

Generalized Craig Interpolation for Stochastic Boolean Satisfiability Problems^{*}

Tino Teige and Martin Fränzle

Carl von Ossietzky Universität, Oldenburg, Germany
{teige, fraenzle}@informatik.uni-oldenburg.de

Abstract. The stochastic Boolean satisfiability (SSAT) problem has been introduced by Papadimitriou in 1985 when adding a probabilistic model of uncertainty to propositional satisfiability through randomized quantification. SSAT has many applications, among them bounded model checking (BMC) of symbolically represented Markov decision processes. This paper identifies a notion of *Craig interpolant* for the SSAT framework and develops an algorithm for computing such interpolants based on SSAT resolution. As a potential application, we address the use of interpolation in SSAT-based BMC, turning the falsification procedure into a verification approach for probabilistic safety properties.

1 Introduction

Papadimitriou [1] has proposed the idea of modeling uncertainty within propositional satisfiability (SAT) by adding *randomized* quantification to the problem description. The resultant *stochastic Boolean satisfiability* (SSAT) problems consist of a quantifier prefix followed by a propositional formula. The quantifier prefix is an alternating sequence of existentially quantified variables and variables bound by randomized quantifiers. The meaning of a randomized variable x is that x takes value **true** with a certain probability p and value **false** with the complementary probability $1 - p$. Due to the presence of such probabilistic assignments, the semantics of an SSAT formula Φ no longer is qualitative in the sense that Φ is satisfiable or unsatisfiable, but rather *quantitative* in the sense that we are interested in the *maximum probability of satisfaction* of Φ . Intuitively, a solution of Φ is a strategy for assigning the existential variables, i.e. a tree of assignments to the existential variables depending on the probabilistically determined values of preceding randomized variables, such that the assignments maximize the probability of satisfying the propositional formula.

In recent years, the SSAT framework has attracted interest within the Artificial Intelligence community, as many problems from that area involving uncertainty have concise descriptions as SSAT problems, in particular probabilistic planning problems [2,3,4]. Inspired by that work, other communities have started

^{*} This work has been supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center “Automatic Verification and Analysis of Complex Systems” (SFB/TR 14 AVACS, www.avacs.org).

to exploit SSAT and closely related formalisms within their domains. The Constraint Programming community is working on *stochastic constraint satisfaction* problems [5,6] to address, a.o., multi-objective decision making under uncertainty [7]. Recently, a technique for the symbolic analysis of probabilistic hybrid systems based on stochastic satisfiability has been suggested by the authors [8,9]. To this end, SSAT has been extended by embedded theory reasoning over arithmetic theories, as known from *satisfiability modulo theories* (SMT) [10], which yields the notion of *stochastic satisfiability modulo theories* (SSMT). By the expressive power of SSMT, bounded probabilistic reachability problems of uncertain hybrid systems can be phrased symbolically as SSMT formulae yielding the same probability of satisfaction. As this bounded model checking approach yields valid lower bounds lb for the probability of reaching undesirable system states along unbounded runs, it is able to *falsify* probabilistic safety requirements of shape “a system error occurs with probability at most 0.1%”.

Though the general SSAT problem is PSPACE-complete, the plethora of real-world applications calls for practically efficient algorithms. The first SSAT algorithm, suggested by Littman [11], extends the Davis-Putnam-Logemann-Loveland (DPLL) procedure [12,13] for SAT with appropriate quantifier handling and algorithmic optimizations like *thresholding*. Majercik improved the DPLL-style SSAT algorithm by *non-chronological backtracking* [14]. Unlike these explicit tree-traversal approaches and motivated by work on *resolution* for propositional and first-order formulae [15] and for QBF formulae [16], the authors have recently developed an alternative SSAT procedure based on resolution [17].

In this paper, we investigate the concept of Craig interpolation for SSAT. Given two formulae A and B for which $A \Rightarrow B$ is true, a *Craig interpolant* [18] \mathcal{I} is a formula over variables common to A and B that “lies in between” A and B in the sense that $A \Rightarrow \mathcal{I}$ and $\mathcal{I} \Rightarrow B$. In the automatic hardware and software verification communities, Craig interpolation has found widespread use in model checking algorithms, both as a means of extracting reasons for non-concretizability of a counterexample obtained on an abstraction as well as for obtaining a symbolic description of reachable state sets. In McMillan’s approach [19,20], interpolants are used to symbolically describe an overapproximation of the step-bounded reachable state set. If the sequence of interpolants thus obtained stabilizes eventually, i.e. no additional state is found to be reachable, then the corresponding state-set predicate R has all reachable system states as its models. The safety property that states satisfying B (where B is a predicate) are never reachable, is then verified by checking $R \wedge B$ for unsatisfiability.

Given McMillan’s verification procedure for non-probabilistic systems, it is natural to ask whether a corresponding probabilistic counterpart can be developed, i.e. a *verification procedure for probabilistic systems based on Craig interpolation for stochastic SAT*. Such an approach would complement the aforementioned falsification procedure for probabilistic systems based on SSAT/SSMT. In this paper, we suggest a solution to the issue above. After a formal introduction of SSAT in Section 2, we recall (and adapt slightly) the resolution calculus for SSAT from [17] in Section 3. We suggest a generalization of the notion of Craig interpolants

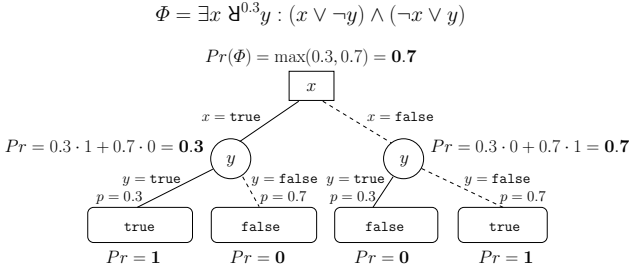


Fig. 1. Semantics of an SSAT formula depicted as a tree

suitable for SSAT as well as an algorithm based on SSAT resolution to compute such generalized interpolants (Section 4). Finally, we propose an interpolation-based approach to probabilistic model checking that is able to *verify* probabilistic safety requirements (Section 5) and illustrate the applicability of this verification procedure on a small example.

2 Stochastic Boolean Satisfiability

A *stochastic Boolean satisfiability* (SSAT) formula is of the form $\Phi = \mathcal{Q} : \varphi$ with a prefix $\mathcal{Q} = Q_1 x_1 \dots Q_n x_n$ of quantified propositional variables x_i , where Q_i is either an existential quantifier \exists or a randomized quantifier \forall^{p_i} with a rational constant $0 < p_i < 1$, and a propositional formula φ s.t. $Var(\varphi) \subseteq \{x_1, \dots, x_n\}$, where $Var(\varphi)$ denotes the set of all (necessarily free) variables occurring in φ . W.l.o.g., we can assume that φ is in *conjunctive normal form* (CNF), i.e. a conjunction of disjunctions of propositional literals. A *literal* ℓ is a propositional variable, i.e. $\ell = x_i$, or its negation, i.e. $\ell = \neg x_i$. A *clause* is a disjunction of literals. Throughout the paper and w.l.o.g., we require that a clause does not contain the same literal more than once as $\ell \vee \ell \equiv \ell$. Consequently, we may also identify a clause with its set of literals. The semantics of Φ , as illustrated in Fig. 1, is defined by the *maximum probability of satisfaction* $Pr(\Phi)$ as follows.

$$Pr(\varepsilon : \varphi) = \begin{cases} 0 & \text{if } \varphi \text{ is logically equivalent to false} \\ 1 & \text{if } \varphi \text{ is logically equivalent to true} \end{cases}$$

$$Pr(\exists x \mathcal{Q} : \varphi) = \max(Pr(\mathcal{Q} : \varphi[\text{true}/x]), Pr(\mathcal{Q} : \varphi[\text{false}/x]))$$

$$Pr(\forall^p x \mathcal{Q} : \varphi) = p \cdot Pr(\mathcal{Q} : \varphi[\text{true}/x]) + (1 - p) \cdot Pr(\mathcal{Q} : \varphi[\text{false}/x])$$

Note that the semantics is well-defined as Φ has no free variables s.t. all variables have been substituted by the constants **true** and **false** when reaching the quantifier-free base case.

3 Resolution for SSAT

As basis of the SSAT interpolation procedure introduced in Section 4, we recall the sound and complete resolution calculus for SSAT from [17], subsequently called *S-resolution*. In contrast to SSAT algorithms implementing a DPLL-based

backtracking procedure, thereby explicitly traversing the tree given by the quantifier prefix and recursively computing the individual satisfaction probabilities for each subtree by the scheme illustrated in Fig. 1, S-resolution follows the idea of *resolution* for propositional and first-order formulae [15] and for QBF formulae [16] by deriving new clauses c^p annotated with probabilities $0 \leq p \leq 1$. S-resolution differs from non-stochastic resolution, as such derived clauses c^p need not be implications of the given formula, but are just entailed with some probability. Informally speaking, the derivation of a clause c^p means that under SSAT formula $\mathcal{Q} : \varphi$, the clause c is violated with a maximum probability at most p , i.e. the satisfaction probability of $\mathcal{Q} : (\varphi \wedge \neg c)$ is at most p . More intuitively, the minimum probability that clause c is implied by φ is at least $1 - p$.¹ Once an annotated empty clause \emptyset^p is derived, it follows that the probability of the given SSAT formula is at most p , i.e. $Pr(\mathcal{Q} : (\varphi \wedge \mathbf{false})) = Pr(\mathcal{Q} : \varphi) \leq p$.

The following presentation of S-resolution differs slightly from [17] in order to avoid overhead in interpolant generation incurred when employing the original definition, like the necessity of enforcing particular resolution sequences. For readers familiar with [17], the particular modifications are: 1) derived clauses c^p may also carry value $p = 1$, 2) former rules R.2 and R.5 are joined into the new rule R.2, and 3) former rules R.3 and R.4 are collapsed into rule R.3. These modifications do not affect soundness and completeness of S-resolution (cf. Corollary 1 and Theorem 1). The advantage of the modification is that derivable clauses c^p are forced to have a tight bound p in the sense that under each assignment which falsifies c , the satisfaction probability of the remaining subproblem *exactly* is p (cf. Lemma 1). This fact confirms the conjecture from [17, p. 14] about the existence of such clauses $(c \vee \ell)^p$ and allows for a generalized clause learning scheme to be integrated into DPLL-SSAT solvers: the idea is that under a partial assignment falsifying c , one may directly propagate literal ℓ as the satisfaction probability of the other branch, for which the negation of ℓ holds, is known to be p already.

In the sequel, let $\Phi = \mathcal{Q} : \varphi$ be an SSAT formula with φ in CNF. W.l.o.g., φ contains only non-tautological clauses², i.e. $\forall c \in \varphi : \not\models c$. Let $\mathcal{Q} = Q_1 x_1 \dots Q_n x_n$ be the quantifier prefix and φ be some propositional formula with $Var(\varphi) \subseteq \{x_1, \dots, x_n\}$. The quantifier prefix $\mathcal{Q}(\varphi)$ is defined to be shortest prefix of \mathcal{Q} that contains all variables from φ , i.e. $\mathcal{Q}(\varphi) = Q_1 x_1 \dots Q_i x_i$ where $x_i \in Var(\varphi)$ and for each $j > i : x_j \notin Var(\varphi)$. Let further be $Var(\varphi) \downarrow_k := \{x_1, \dots, x_k\}$ for each integer $0 \leq k \leq n$. For a non-tautological clause c , i.e. if $\not\models c$, we define the unique assignment ff_c that falsifies c as the mapping

$$ff_c : Var(c) \rightarrow \mathbb{B} \text{ such that } \forall x \in Var(c) : ff_c(x) = \begin{cases} \mathbf{true} & ; \neg x \in c, \\ \mathbf{false} & ; x \in c. \end{cases}$$

Consequently, c evaluates to **false** under assignment ff_c .

¹ We remark that $Pr(\mathcal{Q} : \psi) = 1 - Pr(\mathcal{Q}' : \neg\psi)$, where \mathcal{Q}' arises from \mathcal{Q} by replacing existential quantifiers by universal ones, where universal quantifiers call for *minimizing* the satisfaction probability.

² Tautological clauses c , i.e. $\models c$, are redundant, i.e. $Pr(\mathcal{Q} : (\varphi \wedge c)) = Pr(\mathcal{Q} : \varphi)$.

Starting with clauses in φ , *S-resolution* is given by the consecutive application of rules R.1 to R.3 to derive new clauses c^p with $0 \leq p \leq 1$. Rule R.1 derives a clause c^0 from an original clause c in φ . Referring to the definition of $Pr(\Phi)$ in Section 2, R.1 corresponds to the quantifier-free base case where φ is equivalent to **false** under any assignment that falsifies c .

$$(R.1) \quad \frac{c \in \varphi}{c^0}$$

Similarly, R.2 reflects the quantifier-free base case in which φ is equivalent to **true** under any assignment τ' that is conform to the partial assignment τ since $\models \varphi[\tau(x_1)/x_1] \dots [\tau(x_i)/x_i]$. The constructed clause c^1 then encodes the opposite of this satisfying (partial) assignment τ . We remark that finding such a τ in the premise of R.2 is NP-hard (equivalent to finding a solution of a propositional formula in CNF). This strong condition on τ is not essential for soundness and completeness and could be removed³ but, as mentioned above, facilitates a less technical presentation of generalized interpolation in Section 4. Another argument justifying the strong premise of R.2 is a potential integration of S-resolution into DPLL-based SSAT solvers since whenever a satisfying (partial) assignment τ of φ is found by an SSAT solver then τ meets the requirements of R.2.

$$(R.2) \quad \frac{\begin{array}{l} c \subseteq \{x, \neg x \mid x \in \text{Var}(\varphi)\}, \not\models c, \mathcal{Q}(c) = Q_1x_1 \dots Q_ix_i, \\ \text{for each } \tau : \text{Var}(\varphi) \downarrow_i \rightarrow \mathbb{B} \text{ with } \forall x \in \text{Var}(c) : \tau(x) = \text{ff}_c(x) : \\ \models \varphi[\tau(x_1)/x_1] \dots [\tau(x_i)/x_i] \end{array}}{c^1}$$

Rule R.3 finally constitutes the actual resolution rule as known from the non-stochastic case. Depending on whether an existential or a randomized variable is resolved upon, the probability value of the resolvent clause is computed according to the semantics $Pr(\Phi)$ defined in Section 2.

$$(R.3) \quad \frac{(c_1 \vee \neg x)^{p_1}, (c_2 \vee x)^{p_2}, Qx \in \mathcal{Q}, Qx \notin \mathcal{Q}(c_1 \vee c_2), \not\models (c_1 \vee c_2), \quad p = \begin{cases} \max(p_1, p_2) & ; Q = \exists \\ p_x \cdot p_1 + (1 - p_x) \cdot p_2 & ; Q = \mathfrak{R}^{p_x} \end{cases}}{(c_1 \vee c_2)^p}$$

The derivation of a clause c^p by R.1 from c , by R.2, and by R.3 from $c_1^{p_1}, c_2^{p_2}$ is denoted by $c \vdash_{R.1} c^p$, by $\vdash_{R.2} c^p$, and by $(c_1^{p_1}, c_2^{p_2}) \vdash_{R.3} c^p$, respectively. Given rules R.1 to R.3, S-resolution is sound and complete in the following sense.

Lemma 1. *Let clause c^p be derivable by S-resolution and let $\mathcal{Q}(c) = Q_1x_1 \dots Q_ix_i$. For each $\tau : \text{Var}(\varphi) \downarrow_i \rightarrow \mathbb{B}$ with $\forall x \in \text{Var}(c) : \tau(x) = \text{ff}_c(x)$ it holds that $Pr(Q_{i+1}x_{i+1} \dots Q_nx_n : \varphi[\tau(x_1)/x_1] \dots [\tau(x_i)/x_i]) = p$.*

Corollary 1 (Soundness of S-resolution). *If the empty clause \emptyset^p is derivable by S-resolution from a given SSAT formula $\mathcal{Q} : \varphi$ then $Pr(\mathcal{Q} : \varphi) = p$.*

³ Then, Lemma 1 must be weakened (as for original S-resolution [17]) to $Pr(Q_{i+1}x_{i+1} \dots Q_nx_n : \varphi[\tau(x_1)/x_1] \dots [\tau(x_i)/x_i]) \leq p$.

Theorem 1 (Completeness of S-resolution). *If $Pr(Q : \varphi) = p$ for some SSAT formula $Q : \varphi$ then the empty clause \emptyset^p is derivable by S-resolution.*

The formal proofs of above and subsequent results can be found in [21].

Example. Consider the SSAT formula $\Phi = \forall^{0.8}x_1 \exists x_2 \forall^{0.3}x_3 : ((x_1 \vee x_2) \wedge (\neg x_2) \wedge (x_2 \vee x_3))$ with $Pr(\Phi) = 0.24$. Clauses $(x_1 \vee x_2)^0$, $(\neg x_2)^0$, $(x_2 \vee x_3)^0$ are then derivable by R.1. As $x_1 = \mathbf{true}$, $x_2 = \mathbf{false}$, $x_3 = \mathbf{true}$ is a satisfying assignment, $\vdash_{R.2} (\neg x_1 \vee x_2 \vee \neg x_3)^1$. Then, $((\neg x_1 \vee x_2 \vee \neg x_3)^1, (x_2 \vee x_3)^0) \vdash_{R.3} (\neg x_1 \vee x_2)^{0.3}$, $((\neg x_2)^0, (\neg x_1 \vee x_2)^{0.3}) \vdash_{R.3} (\neg x_1)^{0.3}$, $((\neg x_2)^0, (x_1 \vee x_2)^0) \vdash_{R.3} (x_1)^0$, and finally $((\neg x_1)^{0.3}, (x_1)^0) \vdash_{R.3} \emptyset^{0.24}$.

4 Interpolation for SSAT

Craig interpolation [18] is a well-studied notion in formal logics which has several applications in Computer Science, among them model checking [19,20]. Given two formulae φ and ψ such that $\varphi \Rightarrow \psi$ is valid, a *Craig interpolant* for (φ, ψ) is a formula \mathcal{I} which refers only to common variables of φ and ψ , and \mathcal{I} is “intermediate” in the sense that $\varphi \Rightarrow \mathcal{I}$ and $\mathcal{I} \Rightarrow \psi$. Such interpolants do trivially exist in all logics permitting quantifier elimination, e.g. in propositional logic. Using the observation that $\varphi \Rightarrow \psi$ holds iff $\varphi \wedge \neg\psi$ is unsatisfiable, this gives rise to an equivalent definition which we refer to in the rest of the paper:⁴ given an unsatisfiable formula $\varphi \wedge \neg\psi$, formula \mathcal{I} is an interpolant for (φ, ψ) iff both $\varphi \wedge \neg\mathcal{I}$ and $\mathcal{I} \wedge \neg\psi$ are unsatisfiable and \mathcal{I} mentions only common variables.

In the remainder of this section, we investigate the issue of interpolation for stochastic SAT. We propose a generalization of Craig interpolants suitable for SSAT and show the general existence of such interpolants, alongside with an automatic method for computing them based on S-resolution.

4.1 Generalized Craig Interpolants

When approaching a reasonable definition of interpolants for SSAT, the semantics of the non-classical quantifier prefix poses problems: Let $\Phi = Q : (A \wedge B)$ be an SSAT formula. Each variable in $A \wedge B$ is bound by Q , which provides the probabilistic interpretation of the variables that is lacking without the quantifier prefix. This issue can be addressed by considering the quantifier prefix Q as the global setting that serves to interpret the quantifier-free part, and consequently interpreting the interpolant also within the scope of Q , thus reasoning about $Q : (A \wedge \neg\mathcal{I})$ and $Q : (\mathcal{I} \wedge B)$. A more fundamental problem is that a classical Craig interpolant for Φ only exists if $Pr(\Phi) = 0$, since $A \wedge B$ has to be unsatisfiable by definition of a Craig interpolant which applies iff $Pr(Q : (A \wedge B)) = 0$.

⁴ This is of technical nature as SSAT formulae are interpreted by the maximum probability of satisfaction. As the *maximum* probability that an implication $\varphi \Rightarrow \psi$ holds is inappropriate for our purpose, we reason about the maximum satisfaction probability p of the negated implication, i.e. of $\varphi \wedge \neg\psi$, instead. The latter coincides with the *minimum* probability $1 - p$ that $\varphi \Rightarrow \psi$ holds, which is the desired notion.

The precondition that $Pr(Q : (A \wedge B)) = 0$ would be far too restrictive for application of interpolation, as the notion of unsatisfiability of $A \wedge B$ is naturally generalized to satisfiability with insufficient probability, i.e. $Pr(Q : (A \wedge B))$ being “sufficiently small”, in the stochastic setting. Such relaxed requirements actually appear in practice, e.g. in probabilistic verification where safety properties like “a fatal system error is never reachable” are frequently replaced by probabilistic ones like “a fatal system error is reachable only with (sufficiently small) probability of at most 0.1%”. Motivated by above facts, interpolants for SSAT should also exist when $A \wedge B$ is satisfiable with reasonably low probability.

The resulting notion of interpolation, which is to be made precise in Definition 1, matches the following intuition. In classical Craig interpolation, when performed in logics permitting quantifier elimination, the Craig interpolants of $(A, \neg B)$ form a lattice with implication as its ordering, $A^\exists = \exists a_1, \dots, a_\alpha : A$ as its bottom element and $\overline{B}^\forall = \neg \exists b_1, \dots, b_\beta : B$ as its top element, where the a_i and b_i are the local variables of A and of B , respectively. In the generalized setting required for SSAT⁵, $A \Rightarrow \neg B$ and thus $A^\exists \Rightarrow \overline{B}^\forall$ may no longer hold such that the above lattice can collapse to the empty set. To preserve the overall structure, it is however natural to use the lattice of propositional formulae “in between” $A^\exists \wedge \overline{B}^\forall$ as bottom element and $A^\exists \vee \overline{B}^\forall$ as top element instead. This lattice is non-empty and coincides with the classical one whenever $A \wedge B$ is unsatisfiable.

Definition 1 (Generalized Craig interpolant). *Let A, B be propositional formulae and $V_A := \text{Var}(A) \setminus \text{Var}(B) = \{a_1, \dots, a_\alpha\}$, $V_B := \text{Var}(B) \setminus \text{Var}(A) = \{b_1, \dots, b_\beta\}$, $V_{A,B} := \text{Var}(A) \cap \text{Var}(B)$, $A^\exists = \exists a_1, \dots, a_\alpha : A$, and $\overline{B}^\forall = \neg \exists b_1, \dots, b_\beta : B$. A propositional formula \mathcal{I} is called generalized Craig interpolant for (A, B) iff $\text{Var}(\mathcal{I}) \subseteq V_{A,B}$, $(A^\exists \wedge \overline{B}^\forall) \Rightarrow \mathcal{I}$, and $\mathcal{I} \Rightarrow (A^\exists \vee \overline{B}^\forall)$.*

Given any two propositional formulae A and B , the four quantifier-free propositional formulae equivalent to $A^\exists \wedge \overline{B}^\forall$, to A^\exists , to \overline{B}^\forall , and to $A^\exists \vee \overline{B}^\forall$, are generalized Craig interpolants for (A, B) . These generalized interpolants always exist since propositional logic has quantifier elimination.

While Definition 1 motivates the generalized notion of Craig interpolant from a model-theoretic perspective, we will state an equivalent definition of generalized Craig interpolants in Lemma 2 that substantiates the intuition of generalized interpolants and allows for an illustration of their geometric shape. Given two formulae A and B , the idea of generalized Craig interpolant is depicted in Fig. 2. The set of solutions of A is defined by the rectangle on the $V_A, V_{A,B}$ -plane with a cylindrical extension in V_B -direction as A does not contain variables in V_B . Similarly, the solution set of B is given by the triangle on the $V_B, V_{A,B}$ -plane and its cylinder in V_A -direction. The solution set of $A \wedge B$ is then determined by the intersection of both cylinders. Since $A \wedge B \wedge \neg(A \wedge B)$ is unsatisfiable, the sets $A \wedge \neg(A \wedge B)$ and $B \wedge \neg(A \wedge B)$ are disjoint. This gives us the possibility to talk about interpolants wrt. these sets. However, a formula \mathcal{I} over only common variables in $V_{A,B}$ may not exist when demanding $A \wedge \neg(A \wedge B) \wedge \neg \mathcal{I}$ and

⁵ Though the concept seems to be more general, this paper addresses SSAT only.




Fig. 2. Geometric interpretation of a generalized Craig interpolant \mathcal{I} . V_A -, V_B -, and $V_{A,B}$ -axes denote assignments of variables occurring only in A , only in B , and in both A and B , respectively.

$\mathcal{I} \wedge B \wedge \neg(A \wedge B)$ to be unsatisfiable. This is indicated by Fig. 2 and proven by the simple example $A = (a)$, $B = (b)$. As $V_{A,B} = \emptyset$, \mathcal{I} is either **true** or **false**. In first case, $\mathbf{true} \wedge (b) \wedge \neg(a \wedge b)$ is satisfiable, while $(a) \wedge \neg(a \wedge b) \wedge \mathbf{false}$ is in second case. If we however project the solution set of $A \wedge B$ onto the $V_{A,B}$ -axis and subtract the resulting hyperplane $\mathcal{S}_{A,B}$ from A and B then such a formula \mathcal{I} over $V_{A,B}$ -variables exists. The next lemma formalizes such generalized interpolants \mathcal{I} and shows their equivalence to the ones from Definition 1.

Lemma 2 (Generalized Craig interpolant for SSAT). *Let $\Phi = \mathcal{Q} : (A \wedge B)$ be some SSAT formula, $V_A, V_B, V_{A,B}$ be defined as in Definition 1, and $\mathcal{S}_{A,B}$ be a propositional formula with $\text{Var}(\mathcal{S}_{A,B}) \subseteq V_{A,B}$ s.t. $\mathcal{S}_{A,B} \equiv \exists a_1, \dots, a_\alpha, b_1, \dots, b_\beta : (A \wedge B)$. Then, a propositional formula \mathcal{I} is a generalized Craig interpolant for (A, B) iff the following properties are satisfied.*

1. $\text{Var}(\mathcal{I}) \subseteq V_{A,B}$
2. $\text{Pr}(\mathcal{Q} : (A \wedge \neg \mathcal{S}_{A,B} \wedge \neg \mathcal{I})) = 0$
3. $\text{Pr}(\mathcal{Q} : (\mathcal{I} \wedge B \wedge \neg \mathcal{S}_{A,B})) = 0$

We remark that the concept of generalized Craig interpolants is a generalization of Craig interpolants in the sense that whenever $A \wedge B$ is unsatisfiable, i.e. when

$Pr(Q : (A \wedge B)) = 0$, then each generalized Craig interpolant \mathcal{I} for (A, B) actually is a Craig interpolant for A and B since $\mathcal{S}_{A,B} \equiv \text{false}$.

4.2 Computation of Generalized Craig Interpolants

In this subsection, we proceed to the efficient computation of generalized Craig interpolants. The remark following Definition 1 shows that generalized interpolants can in principle be computed by *explicit* quantifier elimination methods, like Shannon’s expansion or BDDs. We aim at a more efficient method based on S-resolution, akin to resolution-based Craig interpolation for propositional SAT by Pudlák [22], as has been integrated into DPLL-based SAT solvers featuring conflict analysis and successfully applied to symbolic model checking [19,20].

Observe that on SSAT formulae $Q : (A \wedge B)$, Pudlák’s algorithm, which has unsatisfiability of $A \wedge B$ as precondition, will not work in general. When instead considering the unsatisfiable formula $A \wedge B \wedge \neg\mathcal{S}_{A,B}$ with $\neg\mathcal{S}_{A,B}$ in CNF then Pudlák’s method would be applicable and would actually produce a generalized interpolant. The main drawback of this approach, however, is the explicit construction of $\neg\mathcal{S}_{A,B}$, again calling for explicit quantifier elimination.

We now propose an algorithm based on S-resolution for computing generalized Craig interpolants which operates directly on $A \wedge B$ without adding $\neg\mathcal{S}_{A,B}$, and thus does not comprise any preprocessing involving quantifier elimination. For this purpose, rules of S-resolution are enhanced to deal with pairs (c^p, I) of annotated clauses c^p and propositional formulae I . Such formulae I are in a certain sense *intermediate* generalized interpolants, i.e. generalized interpolants for subformulae arising from instantiating some variables by partial assignments that falsify c (cf. Lemma 3). Once a pair (\emptyset^p, I) comprising the empty clause is derived, I thus is a generalized Craig interpolant for the given SSAT formula. This augmented S-resolution, which we call *interpolating S-resolution*, is defined by rules R’.1, R’.2, and R’.3. The construction of intermediate interpolants I in R’.1 and R’.3 coincides with the classical rules by Pudlák [22], while R’.2 misses a corresponding counterpart. The rationale is that R’.2 (or rather R.2) refers to satisfying valuations τ of $A \wedge B$, which do not exist in classical interpolation. As $A \wedge B$ becomes a tautology after substituting the partial assignment τ from R.2 into it, its quantified variant $\mathcal{S}_{A,B} = \exists a_1, \dots, b_1, \dots : A \wedge B$ also becomes tautological under the same substitution $\mathcal{S}_{A,B}[\tau(x_1)/x_1, \dots, \tau(x_i)/x_i]$. Consequently, $\neg\mathcal{S}_{A,B}[\tau(x_1)/x_1, \dots, \tau(x_i)/x_i]$ is unsatisfiable, and so are $(A \wedge \neg\mathcal{S}_{A,B})[\tau(x_1)/x_1, \dots, \tau(x_i)/x_i]$ and $(B \wedge \neg\mathcal{S}_{A,B})[\tau(x_1)/x_1, \dots, \tau(x_i)/x_i]$. This implies that the actual intermediate interpolant in R’.2 can be chosen arbitrarily over variables in $V_{A,B}$. This freedom will allow us to control the geometric extent of generalized interpolants within the “don’t care”-region provided by the models of $\mathcal{S}_{A,B}$ (cf. Corollary 3).

$$(R'.1) \quad \frac{c \vdash_{R.1} c^p, I = \begin{cases} \text{false} & ; c \in A \\ \text{true} & ; c \in B \end{cases}}{(c^p, I)}$$

$$(R'.2) \quad \frac{\vdash_{R,2} c^p, I \text{ is any formula over } V_{A,B}}{(c^p, I)}$$

$$(R'.3) \quad \frac{\begin{array}{l} ((c_1 \vee \neg x)^{p_1}, I_1), ((c_2 \vee x)^{p_2}, I_2), \\ ((c_1 \vee \neg x)^{p_1}, (c_2 \vee x)^{p_2}) \vdash_{R,3} (c_1 \vee c_2)^p, \\ I = \begin{cases} I_1 \vee I_2 & ; x \in V_A \\ I_1 \wedge I_2 & ; x \in V_B \\ (\neg x \vee I_1) \wedge (x \vee I_2) & ; x \in V_{A,B} \end{cases} \end{array}}{(c_1 \vee c_2)^p, I}$$

The following lemma establishes the theoretical foundation of computing generalized Craig interpolants by interpreting the derived pairs (c^p, I) .

Lemma 3. *Let $\Phi = \mathcal{Q} : (A \wedge B)$ with $\mathcal{Q} = Q_1x_1 \dots Q_nx_n$ be some SSAT formula, and the pair (c^p, I) be derivable from Φ by interpolating S-resolution, where $\mathcal{Q}(c) = Q_1x_1 \dots Q_ix_i$. Then, for each $\tau : \text{Var}(A \wedge B) \downarrow_i \rightarrow \mathbb{B}$ with $\forall x \in \text{Var}(c) : \tau(x) = \text{ff}_c(x)$ it holds that*

1. $\text{Var}(I) \subseteq V_{A,B}$,
2. $\text{Pr}(Q_{i+1}x_{i+1} \dots Q_nx_n : (A \wedge \neg \mathcal{S}_{A,B} \wedge \neg I)[\tau(x_1)/x_1] \dots [\tau(x_i)/x_i]) = 0$, and
3. $\text{Pr}(Q_{i+1}x_{i+1} \dots Q_nx_n : (I \wedge B \wedge \neg \mathcal{S}_{A,B})[\tau(x_1)/x_1] \dots [\tau(x_i)/x_i]) = 0$.

Completeness of S-resolution, as stated in Theorem 1, together with the above Lemma, applied to the derived pair (\emptyset^p, I) , yields

Corollary 2 (Generalized Craig interpolants computation). *If $\Phi = \mathcal{Q} : (A \wedge B)$ is an SSAT formula then a generalized Craig interpolant for (A, B) can be computed by interpolating S-resolution.*

Note that computation of generalized interpolants does not depend on the actual truth state of $A \wedge B$. The next observation facilitates to effectively control the geometric extent of generalized Craig interpolants within the “don’t care”-region $\mathcal{S}_{A,B}$. This result will be useful in probabilistic model checking in Section 5.

Corollary 3 (Controlling generalized Craig interpolants computation). *If $I = \text{true}$ is used within each application of rule R'.2 then $\text{Pr}(\mathcal{Q} : (A \wedge \neg \mathcal{I})) = 0$. Likewise, if $I = \text{false}$ is used in rule R'.2 then $\text{Pr}(\mathcal{Q} : (\mathcal{I} \wedge B)) = 0$.*

Observe that the special interpolants \mathcal{I} from Corollary 3 relate to the classical strongest and weakest Craig interpolants A^\exists and \overline{B}^\forall , resp., in the following sense: $\text{Pr}(\mathcal{Q} : (A \wedge \neg \mathcal{I})) = 0$ iff $\models A \Rightarrow \mathcal{I}$ iff $\models \forall a_1, \dots, a_\alpha : (A \Rightarrow \mathcal{I})$ iff $\models (A^\exists \Rightarrow \mathcal{I})$, as a_1, \dots, a_α do not occur in \mathcal{I} . Analogously, $\text{Pr}(\mathcal{Q} : (\mathcal{I} \wedge B)) = 0$ iff $\models \mathcal{I} \Rightarrow \neg B$ iff $\models \forall b_1, \dots, b_\beta : (\mathcal{I} \Rightarrow \neg B)$ iff $\models \mathcal{I} \Rightarrow \overline{B}^\forall$.

5 Interpolation-Based Probabilistic Model Checking

In this section, we investigate the application of generalized Craig interpolation in probabilistic model checking. As a model we consider symbolically represented finite-state Markov decision processes (MDPs), which we check wrt. to

probabilistic state reachability properties. That is, given a set of target states T in an MDP \mathcal{M} , the goal is to maximize the probability $P_{\mathcal{M}}^{\pi}(T)$ of reaching T over all policies π resolving the non-determinism in \mathcal{M} . When considering T as *bad* states of \mathcal{M} , e.g. as fatal system errors, this maximum probability $P_{\mathcal{M}}^{\max}(T) := \max_{\pi} P_{\mathcal{M}}^{\pi}(T)$ reveals the *worst-case* probability of bad system behavior. A *safety property* for \mathcal{M} requires that this worst-case probability does not exceed some given threshold value $\theta \in [0, 1)$, i.e. $P_{\mathcal{M}}^{\max}(T) \leq \theta$.

In [8], we proposed a symbolic *falsification* procedure for such safety properties. Though the approach in [8] is based on SSMT (arithmetic extension of SSAT) and considers the more general class of discrete-time probabilistic hybrid systems, which roughly are MDPs with arithmetic-logical transition guards and actions, the same procedure restricted to SSAT is applicable for finite-state MDPs. The key idea here is to adapt *bounded model checking* (BMC) [23] to the probabilistic case by encoding step-bounded reachability as an SSAT problem: like in classical BMC, the initial states, the transition relation, and the target states of an MDP are symbolically encoded by propositional formulae in CNF, namely by $Init(\mathbf{s})$, $Trans(\mathbf{s}, \mathbf{nt}, \mathbf{pt}, \mathbf{s}')$, and $Target(\mathbf{s})$, resp., where the propositional variable vector \mathbf{s} represents the system state before and \mathbf{s}' after a transition step. To keep track of the *non-deterministic* and *probabilistic* selections of transitions in $Trans(\mathbf{s}, \mathbf{nt}, \mathbf{pt}, \mathbf{s}')$, we further introduce propositional variables \mathbf{nt} and \mathbf{pt} to encode non-deterministic and probabilistic transition choices, respectively. Assignments to these variables determine which of possibly multiple available transitions departing from \mathbf{s} is taken. In contrast to traditional BMC, all variables are quantified: all state variables \mathbf{s} and \mathbf{s}' are existentially quantified in the prefixes $\mathcal{Q}_{\mathbf{s}}$ and $\mathcal{Q}_{\mathbf{s}'}$. The transition-selection variables \mathbf{nt} encoding non-deterministic choice are *existentially quantified* by $\mathcal{Q}_{\mathbf{nt}}$, while the probabilistic selector variables \mathbf{pt} are bound by *randomized quantifiers* in $\mathcal{Q}_{\mathbf{pt}}$.⁶ Let be $\mathbf{t} := \mathbf{nt} \cup \mathbf{pt}$ and $\mathcal{Q}_{\mathbf{t}} := \mathcal{Q}_{\mathbf{nt}} \mathcal{Q}_{\mathbf{pt}}$. Due to [8, Proposition 1], the maximum probability of reaching the target states in a given MDP from the initial states within k transition steps is equivalent to the satisfaction probability

$$(1) \quad lb_k := Pr\left(\mathcal{Q}(k) : \left(\overbrace{Init(\mathbf{s}_0) \wedge \bigwedge_{i=1}^k Trans(\mathbf{s}_{i-1}, \mathbf{t}_i, \mathbf{s}_i)}^{\text{states reachable within } k \text{ steps}} \wedge \left(\overbrace{\bigvee_{i=0}^k Target(\mathbf{s}_i)}^{\text{hit target states}} \right) \right) \right)$$

with $\mathcal{Q}(k) := \mathcal{Q}_{\mathbf{s}_0} \mathcal{Q}_{\mathbf{t}_1} \mathcal{Q}_{\mathbf{s}_1} \dots \mathcal{Q}_{\mathbf{s}_{k-1}} \mathcal{Q}_{\mathbf{t}_k} \mathcal{Q}_{\mathbf{s}_k}$. The probability lb_k can be determined by SSAT solvers. This symbolic approach, called *probabilistic bounded model checking* (PBMC), produces valid *lower* bounds lb_k for $P_{\mathcal{M}}^{\max}(T)$. Thus, PBMC is able to *falsify* a safety property once $lb_k > \theta$ for some k . However, the development of a corresponding counterpart based on SSAT that is able to compute safe *upper* bounds ub_k for $P_{\mathcal{M}}^{\max}(T)$ was left as an open question. Such an approach would permit to *verify* a safety property once $ub_k \leq \theta$ for some k .

We now propose such a verification procedure based on generalized Craig interpolation. The algorithm proceeds in two phases. Phase 1 computes a symbolic representation of an *overapproximation of the backward reachable state set*. This can be integrated into PBMC, as used to falsify the probabilistic safety property.

⁶ Non-deterministic branching of n alternatives can be represented by a binary tree of depth $\lceil \log_2 n \rceil$ and probabilistic branching by a sequence of at most $n - 1$ binary branches, yielding $\lceil \log_2 n \rceil$ existential and $n - 1$ randomized quantifiers, respectively.

Whenever this falsification fails for a given step depth k , we apply generalized Craig interpolation to the (just failed) PBMC proof to compute a *symbolic overapproximation of the backward reachable state set* at depth k and then proceed to PBMC at some higher depth $k' > k$. As an alternative to the integration into PBMC, interpolants describing the backward reachable state sets can be successively extended by “stepping” them by prepending another transition, as explained below. In either case, phase 1 ends when the backward reachable state set becomes stable, in which case we have computed a symbolic overapproximation of the whole backward-reachable state set. In the second phase, we construct an SSAT formula with parameter k that forces the system to *stay within the backward reachable state set* for k steps. The maximum satisfaction probability of that SSAT formula then gives an upper bound on the maximum probability of reaching the target states. The rationale is that system executions leaving the backward reachable states will never reach the target states.

Phase 1. Given an SSAT encoding of an MDP \mathcal{M} as above, the state-set predicate $\mathcal{B}^k(\mathbf{s})$ for $k \in \mathbb{N}_{\geq 0}$ over state variables \mathbf{s} is inductively defined as $\mathcal{B}^0(\mathbf{s}) := \text{Target}(\mathbf{s})$, and $\mathcal{B}^{k+1}(\mathbf{s}) := \mathcal{B}^k(\mathbf{s}) \vee \mathcal{I}^{k+1}(\mathbf{s}_{j-1})$ where $\mathcal{I}^{k+1}(\mathbf{s}_{j-1})$ is a generalized Craig interpolant for $(\text{Trans}(\mathbf{s}_{j-1}, \mathbf{t}_j, \mathbf{s}_j) \wedge \mathcal{B}^k(\mathbf{s}_j), \text{Init}(\mathbf{s}_0) \wedge \bigwedge_{i=1}^{j-1} \text{Trans}(\mathbf{s}_{i-1}, \mathbf{t}_i, \mathbf{s}_i))$ with $j \geq 1$ wrt. to SSAT formula

$$(2) \quad \mathcal{Q}(j) : \left(\overbrace{\text{Init}(\mathbf{s}_0) \wedge \bigwedge_{i=1}^{j-1} \text{Trans}(\mathbf{s}_{i-1}, \mathbf{t}_i, \mathbf{s}_i)}^{j-1 \text{ steps "forward" (=B)}} \wedge \overbrace{\text{Trans}(\mathbf{s}_{j-1}, \mathbf{t}_j, \mathbf{s}_j) \wedge \mathcal{B}^k(\mathbf{s}_j)}^{\text{one step "backward" (=A)}} \right)$$

Observe that each $\mathcal{I}^{k+1}(\mathbf{s})$ can be computed by interpolating S-resolution if we rewrite $\mathcal{B}^k(\mathbf{s})$ into CNF (which is always possible in linear time by adding auxiliary V_A -variables). When doing so, we take $I = \text{true}$ in each application of R'.2 such that $\mathcal{B}^k(\mathbf{s})$ overapproximates all system states backward reachable from target states within k steps due to Corollary 3. Whenever $\mathcal{B}^k(\mathbf{s})$ stabilizes, i.e. $\mathcal{B}^{k+1}(\mathbf{s}) \Rightarrow \mathcal{B}^k(\mathbf{s})$, we can be sure that $\widehat{\text{BReach}}(\mathbf{s}) := \mathcal{B}^k(\mathbf{s})$ overapproximates all backward reachable states.

Note that parameter $j \geq 1$ can be chosen arbitrarily, i.e. the system may execute any number of transitions until state \mathbf{s}_{j-1} is reached since this does not destroy the “backward-overapproximating” property of $\mathcal{B}^k(\mathbf{s})$. The rationale of having parameter j is the additional freedom in constructing $\mathcal{I}^k(\mathbf{s})$ as j may influence the shape of $\mathcal{I}^k(\mathbf{s})$ (cf. example below).

Phase 2. Having symbolically described all backward reachable states by $\widehat{\text{BReach}}(\mathbf{s})$, we are able to compute *upper* bounds ub_k of the maximum probability $P_{\mathcal{M}}^{\max}(T)$ of reaching target states T by SSAT solving applied to

$$(3) \quad ub_k := Pr \left(\mathcal{Q}(k) : \left(\overbrace{\text{Init}(\mathbf{s}_0) \wedge \bigwedge_{i=1}^k \text{Trans}(\mathbf{s}_{i-1}, \mathbf{t}_i, \mathbf{s}_i)}^{\text{states reachable within } k \text{ steps}} \wedge \overbrace{\bigwedge_{i=0}^k \widehat{\text{BReach}}(\mathbf{s}_i)}^{\text{stay in back-reach set}} \right) \right)$$

First observe that the formula above excludes all system runs that leave the set of backward reachable states. This is sound since leaving $\widehat{\text{BReach}}(\mathbf{s})$ means to never reach the $\text{Target}(\mathbf{s})$ states. Second, the system behavior becomes more and more constrained for increasing k , i.e. the ub_k 's are monotonically decreasing.

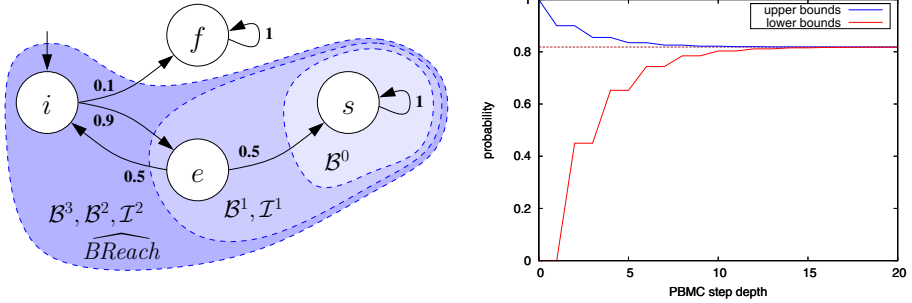


Fig. 3. A simple DTMC \mathcal{M} (left) and lower and upper bounds on probability of reaching s over the PBMC step depth (right)

Regarding model checking, a safety property $P_{\mathcal{M}}^{\max}(T) \leq \theta$ is verified by the procedure above once $ub_k \leq \theta$ is computed for some k .

Example. Consider the simple *discrete-time Markov chain* (DTMC)⁷ \mathcal{M} on the left of Fig. 3, with s as the only target state. The (maximum) probability of reaching s from the initial state i clearly is $P = 0.9 \cdot (0.5 + 0.5 \cdot P) = \frac{9}{11}$. Applying the generalized interpolation scheme (2) for $k = 0$ with parameter $j = 1$ yields the interpolant $\mathcal{I}^1(s) = \neg i$. If we proceed then the set $\widehat{BReach}(s)$ covers all states though state f is not backward reachable. Thus, $j = 1$ is not a good choice as each ub_k in scheme (3) will then be 1. For $j = 2$, scheme (2) produces: $\mathcal{B}^0(s) = Target(s) = s$, $\mathcal{I}^1(s) = \neg i \wedge \neg f$, $\mathcal{B}^1(s) = \mathcal{B}^0(s) \vee \mathcal{I}^1(s) = (\neg i \vee s) \wedge (\neg f \vee s)$, $\mathcal{I}^2(s) = (\neg i \vee \neg e) \wedge \neg f$, $\mathcal{B}^2(s) = \mathcal{B}^1(s) \vee \mathcal{I}^2(s) = (\neg i \vee \neg e \vee s) \wedge (\neg f \vee s)$, $\mathcal{I}^3(s) = (\neg i \vee \neg e) \wedge \neg f$, $\mathcal{B}^3(s) = \mathcal{B}^2(s) \vee \mathcal{I}^3(s) = \mathcal{B}^2(s)$ (cf. left of Fig. 3). Thus, $\widehat{BReach}(s) = \mathcal{B}^2(s)$.

We are now able to compute upper bounds of the reachability probability $P = \frac{9}{11}$ by scheme (3). The results are shown on the right of Fig. 3 where the lower bounds are computed according to scheme (1). The figure indicates a fast convergence of the lower and upper bounds: for depth $k = 20$ the difference $ub_k - lb_k$ is below 10^{-3} and for $k = 100$ below 10^{-15} .

All 200 SSAT formulae were solved by the SSMT tool SiSAT [9] in 41.3sec on a 1.83 GHz Intel Core 2 Duo machine with 1 GByte physical memory running Linux. For the moment, the SSAT encoding of \mathcal{M} and all computations of generalized interpolants from PBMC proofs were performed manually. Details can be found in [21]. While the runtime of this first prototype does not compare favorably to value or policy iteration procedures, it should be noted that this is a first step towards a procedure embedding the same interpolation process into SSMT [8] and thereby directly addressing probabilistic hybrid systems, where the iteration procedures are only applicable after finite-state abstraction. It should also be noted that the symbolic procedures provided by SSAT and SSMT support compact representations of concurrent systems, thus alleviating the state explosion problem [9].

⁷ A DTMC is an MDP without non-determinism.

6 Conclusion and Future Work

In this paper, we elaborated on the idea of Craig interpolation for stochastic SAT. In consideration of the difficulties that arise in this stochastic extension of SAT, we first proposed a suitable definition of a generalized Craig interpolant and second presented an algorithm to automatically compute such interpolants. For the latter purpose, we enhanced the SSAT resolution calculus by corresponding rules for the construction of generalized interpolants. We further demonstrated an application of generalized Craig interpolation in probabilistic model checking. The resulting procedure is able to verify probabilistic safety requirements of the form “the worst-case probability of reaching undesirable system states is smaller than a given threshold”. This complements the existing SSAT-based bounded model checking approach, which mechanizes falsification of such properties.

An essential issue for future work is the practical evaluation of interpolation-based probabilistic model checking on realistic case studies. This involves, a.o., the integration of interpolating S-resolution into DPLL-based SSAT solvers and a thorough empirical study of the results obtained from the interpolation scheme (2). The latter includes the size and shape of generalized Craig interpolants as well as the computational effort for computing them in practice.

Another interesting future direction is the adaptation of generalized Craig interpolation to SSMT [8], i.e. the extension of SSAT with arithmetic theories. Computing Craig interpolants for SSMT would lift schemes (2) and (3) to SSMT problems, thus establishing a symbolic verification procedure for discrete-time probabilistic hybrid systems.

References

1. Papadimitriou, C.H.: Games against nature. *J. Comput. Syst. Sci.* 31(2), 288–301 (1985)
2. Littman, M.L., Majercik, S.M., Pitassi, T.: Stochastic Boolean Satisfiability. *Journal of Automated Reasoning* 27(3), 251–296 (2001)
3. Majercik, S.M., Littman, M.L.: MAXPLAN: A New Approach to Probabilistic Planning. In: *Artificial Intelligence Planning Systems*, pp. 86–93 (1998)
4. Majercik, S.M., Littman, M.L.: Contingent Planning Under Uncertainty via Stochastic Satisfiability. *Artificial Intelligence Special Issue on Planning With Uncertainty and Incomplete Information* 147(1-2), 119–162 (2003)
5. Walsh, T.: Stochastic constraint programming. In: *Proc. of the 15th European Conference on Artificial Intelligence (ECAI 2002)*. IOS Press, Amsterdam (2002)
6. Balafoutis, T., Stergiou, K.: Algorithms for Stochastic CSPs. In: Benhamou, F. (ed.) *CP 2006*. LNCS, vol. 4204, pp. 44–58. Springer, Heidelberg (2006)
7. Bordeaux, L., Samulowitz, H.: On the stochastic constraint satisfaction framework. In: *SAC*, pp. 316–320. ACM, New York (2007)
8. Fränzle, M., Hermanns, H., Teige, T.: Stochastic satisfiability modulo theory: A novel technique for the analysis of probabilistic hybrid systems. In: Egerstedt, M., Mishra, B. (eds.) *HSCC 2008*. LNCS, vol. 4981, pp. 172–186. Springer, Heidelberg (2008)

9. Teige, T., Eggers, A., Fränzle, M.: Constraint-based analysis of concurrent probabilistic hybrid systems: An application to networked automation systems. *Nonlinear Analysis: Hybrid Systems* (to appear, 2011)
10. Barrett, C., Sebastiani, R., Seshia, S.A., Tinelli, C.: Satisfiability modulo theories. In: Biere, A., Heule, M.J.H., van Maaren, H., Walsh, T. (eds.) *Handbook of Satisfiability. Frontiers in Artificial Intelligence and Applications*, vol. 185, pp. 825–885. IOS Press, Amsterdam (2009)
11. Littman, M.L.: Initial Experiments in Stochastic Satisfiability. In: *Proc. of the 16th National Conference on Artificial Intelligence*, pp. 667–672 (1999)
12. Davis, M., Putnam, H.: A Computing Procedure for Quantification Theory. *Journal of the ACM* 7(3), 201–215 (1960)
13. Davis, M., Logemann, G., Loveland, D.: A Machine Program for Theorem Proving. *Communications of the ACM* 5, 394–397 (1962)
14. Majercik, S.M.: Nonchronological backtracking in stochastic Boolean satisfiability. In: *ICTAI*, pp. 498–507. IEEE Computer Society, Los Alamitos (2004)
15. Robinson, J.A.: A machine-oriented logic based on the resolution principle. *J. ACM* 12(1), 23–41 (1965)
16. Büning, H.K., Karpinski, M., Flögel, A.: Resolution for quantified Boolean formulas. *Inf. Comput.* 117(1), 12–18 (1995)
17. Teige, T., Fränzle, M.: Resolution for stochastic Boolean satisfiability. In: Fermüller, C.G., Voronkov, A. (eds.) *LPAR-17. LNCS*, vol. 6397, pp. 625–639. Springer, Heidelberg (2010)
18. Craig, W.: Linear Reasoning. A New Form of the Herbrand-Gentzen Theorem. *J. Symb. Log.* 22(3), 250–268 (1957)
19. McMillan, K.L.: Interpolation and SAT-based model checking. In: Hunt Jr., W.A., Somenzi, F. (eds.) *CAV 2003. LNCS*, vol. 2725, pp. 1–13. Springer, Heidelberg (2003)
20. McMillan, K.L.: Applications of Craig interpolants in model checking. In: Halbwachs, N., Zuck, L.D. (eds.) *TACAS 2005. LNCS*, vol. 3440, pp. 1–12. Springer, Heidelberg (2005)
21. Teige, T., Fränzle, M.: Generalized Craig interpolation for stochastic Boolean satisfiability problems. *Reports of SFB/TR 14 AVACS 67, SFB/TR 14 AVACS* (2011) ISSN: 1860-9821, <http://www.avacs.org>
22. Pudlák, P.: Lower Bounds for Resolution and Cutting Plane Proofs and Monotone Computations. *Journal of Symbolic Logic* 62(3), 981–998 (1997)
23. Biere, A., Cimatti, A., Clarke, E.M., Zhu, Y.: Symbolic model checking without BDDs. In: Cleaveland, W.R. (ed.) *TACAS 1999. LNCS*, vol. 1579, pp. 193–207. Springer, Heidelberg (1999)