

Model Checking: Software and Beyond

Edmund M. Clarke

Computer Science Department, Carnegie Mellon University
Pittsburgh, PA 15213 email: emc+@cs.cmu.edu

Flavio Lerda

Computer Science Department, Carnegie Mellon University
Pittsburgh, PA 15213 email: flerda+@cs.cmu.edu

I. INTRODUCTION

Temporal logic model checking, first developed by Clarke and Emerson [1] and independently discovered by Queille and Sifakis [2], is an automated technique for the verification of finite-state systems. The specification is expressed as a logical formula. Model checking aims at establishing whether a system is a model for a given formula, i.e., if it satisfies its specification. Model checking has had a big impact on formal verification over the past twenty five years [3], [4].

Section II describes the basic algorithm for temporal logic model checking, as well as some of the breakthroughs in this area. Section III introduces some recent developments and ideas for future research in this area.

II. MODEL CHECKING

The aim of formal methods is to prove that a given system satisfies its specification by formal means, e.g., mathematical proofs. In order to do so, the system and the specification need to be formally described. In model checking, the semantics of a system is usually given by means of a Kripke structure, a labeled graph that represents the possible states of a system and the transitions between them. The specification, instead, is expressed using temporal logic [5], an extension of propositional logic that allows reasoning about the relative timing of events. In the following, we will describe Kripke structures and temporal logics.

Kripke Structures

A Kripke structure is a directed graph where vertices are labeled by sets of atomic propositions. Vertices are called *states* and edges are called *transitions*. A subset of the states is designated as initial. An example of a Kripke structure, modeling the behavior of an alarm clock, is shown in Figure 1.

Given a set of *atomic proposition* AP , a **Kripke structure** is a tuple $M = (S, R, I, L)$ where S is a finite *set of states*, $R \subseteq S \times S$ is a *set of transitions* between states, $I \subseteq S$ is a *set of initial states*, and $L : S \rightarrow 2^{AP}$ is a *labeling function*. A *path* of a Kripke structure M is an infinite sequence of states $\pi = s_0, s_1, \dots$ such that $s_0 \in I$ and, for every $i \geq 0$, $(s_i, s_{i+1}) \in R$. Given a path π , π^i indicates the infinite path s_i, s_{i+1}, \dots and $\pi(i)$ indicates s_i . We assume that the transition relation R is total, i.e., every state has a valid successor and each finite path is a prefix of some infinite path.

The states of a Kripke structure represent the different states of a system. For example, the Kripke structure in Figure 1 has four states S_0 , S_1 , S_2 , and S_3 , corresponding to the alarm




Fig. 1. The Kripke structure of an alarm clock

clock being *off*, *on*, *ringing*, and *snoozed* respectively. Transitions of a Kripke structure represent the possible evolutions of a system. For example, the Kripke structure in Figure 1 has a transition from state S_2 to state S_3 , meaning that it is possible for the alarm clock to go from the state where it is ringing to the state where the alarm has been snoozed. However, there is no transition from state S_1 to state S_3 , because the alarm cannot be snoozed if it is not ringing.

Temporal Logic

Temporal logic is a formalism for reasoning about time without introducing time explicitly. It is an extension of propositional logic where temporal operators are introduced to reason about the timing of events. In model checking, two alternative temporal logics are commonly used: Computation Tree Logic (CTL) [6] and Linear-Time Temporal Logic (LTL) [7]. The former uses a branching notion of time and can reason about multiple paths at once. The latter uses a linear notion of time and considers a single path at a time. Both CTL and LTL can be expressed in terms of the temporal logic CTL* [8]. CTL and LTL formulas represent two overlapping, but different subsets of CTL* formulas.

The temporal logic CTL* defines two *path quantifiers* (the universal path quantifier **A** and the existential path quantifier **E**) and four *temporal operators* (next **X**, eventually **F**, globally **G**, and until **U**). There are two types of formulas: *state formulas* and *path formulas*. State formulas are defined as $f ::= p \mid f \vee f \mid f \wedge f \mid \neg f \mid \mathbf{A}g \mid \mathbf{E}g$ and path formulas are defined as $g ::= f \mid g \vee g \mid g \wedge g \mid \neg g \mid \mathbf{X}g \mid \mathbf{F}g \mid \mathbf{G}g \mid g \mathbf{U}g$ where $p \in AP$ is an atomic proposition. CTL* formulas are state formulas according to the above definition.

A state s of M satisfies the formula p if s is labeled by p ; it satisfies $\mathbf{A}g$ if every path π of M starting at s satisfies g ; and it satisfies $\mathbf{E}g$ if there exists a path π of M starting at s that satisfies g . A path π of M satisfies the state formula f if $\pi(0)$ satisfies f ; it satisfies $\mathbf{X}g$ if π^1 satisfies g ; it satisfies $\mathbf{F}g$ if there exists an $i \geq 0$ such that π^i satisfies g ; it satisfies $\mathbf{G}g$

if for every $i \geq 0$ path π^i satisfies g ; and it satisfies $g_1 \mathbf{U} g_2$ if there exists an $i \geq 0$ such that π^i satisfies g_2 and for every $0 \leq j < i$ path π^j satisfies g_1 . Given a Kripke structure M , a state s , and a path π , we write $M, s \models f$ if state s of M satisfies state formula f , and $M, \pi \models g$ if path π of M satisfies path formulas g .

Consider again the Kripke structure of Figure 1. The initial state satisfies the formula $\mathbf{AG} \neg(\text{ring} \wedge \text{snooze})$, since no state is both *ringing* and *snoozing*. On the other hand, it does not satisfy the formula $\mathbf{E}[\neg \text{ring} \mathbf{U} \text{snooze}]$, since every path that reaches a state labeled by *snooze* has to go through a state labeled by *ring*.

CTL and **LTL** are sub-logics of **CTL***. **CTL** is obtained by requiring every temporal operator to be immediately preceded by a path quantifier and vice-versa. An **LTL** formula is a **CTL*** formula of the form $\mathbf{A} g$ where no path quantifiers appear in g .

Model Checking Algorithm

The *model checking problem* can be stated as follows: given a Kripke structure M and a temporal logic formula f , determine if $M, s \models f$ holds for every initial state $s \in I$. Different model checking algorithms use specific techniques to answer this question in an efficient way. In this section, we present two model checking algorithms, one for specifications expressed using **CTL** and one for specifications expressed using **LTL**.

CTL Model Checking. The algorithm [1], [9] assumes that the specification f is a **CTL** formula. Every **CTL** formula can be rewritten in terms of only **EX**, **EG**, **EU**, \neg , and \wedge . The algorithm labels each state with the sub-formulas of f that hold at that state. It is applied recursively on the structure of the formula. The base case corresponds to atomic propositions: states that satisfy an atomic proposition are labeled by it. If the formula is of the form $\neg f$, all states that are not labeled by f are labeled by $\neg f$. If the formula is of the form $f_1 \wedge f_2$, all states that are labeled by f_1 and f_2 are labeled by $f_1 \wedge f_2$. If the formula is of the form **EX** f , each predecessor of a state labeled by f is labeled by **EX** f . If the formula is of the form $\mathbf{E}[f_1 \mathbf{U} f_2]$, then all states that are labeled by f_2 and all predecessor of states labeled by f_2 or by $\mathbf{E}[f_1 \mathbf{U} f_2]$ that are themselves labeled by f_1 are labeled by $\mathbf{E}[f_1 \mathbf{U} f_2]$. In order to determine the states that satisfy **EG** f , the algorithm considers the sub-graph G^f made of the states of M labeled by f . Each state in a non-trivial strongly connected component of G^f is labeled by **EG** f . Moreover, every predecessor of a state labeled by **EG** f that is itself labeled by f is also labeled by **EG** f .

LTL Model Checking. The algorithm [10] assumes that the specification is an **LTL** formula, i.e., of the form $\mathbf{A} g$. Checking that the system satisfies the formula $\mathbf{A} g$ is equivalent to checking that it satisfies $\neg \mathbf{E} \neg g$. The algorithm constructs an automaton $B_{\neg g}$ that accepts the traces that satisfy $\neg g$. The automaton is composed with an automaton that accepts the traces of M . If the composition is empty, then M satisfies the specification $\mathbf{A} g$. Otherwise, any of the traces of the

composition is a counterexample.

The **LTL** model checking algorithm constructs a *Büchi automaton* [11] $B_{\neg g}$ that accepts the traces that satisfy $\neg g$. Different algorithms for doing this have been proposed [10], [12]. Given the Kripke structure M corresponding to the system, an equivalent *Büchi automaton* B_M is constructed. The model checking algorithm constructs, on-the-fly, the synchronous composition of the two automata and checks for the presence of accepting runs. In order to detect accepting runs, the algorithm performs a depth first search (DFS) on the graph corresponding to the composition. Every time an accepting state q is reached, a second DFS is started from that state to determine if state q can be reached from itself. If this is the case, the run, made of the states needed to reach state q from the initial state followed by an infinite number of repetitions of the states needed to reach state q from itself, is accepting. Therefore, it corresponds to a counterexample to the specification. If no accepting runs exist, we can conclude that every initial state satisfies the specification.

Symbolic Model Checking

One of the main limitations of model checking is the size of systems that can be verified. If the system is too large or the specification too complex, model checking might not terminate due to insufficient resources, e.g., running time or memory. This is known as the *state explosion problem*. One of the major breakthroughs in model checking has been the development of **symbolic model checking**, a technique that uses binary decision diagrams (BDDs) to represent sets of states and transitions. Model checking is performed directly on the BDD representations.

Reduced Ordered Binary Decision Diagrams, or ROBDDs, are an efficient data structure to represent Boolean functions [13]. Given a set of variables $V = \{x_1, \dots, x_n\}$, a *binary decision diagram* is a directed acyclic graph where each non-terminal vertex v has two successors $true(v)$ and $false(v)$ and is labeled by a variable $var(v)$, each terminal vertex is labeled by either *true* or *false*, and a vertex r is designed as the root of the binary decision diagram. Given a variable order \prec , a total order over the variables in V , a binary decision diagram is ordered if, given any two nonterminal vertices u and v such that v is a successor of u , we have that $var(u) \prec var(v)$. An ordered binary decision diagram is reduced if: (i) there is only a single terminal vertex for each label; (ii) there exists no two nonterminal vertices u and v such that $var(u) = var(v)$, $false(u) = false(v)$, and $true(u) = true(v)$; and (iii) for every nonterminal vertex v , $false(v) \neq true(v)$. In the following, we will use the term BDDs to refer to ROBDDs unless otherwise noted. BDDs provide a canonical representation of Boolean functions: given a set of variables $V = \{x_1, \dots, x_n\}$, a BDD represents a Boolean function $f(x_1, \dots, x_n)$. To determine the value of f for a given value of the inputs, start from the root node r . When at a nonterminal vertex v , if $var(v)$ is assigned the value *true*, move to the vertex corresponding to $true(v)$ otherwise move to $false(v)$. Proceed until a terminal vertex u is reached:

the value of the function is the label of u . Boolean operations, e.g., \wedge , \neg , and \exists , can be applied directly to BDDs.

Boolean functions can be used to represent Kripke structures. Given a Kripke structure M , we can define a boolean encoding of its states using a finite set of variables $V = \{x_1, \dots, x_n\}$. A set of states S can then be defined by its characteristic function $S(s)$ that evaluates to *true* if state $s \in S$ and *false* otherwise. We define a set of variables $V' = \{x'_1, \dots, x'_n\}$, called next-state variables. The transition relation of M can then be represented by the Boolean function $R(s, s')$ over $V \cup V'$ that evaluates to *true* if a transition from s to s' is possible.

A CTL model checking algorithm that operates on sets of states represented as BDDs was proposed in [14], [15]. As before, we need to consider only **EX**, **EG**, **EU**, \wedge , and \neg , as formulas can be rewritten to contain only these operators and atomic propositions. The set of states labeled by an atomic proposition can be represented as a BDD. Boolean operators can be represented by the corresponding Boolean operations on BDDs. Given a BDD representing the set of states satisfying the formula f , the BDD corresponding to **EX** f can be obtained by computing $ex(s) = \exists s' : f(s') \wedge R(s, s')$, where each operation can be performed directly on BDDs. The **EG** and **EU** operators cannot be computed directly. However, **EG** f can be defined as the greatest fixed point of $eg(Z) = f \wedge \mathbf{EX} Z$ and $\mathbf{E}[f_1 \mathbf{U} f_2]$ as the least fixed point of $eu(Z) = f_2 \vee (f_1 \wedge \mathbf{EX} Z)$. Since the set of states is finite, and $eg(Z)$ and $eu(Z)$ are monotonic functions, the least and greatest fixed point of these functions can be computed iteratively. For instance, for computing $\mathbf{E}[f_1 \mathbf{U} f_2]$, we start with Q being equal to the empty set of states. At each iteration, the states that satisfy f_2 and the states that have a successor belonging to Q and also satisfy f_1 are added to Q . This process continues until no new states can be added to Q . The set Q at the last iteration is equal to the set of states that satisfy $\mathbf{E}[f_1 \mathbf{U} f_2]$. If the BDDs representing the set of states satisfying f_1 and f_2 are given, all operations can be performed directly using BDDs.

The main advantage of symbolic model checking is that the BDD representing a set of states can be much smaller than the set it represents. BDD-based symbolic model checkers have been used to check systems with a number of states that is beyond the reach of ordinary model checkers [14]. However, BDD-based symbolic model checking is not the definitive solution to the state explosion problem: while on average BDDs provide a compact representation of Boolean functions, there are Boolean functions whose representations as BDDs are exponential in the number of variables, which may make symbolic model checking impractical.

Bounded Model Checking

Boolean satisfiability (SAT), the typical example of an NP-complete problem, has been the focus of a lot of attention in recent years. SAT solvers have been developed that are able to handle problems with a large number of Boolean variables [16]. Most of modern SAT solvers take as input

a Boolean formula in conjunctive normal form (CNF) and are based on the Davis-Putnam-Logemann-Loveland (DPLL) algorithm [17], [18]. If the formula is satisfiable, they produce a satisfying assignment, if it is not, they produce a proof of unsatisfiability.

In the context of model checking, SAT solvers have become popular since the introduction of bounded model checking (BMC) [19]. The main idea of BMC is to encode the possible counterexamples to a given specification as a SAT formula. Since the formula must be finite, the counterexample needs to be finite as well. Specifically, BMC generates a formula that encodes the existence of a counterexample of a given length k .

For simplicity, we will consider only safety specifications of the form **AG** p , where p is an atomic proposition. Extensions that are able to handle more complex specifications are present in the literature [19]. Given an encoding of the states, we define three boolean predicates: $I(s)$, $R(s, s')$, and $F(s)$. Predicate $I(s)$ is *true* if s is an initial state. Predicate $R(s, s')$ is *true* if s' is a successor of s . Predicate $F(s)$ is *true* if state s is labeled by atomic proposition p . The formula corresponding to a counterexample of length k is:

$$I(s_0) \wedge R(s_0, s_1) \wedge \dots \wedge R(s_{k-1}, s_k) \wedge F(s_k) \quad (1)$$

which is satisfiable if and only if there is a counterexample of length k . By using a SAT solver it is possible to check if a counterexample exists. Moreover a counterexample can be obtained by analyzing the satisfying assignment produced by the SAT solver. If the formula is unsatisfiable, then there are no counterexamples of length k .

If no counterexamples of length k are present, it is necessary to repeat the check for a greater length. The procedure continues until either a counterexample is found or the completeness threshold is reached. The completeness threshold [20] is the minimum length such that, if the specification is violated, there exists a counterexample shorter than that length. Bounds on the completeness threshold of various classes of specifications have been given in the literature [19], [21], [22]. However, in practice, the computed bounds are often quite large. In this case, the verification terminates when the problem becomes intractable, without being able to prove that the system satisfies the specification.

The approach as presented is, therefore, inherently incomplete, unless tight bounds on the completeness threshold can be determined. However, complete methods based on bounded model checking that rely on alternative methods to determine termination have been proposed [23], [24], [25].

III. SOFTWARE AND BEYOND

The previous section described some of the seminal work in model checking. One of the application domains in which model checking has seen most successes is hardware: model checkers are currently used by many semiconductor manufacturers [3]. One of the reasons is that hardware designs are well suited for model checking: they are defined by using

Boolean gates and their semantics is straightforward. At the same time, historic events have provided the circumstances for technology transfer. For instance, soon after Intel had to recall a large number of Pentium processors because of a design bug, researchers were able to show that the bug could have been detected by using formal verification [26]. Since then, many chip design companies have been using model checking and other formal verification techniques.

Software Model Checking. More recently, software has been the focus of much effort in the model checking community. Tools like SLAM [27], BLAST [28], MAGIC [29], and CBMC [30], just to cite a few, have been developed that are aimed at the verification of software. Some of the techniques for software model checking that have been very successful are *predicate abstraction* [31] and *counterexample-guided abstraction refinement* [32], [33]. These techniques have made model checking of software feasible.

Infinite-State Systems. While the state space of a software may be very large, under certain conditions, it is still finite. This is not true of all types of systems. In recent years, efforts have been made to address formal verification of *infinite state systems*. In general, model checking cannot be applied to an infinite state system directly, because model checking, in its most basic form, would enumerate the states of the system and therefore never terminate. Techniques have been developed that allow model checking infinite state systems. In the following, we will briefly describe two important examples: (i) *timed systems*; and (ii) *hybrid systems*.

Timed Systems. In some cases, correctness of a system depends on the *exact timing of events*. As a consequence, models must include the time at which events occur. A commonly used formalism to model and reason about timed systems is *timed automata* [34], [35]. Timed automata are an extension of finite state automata that define a set of *real-valued clock variables*. The state space of a timed automaton can be infinite as the clocks assume values from the reals. Specialized algorithms and data structures have been developed that enable model checking of timed automata [36], [37], [38].

Hybrid Systems. Another example of infinite state systems, *hybrid systems* are characterized by the presence of *discrete and continuous components*. The continuous components are usually defined using differential equations. Most techniques for model checking hybrid systems represent sets of continuous states using specialized data structures. Polyhedra, defined as the intersection of a set of half-spaces, are a typical example of such data structures. Operations on these data structures are usually quite expensive and they introduce over-approximations in order to guarantee that the sets that are obtained can be represented by the chosen data structure.

Timed and Hybrid Software. Most of the existing approaches for model checking of timed and hybrid systems focus on the infinite-state components. However, when looking at a complex piece of software that includes this type of component, software model checking techniques are needed to

make verification feasible. We believe that techniques that apply existing software model checking to models that include time and continuous dynamics are crucial in making model checking scale to large, complex systems.

In the following, we will present two approaches, one for timed systems and one for hybrid systems, that allow applying software verification techniques to systems that include time or continuous dynamics.

Timed Automata Verification

While real-time systems are, in general, infinite-state systems because time is a continuous variable, under certain assumptions, it is possible to reduce the problem to finite-state model checking. In recent work [39], we have investigated techniques that, by using a well known mapping from infinite-state timed automata to finite-state *region automata*, can leverage the recent advances in model checking. In particular, we have developed abstraction and optimization techniques to reduce the size of the state space that needs to be visited by the model checker. We introduced a new abstraction technique, called GOABSTRACTION, which aims at reducing the size of the state space while preserving the validity of interesting specifications.

As for future work, we are developing technique to handle some of the other sources of the state explosion problem for this class of systems. We are currently looking at techniques that deal with large constants in the constraints and at ways of applying counterexample-guided abstraction refinement [33] to find the best abstraction that is able to prove the specification at hand.

Hybrid Systems Verification

Another area of active research is the verification of hybrid systems, i.e., systems in which discrete and continuous components coexist. One of the challenges is that the continuous components give rise to an infinite set of possible states.

The approach we proposed in [40] focuses on control software, a particular kind of software that interacts with a continuous environment. Very often such software is made of a set of periodic tasks, that are executed on a fixed schedule. In our approach, instead of looking at sets of continuous states, we look at continuous states individually. Given the amount of time during which the continuous system is evolving, the period of the tasks, the continuous state will evolve following the differential equations. If the differential equations are deterministic and time invariant, we can use numerical simulation algorithms to compute the continuous state that is reached after time elapses. One of the advantages is that numerical simulation algorithms are quite efficient and can be applied to a large class of differential equations. It is then possible to alternate continuous transitions —implemented by numerical simulation— and discrete transitions —corresponding to statements in the software— to perform model checking. Experimental results on this technique are promising. This approach, however, introduces approximations in two places: the numerical simulation introduces some numerical error;

and, in order to guarantee termination, it is necessary to introduce a maximum precision that is used to compare the continuous parts of states.

As for future work, we would like to determine in which cases we are able to guarantee correctness, i.e., avoid the approximation needed to guarantee termination. We would like to be able to evaluate the amount of imprecision introduced, as well as define a refinement step that increases the precision where it is most useful.

Acknowledgments: This research was sponsored by the National Science Foundation under grant nos. CNS-0411152, CCF-0429120, CCR-0121547, and CCR-0098072, the US Army Research Office under grant no. DAAD19-01-1-0485, the Office of Naval Research under grant no. N00014-01-1-0796, the Defense Advanced Research Projects Agency under subcontract no. SA423679952, the General Motors Corporation, and the Semiconductor Research Corporation. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government, or any other entity.

REFERENCES

- [1] E. M. Clarke and E. A. Emerson, "Synthesis of Synchronization Skeletons for Branching Time Temporal Logic," in *Proc. of Workshop on Logic of Programs*, 1981.
- [2] J. P. Queille and J. Sifakis, "Specification and Verification of Concurrent Systems in CESAR," in *Proc. of the 5th International Symposium on Programming*, 1982.
- [3] E. M. Clarke and J. M. Wing, "Formal Methods: State of the Art and Future Directions," *ACM Computing Surveys*, vol. 28, no. 4, pp. 626–643, 1996.
- [4] E. M. Clarke, "The Birth of Model Checking," Carnegie Mellon University, Tech. Rep. CMU-CS-TR-110, 2007.
- [5] A. Pnueli, "The Temporal Logic of Programs," in *Proc. of the 18th Annual Symposium on Foundations of Computer Science*, 1977.
- [6] M. Ben-Ari, A. Pnueli, and Z. Manna, "The Temporal Logic of Branching Time," *Acta Informatica*, vol. 20, pp. 207–226, 1983.
- [7] A. Pnueli, "The Temporal Semantics of Concurrent Programs," *Theoretical Computer Science*, vol. 13, pp. 45–60, 1981.
- [8] E. A. Emerson and J. Y. Halpern, "'Sometimes' and 'Not Never' Revisited: On Branching versus Linear Time Temporal Logic," *Journal of the ACM*, vol. 33, no. 1, pp. 151–178, 1986.
- [9] E. Clarke, E. A. Emerson, and A. P. Sistla, "Automatic Verification of Finite-State Concurrent Systems Using Temporal Logic Specifications," *ACM Transactions on Programming Languages and Systems*, vol. 8, no. 2, pp. 244–263, 1986.
- [10] M. Y. Vardi and P. Wolper, "An Automata-Theoretic Approach to Automatic Program Verification," in *Proc. of the 1st Symposium on Logic in Computer Science*, 1986.
- [11] W. Thomas, "Automata on Infinite Objects," *Handbook of Theoretical Computer Science*, vol. B, pp. 133–191, 1990.
- [12] R. Gerth, D. Peled, M. Y. Vardi, and P. Wolper, "Simple On-The-Fly Automatic Verification of Linear Temporal Logic," in *Proc. of the 15th International Symposium on Protocol Specification, Testing and Verification*, 1995.
- [13] R. E. Bryant, "Graph-Based Algorithms for Boolean Function Manipulation," *IEEE Transactions on Computers*, vol. 35, no. 8, pp. 667–691, 1986.
- [14] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang, "Symbolic Model Checking: 10^20 States and Beyond," in *Proc. of the 5th Annual Symposium on Logic in Computer Science*, 1990.
- [15] K. L. McMillan, "Symbolic Model Checking: An Approach to the State Explosion Problem," Ph.D. dissertation, Carnegie Mellon University, 1992.
- [16] M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik, "Chaff: Engineering an Efficient SAT Solver," in *Proc. of the 38th Design Automation Conference*, 2001.
- [17] M. Davis and H. Putnam, "A Computing Procedure for Quantification Theory," *Journal of the ACM*, vol. 7, no. 3, pp. 201–215, 1960.
- [18] M. Davis, G. Logemann, and D. Loveland, "A Machine Program for Theorem-Proving," *Communications of the ACM*, vol. 5, no. 7, pp. 394–397, 1962.
- [19] A. Biere, A. Cimatti, E. M. Clarke, O. Strichman, and Y. Zhu, "Bounded Model Checking," *Advances in Computers*, vol. 58, pp. 118–149, 2003.
- [20] A. Biere, A. Cimatti, E. M. Clarke, and Y. Zhu, "Symbolic Model Checking without BDDs," in *Proc. of the 5th International Conference on Tools and Algorithms for Construction and Analysis of Systems*, 1999.
- [21] D. Kroening and O. Strichman, "Efficient Computation of Recurrence Diameters," in *Proc. of the 4th International Conference on Verification, Model Checking, and Abstract Interpretation*, 2003.
- [22] E. M. Clarke, D. Kroening, J. Ouaknine, and O. Strichman, "Completeness and Complexity of Bounded Model Checking," in *Proc. of the 5th International Conference on Verification, Model Checking, and Abstract Interpretation*, 2004.
- [23] M. Sheeran, S. Singh, and G. Stalmarck, "Checking Safety Properties Using Induction and a SAT-Solver," in *Proc. of 3rd International Conference on Formal Methods in Computer-Aided Design*, 2000.
- [24] K. L. McMillan, "Interpolation and SAT-Based Model Checking," in *Proc. of the 15th International Conference Computer Aided Verification*, 2003.
- [25] M. R. Prasad, A. Biere, and A. Gupta, "A Survey of Recent Advances in SAT-based Formal Verification," *International Journal on Software Tools for Technology Transfer*, vol. 7, no. 2, pp. 156–173, 2005.
- [26] E. M. Clarke, S. M. German, and X. Zhao, "Verifying the SRT Division Algorithm using Theorem Proving Techniques," in *Proc. of the 8th International Conference on Computer Aided Verification*, 1996.
- [27] T. Ball and S. K. Rajamani, "Bebop: A Symbolic Model Checker for Boolean Programs," in *Proc. of the 7th International SPIN Workshop*, 2000.
- [28] T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre, "Lazy Abstraction," in *Proc. of the 29th Symposium on Principles of Programming Languages*, 2002.
- [29] S. Chaki, E. M. Clarke, A. Groce, S. Jha, and H. Veith, "Modular Verification of Software Components in C," in *Proc. of the 25th International Conference on Software Engineering*, 2003.
- [30] E. M. Clarke, D. Kroening, and K. Yorav, "Behavioral Consistency of C and Verilog Programs using Bounded Model Checking," in *Proc. of the 40th Design Automation Conference*, 2003.
- [31] S. Graf and H. Saidi, "Construction of Abstract State Graphs with PVS," in *Proc. of the 9th International Conference on Computer Aided Verification*, 1997.
- [32] R. P. Kurshan, *Computer-Aided Verification of Coordinating Processes: The Automata-Theoretic Approach*. Princeton University Press, 1994.
- [33] E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith, "Counterexample-Guided Abstraction Refinement," in *Proc. of the 12th International Conference Computer Aided Verification*, 2000.
- [34] R. Alur, C. Courcoubetis, and D. L. Dill, "Model-Checking for Real-Time Systems," in *Proc. of the 5th Annual IEEE Symposium on Logic in Computer Science*, 1990.
- [35] R. Alur and D. L. Dill, "A Theory of Timed Automata," *Theoretical Computer Science*, vol. 126, no. 2, pp. 183–235, April 1994.
- [36] S. Yovine, "KRONOS: a verification tool for real-time systems," *International Journal on Software Tools for Technology Transfer*, vol. 1, no. 1-2, pp. 123–133, December 1997.
- [37] F. Wang, "RED: Model-Checker for Timed Automata with Clock-Restriction Diagram," in *Proc. of Workshop on Real-Time Tools*, 2001.
- [38] K. G. Larsen, P. Pettersson, and W. Yi, "Compositional and Symbolic Model-Checking of Real-Time Systems," in *Proc. of the 16th IEEE Real-Time Systems Symposium*, 1995.
- [39] E. M. Clarke, F. Lerda, and M. Talupur, "An Abstraction Technique for Real-Time Verification," in *Proc. of the GM R&D Workshop on Next Generation Design and Verification Methodologies for Distributed Embedded Control System*, 2007.
- [40] S. Scherer, F. Lerda, and E. Clarke, "Model Checking of Robotic Control Systems," in *Proc. of the 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, 2005.