

# 1 Synthesizing Skolem functions: A view from theory and 2 practice

3 S. Akshay · Supratik Chakraborty

4  
5 (Pre-print)

6 **Abstract** Skolem functions play a central role in logic, from helping eliminate  
7 quantifiers in first order logic formulas to providing functional implementations of  
8 relational specifications. While their existence follows from classical results in logic,  
9 less is known about how to compute them effectively and efficiently (whenever  
10 such computation is possible). The problem of computing or synthesizing Skolem  
11 functions from relational specifications, however, continues to show up in many  
12 interesting applications. Recently, a rich line of work has considered theoretical  
13 and practical aspects of the problem in a restricted setting, namely synthesis of  
14 Boolean Skolem functions from Boolean relational specifications. In this article we  
15 take an indepth look into this fascinating problem and its various implications,  
16 from general theoretical and complexity results to practical algorithms, and also  
17 draw interesting connections to the knowledge representation literature.

18 **Keywords** Boolean functional synthesis, Skolem functions, expansion-based  
19 algorithms

## 20 1 Introduction

21 The genesis of Skolem functions dates back to 1920, when the Norwegian mathe-  
22 matician, Thoralf Albert Skolem, gave a simplified proof of a landmark result in  
23 logic, now known as the *Löwenheim-Skolem* theorem. Leopold Löwenheim had al-  
24 ready proved this theorem in 1915. However, Skolem’s 1920 proof was significantly  
25 simpler and made use of a key observation that can be summarized as follows<sup>1</sup>.  
26 *For every first order logic formula  $\exists y \varphi(x, y)$ , the choice of  $y$  that makes  $\varphi(x, y)$  true*  
27 *(if at all) depends on  $x$  in general. This dependence can be thought of as implicitly*  
28 *defining a function that gives the “right” value of  $y$  for every value of  $x$ . If  $F$  denotes*  
29 *a fresh function symbol, the second order sentence  $\exists F \varphi(x, F(x))$  formalizes this de-*  
30 *pendence explicitly. Thus, the second order sentence  $\exists F \forall x (\exists y \varphi(x, y) \Rightarrow \varphi(x, F(x)))$*   
31 *always holds. Since the implication trivially holds in the other direction too, we*  
32 *have  $\exists F \forall x (\exists y \varphi(x, y) \Leftrightarrow \varphi(x, F(x)))$ .*

Indian Institute of Technology Bombay, India

<sup>1</sup> We assume the reader is familiar with basic notation and terminology of first order logic.

Let  $\xi_1$  and  $\xi_2$  denote the first order formulas  $\exists y \varphi(x, y)$  and  $\varphi(x, F(x))$  respectively referred to above. The following points are worth noting.

- While  $\xi_2$  has one less existential quantifier than  $\xi_1$ , the signature of  $\xi_2$  has one more function symbol than the signature of  $\xi_1$ . Thus, an existential quantifier has been traded off, so to say, for a function symbol.
- Although  $\xi_1$  and  $\xi_2$  are not semantically equivalent, there is an interpretation of  $F$  such that for every assignment of the free variable  $x$ , the formula  $\xi_1$  is satisfiable iff  $\xi_2$  is.
- Every model  $\mathfrak{M}$  of  $\forall x \xi_1$  can be augmented with an interpretation of  $F$  to yield a model  $\mathfrak{M}'$  of  $\forall x \xi_2$ . Similarly, for every model  $\mathfrak{M}'$  of  $\forall x \xi_2$ , restricting  $\mathfrak{M}'$  to the signature of  $\xi_1$  yields a model  $\mathfrak{M}$  of  $\forall x \xi_1$ .

The process of transforming  $\xi_1$  to  $\xi_2$  by eliminating  $\exists y$  and substituting  $F(x)$  for  $y$  is an instance of Skolemization. The fresh function symbol  $F$  introduced in the process is called a Skolem function. Skolem functions play a very important role in logic – both in theoretical investigations and in practical applications. The model theory of Skolemization in first order logic is rich: for instance, the Skolem expansion of a complete theory need no longer be complete, thus inviting further characterizations of Skolem hulls and indiscernibles. The extension of Skolemization to higher order logic is problematic and challenging (but needed, for instance, in automatic theorem proving).

While it suffices in some studies to simply know that a Skolem function  $F$  exists, in other cases (see Section 3 for such examples), we require an algorithm that effectively computes  $F(x)$  for every  $x$ . It turns out that obtaining such an algorithm is impossible in general, and even for the subcases where it is possible, the computational complexity is often very high. The purpose of this article is to discuss these computational challenges, and to survey some techniques for computing Skolem functions that have been proposed in recent years in the context of a significantly restricted yet practically useful logic, viz. quantified propositional logic.

Before delving further, it is important to formally define some notation and terminology. We use lower case English letters, viz.  $x, y, z$ , possibly with subscripts, to denote first order variables, and bold-faced upper case English letters, viz.  $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ , to denote sequences of first order variables. We use lower case Greek letters, viz.  $\varphi, \xi, \alpha$ , possibly with subscripts, to denote formulas. For a sequence  $\mathbf{X}$ , we use  $|\mathbf{X}|$  to denote the count of variables in  $\mathbf{X}$ , and  $x_1, \dots, x_{|\mathbf{X}|}$  to denote the individual variables in the sequence. With abuse of notation, we also use  $|\varphi|$  to denote the size of the formula  $\varphi$ , represented using a suitable format (viz. as a string, syntax tree, directed acyclic graph etc.), when there is no confusion. Let  $Q$  denote a quantifier in  $\{\exists, \forall\}$ . For notational convenience, we use  $Q\mathbf{X}$  to denote the block of quantifiers  $Qx_1 \dots Qx_{|\mathbf{X}|}$ . It is a standard exercise in logic to show that every well-formed first order logic formula can be transformed to a semantically equivalent *prenex normal form*, in which all quantifiers appear to the left of the quantifier-free part of the formula. Without loss of generality, let  $\xi(\mathbf{X}) \equiv \exists \mathbf{Y} \forall \mathbf{Z} \exists \mathbf{U} \dots \forall \mathbf{V} \exists \mathbf{W} \varphi(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{U}, \dots, \mathbf{V}, \mathbf{W})$  be such a formula in prenex normal form, where  $\mathbf{X}$  is a sequence of free variables and  $\varphi$  is a quantifier-free formula. In case the leading (resp. trailing) quantifier in  $\xi$  is universal, we consider  $\mathbf{Y}$  (resp.  $\mathbf{W}$ ) to be the empty sequence. Given such a formula  $\xi$ , *Skolemization* refers to the process of transforming  $\xi$  to a new (albeit related) formula  $\xi^*$  without any

81 existential quantifiers via the following steps: (i) for every existentially quantified  
 82 variable, say  $a$ , in  $\xi$ , substitute  $F_a(\mathbf{X}, \mathbf{S}_a)$  for  $a$  in the quantifier-free formula  $\varphi$ ,  
 83 where  $F_a$  is a new function symbol and  $\mathbf{S}_a$  is a sequence of universally quantified  
 84 variables that appear to the left of  $a$  in the quantifier prefix of  $\xi$ , and (ii) remove all  
 85 existential quantifiers from  $\xi$ . The functions  $F_a$  introduced above are called *Skolem*  
 86 *functions*. In case  $\xi$  has no free variables, i.e.  $\mathbf{X}$  is empty, the Skolem functions for  
 87 variables  $y_i$  in the leftmost existential quantifier block of  $\xi$  have no arguments  
 88 (i.e. are nullary functions), and are also called *Skolem constants*. The sentence  $\xi^*$  is  
 89 said to be in *Skolem normal form* if the quantifier-free part of  $\xi^*$  is in conjunctive  
 90 normal form. For notational convenience, let  $\exists\mathfrak{F}$  denote the second order quan-  
 91 tifier block  $\exists F_{y_1} \dots \exists F_{y_{|Y|}} \dots \exists F_{w_1} \dots \exists F_{w_{|W|}}$  that existentially quantifies over all  
 92 Skolem functions introduced above. The key guarantee of Skolemization is that  
 93 the second order sentence  $\exists\mathfrak{F} \forall \mathbf{X} (\xi \Leftrightarrow \xi^*)$  always holds. Note that substituting  
 94 Skolem functions for existentially quantified variables need not always make the  
 95 quantifier-free part of  $\xi$ , i.e.  $\varphi$ , evaluate to true. This can happen, for example, if  
 96 there are valuations of universally quantified variables for which no assignment of  
 97 existentially qualified variables renders  $\varphi$  true. For every other valuation of univer-  
 98 sally quantified variables, the Skolem functions indeed provide the “right” values  
 99 of existentially quantified variables so that  $\varphi$  evaluates to true.

100 *Example 1* Consider  $\xi \equiv \exists y \forall x \exists z \forall u \exists v \varphi(x, y, z, u, v)$ . On Skolemizing, we get  $\xi^* \equiv$   
 101  $\forall x \forall u \varphi(x, C_y, F_z(x), u, F_v(x, u))$ , where  $C_y$  is a Skolem constant for  $y$ , and  $F_z(x)$   
 102 and  $F_v(x, u)$  are Skolem functions for  $z$  and  $v$  respectively.

103 As mentioned earlier, the focus of this article is on effective computation of  
 104 Skolem functions. It is well known (see e.g. [31]) that there exist functions that  
 105 cannot be computed by any halting Turing machine, or equivalently, by any algo-  
 106 rithm. Therefore, it is interesting to ask: *Can every Skolem function be computed?*  
 107 In other words, given a first order formula  $\xi$ , does there always exist a halting  
 108 Turing machine that computes each Skolem function appearing in a Skolemized  
 109 version of  $\xi$ ? In general, such a Turing machine (or algorithm) may need to evalu-  
 110 ate predicate and function symbols that appear in the signature of  $\xi$  as part of its  
 111 computation. Therefore, the most appropriate notion of computation in our con-  
 112 text is that of *relative computation* or *computation by oracle machines*<sup>2</sup>. Formally,  
 113 let  $\mathcal{P}_\xi$  and  $\mathcal{F}_\xi$  denote the set of predicate and function symbols respectively in the  
 114 signature of  $\xi$ . Given oracles for interpretations of predicate symbols in  $\mathcal{P}_\xi$  and of  
 115 function symbols in  $\mathcal{F}_\xi$ , we ask if every Skolem function  $F$  in a Skolemized version  
 116 of  $\xi$  can be computed by a halting Turing machine, say  $M_\xi^F$ , with access to these  
 117 oracles. Note that we require  $M_\xi^F$  to depend only on  $\xi$  and  $F$ . However, the oracles  
 118 that  $M_\xi^F$  accesses can depend on specific interpretations of predicate and function  
 119 symbols.

120 Unfortunately, it has been shown in [1] that  $M_\xi^F$  does not always exist for  
 121 every  $\xi$  and  $F$ . In other words, Skolem functions cannot be effectively computed  
 122 in general, even in the relative sense mentioned above [1]. In fact, it doesn’t take  
 123 much to hit the uncomputability frontier. As shown in [1], uncomputability arises  
 124 even if we allow a single unary uninterpreted predicate in the signature. What  
 125 happens if all predicates and functions are interpreted, viz. in the theory of natural  
 126 numbers with multiplication and addition? It turns out that Skolem functions

<sup>2</sup> See [7] for a detailed exposition on relative computability.

cannot be computed in general in this case too [1]. The proof in this case [1] appeals to the Matiyasevich-Robinson-Davis-Putnam (MRDP) theorem [23] that equates Diophantine sets with recursively enumerable sets.

Not all hope is lost however. As shown in [1] again, Skolem functions can indeed be computed for formulas in several interesting first order theories. For example, every first order theory that is (i) decidable, (ii) has a recursively enumerable domain, and (iii) has computable interpretations of predicates and functions, admits effective computation of Skolem functions. Such theories include Presburger arithmetic, linear rational arithmetic, countable dense linear order without endpoints, theory of evaluated trees, first order theories with bounded domain etc. Whenever Skolem functions are computable, one can further ask: *Can Skolem functions be represented as terms in the underlying logical theory?* It is easy to see that a positive answer to this question implies an effective procedure for quantifier elimination. We also know that some theories, viz. Presburger logic without divisibility predicates, do not admit quantifier elimination. Therefore, there exist first order theories for which Skolem functions can be effectively computed, but are not expressible as terms in the underlying logical theory. The study of algorithmic computation of Skolem functions is therefore highly nuanced.

Given the above discussion, perhaps the simplest theories for which we can compute Skolem functions are those with bounded domains. Consider a formula  $\xi$  in such a theory where the domain  $\mathcal{D}$  has  $\kappa$  ( $\in \mathbb{N}$ ) elements. Since the elements of  $\mathcal{D}$  can be encoded as  $\lceil \log_2 \kappa \rceil$ -tuples of 0's and 1's, reasoning about  $\xi$  can be reduced to reasoning about a quantified propositional formula  $\hat{\xi}$ , where  $|\hat{\xi}| \leq \lceil \log_2 \kappa \rceil \cdot |\xi|$ . While this reduction does not affect the computational complexity results (in terms of complexity classes) that we study later, it can have an impact on the practical performance of algorithms, especially if  $\log_2 \kappa$  is large.

One may argue that over bounded domains, we can replace quantifiers by conjunctions or disjunctions and thus work only with propositional logic. This leads to an exponential blow-up in the size of the formula, which is undesirable. Hence, we are motivated to consider quantified propositional formulas directly. This is analogous to satisfiability for quantified Boolean formulas, which is a well-studied problem with dedicated techniques and implementations, even though it can be reduced to satisfiability for propositional formulas (with an exponential blow-up).

Despite the expressive limitations of *quantified propositional logic*, there are many important applications where quantified propositional formulas play an important role [57]. Furthermore, not only can we effectively compute Skolem functions for formulas in this logic, we can also represent them as Boolean functions. We therefore focus on the algorithmic computation of Skolem functions for quantified propositional logic in the remainder of the article.

## 2 Boolean Skolem functions, synthesis and unification

We use *Quantified Propositional Logic* (henceforth, QPL) to refer to propositional logic augmented with existential and universal quantifiers. Without loss of generality, we assume that formulas in quantified propositional logic (QPL) are given in prenex normal form. Prenex normal form sentences in this logic with the quantifier-

free part expressed in conjunctive normal form (CNF) are also called *quantified Boolean formulas* (QBF).

We introduce some additional notation for clarity of exposition. Given a propositional formula  $\varphi$ , its *support*, denoted  $\text{sup}(\varphi)$ , is the set of variables that appear in  $\varphi$ . As mentioned earlier, we use bold-faced upper case English letters to denote sequences of variables. To reduce notational clutter, we use the same letter to denote the set underlying a sequence as well, when there is no confusion. For example, we speak of a propositional formula  $\varphi(\mathbf{X})$  having support  $\mathbf{X}$ . If  $\mathbf{Y} = (y_1, \dots, y_r)$  is a sequence of variables appearing in  $\varphi$ , and if  $\Psi = (\psi_1, \dots, \psi_r)$  is a sequence of propositional formulas such that no formula  $\psi_i$  has any variable in  $\mathbf{Y}$  in its support, we use  $\varphi[\mathbf{Y} \mapsto \Psi]$  to denote the propositional formula obtained by substituting  $\psi_i$  for each  $y_i$  in  $\varphi$ . If  $\mathbf{Y} = (y)$  and  $\Psi = (\psi)$  are singleton sequences, we simply use  $\varphi[y \mapsto \psi]$  to denote the propositional formula resulting from substituting  $\psi$  for  $y$  in  $\varphi$ .

Let  $\xi(\mathbf{X}) \equiv \exists \mathbf{Y} \forall \mathbf{Z} \exists \mathbf{U} \dots \forall \mathbf{V} \exists \mathbf{W} \varphi(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{U}, \dots, \mathbf{V}, \mathbf{W})$  be a formula in QPL, where  $\varphi$  is a purely propositional formula. We wish to find Skolem functions for all existentially quantified variables in  $\xi$ . Since the domain of variables is  $\{\text{true}, \text{false}\}$ , each Skolem function is a mapping from  $\{\text{true}, \text{false}\}^k$  to  $\{\text{true}, \text{false}\}$ , for some  $k > 0$ . Such a Skolem function can also be viewed as defining the truth semantics of a propositional formula over  $k$  variables. We therefore represent every Skolem function, say  $F$ , in QPL by a propositional formula, say  $\psi^{(F)}$ , such that  $F$  gives the truth semantics of  $\psi^{(F)}$ . Although the distinction between  $F$  and  $\psi^{(F)}$  is significant (one is a function, the other is a formula), for notational convenience, we use the formula  $\psi^{(F)}$  to refer to the Skolem function  $F$ , when there is no confusion. When  $F$  is implicit from the context, we simply use  $\psi$  instead of  $\psi^{(F)}$ .

Although the quantifier prefix of the formula  $\xi$  mentioned above has multiple quantifier alternations, it suffices to know how to generate Skolem functions for QPL formulas with *only a single block* of existential quantifiers. To see why this is so, suppose  $\Psi_{\mathbf{W}}$  is a sequence of propositional formulas (representing Skolem functions), one for each variable  $w_i$  in  $\exists \mathbf{W} \varphi$ . By definition of Skolem functions, we have  $\exists \mathbf{W} \varphi \Leftrightarrow \varphi[\mathbf{W} \mapsto \Psi_{\mathbf{W}}]$ . Let  $\varphi'$  denote  $\exists \mathbf{W} \varphi$ . Since  $\forall \mathbf{V} \exists \mathbf{W} \varphi \Leftrightarrow \forall \mathbf{V} \varphi' \Leftrightarrow \neg \exists \mathbf{V} \neg \varphi'$ , if  $\Psi_{\mathbf{V}}$  represents a sequence of Skolem functions for  $\mathbf{V}$  in  $\exists \mathbf{V} \neg \varphi'$ , then  $\forall \mathbf{V} \exists \mathbf{W} \varphi \Leftrightarrow \neg(\neg \varphi'[\mathbf{V} \mapsto \Psi_{\mathbf{V}}]) \Leftrightarrow \varphi'[\mathbf{V} \mapsto \Psi_{\mathbf{V}}] \Leftrightarrow (\varphi[\mathbf{W} \mapsto \Psi_{\mathbf{W}}])[\mathbf{V} \mapsto \Psi_{\mathbf{V}}]$ . By repeating the above steps, it is possible to successively eliminate all quantifiers in  $\xi$ . This also yields a sequence of Skolem functions  $\Psi_{\mathbf{Y}}, \Psi_{\mathbf{U}}, \dots, \Psi_{\mathbf{W}}$  for the existentially quantified variables in  $\xi \equiv \exists \mathbf{Y} \forall \mathbf{Z} \exists \mathbf{U} \dots \forall \mathbf{V} \exists \mathbf{W} \varphi(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{U}, \dots, \mathbf{V}, \mathbf{W})$ . Note that the Skolem functions in  $\Psi_{\mathbf{Y}}$  (for variables in  $\mathbf{Y}$ ) have only the free variables  $\mathbf{X}$  as arguments. Similarly, the Skolem functions in  $\Psi_{\mathbf{U}}$  (for variables in  $\mathbf{U}$ ) have only the variables in  $\mathbf{X}, \mathbf{Y}$  and  $\mathbf{Z}$  as arguments. By substituting  $\Psi_{\mathbf{Y}}$  for  $\mathbf{Y}$  in  $\Psi_{\mathbf{U}}$ , we obtain Skolem functions for variables in  $\mathbf{U}$  in terms of only  $\mathbf{X}$  and  $\mathbf{Z}$ , i.e. universally quantified variables appearing to the left of  $\mathbf{U}$  in the quantifier prefix of  $\xi$ . It is easy to see that by repeating this process, we obtain Skolem functions for every existentially quantified variable in terms of (i) free variables  $\mathbf{X}$ , and (ii) universally quantified variables appearing to its left in the quantifier prefix of  $\xi$ .

In light of the above discussion, it makes sense to focus only on QPL formulas of the form  $\exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y})$  or  $\forall \mathbf{X} \exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y})$  for purposes of computing Skolem functions. Interestingly, with this restriction on the quantifier prefix, the problem of computing Skolem functions can also be viewed as one of *synthesis*. We elaborate more on this connection below.

## 2.1 The synthesis connection

Automatically and efficiently synthesizing correct systems from logical specifications is one of the holy grails of computer science. Suppose we wish to design a system with inputs  $\mathbf{X}$  and outputs  $\mathbf{Y}$ . To avoid notational confusion, we call  $\mathbf{X}$  as *system inputs*, and  $\mathbf{Y}$  as *system outputs* to distinguish them from inputs and outputs of Skolem functions/formulas. A *relational specification*  $\varphi(\mathbf{X}, \mathbf{Y})$  is a logical formula that implicitly relates desired values of system outputs with values of system inputs. Thus, every model of  $\varphi(\mathbf{X}, \mathbf{Y})$  gives values of  $\mathbf{X}$  and  $\mathbf{Y}$  that corresponds to a desired output in response to a specific input. Monadic second order logic, temporal logic and several variants of these logics [32] have been widely used to specify desirable system behaviour. In general, the specification  $\varphi(\mathbf{X}, \mathbf{Y})$  may permit multiple behaviours of the system outputs in response to a given input. A correct system design is required to produce any one of these allowed behaviours. It is also possible that for some values of the system inputs  $\mathbf{X}$ , there are no values of the system outputs  $\mathbf{Y}$  that render  $\varphi(\mathbf{X}, \mathbf{Y})$  true. In such cases, the specification cannot always be satisfied, no matter how we design the system. Such specifications are also called *unrealizable*. A correct synthesis procedure generates the system outputs  $\mathbf{Y}$  as a function  $\mathfrak{F}$  of the system inputs  $\mathbf{X}$ , such that  $\forall \mathbf{X} (\exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y}) \Leftrightarrow \varphi(\mathbf{X}, \mathfrak{F}(\mathbf{X})))$ . If a specification is *realizable*,  $\forall \mathbf{X} \exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y})$  is identically true; hence the requirement for synthesis simplifies to designing  $\mathfrak{F}(\mathbf{X})$  such that it renders  $\forall \mathbf{X} \varphi(\mathbf{X}, \mathfrak{F}(\mathbf{X}))$  identically true as well. Interestingly, even if a specification is unrealizable, it may be perfectly meaningful to synthesize  $\mathfrak{F}(\mathbf{X})$  such that  $\forall \mathbf{X} (\exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y}) \Leftrightarrow \varphi(\mathbf{X}, \mathfrak{F}(\mathbf{X})))$  holds. Indeed, as long as there is at least one way to generate system outputs in response to a given input such that the specification  $\varphi$  is satisfied, we want the system outputs generated by the synthesized system to satisfy the specification. In other cases, there are effectively no requirements on the system outputs.

Deciding realizability of a specification, and synthesizing a realizable specification are computationally hard problems in general. A relatively simpler cousin of the general synthesis problem, called *Boolean Functional Synthesis*, has recently received a lot of attention [36, 47, 38, 26, 53, 65, 2, 3, 54, 4, 52, 5, 29]. This problem is "simpler" in the sense that it concerns synthesis of Boolean functions, represented as Boolean circuits with AND, OR and NOT gates, from propositional logic specifications. Since every Boolean circuit corresponds to a propositional formula and vice versa, Boolean functional synthesis for  $\varphi(\mathbf{X}, \mathbf{Y})$  with system inputs  $\mathbf{X}$  and system outputs  $\mathbf{Y}$  can be seen to be equivalent to computing Skolem functions for the QPL formula  $\forall \mathbf{X} \exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y})$ . Therefore, we refer to the problem of computing Skolem functions for QPL and that of Boolean functional synthesis interchangeably. For notational convenience, we use *Boolean Skolem function synthesis*, or *BoolSkFnSyn* for short, to refer to either problem in the remainder of this article. It is worth emphasizing here that Boolean Skolem function synthesis is distinct from the problem of *combinational logic synthesis and optimization* [24]. In the former, we start from a relational specification that doesn't necessarily give the system outputs explicitly as functions of system inputs, and our primary task is to synthesize these outputs as Boolean functions of system inputs. In contrast, in combinational logic synthesis and optimization, we are given system outputs as explicit Boolean functions of system inputs, and our goal is to implement these

268 functions optimally as Boolean circuits with specified gate types (viz. NAND,  
269 NOR, XOR, etc.).

270 In the context of QPL, the specification  $\varphi(\mathbf{X}, \mathbf{Y})$  and the Skolem functions for  
271  $\mathbf{Y}$  can be represented in several ways. Some commonly used representations in-  
272 clude lists of clauses for propositional formulas in conjunctive normal form (CNF),  
273 Boolean circuits, reduced ordered binary decision diagrams (ROBDDs) [62], and  
274 inverter graphs (AIGs) [40], decision lists, decision trees etc. The choice of repre-  
275 sentation has a bearing on the computational complexity of BoolSkFnSyn; hence it  
276 is important to spell out the representation clearly when discussing a solution to  
277 the problem. Interestingly, all the representations mentioned above can be trans-  
278 lated to Boolean circuits with AND, OR and NOT gates with at most a linear  
279 blow-up. Hence, we consider Boolean circuits with AND, OR and NOT gates as  
280 a unifying representation for both relational specifications and for Skolem func-  
281 tions. Computational hardness (lower bound) results based on Boolean circuit  
282 representations naturally hold when the other representations are used as well. A  
283 particularly convenient form of Boolean circuits are those in which every NOT  
284 gate is immediately fed by a circuit input (labeled by a variable). Such circuits are  
285 also called *Negation Normal Form (or NNF)* circuits. For notational convenience,  
286 we treat every NOT gate fed by a circuit input labeled  $v$  in a NNF circuit as a new  
287 circuit input labeled  $\neg v$ . Thus, an NNF circuit can be viewed as one containing  
288 only AND and OR gates, with the circuit inputs labeled by *literals* over the set of  
289 variables, i.e. variables and their negations. It is easy to see that every Boolean  
290 circuit can be compiled to a NNF circuit that computes the same function as the  
291 original circuit, and is at most twice its size.

## 292 2.2 The unification connection

293 The BoolSkFnSyn problem is related to that of *Boolean unification* – a classical prob-  
294 lem studied by George Boole [13] and Leopold Löwenheim [44] much before Alan  
295 Turing and Alonzo Church formalized the notion of computation. The interested  
296 reader is referred to an excellent (albeit, dated) survey by Martin and Nipkow [48]  
297 for details about the Boolean unification problem. For our purposes, Boolean unifi-  
298 cation may be viewed as asking the following question: *Given two Boolean functions*  
299  $F, G : \{\text{true}, \text{false}\}^n \rightarrow \{\text{true}, \text{false}\}$ , *find a map*  $\mathfrak{F} : \{\text{true}, \text{false}\}^m \rightarrow \{\text{true}, \text{false}\}^n$ , *where*  
300  $m \geq 0$  *such that*  $F(\mathfrak{F}(\sigma)) = G(\mathfrak{F}(\sigma))$  *for all*  $\sigma \in \{\text{true}, \text{false}\}^m$ , *or report that no such*  
301 *map exists.* The map  $\mathfrak{F}$ , if it exists, is called a *unifier* of  $F$  and  $G$ . In general, there  
302 can be zero, one or multiple unifiers of  $F$  and  $G$ . A unifier  $\mathfrak{F} : \{\text{true}, \text{false}\}^m \rightarrow$   
303  $\{\text{true}, \text{false}\}^n$  is said to be *more general* than unifier  $\mathfrak{G} : \{\text{true}, \text{false}\}^\ell \rightarrow \{\text{true}, \text{false}\}^n$   
304 if there exists a map  $\mathfrak{H} : \{\text{true}, \text{false}\}^\ell \rightarrow \{\text{true}, \text{false}\}^m$  such that  $\mathfrak{F}(\mathfrak{H}(\hat{\sigma})) = \mathfrak{G}(\hat{\sigma})$   
305 for all  $\hat{\sigma} \in \{\text{true}, \text{false}\}^\ell$ . A *most general unifier* of  $F$  and  $G$  is a unifier that is more  
306 general than all unifiers of  $F$  and  $G$ . By a result due to Boole [13], we know that  
307 if two Boolean functions  $F$  and  $G$  are unifiable, there exists a most general unifier  
308 of  $F$  and  $G$ .

309 To see the connection of Boolean unification with BoolSkFnSyn, let  $\varphi(\mathbf{X}, \mathbf{Y})$  be  
310 a propositional relational specification such that (i)  $|\mathbf{X}| + |\mathbf{Y}| = n$ , and (ii) the  
311 truth semantics of  $\varphi$  is given by  $F(\mathbf{X}, \mathbf{Y})$ . Let  $G(\mathbf{X}, \mathbf{Y})$  denote the truth semantics  
312 of  $\exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y})$ , viewed as a function of  $\mathbf{X}$  and (redundantly) of  $\mathbf{Y}$ . A solution to  
313 the BoolSkFnSyn problem for  $\varphi$  yields a vector  $\Psi$  of propositional formulas (repre-

314 sending Skolem functions), one for each variable in  $\mathbf{Y}$ , that can be converted to a  
 315 unifier of  $F$  and  $G$  as follows. Note that  $\Psi$  represents a mapping from  $\{\text{true}, \text{false}\}^{|\mathbf{X}|}$   
 316 to  $\{\text{true}, \text{false}\}^{|\mathbf{Y}|}$ . Let  $\text{Id}_{|\mathbf{X}|}$  be the identity mapping on  $\{\text{true}, \text{false}\}^{|\mathbf{X}|}$ . The conca-  
 317 tentation of  $\text{Id}_{|\mathbf{X}|}$  and  $\Psi$ , denoted  $(\text{Id}_{|\mathbf{X}|}, \Psi)$ , gives a vector of functions mapping  
 318  $\{\text{true}, \text{false}\}^{|\mathbf{X}|}$  to  $\{\text{true}, \text{false}\}^{|\mathbf{X}|+|\mathbf{Y}|}$ , such that  $F(\text{Id}_{|\mathbf{X}|}(\sigma), \Psi(\sigma)) = G(\text{Id}_{|\mathbf{X}|}(\sigma), \Psi(\sigma))$   
 319 for all  $\sigma \in \{\text{true}, \text{false}\}^{|\mathbf{X}|}$ . The above discussion shows that given  $\varphi(\mathbf{X}, \mathbf{Y})$ , if  $F$  and  
 320  $G$  are chosen appropriately, then specific *unifiers* for  $F$  and  $G$  correspond to Skolem  
 321 functions for  $\mathbf{Y}$  in  $\varphi(\mathbf{X}, \mathbf{Y})$ . Indeed, if the unifier is a most general unifier, then  
 322 the Skolem functions turn out to be specific instantiations of this most general  
 323 unifier. Interestingly, algorithms for finding the most general unifier in Boolean  
 324 unification were given by both Boole [13] and Lowenheim [44] in their early work.  
 325 These and other variant algorithms for finding most general unifiers in Boolean  
 326 unification were experimentally evaluated in [45]. Applications of Boolean unifica-  
 327 tion have also been reported in [12, 16, 59, 46]. Unfortunately, solving `BoolSkFnSyn`  
 328 using the Boolean unification approach turns out to be too inefficient for use in  
 329 practical applications with thousands of variables and beyond.

### 330 3 Applications of Boolean Skolem function synthesis

331 Before delving deeper into the computational aspects of `BoolSkFnSyn`, let us look at  
 332 a few interesting applications of the problem. These applications provide strong  
 333 motivation for developing algorithms for `BoolSkFnSyn` that work well in practice,  
 334 despite non-trivial worst-case complexity-theoretic lower bounds.

335 We start with a particularly challenging application that illustrates why an  
 336 efficient algorithmic solution of `BoolSkFnSyn` can have far-reaching implications in  
 337 practice. Consider a system with a single  $2n$ -bit unsigned integer input  $\mathbf{X}$ , and two  
 338  $n$ -bit unsigned integer outputs  $\mathbf{Y}_1$  and  $\mathbf{Y}_2$ . Suppose the relational specification is  
 339 given as  $F_{\text{fact}}(\mathbf{X}, \mathbf{Y}_1, \mathbf{Y}_2) \equiv ((\mathbf{X} = \mathbf{Y}_1 \times_{[n]} \mathbf{Y}_2) \wedge (\mathbf{Y}_1 \neq 1) \wedge (\mathbf{Y}_2 \neq 1))$ , where  $\times_{[n]}$   
 340 denotes  $n$ -bit unsigned integer multiplication. This specification requires that  $\mathbf{Y}_1$   
 341 and  $\mathbf{Y}_2$  are non-trivial factors of  $\mathbf{X}$ . Note, however, that if  $\mathbf{X}$  represents a prime  
 342 number, there are no values of  $\mathbf{Y}_1$  and  $\mathbf{Y}_2$  that satisfy the specification. Technically,  
 343 the specification is unrealizable. Nevertheless, we are interested in obtaining values  
 344 of  $\mathbf{Y}_1$  and  $\mathbf{Y}_2$  that satisfy the specification, whenever possible. Significantly, the  
 345 above specification can be encoded as a Boolean formula of size  $\mathcal{O}(n^2)$  over the  
 346 individual bits of  $\mathbf{X}$ ,  $\mathbf{Y}_1$  and  $\mathbf{Y}_2$ . However, if we want to express  $\mathbf{Y}_1$  and  $\mathbf{Y}_2$  directly  
 347 as Boolean functions of  $\mathbf{X}$ , our task turns out to be significantly harder. In fact,  
 348 there are no known polynomial-sized Boolean functions (represented as circuits of  
 349 AND, OR and NOT gates) that can express individual bits of  $\mathbf{Y}_1$  and  $\mathbf{Y}_2$  directly in  
 350 terms of the individual bits of  $\mathbf{X}$ . Otherwise, we could efficiently factorize products  
 351 of  $n$ -bit prime numbers, rendering cryptographic systems vulnerable to attacks.  
 352 This application also illustrates how relational specifications can be more natural  
 353 and succinct than expressing outputs directly as functions of inputs.

354 As another application, we consider satisfiability checking of quantified boolean  
 355 sentences (also called QBF-SAT), which is increasingly being used in diverse ap-  
 356 plications such as planning, model checking, non-monotonic reasoning, reactive  
 357 synthesis, games, equivalence checking, circuit repair, program synthesis etc. An  
 358 excellent survey of such applications can be found in [57]. Given the sophistication  
 359 of modern QBF-SAT solvers, it is hard to rule out bugs in solver implementations. It



is therefore desirable that when a QBF-SAT solver is invoked, it not only produces a “Yes”/”No” answer to the decision problem, but also a certificate that can be independently (machine-)checked to validate the correctness of the answer. Multiple notions of certificates have been used in the literature [57, 8, 50], including the use of Skolem functions for existentially quantified variables in valid QBFs, and the use of Herbrand functions<sup>3</sup> for universally quantified variables in unsatisfiable QBFs. In addition to their use as certificates of QBF-SAT results, Skolem function based certificates also have independent value as they can be used for other objectives, such as, to extract a feasible plan in a robotic planning problem, a replacement sub-circuit in a circuit repair problem, a program fragment in automated program synthesis, a winning strategy in a game. As discussed earlier, knowing how to synthesize Skolem functions for QBF formulas of the form  $\forall \mathbf{X} \exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y})$  suffices to generate Skolem functions (resp. Herbrand functions) for all existentially (resp. universally) quantified variables in a QBF. This underscores the importance of the BoolSkFnSyn problem.

Talking of synthesis, recall that BoolSkFnSyn can be viewed as a simpler version of the more general reactive synthesis problem (see [25] for a survey). It turns out that several algorithmic approaches to reactive synthesis use BoolSkFnSyn as a key step (see e.g [14, 35]). Hence, a practically efficient algorithmic solution to BoolSkFnSyn benefits reactive synthesis as well.

#### 4 Boolean Skolem function synthesis through lens of computation

Recall the definition of BoolSkFnSyn from Section 2. We are given a propositional formula  $\varphi(\mathbf{X}, \mathbf{Y})$ , specifying a relation between system inputs  $\mathbf{X}$  and system outputs  $\mathbf{Y}$ . For notational convenience, we use  $m$  to denote  $|\mathbf{X}|$  and  $n$  to denote  $|\mathbf{Y}|$ . The BoolSkFnSyn problem requires us to find a vector of propositional formulas (representing Boolean functions)  $\Psi(\mathbf{X}) = (\psi_1(\mathbf{X}), \dots, \psi_n(\mathbf{X}))$  such that  $\forall \mathbf{X} (\exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y}) \Leftrightarrow \varphi(\mathbf{X}, \Psi(\mathbf{X})))$  is true. The formula  $\psi_i(\mathbf{X})$  represents a Skolem function for  $y_i$  in  $\varphi$ , and  $\Psi(\mathbf{X})$  is called a *Skolem function vector* for  $\mathbf{Y}$  in  $\varphi$ . As discussed earlier, we represent all Skolem functions and propositional formulas by Boolean circuits comprised of AND, OR and NOT gates.

*Example 2* Consider the relational specification  $\varphi(\mathbf{X}, \mathbf{Y}) \equiv (x_1 \vee y_2) \wedge (\neg x_2 \vee \neg x_1 \vee y_1)$ . A few (among many possible) Skolem function vectors for  $\mathbf{Y}$  in  $\varphi$  are  $(\text{true}, \text{true})$ ,  $(\text{true}, \neg x_1)$ ,  $(x_1, \neg x_1)$ ,  $(x_2, \neg x_1)$ , where each tuple represents  $(\psi_1(\mathbf{X}), \psi_2(\mathbf{X}))$ .

While a given problem instance may admit multiple Skolem function vectors, a solution to BoolSkFnSyn seeks only one such vector. Thus, there may not be a unique solution to an instance of BoolSkFnSyn.

It is not hard to see that BoolSkFnSyn can be solved in time (and space) exponential in  $|\varphi|$  in the worst-case, simply by brute-force enumeration of all possible values of  $\mathbf{X}$  and  $\mathbf{Y}$ . However, does the problem admit more efficient solutions? If  $|\mathbf{Y}| = n = 1$ , it turns out that there is a surprisingly efficient solution. To understand this, we need some additional notation. Let  $\alpha$  be a propositional formula and  $v \in \text{sup}(\alpha)$ . We use  $\alpha|_v$  (resp.  $\alpha|_{\neg v}$ ) to denote the positive (resp. negative)

<sup>3</sup> A Herbrand function for universally quantified variables in a quantified propositional sentence  $\varphi$  may be thought of as Skolem functions for existentially quantified variables in  $\neg\varphi$ .

402 *co-factor* of  $\alpha$  with respect to  $v$ , i.e.  $\alpha$  with  $v$  set to **true** (resp. **false**). It can now be  
 403 verified that if  $\varphi(\mathbf{X}, y)$  is a specification with a single system output  $y$ , then both  
 404  $\varphi|_y$  and  $\neg(\varphi|_{\neg y})$  serve as Skolem functions for  $y$  in  $\varphi$ . This technique for obtaining  
 405 a Skolem function for a single system output is also called *self-substitution*, and has  
 406 been used in several prior works [66, 36, 26, 38, 2, 29]. In fact, if  $\beta(\mathbf{X})$  denotes  $\varphi|_y$   
 407 and  $\gamma(\mathbf{X})$  denotes  $\varphi|_{\neg y}$ , then the entire set of Skolem functions for  $y$  in  $\varphi$  can be  
 408 parametrically represented as  $(\neg\gamma(\mathbf{X}) \wedge \beta(\mathbf{X})) \vee ((\beta(\mathbf{X}) \Leftrightarrow \gamma(\mathbf{X})) \wedge \delta(\mathbf{X}))$ , where  
 409  $\delta(\mathbf{X})$  is *any* Boolean function on  $\mathbf{X}$  [66, 36].

410 An obvious question to ask at this point is whether the simple solution for  
 411  $|\mathbf{Y}| = 1$  can be extended to the case where  $|\mathbf{Y}| > 1$ . Unfortunately, this turns  
 412 out to be more difficult, and there are complexity-theoretic barriers along the  
 413 way. Nevertheless, the underlying idea for the  $|\mathbf{Y}| = 1$  case can be generalized  
 414 to obtain some insights. Towards this end, let  $y_1 \prec y_2 \cdots \prec y_n$  be a (arbitrary)  
 415 linear ordering of the system outputs, and let  $\mathbf{Y}_i^j$  denote the subsequence  $(y_i, \dots, y_j)$   
 416 of  $\mathbf{Y}$ , for  $1 \leq i \leq j \leq n$ . Furthermore, let  $\varphi^{(i-1)}(\mathbf{X}, \mathbf{Y}_i^n)$  denote  $\exists \mathbf{Y}_1^{i-1} \varphi(\mathbf{X}, \mathbf{Y})$ ,  
 417 where  $\varphi^{(0)}$  is defined to be  $\varphi$ . For every  $i$  in 1 to  $n$  in that order, suppose we view  
 418 the formula  $\varphi^{(i-1)}(\mathbf{X}, y_i, \mathbf{Y}_{i+1}^n)$  as a specification with system inputs  $\mathbf{X} \cup \mathbf{Y}_{i+1}^n$   
 419 and a single system output  $y_i$ . We can now apply the reasoning for synthesizing  
 420 a single Skolem function, as discussed above, to obtain a Skolem function for  
 421  $y_i$  in terms of  $\mathbf{X} \cup \mathbf{Y}_{i+1}^n$ . Let  $\psi_i(\mathbf{X}, \mathbf{Y}_{i+1}^n)$  be such a Skolem function for  $y_i$ , i.e.  
 422  $\varphi^{(i-1)}(\mathbf{X}, \psi_i, \mathbf{Y}_{i+1}^n) \Leftrightarrow \exists y_i \varphi^{(i-1)}(\mathbf{X}, y_i, \mathbf{Y}_{i+1}^n)$ . Once we have computed  $\psi_i$  for  $i \in$   
 423  $\{1, \dots, n\}$  in this manner, we can substitute  $\psi_{i+1}$  through  $\psi_n$  for  $y_{i+1}$  through  $y_n$   
 424 respectively, in the definition of  $\psi_i$  to obtain a Skolem function for  $y_i$  as a function  
 425 of only  $\mathbf{X}$ . This approach is widely used in the BoolSkFnSyn literature [36, 37, 38, 26,  
 426 2, 4, 29], and we follow it for the rest of our discussion. Note that this allows us to  
 427 focus on synthesizing  $\psi_i$  in terms of  $\mathbf{X}$  and  $\mathbf{Y}_{i+1}^n$ , instead of synthesizing it directly  
 428 in terms of  $\mathbf{X}$ . Generalizing the idea of the solution when we have a single system  
 429 output, it can be shown that both  $\neg(\varphi^{(i-1)}|_{\neg y_i})$  and  $\varphi^{(i-1)}|_{y_i}$  serve as Skolem  
 430 functions for  $y_i$  (in terms of  $\mathbf{X}$  and  $\mathbf{Y}_{i+1}^n$ ). Furthermore, if  $\beta_i(\mathbf{X}, \mathbf{Y}_{i+1}^n)$  denotes  
 431  $\varphi^{(i-1)}|_{y_i}$  and  $\gamma_i(\mathbf{X}, \mathbf{Y}_{i+1}^n)$  denotes  $\varphi^{(i-1)}|_{\neg y_i}$ , then every Skolem function for  $y_i$  can  
 432 be parametrically represented as  $(\neg\gamma_i(\mathbf{X}, \mathbf{Y}_{i+1}^n) \wedge \beta_i(\mathbf{X}, \mathbf{Y}_{i+1}^n)) \vee ((\beta_i(\mathbf{X}, \mathbf{Y}_{i+1}^n) \Leftrightarrow$   
 433  $\gamma_i(\mathbf{X}, \mathbf{Y}_{i+1}^n)) \wedge \delta_i(\mathbf{X}, \mathbf{Y}_{i+1}^n))$ , where  $\delta_i(\mathbf{X}, \mathbf{Y}_{i+1}^n)$  is *any* Boolean function on  $\mathbf{X}$  and  
 434  $\mathbf{Y}_{i+1}^n$ .

435 While the above discussion may seem to imply that there is an easy way to  
 436 solve BoolSkFnSyn in general, the difficulty in the above approach lies in com-  
 437 puting a good linear ordering of  $y_i$ 's and also in computing  $\varphi^{(i-1)}$  for  $1 \leq i \leq n$ .  
 438 Experiments, e.g. from [38, 2, 29], show that using different linear orderings affects  
 439 the time taken for synthesizing functions considerably. For values of  $|\mathbf{X}| = m$  and  
 440  $|\mathbf{Y}| = n$  running into thousands, these issues can pose enormous scalability chal-  
 441 lenges in practice. However, the computational hurdles are not restricted to only  
 442 the approach discussed above. It turns out that any other algorithmic technique to  
 443 solve BoolSkFnSyn must also encounter scalability hurdles in the worst-case. Com-  
 444 putational complexity theory provides the tools necessary to reason about these  
 445 challenges, by allowing us to derive lower bounds on computational resources (viz.  
 446 space and time) needed to solve BoolSkFnSyn in general. We elaborate on this in  
 447 the next couple of sections.

## 448 4.1 A quick primer on the polynomial hierarchy and related complexity classes

449 In computational complexity theory, a *decision problem* is one that has a “Yes”/”No”  
 450 answer. An example of such a problem is: *Given a propositional formula  $\varphi$ , is  $\varphi$  sat-*  
 451 *isfiable?* A *function problem* generalizes a decision problem by allowing the answer  
 452 to be more general than “Yes”/”No”. For example, we could ask: *Given a proposi-*  
 453 *tional formula  $\varphi$  in conjunctive normal form, what is the maximum number of clauses*  
 454 *of  $\varphi$  that can be simultaneously satisfied?* For a large class of function problems,  
 455 an efficient solution to an appropriately defined decision version of the problem  
 456 implies an efficient solution to the function problem itself. Studying the complex-  
 457 ity of decision problems has therefore been a major focus of complexity theoretic  
 458 investigations. A decision problem can also be viewed as a *language recognition*  
 459 problem, where the input is presented as a finite string over the alphabet  $\{0, 1\}$ ,  
 460 and the set of all input strings that yield a “Yes” answer comprises the language  
 461  $L$  corresponding to the problem. Thus, given an input string  $\text{str}$  representing an  
 462 instance of the problem, the decision problem effectively asks if  $\text{str} \in L$ . This is  
 463 equivalent to asking if the problem instance has a “Yes” answer.

464 The complexity class  $P$  (resp.  $NP$ ) consists of the set of all languages accepted  
 465 by deterministic (resp. non-deterministic) Turing machines in time that grows at  
 466 most polynomially in the size of the input. The class  $coNP$  is the set of all languages,  
 467 the complement of which are in  $NP$ . The polynomial hierarchy generalizes these  
 468 classes by defining two inter-related sub-hierarchies – the  $\Sigma^P$ -hierarchy and the  
 469  $\Pi^P$ -hierarchy. We start by defining  $\Sigma_0^P = \Pi_0^P = P$ . For every  $n \in \mathbb{N} \setminus \{0\}$ , we then  
 470 define  $\Sigma_n^P$  and  $\Pi_n^P$  inductively as follows, where  $\{0, 1\}^*$  denotes the set of all finite  
 471 strings over  $\{0, 1\}$ , and  $|\text{str}|$  denotes the length of the string  $\text{str}$ .

- $\Sigma_n^P$  consists of all languages/problems  $L$  such that there exists a language  
 $L' \in \Pi_{n-1}^P$  and a polynomial  $q$  such that

$$\forall x \in \{0, 1\}^* x \in L \Leftrightarrow \exists y \in \{0, 1\}^*, |y| \leq q(|x|) \text{ and } (x, y) \in L'.$$

- $\Pi_n^P$  consists of all languages/problems  $L$  such that there exists a language  
 $L' \in \Sigma_{n-1}^P$  and a polynomial  $q$  such that

$$\forall x \in \{0, 1\}^* x \in L \Leftrightarrow \forall y \in \{0, 1\}^*, |y| \leq q(|x|) \Rightarrow (x, y) \in L'.$$

472 It is easy to see from the definitions that  $NP = \Sigma_1^P$  and  $coNP = \Pi_1^P$ . The hier-  
 473 archy of complexity classes defined above is known as the *Polynomial Hierarchy*  
 474 (henceforth, PH). The PH is said to *collapse to level  $i \in \mathbb{N}$*  if  $\Sigma_i^P = \Sigma_{i+1}^P$ . Notice  
 475 that if PH collapses to level 0, then  $P = NP$ . It is widely believed that PH is a  
 476 strict infinite hierarchy and does not collapse to any finite level. However, this is  
 477 only a conjecture; the question of whether PH indeed collapses to any finite level  
 478 has remained open for decades, and is one of the outstanding open problems in  
 479 computational complexity theory.

480 The classes in PH are also related to the notion of *oracle computation* or *relative*  
 481 *computation*, referred to in Section 1. Recall that an oracle machine is a Turing ma-  
 482 chine with access to a “black-box” (oracle) that can provide “Yes”/”No” answers  
 483 to a specific class of decision problem in a single step. If oracles are restricted to be  
 484 Turing machines themselves with well-defined resource constraints, we obtain an  
 485 alternative characterization of the complexity classes in PH. The interested reader

is referred to [7] for details. For our purposes, it suffices to note that  $P^{NP}$  is one such complexity class obtained by considering polynomial-time Turing machines with access to an NP oracle. That is, any problem in this class can be solved by a deterministic Turing machine in polynomially many steps, if it is allowed to make at most polynomially many calls to an NP oracle. In fact, the complexity class  $P^{NP}$  can be shown to coincide with  $\Sigma_2^P \cap \Pi_2^P$ , and hence is within the second level of the polynomial hierarchy!

Just as P is the class of languages accepted by deterministic Turing machines running for at most polynomial time, PSPACE denotes the class of languages accepted by deterministic Turing machines that use at most polynomial space. It is known that non-determinism does not add power in this case, i.e.,  $NPSPACE = PSPACE$ . Also it is known that  $PH \subseteq PSPACE$ , i.e., the entire polynomial hierarchy is contained in the class PSPACE, thereby making this a very expressive class. Notice, however, that if a Turing machine can run for exponential time, then it can indeed simulate a Turing machine that is allowed to use only polynomial space. The class of languages accepted by deterministic Turing machines running for exponential time is denoted EXP, and we immediately see that  $PSPACE \subseteq EXP$ . We refer the interested reader to excellent textbooks, e.g., [7], in this area for more information about complexity classes and their relations.

#### 4.2 Computational hardness of Boolean Skolem Function Synthesis

With the above notations, we can now present complexity-theoretic hardness results for BoolSkFnSyn. As mentioned earlier, we assume the input and output of BoolSkFnSyn are represented as Boolean circuits. It turns out that three conditional results can be shown, two of which are related to the collapse of the polynomial hierarchy defined above.

The first result is about time-complexity. Specifically, any algorithm that solves BoolSkFnSyn must take *super-polynomial* (i.e., asymptotic growth greater than that of any polynomial) time in the worst case, unless the polynomial hierarchy collapses to the first level (i.e.,  $P = NP$ ). Since the question of whether  $P = NP$  has remained open for decades, with the general wisdom being  $P \neq NP$ , it is highly unlikely that all instances of BoolSkFnSyn can be solved in polynomial time. This easily follows from the observation that propositional satisfiability can be reduced to BoolSkFnSyn where we have no system inputs  $\mathbf{X}$ .

Next, we inquire about the space complexity of BoolSkFnSyn, and ask if it is possible to solve BoolSkFnSyn *compactly*. More precisely, do there always exist polynomial-sized Skolem functions for instances of BoolSkFnSyn, even if it takes exponential time to synthesize them? Again, the answer turns out to be negative, but with a stronger condition. It is shown in [3, 5] that unless the polynomial hierarchy collapses to the second level, there must exist instances of BoolSkFnSyn for which any algorithm must generate super-polynomial sized Skolem functions.

The above results provide conditional super-polynomial time and space lower bounds for BoolSkFnSyn. On the other hand, a trivial upper bound was mentioned earlier, namely, BoolSkFnSyn can be solved in exponential time and space. A naive exponential time algorithm would be to enumerate all possible values of system inputs  $\mathbf{X}$ , and for each such valuation, check by enumeration again if there exists a valuation of the system outputs  $\mathbf{Y}$  that satisfies the given specification. Since we

532 are concerned about Boolean specifications, this can be done in time exponential  
 533 in  $|\mathbf{X}|$  and  $|\mathbf{Y}|$ ; of course, in doing so, it may produce Skolem functions of at most  
 534 exponential size.

535 Given the large gap between a polynomial lower bound and an exponential  
 536 upper bound, a natural question is whether this gap can be narrowed or bridged.  
 537 In [3, 5], it is shown that under a stronger hypothesis, this gap can in fact be  
 538 completely eliminated giving us optimal and tight (albeit conditional) complex-  
 539 ity bounds. To understand this result, let us start by considering two unproven  
 540 complexity-theoretic conjectures. The exponential-time hypothesis ETH [34] and  
 541 its non-uniform variant,  $\text{ETH}_{\text{nu}}$  [18], are unproven computational hardness con-  
 542 jectures that have been used to show that several classical decision, functional  
 543 and parametrized NP-complete problems are unlikely to have sub-exponential al-  
 544 gorithms. These conjectures are also widely believed to be true. Formally,  $\text{ETH}_{\text{nu}}$  –  
 545 the variant that we need – states that there is no family of algorithms (one for each  
 546 input-size  $n$ ) that can solve the  $n$ -variable instance of the propositional satisfia-  
 547 bility problem (the canonical NP-complete problem) in sub-exponential time (i.e.,  
 548 in time that is lower than any exponential function of  $n$ , also written  $2^{o(n)}$ ). By  
 549 adapting the earlier result, one can now show that, unless the non-uniform expo-  
 550 nential time hypothesis  $\text{ETH}_{\text{nu}}$  fails, there exist instances of  $\text{BoolSkFnSyn}$  for which  
 551 any algorithm for must generate exponential-sized Skolem functions. Notice that  
 552 this immediately implies exponential time complexity as well, since generating an  
 553 output of size  $f(n)$  requires at least  $f(n)$  time.

554 Summarizing, we obtain the following theorem, whose details and proof can be  
 555 found in [3, 5].

- 556 **Theorem 1** 1. *BoolSkFnSyn can be solved in exponential time and space.*  
 557 2. *There exists no algorithm for BoolSkFnSyn that*  
 558 (a) *always takes polynomial time on all inputs, unless PH collapses to level 0.*  
 559 (b) *always generates polynomial sized Skolem functions, unless PH collapses to the*  
 560 *second level.*  
 561 (c) *always generates sub-exponential sized Skolem functions (and takes sub-exponential*  
 562 *time), unless the non-uniform exponential-time hypothesis fails.*

563 Together these results imply that  $\text{BoolSkFnSyn}$  is unlikely to have polynomial-  
 564 time or polynomial-space algorithms in general. Any such efficient algorithm must  
 565 necessarily falsify one of the above well-regarded and intensely researched conjec-  
 566 tures in complexity theory.

### 567 4.3 Exploiting the structure of the specification

568 Given a Boolean relational specification as a circuit, we now ask if there are con-  
 569 ditions on the structure/representation of the circuit that can be exploited to  
 570 efficiently synthesize Skolem functions. Indeed, this turns out to be the case, and  
 571 we discuss some such cases below.

#### 572 4.3.1 Unate variables

573 Recall that  $\text{BoolSkFnSyn}$  requires us to synthesize the entire Skolem function vector,  
 574 i.e., Skolem functions for all system outputs in  $\mathbf{Y}$ . However, synthesizing Skolem

575 functions for some system output variables may be easier than that for others. For  
 576 example, consider the case of *unate* variables. The formula  $\varphi$  is said to be *positive*  
 577 *unate* in  $v \in \text{sup}(\varphi)$  iff  $\varphi|_{\neg v} \Rightarrow \varphi|_v$ . Similarly,  $\varphi$  is said to be *negative unate* in  $v$  iff  
 578  $\varphi|_v \Rightarrow \varphi|_{\neg v}$ . Finally,  $\varphi$  is *unate* in  $v$  if it is either positive unate or negative unate in  
 579  $v$ . If  $\varphi$  is positive unate in  $v$ , it immediately follows that  $\exists v \varphi \Leftrightarrow (\varphi|_v \vee \varphi|_{\neg v}) \Leftrightarrow \varphi|_v$ .  
 580 As a result, if  $v$  is a system output, the constant function `true` serves as a correct  
 581 Skolem function for  $v$  in  $\varphi$ . Similarly `false` serves a correct Skolem function for  $v$   
 582 in  $\varphi$  if  $\varphi$  is negative unate in  $\varphi$ . Thus, we obtain,

583 **Proposition 1** *If a specification  $\varphi(\mathbf{X}, \mathbf{Y})$  is unate in  $y_i \in \mathbf{Y}$ , one can generate*  
 584 *constant-sized Skolem functions for  $y_i$  in  $\varphi$  in constant time.*

585 Substituting a constant Skolem function for  $y_i \in \mathbf{Y}$  in the specification  $\varphi(\mathbf{X}, \mathbf{Y})$   
 586 and simplifying it may, in turn, reveal that the simplified specification is unate  
 587 in  $y_j$  (distinct from  $y_i$ ), even if the original specification was not unate in  $y_j$ . It  
 588 is therefore beneficial to iterate through this process of detecting if a specification  
 589 is unate in a system output variable and substituting a constant Skolem function  
 590 for the variable to simplify the specification.

591

592 *Example 3* Consider the specification  $\varphi \equiv (\neg x_1 \vee y_1) \wedge (x_1 \vee \neg x_2 \vee y_1 \vee \neg y_2) \wedge (\neg x_1 \vee$   
 593  $\neg x_2 \vee y_2 \vee y_3) \wedge (x_2 \vee \neg y_3 \vee y_2)$ . Applying the checks for positive and negative  
 594 unateness described above, it is easy to verify that  $\varphi$  is only positive unate in  
 595  $y_1$ , and neither positive nor negative unate in  $y_2$  or  $y_3$ . If we now set the Skolem  
 596 function for  $y_1$  to the constant `true`, the specification simplifies to  $\varphi|_{y_1} \equiv (\neg x_1 \vee$   
 597  $\neg x_2 \vee y_2 \vee y_3) \wedge (x_2 \vee \neg y_3 \vee y_2)$ . Using the unateness checks again, we now find  
 598 that  $\varphi|_{y_1}$  is positive unate in  $y_2$ , but neither positive nor negative unate in  $y_3$ .  
 599 Setting the Skolem function for  $y_2$  to `true`, the specification further simplifies to  
 600  $(\varphi|_{y_1})|_{y_2} \equiv \text{true}$ . Hence, any Skolem function for  $y_3$  suffices; in particular, we choose  
 601  $y_3 \equiv \text{false}$ . We have thus solved the `BoolSkFnSyn` problem for the given specification,  
 602 obtaining the constant Skolem functions  $\psi_1 \equiv \psi_2 \equiv \text{true}$  and  $\psi_3 \equiv \text{false}$ .

603 From the definition of unateness, we can see that checking unateness can be  
 604 reduced to checking (un)satisfiability of a propositional formula:  $\varphi$  is positive (resp.  
 605 negative) unate in  $v$  iff the formula  $\varphi|_{\neg v} \wedge \neg \varphi|_v$  (resp.  $\varphi|_v \wedge \neg \varphi|_{\neg v}$ ) is unsatisfiable.  
 606 A variant of this unateness check is used in [6] and other recent approaches to  
 607 `BoolSkFnSyn` (e.g., [4, 5]). In the other direction, checking validity of an arbitrary  
 608 formula  $\varphi$  can be reduced to checking if the formula  $z \vee \varphi$  is positive unate in  $z$ ,  
 609 where  $z \notin \text{sup}(\varphi)$ . Thus, unateness checking is `coNP`-hard, and cannot be done in  
 610 polynomial time unless  $\mathbf{P} = \mathbf{NP}$ . However, we can have sufficient conditions for  
 611 unateness that are checkable in polynomial time. For example, if  $v$  (resp.  $\neg v$ ) is a  
 612 *pure literal* in  $\varphi$ , i.e., the negation of the literal does not appear as the label of any  
 613 leaf in a NNF circuit representation of  $\varphi$ , then  $\varphi$  is positive (resp. negative) unate  
 614 in  $v_i$ . The above structural condition can clearly be checked in time linear in the  
 615 size of the NNF circuit representing  $\varphi$ .

#### 616 4.3.2 Functionally determined or implicitly defined variables

617 Suppose the specification  $\varphi$  *uniquely* defines a system output variable as a func-  
 618 tion of system input variables and other system output variables. We call such

619 a variable *functionally determined* or *implicitly defined* in  $\varphi$ . For example, if  $\varphi \equiv$   
620  $(\neg y_i \vee y_j) \wedge (\neg y_i \vee x_k) \wedge (y_i \vee \neg y_j \vee \neg x_k) \wedge \dots$ , then we can infer  $(y_i \Leftrightarrow (y_j \wedge x_k))$   
621 and hence,  $y_i$  is functionally determined (henceforth called FD) in  $\varphi$ . The implied  
622 functional dependencies like  $(y_i \Leftrightarrow (y_j \wedge x_k))$  are called *functional definitions* of FD  
623 variables. Given a set  $\mathbf{T} \subseteq \mathbf{Y}$  of FD system outputs in  $\varphi$ , we let  $\text{Fun}_{\mathbf{T}}$  denote the  
624 conjunction of functional definitions of all variables in  $\mathbf{T}$ . We say that  $(\mathbf{T}, \text{Fun}_{\mathbf{T}})$   
625 is an *acyclic system of functional definitions* if no variable in  $\mathbf{T}$  transitively depends  
626 on itself via the functional definitions in  $\text{Fun}_{\mathbf{T}}$ . The main observation is that for a  
627 given acyclic system  $(\mathbf{T}, \text{Fun}_{\mathbf{T}})$  obtained from  $\varphi$ , we can simply replace each of the  
628 output variables in  $\mathbf{T}$  by their functional definitions. Recall that these functional  
629 definitions are in terms of system inputs and other system outputs. Thus, once  
630 Skolem functions for all system outputs other than those in  $\mathbf{T}$  are generated, we  
631 can generate Skolem functions for those in  $\mathbf{T}$  simply by substituting the already  
632 generated Skolem functions in the functional definitions in  $\text{Fun}_{\mathbf{T}}$ . This can be done  
633 in polynomial time by effectively connecting the outputs of sub-circuits represent-  
634 ing already generated Skolem functions to corresponding inputs of sub-circuits  
635 representing functional definitions in  $\text{Fun}_{\mathbf{T}}$ .

636 The above idea is remarkably simple and results in considerable simplifica-  
637 tion in practical benchmarks. The reason is that functionally determined variables  
638 occur widely in practice and are often easy to identify. For instance, specifica-  
639 tions containing functionally determined variables arise naturally when a non-CNF  
640 Boolean formula is converted to CNF via Tseitin encoding [67], and are easily iden-  
641 tifiable as patterns in the formula. Given the widespread use of Tseitin encoding in  
642 obtaining CNF formulas, such variables have a surprisingly large impact on bench-  
643 marks. As a result many practical tools for BoolSkFnSyn, (including [53, 5, 4, 29])  
644 first identify and eliminate (at least some!) functionally determined variables be-  
645 fore processing the formulas.

646 A note about Beth definability, as applied to quantified propositional formulas,  
647 is pertinent here. By a celebrated theorem of Beth [10], a system output  $y_i$  that  
648 is implicitly defined by a specification  $\varphi$  also has an explicit definition in terms  
649 of the system inputs and other system outputs. Such an explicit definition can  
650 indeed serve as the functional definition for  $y_i$ . However, Beth's theorem doesn't  
651 immediately give us an explicit definition of  $y_i$ ; indeed, it can be computationally  
652 expensive to extract an explicit definition of  $y_i$  from  $\varphi$  in general. Practical tools  
653 therefore often use a range of heuristics to efficiently extract explicit definitions of  
654 implicitly defined system output variables. Fortunately, for variables introduced  
655 by Tseitin encoding, this can be done easily by matching patterns of clauses in a  
656 given CNF formula, as was illustrated in the example above. Such techniques, also  
657 called *syntactic gate extraction* (see e.g. [27]), are incomplete in general, but can be  
658 very effective in practice when reasoning about specifications containing Tseitin  
659 variables. In a recent work [60], a practically efficient, sound and complete *semantic*  
660 *gate extraction* technique for extracting explicit definitions of all implicitly defined  
661 variables, has been proposed. Incorporation of such techniques in Boolean Skolem  
662 function synthesis tools is likely to result in improved performance of such tools  
663 in practice.

### 664 4.3.3 Using maximal falsifiable sets of input clauses

665 Yet another class of specifications that admit relatively efficient synthesis in prac-  
 666 tice, follows from the work of [17]. Consider a specification  $\varphi(\mathbf{X}, \mathbf{Y})$  given in  
 667 CNF as a set of implicitly conjoined clauses  $C = \{C_1, \dots, C_k\}$ . Each clause po-  
 668 tentially has some literals over system inputs  $\mathbf{X}$ , and some literals over system  
 669 outputs  $\mathbf{Y}$ . Such a specification can of course be represented as a 3-level NNF  
 670 circuit. For all  $i \in \{1, \dots, k\}$ , let  $C_i|_{\mathbf{X}}$  denote the clause formed by taking the  
 671 disjunction of all literals over  $\mathbf{X}$  in  $C_i$ . Similarly, let  $C_i|_{\mathbf{Y}}$  be the clause formed  
 672 by disjoining all literals over  $\mathbf{Y}$  in  $C_i$ . The set of *input clauses* of  $\varphi$  is then de-  
 673 fined to be  $S_{in} = \{C_1|_{\mathbf{X}}, \dots, C_k|_{\mathbf{X}}\}$ . Similarly, the set of *output clauses* of  $\varphi$  is  
 674  $S_{out} = \{C_1|_{\mathbf{Y}}, \dots, C_k|_{\mathbf{Y}}\}$ . Note that if a clause has no system input (resp. system  
 675 output) literal, then the corresponding clause in  $S_{in}$  (resp.  $S_{out}$ ) is the empty  
 676 clause, representing false.

677 Let  $S$  be a subset of clauses in  $S_{in}$ . We say  $S$  is a *maximal falsifiable subset*  
 678 (MFS) of  $S_{in}$  if (i) there exists an assignment  $\pi$  that makes all clauses in  $S$  false,  
 679 and (ii) for every set  $S'$  such that  $S \subset S' \subseteq S_{in}$ , there exists no assignment that  
 680 makes all clauses in  $S'$  false. In a similar manner,  $\widehat{S} \subseteq S_{out}$  is said to a *maximal*  
 681 *satisfiable subset* (MSS) of  $S_{out}$  if (i) there exists an assignment  $\pi$  that makes all  
 682 clause in  $\widehat{S}$  true, and (ii) for every  $S''$  such that  $\widehat{S} \subset S'' \subseteq S_{out}$ , it is not possible  
 683 to find an assignment that renders all clauses in  $S''$  true.

684 With the above notation, the following results follow from the work of [17].

- 685 **Proposition 2** (a) Let  $\text{MFS}(S_{in})$  be the set of all MFS of  $S_{in}$ . Given  $\text{MFS}(S_{in})$ , the  
 686 `BoolSkFnSyn` problem for  $\varphi(\mathbf{X}, \mathbf{Y})$  can be solved in time linear in  $|\text{MFS}(S_{in})| \cdot$   
 687  $|\varphi(\mathbf{X}, \mathbf{Y})|$ , given access to an NP-oracle.
- 688 (b) Let  $\text{MSS}(S_{out})$  be the set of all MSS of  $S_{out}$ . Given  $\text{MSS}(S_{out})$ , the `BoolSkFnSyn`  
 689 problem for  $\varphi(\mathbf{X}, \mathbf{Y})$  can be solved in time linear in  $|\text{MSS}(S_{out})| \cdot |\varphi(\mathbf{X}, \mathbf{Y})|$ , given  
 690 access to an NP-oracle.

691 The intuition behind Proposition 2 can be informally stated as follows. For every  
 692 assignment  $\pi_{\mathbf{X}}$  of  $\mathbf{X}$ , consider the set of input clauses not satisfied by  $\pi_{\mathbf{X}}$ . By  
 693 definition, this set is included in some MFS, say  $S'$ , of  $S_{in}$ , and  $\pi_{\mathbf{X}}$  satisfies all  
 694 input clauses in  $S_{in} \setminus S'$ . Clearly, for each input clause in  $S_{in} \setminus S'$ , the corresponding  
 695 clause in the specification  $\varphi$  is also satisfied by  $\pi_{\mathbf{X}}$ , regardless of what we assign to  
 696  $\mathbf{Y}$ . Therefore, if we assign values to  $\mathbf{Y}$  such that all output clauses corresponding  
 697 to input clauses in  $S'$  are satisfied, the overall specification is satisfied. This gives a  
 698 way to solve `BoolSkFnSyn` by considering each MFS of  $S_{in}$  and by finding a satisfying  
 699 assignment of the corresponding subset of  $S_{out}$ . To see how `BoolSkFnSyn` can be  
 700 solved using MSS of  $S_{out}$ , let  $\pi_{\mathbf{Y}}$  be an assignment of  $\mathbf{Y}$  that satisfies an MSS, say  
 701  $S''$ , of  $S_{out}$ . Since  $S''$  is an MSS,  $\pi_{\mathbf{Y}}$  must falsify all clauses in  $S_{out} \setminus S''$ . Therefore,  
 702 if the assignment of  $\mathbf{X}$  satisfies all input clauses corresponding to output clauses  
 703 in  $S_{out} \setminus S''$ , the overall specification  $\varphi$  is again satisfied. Thus, `BoolSkFnSyn` can  
 704 be solved by considering satisfying assignments of every MSS of  $S_{out}$ .

705 In order to use Proposition 2 effectively, we must, of course, find ways to  
 706 compute  $\text{MFS}(S_{in})$  or  $\text{MSS}(S_{out})$  efficiently in practice. Fortunately, finding an  
 707 MFS of a given set of clauses, viz.  $S_{in}$ , is not hard. One way of doing this is by  
 708 analyzing the *consensus graph* [28] of  $S_{in}$ . This is an undirected graph with a node  
 709 for each clause in  $S_{in}$ , and an edge between two nodes iff the corresponding clauses  
 710 have no literal  $\ell$  that appear with opposite polarities in the two clauses. It is easy



711 to see that two clauses of  $S_{in}$  can be falsified at the same time iff there is an  
 712 edge between the corresponding nodes in the consensus graph. Thus, there is a  
 713 one-to-one correspondence between the MFS of  $S_{in}$  and the maximal cliques in its  
 714 consensus graph. The set of all MFS can therefore be enumerated by enumerating  
 715 the maximal cliques in the consensus graph. Finding a maximal clique in a graph  
 716 can be achieved by a greedy algorithm in time polynomial in the size of the graph.  
 717 This yields an algorithm for enumerating all MFS of  $S_{in}$  that takes time polynomial  
 718 in  $|S_{in}|$  and in the number of maximal cliques in the consensus graph of  $S_{in}$  [33].  
 719 The following result, derived from [17], is an immediate consequence of the above  
 720 observations.

721 **Proposition 3** [17] *Let  $\mathcal{C}$  be a class of CNF specifications such that the consensus*  
 722 *graphs of input clauses of specifications in  $\mathcal{C}$  have polynomially many maximal cliques.*  
 723 *This is the case, for example, if the consensus graphs are planar or chordal. Then, the*  
 724 *BoolSkFnSyn problem for class  $\mathcal{C}$  of specifications is in  $\text{P}^{\text{NP}}$  (i.e. solvable in polynomial*  
 725 *time by a Turing machine with access to an NP oracle).*

726 In practice, when implementing an algorithm for solving BoolSkFnSyn, a proposi-  
 727 tional satisfiability solver must be used in place of an NP-oracle. Given the signifi-  
 728 cant advances made in propositional satisfiability solving over the last few decades,  
 729 Proposition 3 allows us to identify a class of specifications for which BoolSkFnSyn  
 730 can be solved efficiently in practice.

731 Unlike in the case of finding MFS, however, we do not know of any polynomial-  
 732 time algorithm for finding an MSS of a given set of clauses. Indeed, finding an  
 733 MSS requires solving an instance of the MaxSAT problem, which is known to be  
 734 NP-complete. Therefore, Proposition 2(b) does not yield an easily identifiable class  
 735 of specifications for which BoolSkFnSyn can be solved efficiently in practice.

736 *Example 4* Consider the specification  $\varphi \equiv (x_1 \vee y_1) \wedge (x_2 \vee \neg y_1 \vee \neg y_2) \wedge (x_2 \vee \neg x_3 \vee$   
 737  $\neg y_2) \wedge (\neg x_1 \vee \neg y_1 \vee y_2)$ . Clearly,  $S_{in} = \{(x_1), (x_2), (x_2 \vee \neg x_3), (\neg x_1)\}$ , and  $S_{out} =$   
 $\{(y_1), (\neg y_1 \vee \neg y_2), (\neg y_2), (\neg y_1 \vee y_2)\}$ . The consensus graph of  $S_{in}$  is shown in Fig. 1.




Fig. 1: Consensus graph of  $S_{in}$

738 Notice that there are two maximal cliques in this graph, corresponding to two MFS  
 739 of  $S_{in}$ , i.e.  $\{(x_1), (x_2), (x_2 \vee \neg x_3)\}$  and  $\{(x_2), (x_2 \vee \neg x_3), (\neg x_1)\}$ . The corresponding  
 740 subsets of  $S_{out}$  are  $\{(y_1), (\neg y_1 \vee \neg y_2), (\neg y_2)\}$  and  $\{(\neg y_1 \vee \neg y_2), (\neg y_2), (\neg y_1 \vee y_2)\}$ ,  
 741 with satisfying assignments  $(y_1, y_2) = (\text{true}, \text{false})$  and  $(\text{false}, \text{false})$  respectively. Fur-  
 742 thermore, the subsets of input clauses not included in the MFS are  $\{(\neg x_1)\}$  and  
 743  $\{(x_1)\}$  respectively. Therefore, using the idea sketched above in the intuition be-  
 744 hind Proposition 2, we can obtain a Skolem function vector  $(\psi_1, \psi_2)$  that evaluates  
 745 as follows:  
 746

747 
$$\text{if } (\neg x_1) \text{ then } (\psi_1, \psi_2) = (\text{true}, \text{false}) \text{ else } (\psi_1, \psi_2) = (\text{false}, \text{false})$$

748 For more details of the technique, and also to see how a Skolem function vector  
749 can be obtained from the MSS of  $S_{out}$ , the reader is referred to [17].

## 750 5 Knowledge representation for Boolean Skolem function synthesis

751 The representation of the relational specification  $\varphi(\mathbf{X}, \mathbf{Y})$  has an important bearing  
752 on the computational complexity of solving **BoolSkFnSyn**. In the previous sections,  
753 we assumed that the specification is given by a NNF Boolean circuit, represented  
754 as a directed acyclic graph (DAG). It turns out that if this circuit has special struc-  
755 tural and functional properties, **BoolSkFnSyn** can indeed be solved efficiently. Of  
756 course, compiling an arbitrary specification to a circuit representation with these  
757 properties isn't always easy. Given the hardness results of Section 4.2, such compi-  
758 lation must necessarily require super-polynomial time and space in the worst-case,  
759 unless long-standing complexity theoretic conjectures are falsified. Nevertheless, it  
760 is interesting to study normal forms of circuit-based representations of relational  
761 specifications that allow efficient synthesis of Boolean Skolem functions.

762 We start by considering some circuit (and related) representations of Boolean  
763 formulas that have been studied extensively in the context of hardware verification,  
764 model counting, artificial intelligence etc. Consider an NNF circuit representing a  
765 Boolean formula  $\varphi$ . For every node  $N$  in a DAG representation of the circuit, let  
766  $lits(N)$  (resp.  $vars(N)$ ) denote the set of literals (resp. variables) labeling leaves  
767 that have a path from  $N$  in the DAG. Suppose for each AND-labeled node with  
768 children  $c_1, \dots, c_k$  in the DAG, we have  $vars(c_r) \cap vars(c_s) = \emptyset$  for all distinct  
769  $r, s \in \{1, \dots, k\}$ . The circuit is then said to be in *decomposable negation normal form*  
770 or **DNNF** [20]. **DNNF** is a popular representation form used in artificial intelligence  
771 applications, and enjoys many nice properties [20]. Similarly, free/reduced ordered  
772 binary decision diagrams (collectively, **BDDs**) [15] is a representation form for  
773 Boolean formulas that is widely used in hardware verification, symbolic model  
774 checking etc. As shown in [20], every such **BDD** can be converted to **DNNF** in  
775 linear time [20]. In [4], a slight generalization of **DNNFs**, called *weak decomposable*  
776 *negation normal form*, or **wDNNF**, was introduced. In **wDNNF**, for each AND-labeled  
777 internal node with children  $c_1, \dots, c_k$  in an NNF circuit, we have  $lits(c_r) \cap \{\neg \ell \mid \ell \in$   
778  $lits(c_s)\} = \emptyset$  for every distinct  $r, s \in \{1, \dots, k\}$ . Note that every **DNNF** circuit is also  
779 a **wDNNF** circuit.

780 We now have the following result from [4], which says that for all the above  
781 normal forms **BoolSkFnSyn** is easy, i.e., solvable in polynomial time and size.

782 **Theorem 2** ([4]) *Given an input specification  $\varphi(\mathbf{X}, \mathbf{Y})$  as a **DNNF** or **wDNNF** cir-*  
783 *cuit, or as a **BDD**, **BoolSkFnSyn** can be solved in time polynomial in the size of the*  
784 *representation. This yields a polynomial-sized Skolem function vector.*

*Example 5* Consider the following Boolean formulas in NNF over the set of variables  
 $x_1, x_2, x_3, y_1, y_2, y_3$ :

$$\varphi_1 \equiv (x_1 \vee x_2) \wedge (x_3 \vee \neg y_1) \wedge (\neg y_2 \vee y_3) \quad (1)$$

$$\varphi_2 \equiv (x_1 \vee x_2) \wedge (x_2 \vee \neg y_1) \wedge (\neg y_1 \vee y_2) \quad (2)$$

$$\varphi_3 \equiv (\neg x_1 \vee x_2) \wedge (x_1 \vee \neg y_2) \wedge (y_1 \vee y_2) \quad (3)$$

Each of these formulas is naturally represented as a 3-level NNF circuit with an AND-labeled root node having three OR-labeled children, and leaves labeled by literals as shown in Figure 2. Note that the representation of  $\varphi_1$  is in DNNF, and hence also in wDNNF. However, the representation of  $\varphi_2$  is not in DNNF, although it is in wDNNF. Indeed, in the circuit representing  $\varphi_2$ , the label  $\neg y_1$  appears in a leaf reachable from two distinct children of the AND-labeled root. However, there is no literal  $\ell$  such that a leaf labeled  $\ell$  is reachable from one child of the AND-labeled root, and a literal labeled  $\neg\ell$  is reachable from another child of the root. Hence, the requirement for wDNNF is satisfied by the representation of  $\varphi_2$ . Finally, the representation of  $\varphi_3$  is not in wDNNF since the AND-labeled root has two distinct children such that leaves labeled  $y_2$  and  $\neg y_2$  are reachable from these children. Of course, this also means that the representation of  $\varphi_3$  is not in DNNF either.

By Theorem 2, it is “easy” to synthesize Skolem functions for  $\varphi_1$  and  $\varphi_2$ , as given in Example 5. Importantly, the above theorem only gives a sufficient, but not necessary condition for efficient Boolean Skolem function synthesis. Indeed, it turns out that even for  $\varphi_3$  given in Example 5, Boolean Skolem functions can be synthesized efficiently. It is therefore interesting to ask if we can weaken the representational requirements beyond that of wDNNF, while ensuring polynomial time synthesis of Boolean Skolem functions. One easy way is to require the wDNNF condition only on literals corresponding to system outputs. This captures NNFs that are decomposable except on a set of atoms [20]. It can be seen that Theorem 2 applies in this setting as well. However, it turns out that we can go significantly beyond this, as we discuss in the next section.

### 5.1 A representation for efficient synthesis

Recall from the discussion in the initial part of Section 4 that if we can efficiently compute  $\varphi^{(i-1)}(\mathbf{X}, \mathbf{Y}_i^n)$ , i.e.  $\exists y_1, \dots, y_{i-1} \varphi(\mathbf{X}, \mathbf{Y})$ , for all  $i \in \{2, \dots, n\}$ , then we can solve BoolSkFnSyn efficiently. We will therefore try to arrive at a representational requirement weaker than that of wDNNF and that allows us to compute  $\varphi^{(i-1)}(\mathbf{X}, \mathbf{Y}_i^n)$  for all  $i \in \{2, \dots, n\}$ .

Consider an NNF circuit representing the formula  $\varphi(\mathbf{X}, \mathbf{Y})$ . The *output-positive form* of  $\varphi$ , denoted  $\hat{\varphi}$ , is obtained by replacing all leaves labeled  $\neg y_i$  by new variables  $\bar{y}_i$  in the NNF circuit representation of  $\varphi(\mathbf{X}, \mathbf{Y})$ . Thus,  $\hat{\varphi}$  is a formula with support  $\mathbf{X} \cup \mathbf{Y} \cup \bar{\mathbf{Y}}$ , where  $\bar{\mathbf{Y}}$  denotes the sequence (or set, depending on the




Fig. 2: NNF circuit representations of formula  $\varphi_1, \varphi_2, \varphi_3$  from Example 5.

819 context)  $(\overline{y_1}, \dots, \overline{y_n})$ . It is easy to see that  $\varphi(\mathbf{X}, \mathbf{Y}) \Leftrightarrow (\widehat{\varphi})[\overline{\mathbf{Y}} \mapsto \neg \mathbf{Y}]$ , where  $\neg \mathbf{Y}$   
 820 denotes the sequence  $(\neg y_1, \dots, \neg y_n)$ . Since the output-positive form, represented as  
 821 a NNF circuit, does not have any leaf labeled  $\neg y_i$  or  $\neg \overline{y_i}$  for any  $i \in \{1, \dots, n\}$ , it  
 822 follows that  $\widehat{\varphi}$  is monotone with respect to every such  $y_i$  and  $\overline{y_i}$ .

An immediate consequence of the above monotonicity is that we have

$$\exists y_1 \varphi(\mathbf{X}, \mathbf{Y}) \Leftrightarrow (\varphi|_{y_1} \vee \varphi|_{\neg y_1}) \Rightarrow (\widehat{\varphi}|_{y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n], \quad (4)$$

823 where we have used  $\widehat{\varphi}|_{y_1, \overline{y_1}}$  to denote  $(\widehat{\varphi}[y_1 \mapsto \text{true}])[\overline{y_1} \mapsto \text{true}]$ , and  $\overline{\mathbf{Y}}_2^n$  and  $\neg \mathbf{Y}_2^n$   
 824 to denote the sequences  $(\overline{y_2}, \dots, \overline{y_n})$  and  $(\neg y_2, \dots, \neg y_n)$ , respectively. In general, the  
 825 converse of the above implication, i.e.  $(\widehat{\varphi}|_{y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n] \Rightarrow (\varphi|_{y_1} \vee \varphi|_{\neg y_1})$ ,  
 826 doesn't always hold. However, if we can ensure (for example, by imposing restric-  
 827 tions on the representation of  $\varphi$ ) that the converse implication also holds, then we  
 828 will have  $\exists y_1 \varphi(\mathbf{X}, \mathbf{Y}) \Leftrightarrow (\widehat{\varphi}|_{y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$ . This will immediately give us an  
 829 efficient way to obtain  $\exists y_1 \varphi(\mathbf{X}, \mathbf{Y})$ . Specifically, we can simply set  $y_1$  and  $\overline{y_1}$  to true  
 830 in  $\widehat{\varphi}$ , and set all other  $\overline{y_i}$  to  $\neg y_i$ , in order to obtain  $\exists y_1 \varphi(\mathbf{X}, \mathbf{Y})$ . As already seen  
 831 earlier, efficient existential quantification of system output variables from  $\varphi(\mathbf{X}, \mathbf{Y})$   
 832 directly leads to an efficient way of computing Skolem functions. Hence, it is mean-  
 833 ingful to investigate what restrictions on the representation of  $\varphi$  ensure that the  
 834 converse of implication (4) holds.

835 We start by asking: *when is implication (4) given above strict, i.e. when does its*  
 836 *converse not hold?* Clearly, this happens iff there is an assignment  $\pi$  of  $\mathbf{X}$  and  $\mathbf{Y}_2^n$   
 837 that renders  $(\widehat{\varphi}|_{y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$  true and also simultaneously renders  $\exists y_1 \varphi(\mathbf{X}, \mathbf{Y})$   
 838 false. It follows from the definitions of  $\exists y_1 \varphi(\mathbf{X}, \mathbf{Y})$  and  $\widehat{\varphi}(\mathbf{X}, \mathbf{Y}, \overline{\mathbf{Y}})$  that assignment  
 839  $\pi$  must cause both  $(\widehat{\varphi}|_{y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$  and  $(\widehat{\varphi}|_{\neg y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$  to evalu-  
 840 ate to false. Since  $\varphi$  is monotone with respect to  $y_1$  and  $\overline{y_1}$ , it also follows that  
 841  $(\widehat{\varphi}|_{\neg y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$  evaluates to false under assignment  $\pi$ . Thus, assignment  
 842  $\pi$  causes  $\widehat{\varphi}[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$  to “semantically behave like”  $y_1 \wedge \overline{y_1}$ .

The above discussion yields the important intuition that  $\exists y_1 \varphi(\mathbf{X}, \mathbf{Y})$  is se-  
 mantically equivalent to  $(\widehat{\varphi}|_{y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$  iff  $\widehat{\varphi}[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$  can never be  
 made to behave like  $y_1 \wedge \overline{y_1}$  under any assignment of  $\mathbf{X}$  and  $\mathbf{Y}_2^n$ . In other words,  
 $\forall y_1 \forall \overline{y_1} (\widehat{\varphi}[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n] \Leftrightarrow (y_1 \wedge \overline{y_1}))$  must be unsatisfiable. By virtue of the mono-  
 tonicity properties of  $\widehat{\varphi}$ , the above condition simplifies to the requirement that  
 $(\widehat{\varphi}|_{y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n] \wedge \neg(\widehat{\varphi}|_{\neg y_1, \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n] \wedge \neg(\widehat{\varphi}|_{y_1, \neg \overline{y_1}})[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$   
 is unsatisfiable. This intuition can now be inductively lifted to the general case.

Towards this end, let  $\overbrace{(\text{true} \dots \text{true})}^t$  denote a sequence of  $t$  Boolean constants, each  
 being true. For each  $i \in \{1, \dots, n\}$ , we now define a formula  $[\widehat{\varphi}]_i$ , also called the  $i^{\text{th}}$   
 reduct of  $\varphi$ , as follows.

$$[\widehat{\varphi}]_i \equiv ((\widehat{\varphi}[\mathbf{Y}_1^{i-1} \mapsto \overbrace{(\text{true} \dots \text{true})}^{i-1}]))[\overline{\mathbf{Y}}_1^{i-1} \mapsto \overbrace{(\text{true} \dots \text{true})}^{i-1}][\overline{\mathbf{Y}}_{i+1}^n \mapsto \neg \mathbf{Y}_{i+1}^n]. \quad (5)$$

843 Thus, we take  $\widehat{\varphi}$  and set all  $y_j$  and  $\overline{y_j}$  for  $j \in \{1, \dots, i-1\}$  to true, and all  $\overline{y_k}$  for  
 844  $k \in \{i+1, \dots, n\}$  to  $\neg y_k$ , in order to get  $[\widehat{\varphi}]_i$ . The reduct  $[\widehat{\varphi}]_1$  is simply defined as  
 845  $\widehat{\varphi}[\overline{\mathbf{Y}}_2^n \mapsto \neg \mathbf{Y}_2^n]$ . Note that the support of  $[\widehat{\varphi}]_i$  includes  $\overline{y_i}$  in addition to  $\mathbf{X} \cup \mathbf{Y}_i^n$ .

846 Using arguments similar to that used above, we can now show that  $\exists \mathbf{Y}_1^i \varphi(\mathbf{X}, \mathbf{Y})$   
 847  $\Rightarrow ([\widehat{\varphi}]_i)|_{y_i, \overline{y_i}}$ . Furthermore, the converse implication holds iff  $[\widehat{\varphi}]_i$  cannot be made  
 848 to semantically behave like  $y_i \wedge \overline{y_i}$  for any assignment of  $\mathbf{X}$  and  $\mathbf{Y}_{i+1}^n$ , i.e. iff  
 849  $([\widehat{\varphi}]_i)|_{y_i, \overline{y_i}} \wedge \neg([\widehat{\varphi}]_i)|_{\neg y_i, \overline{y_i}} \wedge \neg([\widehat{\varphi}]_i)|_{y_i, \neg \overline{y_i}}$  is unsatisfiable. Referring back to

850 the discussion in the initial part of Section 4, it follows that if the above unsatisfi-  
 851 ability condition holds, then both  $([\widehat{\varphi}]_{i+1})|_{y_{i+1}, \neg \overline{y_{i+1}}}$  and  $\neg([\widehat{\varphi}]_{i+1})|_{\neg y_{i+1}, \overline{y_{i+1}}}$  serve  
 852 as Skolem functions for  $y_{i+1}$  (in terms of  $\mathbf{X} \cup \mathbf{Y}_{i+2}^n$ ) in  $\varphi(\mathbf{X}, \mathbf{Y})$ . Specifications that  
 853 satisfy the above unsatisfiability condition for all reducts  $[\widehat{\varphi}]_i$  are said to be in *Syn-*  
 854 *thesis Negation Normal Form* or *SynNNF*, and the corresponding Skolem functions  
 855 alluded to above are called *GACKS* functions, following the terminology of [4]. Note  
 856 that if  $\varphi(\mathbf{X}, \mathbf{Y})$  is in *SynNNF*, then computing the *GACKS* functions is easy, i.e.,  
 857 can be done in polynomial time. Formally, we have the following definition.

858 **Definition 1** [4] An NNF circuit representing a specification  $\varphi(\mathbf{X}, \mathbf{Y})$  is said to  
 859 be in *SynNNF* with respect to the sequence  $\mathbf{Y}$  of system outputs iff the formula  
 860  $([\widehat{\varphi}]_i)|_{y_i, \overline{y_i}} \wedge \neg([\widehat{\varphi}]_i)|_{\neg y_i, \overline{y_i}} \wedge \neg([\widehat{\varphi}]_i)|_{y_i, \neg \overline{y_i}}$  is unsatisfiable for all  $i \in \{1, \dots, n\}$

861 *Example 6* Consider again  $\varphi_3 \equiv (\neg x_1 \vee x_2) \wedge (x_1 \vee \neg y_2) \wedge (y_1 \vee y_2)$  from Example 5,  
 862 represented as the rightmost circuit in Fig. 2. We have seen that this representation  
 863 is neither in *wDNNF* nor in *DNNF*. However, with respect to the sequence of system  
 864 outputs  $(y_1, y_2)$ , it is in *SynNNF*. To see this, note that  $[\widehat{\varphi}_3]_1$  cannot be equivalent  
 865 to  $y_1 \wedge \overline{y_1}$  for any assignment of the other variables as  $y_1$  does not occur negatively  
 866 at all. Furthermore, in obtaining  $[\widehat{\varphi}_3]_2$ , we must assign *true* to  $y_1$ ; hence the clause  
 867  $y_1 \vee y_2$  becomes *true*. As a result,  $[\widehat{\varphi}_3]_2$  cannot evaluate to  $y_2 \wedge \overline{y_2}$  for any assignment  
 868 of  $x_1$  and  $x_2$ . Hence, we conclude that the representation of  $\varphi_3$  as the rightmost  
 869 circuit in Fig. 2 is in *SynNNF*.

870 Note that the definition of *SynNNF* makes crucial reference to a sequence (or  
 871 ordering) of variables in  $\mathbf{Y}$ . Indeed, if we change the ordering of system output  
 872 variables, say from  $(y_1, y_2)$  to  $(y_2, y_1)$  in the example of  $\varphi_3$  discussed above, then  
 873  $\varphi_3$  is no longer in *SynNNF* with respect to this new ordering. Specifically, for  
 874 the assignment in which  $x_1 = \text{false}$  and  $y_1 = \text{false}$ ,  $[\widehat{\varphi}_3]_1$  becomes semantically  
 875 equivalent to  $y_2 \wedge \overline{y_2}$ .

876 In [4], it is also shown that *SynNNF* strictly subsumes ( $\varphi_3$  being an example!)  
 877 previously considered normal forms including *wDNNF*, *DNNF* and *BDDs*. In fact,  
 878 we can say more. In the following theorem, sizes and times are in terms of the  
 879 number of system input and system output variables, i.e.  $|\mathbf{X}| + |\mathbf{Y}|$ .

880 **Proposition 4** ([4]) *Every specification in BDD, DNNF or wDNNF form is either*  
 881 *already in SynNNF or can be compiled in linear time to SynNNF. Moreover, there exist*  
 882 *polynomial-sized SynNNF specifications that only admit*

- 883 (i) *exponential sized BDD representations*
- 884 (ii) *super-polynomial sized wDNNF and DNNF representations, unless P = NP.*

885 Finally, we come to the practical utility of *SynNNF*, which is formalized in the  
 886 following result.

887 **Theorem 3** ([4]) *If a relational specification  $\varphi(\mathbf{X}, \mathbf{Y})$  is given in SynNNF, the GACKS*  
 888 *functions serve as polynomial sized Skolem functions for  $\varphi$ , and can be computed in*  
 889 *polynomial time. Hence BoolSkFnSyn is solvable in polynomial time for SynNNF speci-*  
 890 *fications.*

891 From Theorem 1 and Theorem 3, it follows that it is not possible to compile  
 892 an arbitrary relational specifications to *SynNNF* in polynomial time, unless some  
 893 long-standing complexity-theoretic conjectures are falsified. Such hardness results

894 for knowledge compilation are not uncommon in Computer Science, and similar  
 895 results are known for other important problems like model counting, satisfiability  
 896 checking, consistency checking and the like. Nevertheless, this has motivated  
 897 researchers to build compilers that work well in practice, thereby facilitating ef-  
 898 ficient solutions for important classes of problems. For example, several compil-  
 899 ers for converting an arbitrary formula into DNNF and its variants are presented  
 900 in [20, 22, 51, 43, 49]. Similarly, there are several mature tools (viz. [30, 63, 11])  
 901 that can be used to compile a propositional formula into a BDD. This approach of  
 902 converting a given specification into a BDD and then generating Skolem functions  
 903 is used, for instance, in [26] and also in one of the experimental pipelines reported  
 904 in [5]. In [4], a compiler called C2SYN was described that converts a relational  
 905 specification given in CNF directly to SynNNF. We refer the interested reader to [4]  
 906 for more details of C2SYN.

907 To complete the discussion on SynNNF, we note that SynNNF captures a seman-  
 908 tic requirement. This is unlike BDD, DNNF and wDNNF, all of which impose purely  
 909 syntactic requirements on the structure of the representation, that can be checked  
 910 in time polynomial in the size of the representation. Normal forms defined by se-  
 911 mantic conditions are however not new, e.g., the disjoint decomposable negation  
 912 normal form (dDNNF) uses a semantic condition in its definition (see [21]). The  
 913 semantic condition does, however, mean that the problem of checking if a circuit  
 914 is in SynNNF is not always easy.

915 **Proposition 5 ([56])** *Checking whether a given formula is in SynNNF w.r.t a given*  
 916 *ordering on the variables is coNP-complete. Further, checking whether it is in SynNNF*  
 917 *w.r.t any ordering is in  $\Sigma_2^P$ .*

918 In [56], the above result was established for a more general normal form. In fact,  
 919 the normal form considered in [56] not only generalizes SynNNF but also precisely  
 920 characterizes polynomial time and polynomial sized Boolean Skolem function syn-  
 921 thesis. We refer interested readers to [56] for more details regarding this form.

## 922 6 Algorithmic Paradigms for Boolean Skolem function synthesis

923 We have seen earlier that efficient algorithms for BoolSkFnSyn are unlikely, due to  
 924 the hardness results given in Theorem 1. However, this refers to the “worst-case  
 925 complexity” or efficiency for all inputs, which does not always translate to use-case  
 926 hardness. Given the practical relevance of the problem, different approaches have  
 927 been tried to design algorithms and build software tools that work well for real-  
 928 life benchmarks. Indeed, these tools have also been shown to work well in several  
 929 practical instances. In this section, we discuss in some detail one such approach,  
 930 that we call the *guess-check-repair paradigm* for Boolean Skolem function synthesis.  
 931 Before that, let us quickly survey other (mostly orthogonal) approaches that have  
 932 been explored for algorithmic solutions to BoolSkFnSyn.

933 – *Proof systems and proof rules.* This approach is mostly applicable to specifica-  
 934 tions  $\varphi(\mathbf{X}, \mathbf{Y})$  that are realizable, i.e.  $\forall \mathbf{X} \exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y})$  is valid. In [47, 9, 39, 8],  
 935 special proof systems for quantified Boolean formulas have been proposed,  
 936 and then Skolem functions have been extracted from a proof of validity of  
 937  $\forall \mathbf{X} \exists \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y})$ . While this works well with short proofs of validity, there are

938 challenges when such proofs are long or when no such proof exists, e.g. if the  
939 specification is unrealizable. As the factorization example in Section 3, it is of-  
940 ten important and useful to synthesize Skolem functions even for unrealizable  
941 specifications.

942 – *Incremental determinization.* A relational specification may functionally deter-  
943 mine some system outputs, as explained in Section 4.3.2. However, there may  
944 be other system outputs that are constrained but not completely functionally  
945 determined. In [53], a technique for incrementally determinizing such system  
946 outputs is described. The technique makes us of highly effective strategies used  
947 in modern conflict-driven clause learning (CDCL) based propositional satisfi-  
948 ability solvers to yield a practically efficient algorithm for Boolean Skolem  
949 function synthesis. The interested reader is referred to [58] for details of CDCL  
950 satisfiability solvers. The incremental determinization technique of [53] was  
951 further developed as a system of proof rules in [54, 52].

952 – *Synthesis via functional composition of circuits.* A completely different approach  
953 to `BoolSkFnSyn` is considered in [36, 37, 66], where iterated compositions (or  
954 substitutions) of Boolean circuits are used to synthesize Skolem functions.  
955 Given  $\varphi(\mathbf{X}, \mathbf{Y})$ , the basic idea here is to express one system output, say  $y_1$ ,  
956 as a Skolem function in terms of other system outputs and system inputs.  
957 While techniques similar to self-substitution have been used to generate such  
958 a Skolem function in [36, 66], interpolation based techniques have been used  
959 in [37]. Once such a Skolem function is obtained, it is composed with (or sub-  
960 stituted in)  $\varphi(\mathbf{X}, \mathbf{Y})$  to effectively existentially quantify  $y_1$  from  $\varphi(\mathbf{X}, \mathbf{Y})$ . This  
961 yields a simplified specification with one less system output. By repeating this  
962 process, we can eventually obtain a Skolem function for  $y_n$  in terms of only  
963 the system inputs. Subsequently, the Skolem function for  $y_n$  (in terms of only  
964 system inputs) can be substituted in the Skolem function for  $y_{n-1}$  (in terms  
965 of  $y_n$  and system inputs) to obtain a Skolem function for  $y_{n-1}$  in terms of  
966 only system inputs. By continuing this process, Skolem functions for all sys-  
967 tem outputs in terms of system inputs can be obtained. While this approach is  
968 simple to understand, it suffers from the drawback that iterated composition  
969 (or substitution) can result in an exponential blow-up in the representation  
970 of Boolean formulas. Hence, tools using this approach have been empirically  
971 found not to scale well to large benchmarks.

972 – *ROBDD-based techniques.* ROBDDs are widely used as compact representations  
973 of complex Boolean formulas. Researchers have therefore developed techniques  
974 for synthesizing Boolean Skolem functions from relational specifications given  
975 as ROBDDs. In [41], Kukula and Shiple presented one such technique in which  
976 a circuit that is structurally similar to the ROBDD representation of the spec-  
977 ification is generated to implement Boolean Skolem functions. In Kunčák et  
978 al [42], a generic framework for functional synthesis with unbounded domains  
979 like integers is described. As part of their exposition, the authors of [42] also  
980 suggest using ROBDDs with *input-first ordering* of variables. This approach has  
981 been developed further in [26], where a new algorithm called *TrimSubstitute* was  
982 proposed that optimizes the application of the self-substitution technique (see  
983 Section 4) to ROBDDs with input-first variable ordering. For factored specifica-  
984 tions, i.e. specifications that are conjunctions of sub-specifications, ideas from  
985 symbolic model checking using implicitly conjoined ROBDDs have been used  
986 to enhance the scalability of ROBDD-based synthesis further in [65]. Note that

the works of [42, 26, 65] attempt to synthesize Skolem functions directly as ROBDDs. This can be significantly more difficult than generating Skolem functions as Boolean circuits from ROBDD specifications. Indeed, we know from Proposition 4 and Theorem 3 that it is possible to generate Boolean circuits representing Skolem functions in polynomial-time from specifications given as ROBDDs. This holds regardless of the variable order used in the ROBDD representing the specification. Note, however, that the Skolem functions generated by application of Theorem 3 may not be compactly representable as ROBDDs. Interestingly, the requirement of having input-first ordering of variables when representing specifications as ROBDDs, as in the works of [42, 26, 65], may result in significantly larger ROBDDs compared to the case when there are no restrictions on the variable ordering. This may be viewed as the price that has to be paid in order to obtain the Skolem functions as ROBDDs themselves.

- *Input-output separation.* We have already discussed in Section 4.3.3 how literals in the clauses of a CNF specification can be partitioned to yield a set of input clauses and a set of output clauses. We also discussed in the same section specific conditions under which either the set of input clauses can be processed to obtain Skolem functions efficiently in practice. This idea has been developed further in [17], yielding a *back-and-forth* algorithm that alternates between processing of input clauses and output clauses to generate Skolem functions as decision lists [55]. This approach has been shown to work on some difficult classes of benchmarks, for which several other state-of-the-art techniques run out of steam.
- *Template/sketch-based techniques.* In addition to the above algorithmic techniques, template-based [64] and sketch-based [61] approaches have been developed, when we have information about the set of candidate Skolem functions. In the absence of such information, however, these techniques are not very effective.

We wish to emphasize that despite the diversity of techniques, there is no single technique that dominates others when solving `BoolSkFnSyn`. Furthermore, it is still largely unclear which technique would perform best for a given benchmark. This suggests the use of a portfolio solver, in which we can try multiple techniques and choose the one that best suits a given problem instance. On a related note, the knowledge representation approach presented earlier allows us to understand what input representations make the problem easy to solve, without providing an efficient technique to compile a given specification into a desired normal form. Coming up with better compilation algorithms and insights into which tool performs well on which benchmark, are part of ongoing and future work.

## 6.1 A guess, check and repair paradigm for synthesis

In the rest of this section, we focus on *one* specific algorithmic paradigm for solving `BoolSkFnSyn`, that has been developed recently in a series of papers [38, 2, 3, 5] and further augmented in [29]. Let us start by recalling that, sometimes we may get “lucky” in that the representation of the relational specification may already have structure (as explained in the previous sections) that permits efficient Boolean Skolem function synthesis. However, this raises three questions: (i) how do we get lucky? (ii) how easy is it to check if we have been lucky and (iii) what do we do when



1033 we are not lucky? Indeed, in practical applications, there is no guarantee that the  
1034 representation of the relational specification has structure that makes it amenable  
1035 to efficient synthesis. The *guess, check and repair* paradigm, that lies (sometimes  
1036 implicitly) at the heart of several existing works on `BoolSkFnSyn`, address these  
1037 questions very elegantly. In this section, we elucidate this generic paradigm as  
1038 well as show how it is instantiated in practice. The paradigm can be broken into  
1039 three key steps.

- 1040 – The first step runs efficiently in practice (viz. polynomial time relative to an  
1041 NP-oracle) and generates polynomial-sized guesses (or candidates) for Skolem  
1042 functions. If the representation of the relational specification has desirable  
1043 properties (such as those mentioned in previous sections), then these candi-  
1044 dates are often good enough to serve as Skolem functions themselves.
- 1045 – Even if the representation of the relational specification does not satisfy re-  
1046 strictions that guarantee correctness of the guesses made above, the guessed  
1047 Skolem functions may still be correct. We must therefore check if the guessed  
1048 Skolem functions can indeed serve as correct Skolem functions. As we show be-  
1049 low, this requires a single call to an NP-oracle, practically implemented using  
1050 a propositional satisfiability solver.
- 1051 – Finally, if the above check results in a negative answer (i.e. not all the guessed  
1052 Skolem functions are correct), we need to repair the guesses to obtain correct  
1053 Skolem functions. This is the third step of the paradigm, and can be done in  
1054 several ways. Given the computational hardness results, we know that in the  
1055 worst case, this phase may take exponential time. However, in practice, we are  
1056 often able to do much better!

1057 The reason we call this a paradigm, rather than an algorithm, is that one can  
1058 take different algorithms for solving each of the above steps and put them together  
1059 to obtain an overall algorithm that solves `BoolSkFnSyn`. We describe each of these  
1060 steps in more detail, along with some algorithms for implementing the steps, in  
1061 the next three subsections.

### 1062 6.1.1 Science of Guessing

1063 It is not surprising that the initial guesses of Skolem functions play an important  
1064 role in the guess-check-repair paradigm of solving `BoolSkFnSyn`. As mentioned ear-  
1065 lier, if the representation of the relational specification has desirable properties  
1066 (viz. being in `SynNNF`), then the initial guesses (viz. the GACKS functions alluded  
1067 to in Section 5) already serve as correct Skolem functions without any need for  
1068 further checking. Note, however, that Theorem 2 only asserts that a specification  
1069 being in `SynNNF` is a sufficient, not necessary, condition for the GACKS functions to  
1070 be correct Skolem functions. So, if GACKS functions are used as the initial guesses  
1071 for Skolem function, they may work for more general specifications (that are not  
1072 in `SynNNF`) too! This is indeed what was empirically observed in [3, 5], where  
1073 GACKS functions were found to be correct Skolem functions for a large collection  
1074 of benchmarks, not all of which were in `SynNNF`. In the works of [47, 53], coming  
1075 up with good initial candidates for Skolem functions from appropriate representa-  
1076 tions of the specification (or from a proof of its realizability), has often been called  
1077 *preprocessing*, or *initialization*. It turns out that this is not only a crucial step for  
1078 effective Boolean Skolem function synthesis, but also has deep connections with

1079 the area of knowledge representation and compilation. Indeed, in [4], this aspect  
 1080 has been explored in detail, and an algorithm presented to compile a specifica-  
 1081 tion given in CNF to a representational form (SynNNF) where the initial guesses of  
 1082 Skolem functions can always be correctly made.

1083 Another important consideration when guessing candidate Skolem functions  
 1084 is the kind of “errors” that are allowed in the guessed functions. For example,  
 1085 the work of [38, 5] requires the guessed Skolem functions to either be under-  
 1086 approximations or over-approximations of correct Skolem functions. Thus, the  
 1087 error in a candidate Skolem function is always one-sided in these approaches.  
 1088 While this allows for easier proofs of soundness and termination (when applied in  
 1089 conjunction with appropriate techniques for repair), the repair of guessed Skolem  
 1090 functions with one-sided error may take longer in practice. Other more recent  
 1091 approaches, e.g. [29], have relaxed the restriction of one-sided errors, and used  
 1092 machine-learning based heuristics for arriving at good initial guesses of Skolem  
 1093 functions, albeit with two-sided errors.

#### 1094 6.1.2 Checking the guess

1095 This step involves deciding whether a guessed Skolem function vector suffices to  
 1096 serve as a correct Skolem function vector for the given relational specification. If  
 1097 the answer turns out to be in the negative, it is also useful to obtain a valuation  
 1098 of the system inputs  $\mathbf{X}$  for which at least one of the guessed Skolem functions  
 1099 generates an incorrect value for the corresponding system output. It turns out  
 1100 that this problem can be easily reduced to checking the unsatisfiability of an  
 1101 appropriately constructed propositional formula, called the error formula in [38].

Given the relational specification  $\varphi(\mathbf{X}, \mathbf{Y})$ , suppose the vector of guessed Skolem  
 functions for the system outputs  $\mathbf{Y}$  is  $\Psi = (\psi_1, \dots, \psi_n)$ . Following [38], the error  
 formula for  $\varphi$  with respect to this guess is defined as:

$$\varepsilon_{\varphi, \Psi}(\mathbf{X}, \mathbf{Y}, \mathbf{Y}') \equiv \varphi(\mathbf{X}, \mathbf{Y}') \wedge \bigwedge_{i=1}^n (y_i \Leftrightarrow \psi_i) \wedge \neg\varphi(\mathbf{X}, \mathbf{Y})$$

1102 Note that the first sub-formula in  $\varepsilon_{\varphi, \Psi}$  has free variables from  $\mathbf{Y}' = (y'_1, \dots, y'_n)$ ,  
 1103 where each  $y'_i$  is a fresh variable, not originally present in  $\varphi(\mathbf{X}, \mathbf{Y})$ . This sub-  
 1104 formula asserts that there exists some valuation of  $\mathbf{Y}$  that renders  $\varphi(\mathbf{X}, \mathbf{Y})$  true.  
 1105 This is needed in order to focus only on those assignments of  $\mathbf{X}$  for which  $\varphi(\mathbf{X}, \mathbf{Y})$   
 1106 is satisfiable. The second sub-formula in  $\varepsilon_{\varphi, \Psi}$  assigns variables in  $\mathbf{Y}$  to the values  
 1107 given by the corresponding guessed Skolem functions in  $\Psi$ , and the third sub-  
 1108 formula checks if this assignment falsifies the specification  $\varphi$ . As proved in [38, 5],  
 1109 the formula  $\varepsilon_{\varphi, \Psi}$  is unsatisfiable iff  $\Psi$  is a correct Skolem function vector for the  
 1110 specification  $\varphi(\mathbf{X}, \mathbf{Y})$ .

1111 Thus, checking if a candidate Skolem function vector suffices to serve as a cor-  
 1112 rect Skolem function vector can be done using a single call to an NP-oracle. In  
 1113 practice, a propositional satisfiability solver is used for this purpose, and this has  
 1114 its own advantages. Unlike an NP-oracle that simply yields a “Yes”/”No” answer,  
 1115 an invocation of a propositional satisfiability solver also generates a satisfying as-  
 1116 signment, say  $\pi$ , of  $\varepsilon_{\varphi, \Psi}(\mathbf{X}, \mathbf{Y}, \mathbf{Y}')$  in case the candidate Skolem function vector  
 1117 is incorrect. From the definition of  $\varepsilon_{\varphi, \Psi}$ , it is easy to see that in such a case, the  
 1118 projection of  $\pi$  on  $\mathbf{X}$  gives an assignment of system inputs for which at least one

1119 guessed Skolem function in  $\Psi$  generates an incorrect value for the corresponding  
 1120 system output. Indeed, there exists an assignment of system outputs (viz. projec-  
 1121 tion of  $\pi$  on  $\mathbf{Y}'$ ) that satisfies the specification  $\varphi$  for the above assignment of  $\mathbf{X}$ ,  
 1122 and yet the values given by the guessed Skolem function vector (viz. projection of  
 1123  $\pi$  on  $\mathbf{Y}$ ) fail to satisfy the specification with the same assignment of  $\mathbf{X}$ .

### 1124 6.1.3 The Art of Repairing

1125 Finally, if the above check reports that the guessed Skolem function vector is incor-  
 1126 rect, we need a way to repair the guess. As mentioned above, using a propositional  
 1127 satisfiability solver to check the satisfiability of the error formula also gives us an  
 1128 assignment of  $\mathbf{X}$ ,  $\mathbf{Y}$  and  $\mathbf{Y}'$  that demonstrates why the guessed Skolem function  
 1129 vector  $\Psi$  is not correct. This information is crucial in repairing the incorrect guess.  
 1130 Indeed, multiple approaches have been used in the literature to repair incorrect  
 1131 guesses of Skolem functions.

1132 – In [38, 3, 5], the authors use an approach called *expansion based repair*. This  
 1133 works when the guessed Skolem functions always have one-sided error. Intu-  
 1134 itively, if a guessed Skolem function is an under-approximation of a correct  
 1135 Skolem function, the set of assignments on which it evaluates to true must be  
 1136 “expanded” to repair the guess. Similarly, if a guess Skolem function is an  
 1137 over-approximation of a correct Skolem function, the set of assignments on  
 1138 which it evaluates to false must be “expanded” to effect the repair. For every  
 1139 Skolem function in error, the repair strategy ensures that errors, if any, of the  
 1140 repaired Skolem function are of the same nature (i.e. under-approximation er-  
 1141 ror or over-approximation error) as in the original erroneous Skolem function.  
 1142 Thus, the erroneous Skolem function vector monotonically approaches a cor-  
 1143 rect Skolem function vector, with at least one erroneous Skolem function in the  
 1144 vector being changed in each iteration of repair. The actual repair is obtained  
 1145 by examining the satisfying assignment returned by the (un)satisfiability check  
 1146 of  $\varepsilon_{\varphi, \Psi}$  to determine which Skolem functions in  $\Psi$  need to be repaired. In ad-  
 1147 dition, the satisfying assignment is “generalized” to obtain a set of (instead of  
 1148 a single) assignments of  $\mathbf{X}$  for which the same expansion-based repair must be  
 1149 applied. This helps in reducing the number of repair iterations, since a good  
 1150 “generalization” may address problems that can arise with multiple valuations  
 1151 of  $\mathbf{X}$ . After each iteration of repair, the error formula is reconstructed for  
 1152 the repaired Skolem function vector, and its (un)satisfiability checked again.  
 1153 Since there are only finitely many valuations of  $\mathbf{X}$  and finitely many Skolem  
 1154 functions to repair, it is not hard to show that expansion based repair is guar-  
 1155 anteed to terminate with a correct Skolem function vector. However, the way  
 1156 in which the expansion is done crucially determines how fast and effective the  
 1157 repair algorithm is. The interested reader is referred to [5] for more details of  
 1158 expansion-based repair techniques.

1159 – In [2], the authors use the circuit structure of the input specification to paral-  
 1160 lelize the task of repairing an incorrectly guessed Skolem function vector. While  
 1161 the basic approach remains one of expansion-based repair, the added benefit of  
 1162 parallelization shows in significantly reduced synthesis times, as demonstrated  
 1163 in [2].

1164 – In a recent work [29], a new and powerful idea of repair has been used in a  
 1165 guess-check-repair tool for solving BoolSkFnSyn. Specifically, the authors of [29]

1166 delve deeper into the reason why an assignment of  $\mathbf{X}$  leads some candidate  
 1167 Skolem functions in  $\Psi$  to evaluate to the wrong values for the corresponding  
 1168 system outputs. Using powerful techniques based on minimal unsatisfiable core  
 1169 extraction, they are able to obtain significant generalizations starting from a  
 1170 single satisfying assignment of  $\varepsilon_{\varphi, \Psi}$ . This technique has the advantage that it  
 1171 can repair initial guesses of Skolem functions that even have two-sided errors  
 1172 (i.e. the guessed Skolem function is neither an under-approximation nor an  
 1173 over-approximation of a correct Skolem function). As shown by an extensive  
 1174 set of experiments in [29], allowing two-sided errors in the initial guesses of  
 1175 Skolem functions chosen by means of machine learning techniques, followed  
 1176 by powerful unsatisfiable core based repair techniques can be very effective in  
 1177 synthesizing Boolean Skolem functions for a large set of benchmarks.

1178 While we have given a high-level overview of some algorithms that implement  
 1179 the guess-check-repair paradigm of solving `BoolSkFnSyn`, there appears to be a  
 1180 lot of uncharted territory, and the last word on the topic of practically efficient  
 1181 algorithm for `BoolSkFnSyn` is yet to be said. Our primary focus in this article has  
 1182 been on the theory behind the algorithms. However, the proof of the pudding is  
 1183 indeed in the eating, and we strongly recommend the interested reader to go through  
 1184 the relevant papers to see the practical performance of the ideas and algorithms  
 1185 sketched above.

## 1186 7 Conclusion

1187 In this article, we have explained how Skolem function synthesis lies at the heart  
 1188 of several lines of research. These have spanned from theoretical questions, both  
 1189 about existence and explicit construction of Skolem functions in the general setting  
 1190 of first order logic, to more practical questions about the computational hardness  
 1191 and efficient algorithms in simpler settings. In the simplest case of the proposi-  
 1192 tional setting, we have presented a deeper insight into computational hardness  
 1193 issues, and also how specific properties of the representation of the specification  
 1194 can be exploited to design practically efficient algorithms. Finally, we have dis-  
 1195 cussed a powerful paradigm, called guess-check-repair, that has been instantiated  
 1196 in multiple tools to obtain practically efficient strategies to solve the `BoolSkFnSyn`  
 1197 problem on a large suite of benchmarks.

1198 Multiple lines of research emerge most naturally from the results discussed  
 1199 here. One immediate question is whether structural (or even functional) proper-  
 1200 ties for representations of specifications can be identified for non-Boolean settings,  
 1201 such that they allow efficient synthesis of Skolem functions. Furthermore, can we  
 1202 lift the ideas and techniques for synthesis beyond Boolean specifications, to say  
 1203 specifications in temporal logics? Similarly, the synthesis question discussed in  
 1204 this article does not take into account dependency information for existentially  
 1205 quantified variables. Finding Skolem functions for dependency quantified Boolean  
 1206 formulas is an important problem, and it would be interesting to consider exten-  
 1207 sions of existing `BoolSkFnSyn` techniques to solve this problem. Overall, given its  
 1208 central importance, we hope researchers will be encouraged to pursue research on  
 1209 synthesis of Skolem functions for richer classes of specifications, both from theo-  
 1210 retical and practical points of view.

## References

- 1211 1. Akshay S, Chakraborty S (2021) On synthesizing Skolem func-  
1212 tions for first-order logic formulae. CoRR Identifier: 2102.07463,  
1213 (<https://arxiv.org/abs/2102.07463>)  
1214
- 1215 2. Akshay S, Chakraborty S, John AK, Shah S (2017) Towards parallel boolean  
1216 functional synthesis. In: TACAS 2017 Proceedings, Part I, pp 337–353, URL  
1217 [https://doi.org/10.1007/978-3-662-54577-5\\_19](https://doi.org/10.1007/978-3-662-54577-5_19)
- 1218 3. Akshay S, Chakraborty S, Goel S, Kulal S, Shah S (2018) How hard is boolean  
1219 functional synthesis. In: In CAV 2018 Proceedings, URL [https://doi.org/10.](https://doi.org/10.1007/978-3-662-54577-5_19)  
1220 [1007/978-3-662-54577-5\\_19](https://doi.org/10.1007/978-3-662-54577-5_19)
- 1221 4. Akshay S, Arora J, Chakraborty S, Krishna S, Raghunathan D, Shah S (2019)  
1222 Knowledge compilation for boolean functional synthesis. In: Proc. of Formal  
1223 Methods in Computer Aided Design (FMCAD)
- 1224 5. Akshay S, Chakraborty S, Goel S, Kulal S, Shah S (2020) Boolean functional  
1225 synthesis: hardness and practical algorithms. Form Methods Syst Des DOI  
1226 <https://doi.org/10.1007/s10703-020-00352-2>
- 1227 6. Andersson G, Bjesse P, Cook B, Hanna Z (2002) A proof engine approach to  
1228 solving combinational design automation problems. In: Proceedings of the 39th  
1229 Annual Design Automation Conference, ACM, New York, NY, USA, DAC '02,  
1230 pp 725–730, DOI 10.1145/513918.514101, URL [http://doi.acm.org/10.1145/](http://doi.acm.org/10.1145/513918.514101)  
1231 [513918.514101](http://doi.acm.org/10.1145/513918.514101)
- 1232 7. Arora S, Barak B (2009) Computational Complexity: A Modern Approach,  
1233 1st edn. Cambridge University Press, USA
- 1234 8. Balabanov V, Jiang JHR (2012) Unified QBF certification and its applications.  
1235 Form Methods Syst Des 41(1):45–65, DOI 10.1007/s10703-012-0152-6, URL  
1236 <http://dx.doi.org/10.1007/s10703-012-0152-6>
- 1237 9. Benedetti M (2005) sKizzo: A Suite to Evaluate and Certify QBFs. In: Proc.  
1238 of CADE, Springer-Verlag, pp 369–376
- 1239 10. Beth E (1953) On Padoa’s method in the theory of definition.  
1240 Indagationes Mathematicae (Proceedings) 56:330–339, DOI [https://doi.](https://doi.org/10.1016/S1385-7258(53)50042-3)  
1241 [org/10.1016/S1385-7258\(53\)50042-3](https://doi.org/10.1016/S1385-7258(53)50042-3), URL [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S1385725853500423)  
1242 [science/article/pii/S1385725853500423](https://www.sciencedirect.com/science/article/pii/S1385725853500423)
- 1243 11. Biere A (1998) ABCD. <http://fmv.jku.at/abcd/>
- 1244 12. Bockmayr A (1993) Logic Programming with Pseudo-Boolean Constraints,  
1245 MIT Press, Cambridge, MA, USA, pp 327–350
- 1246 13. Boole G (1847) The Mathematical Analysis of Logic. Philosophical Library,  
1247 URL <https://books.google.co.in/books?id=zv4YAQAATAAJ>
- 1248 14. Brenguier R, Pérez GA, Raskin JF, Sankur O (2014) Absynthe: abstract  
1249 synthesis from succinct safety specifications. In: Proceedings 3rd Workshop on  
1250 Synthesis (SYNT’14), Open Publishing Association, Electronic Proceedings in  
1251 Theoretical Computer Science, vol 157, pp 100–116, DOI 10.4204/EPTCS.157.  
1252 11, URL <http://arxiv.org/abs/1407.5961v1>
- 1253 15. Bryant RE (1986) Graph-based algorithms for boolean function manipu-  
1254 lation. IEEE Trans Comput 35(8):677–691, DOI 10.1109/TC.1986.1676819,  
1255 URL <http://dx.doi.org/10.1109/TC.1986.1676819>
- 1256 16. Buttner W, Simonis H (1987) Embedding boolean expressions into logic pro-  
1257 gramming. Journal of Symbolic Computation 4(2):191–205

- 1258 17. Chakraborty S, Fried D, Tabajara LM, Vardi MY (2018) Functional synthesis  
1259 via input-output separation. In: 2018 Formal Methods in Computer Aided  
1260 Design, FMCAD 2018, Austin, TX, USA, October 30 - November 2, 2018, pp  
1261 1–9
- 1262 18. Chandrasekaran V, Srebro N, Harsha P (2008) Complexity of inference in  
1263 graphical models. In: UAI 2008, Proceedings of the 24th Conference in Uncer-  
1264 tainty in Artificial Intelligence, Helsinki, Finland, July 9-12, 2008, pp 70–78
- 1265 19. Dao TBH, Djelloul K (2006) Solving first-order constraints in the theory of  
1266 the evaluated trees. In: Proceedings of the Constraint Solving and Constraint  
1267 Logic Programming 11th Annual ERCIM International Conference on Recent  
1268 Advances in Constraints, Springer-Verlag, Berlin, Heidelberg, CSCLP'06, p  
1269 108–123
- 1270 20. Darwiche A (2001) Decomposable negation normal form. *J ACM* 48(4):608–  
1271 647
- 1272 21. Darwiche A (2001) On the tractable counting of theory models and its applica-  
1273 tion to truth maintenance and belief revision. *Journal of Applied Non-Classical*  
1274 *Logics* 11(1-2):11–34
- 1275 22. Darwiche A (2002) A compiler for deterministic, decomposable negation nor-  
1276 mal form. In: Proceedings of the Eighteenth National Conference on Artificial  
1277 Intelligence (AAAI), AAAI Press, Menlo Park, California, pp 627–634
- 1278 23. Davis M, Matijasevic Y, Robinson J (1976) Hilbert’s tenth problem. diophan-  
1279 tine equations: positive aspects of a negative solution. In: Proceedings of sym-  
1280 posia in pure mathematics, vol 28, pp 323–378
- 1281 24. De Micheli G (1994) *Synthesis and Optimization of Digital Circuits*, 1st edn.  
1282 McGraw-Hill Science/Engineering/Math, USA
- 1283 25. Finkbeiner B (2016) Synthesis of reactive systems. In: Esparza J, Grumberg  
1284 O, Sickert S (eds) *Dependable Software Systems Engineering*, NATO Science  
1285 for Peace and Security Series - D: Information and Communication Security,  
1286 vol 45, IOS Press, pp 72–98, DOI 10.3233/978-1-61499-627-9-72, URL <https://doi.org/10.3233/978-1-61499-627-9-72>
- 1287 26. Fried D, Tabajara LM, Vardi MY (2016) BDD-based boolean functional syn-  
1288 thesis. In: *Computer Aided Verification - 28th International Conference, CAV*  
1289 *2016*, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part II, pp 402–  
1290 421
- 1291 27. Fu Z, Malik S (2007) Extracting logic circuit structure from conjunctive normal  
1292 form descriptions. In: 20th International Conference on VLSI Design (VLSI  
1293 Design 2007), Sixth International Conference on Embedded Systems (ICES  
1294 2007), 6-10 January 2007, Bangalore, India, IEEE Computer Society, pp 37–  
1295 42
- 1296 28. Ganian R, Szeider S (2017) New width parameters for model counting. In:  
1297 *Theory and Applications of Satisfiability Testing – SAT 2017*, Springer Inter-  
1298 national Publishing, pp 38–52
- 1299 29. Golia P, Roy S, Meel KS (2020) Manthan: A data-driven approach for boolean  
1300 function synthesis. In: *Proceedings of International Conference on Computer-*  
1301 *Aided Verification (CAV)*
- 1302 30. Group BLV (2008) *ABC: A system for sequential synthesis and verification*
- 1303 31. Hopcroft JE, Motwani R, Ullman JD (2006) *Introduction to Automata The-*  
1304 *ory, Languages, and Computation (3rd Edition)*. Addison-Wesley Longman  
1305 Publishing Co., Inc., USA  
1306

- 1307 32. Huth M, Ryan M (2004) *Logic in Computer Science: Modelling and Reasoning*  
1308 *about Systems*. Cambridge University Press, USA
- 1309 33. Ignatiev A, Morgado A, Planes J, Marques-Silva J (2013) Maximal falsifiabil-  
1310 ity. In: *Logic for Programming, Artificial Intelligence, and Reasoning*, Springer  
1311 Berlin Heidelberg, Berlin, Heidelberg, pp 439–456
- 1312 34. Impagliazzo R, Paturi R (2001) On the complexity of k-SAT. *J Comput Syst*  
1313 *Sci* 62(2):367–375
- 1314 35. Jacobs S, Bloem R, Brenguier R, Könighofer R, Pérez GA, Raskin J, Ryzhyk  
1315 L, Sankur O, Seidl M, Tentrup L, Walker A (2015) The second reactive syn-  
1316 thesis competition (SYNTCOMP 2015). In: *Proceedings Fourth Workshop on*  
1317 *Synthesis, SYNT 2015, San Francisco, CA, USA, 18th July 2015.*, pp 27–57
- 1318 36. Jiang JHR (2009) Quantifier elimination via functional composition. In: *Proc.*  
1319 *of CAV*, Springer, pp 383–397
- 1320 37. Jiang JR, Lin H, Hung W (2009) Interpolating functions from large boolean  
1321 relations. In: *2009 International Conference on Computer-Aided Design, IC-*  
1322 *CAD 2009, San Jose, CA, USA, November 2-5, 2009*, pp 779–784
- 1323 38. John A, Shah S, Chakraborty S, Trivedi A, Akshay S (2015) Skolem functions  
1324 for factored formulas. In: *FMCAD*, pp 73–80
- 1325 39. Jussila T, Biere A, Sinz C, Kröning D, Wintersteiger C (2007) A First Step  
1326 Towards a Unified Proof Checker for QBF. In: *Proc. of SAT, LNCS*, vol 4501,  
1327 Springer, pp 201–214
- 1328 40. Kuehlmann A, Paruthi V, Krohm F, Ganai MK (2002) Robust boolean rea-  
1329 soning for equivalence checking and functional property verification. *IEEE*  
1330 *Trans on CAD of Integrated Circuits and Systems* 21(12):1377–1394, URL  
1331 <http://dblp.uni-trier.de/db/journals/tcad/tcad21.html#KuehlmannPKG02>
- 1332 41. Kukula JH, Shiple TR (2000) Building circuits from relations. In: *Computer*  
1333 *Aided Verification, 12th International Conference, CAV 2000, Chicago, IL,*  
1334 *USA, July 15-19, 2000, Proceedings*, pp 113–123
- 1335 42. Kuncak V, Mayer M, Piskac R, Suter P (2010) Complete functional synthesis.  
1336 *SIGPLAN Not* 45(6):316–329
- 1337 43. Lagniez JM, Marquis P (2017) An improved decision-DNNF compiler. In: *Pro-*  
1338 *ceedings of the 24th International Joint Conference on Artificial Intelligence*  
1339 *(IJCAI)*, pp 667–673
- 1340 44. Löwenheim L (1910) Über die Auflösung von Gleichungen in Logischen Gebi-  
1341 etkalkul. *Math Ann* 68:169–207
- 1342 45. Macii E, Odasso G, Poncino M (2006) Comparing different boolean unification  
1343 algorithms. In: *Proc. of 32nd Asilomar Conference on Signals, Systems and*  
1344 *Computers*, pp 17–29
- 1345 46. Madsen M, van de Pol J (2020) Polymorphic types and effects with boolean  
1346 unification. *Proceedings of the ACM on Programming Languages* 4(OOPSLA)
- 1347 47. Marijn Heule MS, Biere A (2014) Efficient Extraction of Skolem Functions  
1348 from QRAT Proofs. In: *Formal Methods in Computer-Aided Design, FMCAD*  
1349 *2014, Lausanne, Switzerland, October 21-24, 2014*, pp 107–114
- 1350 48. Martin U, Nipkow T (1989) Boolean unification - the story so far. *J Symb*  
1351 *Comput* 7(3-4):275–293, DOI 10.1016/S0747-7171(89)80013-6, URL [http://](http://dx.doi.org/10.1016/S0747-7171(89)80013-6)  
1352 [dx.doi.org/10.1016/S0747-7171\(89\)80013-6](http://dx.doi.org/10.1016/S0747-7171(89)80013-6)
- 1353 49. Muise C, McIlraith SA, Beck JC, Hsu E (2012) DSHARP: Fast d-DNNF Com-  
1354 pilation with sharpSAT. In: *Canadian Conference on Artificial Intelligence*

- 1355 50. Niemetz A, Preiner M, Lonsing F, Seidl M, Biere A (2012) Resolution-based  
1356 certificate extraction for QBF - (tool presentation). In: Theory and Appli-  
1357 cations of Satisfiability Testing - SAT 2012 - 15th International Conference,  
1358 Trento, Italy, June 17-20, 2012. Proceedings, pp 430–435
- 1359 51. Oztok U, Darwiche A (2015) A top-down compiler for sentential decision dia-  
1360 grams. In: Proceedings of the 24th International Joint Conference on Artificial  
1361 Intelligence (IJCAI), pp 3141–3148
- 1362 52. Rabe MN (2019) Incremental determinization for quantifier elimination and  
1363 functional synthesis. In: Computer Aided Verification - 31st International Con-  
1364 ference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings,  
1365 Part II, pp 84–94
- 1366 53. Rabe MN, Seshia SA (2016) Incremental determinization. In: Theory and Ap-  
1367 plications of Satisfiability Testing - SAT 2016 - 19th International Conference,  
1368 Bordeaux, France, July 5-8, 2016, Proceedings, pp 375–392, DOI 10.1007/978-  
1369 3-319-40970-2\_23, URL [https://doi.org/10.1007/978-3-319-40970-2\\_23](https://doi.org/10.1007/978-3-319-40970-2_23)
- 1370 54. Rabe MN, Tentrup L, Rasmussen C, Seshia SA (2018) Understanding and ex-  
1371 tending incremental determinization for 2QBF. In: Computer Aided Verifica-  
1372 tion - 30th International Conference, CAV 2018, Held as Part of the Federated  
1373 Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part  
1374 II, pp 256–274
- 1375 55. Rivest R (1987) Learning decision lists. *Machine Learning* 2(3):229–246
- 1376 56. Shah P, Bansal A, Akshay S, Chakraborty S (2021) A normal form character-  
1377 ization for efficient boolean skolem function synthesis. *CoRR* abs/2104.14098,  
1378 URL <https://arxiv.org/abs/2104.14098>, 2104.14098
- 1379 57. Shukla A, Bierre A, Siedl M, Pulina L (2019) A survey on applications of  
1380 quantified boolean formula. In: Proceedings of the Thirty-First International  
1381 Conference on Tools with Artificial Intelligence (ICTAI), pp 78–84
- 1382 58. Silva JPM, Lynce I, Malik S (2021) Conflict-driven clause learning sat solvers.  
1383 In: Biere A, Heule M, van Maaren H, Walsch T (eds) *Handbook of Satisfia-  
1384 bility*, IOS Press, chap 4, pp 131–153
- 1385 59. Simonis H, Dincbas M (1987) Using an extended Prolog for digital circuit  
1386 design. In: *IEEE International Workshop on AI Applications to CAD Systems  
1387 for Electronics*, Springer International Publishing, pp 165–188
- 1388 60. Slivovsky F (2020) Interpolation-based semantic gate extraction and its ap-  
1389 plications to QBF preprocessing. In: *Computer Aided Verification - 32nd In-  
1390 ternational Conference, CAV 2020, Los Angeles, CA, USA, July 21-24, 2020,  
1391 Proceedings, Part I*, Springer, *Lecture Notes in Computer Science*, vol 12224,  
1392 pp 508–528
- 1393 61. Solar-Lezama A, Rabbah RM, Bodík R, Ebcioğlu K (2005) Programming  
1394 by sketching for bit-streaming programs. In: *Proceedings of the ACM SIG-  
1395 PLAN 2005 Conference on Programming Language Design and Implementa-  
1396 tion*, Chicago, IL, USA, June 12-15, 2005, pp 281–294
- 1397 62. Somenzi F (1999) Binary decision diagrams. In: *Calculational System Design*,  
1398 vol. 173 of *NATO Science Series F*, IOS Press, pp 303–366
- 1399 63. Somenzi F (2008) CUDD: CU decision diagram package release 2.5.0. <http://vlsi.colorado.edu/~fabio/CUDD/>
- 1400  
1401 64. Srivastava S, Gulwani S, Foster JS (2013) Template-based program verification  
1402 and program synthesis. *STTT* 15(5-6):497–518



- 
- 1403 65. Tabajara LM, Vardi MY (2017) Factored boolean functional synthesis. In: 2017  
1404 Formal Methods in Computer Aided Design, FMCAD 2017, Vienna, Austria,  
1405 October 2-6, 2017, pp 124–131
- 1406 66. Trivedi A (2003) Techniques in symbolic model checking. Master’s thesis, In-  
1407 dian Institute of Technology Bombay, Mumbai, India
- 1408 67. Tseitin GS (1968) On the complexity of derivation in propositional calculus.  
1409 Structures in Constructive Mathematics and Mathematical Logic, Part II,  
1410 Seminars in Mathematics pp 115–125