to be encoded in the types but require the user to help discharge the subtyping obligations. In contrast, work on *dependent ML* [Xi and Pfenning 1999] shows how to extend ML with a restricted form of dependent types over a (parameterized) constraint domain C . Type checking is shown to be decidable modulo the decidability of the domain, thus allowing the programmer to specify and verify pre- and post-conditions that can be expressed over C . It is still assumed that the dependent types (which correspond to pre- and post-conditions and loop invariants in deductive verification) are provided by the programmer. Work on *Hoare type theory* [Nanevski et al. 2008] makes the connection between Hoare-style verification and types more explicit. The above approaches require that programmers provide type annotations corresponding to pre- and post- conditions for all functions. [Rondon et al. 2008] shows how the machinery of software model checking (predicate abstraction to be precise) can be brought to bear on the problem of *inferring* dependent type annotations, in much the same way as model checking can be viewed as a mechanism for synthesizing invariants. As a result one can combine the complementary strengths of model checking (local path- and value-information) and type systems (higher-order functions, recursive data, polymorphism) to verify properties that are well-beyond the abilities of either technique in isolation. Large scale combinations of software model checking algorithms for safety verification, which typically target imperative first-order programs, and dependent type systems, which typically target higher-order functional programs, is an interesting, but under-explored problem.

Hybrid type checking [Flanagan 2006] provides a pragmatic compromise: one starts with an expressive dependent type system and tries to prove as many assertions as possible statically, inserting dynamic checks in the program for those assertions that cannot be proved statically.

10. CONCLUSION

Software model checkers and related algorithmic verification tools hold the potential to close the gap between the programmer's intent and the actual code. However, the current generation of software model checking tools work best only for control-dominated protocol properties, and we are still far away from proving functional properties of complex software systems, such as data invariants. There are many remaining problems, both in scaling current techniques to large programs, and in devising algorithmic analyses for modern software systems. For example, scaling verification techniques in the presence of expressive heap abstractions and concurrent interactions remain outstanding open problems.

Many modern programming language features, such as object-orientation and dynamic dispatch, abstract data types, higher-order control flow and continuations, etc. are skirted in current algorithms and tools, and we would like to see verification tools exploiting language-level features. Similarly, common practice in large-scale software engineering, such as design patterns, the use of information hiding and layering, incremental development with regression tests, and design and architectural information is not exploited by current tools, but could be crucial in scaling tools to large scale software projects. An associated problem is to integrate software model checking into the general software engineering process. This leads to quite a few technical challenges, for example, how to model the environment of a software module, how to represent libraries and other third-party components for which code is not available, and how to make the verification process incremental. Finally, many tools make simplifying assumptions about the low-level machine semantics and data layout. While in principle, one can model the exact machine-level semantics, it is not clear if this level of modeling will preserve the scalability of tools. The problem of language-level semantics is exacerbated by software written in multiple programming languages communicating through inter-language APIs. Often, these APIs are brittle and a source of potential bugs. However, analyzing multi-language software has to model the semantics of the API precisely, and ensure that inter-language API calls do not break program invariants on either side. Each of these directions would form excellent research topics.

Despite the shortcomings, we believe software model checking has made excellent progress in the past decade by selecting winning combinations of ideas from many disciplines, and in several settings, verification techniques can complement or outperform more traditional quality assurance processes based on testing and code inspection in terms of cost and effectiveness.

On the whole, it is unlikely that just software model checking tools will turn software development into a routine effort. Developing reliable software is too complex a problem, and has social aspects in addition to technical ones. However, we believe that the emergence of automatic tools and their use in the development process will help amplify programmer productivity by checking for partial properties of code, leaving the programmer more time to focus on more complex issues.

ACKNOWLEDGMENTS

We thank Jay Misra and Tony Hoare for encouraging us to write this survey and several useful comments along the way. We thank Michael Emmi, Aarti Gupta, Roman Manevich, Andreas Podelski, Sriram Rajamani,

Andrey Rybalchenko, Natarajan Shankar, Ashish Tiwari, and the anonymous reviewers for feedback on earlier drafts.

REFERENCES

- ABADI, M. AND LAMPORT, L. 1993. Composing specifications. *ACM Trans. Prog. Lang. Syst.* 15, 1, 73–132.
- ABADI, M. AND LAMPORT, L. 1995. Conjoining specifications. *ACM Trans. Prog. Lang. Syst.* 17, 3, 507–534.
- AGERWALA, T. AND MISRA, J. 1978. Assertion graphs for verifying and synthesizing programs. Tech. Rep. 83, University of Texas, Austin, TX.
- ALPERN, B. AND SCHNEIDER, F. 1987. Recognizing safety and liveness. *Distributed Computing* 3, 3, 117–126.
- ALUR, R., DANG, T., AND IVANČIĆ, F. 2006. Counterexample-guided predicate abstraction of hybrid systems. *Theoret. Comput. Sci.* 354, 2, 250–271.
- ALUR, R. AND HENZINGER, T. 1999. Reactive modules. *Form. Meth. Syst. Des.* 15, 1, 7–48.
- ALUR, R., HENZINGER, T., MANG, F., QADEER, S., RAJAMANI, S., AND TASIRAN, S. 1998. MOCHA: Modularity in model checking. In *CAV 98: Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 1427. Springer-Verlag, Berlin, Germany, 521–525.
- ALUR, R., ITAI, A., KURSHAN, R., AND YANNAKAKIS, M. 1995. Timing verification by successive approximation. *Inf. Comput.* 118, 1, 142–157.
- ANDERSEN, L. 1994. Program analysis and specialization for the c programming language. Ph.D. dissertation. DIKU, University of Copenhagen, Copenhagen.
- APT, K. AND OLDEROG, E.-R. 1991. *Verification of Sequential and Concurrent Programs*. Springer-Verlag, Berlin, Germany.
- ARMANDO, A., MANTOVANI, J., AND PLATANIA, L. 2006. Bounded model checking of software using SMT solvers instead of SAT solvers. In *Model Checking Software: SPIN Workshop*. Lecture Notes in Computer Science, vol. 3925. Springer-Verlag, Berlin, Germany, 146–162.
- BABIC, D. AND HU, A. 2008. Calyso: scalable and precise extended static checking. In *ICSE 08: Proceedings of the International Conference on Software Engineering*. ACM, 211–220.
- BAGNARA, R., HILL, P. M., MAZZI, E., AND ZAFFANELLA, E. 2005. Widening operators for weakly-relational numeric abstractions. In *SAS 05: Proceedings of the Static Analysis Symposium*. Lecture Notes in Computer Science, vol. 3672. Springer-Verlag, New York, 3–18.
- BAGNARA, R., HILL, P. M., AND ZAFFANELLA, E. 2008. The Parma Polyhedra Library: Toward a complete set of numerical abstractions for the analysis and verification of hardware and software systems. *Sci. Comput. Prog.* 72, 1-2, 3–21.
- BALL, T., BOUNIMOVA, E., COOK, B., LEVIN, V., LICHTENBERG, J., MCGARVEY, C., ONDRUSEK, B., RAJAMANI, S. K., AND USTUNER, A. 2006. Thorough static analysis of device drivers. *EuroSys*. 73–85.
- BALL, T., MAJUMDAR, R., MILLSTEIN, T., AND RAJAMANI, S. K. 2001. Automatic predicate abstraction of C programs. In *PLDI 01: Proceedings of the Symposium on Programming Languages Design and Implementation*. ACM, New York, 203–213.
- BALL, T., MILLSTEIN, T. D., AND RAJAMANI, S. K. 2005. Polymorphic predicate abstraction. *ACM Trans. Prog. Lang. Syst.* 27, 2, 314–343.
- BALL, T., PODELSKI, A., AND RAJAMANI, S. K. 2001. Boolean and Cartesian abstractions for model checking C programs. In *TACAS 01: Proceedings of the Symposium on Tools and Algorithms for Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 2031. Springer-Verlag, New York, 268–283.
- BALL, T., PODELSKI, A., AND RAJAMANI, S. K. 2002. Relative completeness of abstraction refinement for software model checking. In *TACAS 02: Proceedings of the Symposium on Tools and Algorithms for Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 2280. Springer-Verlag, New York, 158–172.
- BALL, T. AND RAJAMANI, S. 2002a. Generating abstract explanations of spurious counterexamples in C programs. Tech. Rep. MSR-TR-2002-09, Microsoft Research.
- BALL, T. AND RAJAMANI, S. 2002b. The SLAM project: debugging system software via static analysis. In *POPL 02: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 1–3.
- BALL, T. AND RAJAMANI, S. K. 2000a. Bebop: A symbolic model checker for Boolean programs. In *SPIN 00: Proceedings of the SPIN Workshop*. Lecture Notes in Computer Science 1885. Springer-Verlag, 113–130.
- BALL, T. AND RAJAMANI, S. K. 2000b. Boolean programs: a model and process for software analysis. Tech. Rep. MSR Technical Report 2000-14, Microsoft Research.

- BECKMAN, N., NORI, A. V., RAJAMANI, S. K., AND SIMMONS, R. J. 2008. Proofs from tests. In *ISSTA 08: Proceedings of the International Symposium on Software Testing and Analysis*. ACM, New York, 3–14.
- BEYER, D., CHLIPALA, A. J., HENZINGER, T. A., JHALA, R., AND MAJUMDAR, R. 2004. Generating tests from counterexamples. In *ICSE 04: Proceedings of the International Conference on Software Engineering*. ACM, New York, 326–335.
- BEYER, D., HENZINGER, T., MAJUMDAR, R., AND RYBALCHENKO, A. 2007a. Invariant synthesis in combination theories. In *VMCAI 07: Proceedings of the Symposium on Verification, Model Checking, and Abstract Interpretation*. Lecture Notes in Computer Science, vol. 4349. Springer-Verlag, Berlin, Germany, 378–394.
- BEYER, D., HENZINGER, T., MAJUMDAR, R., AND RYBALCHENKO, A. 2007b. Path invariants. In *PLDI 07: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 300–309.
- BEYER, D., HENZINGER, T. A., JHALA, R., AND MAJUMDAR, R. 2007c. The software model checker blast. *Softw. Tools Tech. Trans.* 9, 5-6, 505–525.
- BEYER, D., HENZINGER, T. A., AND THÉODULOZ, G. 2007. Configurable software verification: Concretizing the convergence of model checking and program analysis. In *CAV 07: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 4590. Springer-Verlag, Berlin, Germany, 504–518.
- BIERE, A., CIMATTI, A., CLARKE, E. M., FUJITA, M., AND ZHU, Y. 1999. Symbolic model checking using SAT procedures instead of BDDs. In *DAC 99: Proceedings of the Design Automation Conference*. ACM, New York, 317–320.
- BLANCHET, B., COUSOT, P., COUSOT, R., FERET, J., MAUBORGNE, L., MINE, A., MONNIAUX, D., AND RIVAL, X. 2002. Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software. In *The Essence of Computation, Complexity, Analysis, Transformation: Essays Dedicated to Neil D. Jones*. Lecture Notes in Computer Science, vol. 2566. Springer-Verlag, Berlin, Germany, 85–108.
- BLANCHET, B., COUSOT, P., COUSOT, R., FERET, J., MAUBORGNE, L., MINE, A., MONNIAUX, D., AND RIVAL, X. 2003. A static analyzer for large safety-critical software. In *PLDI 03: Proceedings of the Symposium on Programming Languages Design and Implementation*. ACM, New York, 196–207.
- BOUAIJANI, A., ESPARZA, J., AND MALER, O. 1994. Reachability analysis of pushdown automata: application to model checking. In *CONCUR 97: Proceedings of the Symposium on Concurrency Theory*. Lecture Notes in Computer Science, vol. 1243. Springer-Verlag, Berlin, Germany, 135–150.
- BOUAIJANI, A., ESPARZA, J., AND TOULI, T. 2003. A generic approach to the static analysis of concurrent programs with procedures. In *POPL 03: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 62–73.
- BOUAIJANI, A., FERNANDEZ, J.-C., AND HALBWACHS, N. 1990. Minimal model generation. In *CAV 90: Proceedings of the Symposium on Computer-aided Verification*. Lecture Notes in Computer Science, vol. 531. Springer-Verlag, Berlin, Germany, 197–203.
- BRADLEY, A., MANNA, Z., AND SIPMA, H. 2005. The polyranking principle. In *ICALP 05: Proceedings of the International Colloquium on Automata, Languages, and Programming*. Lecture Notes in Computer Science, vol. 3580. Springer-Verlag, Berlin, Germany, 1349–1361.
- BRAT, G., DRUSINSKY, D., GIANNAKOPOLOU, D., GOLDBERG, A., HAVELUND, K., LOWRY, M., PASAREANU, C., VENET, A., WASHINGTON, R., AND VISSER, W. 2004. Experimental evaluation of verification and validation tools on Martian rover software. *Form. Meth. Syst. Des.* 25.
- BRUTTOMESSO, R., CIMATTI, A., FRANZÉN, A., GRIGGIO, A., AND SEBASTIANI, R. 2008. The MathSAT 4 SMT solver. In *CAV 08: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 5123. Springer-Verlag, Berlin, Germany, 299–303.
- BRYANT, R. 1986. Graph-based algorithms for Boolean function manipulation. *IEEE Trans. Comput.* C-35, 8, 677–691.
- BURCH, J., CLARKE, E., McMILLAN, K., DILL, D., AND HWANG, L. 1992. Symbolic model checking: 10^{20} states and beyond. *Inf. Comput.* 98, 2, 142–170.
- BUSTAN, D. AND GRUMBERG, O. 2003. Simulation-based minimization. *ACM Trans. Comput. Logic* 4, 181–206.
- CADAR, C., GANESH, V., PAWLOWSKI, P., DILL, D., AND ENGLER, D. 2006. EXE: automatically generating inputs of death. In *CCS 02: Proceedings of the Conference on Computer and Communications Security*. ACM, New York.
- CHAKI, S., CLARKE, E., GROCE, A., OUAKNINE, J., STRICHMAN, O., AND YORAV, K. 2004. Efficient verification of sequential and concurrent C programs. *Form. Meth. Syst. Des.* 25, 2-3, 129–166.

- CHAKI, S., CLARKE, E., KIDD, N., REPS, T., AND TOULI, T. 2006. Verifying concurrent message-passing C programs with recursive calls. In *TACAS 06: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 3920. Springer-Verlag, Berlin, Germany, 334–349.
- CHAKI, S., CLARKE, E. M., GROCE, A., AND STRICHMAN, O. 2003. Predicate abstraction with minimum predicates. In *CHARME*. 19–34.
- CHANDRA, S., GODEFROID, P., AND PALM, C. 2002. Software model checking in practice: an industrial case study. In *ICSE 02: Proceedings of the International Conference on Software Engineering*. ACM, New York, 431–441.
- CHANG, B. E. AND RIVAL, X. 2008. Relational inductive shape analysis. In *POPL 08: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 247–260.
- CHASE, D., WEGMAN, M., AND ZADECK, F. 1990. Analysis of pointers and structures. In *PLDI 90: Proceedings of the Symposium on Programming Languages Design and Implementation*. ACM, New York, 296–310.
- CHAUDHURI, S. 2008. Sub-cubic algorithms for recursive state machines. In *POPL 08: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 159–169.
- CHEN, H. AND WAGNER, D. 2002. MOPS: an infrastructure for examining security properties of software. In *Proceedings of the ACM Conference on Computer and Communications Security 2002*. 235–244.
- CLARKE, E. M. AND EMERSON, E. A. 1981. Synthesis of synchronization skeletons for branching time temporal logic. In *Logic of Programs*. Lecture Notes in Computer Science, vol. 131. Springer-Verlag, Berlin, Germany, 52–71.
- CLARKE, E. M., FILKORN, T., AND JHA, S. 1993. Exploiting symmetry in temporal logic model checking. In *CAV 93: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 697. Springer-Verlag, Berlin, Germany, 450–462.
- CLARKE, E. M., GRUMBERG, O., JHA, S., LU, Y., AND VEITH, H. 2000. Counterexample-guided abstraction refinement. In *CAV 00: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 1855. Springer-Verlag, Berlin, Germany, 154–169.
- CLARKE, E. M., GRUMBERG, O., AND LONG, D. 1992. Model checking and abstraction. In *POPL 92: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 343–354.
- CLARKE, L. 1976. A system to generate test data and symbolically execute programs. *IEEE Trans. Softw. Eng.* 2, 2, 215–222.
- COLÓN, M. AND SIPMA, H. 2001. Synthesis of linear ranking functions. In *TACAS 01: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 2031. Springer-Verlag, Berlin, Germany, 67–81.
- COLÓN, M. AND SIPMA, H. 2002. Practical methods for proving program termination. In *CAV 02: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 2404. Springer-Verlag, Berlin, Germany, 442–454.
- CONSTABLE, R. 1986. *Implementing Mathematics with the Nuprl Proof Development System*. Prentice-Hall, Englewood Cliffs, NJ.
- COOK, B., PODELSKI, A., AND RYBALCHENKO, A. 2006. Termination proofs for systems code. In *PLDI 06: Proceedings of the Symposium on Programming Languages Design and Implementation*. ACM, New York, 415–426.
- COOK, S. A. 1978. Soundness and completeness of an axiom system for program verification. *SIAM J. Comput.* 7, 1, 70–90.
- CORBETT, J., DWYER, M., HATCLIFF, J., PASAREANU, C., ROBBY, LAUBACH, S., AND ZHENG, H. 2000. Bandera: Extracting finite-state models from Java source code. In *ICSE 00: Proceedings of the International Conference on Software Engineering*. 439–448.
- COURCOUBETIS, C., VARDI, M., WOLPER, P., AND YANNAKAKIS, M. 1992. Memory-efficient algorithms for the verification of temporal properties. *Form. Meth. Syst. Des.* 1, 275–288.
- COUSOT, P. AND COUSOT, R. 1976. Static determination of dynamic properties of programs. In *ISOP*. 106–130.
- COUSOT, P. AND COUSOT, R. 1977. Abstract interpretation: a unified lattice model for the static analysis of programs. In *POPL 77: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 238–252.
- COUSOT, P. AND COUSOT, R. 1979. Systematic design of program analysis frameworks. In *POPL 79: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 269–282.
- COUSOT, P. AND COUSOT, R. 2000. Temporal abstract interpretation. In *POPL 00: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 12–25.

- COUSOT, P. AND HALBWACHS, N. 1978. Automatic discovery of linear restraints among variables of a program. In *POPL 78: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York.
- CRAIG, W. 1957. Linear reasoning. *J. Symbol. Logic* 22, 250–268.
- DAS, M., LERNER, S., AND SEIGLE, M. 2002. ESP: Path-sensitive program verification in polynomial time. In *PLDI 02: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 57–68.
- DAS, S., DILL, D. L., AND PARK, S. 1999. Experience with predicate abstraction. In *CAV 99: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 1633. Springer-Verlag, Berlin, Germany, 160–171.
- DE MILLO, R., LIPTON, R., AND PERLIS, A. 1979. Social processes and proofs of theorems and programs. *Commun. ACM* 22, 271–280.
- DE MOURA, L. AND BJØRNER, N. 2008. Z3: An efficient SMT solver. In *TACAS 08: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 4963. Springer-Verlag, Berlin, Germany, 337–340.
- DE MOURA, L. AND RUESS, H. 2003. Bounded model checking and induction: From refutation to verification. In *CAV 03: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 2725. Springer-Verlag, Berlin, Germany, 14–26.
- DE MOURA, L., RUESS, H., AND SOREA, M. 2002. Lazy theorem proving for bounded model checking over infinite domains. In *CADE 02: Proceedings of the Symposium on Automated Deduction*. Lecture Notes in Computer Science, vol. 2392. Springer-Verlag, Berlin, Germany, 438–455.
- DIJKSTRA, E. 1976. *A Discipline of Programming*. Prentice-Hall, Englewood Cliffs, NJ.
- DILL, D. 1996. The Murphi verification system. In *CAV 96: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 1102. Springer-Verlag, Berlin, Germany, 390–393.
- DIMITROVA, R. AND PODELSKI, A. 2008. Is lazy abstraction a decision procedure for broadcast protocols? In *VMCAI 08: Proceedings of the Symposium on Verification, Model Checking, and Abstract Interpretation*. Lecture Notes in Computer Science, vol. 4905. Springer-Verlag, Berlin, Germany, 98–111.
- DISTEFANO, D., O’HEARN, P. W., AND YANG, H. 2006. A local shape analysis based on separation logic. In *TACAS 06: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 3920. Springer-Verlag, Berlin, Germany, 287–302.
- DONER, J. E. 1965. Decidability of the weak second-order theory of two successors. *Notices Amer. Math. Soc.* 12, 365–468.
- DUTERTE, B. AND DE MOURA, L. 2006. A fast linear-arithmetic solver for DPLL(T). In *CAV 06: Proceedings of the International Conference on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 4144. Springer-Verlag, Berlin, Germany, 81–94.
- DWYER, M. AND CLARKE, L. 1994. Data flow analysis for verifying properties of concurrent programs. In *FSE 94: Proceedings of the Symposium on Foundations of Software Engineering*. ACM, New York, 62–75.
- EDELKAMP, S., LEUE, S., AND LLUCH-LAFUENTE, A. 2004. Directed explicit-state model checking in the validation of communication protocols. *Softw. Tools Tech. Trans.* 5, 247–267.
- EEN, N. AND SORENSSON, N. 2003. An extensible SAT solver. In *SAT 2003: Proceedings of the 6th International Conference on Theory and Applications of Satisfiability Testing*. Lecture Notes in Computer Science, vol. 2919. Springer-Verlag, Berlin, Germany, 502–518.
- EMERSON, E. 1990. Temporal and modal logic. In *Handbook of Theoretical Computer Science*, J. van Leeuwen, Ed. vol. B. Elsevier Science Publishers, Amsterdam, the Netherlands, 995–1072.
- EMERSON, E. AND LEI, C. 1986. Efficient model checking in fragments of the propositional μ -calculus. In *Proceedings of the 1st Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, Los Alamitos, CA, 267–278.
- EMERSON, E. AND SISTLA, A. 1996. Symmetry and model checking. *Form. Meth. Syst. Des.* 9, 105–131.
- ESPARZA, J. AND SCHWOON, S. 2001. A BDD-based model checker for recursive programs. In *CAV*. 324–336.
- FAHNDRICH, M. AND DELINE, R. 2004. Typestates for objects. In *ECOOP 04: Proceedings of the Symposium on Object-Oriented Programming*. Lecture Notes in Computer Science, vol. 3086. Springer-Verlag, Berlin, Germany, 465–490.
- FISCHER, J., JHALA, R. AND MAJUMDAR, R. 2005. Joining dataflow with predicates. In *ESEC/FSE 2005: Proceedings of the Symposium on Foundations of Software Engineering*. ACM, New York, 227–236.
- FLANAGAN, C. 2006. Hybrid type checking. In *POPL 06: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York.

- FLANAGAN, C., FREUND, S. AND QADEER, S. 2002. Thread-modular verification for shared-memory programs. In *ESOP 02: Proceedings of the European Symposium on Programming*. Lecture Notes in Computer Science, vol. 2305. Springer-Verlag, Berlin, Germany, 262–277.
- FLANAGAN, C., FREUND, S., QADEER, S., AND SESHIA, S. 2005. Modular verification of multithreaded programs. *Theoret. Comput. Sci.* 338, 153–183.
- FLANAGAN, C., JOSHI, R., AND LEINO, K. R. M. 2001. Annotation inference for modular checkers. *Inf. Proc. Lett.* 77, 2-4, 97–108.
- FLANAGAN, C., LEINO, K., LILLIBRIDGE, M., NELSON, G., SAXE, J. B., AND STATA, R. 2002. Extended static checking for Java. In *PLDI 02: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 234–245.
- FLANAGAN, C. AND QADEER, S. 2002. Predicate abstraction for software verification. In *POPL 02: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 191–202.
- FLANAGAN, C. AND SAXE, J. 2000. Avoiding exponential explosion: generating compact verification conditions. In *POPL 00: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 193–205.
- FLOYD, R. 1967. Assigning meanings to programs. In *Mathematical Aspects of Computer Science*. American Mathematical Society, 19–32.
- FOSTER, J., TERAUCHI, T., AND AIKEN, A. 2002. Flow-sensitive type qualifiers. In *PLDI 02: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 1–12.
- FRANCEZ, N. 1986. *Fairness*. Springer-Verlag, Berlin, Germany.
- FRASER, R., KAMHI, G., ZIV, B., VARDI, M., AND FIX, L. 2000. Prioritized traversal: efficient reachability analysis for verification and falsification. In *CAV 00: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 1855. Springer-Verlag, Berlin, Germany, 389–402.
- GANAI, M. AND GUPTA, A. 2006. Accelerating high-level bounded model checking. In *ICCAD 06: Proceedings of the International Conference on Computer-Aided Design*. ACM, New York, 794–801.
- GHIYA, R. AND HENDREN, L. J. 1996. Is it a tree, a DAG, or a cyclic graph? A shape analysis for heap-directed pointers in C. In *POPL 96: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 1–15.
- GIESL, J. AND KAPUR, D. 2001. Decidable classes of inductive theorems. In *IJCAR 2001: Proceedings of the International Joint Conference on Automated Reasoning*. Lecture Notes in Computer Science, vol. 2083. Springer-Verlag, Berlin, Germany, 469–484.
- GODEFROID, P. 1996. *Partial-Order Methods for the Verification of Concurrent Systems—An Approach to the State-Explosion Problem*. Lecture Notes in Computer Science, vol. 1032. Springer-Verlag, Berlin, Germany.
- GODEFROID, P. 1997. Model checking for programming languages using Verisoft. In *POPL 97: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 174–186.
- GODEFROID, P., KLARLUND, N., AND SEN, K. 2005. DART: Directed Automated Random Testing. In *PLDI 05: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 213–223.
- GOPAN, D., REPS, T. W., AND SAGIV, S. 2005. A framework for numeric analysis of array operations. In *POPL 05: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 338–350.
- GRAF, S. AND SAÏDI, H. 1997. Construction of abstract state graphs with PVS. In *CAV*. Lecture Notes in Computer Science, vol. 1254. Springer-Verlag, Berlin, Germany, 72–83.
- GULAVANI, B. S., CHAKRABORTY, S., NORI, A. V., AND RAJAMANI, S. K. 2008. Automatically refining abstract interpretations. In *TACAS 08: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 4963. Springer-Verlag, Berlin, Germany, 443–458.
- GULAVANI, B. S., HENZINGER, T. A., KANNAN, Y., NORI, A. V., AND RAJAMANI, S. K. 2006. Synergy: A new algorithm for property checking. In *FSE 06: Proceedings of the Symposium on Foundations of Software Engineering*. ACM, New York, 117–127.
- GULWANI, S., MCCLOSKEY, B., AND TIWARI, A. 2008. Lifting abstract interpreters to quantified logical domains. In *POPL 08: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 235–246.
- GULWANI, S. AND TIWARI, A. 2006. Combining abstract interpreters. In *PLDI 2006: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 376–386.
- GUPTA, A., HENZINGER, T., MAJUMDAR, R., RYBALCHENKO, A., AND XU, R. 2008. Proving non-termination. In *POPL 08: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 147–158.

- GUPTA, A., MAJUMDAR, R., AND RYBALCHENKO, A. 2009. From tests to proofs. In *TACAS 09: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 5505. Springer-Verlag, Berlin, Germany, 262–276.
- HACKETT, B. AND AIKEN, A. 2006. How is aliasing used in systems software? In *FSE 06: Proceedings of the Symposium on Foundations of Software Engineering*. ACM, New York, 69–80.
- HARDEKOPF, B. AND LIN, C. 2007. The ant and the grasshopper: Fast and accurate pointer analysis for millions of lines of code. In *PLDI 07: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 290–299.
- HART, P., NILSSON, N., AND RAPHAEL, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. Sci. Cybernetics SSC4 2*, 100–107.
- HAVELUND, K. AND PRESSBURGER, T. 2000. Model checking Java programs using Java Pathfinder. *Softw. Tools Tech. Trans. (STTT) 2*, 4, 72–84.
- HENZINGER, T., JHALA, R., MAJUMDAR, R., AND MCMILLAN, K. 2004. Abstractions from proofs. In *POPL 04: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 232–244.
- HENZINGER, T., JHALA, R., MAJUMDAR, R., AND QADEER, S. 2003. Thread-modular abstraction refinement. In *CAV 03: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany.
- HENZINGER, T., JHALA, R., MAJUMDAR, R., AND SUTRE, G. 2002. Lazy abstraction. In *POPL 02: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 58–70.
- HENZINGER, T., QADEER, S., AND RAJAMANI, S. 1998. You assume, we guarantee: methodology and case studies. In *CAV 98: Proceedings of the Symposium on Computer-Aided Verification*, Lecture Notes in Computer Science, vol. 1427. Springer-Verlag, Berlin, Germany, 440–451.
- HENZINGER, T. A., JHALA, R., AND MAJUMDAR, R. 2004. Race checking by context inference. In *PLDI 2004: Proceedings of the Symposium on Programming Languages Design and Implementation*. ACM, New York, 1–12.
- HIND, M. 2001. Pointer analysis: Haven't we solved this problem yet? In *Proceedings of the ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering (PASTE'01)*. ACM, New York, 54–61.
- HOARE, C. 1969. An axiomatic basis for computer programming. *Commun. ACM 12*, 576–580.
- HOLZMANN, G. 1997. The Spin model checker. *IEEE Trans. Softw. Eng.* 23, 5 (May), 279–295.
- IMMERMAN, N., RABINOVICH, A., REPS, T., SAGIV, M., AND YORSH, G. 2004. The boundary between decidability and undecidability for transitive-closure logics. In *CSL*. 160–174.
- IP, C. AND DILL, D. 1996. Better verification through symmetry. *Form. Meth. Syst. Des.* 9, 41–75.
- IVANCIC, F., YANG, Z., GANAI, M. K., GUPTA, A., AND ASHAR, P. 2008. Efficient SAT-based bounded model checking for software verification. *Theoret. Comput. Sci.* 404, 3, 256–274.
- IVANCIC, F., YANG, Z., GANAI, M. K., GUPTA, A., SHLYAKHTER, I., AND ASHAR, P. 2005. F-soft: Software verification platform. In *CAV 05: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 301–306.
- JAIN, H., IVANČIĆ, F., GUPTA, A., SHLYAKHTER, I., AND WANG, C. 2006. Using statically computed invariants inside the predicate abstraction and refinement loop. In *CAV 06: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 137–151.
- JHA, S. K., KROGH, B. H., WEIMER, J. E., AND CLARKE, E. M. 2007. Reachability for linear hybrid automata using iterative relaxation abstraction. In *Proceedings of the 12th ACM International Conference on Hybrid Systems: Computation and Control*. ACM, New York, 287–300.
- JHALA, R. AND MAJUMDAR, R. 2005. Path slicing. In *PLDI 05: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 38–47.
- JHALA, R. AND MCMILLAN, K. 2006. A practical and complete approach to predicate refinement. In *TACAS 06: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 2987. Springer-Verlag, Berlin, Germany, 298–312.
- JHALA, R. AND MCMILLAN, K. L. 2005. Interpolant-based transition relation approximation. In *CAV 05: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 39–51.
- JONES, C. 1983. Tentative steps toward a development method for interfering programs. *ACM Trans. Programming Languages Systems* 5, 4, 596–619.
- KAHLON, V. AND GUPTA, A. 2007. On the analysis of interacting pushdown systems. In *POPL 07: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 303–314.

- KATZ, S. AND PELED, D. 1992. Verification of distributed programs using representative interleaving sequences. *Distrib. Comput.* 6, 2, 107–120.
- KILLIAN, C. E., ANDERSON, J. W., JHALA, R., AND VAHDAT, A. 2007. Life, death, and the critical transition: Finding liveness bugs in systems code (awarded best paper). In *Proceedings of the 4th Symposium on Networked Systems Design and Implementation (NSDI)*. ACM, New York, 243–256.
- KING, J. 1976. Symbolic execution and program testing. *Commun. ACM* 19, 7, 385–394.
- KORF, R. 1985. Depth-first iterative deepening: an optimal admissible tree search. *Artif. Intell.* 27, 97–109.
- KROENING, D., CLARKE, E., AND YORAV, K. 2003. Behavioral consistency of C and Verilog programs using bounded model checking. In *DAC 03: Proceedings of the Symposium on Design Automation Conference*. ACM, New York, 368–371.
- KURSHAN, R. 1994. *Computer-Aided Verification of Coordinating Processes*. Princeton University Press, Princeton, NJ.
- LAHIRI, S. K. AND BRYANT, R. E. 2004. Constructing quantified invariants via predicate abstraction. In *Proceedings of the 5th International Conference on Verification Model Checking, and Abstract Interpretation (VMCAI)*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 267–281.
- LAHIRI, S. K. AND QADEER, S. 2008. Back to the future: revisiting precise program verification using SMT solvers. In *POPL 08: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 171–182.
- LAL, A., TOULI, T., KIDD, N., AND REPS, T. W. 2008. Interprocedural analysis of concurrent programs under a context bound. In *TACAS 08: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 4963. Springer-Verlag, Berlin, Germany, 282–298.
- LAMPART, L. 1983. Specifying concurrent program modules. *ACM Trans. Prog. Lang. Syst.* 5, 2, 190–222.
- LEE, D. AND YANNAKAKIS, M. 1992. Online minimization of transition systems. In *Proceedings of the 24th Annual Symposium on Theory of Computing*. ACM, New York, 264–274.
- LEHMANN, D., PNUELI, A., AND STAVI, J. 1982. Impartiality, justice, and fairness: The ethics of concurrent termination. In *ICALP 81: Proceedings of the International Conference on Automata, Languages, and Programming*. Lecture Notes in Computer Science 115. Springer-Verlag, Berlin, Germany, 264–277.
- LEINO, K. R. M. AND NELSON, G. 1998. An extended static checker for Modula-3. In *CC 98: Proceedings of the Symposium on Compiler Construction*. Lecture Notes in Computer Science, vol. 1383. Springer-Verlag, Berlin, Germany, 302–305.
- LEV-AMI, T. AND SAGIV, S. 2000. TVLA: A system for implementing static analyses. In *Proceedings of the 5th International Symposium on Static Analysis (SAS)*. Lecture Notes in Computer Science, vol. 1824. Springer-Verlag, Berlin, Germany, 280–301.
- LLUCH-LAFUENTE, A. 2003. Directed search for the verification of communication protocols. Ph.D. dissertation, Univ., Freiburg, Freiburg, Germany.
- LOISEAUX, C., GRAF, S., SIFAKIS, J., BOUAIJANI, A., AND BENSALÉM, S. 1995. Property-preserving abstractions for the verification of concurrent systems. *Form. Meth. Syst. Des.* 6, 11–44.
- MAGILL, S., BERDINE, J., CLARKE, E. M., AND COOK, B. 2007. Arithmetic strengthening for shape analysis. In *Proceedings of the 14th International Symposium on Static Analysis*. Lecture Notes in Computer Science, Springer-Verlag, Berlin, Germany, 419–436.
- MANNA, Z. AND PNUELI, A. 1992. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, Berlin, Germany.
- MARTIN-LÖF, P. 1984. Constructive mathematics and computer programming. *Roy. Soc. London Philos. Trans. Ser. A* 312, 501–518.
- McMILLAN, K. 1993. *Symbolic Model Checking: An Approach to the State-Explosion Problem*. Kluwer Academic Publishers.
- McMILLAN, K. L. 2004. An interpolating theorem prover. In *TACAS: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 16–30.
- McMILLAN, K. L. 2006. Lazy abstraction with interpolants. In *CAV 2006: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 123–136.
- McMILLAN, K. L. 2008. Quantified invariant generation using an interpolating saturation prover. In *Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 413–427.

- MEYER, R., FABER, J., AND RYBALCHENKO, A. 2006. Model checking duration calculus: A practical approach. In *ICTAC 06: Proceedings of the 3rd International Colloquium on Theoretical Aspects of Computing*. Lecture Notes in Computer Science, vol. 4281. Springer-Verlag, Berlin, Germany, 332–346.
- MINÉ, A. 2006. The octagon abstract domain. *Higher-Order Symb. Comput.* 19, 1, 31–100.
- MISRA, J. AND CHANDY, K. 1981. Proofs of networks of processes. *IEEE Trans. Softw. Eng. SE-7*, 4, 417–426.
- MØLLER, A. AND SCHWARTZBACH, M. I. 2001. The Pointer assertion logic engine. In *PLDI: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 221–231.
- MOSKEWICZ, M., MADIGAN, C., ZHAO, Y., ZHANG, L., AND MALIK, S. 2001. Chaff: Engineering an efficient SAT solver. In *DAC 01: Proceedings of the Design Automation Conference*. ACM, New York, 530–535.
- MUCHNICK, S. 1997. *Advanced Compiler Design and Implementation*. Morgan-Kaufman, San Francisco, CA.
- MUSUVATHI, M. AND ENGLER, D. R. 2004. Model checking large network protocol implementations. In *NSDI*. 155–168.
- MUSUVATHI, M. AND QADEER, S. 2007. Iterative context bounding for systematic testing of multithreaded programs. In *PLDI: Proceedings of the Symposium on Programming Languages Design and Implementation*. ACM, New York, 446–455.
- NAMJOSHI, K. S. AND KURSHAN, R. P. 2000. Syntactic program transformations for automatic abstraction. In *CAV 00: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 435–449.
- NANEVSKI, A., MORRISSET, G., SHINNAR, A., GOVEREAU, P., AND BIRKEDAL, L. 2008. Ynot: Reasoning with the awkward squad. In *Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming, ICFP*. ACM, New York.
- NECULA, G. C. AND LEE, P. 2000. Proof generation in the Touchstone theorem prover. In *CADE 00: Proceedings of the Symposium on Computer-Aided Deduction*. Lecture Notes in Computer Science, vol. 1831. Springer-Verlag, Berlin, Germany, 25–44.
- NELSON, G. 1981. Techniques for program verification. Tech. Rep. CSL81-10, Xerox Palo Alto Research Center, Palo Alto, CA.
- NELSON, G. 1983. Verifying reachability invariants of linked structures. In *POPL 83: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 38–47.
- NELSON, G. AND OPPEN, D. 1980. Fast decision procedures based on congruence closure. *J. ACM* 27, 2, 356–364.
- OWRE, S., RAJAN, S., RUSHBY, J., SHANKAR, N., AND SRIVAS, M. 1996. PVS: Combining specification, proof checking, and model checking. In *CAV 96: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 1102. Springer-Verlag, Berlin, Germany, 411–414.
- PASAREANU, C. S., GIANNAKOPOULOU, D., BOBARU, M. G., COBLEIGH, J. M. AND BARRINGER, H. 2008. Learning to divide and conquer: applying the L^+ algorithm to automate assume-guarantee reasoning. *Form. Meth. Syst. Des.* 32, 3, 175–205.
- PENIX, J., VISSER, W., PARK, S., PASAREANU, C., ENGSTROM, E., LARSON, A., AND WEININGER, N. 2005. Verifying time partitioning in the DEOS scheduling kernel. *Form. Meth. Syst. Des.* 26.
- PNUELI, A. 1977. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, CA, 46–57.
- PNUELI, A., PODELSKI, A., AND RYBALCHENKO, A. 2005. Separating fairness and well-foundedness for the analysis of fair discrete systems. In *TACAS 05: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, vol. 3440. Springer-Verlag, Berlin, Germany, 124–139.
- PODELSKI, A. AND RYBALCHENKO, A. 2004a. A complete method for the synthesis of linear ranking functions. In *VMCAI*. 239–251.
- PODELSKI, A. AND RYBALCHENKO, A. 2004b. Transition invariants. In *LICS 04: Proceedings of the Symposium on Logic in Computer Science*. IEEE, Computer Society Press, Los Alamitos, CA.
- PODELSKI, A. AND RYBALCHENKO, A. 2007. Armc: The logical choice for software model checking with abstraction refinement. In *PADL 07: Proceedings of the Symposium on Practical Aspects of Declarative Programming*. Lecture Notes in Computer Science, vol. 4354. Springer-Verlag, Berlin, Germany, 245–259.
- PODELSKI, A. AND WIES, T. 2005. Boolean Heaps. In *SAS: Proceedings of the Static Analysis Symposium*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 267–282.

- QADEER, S., RAJAMANI, S. K., AND REHOF, J. 2004. Summarizing procedures in concurrent programs. In *POPL 04: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 245–255.
- QADEER, S. AND REHOF, J. 2005. Context-bounded model checking of concurrent software. In *TACAS: Proceedings of the Symposium on Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 93–107.
- QADEER, S. AND WU, D. 2004. Kiss: Keep It Simple and Sequential. In *PLDI: Proceedings of the Symposium on Programming Languages Design and Implementation*. ACM, New York, 14–24.
- QUEILLE, J. AND SIFAKIS, J. 1981. Specification and verification of concurrent systems in CESAR. In *Proceedings of the 5th International Symposium on Programming*, M. Dezani-Ciancaglini and U. Montanari, Eds. Lecture Notes in Computer Science, vol. 137. Springer-Verlag, Berlin, Germany, 337–351.
- RAMALINGAM, G. 2000. Context-sensitive synchronization-sensitive analysis is undecidable. *ACM Trans. Prog. Lang. Syst.* 22, 2, 416–430.
- REPS, T., HORWITZ, S., AND SAGIV, M. 1995. Precise interprocedural dataflow analysis via graph reachability. In *POPL 95: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 49–61.
- REPS, T. W., SCHWON, S., JHA, S., AND MELSKI, D. 2005. Weighted pushdown systems and their application to interprocedural dataflow analysis. *Sci. Comput. Prog.* 58, 1-2, 206–263.
- REYNOLDS, J. C. 2002. Separation logic: A logic for shared mutable data structures. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 55–74.
- RONDON, P., KAWAGUCHI, M., AND JHALA, R. 2008. Liquid types. In *PLDI: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 158–169.
- RUSSELL, S. AND NORVIG, P. 2003. *Artificial Intelligence: A Modern Approach*, 2nd ed. Prentice-Hall, Englewood Cliffs, NJ.
- RYBALCHENKO, A. AND SOFRONIE-STOKKERMANS, V. 2007. Constraint solving for interpolation. In *VMCAI*. 346–362.
- SAGIV, S., REPS, T. W., AND WILHELM, R. 2002. Parametric shape analysis via 3-valued logic. *ACM Trans. Prog. Lang. Syst.* 24, 3, 217–298.
- SAIDI, H. 2000. Model checking guided abstraction and analysis. In *SAS 00: Proceedings of the Static Analysis Symposium*. Lecture Notes in Computer Science, vol. 1824, Springer-Verlag, Berlin, Germany, 377–396.
- SAÏDI, H. AND SHANKAR, N. 1999. Abstract and model check while you prove. In *CAV 99: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 1633. Springer-Verlag, Berlin, Germany, 443–454.
- SANKARANARAYANAN, S., SIPMA, H. B., AND MANNA, Z. 2005. Scalable analysis of linear systems using mathematical programming. In *VMCAI*. 25–41.
- SCHMIDT, D. 1998. Data flow analysis is model checking of abstract interpretation. In *POPL 98: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 38–48.
- SCHMIDT, D. A. AND STEFFEN, B. 1998. Program analysis as model checking of abstract interpretations. In *SAS 98: Proceedings of the Static Analysis Symposium*. Lecture Notes in Computer Science, vol. 1503. Springer-Verlag, ACM, New York, 351–380.
- SEN, K., MARINOV, D., AND AGHA, G. 2005. CUTE: A concolic unit testing engine for C. In *FSE*.
- SHARIR, M. AND PNUELI, A. 1981. Two approaches to interprocedural data flow analysis. In *Program Flow Analysis: Theory and Applications*. Prentice-Hall, Englewood Cliffs, NJ, 189–233.
- SHEERAN, M., SINGH, S., AND STALMARCK, G. 2000. Checking safety properties using induction and a SAT-solver. In *FMCAD 00: Proceedings of the Symposium on Formal Methods in Computer-Aided Design*. Lecture Notes in Computer Science, vol. 1954. Springer-Verlag, Berlin, Germany, 108–125.
- SHOSTAK, R. 1984. Deciding combinations of theories. *J. ACM* 31, 1, 1–12.
- SILVA, J. P. M. AND SAKALLAH, K. A. 1996. Grasp—A new search algorithm for satisfiability. In *ICCAD 96: Proceedings of the International Conference on Computer-Aided Design*. ACM, New York, 220–227.
- SISTLA, A., GYURIS, V., AND EMERSON, E. 2000. SMC: A symmetry-based model checker for verification of safety and liveness properties. *ACM Trans. Softw. Eng. Method.* 9, 133–166.
- SOMENZI, F. 1998. Colorado University decision diagram package. <http://vlsi.colorado.edu/pub/>.
- STARK, E. 1985. A proof technique for rely/guarantee properties. In *FSTTCS 1985: Foundations of Software Technology and Theoretical Computer Science*, S. N. Maheshwari, Ed. Lecture Notes in Computer Science, vol. 206. Springer-Verlag, Berlin, Germany, 369–391.

- STEENSGARD, B. 1996. Points-to analysis in almost linear time. In *POPL 96: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 32–41.
- STEFFEN, B. 1991. Data flow analysis as model checking. In *TACS 91: Proceedings of the Symposium on Theoretical Aspects of Computer Science*. Lecture Notes in Computer Science, vol. 536. Springer-Verlag, Berlin, Germany, 346–365.
- STERN, U. AND DILL, D. L. 1998. Using magnetic disk instead of main memory in the murhi verifier. In *CAV 98: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 172–183.
- STROM, R. AND YEMINI, S. 1986. Typestate: A programming language concept for enhancing software reliability. *IEEE Trans. Softw. Eng.* 12, 1, 157–171.
- SUNSHINE, C. 1978. Survey of protocol definition and verification techniques. *Comput. Netw.* 2, 346–350.
- SUNSHINE, C., THOMPSON, D., ERICKSON, R., GERHART, S., AND SCHWABE, D. 1982. Specification and verification of communication protocols in AFFIRM using state transition models. *IEEE Trans. Softw. Eng.* 8, 5, 460–489.
- SUWIMONTEERABUTH, D., ESPARZA, J., AND SCHWOON, S. 2008. Symbolic context-bounded analysis of multi-threaded Java programs. In *SPIN*. 270–287.
- TIWARI, A. 2004. Termination of linear programs. In *CAV 04: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 3114. Springer-Verlag, Berlin, Germany, 70–82.
- TURING, A. M. 1936. On computable numbers, with an application to the eintscheidungsproblem. In *Proceedings of the London Mathematical Society*. 230–265.
- VALMARI, A. 1992. A stubborn attack on state explosion. *Form. Meth. Syst. Des.* 1, 4, 297–322.
- VARDI, M. 1991. Verification of concurrent programs—the automata-theoretic framework. *Ann. Pure Appl. Logic* 51, 79–98.
- VARDI, M. 1995. An automata-theoretic approach to linear temporal logic. In *Logics for Concurrency—Structure versus Automata (8th Banff Higher Order Workshop Proceedings)*. Lecture Notes in Computer Science, vol. 1043. Springer-Verlag, Berlin, Germany, 238–266.
- VARDI, M. AND WOLPER, P. 1986. Automata-theoretic techniques for modal logics of programs. *J. Comput. Syst. Sci.* 32, 183–221.
- VARDI, M. AND WOLPER, P. 1994. Reasoning about infinite computations. *Inf. Comput.* 115, 1, 1–37.
- VELROYEN, H. AND RÜMMER, P. 2008. Non-termination checking for imperative programs. In *TAP: Tests and Proofs*. Lecture Notes in Computer Science, vol. 4966. Springer-Verlag, Berlin, Germany, 154–170.
- VISSER, W., HAVELUND, K., BRAT, G., PARK, S., AND LERDA, F. 2003. Model checking programs. *Automat. Softw. Eng. J.* 10.
- WALUKIEWICZ, I. 1996. Pushdown processes: Games and model checking. In *CAV 96: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 1102. Springer-Verlag, Berlin, Germany, 62–74.
- WANG, C., YANG, Z., GUPTA, A., AND IVANCIC, F. 2007. Using counterexamples for improving the precision of reachability computation with polyhedra. In *CAV 07: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 352–365.
- WHALEY, J. AND LAM, M. S. 2004. Cloning-based context-sensitive pointer alias analysis using binary decision diagrams. In *PLDI 04: Proceedings of the Symposium on Programming Language Design and Implementation*. ACM, New York, 131–144.
- XI, H. AND PFENNING, F. 1999. Dependent types in practical programming. In *POPL 99: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 214–227.
- XIA, S., VITO, B. D., AND MUÑOZ, C. 2005. Automated test generation for engineering applications. In *ASE 05: Proceedings of the Symposium on Automated Software Engineering*. ACM, New York, 283–286.
- XIE, Y. AND AIKEN, A. A. 2005. Scalable error detection using Boolean satisfiability. In *POPL 05: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York.
- YAHAV, E. 2001. Verifying safety properties of concurrent Java programs using 3-valued logic. In *POPL 01: Proceedings of the Symposium on Principles of Programming Languages*. ACM, New York, 27–40.
- YANG, C. H. AND DILL, D. L. 1998. Validation with guided search of the state space. In *DAC 98: Proceedings of the Symposium on Design Automation Conference*. ACM, New York, 599–604.
- YANG, H., LEE, O., BERDINE, J., CALCAGNO, C., COOK, B., DISTEFANO, D., AND O’HEARN, P. W. 2008. Scalable shape analysis for systems code. In *CAV 08: Proceedings of the Symposium on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 5123. Springer-Verlag, Berlin, Germany, 385–398.

- YANG, J., TWOHEY, P., ENGLER, D., AND MUSUVATHI, M. 2004. Using model checking to find serious file system errors. In *OSDI 04: Proceedings of the Symposium on Operating System Design and Implementation*. Usenix Association.
- YANG, Z., WANG, C., GUPTA, A., AND IVANCIC, F. 2006. Mixed symbolic representations for model checking software programs. In *MEMOCODE*. 17–26.
- YANNAKAKIS, M. 1990. Graph theoretic methods in database theory. In *Proceedings of the 9th ACM Symposium on Principles of Database Systems*. ACM, New York, 203–242.

Received December 2008; revised May 2009; accepted June 2009