

Top-Down Knowledge Compilation for Counting Modulo Theories

Vincent Derkinderen^{1,2}, Pedro Zuidberg Dos Martires³, Samuel Kolb^{1,2}, and
Paolo Morettin⁴

¹ KU Leuven, Leuven, Belgium
{vincent.derkinderen, samuel.kolb}@kuleuven.be

² Leuven.AI, Leuven, Belgium

³ Örebro University, Örebro, Sweden
pedro.zuidberg-dos-martires@oru.se

⁴ University of Trento, Trento, Italy
paolo.morettin@unitn.it

Abstract

Propositional model counting ($\#SAT$) can be solved efficiently when the input formula is in deterministic decomposable negation normal form (d-DNNF). Translating an arbitrary formula into a representation that allows inference tasks, such as counting, to be performed efficiently, is called knowledge compilation. Top-down knowledge compilation is a state-of-the-art technique for solving $\#SAT$ problems that leverages the traces of exhaustive DPLL search to obtain d-DNNF representations. While knowledge compilation is well studied for propositional approaches, knowledge compilation for the (quantifier free) counting modulo theory setting ($\#SMT$) has been studied to a much lesser degree. In this paper, we discuss compilation strategies for $\#SMT$. We specifically advocate for a top-down compiler based on the traces of exhaustive $DPLL(T)$ search.

1 Introduction

The key motivation of knowledge compilation research is to compile a logical formula into a target language whose properties allow for certain tasks to be performed in time polynomial in the representation size [12]. For propositional model counting ($\#SAT$) problems, this representation is called deterministic decomposable negation normal form (d-DNNF) [8]. Huang and Darwiche realised that d-DNNF compilation can be achieved through exhaustive DPLL search by simply storing the search traces [9, 21]. Several state-of-the-art d-DNNF compilers are based on this principle, e.g. c2d [9], dSharp [29], miniC2D [31], and D4 [25].

The advantage of compiling to a d-DNNF formula instead of directly computing the (weighted) model count, arises with repeated inference. As the d-DNNF only depends on the logical formula, any weight change does not require recompiling. This splits model counting into a hard offline compilation step, the cost of which is amortized over (cheap) repeated online steps. This has for example been exploited in probabilistic logic programming for parameter learning and efficiently answering multiple queries [19, 27].

Top-down knowledge compilers, which are based on exhaustive DPLL, have so far mainly focused on $\#SAT$ while ignoring the domain of modulo theories [3]. In this paper we consider counting modulo theories ($\#SMT$) and discuss compilation strategy choices for the quantifier free setting (Section 4). Most importantly, we draw parallels with $\#SAT$ and identify a $DPLL(T)$ [30] based approach as a promising direction for a future knowledge compiler.

2 Background

DPLL traces & d-DNNF. The DPLL algorithm, introduced to solve the satisfiability problem (SAT), repeatedly branches on variables in order to find a satisfying assignment to all propositional Boolean variables [13, 14]. It is clear that any \vee -node following from a branching decision satisfies the *determinism* property: the children of an \vee -node do not share any models. This is important for model counting and model enumeration as to avoid considering a model more than once. The exhaustive version of the DPLL algorithm ($\#$ DPLL) is an extension that does not stop when a model is found, instead continuing to find all models. This version, which is used for model counting, enumeration and compilation problems, often includes *component decomposition*. This algorithmic improvement detects when there are sets of clauses that do not share any variables, i.e., *components*, and solves those independently [1]. The process results in an \wedge -node that satisfies the *decomposability* property: the children of \wedge -nodes do not share any variables. This property is important for task decomposability, for example in weighted model counting it ensures that each literal weight is considered at most once. Decomposability also avoids having \wedge -branches that could disagree on literals, which is important to avoid enumerating or considering logically inconsistent assignments. Additionally, the traces are also in negation normal form (NNF) because negation only occurs on literals and no other operations than \vee and \wedge take place. Thus, the trace of an exhaustive DPLL algorithm naturally yields a d-DNNF formula [9, 21].

Counting Modulo Theories. The satisfiability modulo theory (SMT) problem extends propositional satisfiability (SAT) with respect to a certain decidable background theory T [4]. Notable examples of theories T include linear real (\mathcal{LRA}) or integer (\mathcal{LIA}) arithmetic and fixed-size bit vectors (\mathcal{BV}). In SMT- \mathcal{LRA} , formulas include propositional and \mathcal{LRA} atoms of the form $\sum_i c_i x_i \bowtie c$ with $c_i, c \in \mathbb{Q}$, x_i a real variable, and $\bowtie \in \{<, >, \leq, \geq, =, \neq\}$. An example \mathcal{LRA} formula F would be $[(x \leq 0) \vee (x \geq 1)] \wedge [A \vee (x \leq 0)]$. The task is determining whether there exists an assignment to the variables such that both the input formula F and the background theory T are satisfied. For instance, in the example above, $(x \leq 1)$ entails A . Just as SMT generalizes SAT with additional theories, counting modulo theories ($\#$ SMT) generalizes model counting ($\#$ SAT) with additional theories. In stark contrast with the purely propositional setting, formulas involving numerical theories may have infinitely many models. While in this paper we mainly consider the quantifier-free \mathcal{LRA} setting, we expect the proposed ideas to apply to a broader set of quantifier-free theories.

$\text{DPLL}(T)$ is a generalisation of the DPLL algorithm that is designed to solve SMT problems [30]. The key difference with the classic DPLL algorithm is that it involves a theory specific solver that interacts with the DPLL algorithm, using the background theory T to propagate additional atoms when possible. For instance in the previous example, when a branch assigns $(x \geq 1)$ to true, the theory solver would additionally propagate $(x \leq 0)$ to be false. We note that most current implementations are actually extensions of the conflict driven clause learning (CDCL) algorithm, which augments the DPLL algorithm with conflict clause learning and backjumping [15, 23, 30]. For the purpose of our discussion we will use $\text{DPLL}(T)$ to refer to both $\text{DPLL}(T)$ and its $\text{CDCL}(T)$ augmentation.

3 d-DNNF for modulo theory

We now discuss in more detail the d-DNNF properties [12] in the context of $\#$ SMT.

A propositional formula is in **negation normal form** iff negation only occurs on literals, and the only other operations are \vee and \wedge . Similar to the DPLL algorithm, the traces of a $\text{DPLL}(T)$ algorithm produce NNF formulas. The fact that atoms can be associated with additional logic (e.g. \mathcal{LRA} atoms such as $x < 5$), and that a theory specific solver aids the propagation process, does not change this. The traces of a $\text{DPLL}(T)$ algorithm are also **deterministic**, i.e., the children of each \vee -node do not share any models, because every decision associated with \vee still partitions the models. The **decomposability** property is defined as \wedge -node branches not sharing any variables. In the modulo theory setting this needs further specification, because different atoms can still be linked together through their inner logic (e.g. shared continuous variables).

Decomposability on the level of atoms is insufficient to ensure the decomposability of counting tasks on the formula. Consider the following example where the \wedge -branches do not share any atoms, but do share variables: $[(x < 5) \vee (x < y)] \wedge [(x > 10) \vee (y > 10)]$. Here, $\{(x < 5), (x > y)\}$ and $\{(x > 10), (y > 10)\}$ are satisfying truth assignments for respectively the left and right \wedge -branch. Together however, they do not form a theory satisfying truth assignment. The benefit of decomposing on the level of variables, is that the previous situation is impossible, making it easier to ensure that each truth assignment is indeed theory consistent. Unfortunately such decomposability can not be achieved in general. For the purpose of our discussion on modulo theory compilation, we will therefore use the d in $d\text{-DNNF}$ to refer to decomposability on an atomic level rather than on a variable level. We will however expect that all truth assignments captured by the compiled formula are theory satisfiable. This is not necessarily ensured when using a theory agnostic compilation approach, which we discuss in more detail in the next section. For now, note that the traces of an exhaustive $\text{DPLL}(T)$ algorithm will produce exactly such $d\text{-DNNF}$ formulas where each truth assignment is theory satisfiable.

4 Compilation strategies

4.1 Theory aware versus theory agnostic

Approaches for compiling a formula with respect to a background theory T can be divided into two categories: theory aware versus theory agnostic approaches. In case of the latter, a Boolean abstraction of the input formula F is first created, i.e. the formula that is obtained by replacing every theory atom in F with a fresh Boolean atom. Afterwards, the Boolean abstraction can be compiled into a $d\text{-DNNF}$ formula using any off-the-shelf propositional compiler. As a consequential benefit, advancements made to those compilers are automatically inherited. Evidently, theory agnostic approaches also have a downside compared to approaches that are more theory aware. Namely, since the compiler is unaware of T , the resulting $d\text{-DNNF}$ formula may contain truth assignments that are not consistent with T . For example, when abstracting $F = (x \leq 0) \vee (x \geq 1)$, both literals could be assigned to true, forming a T -unsatisfiable truth assignment. It is then up to the down stream inference algorithm that uses the resulting $d\text{-DNNF}$, to deal with the inconsistent models, typically impacting negatively the run time of the downstream task. For example, this theory agnostic approach was adopted by [Kolb et al.](#) for solving weighted model integration (WMI) problems, i.e. weighted $\#SMT$ problems over \mathcal{LRA} formulas [24]. As a consequence of the theory-agnostic offline phase, a potentially large number of intermediate computations during the online integration procedure would result in regions with no models and would be discarded.

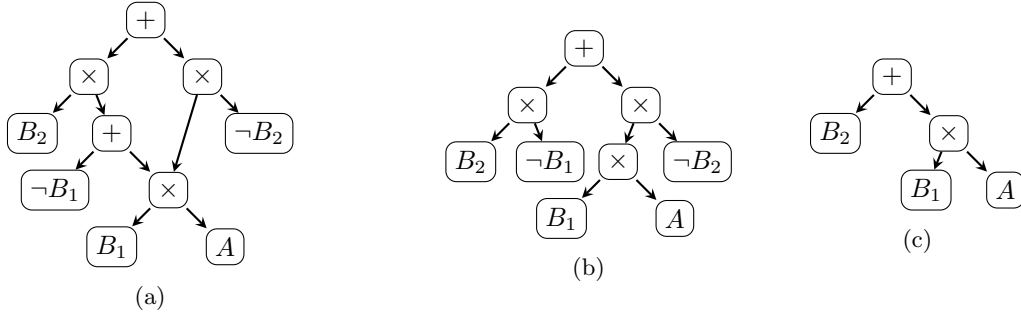


Figure 1: Different representations of $(B_1 \vee B_2) \wedge (\neg B_1 \vee A)$, with abbreviations $B_1 = (x < y - 1)$, $B_2 = (x > y + 1)$, and $A = (x > 20)$.

Figure 1 shows an example of a formula representation that may result from a theory agnostic approach (1a), compared to a theory aware approach (1b and 1c). In this example, the formula representation shrank once the theory was considered. This is not always necessarily the case: the effect that theory awareness has on the representation size heavily depends on the exploitable structure, which may increase or decrease. Depending on the application requirements, we could even consider a more condensed representation that omits theory implied atoms (Figure 1c).

4.2 Eager versus lazy solving

It is possible to exploit the existing purely propositional tools while still achieving theory awareness, by first adapting the Boolean abstracted input formula F in a way that no theory unsatisfiable truth assignments arise. This idea was used by early SMT solvers, where it is called ‘eager solving’ [30]. For example, if F includes both $(x \leq 0)$ and $(x \geq 1)$, which were respectively abstracted into A_0 and A_1 , then also add $A_0 \implies \neg A_1$ or $A_1 \implies \neg A_0$ to F . Afterwards, because F is purely propositional, any SAT solver (or in our case compiler) can be used. Since the adapted abstraction step is also present in eager SMT solvers, we can simply implement an eager theory compiler by using an existing eager SMT solver as foundation, swapping their SAT solver for a propositional d-DNNF compiler.

A major disadvantage of the eager approach is that depending on the theory T , it may introduce many additional variables, and greatly increase the formula size of the abstracted input formula F . For this reason, numerical theory SMT solvers have instead progressed towards using a so called ‘lazy solving’ [30] approach that we discuss next. This does not mean, however, that the eager approach has become irrelevant. It still forms the state-of-the-art SMT approach for certain theories such as \mathcal{BV}^1 . We hypothesize that similar conclusions can be drawn for the compilation setting, i.e., that a lazy approach is most suitable for numeric background theories.

In the most basic form, the lazy approach involves purely propositional reasoning over F , and verifying (partial) assignments with a theory solver to make sure that every model is indeed theory satisfiable. Enhancements of this approach have moved to a more proactive design where the theory solver instead helps propagating atoms and learning conflict clauses. This approach is the DPLL(T) (or CDCL(T)) algorithm discussed in Section 2 and is used by most state-of-the-art SMT solvers [2, 6, 7, 15, 17, 22].

¹The 2022 SMT competition winner of the incremental quantifier free \mathcal{BV} track was Yices 2 [17], which uses bit blasting for \mathcal{BV} , an eager solving approach.

4.3 Top-down versus bottom-up compilation

DPLL-based compilers are so called top-down compilers: they start the compilation process from the whole input formula and work their way down (by branching on literals). In contrast, bottom-up compilers process an input formula F from the bottom of the expression to the top. For example, if F is $(A \vee \neg B) \wedge (\neg A \vee C)$, then the two conjuncts are separately compiled before processing the conjunction itself. This design implies an apply-operation that performs conjunction, disjunction and negation. The apply-operation enables the incremental construction of (multiple) formulas, while a top-down compiler requires the complete formula to be known up front. To achieve an efficient apply-operation, however, the target language of bottom-up compilers is usually restricted to strongly deterministic and structured decomposable NNF, such as OBDD and the more general SDD class [10, 32], strict subsets of d-DNNF. In comparison, DPLL based top-down compilers instead produce Decision-DNNF, a strict superset of OBDD [11] (but not of SDD with which it partially overlaps).

An advantage of d-DNNF bottom-up compilation is that it does not restrict the input form. In contrast, top-down compilation is commonly restricted to formulas in conjunctive normal form (CNF). While it is possible to translate any propositional formula into CNF, this costs time and may result in an exponential expression (when using De Morgan’s law), or requires introducing additional variables [35].

An advantage of top-down compilation is that it has available the full information contained within F . This assumption allows more informed decisions, in turn leading to more compact compiled formulas. Oppositely, bottom-up compilation may have a pre-analysis step that considers the whole formula F to determine a suitable guiding structure (cf. static variable ordering heuristics for OBDD [33] or vtree heuristics for SDD [10]), but that structure can be suboptimal for intermediate formula representations. This naturally leads to the second point, which is that even when the final d-DNNF produced by bottom-up compilation is small, the intermediate results of the apply-operation may grow to be very large, severely impacting the run time and memory requirements [20].

In terms of compilation in the modulo theory domain, only bottom-up approaches have been considered so far. For example, the work on extended algebraic decision diagrams (XADD) [34] which uses a bottom-up approach. Within their implementation, they used a feasibility checker of a linear programming solver to prune theory unsatisfiable paths during- or after bottom-up compilation, observing impressive XADD size reductions. Other bottom-up compilers include work on difference decision diagrams [28] and linear decision diagrams [5].

5 Traces of an exhaustive DPLL(T) algorithm

In line with top-down compilation for the propositional setting, we propose to compile modulo theory formulas by storing the search traces generated by an exhaustive DPLL(T) algorithm. This is a theory aware, top-down, lazy approach, which we hypothesize to be very suitable for compiling formulas with numerical background theories such as \mathcal{LRA} (cf. Section 4.2).

Existing SMT solvers based on the DPLL(T) framework form a natural implementation starting point. As explained in Section 2, the traces of such an approach form d-DNNF formulas. To further augment the implementation, it is worth investigating optimisations developed for #DPLL. For example, component decomposition mentioned in Section 2 is rarely used for satisfiability problems but has shown to be highly beneficial for model counting and compiling, especially in combination with caching [1]. The augmentations are not straightforward however, because of the implicit interactions between atoms: components must be decomposed

on a variable level (rather than atom), and previous decisions must be properly considered. Especially the latter is a novel challenge for caching. Consider the following example formula F' : $((x < 3) \vee (x > 5)) \wedge ((y < 0) \vee (y > 4))$. Generally, branching on $(x > 5)$ does not imply $(y < 0)$. However, if $(x + y) < 5$ is a decision that has led to the intermediate formula F' , then branching on $(x > 5)$ while compiling F' does imply $(y < 0)$. This shows that the compiled form of F' is indeed influenced by previous decisions. This is in contrast to the propositional setting, where after making a decision l and adapting formula F , the remaining formula $F|_l$ (here F') is independent of previous decisions.

Related work To the best of our knowledge, we are the first to propose a theory aware top-down lazy compiler. The closest related works are [16, 26], who use a DPLL(T) like approach to count over modulo theories, but do not consider storing the traces for compilation. [18], in contrast, does consider compilation but is positioned further from DPLL(T), excluding optimizations such as conflict clause learning. Moreover, there is limited background theory support, for example atoms with arithmetic expressions such as $x + y < 2$ are not possible.

6 Conclusion & Future Work

We have provided a discussion on compilation strategies for (top-down) d-DNNF compilation for quantifier free modulo theories. We specifically propose compilation through an exhaustive version of the DPLL(T) algorithm, drawing parallels with the propositional setting’s evolution from SAT to #SAT. In future work, we aim to develop a tool based on these ideas, and plan to further investigate component decomposition and caching for modulo theories. Additionally, refining the d-DNNF properties based on application requirements, is of interest.

Acknowledgments

This work was partially supported by the Research Foundation-Flanders (FWO, 1SA5520N and S007318N), the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme, the EU H2020 ICT48 project “TAILOR” under contract #952215, and the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. PM is supported by the PNRR project FAIR - Future AI Research (PE00000013), under the NRRP MUR program funded by the NextGenerationEU.

References

- [1] Fahiem Bacchus, Shannon Dalmao, and Toniann Pitassi. DPLL with caching: A new algorithm for #SAT and Bayesian inference. *Electron. Colloquium Comput. Complex.*, 10(003), 2003.
- [2] Haniel Barbosa, Clark W. Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Mudathir Mohamed, Aina Niemetz, Andres Nötzli, Alex Ozdemir, Mathias Preiner, Andrew Reynolds, Ying Sheng, Cesare Tinelli, and Yoni Zohar. cvc5: A versatile and industrial-strength SMT solver. In *TACAS (1)*, volume 13243 of *Lecture Notes in Computer Science*, pages 415–442. Springer, 2022.

- [3] Clark Barrett and Cesare Tinelli. *Satisfiability modulo theories*. Springer, 2018.
- [4] Clark Barrett, Roberto Sebastiani, Sanjit A Seshia, Cesare Tinelli, A Biere, M Heule, H van Maaren, and T Walsh. Handbook of satisfiability. *Satisfiability modulo theories*, 185:825–885, 2009.
- [5] Sagar Chaki, Arie Gurfinkel, and Ofer Strichman. Decision diagrams for linear arithmetic. In *FMCAD*, pages 53–60. IEEE, 2009.
- [6] Jürgen Christ, Jochen Hoenicke, and Alexander Nutz. SMTInterpol: an interpolating SMT solver. In *SPIN*, volume 7385 of *Lecture Notes in Computer Science*, pages 248–254. Springer, 2012.
- [7] Alessandro Cimatti, Alberto Griggio, Bastiaan Joost Schaafsma, and Roberto Sebastiani. The MathSAT5 SMT solver. In *TACAS*, volume 7795 of *Lecture Notes in Computer Science*, pages 93–107. Springer, 2013.
- [8] Adnan Darwiche. On the tractable counting of theory models and its application to truth maintenance and belief revision. *Journal of Applied Non-Classical Logics*, 11(1-2):11–34, 2001.
- [9] Adnan Darwiche. New advances in compiling CNF to decomposable negation normal form. In *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI’04*, page 318–322, NLD, 2004. IOS Press. ISBN 9781586034528.
- [10] Adnan Darwiche. SDD: A new canonical representation of propositional knowledge bases. In *IJCAI*, 2011.
- [11] Adnan Darwiche and Auguste Hirth. On the reasons behind decisions. In *ECAI*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 712–720. IOS Press, 2020.
- [12] Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *J. Artif. Int. Res.*, 17(1):229–264, September 2002. ISSN 1076-9757.
- [13] Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, 1960.
- [14] Martin Davis, George Logemann, and Donald W. Loveland. A machine program for theorem-proving. *Commun. ACM*, 5(7):394–397, 1962.
- [15] Leonardo Mendonça de Moura and Nikolaj S. Bjørner. Z3: an efficient SMT solver. In *TACAS*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008.
- [16] Rodrigo De Salvo Braz, Ciaran O’Reilly, Vibhav Gogate, and Rina Dechter. Probabilistic inference modulo theories. In *IJCAI*, 2016.
- [17] Bruno Dutertre. Yices 2.2. In *CAV*, volume 8559 of *Lecture Notes in Computer Science*, pages 737–744. Springer, 2014.
- [18] Jonathan Feldstein and Vaishak Belle. Lifted reasoning meets weighted model integration. In *UAI*, volume 161 of *Proceedings of Machine Learning Research*, pages 322–332. AUAI Press, 2021.

- [19] Daan Fierens, Guy Van den Broeck, Joris Renkens, Dimitar Sht. Shterionov, Bernd Gutmann, Ingo Thon, Gerda Janssens, and Luc De Raedt. Inference and learning in probabilistic logic programs using weighted boolean formulas. *Theory Pract. Log. Program.*, 15(3):358–401, 2015.
- [20] Jinbo Huang and Adnan Darwiche. Using DPLL for efficient OBDD construction. In *SAT*, 2004.
- [21] Jinbo Huang and Adnan Darwiche. DPLL with a trace: From SAT to knowledge compilation. In *IJCAI*, pages 156–162. Professional Book Center, 2005.
- [22] Antti E. J. Hyvärinen, Matteo Marescotti, Leonardo Alt, and Natasha Sharygina. OpenSMT2: an SMT solver for multi-core and cloud computing. In *SAT*, volume 9710 of *Lecture Notes in Computer Science*, pages 547–553. Springer, 2016.
- [23] Roberto J. Bayardo Jr. and Robert Schrag. Using CSP look-back techniques to solve real-world SAT instances. In *AAAI/IAAI*, pages 203–208. AAAI Press / The MIT Press, 1997.
- [24] Samuel Kolb, Pedro Zuidberg Dos Martires, and Luc De Raedt. How to exploit structure while solving weighted model integration problems. In *UAI*, volume 115 of *Proceedings of Machine Learning Research*, pages 744–754. AUAI Press, 2019.
- [25] Jean-Marie Lagniez and Pierre Marquis. An improved Decision-DNNF compiler. In *IJCAI*, pages 667–673. ijcai.org, 2017.
- [26] Feifei Ma, Sheng Liu, and Jian Zhang. Volume computation for boolean combination of linear arithmetic constraints. In *CADE*, volume 5663 of *Lecture Notes in Computer Science*, pages 453–468. Springer, 2009.
- [27] Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Neural probabilistic logic programming in DeepProbLog. *Artif. Intell.*, 298:103504, 2021.
- [28] Jesper B. Møller, Jakob Lichtenberg, Henrik Reif Andersen, and Henrik Hulgaard. Difference decision diagrams. In *CSL*, volume 1683 of *Lecture Notes in Computer Science*, pages 111–125. Springer, 1999.
- [29] Christian J. Muise, Sheila A. McIlraith, J. Christopher Beck, and Eric I. Hsu. Dsharp: Fast d-DNNF compilation with sharpSAT. In *Canadian Conference on AI*, volume 7310 of *Lecture Notes in Computer Science*, pages 356–361. Springer, 2012.
- [30] Robert Nieuwenhuis, Albert Oliveras, and Cesare Tinelli. Solving SAT and SAT modulo theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL(T). *J. ACM*, 53(6):937–977, November 2006. ISSN 0004-5411. doi: 10.1145/1217856.1217859.
- [31] Umut Oztok and Adnan Darwiche. A top-down compiler for sentential decision diagrams. In *IJCAI*, pages 3141–3148. AAAI Press, 2015.
- [32] Knot Pipatsrisawat and Adnan Darwiche. New compilation languages based on structured decomposability. In *AAAI*, volume 8, pages 517–522, 2008.
- [33] Michael Rice and Sanjay Kulhari. A survey of static variable ordering heuristics for efficient BDD/MDD construction. *University of California, Tech. Rep*, page 130, 2008.

- [34] Scott Sanner, Karina Valdivia Delgado, and Leliane Nunes de Barros. Symbolic dynamic programming for discrete and continuous state MDPs. In *UAI*, pages 643–652. AUAI Press, 2011.
- [35] Grigori S Tseitin. On the complexity of derivation in propositional calculus. *Automation of reasoning: 2: Classical papers on computational logic 1967–1970*, pages 466–483, 1983.