Efficient Proof Engines for Bounded Model Checking of Hybrid Systems

Martin Fränzle [1,3]

Informatics and Mathematical Modelling, The Technical University of Denmark,
Richard Petersens Plads, Bldg. 322, DK-2800 Kgs. Lyngby, Denmark

Christian Herde [2,3]

Department of Computing Science, Carl-von-Ossietzky Universität Oldenburg,
P.O. Box 2503, D-26111 Oldenburg, Germany

Abstract

In this paper we present HySat, a new bounded model checker for linear hybrid systems, incorporating a tight integration of a DPLL–based pseudo–Boolean SAT solver and a linear programming routine as core engine. In contrast to related tools like MathSAT, ICS, or CVC, our tool exploits all of the various optimizations that arise naturally in the bounded model checking context, e.g. isomorphic replication of learned conflict clauses or tailored decision strategies, and extends them to the hybrid domain. We demonstrate that those optimizations are crucial to the performance of the tool.

> *Key words:* verification, bounded model checking, hybrid systems, infinite-state systems, decision procedures, satisfiability.

1 Introduction

During the last ten years, formal verification of digital systems has evolved from an academic subject to an approach accepted by the industry, with dozens of commercial tools now available and used by major companies. Among the most successful methods in formal verification of discrete systems is bounded model checking (BMC), as suggested by Groote et al. in [17] and by Biere et al. in [9]. The idea of BMC is to encode the next–state relation of a system as

---

a propositional formula, unroll this to some given finite depth k, and to augment it with a corresponding finite unravelling of the tableau of (the negation of) a temporal formula in order to obtain a propositional SAT problem which is satisfiable iff an error trace of length k exists. Enabled by the impressive gains in performance of propositional SAT checkers in recent years, BMC can now be successfully applied even to very large finite-state designs.

Though originally formulated for discrete transition systems only, the basic idea of BMC to reduce the search for an error path to a satisfiability problem of a formula also applies to hybrid discrete–continuous systems. However, the BMC formulae arising from such systems are no longer purely propositional, but usually comprise complex Boolean combinations of arithmetic constraints over real-valued variables, thus entailing the need for new decision procedures to solve them.

Our tool HySat provides a decision procedure that is tailored to fit the needs of BMC of infinite–state systems with piecewise linear variable updates, e.g. of linear hybrid automata. HySat tightly integrates a state–of–the–art Davis–Putnam style SAT solver for pseudo–Boolean constraints with a linear programming routine, combining the virtues of both methods: Linear programming adds the capability of solving large conjunctive systems of linear inequalities over the reals, whereas the SAT solver accounts for fast Boolean search and efficient handling of disjunctions.

The idea to combine algorithms for SAT with decision procedures for conjunctions of numerical constraints in order to solve arbitrary Boolean combinations thereof has been pursued by several groups. A tight integration of a resolution based SAT checker with linear programming has first been proposed and successfully applied to planning problems by Wolfman and Weld [30]. More recently, Audemard et al. [3] have followed up with MathSAT, a tool combining SAT solving with a Bellman–Ford algorithm for difference logic constraints and a simplex algorithm for general linear constraints, used for applications in the context of temporal reasoning and model checking of timed automata. Tools supporting a more general class of formulae are CVC [6] and ICS [15], both integrating decision procedures for various theories, including Boolean logic, linear real arithmetic, uninterpreted function symbols, functional arrays, and abstract data types.

However, except for HySat, all tools mentioned above lack some or all of the particular optimizations that arise naturally in the bounded model checking context. As observed by Shtrichman [26], BMC yields SAT instances that are highly symmetric as they comprise a k–fold unrolling of the systems transition relation. This special structure can be exploited to accelerate solving, e.g. by copying the explanation for a conflict which was encountered during the backtrack search performed by the SAT solver, to all isomorphic parts of the formula in order to prune similar conflicts from the search tree. This technique, in the following referred to as *isomorphy inference*, has been shown to yield considerable performance gains when performing BMC with propo-

sitional SAT engines. To the best of our knowledge, HySat is the first solver that extends isomorphy inference accross transitions, as well as other domain–specific optimizations described in [26], to the hybrid domain. We will show that, compared to purely propositional BMC, similar or even higher performance gains can be accomplished within this context. The reason is that an inference step in the hybrid domain is computationally much more expensive than in propositional logic, as now richer logics have to be dealt with.

The paper is organized as follows. In the following two sections we explain the logical language solved by our SAT checker and review briefly how a linear hybrid automaton can be translated into a predicative formula suitable for bounded model checking. In section 4 we explain in detail the algorithmic ingredients of HySat. In particular, we discuss the BMC–specific optimizations implemented in our tool. In section 5 we report some experimental results, and section 6 draws conclusions and describes directions for future research.

2   The logics

As we are aiming at automated state-exploratory analysis of linear hybrid automata [20,19] without prior finite-state abstraction, HySat addresses satisfiability problems in a two-sorted logics entailing Boolean-valued and real-valued variables. When encoding properties of linear hybrid automata, the Boolean variables are used for encoding the discrete state components, while the real variables represent the continuous state components.

The formulae are actually propositional, being conjunctions of *linear zero-one constraints* [16] (also known as *pseudo-Boolean constraints* [7]) for the Boolean part and of *guarded linear constraints* [30] for the real-valued part:

$$formula ::= \{\, clause \land \}^* clause$$
$$clause ::= linear\_ZO\_constraint \mid boolean\_var \implies linear\_constraint$$

Here, *linear_constraint* denotes a conjunction of linear inequalities over *real-valued* variables, i.e. the constraint part of an arbitrary linear program, while *linear_ZO_constraint* denotes a linear inequality over *Boolean-valued* variables. The reason for using linear zero-one constraint clauses instead of, e.g., disjunctive clauses (like in conjunctive normal forms) is that linear zero-one constraints are much more concise than disjunctive clauses and that we have a very efficient SAT solver —called "Goblin" [16]— for such constraint systems, yielding the base engine for HySat.

2.1   Zero-one linear constraints

Rewriting arbitrary propositional formulae to conjunctive normal form (CNF) yields a worst-case exponential blowup in formula size if the number of propositional variables is to be preserved. To avoid this, all practical verification environments take advantage of satisfiability-preserving transformations that yield linear-size encodings through introduction of a linear number of auxil-

iary variables [27,25,28]. The price for introducing a linear number of auxiliary variables is, however, a worst-case exponential blow-up in the size of the search tree upon backtrack search. Yet, it has been observed that both causes of blow-up can often be avoided, as the Davis-Putnam-Loveland-Logemann search procedure for satisfying valuations generalizes smoothly to zero-one linear constraint systems (ZOLCS), which are the constraint parts of zero-one linear programs [7,29,2,16]. Zero-one linear constraint systems are expressive enough to facilitate a linear-size encoding of, e.g., gate-level netlists without use of auxiliary variables.

In a *zero-one linear constraint system* or *linear pseudo-Boolean constraint systems*, formulae are conjunctions of linear zero-one constraints. A *linear zero-one constraint* is of the form $a_1 x_1 + a_2 x_2 + \ldots a_n x_n \geq k$, where the x_i are *literals*, i.e. positive or negated *propositional variables*, the a_i are natural numbers, called the *weights* of the individual literals, and $k \in \mathbb{N}$ is the *threshold*.

Given a Boolean valuation of the propositional variables, a zero-one constraint is satisfied iff its left hand side evaluates to a value exceeding the threshold when the truth values `false` and `true` of the literals are identified with 0 and 1, respectively. Zero-one constraints can represent a wide class of monotonic Boolean functions, e.g. $1a + 1b + 1\overline{c} + 1d \geq 1$ is equivalent to $a \vee b \vee \overline{c} \vee d$, $1a + 1b + 1\overline{c} + 1d \geq 4$ is equivalent to $a \wedge b \wedge \overline{c} \wedge d$, and $1a + 1b + 3\overline{c} + 1d \geq 3$ is equivalent to $c \implies (a \wedge b \wedge d)$. Consequently, ZOLCS can be exponentially more concise than CNF: a CNF expressing that at least n out of k variables should be true requires $\binom{n}{k}$ disjunctive clauses of length n each, i.e. is of size $O\left(\binom{n}{k}n\right)$, whereas the corresponding ZOLCS has size linear in k and logarithmic in n.

Formally, the syntax of linear zero-one constraints is

$$
\begin{aligned}
linear\_ZO\_constraint &::= linear\_term \geq threshold \\
linear\_term &::= \{weight\ literal\ +\}^*weight\ literal \\
weight &::\in \mathbb{N} \\
literal &::= boolean\_var \mid \overline{boolean\_var} \\
boolean\_var &::\in BV \\
threshold &::\in \mathbb{N}
\end{aligned}
$$

where BV is a countable set of Boolean variable names.

Zero-one constraints are interpreted over *Boolean valuations* $\sigma_\mathbb{B} : BV \xrightarrow{\text{total}} \mathbb{B}$ of the propositional variables. $\sigma_\mathbb{B}$ satisfies a constraint $a_1 x_1 + a_2 x_2 + \ldots a_n x_n \geq k$ iff $a_1 \chi_{\sigma_\mathbb{B}}(x_1) + a_2 \chi_{\sigma_\mathbb{B}}(x_2) + \ldots a_n \chi_{\sigma_\mathbb{B}}(x_n) \geq k$, where

$$
\chi_{\sigma_\mathbb{B}}(x) = \begin{cases} 0 & \text{if } x \in V \text{ and } \sigma_\mathbb{B}(x) = \texttt{false}, \\ 1 & \text{if } x \in V \text{ and } \sigma_\mathbb{B}(x) = \texttt{true}, \\ 1 - \chi_{\sigma_\mathbb{B}}(y) & \text{if } x \equiv \overline{y} \text{ for some } y \in V. \end{cases}
$$

2.2 Guarded linear constraints

Zero-one constraints can only express constraints on Boolean variables. A second kind of clauses in our logics is *Boolean-guarded linear constraints* which express (linear) constraints between real-valued variables, as well as their interdependence with the Boolean valuation. A guarded linear constraint simply is an implication

$$boolean\_var \implies linear\_constraint$$

between a Boolean variable and a linear constraint over real-valued variables, i.e. a conjunction of linear inequations. Such a guarded linear constraint is interpreted over a valuation $\sigma = (\sigma_\mathbb{B}, \sigma_\mathbb{R}) \in (BV \xrightarrow{\text{total}} \mathbb{B}) \times (RV \xrightarrow{\text{total}} \mathbb{R})$, where RV is the set of real variables occurring in linear constraints. The guarded linear constraint $v \implies c$ is satisfied by $\sigma = (\sigma_\mathbb{B}, \sigma_\mathbb{R})$ iff $\sigma_\mathbb{R}$ satisfies the linear constraint c or if $\sigma_\mathbb{B}(v) = \texttt{false}$.

2.3 Satisfaction of formulae

A *formula* ϕ is a conjunction of linear zero-one constraints and of guarded linear constraints and is thus interpreted over valuations

$$\sigma = (\sigma_\mathbb{B}, \sigma_\mathbb{R}) \in (BV \xrightarrow{\text{total}} \mathbb{B}) \times (RV \xrightarrow{\text{total}} \mathbb{R}) \ .$$

Obviously, ϕ is satisfied by $\sigma = (\sigma_\mathbb{B}, \sigma_\mathbb{R})$, denoted $\sigma \models \phi$, iff all linear zero-one constraints in ϕ are satisfied by $\sigma_\mathbb{B}$ and all guarded linear constraints in ϕ are satisfied by $(\sigma_\mathbb{B}, \sigma_\mathbb{R})$.

When solving satisfiability problems of formulae with Davis-Putnam-like procedures, we will build valuations incrementally such that we have to reason about *partial valuations* $\rho \in (BV \xrightarrow{\text{part.}} \mathbb{B}) \times (RV \xrightarrow{\text{part.}} \mathbb{R})$ of variables. We say that a variable $v \in BV \cup RV$ is *unassigned in* ρ iff $v \notin \text{dom}(\rho_\mathbb{B}) \cup \text{dom}(\rho_\mathbb{R})$. A partial valuation ρ is called *consistent for a formula* ϕ iff there exists a total extension $\sigma : (BV \xrightarrow{\text{total}} \mathbb{B}) \times (RV \xrightarrow{\text{total}} \mathbb{R})$ of ρ that satisfies ϕ. Otherwise, we call ρ *inconsistent for* ϕ. Furthermore, a partial valuation ρ is said to *satisfy* ϕ iff all its total extensions satisfy ϕ. As this definition of satisfaction agrees with the previous one on total valuations, we will use the same notation $\rho \models \phi$ for satisfaction by partial and by total valuations.

3   Predicative encoding of linear hybrid automata

A *linear hybrid automaton* $A = (\Sigma, T, R, inv, l, u, m, g, ass, init)$, as depicted in Fig. 1, consists of

- a finite set Σ of *locations*,
- a finite set T of *transitions*,
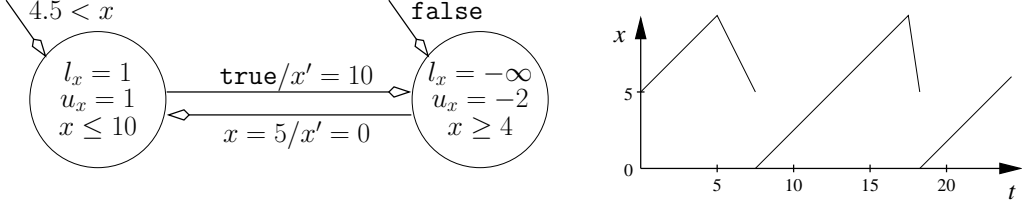- a finite set R of continuous state components,

Fig. 1. A linear hybrid automaton and a sample trajectory. l_x and u_x denote the lower and upper bounds on the slope of x in the corresponding states, while $x \leq 10$ and $x \geq 4$ are state invariants constraining x itself.

- a family $inv = (inv_\sigma)_{\sigma \in \Sigma}$ of *state invariants*, where each state invariant inv_σ is a linear predicate over R which constrains the valuations of the continuous state components when control resides in the discrete location σ,
- two families $l = (l_{\sigma,x})_{\sigma \in \Sigma, x \in R}$ and $u = (u_{\sigma,x})_{\sigma \in \Sigma, x \in R}$ assigning two each location $\sigma \in \Sigma$ and each continuous state component $x \in R$ the *minimum and maximum slope* of x while control resides in location σ. The individual $l_{\sigma,x}$ are constants in $\mathbb{Q} \cup \{-\infty\}$ and similarly $u_{\sigma,x} \in \mathbb{Q} \cup \{\infty\}$.
- a mapping $m : T \xrightarrow{\text{total}} \Sigma^2$ assigning to each transition the pair of source and sink state of the transition,
- a family $g = (g_t)_{t \in T}$ assigning to each transition a *transition guard* enabling that transition, where the transition guard is a linear predicate over R,
- a family $ass = (ass_t)_{t \in T}$ assigning to each transition a (possibly nondeterministic) *assignment* which is a linear predicate over R and R', where R' denotes primed variants of the state components in R. The interpretation is that undecorated state components $x \in R$ refer to the state immediately before the transition, while the primed variant $x' \in R'$ refers to the state immediately thereafter.
- a family $init = (init_\sigma)_{\sigma \in \Sigma}$ of *initial state predicates*, where each $init_\sigma$ is a linear predicate over R which constrains the valuations of the continuous state components when control resides *initially* in the discrete location σ.[4]

Hybrid automata engage in an alternation of continuous evolutions and discrete transitions. A *continuous evolution* of $A = (\Sigma, T, R, inv, l, u, m, g, ass, init)$ can be represented by a tuple $(\sigma, \boldsymbol{x}, \delta, \boldsymbol{x}')$ consisting of a discrete state $\sigma \in \Sigma$ the automaton resides in, a source continuous state $\boldsymbol{x} \in (R \xrightarrow{\text{total}} \mathbb{R})$ and a target continuous state $\boldsymbol{x}' \in (R \xrightarrow{\text{total}} \mathbb{R})$, as well as a duration $\delta \in \mathbb{R}_{\geq 0}$. Such a tuple is a continuous evolution of A iff for each $y \in R$ it holds that $\boldsymbol{x}'(y) \geq \boldsymbol{x}(y) + l_{\sigma_i,y} \cdot \delta$ and $\boldsymbol{x}'(y) \leq \boldsymbol{x}(y) + u_{\sigma_i,y} \cdot \delta$, and both $\boldsymbol{x}$ and $\boldsymbol{x}'$ satisfy inv_σ. Thus, δ represents the duration of A residing in state σ, and all continuous variables y evolve according to their slope bounds, and the invariant is true in the start and the end state (and thus, by convexity, in between). Similarly, an *immediate transition* can be represented by a tuple $(\sigma, \boldsymbol{x}, \sigma', \boldsymbol{x}')$ consisting of a discrete source state $\sigma \in \Sigma$ and a discrete target state σ', plus a continuous

---

[4] A discrete location σ not to be taken initially takes the predicate $init_\sigma = \texttt{false}$.

source state $\boldsymbol{x} \in (R \xrightarrow{\text{total}} \mathbb{R})$ and a continuous target state $\boldsymbol{x}' \in (R \xrightarrow{\text{total}} \mathbb{R})$. Such a tuple is an immediate transition iff there is a transition $t \in T$ with $m(t) = (\sigma, \sigma')$ such that $\boldsymbol{x}$ satisfies g_t and such that ass_t is satisfied if $\boldsymbol{x}$ is substituted for the variables in R and $\boldsymbol{x}'$ is substituted for the variables in R'.

A *run* $r = \langle(\sigma^0, \boldsymbol{x}^0, \delta^0, \boldsymbol{x}'^0), \ldots, (\sigma^n, \boldsymbol{x}^n, \delta^n, \boldsymbol{x}'^n)\rangle \in (\Sigma \times (R \xrightarrow{\text{total}} \mathbb{R}) \times \mathbb{R}_{\geq 0} \times (R \xrightarrow{\text{total}} \mathbb{R}))^*$ is a sequence of continuous evolutions of A linked by immediate transitions and grounded in a viable initial state. I.e., a run r satisfies the following properties:

- *Initialization:* $\boldsymbol{x}^0$ satisfies $init_{\sigma^0}$.
- *Progression by continuous evolution:* for all i, the tuple $(\sigma^i, \boldsymbol{x}^i, \delta^i, \boldsymbol{x}'^i)$ is a continuous evolution of A.
- *Progression by immediate transitions:* the tuple $(\sigma^i, \boldsymbol{x}'^i, \sigma^{i+1}, \boldsymbol{x}^{i+1})$ is an immediate transition of A for all $i < n$.

In order to perform *bounded model checking* (BMC) [9] with HySat, i.e. checking of validity of temporal properties on finite unrollings of a transition system, we need to encode all runs of a given length $k \in \mathbb{N}$ in HySat's logics. There are various ways of doing this, all with specific strengths and weaknesses. Yet all the reasonable ones share the property of featuring a plethora of structurally similar sub-formulae stemming from the iterated application of the transition relation and from the iterated continuous evolution in the k-fold unrolling. In order to exemplify this, we present here one particular form of such an unrolling which is very similar to the one used by Audemard et al. for MathSAT-based BMC of linear hybrid automata [4] and by Bemporad et al. for MILP-based BMC of linear hybrid automata [8].

Let $A = (\Sigma, T, R, inv, l, u, m, g, ass, init)$ be a linear hybrid automaton. In order to encode a transition sequence of A of some given length $k \in \mathbb{N}$, we proceed as follows:

(i) For each discrete state $\sigma \in \Sigma$ we take $k + 1$ Boolean variables σ^i, with $0 \leq i \leq k$. The value of σ^i encodes whether the automaton A is in state σ in step i. Here, we take "one-hot" encoding, i.e. $\sigma^i = \texttt{true}$ iff A is in state σ in step i. With one-hot encoding, there consequently is, for any $i \leq k$, exactly one $\sigma \in \Sigma$ such that σ^i holds, which is enforced in the BMC formula by the $2k + 2$ linear zero-one constraints

$$\bigwedge_{i=0}^{k} \left(\sum_{\sigma \in \Sigma} 1\sigma^i \leq 1 \right) \wedge \bigwedge_{i=0}^{k} \left(\sum_{\sigma \in \Sigma} 1\overline{\sigma^i} \geq |\Sigma| - 1 \right)$$

(ii) For each transition $t \in T$ we take k Boolean variables t^i, with $1 \leq i \leq k$. The value of t^i encodes via one-hot encoding whether the ith move in the run is transition t. Wellformedness of the unrolling in the sense that exactly one transition is taken in each step is guaranteed by conjunctively

adding the $2k$ linear zero-one constraints

$$\bigwedge_{i=1}^{k} \left(\sum_{t \in T} 1 t^i \leq 1 \right) \wedge \bigwedge_{i=1}^{k} \left(\sum_{t \in T} 1 \overline{t^i} \geq |T| - 1 \right)$$

to the formula.

(iii) For each continuous state component $x \in R$ we take $k + 1$ real-valued variables x^i and another $k + 1$ real-valued variables x'^i, with $i \leq k$. The value of x^i encodes the value of x immediately after the ith transition in the run, whereas x'^i represents the value immediately before transition $(i + 1)$. For each $i \leq k$ we do, furthermore, take one real-valued variable δ^i representing the time spent in the ith state of the run. This allows us to formalize the *continuous evolutions* by conjoining the guarded linear constraint

$$\sigma^i \implies (x'^i \geq x^i + l_{\sigma,x}\delta^i \wedge x'^i \leq x^i + u_{\sigma,x}\delta^i)$$

for each $\sigma \in \Sigma$ and each $i \leq k$ to the formula. [5] Furthermore, we have to keep track of the *state invariants*, which are enforced by the guarded linear constraints

$$\sigma^i \implies (inv_{\sigma^i}[x_1^i, \ldots, x_n^i/x_1, \ldots, x_n] \wedge inv_{\sigma^i}[x_1'^i, \ldots, x_n'^i/x_1, \ldots, x_n]) \ ,$$

where $\{x_1, \ldots, x_n\} = R$.

(iv) The interplay between discrete states and transitions requires that t^i implies σ^{i-1} and $\widetilde{\sigma}^i$ for $(\sigma, \widetilde{\sigma}) = m(t)$. With linear zero-one constraints, this can be expressed by a single constraint

$$2t^i + 1\sigma^{i-1} + 1\widetilde{\sigma}^i \geq 2$$

for each $t \in T$ and each $1 \leq i \leq k$. Furthermore, enabledness of the transition, i.e. validity of the *transition guard*, is enforced through the guarded linear constraint

$$t^{i+1} \implies g_t[x_1'^i, \ldots, x_n'^i/x_1, \ldots, x_n] \ .$$

Likewise, *assignments* are dealt with by

$$t^{i+1} \implies ass_t[x_1^i, \ldots, x_n^i/x_1, \ldots, x_n][x_1'^i, \ldots, x_n'^i/x_1', \ldots, x_n']$$

(v) Finally, we have to add constraints describing the allowable initial states through the guarded linear constraint system

$$\bigwedge_{\sigma \in \Sigma} \left(\sigma^0 \implies init_\sigma \right)$$

---

[5] If $l_{\sigma,x} = -\infty$ or $u_{\sigma,x} = \infty$, the corresponding part of the constraint is left out.

Satisfying valuations of the formula thus obtained are in one-to-one correspondence to the runs of A of length k. As in BMC [9], satisfaction of temporal properties on all runs of depth k can thus be checked by adding to the formula the k-fold unrolling of a tableaux of the (negated) property, then checking the resulting formula for unsatisfiability. Using standard techniques from predicative semantics [18], the translation scheme can be extended to both shared variable and synchronous message-passing parallelism, thereby yielding formulae of size linear in the number of parallel components.

Note that, except for step (v) of above encoding scheme, all steps generate multiple copies of the same basic formula, where the k or $k + 1$ individual copies differ just in a consistent renaming of the variables. Therefore, a satisfiability checker tailored towards BMC of hybrid automata should exploit such isomorphies between subformulae for accelerating satisfiability checking, which is the distinguishing feature of HySat. In order to simplify detection of isomorphic copies, HySat is in fact fed with just a single copy of the transition and evolution predicates and performs the unrolling itself.

4   Ingredients of HySat

The predicative encoding outlined above yields formulae which are Boolean combinations of linear arithmetic contraints. To deal with such formulae, HySat's main components are

- the *solver core*, consisting of a tight integration of a SAT solver with a linear programming routine, described in section 4.1, and enhanced with domain-specific optimizations for BMC, as explained in section 4.2,
- an *API* to the solver core, providing methods for formula generation, simplification, common subexpression eliminiation, and for rewriting the resulting formula into a conjunctive form, namely a conjunction of zero-one linear constraints and guarded linear constraints, which is the input format of the solver core,
- a *frontend*, consisting of HySat's input language and a bounded model checker, which performs the unwinding of the transition relation and controls the solver core via API calls.

To fit the needs of BMC, which involves checking the same system on different unrolling depths, the solver core and the API are designed to work in an incremental fashion in the sense that they allow to add (as well as delete) successively sets of constraints to (from) an existing problem and then redo the satisfiability check without starting SAT search from scratch each time.

4.1   Integration of DPLL-SAT and Linear Programming

Before addressing the integration of a propositional SAT solver with linear programming, we first briefly review some basics of the individual methods.

4.1.1 Boolean SAT

The best currently known procedures for deciding Boolean SAT problems implement variants of the classical Davis-Putnam-Loveland-Logemann (DPLL) procedure [12] and are based on backtracking in the space of partial value assignment. Given a Boolean formula Φ in conjunctive normal form (CNF) and a partial valuation ρ, which is empty at the start, the DPLL procedure incrementally extends ρ until either $\rho \models \phi$ holds or ρ turns out to be inconsistent for ϕ, in which case another extension is tried through backtracking. Extensions are constructed by performing *decision steps*, which entail selecting an unassigned variable "blindly" and assigning a truth-value to it, each followed by a *deduction phase*, involving the search for *propagating clauses* that enforce certain assignments to preserve satisfiability, where execution of the latter might cause the need for further such assignments, in this context also referred to as *implications*. However, deduction may also yield a *conflicting clause* which has all its literals assigned false, indicating the need for backtracking.

Like all pure backtracking algorithms, the classical DPLL procedure suffers from thrashing, i.e. repeated failure due to the same reason. To overcome this problem, modern SAT solvers implement a technique called *conflict-driven learning* [31], which attempts to derive sufficiently general reasons for conflicts being encountered and stores them for future guidance of the search. The standard scheme traces the reason back to a small (ideally minimal) number of assignments that triggered the particular conflict, and stores this reason by adding the negation of that assignment as as clause, termed *conflict clause*, to the clause database. Besides *learning*, state-of-the-art SAT solvers, as the one being integrated in HySat, enhance the basic DPLL procedure by sophisticated heuristics for selecting the assignment performed at decision steps [22,24], and add various algorithmic refinements, among them non-chronological backtracking [23,24], random restarts [5] and lazy clause evaluation [24], to accelerate the proof search.

A pecularity of HySat's SAT solver is its ability to directly handle linear zero-one constraint systems, a considerably more concise language than CNF.

4.1.2 Linear programming

Linear programming deals with finding extreme values of a linear function when the variables are constrained by linear (in)equalities, i.e. with problems that can be put in the general form

$$
\begin{aligned}
\text{maximize} \quad & \boldsymbol{c}^T \boldsymbol{x} \\
\text{subject to} \quad & \mathbf{A}\boldsymbol{x} \le \boldsymbol{b}
\end{aligned}
\tag{1}
$$

where $\boldsymbol{x}$ is the vector of variables to be solved for, and $\mathbf{A}$, $\boldsymbol{b}$ and $\boldsymbol{c}$ are given matrices or vectors of known coefficients. The linear expression $\boldsymbol{c}^T \boldsymbol{x}$ is called the *objective function*, (1) is referred to as a *linear program*.

HySat uses LP as a black-box method to decide the feasibility of a set of linear constraints, i.e. to check whether for a given system of inequalities $\mathbf{A}\boldsymbol{x} \leq \boldsymbol{b}$ the set of solutions $\{\boldsymbol{x} \in \mathbb{R} \mid \mathbf{A}\boldsymbol{x} \leq \boldsymbol{b}\}$ is non-empty. Linear programming is known to be polynomial. Commercial codes like CPLEX tackle instances with more than 10^6 variables. In HySat, however, we use the free LP library glpk[6] by Andrew Makhorin which provides a simplex solver, an interior point solver, and a solver supporting mixed integer linear programming (MILP), where some of the variables are required to be integer.

4.1.3   Coupling SAT and LP

The basic idea of the integration is to guard each non-propositional constraint occuring in the input formula with a new Boolean variable and to pass the corresponding constraint to the linear programming routine whenever the SAT solver assigns that variable to true. In turn, constraints are removed from the LP-solver's database when their guard variables are unassigned again due to backtracking.

After each deduction phase in which no Boolean conflict was encountered, the SAT solver checks if new constraints have been added to the linear programm since its last evaluation. If so, the linear programming routine is called to decide the feasibility of the set of constraints residing in its database. If the linear program turns out to be inconsistent, a conflict is reported to the SAT solver. Otherwise the SAT solver can proceed with the next decision step.

In case of a conflict, however, HySat invokes a conflict-analysis routine that extracts an *irreducible infeasible subsystem* (IIS) from the linear program[7], i.e. a subset of constraints which itself is infeasible, but becomes feasible if any one constraint is removed. The IIS, providing a minimal (however in general not unique) reason for the conflict, is communicated back to the SAT solver, which uses the guard variables of the linear constraints involved to construct a conflict clause which prevents that particular combination of constraints to be investigated again. The resulting interaction between DPLL proof search and feasibility check via LP is illustrated in Figure 2.

4.2   Optimizations for BMC

Compared to related tools like ICS which aim at being general-purpose decision procedures suitable for arbitrary formulae, HySat's decision procedure has been tuned to exploit the unique characteristics of BMC formulae.

As observed by Shtrichman [26], the highly symmetric structure of the k-fold unrolling as shown in section 3 as well as the incremental nature of BMC can both be exploited for various optimizations in the underlying decision procedure. Currently, HySat implements three optimizations which are described

---

[6] http://www.gnu.org/software/glpk/glpk.html
[7] See [10] and [11] for surveys of methods for doing so. Our current implementation of HySat employs the *Deletion-Filter* method for isolating IISs.
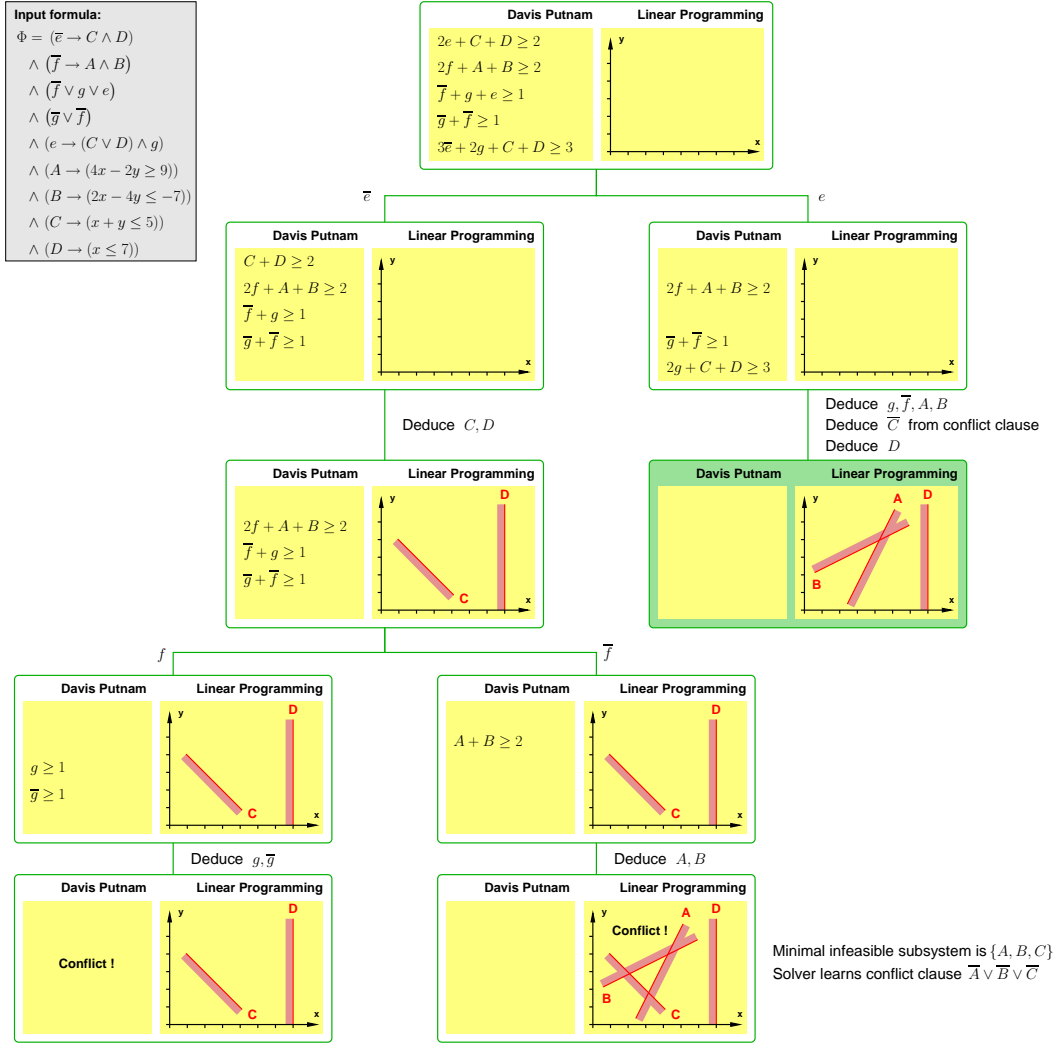
Fig. 2. Backtrack-search tree arising in a tight integration of DPLL proof search with linear programming. x and y are real-valued, while e, f, g and A, B, C, D are Boolean. A, B, C, D are, furthermore, guard variables for arithmetic facts.

below.

4.3 Isomorphy inference

The learning scheme employed in propositional SAT solvers accounts for a substantial fraction of the solver's running time as it entails a non-trivial analysis of the implications that led to an inconsistent valuation. The creation of a conflict clause is in general even considerably more expensive in a combined solver like HySat, as the analysis of a conflict involving non-propositional constraints requires the computationally expensive extraction of an IIS.

Isomorphy inference uses the (almost) symmetric structure of a BMC formula in order to add isomorphic copies of a conflict clause to the problem, thus multiplying the benefit taken from the time-consuming reasoning process

which was required to derive the original conflict clause.

The concept is best illustrated using an example. Suppose that while solving a BMC instance the solver has encountered a conflict which yields the conflict clause $\mathcal{C}^0 = (\overline{x}_3^{j_1} \vee x_4^{j_2} \vee x_9^{j_3})$, relating three variables from cycles j_1, j_2 and j_3. The solver then not only adds $\mathcal{C}^0$ to ϕ^k, but also all possible clauses $\mathcal{C}^i = (\overline{x}_3^{j_1 \pm i} \vee x_4^{j_2 \pm i} \vee x_9^{j_3 \pm i})$, $i = 1, 2, \ldots$, obtained from $\mathcal{C}^0$ simply by index shifting.

Note, however, that BMC is not fully symmetric because of the initialization properties of runs (clause (v) of the translation scheme of section 3) and perhaps the verification goal. This implies that only conflict clauses inferred from facts which are independent from such asymmetric formula parts may be soundly replicated. Such dependency can be traced cheaply by marking initialization/goal predicates and dominantly inheriting such marks upon all inferences, inhibiting isomorphy inference whenever a mark is encountered.

4.3.1 Constraint sharing

When carrying out BMC incrementally for longer and longer unrollings, the consecutive formulae passed to the solver share a large number of clauses. Thus, when moving from the k-instance to the $(k+1)$-instance, we can simply conjoin the conflict clauses derived when solving k-instance to the formula for step $k+1$. However, this is only allowed for conflict clauses that were inferred from clauses which are common to both instances. We do currently decide this based on simple syntactic criteria, namely that the conflict clause was inferred purely from clauses stemming from the automaton. I.e. the inference may not involve the verification goal, which tends to become a weaker predicate on longer instances, as it usually entails reachability or recurrence. More elaborate schemes have, however, been investigated for propositional BMC in [21].

4.3.2 Tailored decision strategy

When applying general-purpose decision strategies to BMC formulae one can observe the phenomenon described in [26] that during the SAT search large sets of constraints belonging to distant cycles of the transition relation are being satisfied independently, until they finally turn out to be incompatible, often entailing the need for backtracking over long distances in the search tree.

In HySat we adopt the solution proposed by Shtrichman [26] to avoid this problem: The heuristics of the SAT solver selects the decision variables in the natural order induced by the variable dependency graph of the BMC formula, i.e. either starting with variables from $\boldsymbol{x}^0$, then from $\boldsymbol{x}^1$, etc., or vice versa. This allows conflicts to be detected and resolved more locally, speeding up the search.

5 Benchmark results

For a first evaluation of HySat we conducted a series of experiments in which we compared our tool with the ICS solver [13] on BMC problems for hybrid automata. The unwindings fed to ICS were obtained through either SRI's infinite-state BMC frontend to ICS as distributed in the SAL tool-set, or through Ábrahám et al.'s corresponding frontend [1], yet the latter without learning across instances, as it was not available when our experiments were performed. Our benchmarks are the *"leaking gas burner"* and *"water-level monitor"* included in the SAL distribution, as well as various instances of an elastic approach to distance control of trains running on the same track, similar to the car platooning system used in the PATH project. Here, trains can accelerate or decelerate freely if they do not violate their mutual safety envelopes, yet an automatic speed control takes authority over a train if another train gets close, thereby controlling acceleration proportional (within physical limits) to the front and/or back proximity of the neighboring trains.

The results are shown in Figure 3, with each dot representing a single BMC instance. As expected, isomorphy inference typically provides performance benefits, with the merits becoming more evident with increasing unrolling depth, corresponding to computationally costly SAT instances. An exception is the extremely deterministic gasburner model, where a strict state alternation is enforced by the discrete part such that learning of infeasible subsystems provides negligible extra information.

6 Conclusion and further work

Even though development of HySat is still in an early stage, first experiments indicate a very competitive performance when used for bounded model checking of linear hybrid systems. A substantial part of this performance gain can be attributed to inheritance of inference results along the temporal axis, called isomorphy inference. The pure base engine of HySat without BMC-specific optimizations exhibits a performance comparable to major other engines like SAL/ICS [13]. The isomorphy inference scheme along the temporal axis was inspired by a similar scheme developed by Shtrichman for finite-state BMC [26]; however such inference-inheritance schemes exhibit an even better payoff on the two-sorted logics used here, as the price for copying inferences increases only marginally while the computational cost of individual inferences grows dramatically in the hybrid-state case.

An interesting aspect of isomorphically copying inference results is that even extremely costly inferences may amortize, provided that their results can be reused sufficiently often. Our next step will thus be to implement more advanced —and computationally more costly— techniques for finding infeasible subsystems of linear constraint systems. In particular, we will try to extract and learn multiple different irreducible infeasible subsystems from
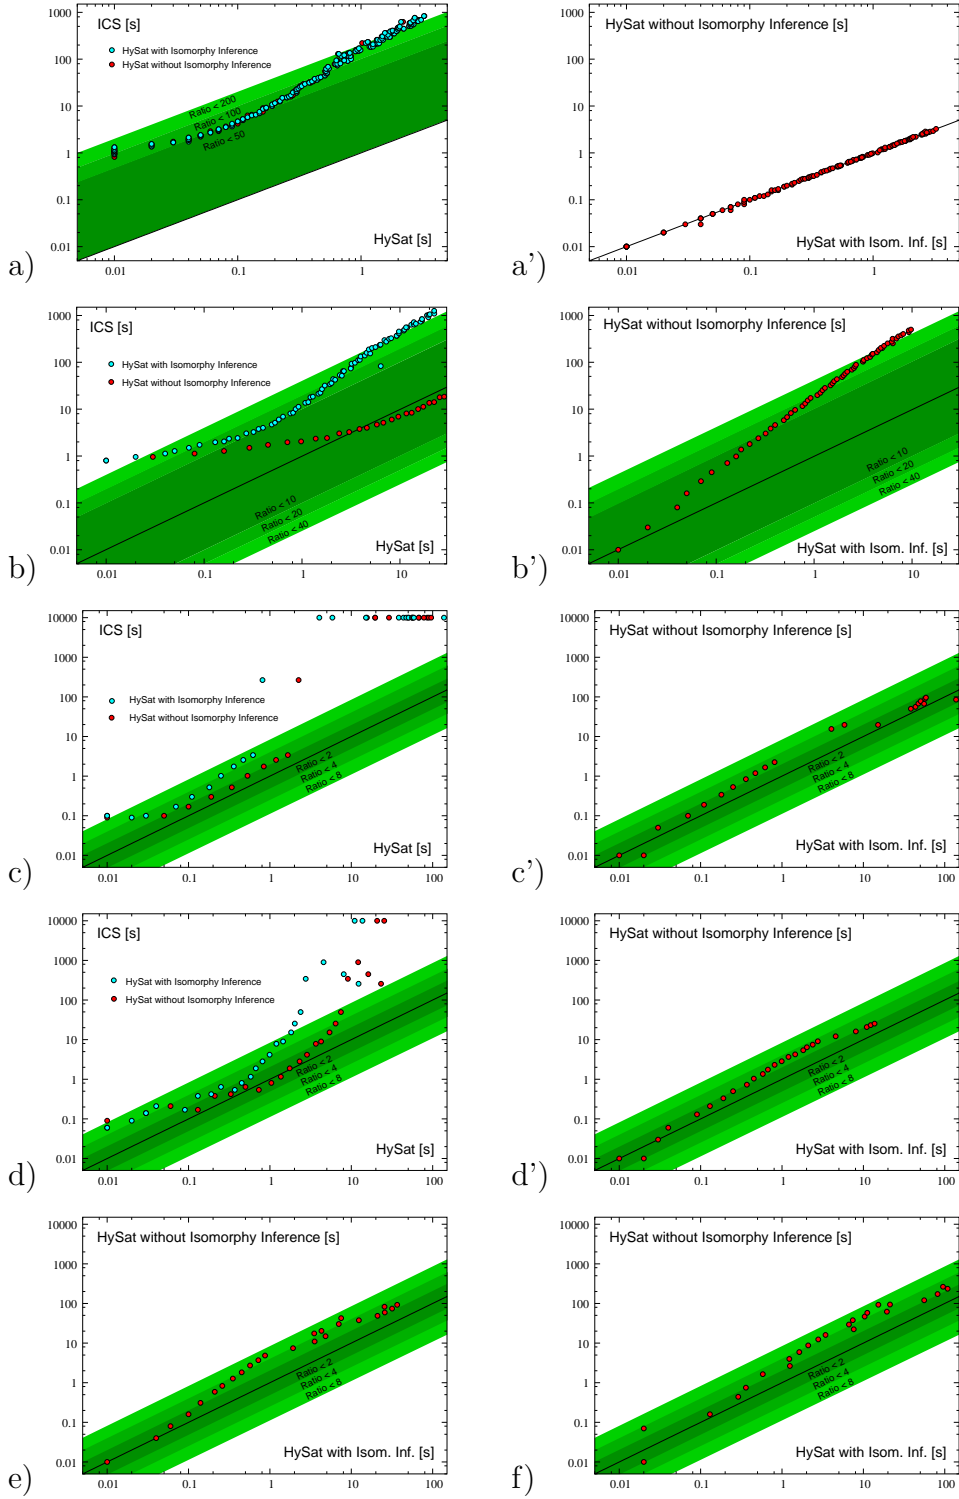
Fig. 3. BMC times for a) gasburner model, b) water-level monitor, and c) – f) different scenarios of the train distance control model, involving 3 (c and d), 4 (e) and 5 (f) trains with different parameters. The graphics show the performance of HySat relative to ICS (a, b, c, d) and the impact of isomorphy inference (a', b', c', d', e, f). Computation times of 10000 denote timeouts.

a single conflict encountered. While the cost of finding an actually irreducible subsystem, and even more so of finding multiple such, is by far the most expensive inference operation in a combined DPLL-plus-LP solver, doing so can provide very aggressive proof-tree pruning.

Another direction for future development will be to add inheritance of inference results across similar components in a multi-component system. While this is in principle similar to inheriting inferences along the temporal axis in k-bounded model checking, the possible forms of symmetry breaks in multi-component ensembles are more diverse and thus harder to detect, as witnessed by the extensive research on symmetry reductions.

Besides exploiting similarity within subformulae for accelerating inference, there obviously also remains ample opportunity for optimizing the underlying decision procedure for HySat's base logic. To this end, we will attack the benchmarks performed by de Moura and Rueß in [14] to obtain a more profound evaluation of the capabilities of HySat's core decision procedure.

References

[1] E. Ábrahám, B. Becker, F. Klaedtke, and M. Steffen. Optimizing bounded model checking for linear hybrid systems. To be submitted to VMCAI '05.

[2] F. A. Aloul, A. Ramani, I. L. Markov, and K. A. Sakallah. Generic ILP versus specialized 0-1 ILP: An update. In *Proc. ACM/IEEE Intl. Conf. Comp.-Aided Design (ICCAD)*, pages 450–457, Nov. 2002.

[3] G. Audemard, P. Bertoli, A. Cimatti, A. Kornilowics, and R. Sebastiani. A SAT-based approach for solving formulas over boolean and linear mathematical propositions. In A. Voronkov, editor, *Proc. of the 18th International Conference on Automated Deduction*, volume 2392 of *Lecture Notes in Artificial Intelligence*, pages 193–208. Springer-Verlag, 2002.

[4] G. Audemard, M. Bozzano, A. Cimatti, and R. Sebastiani. Verifying industrial hybrid systems with MathSAT. *ENTCS*, 89(4), 2004.

[5] L. Baptista, I. Lynce, and J. Marques-Silva. Complete search restart strategies for satisfiability. In *Proc. of the IJCAI'01 Workshop on Stochastic Search Algorithms (IJCAI-SSA)*, August 2001.

[6] C. Barrett, D. Dill, and A. Stump. Checking satisfiability of first-order formulas by incremental translation to SAT. In *14th International Conference on Computer-Aided Verification*, 2002.

[7] P. Barth. A Davis-Putnam based enumeration algorithm for linear pseudo-boolean optimization. Technical Report MPI-I-95-2-003, Max-Planck-Institut für Informatik, Saarbrücken, Germany, 1995.

[8] A. Bemporad and M. Morari. Verification of hybrid systems via mathematical programming. In F. W. Vaandrager and J. H. van Schuppen, editors, *Hybrid Systems: Computation and Control (HSCC'99)*, volume 1569 of *Lecture Notes in Computer Science*, pages 31–45. Springer-Verlag, 1999.

[9] A. Biere, A. Cimatti, and Y. Zhu. Symbolic model checking without BDDs. In *TACAS'99*, volume 1579 of *Lecture Notes in Computer Science*. Springer-Verlag, 1999.

[10] J. W. Chinneck. Finding a useful subset of constraints for analysis in an infeasible linear program. *INFORMS Journal on Computing*, 9(2):164–174, 1997.

[11] J. W. Chinneck and E. W. Dravnieks. Locating minimal infeasible constraint sets in linear programs. *ORSA Journal on Computing*, 3(2):157–168, 1991.

[12] M. Davis, G. Logemann, and D. Loveland. A machine program for theorem proving. *Communications of the ACM*, 5:394–397, 1962.

[13] L. de Moura, S. Owre, H. Ruess, J. Rushby, and N. Shankar. The ICS decision procedures for embedded deduction. In *2nd International Joint Conference on Automated Reasoning (IJCAR)*, volume 3097 of *Lecture Notes in Computer Science*, pages 218–222, Cork, Ireland, July 2004. Springer-Verlag.

[14] L. de Moura and H. Rueß. An experimental evaluation of ground decision procedures. In *Proceedings of CAV'04*, Lecture Notes in Computer Science. Springer-Verlag, July 2004.

[15] L. de Moura, H. Rueß, J. Rushby, and N. Shankar. Embedded deduction with ICS. In B. Martin, editor, *HCSS'03—High Confidence Software and Systems Conference*, Baltimore, MD, 1-3 April 2003.

[16] M. Fränzle and C. Herde. Efficient SAT engines for concise logics: Accelerating proof search for zero-one linear constraint systems. In M. Vardi and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2003)*, volume 2850 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2003.

[17] J. F. Groote, J. W. C. Koorn, and S. F. M. van Vlijmen. The safety guaranteeing system at station hoorn-kersenboogerd. In *Compass '95: 10th Annual Conference on Computer Assurance*, pages 57–68, Gaithersburg, Maryland, 1995. National Institute of Standards and Technology.

[18] E. C. R. Hehner. Predicative programming. *Communications of the ACM*, 27:134–151, 1984.

[19] T. A. Henzinger, P.-H. Ho, and H. Wong-Toi. HyTech: The next generation. In *16th Annual IEEE Real-time Systems Symposium (RTSS 1995)*, pages 56–65. IEEE Computer Society Press, 1995.

[20] T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya. What's decidable about hybrid automata. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on the Theory of Computing*, pages 373–382. ACM, 1995.

[21] H. Jin and F. Somenzi. An incremental algorithm to check satisfiability for bounded model checking. In A. Biere and O. Strichman, editors, *Preliminary Proceeding of BMC'04*. ETH Zürich, 2004. Available from http://bmc04.inf.ethz.ch/JinSomenzi-BMC04-preliminary.pdf.

[22] J. P. Marques-Silva. The impact of branching heuristics in propositional satisfiability algorithms. In *Proc. of the 9th Portuguese Conference on Artificial Intelligence (EPIA)*, Sept. 1999.

[23] J. P. Marques-Silva and K. A. Sakallah. GRASP: A search algorithm for propositional satisfiability. *IEEE Transactions on Computers*, 48(5):506–521, May 1999.

[24] 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 (DAC'01)*, June 2001.

[25] A. Nonnengart and C. Weidenbach. Computing small clause normal forms. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*. Elsevier Science B.V., 1999.

[26] O. Shtrichman. Tuning SAT checkers for bounded model checking. In E. A. Emerson and A. P. Sistla, editors, *Computer Aided Verification (CAV 2000)*, volume 1855 of *Lecture Notes in Computer Science*, pages 480–494. Springer-Verlag, 2000.

[27] G. Tseitin. On the complexity of derivations in propositional calculus. In A. Slisenko, editor, *Studies in Constructive Mathematics and Mathematical Logics*, 1968.

[28] J. P. Warners. A linear-time transformation of linear inequalities into conjunctive normal form. *Information Processing Letters*, 68(2):63–69, 1998.

[29] J. Whittemore, J. Kim, and K. Sakallah. SATIRE: A new incremental satisfiability engine. In *Proc. of the Design Automation Conference (DAC 2001)*, pages 542–545, Las Vegas (Nevada, USA), June 2001.

[30] S. A. Wolfman and D. S. Weld. The LPSAT engine & its application to resource planning. In T. Dean, editor, *Proc. 16th International Joint Conference on i Artificial Intelligence*, pages 310–315. Morgan Kaufmann Publishers, 1999.

[31] L. Zhang, C. F. Madigan, M. W. Moskewicz, and S. Malik. Efficient conflict driven learning in a Boolean satisfiability solver. In *Proc. of the International Conference on Computer-Aided Design (ICCAD01)*, pages 279–285, Nov. 2001.