

# Unifying Decision and Function Queries in Stochastic Boolean Satisfiability

Yu-Wei Fan<sup>1</sup>, Jie-Hong R. Jiang<sup>1,2</sup>

<sup>1</sup>Graduate Institute of Electronics Engineering, National Taiwan University, Taipei, Taiwan

<sup>2</sup>Department of Electrical Engineering, National Taiwan University, Taipei, Taiwan  
{r11943096, jhjiang}@ntu.edu.tw

## Abstract

Stochastic Boolean satisfiability (SSAT) is a natural formalism for optimization under uncertainty. Its decision version implicitly imposes a final threshold quantification on an SSAT formula. However, the single threshold quantification restricts the expressive power of SSAT. In this work, we enrich SSAT with an additional threshold quantifier, resulting in a new formalism  $\text{SSAT}(\Theta)$ . The increased expressiveness allows  $\text{SSAT}(\Theta)$ , which remains in the PSPACE complexity class, to subsume and encode the languages in the counting hierarchy. An  $\text{SSAT}(\Theta)$  solver,  $\text{ClauSSat}(\Theta)$ , is developed. Experiments show the applicability of the solver in uniquely solving complex  $\text{SSAT}(\Theta)$  instances of parameter synthesis and SSAT extension.

## 1 Introduction

Stochastic Boolean satisfiability (SSAT) is a logical formalism, enabling natural characterization for optimizing decisions under uncertainty (Papadimitriou 1985). Due to its powerful expressiveness, recent endeavors have been made in both its efficient solving and potential applications. There are recent developments of specialized solvers (Lee, Wang, and Jiang 2017, 2018) designed for specific fragments of SSAT, and general solvers (Majercik and Boots 2005; Chen, Huang, and Jiang 2021; Wang et al. 2022; Fan and Jiang 2023) that place no restrictions on the SSAT formula. Regarding applications, SSAT has been used in encoding problems, such as contingent planning (Majercik and Littman 2003), partially observable Markov decision processes (POMDPs) (Salmon and Poupart 2020), the fairness analysis of machine learning models (Ghosh, Basu, and Meel 2021), and probabilistic graphical models (Hsieh and Jiang 2022).

Despite its expressiveness and broad applications, certain limitations are inherent to SSAT. For instance, SSAT implicitly imposes a linear ordering upon the dependency sets of existential variables according to the prefix. In (Lee and Jiang 2021), dependency stochastic Boolean satisfiability (DSSAT) (Lee and Jiang 2021) is formulated to allow explicit representation of dependency sets of existential variables. In this work, we tackle the limitations of SSAT from another aspect. Specifically, the decision version of SSAT

enforces a single outermost threshold quantification on an SSAT formula, limiting its expressive power. Inspired by the counting quantifier in the Counting Hierarchy (Wagner 1986), we introduce a new formalism  $\text{SSAT}(\Theta)$ , which enriches SSAT with the threshold quantifier. We prove that  $\text{SSAT}(\Theta)$  is in the PSPACE complexity class. Remarkably, while remaining in the same complexity class, the new formalism subsumes both SSAT and the languages in the Counting Hierarchy. Therefore,  $\text{SSAT}(\Theta)$  can be powerful in encoding problems not succinctly expressible before.

To demonstrate its practical applications, we develop an  $\text{SSAT}(\Theta)$  solver  $\text{ClauSSat}(\Theta)$ , based on the state-of-the-art SSAT solver  $\text{ClauSSat}$  (Chen, Huang, and Jiang 2021). We further provide  $\text{SSAT}(\Theta)$  encodings for probabilistic model-checking and its extension specified in probabilistic computation tree logic (BPCTL). To evaluate the  $\text{SSAT}(\Theta)$  solver, experiments were conducted to study the applicability in solving application benchmarks and to investigate the effect of the threshold quantifier on solving efficiency. We note that since  $\text{SSAT}(\Theta)$  subsumes the languages in the Counting Hierarchy,  $\text{ClauSSat}(\Theta)$  is the first solver to tackle general counting formulas, while previous work has primarily concentrated on certain restricted fragments, such as counting formulas with only one (Chou et al. 2016) or two (Oztok, Choi, and Darwiche 2016) counting quantifiers.

The rest of this paper is organized as follows. The preliminary background is provided in Section 2. The new formalism  $\text{SSAT}(\Theta)$  is presented, and its properties and complexity results are studied in Section 3. The extension of  $\text{ClauSSat}$  for  $\text{SSAT}(\Theta)$  solving is elaborated in Section 4. The encodings of the model-checking problem specified in BPCTL and the parameter synthesis problems are detailed in Section 5. Section 6 investigates the efficiency of the solver on application benchmarks and the effect of threshold quantifier on general  $\text{SSAT}(\Theta)$  instances. Section 7 concludes this work and discusses future work.

## 2 Preliminaries

In the sequel, the symbols “ $\top$ ” and “ $\perp$ ” represent Boolean values TRUE and FALSE, respectively. Boolean connectives “ $\neg$ ,” “ $\vee$ ,” “ $\wedge$ ,” “ $\rightarrow$ ,” and “ $\leftrightarrow$ ” are associated with their conventional meanings. For simplicity, a conjunction  $\wedge$  may be omitted in a Boolean formula. A *literal*, associated with a variable  $v$ , is either the variable itself  $v$  or the negation  $\neg v$ .

A *clause* is a conjunction of literals. A Boolean formula is in the *conjunctive normal form* (CNF) if it is a conjunction of clauses.

For a CNF formula  $\phi$ , we use  $\text{vars}(\phi)$  to denote the set of variables that appear in  $\phi$ . A Boolean function, represented by a Boolean formula  $f$  over variables  $V$ , is a mapping  $f : \mathbb{B}^{|V|} \rightarrow \mathbb{B}$ . An *assignment*  $\sigma$  over the variables  $V$ , is a mapping  $\sigma : V \rightarrow \mathbb{B}$ . The *induced* formula over an assignment  $\sigma$ , denoted as  $f|_\sigma$ , is obtained by replacing each variable  $v$  with its assigned Boolean value  $\sigma(v)$ .

The interval notation  $[l..u]$ , for  $l, u \in \mathbb{Z}^+ \cup \{0\}$  and  $l < u$ , represents the set of integers from  $l$  to  $u$ . We use symbols  $\triangleright$  and  $\bowtie$  to denote one of the predicates in the set  $\{>, \geq\}$  and  $\{>, \geq, <, \leq\}$ , respectively.

### Quantified Boolean and Counting Formulas

A quantified Boolean formula (QBF) in the prenex CNF form is expressed as

$$\mathcal{Q}.\phi, \quad (1)$$

where  $\mathcal{Q} = Q_1, \dots, Q_n$ , for  $Q_i \in \{\exists v_i, \forall v_i\}$ , and  $\phi$  is a quantifier-free CNF formula.

The QBF satisfiability well corresponds to the *Polynomial Hierarchy* (PH) (Stockmeyer 1976). Specifically, the language of QBFs with  $k$  quantifier alternations, denoted  $\text{QBF}_k$ , and  $Q_1 = \exists$  (resp.  $\forall$ ) is complete in the complexity class  $\Sigma_{k+1}^P$  (resp.  $\Pi_{k+1}^P$ ). When  $k$  is unbounded, QBF is PSPACE-complete.

The *Counting Hierarchy* (CH) (Wagner 1986) for the counting problems is a complexity analogy to PH for the decision problems. Similar to QBFs, which well characterize PH, there are *counting formulas* (CFs), which well characterize CH. A counting formula in the prenex form can be expressed as

$$C V_1, \dots, C V_n.\phi, \quad (2)$$

where  $C$  is the *counting quantifier*,  $V_i$  is a set of variables, and  $\phi$  is a quantifier-free CNF formula. The quantification  $C V$  asks *whether there exist at least half of the assignments over  $V$* . For singleton  $V = \{v\}$ ,  $C V$  is equivalent to  $\exists v$ . Quantifier  $\forall v$  can also be expressed by  $C$  by proper formula negation or rewriting. E.g.,  $\forall v.\phi$  can be rewritten as  $\neg C v.\neg\phi$ , or  $C V.\phi \wedge v'$  for  $V = \{v, v'\}$ , where  $v'$  is a fresh new auxiliary variable. Therefore, a counting formula only requires quantifier  $C$ , without  $\exists$  and  $\forall$ . The language of counting formulas of  $k$  levels of  $C$  quantifiers, denoted  $\text{CF}_k$ , is complete in the complexity class  $C_k^P$ . When  $k$  is unbounded, CF is PSPACE-complete.

### Stochastic Boolean Satisfiability

An SSAT formula in the prenex form can be expressed by Eq. (1), but with the *prefix*  $\mathcal{Q} = Q_1, \dots, Q_n$  for  $Q_i \in \{\exists v_i, \mathfrak{P}^{p_i} v_i\}$  and the matrix  $\phi$  being a quantifier-free CNF formula. In quantifier  $Q_i$ , variable  $v_i$  is either existential-quantified, i.e.,  $\exists v_i$ , or random-quantified, i.e.,  $\mathfrak{P}^{p_i} v_i$ , denoting  $v_i = \top$  (resp.  $\perp$ ) with probability  $p_i$  (resp.  $1 - p_i$ ). The semantics of an SSAT formula  $\Phi = Q_1, \dots, Q_n.\phi = Q_1.\Phi'$  is interpreted as its satisfying probability computed recursively by the following rules:

- $\Pr[\top] = 1$ ,

- $\Pr[\perp] = 0$ ,
- $\Pr[\exists v.\Phi'] = \max\{\Pr[\Phi'|_v], \Pr[\Phi'|_{\neg v}]\}$ ,
- $\Pr[\mathfrak{P}^p v.\Phi'] = p \cdot \Pr[\Phi'|_v] + (1 - p) \cdot \Pr[\Phi'|_{\neg v}]$ .

Given an SSAT formula  $\Phi$ , the *decision version* of SSAT, a PSPACE-complete problem, is to determine if  $\Pr[\Phi]$  is greater than a threshold probability, whereas the *optimization (function) version* is to return the probability  $\Pr[\Phi]$ .

Note that the existential quantifier  $\exists$  in SSAT differs from that in QBFs and counting formulas in its function sense, searching for an assignment maximizing the satisfying probability. More precisely, it is a “maximization quantifier.” Nevertheless, we abuse the notation as its meaning should be clear from the context.

We remark that SSAT can incorporate the universal quantifier  $\forall$  serving as the “minimization quantifier” (Littman, Majercik, and Pitassi 2001). Because  $\forall$  quantification can be achieved through  $\exists$  quantification and negation, e.g.,  $\Pr[\neg \exists X.\phi] = \Pr[\forall X.\neg\phi]$ , we omit the universal quantifier in our discussion for simplicity.

### 3 Threshold Quantifier and SSAT

The language  $\text{SSAT}(\Theta)$ , i.e., SSAT augmented with the threshold quantifier, is defined as follows.

**Definition 1** ( $\text{SSAT}(\Theta)$  Syntax). The syntax of  $\text{SSAT}(\Theta)$  is the same as that of SSAT except that  $Q_i \in \{\exists v_i, \mathfrak{P}^{p_i} v_i, \Theta^{\triangleright p_i}\}$ , with the additional *threshold quantifier*  $\Theta^{\triangleright p_i}$ . We omit to specify the variables within the scope of the threshold quantifier  $Q_i = \Theta^{\triangleright p_i}$  by implicitly assuming its inclusion of all the variables involved in  $Q_{i+1}, \dots, Q_n$ .

**Definition 2** ( $\text{SSAT}(\Theta)$  Semantics). Given an  $\text{SSAT}(\Theta)$  formula  $\Phi = Q_1, \dots, Q_n.\phi = Q_1.\Phi'$ , its *satisfying probability* is defined the same as that of SSAT except for the following additional rule for the Boolean interpretation of threshold quantifier:

- $\Theta^{\triangleright p}.\Phi' = \begin{cases} \top, & \text{if } \Pr[\Phi'] \triangleright p, \\ \perp, & \text{otherwise.} \end{cases}$

We note that  $\text{SSAT}(\Theta)$  may incorporate the universal quantifier and threshold quantifier of predicates  $\{<, \leq\}$ . However, as we can rewrite them by negation along with the existential quantifier and threshold quantifier of predicates  $\{>, \geq\}$ , we omit them in our discussion. E.g., the formula  $\Theta^{< p}.\exists x, \mathfrak{P}^{p'} y.\phi$  is equivalent to  $\neg \Theta^{\geq p}.\exists x, \mathfrak{P}^{p'} y.\phi = \Theta^{\geq 1-p}.\forall x, \mathfrak{P}^{p'} y.\neg\phi$ .

An  $\text{SSAT}(\Theta)$  formula is called *closed* if all the variables are quantified, and *opened* otherwise. Also, an  $\text{SSAT}(\Theta)$  formula  $\Phi = \mathcal{Q}.\phi$  is called *pure-counting* if its prefix  $\mathcal{Q}$  consists of only threshold and random quantifiers, and has  $Q_1$  being a threshold quantifier. We also called such prefix  $\mathcal{Q}$  *pure-counting*. Note that when the random quantifiers in a pure-counting  $\text{SSAT}(\Theta)$  formula are of probability 0.5,  $\Phi$  is equivalently a counting formula. E.g., the counting formula  $C X_1, \dots, C X_n.\phi$  is equivalent to the  $\text{SSAT}(\Theta)$  formula  $\Theta^{\geq 0.5}.\mathfrak{P}^{0.5} X_1, \dots, \Theta^{\geq 0.5}.\mathfrak{P}^{0.5} X_n.\phi$ . When the random quantifiers in a pure-counting  $\text{SSAT}(\Theta)$  formula are of

arbitrary probabilities expressed in binary fractional numbers, we can still derive its equivalent counting formula, exploiting the normalization technique that transforms SSAT formulas to have only probability 0.5 (Wang et al. 2022).

Section 5 will show that the BPCTL model-checking problem can be naturally encoded with the pure-counting SSAT( $\Theta$ ) formula.

Note that an opened and pure-counting SSAT( $\Theta$ ) formula  $\Phi$  with a set of free variables  $X$  represents a Boolean function  $f(X)$  such that  $f|_{\sigma_X} = \top$  if and only if  $\Phi|_{\sigma_X} = \top$ , where  $\sigma_X$  is an assignment over  $X$ . Unless otherwise stated, we assume that an SSAT( $\Theta$ ) formula is closed in the sequel.

Note also that, unlike SSAT, SSAT( $\Theta$ ) requires no distinction between the decision and function versions as the threshold-quantifier extension unifies the decision and function specification.

### Properties and Normal Form

We study some properties of the threshold quantifier and exploit them for a normal form conversion.

**Definition 3** (Normal Form of SSAT( $\Theta$ )). A prenex SSAT( $\Theta$ ) formula  $\Phi = Q_1, \dots, Q_n.\phi$  is in a normal form if the following two conditions hold:

1. There are no consecutive threshold quantifiers. I.e.,  $Q_i$  and  $Q_{i+1}$ , for  $i \in [1..n-1]$ , cannot be both threshold quantifiers.
2. A threshold quantifier  $Q_i$ , for  $i \in [1..n-1]$ , cannot be followed by an existential quantifier  $Q_{i+1}$ .

The normal form can be enforced for any prenex SSAT( $\Theta$ ) formula due to Lemmas 1 and 2 stated below.

Given two threshold quantifiers  $\Theta^{\triangleright p_1}$  and  $\Theta^{\triangleright p_2}$ , we say quantifier  $\Theta^{\triangleright p_1}$  *dominates*  $\Theta^{\triangleright p_2}$  if the implication

$$\Theta^{\triangleright p_2}.\Phi \rightarrow \Theta^{\triangleright p_1}.\Phi \quad (3)$$

holds for any SSAT( $\Theta$ ) formula  $\Phi$ . The following lemmas are immediate.

**Lemma 1.** Let  $\Theta^{\triangleright p_1}$  dominate  $\Theta^{\triangleright p_2}$ . Then the following equalities hold.

$$\Theta^{\triangleright p_1}, \Theta^{\triangleright p_2}.\Phi = \Theta^{\triangleright p_2}, \Theta^{\triangleright p_1}.\Phi = \Theta^{\triangleright p_1}.\Phi \quad (4)$$

**Lemma 2.** By treating  $\top$  (resp.  $\perp$ ) as probability value 1 (resp. 0), and vice versa, the following equality holds.

$$\Theta^{\triangleright p}, \exists v.\Phi = \exists v, \Theta^{\triangleright p}.\Phi \quad (5)$$

### Computation Complexity

Just as each level in PH forms a complete class, each level in CH forms a complete class, which can be characterized by adding the majority quantifier to QBF. E.g., E-MAJSAT (Littman, Goldsmith, and Mundhenk 1998) is  $\vee$  C-complete (Wagner 1986), the same complexity class as  $\text{NP}^{\text{PP}}$ .

As discussed, the pure-counting SSAT( $\Theta$ ) formulas subsume the counting formulas. SSAT( $\Theta$ ) can succinctly encode any problem in the counting hierarchy. On the other hand, while CH consists of only decision problems, SSAT( $\Theta$ ) allows the encoding of function or optimization problems. Hence, SSAT( $\Theta$ ) is strictly more expressive.

In the following, we give the computation complexity of the decision version of SSAT( $\Theta$ ).

**Theorem 1.** SSAT( $\Theta$ ) is PSPACE-complete.

## 4 An SSAT( $\Theta$ ) Solver

A naive reference procedure for SSAT( $\Theta$ ) evaluation is shown in Algorithm 1. Although it requires only polynomial space, it may not be effective in run-time efficiency. Therefore, we need a more advanced algorithm to alleviate the intrinsic hardness of SSAT( $\Theta$ ).

Algorithm 1: SSAT( $\Theta$ ) Evaluation

---

```

1: procedure EVALUATE(  $\Phi = Q_1, \dots, Q_n.\phi$  )
2:   if  $\phi = \top$  or  $\phi = \perp$  then
3:     return  $\text{Pr}[\phi]$ 
4:   if  $Q_1$  is  $\Theta^{\triangleright p}$  then
5:     return EVALUATE( $Q_2, \dots, Q_n.\phi$ )  $\triangleright p$ 
6:    $v \leftarrow$  the outermost variable
7:    $p_0 \leftarrow$  EVALUATE( $Q_2, \dots, Q_n.\phi|_{\neg v}$ )
8:    $p_1 \leftarrow$  EVALUATE( $Q_2, \dots, Q_n.\phi|_v$ )
9:   if  $Q_1$  is  $\exists v$  then
10:    return  $\max\{p_0, p_1\}$ 
11:   if  $Q_1$  is  $\forall^p v$  then
12:    return  $p \cdot p_1 + (1 - p) \cdot p_0$ 

```

---

As SSAT( $\Theta$ ) is a generalization of SSAT, existing SSAT solvers could be extended for SSAT( $\Theta$ ). In this work, we consider the state-of-the-art general SSAT solvers SharpSSAT (Fan and Jiang 2023) and ClaSSat (Chen, Huang, and Jiang 2021) for such extension.

For the case of SharpSSAT, a direct extension is not possible unfortunately because the component decomposition property of SSAT (Salmon and Poupart 2020) does not hold for SSAT( $\Theta$ ).

**Lemma 3.** Given an SSAT( $\Theta$ ) formula  $\Phi = Q.\phi$  with matrix  $\phi = \phi_1 \wedge \dots \wedge \phi_k$ , where  $\text{vars}(\phi_i) \cap \text{vars}(\phi_j) = \emptyset$  for  $i \neq j$ , let  $\Phi_i = Q.\phi_i$ . Then the equality  $\text{Pr}[\Phi] = \prod_{i=1}^k \text{Pr}[\Phi_i]$  does not hold in general.

*Proof.* As a counterexample, consider the SSAT( $\Theta$ ) formula

$$\Phi = \forall^{0.5} v_1, v_2, \Theta^{\geq 0.5}, \forall^{0.5} v_3, v_4. (v_1 \leftrightarrow v_3) \wedge (v_2 \leftrightarrow v_4).$$

Its matrix can be decomposed into  $\phi_1 \wedge \phi_2$ , where  $\phi_1 = (v_1 \leftrightarrow v_3)$  and  $\phi_2 = (v_2 \leftrightarrow v_4)$  have disjoint support variables. However,  $\text{Pr}[\Phi] = 0 \neq \text{Pr}[\Phi_1] \cdot \text{Pr}[\Phi_2] = 0.5 \cdot 0.5 = 0.25$ .  $\square$

For the case of ClaSSat, its extension to SSAT( $\Theta$ ) solving is possible, as discussed below.

ClaSSat solves an SSAT formula by partitioning the literals in a clause in the matrix into several groups with respect to the quantification levels according to the prefix.

For the extension to SSAT( $\Theta$ ), given an SSAT( $\Theta$ ) formula  $\Phi$ , if the outermost quantifier is a threshold quantifier  $\Theta^{\triangleright p}$ , we omit it and solve the remaining formula  $\Phi'$ . (Otherwise, we solve  $\Phi$  directly.) Once  $\text{Pr}[\Phi']$  is computed, we simply check if  $\text{Pr}[\Phi'] \triangleright p$  and return the corresponding truth or

falsity. Moreover, we modify the definition of quantification level as follows: Consider the formula  $Q_1, \dots, Q_n.\phi$ . For a randomly or existentially quantified variable at quantification  $Q_i$ , let  $k$  be the number of the encountered alternations in one of the forms  $\exists\forall$ ,  $\forall\exists$ , and  $\Theta\forall$ , when traversing from  $Q_1$  to  $Q_i$ . Then the quantification level of that variable is defined as  $k + 1$ . E.g., for the formula  $\forall x_1, \exists x_2, \forall x_3, \Theta, \forall x_4.\phi$ , the quantification levels of  $x_1$ ,  $x_2$ ,  $x_3$ , and  $x_4$  are 1, 2, 3, and 4, respectively.

The literals in a clause in the matrix are then partitioned into groups with respect to the newly defined quantification levels. The formulas are then solved recursively on the quantification levels in the same way as SSAT, except that for the random quantifier  $Q_i$  with an outer threshold quantifier  $Q_{i-1} = \Theta^{>p}$ , the probability should be mapped to a Boolean value according to  $\Theta^{>p}$ .

To effectively prune the search space, `Claussat` incorporates several pruning techniques for the alternations  $\exists\forall$  and  $\forall\exists$ . To handle these techniques, we follow the same implementation as `Claussat` when encountering the  $\exists\forall$  and  $\forall\exists$  alternations. For the  $\exists\Theta\forall$  alternation, as stated in Lemma 2, since the probability is preserved under the re-ordering between the threshold and existential quantifiers, the techniques for  $\exists\forall$  are applicable in  $\exists\Theta\forall$ . For the  $\forall\Theta\forall$  alternation, we disable all the pruning techniques proposed in `Claussat`.

## 5 Encoding Application Problems

This section first gives some background on discrete-time Markov Chains (DTMCs) and bounded probabilistic computation tree logic (BPCTL). We then show how to encode the model-checking problem with SSAT( $\Theta$ ). Further, we further introduce the parameter synthesis problem and demonstrate its SSAT( $\Theta$ ) encoding.

### DTMCs and BPCTL

**Discrete-Time Markov Chains** A *discrete-time Markov chain* can be viewed as a state transition system where each transition takes place with a transition probability. It is a common model used to describe the behavior of a probabilistic system.

**Definition 4** (Discrete-Time Markov Chain.). A discrete-time Markov Chain is a tuple  $\mathcal{M} = (\mathcal{S}, s_0, \mathbf{P}, \mathcal{AP}, \mathcal{L})$ , where

- $\mathcal{S}$  is a finite set of states and  $s_0 \in \mathcal{S}$  is a initial state.
- $\mathbf{P}$  is a probabilistic matrix representing the transition probabilities, where  $\mathbf{P} : \mathcal{S} \times \mathcal{S} \rightarrow [0, 1]$  and  $\sum_{s' \in \mathcal{S}} \mathbf{P}(s, s') = 1$  for each  $s \in \mathcal{S}$ .
- $\mathcal{AP}$  is a set of atomic propositions and  $\mathcal{L} : \mathcal{S} \rightarrow 2^{\mathcal{AP}}$  is a labelling function for states.

A  $k$ -path  $\pi$  in a DTMC is a finite sequence of states  $(s_1, \dots, s_k)$  with length  $k$ . We use  $\pi(i)$  to denote the state  $s_i$ . The set of all the  $k$ -path starting with a set  $s$  is denoted as  $Path_k(s)$ . For simplicity, we use *path* to refer to a  $k$ -path in the sequel. The probability of  $\pi$ , denoted as  $\Pr[\pi]$ , is the product of the probabilities of the involved transitions  $\prod_{i=1}^{k-1} \mathbf{P}(s_i, s_{i+1})$ .

**Bounded Probabilistic Computation Tree Logic** Given a probabilistic system and a property, the problem of probabilistic model checking is to check if the system satisfies the given property. In this work, the system description we consider is the DTMC and the properties we focus on are specified in *bounded probabilistic computation tree logic* (BPCTL).

BPCTL is a bounded-step fragment of probabilistic computation tree logic. The syntax of BPCTL is defined as follows:

**Definition 5** (BPCTL Syntax). *State formula*, denoted as  $\phi$ , and *path formula*, denoted as  $\psi$ , are two main components of a BPCTL formula with the following definition:

$$\begin{aligned} \phi &:= \top \mid a \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid P^{\bowtie p}[\psi] \\ \psi &:= F^{\leq k}\phi \mid \phi_1 U^{\leq k}\phi_2 \mid X\phi \end{aligned}$$

The  $a$  is an atomic proposition. A BPCTL formula is a state formula.

The semantics of BPCTL can then be defined with respect to a DTMC as follows.

**Definition 6** (BPCTL Semantics). Let the DTMC  $\mathcal{M} = (\mathcal{S}, s_0, \mathbf{P}, \mathcal{AP}, \mathcal{L})$ . The semantics of the satisfaction relation are defined by:

$$\begin{aligned} \mathcal{M}, s &\models \top && \text{for all } s \in \mathcal{S}. \\ \mathcal{M}, s &\models a && \text{if } a \in \mathcal{L}(s). \\ \mathcal{M}, s &\models \neg\varphi && \text{if } \mathcal{M}, s \not\models \varphi. \\ \mathcal{M}, s &\models \varphi_1 \wedge \varphi_2 && \text{if } \mathcal{M}, s \models \varphi_1 \text{ and } \mathcal{M}, s \models \varphi_2. \\ \mathcal{M}, \pi &\models F^{\leq k}\varphi && \text{if } \exists i \in [1..k] : \mathcal{M}, \pi(i) \models \varphi. \\ \mathcal{M}, \pi &\models \varphi_1 U^{\leq k}\varphi_2 && \text{if } \exists i \in [1..k] : \mathcal{M}, \pi(i) \models \varphi_2 \text{ and } \\ &&& \mathcal{M}, \pi(j) \models \varphi_1 \forall j \in [1, i-1]. \\ \mathcal{M}, \pi &\models P^{\bowtie p}[\psi] && \text{if } \sum_{\pi \in Path(s)} \Pr[\pi] \bowtie p. \end{aligned}$$

With Definition 6, we say the DTMC  $\mathcal{M}$  satisfies  $\phi$  if  $\mathcal{M}, s_0 \models \phi$ . In the following, we use *model-checking* to refer to the model-checking problem of a DTMC specified in BPCTL. We note that when there is a single outermost  $P$  operator in  $\varphi$ , the state-of-the-art probabilistic model checkers often allow it to be  $P=?$ , which is to return the probability rather than a Boolean value. The following SSAT( $\Theta$ ) encodings can also handle such an extension.

### BPCTL Model Checking for DTMC

The main idea is to somehow map the  $P$  operator in BPCTL to the threshold quantifier. In consequence, a BPCTL formula nested with  $P$  operators would correspond to a pure-counting SSAT( $\Theta$ ) formula. Before delving into the details, we shall first transform  $\varphi$  into the form where only  $P^{>p}$  and  $P^{\geq p}$  are allowed for the  $P$  operator. This can be done by replacing each  $P^{<p}$  and  $P^{\leq p}$  with  $\neg P^{\geq p}$  and  $\neg P^{>p}$ , respectively. In the following, we assume that a DTMC  $\mathcal{M} = (\mathcal{S}, s_0, \mathbf{P}, \mathcal{AP}, \mathcal{L})$  and a BPCTL formula  $\varphi$  of the mentioned form are given. Also, given an atomic proposition  $a$  and any state  $s$ , let  $F_a$  be the Boolean function such that  $F_a(s) = \top$  if and only if  $a \in \mathcal{L}(s)$ .

**BPCTL Encoding** For readability, we assume the state space  $\mathcal{S} = [0..|\mathcal{S}| - 1]$  and let  $n = \lfloor \log_2 |\mathcal{S}| \rfloor + 1$ . We

refer the *state variables* to a vector of  $n$ -bits Boolean variables  $X = (x_1, \dots, x_n)$ , who takes integer values in  $\mathcal{S}$ . We allocate one random-quantified variable for each transition  $(s, s')$ , denoted as  $x_r(s, s')$ , where  $x_r(s, s') = \top$  with non-zero probability  $\mathbf{P}(s, s')$ . We let the set of the allocated random-quantified variables be  $X_r = \{x_r(s, s') \mid \forall s, s' \in \mathcal{S}, \mathbf{P}(s, s') > 0\}$ .

To derive an equivalent SSAT( $\Theta$ ) formula, we recursively define the *characteristic function* in the form of SSAT( $\Theta$ ) formulas, for both the state formula and path formula. In order for that, we need a *probabilistic transition relation*  $T(X_s, X_{s'})$  over the current state variables  $X_s$  and next state variables  $X_{s'}$ . Intuitively, given a current state  $s$  and a next state  $s'$ , we want the satisfying probability  $\Pr[T(s, s')] = \mathbf{P}(s, s')$ . The following is the considered probabilistic transition relation  $T(X_s, X_{s'})$ :

$$\forall X_r. \bigwedge_{s, s' \in \mathcal{S}} ((X_s = s \wedge X_{s'} = s') \rightarrow x_r(s, s')) \wedge \bigvee_{s, s' \in \mathcal{S}, \mathbf{P}(s, s') > 0} (X_s = s \wedge X_{s'} \leftrightarrow s') \quad (6)$$

With Eq. (6), we can expand time-frame for a  $k$ -path through proper variable-renaming as follows:

$$T^{k-1}(X_s^1, \dots, X_s^k) = \bigwedge_{i=1}^{k-1} T(X_s^i, X_s^{i+1}), \quad (7)$$

where the state variables  $X_s^i$  corresponds to the  $i$ -th state in the path. Given a path  $(s_1, \dots, s_k)$ ,  $\Pr[T^{k-1}(s_1, \dots, s_k)]$  equals  $\prod_{i=1}^{k-1} \mathbf{P}(s_i, s_{i+1})$ .

We use  $C_\psi$  to represent the characteristic function of  $\psi$ , either a state formula or a path formula. Recall that we already transformed the given BPCTL formula with only  $P^{\triangleright p}$  for the  $P$  operator. We give the rules, one for each composition rule in Definition 5, to construct the characteristic function, which is an opened pure-counting formula, as follows:

$$\begin{aligned} C_\top(X_s) &:= \top \\ C_a(X_s) &:= F_a(X_s) \\ C_{\neg\varphi}(X_s) &:= \neg C_\varphi(X_s) \\ C_{\varphi_1 \wedge \varphi_2}(X_s) &:= C_{\varphi_1}(X_s) \wedge C_{\varphi_2}(X_s) \\ C_{P^{\triangleright p}[\psi]}(X_s) &:= \Theta^{\triangleright p/2^t} C_\psi(X_s) \\ C_{X\varphi}(X_s) &:= \mathfrak{A}^{0.5} X'_s.T(X_s, X'_s) \\ C_{F^{\leq k}\varphi}(X_s) &:= \mathfrak{A}^{0.5} X_s^2, \dots, \mathfrak{A}^{0.5} X_s^{k+1}. \\ &\quad T^k(X_s, X_s^2, \dots, X_s^{k+1}) \wedge \\ &\quad (C_\varphi(X_s) \vee \bigvee_{i=2}^{k+1} C_\varphi(X_s^i)) \\ C_{\varphi_1 U^{\leq k}\varphi_2}(X_s) &:= \mathfrak{A}^{0.5} X_s^2, \dots, \mathfrak{A}^{0.5} X_s^{k+1}. \\ &\quad T^k(X_s, X_s^2, \dots, X_s^{k+1}) \wedge \\ &\quad (C_{\varphi_2}(X_s) \vee (C_{\varphi_1}(X_s) \wedge \\ &\quad \bigvee_{i=2}^{k+1} (C_{\varphi_2}(X_s^i) \wedge \bigwedge_{j=2}^{i-1} C_{\varphi_1}(X_s^j)))) \end{aligned}$$

Intuitively, for a state formula  $\varphi$  and a state  $s$ ,  $C_\varphi(s) = \top$  if and only if  $s$  satisfies  $\varphi$ . For a path formula  $\psi$ ,  $\Pr[C_\psi(s)]$  gives the sum of the probabilities the paths in  $Path_k(s)$  satisfying  $\psi$ . In the construction for the  $P$  operator, the factor  $1/2^t$ , with  $t$  being the number of the extra quantified state variables in the construction of the outermost characteristic

functions, is a scaling factor for the original threshold  $p$  since we use probability 0.5 for each state variable in the construction. On the other hand, in the constructions involving multiple characteristic functions, namely,  $C_{\phi_1 \wedge \phi_2}$ ,  $C_{F^{\leq k}\varphi}$ , and  $C_{\varphi_1 U^{\leq k}\varphi_2}$ , the quantified variables for transition relations have to be duplicated among the characteristic functions so that their quantified variables are disjoint.

The entire working flow is as follows: We first derive the characteristic function of the BPCTL formula  $C_\phi(X_s)$  and then transform it into an equivalent SSAT( $\Theta$ ) in the prenex form  $\Theta, \mathfrak{A}, \dots, \Theta, \mathfrak{A}.\phi$  with the rules stated in the following three propositions.

**Proposition 1** (Conjunction). *Let  $\Phi_1 = \Theta^{\triangleright p_1}, \mathfrak{A} Y, Q_1.\phi_1$  and  $\Phi_2 = \Theta^{\triangleright p_2}, Q_2.\phi_2$  be two pure-counting formulas, possibly with a common set of free variables  $X$ .  $Q_1$  is pure-counting. If all the variables in  $Y$  are not quantified in  $\Phi_2$ , then*

$$\Phi_1 \wedge \Phi_2 = \Theta^{\triangleright p_1}, \mathfrak{A} Y.((Q_1.\phi_1) \wedge (\Theta^{\triangleright p_2}, Q_2.\phi_2)).$$

**Proposition 2** (Disjunction). *We are given two pure-counting SSAT( $\Theta$ ) formulas  $\Phi_1 = \Theta^{\triangleright p_1}, \mathfrak{A} Y, Q_1.\phi_1$  and  $\Phi_2 = \Theta^{\triangleright p_2}, Q_2.\phi_2$ , possibly with common set of free variables  $X$ , where  $Y$  is a set of random-quantified variables. If all the variables in  $Y$  are not quantified in  $\Phi_2$ , then*

$$\Phi_1 \vee \Phi_2 = \Theta^{\triangleright p_1}, \mathfrak{A} Y.((Q_1.\phi_1) \vee (\Theta^{\triangleright p_2}, Q_2.\phi_2)).$$

**Proposition 3** (Negation). *Let  $\Phi = \Theta^{\triangleright p} \mathfrak{A}^{p'} V Q.\phi$  be a pure-counting formula, possibly with free variables  $X$ .  $Q$  is pure-counting. Then its negations  $\neg\Phi = \Theta^{\triangleright(1-p)} \mathfrak{A}^{p'} V \neg Q.\phi$ .*

Finally, the resulting SSAT( $\Theta$ ) formula is derived by assigning the state variables  $X_s$  to the initial state  $s_0$  and converting the matrix into a CNF formula by Tseitin transformation:

$$Q, \exists X_D.\phi', \quad (8)$$

where  $\phi'$  is in CNF,  $X_D$  is the set of extra definition variables introduced by the Tseitin transformation, and  $Q$  is pure-counting.

Recall that state-of-the-art probabilistic model checkers allow  $P^{=?}$  for the single outermost  $P$  operator. To handle that extension, we can simply remove the outermost threshold quantifier of  $Q$ , say  $Q'$ , in Eq. (8). The resulting formula is  $Q', \exists X_D.\phi'$ .

## Parameter Synthesis with SSAT( $\Theta$ )

In most applications, the DTMCs are parameterized by several parameters. The probability of a certain BPCTL property is also dependent on the parameters. Take the Crowds (Shmatikov 2004) protocol for message transmission, which provides a probabilistic guarantee of the anonymity of the sender. The protocol assumes there are certain numbers of *good people* and *bad people*, where the good people cannot communicate with each other and the bad people would cooperate with each other. To provide the probabilistic anonymity guarantee, the message is passed to a randomly chosen bad person with probability  $p$  and to a randomly chosen good person with probability  $1 - p$ . The

only method the bad people can acquire information about the real sender is to observe the identity of the one who passed the message to the bad people. One important parameter is the size of the *good people* or called *crowd size*, which affects the probabilistic bound of anonymity — the greater the crowd size the better anonymity *Crowds* can provide. Under such DTMC, we ask the *parameter synthesis* problem: Given a DTMC and a BPCTL property, how to determine the size of the DTMC, in terms of the parameters, so that the property is maximized?

Note that as long as we have the transition relation of the parameterized DTMC, the parameter synthesis problem can be naturally encoded using SSAT( $\Theta$ ). We first construct the transition relation of each instantiated DTMC with a certain parameter value. Suppose the set of variables for the parameter is  $X_p$ , which can take values from the finite set of range  $R$ . Suppose the maximum value in  $R$  is  $r_{\max}$  and the minimum value in  $R$  is  $r_{\min}$ . We allocate  $\lceil \log_2(r_{\max} - r_{\min} + 1) \rceil + 1$  number of variables for  $X_p$ . We use  $T_v$  to denote the transition relation of DTMC with the parameter value parameter being  $v \in R$ . Then the parametric transition relation is constructed as follows:

$$T_p(X_p, X_s, X_{s'}) = \bigwedge_{v \in R} ((X_p = v) \rightarrow T_v). \quad (9)$$

We follow the same BPCTL encoding procedure in the previous subsection except that we use the parametric transition relation in Eq. (9) instead. Suppose we have the SSAT( $\Theta$ ) formula in the form in Eq. Eq. (8). Then the parameter synthesis problem can be encoded as:

$$\exists X_p, \forall X, \mathcal{Q}, \exists X_D. \phi' \wedge \bigwedge_{v \in R} (X_p = v), \quad (10)$$

with  $\mathcal{Q}$  being pure-counting and  $X$  being some outermost state variables.

## 6 Experimental Results

We implemented our SSAT( $\Theta$ ) solver, named *ClauSSat*( $\Theta$ ), in the C++ language by extending *ClauSSat* (Chen, Huang, and Jiang 2021). The experiments were conducted on a Linux machine with 2.2 GHz Intel Xeon CPU and 128 GB RAM. Two benchmark sets were experimented for evaluation. The first set includes instances from the case study of parameter synthesis on the DTMC *Crowds* (Shmatikov 2004) protocol.<sup>1</sup> The second set includes instances converted from SSAT formulas. A 1000-second time limit was imposed on solving each instance.

<sup>1</sup>We note that although BPCTL model checking for DTMC can be encoded in SSAT( $\Theta$ ), the converted instances are often too large to be solved. For this problem, there are dedicated model checkers, such as *Prism* (Kwiatkowska, Norman, and Parker 2002), *Storm* (Dehnert et al. 2017), and *EPMC* (Hahn et al. 2014), for more direct and effective solving. Hence, we focused on evaluating instances with complex queries that cannot be handled by existing model checkers.

Range	$k$	Formula		Run Time (s)
		#Cls	#Vars	
[2..4]	1	6112	2203	9.13
	2	6763	2428	0.07
	3	8697	3113	83.78
	4	9318	3337	554.75
	5	10093	3626	TO
[2..6]	1	6923	2489	25.00
	2	9961	3541	277.54
	3	10561	3761	394.02
	4	11625	4139	658.14
	5	11786	4207	TO

Table 1: Results on instances of parameter synthesis.

## Evaluation on Instances of Parameter Synthesis

To create the benchmark instances of parameter synthesis, we adopted the *Crowds* (Shmatikov 2004) protocol. A DTMC model was first specified in the prism model format (Kwiatkowska, Norman, and Parker 2002), then converted to a transition relation in a bit-vector form of the QF-BV SMT format, then further bit-blasted using the SMT solver *Boolector* (Niemetz, Preiner, and Biere 2014).<sup>2</sup> Given the BPCTL property to be checked and the transition relation, the SSAT( $\Theta$ ) formula was created with the approaches introduced in Section 5.

Recall that *Crowds* is a protocol for providing anonymity of the actual sender. The protocol assumes that there is only one actual sender among  $t$  potential senders, for  $t$  being the crowd size. We define the “safe <sub>$i$</sub> ” property as follows.

$$\text{safe}_i = P^{>0.5}[F^{\leq m}(\text{observe}_i < 1)], \quad (11)$$

where  $\text{observe}_i$  is an integer state variable and  $(\text{observe}_i < 1)$  asserts the actual sender  $i$  is not observed. It asserts that the probability that the adversary does not observe the actual sender  $i$  in the future  $m$ -steps is greater than 0.5.

We say that the system is in a safe state for sender  $i$  if the system state satisfies property  $\text{safe}_i$ .<sup>3</sup> For the BPCTL property to be checked, we extend the  $\text{safe}_i$  property to

$$k\text{-step safe}_i = P^{=?}[F^{\leq k}(\text{safe}_i)]. \quad (12)$$

Thereby, with Eq. (10), the problem is to search for a crowd size value such that the  $k$ -step  $\text{safe}_i$  property is maximized starting from the initial state of the DTMC. In the experiment, we checked “safe<sub>0</sub>” under  $m = 5$ . All the created formulas share the same prefix form  $\exists\text{-}\forall\text{-}\Theta\text{-}\forall\text{-}\exists$ .

The results are shown in Table 1, where columns “Range,” “ $k$ ,” and “Run Time (s)” report the range of the parameter values,  $k$  of the  $k$ -step safe property, and the time

<sup>2</sup>We disabled the *-vs* option in *Boolector* as it may perform aggressive SAT-based encoding that might not be sound in SSAT( $\Theta$ ). The instances were only simplified using Boolean constraint propagation.

<sup>3</sup>The  $\text{safe}_i$  property can be encoded as an  $\exists\text{-}\forall\text{-}\exists$  quantified SSAT formula.

Family	Instance	#R	Formula		Run Time (s)	
			#Cls	#Vars	SSAT	SSAT( $\Theta$ )
tlc	depth-7	2	7491	2809	0.56	0.47
	depth-8	2	8430	3160	0.73	0.61
	depth-9	2	9360	3511	0.92	0.78
gttt	2_2_0010	9	2958	1133	10.05	26.85
	2_2_001020	9	2814	1085	144.16	310.16
	2_2_000111	9	3214	1165	145.36	298.59
Robot	8	39	30444	8892	41.97	61.14
	9	45	34078	9880	50.32	34.35
	10	50	37712	10868	63.78	286.01
stracomp	25.9	25	2278	779	0.45	309.74
	30.9	30	2754	934	0.38	TO
	35.9	35	3227	1089	0.37	TO

Table 2: Results of evaluation on general SSAT( $\Theta$ ) instances converted from SSAT instances.

spent for solving each case, respectively, and the numbers of clauses and variables of the SSAT( $\Theta$ ) formula are reported in columns “#Cls” and “#Vars,” respectively. The cases that failed to be solved within the time limit are denoted with “TO.” Unsurprisingly, the running time increases drastically as the  $k$  increases since the transition relation and the variables have to be duplicated as mentioned in Section 5. When comparing the solving time of the two different configurations of range, we notice that the formulas failed to be solved within the time limit when  $k = 5$  for both configurations. We observed that for the [2..4] configuration, the run time seems to increase more acutely as  $k$  increases than the [2..6] configuration. As two (resp. three) existential variables are allocated (for  $X_p$  mentioned in the subsection of Parameter Synthesis) for the range [2..4] (resp. [2..6]), it would be interesting to investigate the effect of the number of these outermost existential variables on the solving efficiency for the parameter synthesis benchmarks.

Finally, we remark that although the current SSAT( $\Theta$ ) solver can solve most of the generated instances, the step size  $k$  and the range size of the parameter values are relatively small, and the efficiency is sensitive to the increase of the range size. In order to alleviate the high computation complexity, it may be crucial to develop specialized SSAT( $\Theta$ ) solver and preprocessing techniques.

### Evaluation on Instances of SSAT Extension

To study the impact of the threshold quantifier on SSAT instances, we assess the solver’s efficiency by evaluating SSAT( $\Theta$ ) instances generated from existing SSAT benchmarks, taken from (Chen, Huang, and Jiang 2021).<sup>4</sup> The conversion was done by randomly selecting a random quantification block  $\forall V$ , bi-partitioning variables  $V$  into  $V_1$  and  $V_2$ , and inserting a randomly generated threshold quantifier  $\Theta^{>p}$  so that the quantification becomes  $\forall V_1, \Theta^{>p}, \forall V_2$ . The bi-partition of variables  $V$  is also determined randomly.

Specifically, we selected four families of benchmarks where all the cases in each family are solved by ClauSSat within the time limit. To examine the effect of the threshold quantifier, we compare the performance on the SSAT( $\Theta$ ) instances with that on the original SSAT instances. The results are shown in Table 2, where three representative instances for each family are listed due to space limit. The column “#R” represents the original number of random quantified variables in the inserted quantification and the column “SSAT” (resp. “SSAT( $\Theta$ )”) reports the time spent on solving the SSAT (resp. SSAT( $\Theta$ )) instance.

By comparing the performance on SSAT and SSAT( $\Theta$ ) instances, we observed that in most cases within the families, the insertion of threshold quantifier deteriorates the performance as expected. The degradation is expected because the threshold quantifier increases the number of quantification levels, and the current implementation does not equip efficient pruning techniques for the alternation  $\forall$ - $\exists$ - $\forall$ . Also, observe that the number of random-quantified variables may play a role in affecting the run time due to the increased complexity introduced by the threshold quantifier. This effect is especially significant in the family *stracomp*, where all the listed SSAT instances can be solved in less than one second, while their SSAT( $\Theta$ ) counterparts fail to be solved within the time limit as the number of random-quantified variables attains or exceeds 30. However, there is an exception in instance 9 of the *Robot* family, whose SSAT( $\Theta$ ) instance, in contrast, takes less time to be solved. Some other factors contributing to the hardness of computation remain to be further investigated.

## 7 Conclusions and Future Work

This work presents SSAT( $\Theta$ ), which unifies the decision and function queries of SSAT by augmenting SSAT with the threshold quantifier. SSAT( $\Theta$ ) subsumes the counting formulas and can encode problems in the Polynomial and Counting Hierarchies. For a practical case study, we encode BPTCL model checking and the parameter synthesis prob-

<sup>4</sup>Available at <https://github.com/NTU-ALComLab/ClauSSat>.

lem of DTMCs into SSAT( $\Theta$ ).

Experiments demonstrate the feasibility of extending ClauSSat for solving SSAT( $\Theta$ ) instances. For future work, we plan to develop more advanced SSAT( $\Theta$ ) solvers with advanced pruning techniques and explore more applications.

Also, extending DSSAT (Lee and Jiang 2021) with the threshold quantifier could be useful.

## Acknowledgements

This work was supported in part by the National Science and Technology Council of Taiwan under Grant 111-2923-E-002-013-MY3.

## References

- Chen, P.-W.; Huang, Y.-C.; and Jiang, J.-H. R. 2021. A Sharp Leap from Quantified Boolean Formula to Stochastic Boolean Satisfiability Solving. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 3697–3706.
- Chou, Y.-M.; Chen, Y.-C.; Wang, C.-Y.; and Huang, C.-Y. 2016. MajorSat: A SAT Solver to Majority Logic. In *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC)*, 480–485.
- Dehnert, C.; Junges, S.; Katoen, J.-P.; and Volk, M. 2017. A Storm Is Coming: A Modern Probabilistic Model Checker. In *Proceedings of the International Conference on Computer Aided Verification (CAV)*, 592–600.
- Fan, Y.-W.; and Jiang, J.-H. R. 2023. SharpSSAT: A Witness-Generating Stochastic Boolean Satisfiability Solver. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 3949–3958.
- Ghosh, B.; Basu, D.; and Meel, K. S. 2021. Justicia: A Stochastic SAT Approach to Formally Verify Fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 7554–7563.
- Hahn, E. M.; Li, Y.; Schewe, S.; Turrini, A.; and Zhang, L. 2014. iscasMc: A Web-Based Probabilistic Model Checker. In *Proceedings of the International Symposium of Formal Methods (FM)*, 312–317.
- Hsieh, C.-H.; and Jiang, J.-H. R. 2022. Encoding Probabilistic Graphical Models into Stochastic Boolean Satisfiability. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 1834–1842.
- Kwiatkowska, M.; Norman, G.; and Parker, D. 2002. PRISM: Probabilistic symbolic model checker. In *International Conference on Modelling Techniques and Tools for Computer Performance Evaluation*, 200–204.
- Lee, N.-Z.; and Jiang, J.-H. R. 2021. Dependency Stochastic Boolean Satisfiability: A Logical Formalism for NEXP-TIME Decision Problems with Uncertainty. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 3877–3885.
- Lee, N.-Z.; Wang, Y.-S.; and Jiang, J.-H. R. 2017. Solving Stochastic Boolean Satisfiability under Random-Exist Quantification. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 688–694.
- Lee, N.-Z.; Wang, Y.-S.; and Jiang, J.-H. R. 2018. Solving Exist-Random Quantified Stochastic Boolean Satisfiability via Clause Selection. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 1339–1345.
- Littman, M. L.; Goldsmith, J.; and Mundhenk, M. 1998. The Computational Complexity of Probabilistic Planning. *Journal of Artificial Intelligence Research (JAIR)*, 9(1): 1–36.
- Littman, M. L.; Majercik, S. M.; and Pitassi, T. 2001. Stochastic Boolean Satisfiability. *Journal of Automated Reasoning (JAR)*, 27(3): 251–296.
- Majercik, S. M.; and Boots, B. 2005. DC-SSAT: A Divide-and-Conquer Approach to Solving Stochastic Satisfiability Problems Efficiently. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 416–422.
- Majercik, S. M.; and Littman, M. L. 2003. Contingent Planning under Uncertainty via Stochastic Satisfiability. *Artificial Intelligence (AI)*, 147(1-2): 119–162.
- Niemetz, A.; Preiner, M.; and Biere, A. 2014. Boolector 2.0. *Journal on Satisfiability, Boolean Modeling and Computation*, 9(1): 53–58.
- Oztok, U.; Choi, A.; and Darwiche, A. 2016. Solving PP<sup>PP</sup>-Complete Problems Using Knowledge Compilation. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 94–103.
- Papadimitriou, C. H. 1985. Games Against Nature. *Journal of Computer and System Sciences (JSCC)*, 31(2): 288–301.
- Salmon, R.; and Poupart, P. 2020. On the Relationship Between Satisfiability and Markov Decision Processes. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 1105–1115.
- Shmatikov, V. 2004. Probabilistic Model Checking of an Anonymity System. *Journal of Computer Security*, 12(3-4): 355–377.
- Stockmeyer, L. J. 1976. The Polynomial-Time Hierarchy. *Theoretical Computer Science (TCS)*, 3(1): 1–22.
- Wagner, K. W. 1986. The Complexity of Combinatorial Problems with Succinct Input Representation. *Acta Informatica*, 23(3): 325–356.
- Wang, H.-R.; Tu, K.-H.; Jiang, J.-H. R.; and Scholl, C. 2022. Quantifier Elimination in Stochastic Boolean Satisfiability. In *Proceedings of International Conference on Theory and Applications of Satisfiability Testing (SAT)*, 23:1–23:17.



## Proofs

In this appendix, we provide the omitted proofs of propositions and theorems in the paper.

*Proof of Theorem 1.* We shall prove this by showing that  $\text{SSAT}(\Theta)$  is in PSPACE and is PSPACE-hard.

As the SSAT problem is in PSPACE (Papadimitriou 1985), the  $\text{SSAT}(\Theta)$  problem is also in PSPACE as the additional threshold quantifier does not increase space complexity. Consider a polynomial-space bounded algorithm to  $\text{SSAT}(\Theta)$  as shown in Algorithm 1. Let  $S(i)$  be the space needed for the  $i$ -th recursive call, which equals the space needed within the  $i$ -th call plus the space needed for the  $(i+1)$ -th recursive calls. Then the total space needed for the whole formula is  $S(1)$ . Suppose the number of clauses is  $m$  and the number of bits for storing a probability is  $b$ . For the base case in Line 3, we have to evaluate the truth value of  $\phi$  over an assignment so the space needed  $S(n)$  is  $O(mn + b)$ , where  $O(mn)$  is the space needed for the induced formula and  $O(b)$  is the space required for the probability. Now consider the  $i$ -th recursive call. For the case of threshold quantifier in Line 5, the space  $S(i) = S(i+1) + O(b)$ . For the case of existential and randomized quantifiers from Line 6 to Line 12, we reuse the space required for computing  $p_0$  when computing  $p_1$ . Therefore, the space  $S(i) = S(i+1) + O(mn + b)$ . It turns out that the total space required  $S(1)$  is  $O(mn^2 + bn)$ .

$\text{SSAT}(\Theta)$  is PSPACE-hard because it contains SSAT as a special case.  $\square$

*Proof of Proposition 1.* We shall prove this by showing that under any assignment over  $X$ , LHS is equivalent to RHS. Consider the induced formula on an assignment  $\sigma_X$ . Let LHS  $\Phi' = \Phi_1|_{\sigma_X} \wedge \Phi_2|_{\sigma_X}$  and RHS  $\Phi'' = \Theta^{\triangleright p_1}, \forall Y.((\mathcal{Q}_1.\phi_1|_{\sigma_X}) \wedge (\Theta^{\triangleright p_2}, \mathcal{Q}_2.\phi_2|_{\sigma_X}))$ , respectively.  $(\Rightarrow) : \Phi' = \top$  if and only if

$$\Theta^{\triangleright p_1}, \forall Y, \mathcal{Q}_1.\phi_1|_{\sigma_X} = \top \quad (13)$$

and

$$\Theta^{\triangleright p_2}, \mathcal{Q}_2.\phi_2|_{\sigma_X} = \top \quad (14)$$

By Eq. (14), we have  $\Phi'' = \Theta^{\triangleright p_1}, \forall Y, \mathcal{Q}_1.\phi_1|_{\sigma_X}$ . By Eq. (13), we have  $\Phi'' = \top$ .

$(\Leftarrow) : \Phi'' = \top$  only if Eq. (14) holds. Otherwise,

$$\Phi'' = \Theta^{\triangleright p_1}, \forall Y.((\mathcal{Q}_1.\phi_1|_{\sigma_X}) \wedge \perp) = \perp,$$

leading to a contradiction. It follows that  $\Phi'' = \Theta^{\triangleright p_1}, \forall Y, \mathcal{Q}_1.\phi_1|_{\sigma_X}$ . By the assumption, we have that Eq. (13) also holds.  $\square$

*Proof of Proposition 2.* The proof is similar to that in Proposition 1.  $\square$

*Proof of Proposition 3.* We prove the case when there is only one variable  $v$  in  $V$ . The proof can be generalized to a set of variables  $V$ . We shall prove this by showing that under any assignment  $\sigma_X$ , LHS is equivalent to RHS. We only consider the case with the outermost threshold quantifier being

$\Theta^{>p}$  as the case of  $\Theta^{\geq p}$  is similar to  $\Theta^{>p}$ . Let  $\phi|_{\sigma_X} = \phi'$ . Let  $\Phi' = \neg\Phi|_{\sigma_X}$  and  $\Phi'' = \Theta^{>(1-p)} \forall^{p'} v \neg \mathcal{Q}.\phi'$ . We have

$$\begin{aligned} \Phi' &= \neg(\Pr[\forall^{p'} v, \mathcal{Q}.\phi'] > p) \\ &= \Pr[\forall^{p'} v, \mathcal{Q}.\phi'] \leq p \\ &= (\Pr[\mathcal{Q}.\phi'|_v] \cdot p' + \Pr[\mathcal{Q}.\phi'|_{\neg v}] \cdot (1 - p')) \leq p \\ &= ((1 - \Pr[\mathcal{Q}.\phi'|_v]) \cdot p' + (1 - \Pr[\mathcal{Q}.\phi'|_{\neg v}]) \cdot (1 - p')) \\ &> 1 - p \\ &= (\Pr[\neg \mathcal{Q}.\phi'|_v] \cdot p' + \Pr[\neg \mathcal{Q}.\phi'|_{\neg v}] \cdot (1 - p')) > 1 - p \\ &= \Theta^{>(1-p)}, \forall^{p'} v, \neg \mathcal{Q}.\phi' \\ &= \Phi'' \end{aligned}$$

$\square$