

Canonical Decision Diagrams Modulo Theories

Massimo Michelutti^a, Gabriele Masina^a, Giuseppe Spallitta^a and Roberto Sebastiani^a

^aDISI, University of Trento, Italy

Abstract. Decision diagrams (DDs) are powerful tools to represent effectively propositional formulas, which are largely used in many domains, in particular in formal verification and in knowledge compilation. Some forms of DDs (e.g., OBDDs, SDDs) are *canonical*, that is, (under given conditions on the atom list) they univocally represent equivalence classes of formulas. Given the limited expressiveness of propositional logic, a few attempts to leverage DDs to SMT level have been presented in the literature. Unfortunately, these techniques still suffer from some limitations: most procedures are theory-specific; some produce theory DDs ($\mathcal{T}$ -DDs) which do not univocally represent $\mathcal{T}$ -valid formulas or $\mathcal{T}$ -inconsistent formulas; none of these techniques provably produces *theory-canonical* $\mathcal{T}$ -DDs, which (under given conditions on the $\mathcal{T}$ -atom list) univocally represent $\mathcal{T}$ -equivalence classes of formulas. Also, these procedures are not easy to implement, and very few implementations are actually available.

In this paper, we present a novel very-general technique to leverage DDs to SMT level, which has several advantages: it is very easy to implement on top of an AllSMT solver and a DD package, which are used as blackboxes; it works for every form of DDs and every theory, or combination thereof, supported by the AllSMT solver; *it produces theory-canonical $\mathcal{T}$ -DDs* if the propositional DD is canonical. We have implemented a prototype tool for both $\mathcal{T}$ -OBDDs and $\mathcal{T}$ -SDDs on top of OBDD and SDD packages and the MathSAT SMT solver. Some preliminary empirical evaluation supports the effectiveness of the approach.

1 Introduction

In the field of Knowledge Compilation (KC), the aim is to transform a given knowledge base, often represented as a Boolean formula, into a more suitable form that facilitates efficient query answering. This involves shifting the bulk of computational effort to the offline compilation phase, thereby optimizing the efficiency of the online query-answering phase. Many representations are subsets of Negation Normal Form (NNF), and in particular of decomposable, deterministic NNF (d-DNNF) [18]. Among these, decision diagrams (DDs) such as Ordered Binary Decision Diagrams (OBDDs) [6] and Sentential Decision Diagrams (SDDs) [17] are well-established and widely adopted representations in KC. They offer efficient querying and manipulation of Boolean functions and serve as foundational elements in numerous tools across various domains, including planning [27], probabilistic inference [5, 35], probabilistic reasoning [12, 21], and formal verification [8]. Central to KC is the notion of *canonicity*, where two equivalent Boolean formulas yield identical decision diagrams. Under specific conditions, both OBDDs and SDDs can achieve canonicity.

The literature on KC, decision diagrams, and canonicity for Boolean formulas is extensive. However, there is a notable scarcity

of literature addressing scenarios where formulas contain first-order logic theories such as difference logic ($\mathcal{DL}$), two variables per inequality ($\mathcal{TVPI}$), linear and non-linear arithmetic ($\mathcal{LRA}$ and $\mathcal{NLA}$), and equalities ($\mathcal{E}$) with uninterpreted functions ($\mathcal{EUF}$), which requires leveraging decision diagrams for Satisfiability Modulo Theories (SMT).

Related work. Most of the literature has focused on theory-aware OBDDs. The majority of the works are theory-specific, in particular focusing on $\mathcal{E}$ [23, 25, 26, 7], $\mathcal{EUF}$ [34, 1, 2], and fragments of arithmetic, such as $\mathcal{DL}$ [30], $\mathcal{TVPI}$ [10], and $\mathcal{NLA}$ [11]. Some general approaches have been proposed to support arbitrary theories [19, 22, 9, 13]. To the best of our knowledge, the only tentative to extend SDDs to support first-order theories are XSDDs [20, 28] which support $\mathcal{LRA}$.

From the practical point of view, most of the techniques are hard to implement since they require modifying the internals of some SMT solver or DD package, or both. Indeed, all of them, with the only exception of LDDs [10], do not have a public implementation, or are implemented within other tools, making them not directly usable and comparable to our approach. From the theoretical point of view, some techniques allow for theory-inconsistent paths (e.g., LDDs [10] and XSDDs [20, 28]), while others only guarantee theory-semicanonicity, i.e., they map all theory-valid and all theory-inconsistent formulas to the same DD (e.g., DDDs [30]). Notably, none of them has been proven to be theory-canonical.

An extensive and detailed analysis of all these techniques is available in the appendix.

Contributions. In this paper, we investigate the problem of leveraging Boolean decision diagrams (DDs) to SMT level ($\mathcal{T}$ -DDs). We present a general formal framework for $\mathcal{T}$ -DDs. Then, we introduce a novel and highly versatile technique for extending decision diagrams to the realm of SMT, which operates as follows: we perform a total enumeration of the truth assignments satisfying the input SMT formula (AllSMT) [29], extracting a set of theory lemmas, i.e. $\mathcal{T}$ -valid clauses, that rules out all $\mathcal{T}$ -inconsistent truth assignments. These $\mathcal{T}$ -lemmas are then conjoined to the original SMT problem, and its Boolean abstraction is fed to a Boolean DD compiler to generate a theory DD ($\mathcal{T}$ -DD). We formally establish how our proposed framework ensures the generation of $\mathcal{T}$ -canonical decision diagrams, provided the underlying Boolean decision diagram is canonical.

Our technique offers several advantages. Firstly, it is very easy to implement, relying on standard AllSMT solvers and existing DD packages as black boxes, with no need to put the hands inside the code of the AllSMT solver and of the DD package. This simplicity makes it accessible to a wide range of users, regardless of their expertise level in SMT solving and DD compiling. Additionally, our technique is theory-agnostic, accommodating any theory or combi-

nation thereof supported by the AllSMT solver, and DD-agnostic, since it potentially works with any form of DD. Remarkably, if the underlying DD is canonical, it produces theory-canonical $\mathcal{T}$ -DDs, ensuring that two $\mathcal{T}$ -equivalent formulas under the same set of theory atoms share the same $\mathcal{T}$ -DD. Also, it is the first implementation that can be used for #SMT [31], because the $\mathcal{T}$ -DD represents only theory-consistent truth assignments.

We have implemented a prototype of $\mathcal{T}$ -OBDD and $\mathcal{T}$ -SDD compiler based on our algorithm, using the MATHSAT AllSMT solver [16] along with state-of-the-art packages for OBDDs and SDDs. A preliminary empirical evaluation demonstrates the effectiveness of our approach in producing $\mathcal{T}$ -canonical $\mathcal{T}$ -OBDDs and $\mathcal{T}$ -SDDs for several theories.

2 Background

Notation & terminology. We assume the reader is familiar with the basic syntax, semantics, and results of propositional and first-order logics. We adopt the following terminology and notation.

Satisfiability Modulo Theories (SMT) extends SAT to the context of first-order formulas modulo some background theory $\mathcal{T}$, which provides an intended interpretation for constant, function, and predicate symbols. We restrict to quantifier-free formulas. A $\mathcal{T}$ -formula is a combination of theory-specific atoms ($\mathcal{T}$ -atoms) via Boolean connectives. For instance, in $\mathcal{LRA}$ $\mathcal{T}$ -atoms are linear (in)equalities over rational variables.

$\mathcal{T}2\mathcal{B}$ is a bijective function (“theory to Boolean”), called *Boolean (or propositional) abstraction*, which maps Boolean atoms into themselves, $\mathcal{T}$ -atoms into fresh Boolean variables, and is homomorphic w.r.t. Boolean operators and set inclusion. The function $\mathcal{B}2\mathcal{T}$ (“Boolean to theory”), called *refinement*, is the inverse of $\mathcal{T}2\mathcal{B}$. (For instance $\mathcal{T}2\mathcal{B}(\{(x - y \leq 3) \vee (x = z)\}) = \{(A_1 \vee A_2)\}$, A_1 and A_2 being fresh Boolean variables, and $\mathcal{B}2\mathcal{T}(\{\neg A_1, A_2\}) = \{\neg(x - y \leq 3), (x = z)\}$.)

The symbols $\alpha \stackrel{\text{def}}{=} \{\alpha_i\}_i, \beta \stackrel{\text{def}}{=} \{\beta_i\}_i$ denote ground $\mathcal{T}$ -atoms on $\mathcal{T}$ -variables $x \stackrel{\text{def}}{=} \{x_i\}_i$. The symbols $A \stackrel{\text{def}}{=} \{A_i\}_i, B \stackrel{\text{def}}{=} \{B_i\}_i$ denote Boolean atoms, and typically denote also the Boolean abstraction of the $\mathcal{T}$ -atoms in α, β respectively. (Notice that a Boolean atom is also a $\mathcal{T}$ -atom, which is mapped into itself by $\mathcal{T}2\mathcal{B}$.) We represent truth assignments as conjunctions of literals. We denote by 2^α the set of all total truth assignments on α . The symbols φ, ψ denote $\mathcal{T}$ -formulas, and μ, η, ρ denote conjunctions of $\mathcal{T}$ -literals; φ^p, ψ^p denote Boolean formulas, μ^p, η^p, ρ^p denote conjunctions of Boolean literals (i.e., truth assignments) and we use them as synonyms for the Boolean abstraction of φ, ψ, η , and ρ respectively, and vice versa (e.g., φ^p denotes $\mathcal{T}2\mathcal{B}(\varphi)$, η denotes $\mathcal{B}2\mathcal{T}(\eta^p)$). If $\mathcal{T}2\mathcal{B}(\eta) \models \mathcal{T}2\mathcal{B}(\varphi)$, then we say that η *propositionally (or tautologically) satisfies* φ , written $\eta \models_{\mathcal{B}} \varphi$. (Notice that if $\eta \models_{\mathcal{B}} \varphi$ then $\eta \models_{\mathcal{T}} \varphi$, but not vice versa.) The notion of propositional/tautological entailment and validity follow straightforwardly. When both $\varphi \models_{\mathcal{B}} \psi$ and $\psi \models_{\mathcal{B}} \varphi$, we say that φ and ψ are *propositionally/tautologically equivalent*, written “ $\varphi \equiv_{\mathcal{B}} \psi$ ”. When both $\varphi \models_{\mathcal{T}} \psi$ and $\psi \models_{\mathcal{T}} \varphi$, we say that φ and ψ are *$\mathcal{T}$ -equivalent*, written “ $\varphi \equiv_{\mathcal{T}} \psi$ ”. (Notice that if $\eta \models_{\mathcal{B}} \varphi$ then $\eta \equiv_{\mathcal{T}} \varphi$, but not vice versa.) We call a $\mathcal{T}$ -lemma any $\mathcal{T}$ -valid clause.

Decision Diagrams. Knowledge compilation is the process of transforming a formula into a representation that is more suitable for answering queries [18]. Many known representations are subsets of Negation Normal Form (NNF), which requires formulas to be represented by Directed Acyclic Graphs (DAGs) where internal nodes

are labelled with $\wedge$ or $\vee$, and leaves are labelled with literals $A, \neg A$, or constants $\top, \perp$. Other languages are defined as special cases of NNF [18]. In particular, Decision Diagrams (DDs) like OBDDs and SDDs are popular compilation languages.

Ordered Binary Decision Diagrams (OBDDs) [6] are NNFs where the root node is a decision node and a total order “ $<$ ” on the atoms is imposed. A decision node is either a constant $\top, \perp$, or a $\vee$ -node having the form $(A \wedge \varphi) \vee (\neg A \wedge \varphi')$, where A is an atom, and φ, φ' are decision nodes. In every path from the root to a leaf, each atom is tested only once, following the order “ $<$ ”. Figure 1 (left) shows a graphical representation of an OBDD, where each decision node is graphically represented as a node labelled with the atom A being tested; a solid and a dashed edge connect it to the nodes of φ and φ' , representing the cases in which A is true or false, respectively.¹ OBDDs allow performing many operations in polynomial time, such as performing Boolean combinations of OBDDs, or checking for (un)satisfiability or validity.

SDDs [17] are a generalization of OBDDs, in which decisions are not binary and are made on sentences instead of atoms. Formally, and SDD is an NNF that satisfies the properties of *structured decomposability* and *strong determinism*. A v-tree v for atoms A is a full binary tree whose leaves are in one-to-one correspondence with the atoms in A . We denote with v_l and v_r the left and right subtrees of v . A SDD that respects v is either: a constant $\top, \perp$; a literal $A, \neg A$ if v is a leaf labelled with A ; a decomposition $\bigvee_{i=1}^n (\varphi_i \wedge \psi_i)$ if v is internal, $\varphi_1, \dots, \varphi_n$ are SDDs that respect subtrees of $v_l, \psi_1, \dots, \psi_n$ are SDDs that respect subtrees of v_r , and $\varphi \stackrel{\text{def}}{=} \{\varphi_1, \dots, \varphi_n\}$ is a partition. φ is called a partition if each φ_i is consistent, every pair φ_i, φ_j for $i \neq j$ are mutually exclusive, and the disjunction of all φ_i s is valid. φ_i s are called *primes*, and ψ_i s are called *subs*. A pair φ_i, ψ_i is called an *element*. Figure 1 (right) shows a graphical representation of a SDD. Decomposition nodes are represented as circles, with an outgoing edge for each element. Elements are represented as paired boxes, where the left and right boxes represent the prime and the sub, respectively. SDDs maintain many of the properties of OBDDs, with the advantage of being exponentially more succinct [4].

These forms of DDs are *canonical modulo some canonicity condition* $\Lambda(A)$ on the Boolean atoms A : under the assumption that the DDs are built according to the same canonicity condition $\Lambda(A)$, then each formula $\varphi^p[A]$ has a unique DD representation $\text{DD}(\varphi^p[A])$, and $\text{DD}(\varphi^p[A]) = \text{DD}(\varphi'^p[A])$ if and only if $\varphi^p[A] \equiv \varphi'^p[A]$ and hence if and only if $\text{DD}(\varphi^p[A]) \equiv \text{DD}(\varphi'^p[A])$. (E.g., for OBDDs, $\Lambda(A)$ is given total order on A [6]; for SDDs $\Lambda(A)$ is the structure induced by a given v-tree [17].) Canonicity allows easily checking if a formula is a tautology or a contradiction, and if two formulas are equivalent. Also, canonicity allows storing equivalent subformulas only once.

3 A Formal Framework for $\mathcal{T}$ -DDs

In this section, we introduce the theoretical results that will be used in the rest of the paper. For the sake of compactness, all the proofs of the theorems are deferred to the appendix.

We denote by $\varphi[\alpha]$ the fact that α is a superset of the set of $\mathcal{T}$ -atoms occurring in φ whose truth assignments we are interested in. The fact that it is a superset is sometimes necessary for comparing formulas with different sets of $\mathcal{T}$ -atoms: $\varphi[\alpha]$ and $\varphi'[\alpha']$ can be compared only if they are both considered as formulas on $\alpha \cup \alpha'$. (E.g., in order to check that $(A_1 \vee A_2) \wedge (A_1 \vee \neg A_2)$ and $(A_1 \vee A_3) \wedge$

¹ In all examples we use OBDDs only because it is eye-catching to detect the partial assignments which verify or falsify the formula.




Figure 1: OBDD (left) and SDD (right) for $\varphi = (A_1 \wedge A_2) \vee (A_2 \wedge A_3) \vee (A_3 \wedge A_4)$.

$(A_1 \vee \neg A_3)$ are equivalent, we need considering them as formulas on $\{A_1, A_2, A_3\}$.)

Given a set of $\mathcal{T}$ -atoms α and a $\mathcal{T}$ -formula $\varphi[\alpha]$, we denote by $H_\alpha(\varphi) \stackrel{\text{def}}{=} \{\eta_i[\alpha]\}_i$ and $P_\alpha(\varphi) \stackrel{\text{def}}{=} \{\rho_j[\alpha]\}_j$ respectively the set of all $\mathcal{T}$ -consistent and that of all $\mathcal{T}$ -inconsistent total truth assignments on the set of $\mathcal{T}$ -atoms α which propositionally satisfy φ , i.e., s.t.

$$\varphi[\alpha] \equiv_{\mathbb{B}} \bigvee_{\eta_i[\alpha] \in H_\alpha(\varphi)} \eta_i[\alpha] \vee \bigvee_{\rho_j[\alpha] \in P_\alpha(\varphi)} \rho_j[\alpha]. \quad (1)$$

The following facts are straightforward consequences of the definition of $H_\alpha(\varphi)$ and $P_\alpha(\varphi)$.

Proposition 1. Given two $\mathcal{T}$ -formulas $\varphi[\alpha]$ and $\varphi'[\alpha]$, we have that:

- (a) $H_\alpha(\varphi)$, $P_\alpha(\varphi)$, $H_\alpha(\neg\varphi)$, $P_\alpha(\neg\varphi)$ are pairwise disjoint;
- (b) $H_\alpha(\varphi) \cup P_\alpha(\varphi) \cup H_\alpha(\neg\varphi) \cup P_\alpha(\neg\varphi) = 2^\alpha$;
- (c) $\varphi[\alpha] \equiv_{\mathbb{B}} \varphi'[\alpha]$ if and only if $H_\alpha(\varphi) = H_\alpha(\varphi')$ and $P_\alpha(\varphi) = P_\alpha(\varphi')$;
- (d) $\varphi[\alpha] \equiv_{\mathcal{T}} \varphi'[\alpha]$ if and only if $H_\alpha(\varphi) = H_\alpha(\varphi')$.

Example 2. Let $\alpha = \{\alpha_1, \alpha_2\} \stackrel{\text{def}}{=} \{(x \leq 0), (x = 1)\}$. Consider the formulas $\varphi_1[\alpha] \stackrel{\text{def}}{=} (x \leq 0) \vee (x = 1)$ and $\varphi_2[\alpha] \stackrel{\text{def}}{=} \neg(x \leq 0) \leftrightarrow (x = 1)$, so that $\varphi_1^p[\mathbf{A}] \stackrel{\text{def}}{=} A_1 \vee A_2$ and $\varphi_2^p[\mathbf{A}] \stackrel{\text{def}}{=} \neg A_1 \leftrightarrow A_2$. It is easy to see that $\varphi_1[\alpha] \not\equiv_{\mathbb{B}} \varphi_2[\alpha]$ and $\varphi_1[\alpha] \equiv_{\mathcal{T}} \varphi_2[\alpha]$. Then $H_\alpha(\varphi_1[\alpha]) = H_\alpha(\varphi_2[\alpha]) = \{\eta_1, \eta_2\} = \{(x \leq 0) \wedge \neg(x = 1), \neg(x \leq 0) \wedge (x = 1)\}$, whereas $P_\alpha(\varphi_1[\alpha]) = \{(x \leq 0) \wedge (x = 1)\}$ and $P_\alpha(\varphi_2[\alpha]) = \emptyset$.

3.1 Canonicity for $\text{SMT}(\mathcal{T})$ formulas

Definition 1. Given a set of $\mathcal{T}$ -atoms α and its Boolean abstraction $\mathbf{A} \stackrel{\text{def}}{=} \mathcal{T}2\mathcal{B}(\alpha)$, some $\mathcal{T}$ -formula $\varphi[\alpha]$, and some form of DDs with canonicity condition $\Lambda(\mathbf{A})$ (if any), we call “ $\mathcal{T}$ -DD($\varphi[\alpha]$)” with canonicity condition $\Lambda(\alpha)$ an $\text{SMT}(\mathcal{T})$ formula $\Psi[\alpha]$ such that $\Psi[\alpha] \equiv_{\mathcal{T}} \varphi[\alpha]$ and its Boolean abstraction $\Psi^p[\mathbf{A}]$ is a DD.

E.g., “ $\mathcal{T}$ -OBDDs” and “ $\mathcal{T}$ -SDDs” denote $\text{SMT}(\mathcal{T})$ extensions of OBDDs and SDDs respectively.

Theorem 3. Consider some form of $\mathcal{T}$ -DD such that its Boolean abstraction DD is canonical. Then $\mathcal{T}$ -DD($\varphi[\alpha]$) = $\mathcal{T}$ -DD($\varphi'[\alpha]$) if and only if $\mathcal{T}$ -DD($\varphi[\alpha]$) $\equiv_{\mathbb{B}}$ $\mathcal{T}$ -DD($\varphi'[\alpha]$).

There are potentially many possible ways by which DDs can be extended into $\mathcal{T}$ -DDs, depending mainly on how the $\mathcal{T}$ -consistency of branches and subformulas is handled. A straightforward way would




Figure 2: Top: OBDD for $A_1 \vee A_2$ and its refinement for $(x \leq 0) \vee (x = 1)$.

Bottom: OBDD for $\neg A_1 \leftrightarrow A_2$ and its refinement for $\neg(x \leq 0) \leftrightarrow (x = 1)$.

be to define it as the refinement of the DD of the Boolean abstraction, i.e. $\mathcal{T}$ -DD($\varphi[\alpha]$) $\stackrel{\text{def}}{=} \mathcal{B}2\mathcal{T}(\text{DD}(\varphi^p[\mathbf{A}]))$, without pruning $\mathcal{T}$ -inconsistent branches and subformulas. Such $\mathcal{T}$ -DDs, however, would be neither $\mathcal{T}$ -canonical nor $\mathcal{T}$ -semicanonical, as defined below.

Definition 2. Let α be set of $\mathcal{T}$ -atoms, and let $\Lambda(\alpha)$ be some canonicity condition. We say that a form of $\mathcal{T}$ -DD is $\mathbb{B}$ -canonical wrt. $\Lambda(\alpha)$ iff, for every pair of formulas $\varphi[\alpha]$ and $\varphi'[\alpha]$, $\mathcal{T}$ -DD($\varphi[\alpha]$) = $\mathcal{T}$ -DD($\varphi'[\alpha]$) if $\varphi[\alpha] \equiv_{\mathbb{B}} \varphi'[\alpha]$.

We say that a form of $\mathcal{T}$ -DD is $\mathcal{T}$ -canonical wrt. $\Lambda(\alpha)$ iff, for every pair of formulas $\varphi[\alpha]$ and $\varphi'[\alpha]$, $\mathcal{T}$ -DD($\varphi[\alpha]$) = $\mathcal{T}$ -DD($\varphi'[\alpha]$) if and only if $\varphi[\alpha] \equiv_{\mathcal{T}} \varphi'[\alpha]$.

We say that a form of $\mathcal{T}$ -DD is $\mathcal{T}$ -semicanonical wrt. $\Lambda(\alpha)$ iff, for every pair of formulas $\varphi[\alpha]$ and $\varphi'[\alpha]$, if $\varphi[\alpha]$ and $\varphi'[\alpha]$ are both $\mathcal{T}$ -inconsistent or are both $\mathcal{T}$ -valid, then $\mathcal{T}$ -DD($\varphi[\alpha]$) = $\mathcal{T}$ -DD($\varphi'[\alpha]$).

If $\mathcal{T}$ -DD is $\mathcal{T}$ -canonical, then it is also $\mathcal{T}$ -semicanonical, but not vice versa. As a consequence of Theorem 3, $\mathcal{T}$ -DD is $\mathbb{B}$ -canonical if its corresponding DD is canonical, but not vice versa.

Notice the “if” rather than “if and only if” in the definition of $\mathbb{B}$ -canonical: it may be the case that $\mathcal{T}$ -DD($\varphi[\alpha]$) = $\mathcal{T}$ -DD($\varphi'[\alpha]$) even if $\varphi[\alpha] \not\equiv_{\mathbb{B}} \varphi'[\alpha]$ (e.g., if $\varphi[\alpha] \equiv_{\mathcal{T}} \varphi'[\alpha]$, as in the case of $\mathcal{T}$ -canonicity).

Example 4. Let $\mathcal{T}$ -OBDD be defined as $\mathcal{T}$ -OBDD($\varphi_i[\alpha]$) $\stackrel{\text{def}}{=} \mathcal{B}2\mathcal{T}(\text{OBDD}(\varphi_i^p[\mathbf{A}]))$. Consider the formulas $\varphi_1[\alpha], \varphi_2[\alpha]$ in Example 2. Figure 2 shows the OBDDs for $\varphi_1^p[\mathbf{A}]$ and $\varphi_2^p[\mathbf{A}]$ (left) –considering as canonicity condition $\Lambda(\mathbf{A})$ the order $\{A_1, A_2\}$ – and the $\mathcal{T}$ -OBDDs for $\varphi_1[\alpha]$ and $\varphi_2[\alpha]$ (right). Notice that the two $\mathcal{T}$ -OBDDs are different despite the fact that $\varphi_1[\alpha] \equiv_{\mathcal{T}} \varphi_2[\alpha]$. Thus this form of $\mathcal{T}$ -OBDDs is not $\mathcal{T}$ -canonical.

Consider the $\mathcal{T}$ -valid $\mathcal{T}$ -formulas $\varphi_3[\alpha] \stackrel{\text{def}}{=} ((x \leq 0) \vee \neg(x \leq 0)) \wedge ((x = 1) \vee \neg(x = 1))$ and $\varphi_4[\alpha] \stackrel{\text{def}}{=} \neg(x \leq 0) \vee \neg(x = 1)$, so that $\varphi_3^p[\mathbf{A}] = (A_1 \vee \neg A_1) \wedge (A_2 \vee \neg A_2)$ and $\varphi_4^p[\mathbf{A}] = \neg A_1 \vee \neg A_2$. Since $\varphi_3^p[\mathbf{A}]$ is propositionally valid whereas $\varphi_4^p[\mathbf{A}]$ is not, then $\mathcal{T}$ -OBDD($\varphi_3[\alpha]$) reduces to the $\top$ node whereas $\mathcal{T}$ -OBDD($\varphi_4[\alpha]$) does not. Dually, $\neg\varphi_3[\alpha]$ and $\neg\varphi_4[\alpha]$ are both $\mathcal{T}$ -inconsistent, and $\mathcal{T}$ -OBDD($\neg\varphi_3[\alpha]$) reduces to the $\perp$ node whereas $\mathcal{T}$ -OBDD($\neg\varphi_4[\alpha]$) does not. Thus this form of $\mathcal{T}$ -OBDDs is not $\mathcal{T}$ -semicanonical.

Theorem 5. Consider a form of $\mathcal{T}$ -DDs which are $\mathbb{B}$ -canonical wrt. some canonicity condition $\Lambda(\alpha)$. Suppose that, for every $\text{SMT}(\mathcal{T})$ formula $\varphi[\alpha]$, $\mathcal{T}$ -DD($\varphi[\alpha]$) $\equiv_{\mathbb{B}} \bigvee_{\eta_i \in H_\alpha(\varphi)} \eta_i$. Then $\mathcal{T}$ -DD are $\mathcal{T}$ -canonical wrt. $\Lambda(\alpha)$.

Theorem 5 states a sufficient condition to guarantee the $\mathcal{T}$ -canonicity of some form of $\mathcal{T}$ -DD: it should represent all and only $\mathcal{T}$ -consistent total truth assignments propositionally satisfying the formula. Since typically $\mathcal{T}$ -DDs represent *partial* assignments μ_i , the latter ones should not have $\mathcal{T}$ -inconsistent total extensions.

3.2 Canonicity via $\mathcal{T}$ -lemmas

Definition 3. We say that a set $\{C_1[\alpha], \dots, C_K[\alpha]\}$ of $\mathcal{T}$ -lemmas **rules out** a set $\{\rho_1[\alpha], \dots, \rho_M[\alpha]\}$ of $\mathcal{T}$ -inconsistent total truth assignments if and only if, for every $\rho_j[\alpha]$ in the set, there exists a $C_l[\alpha]$ s.t. $\rho_j[\alpha] \models_{\mathbb{B}} \neg C_l[\alpha]$, that is, if and only if $\bigvee_{j=1}^M \rho_j[\alpha] \wedge \bigwedge_{l=1}^K C_l[\alpha] \equiv_{\mathbb{B}} \perp$.

Given α and some $\mathcal{T}$ -formula $\varphi[\alpha]$, we denote as $Cl_{\alpha}(\varphi)$ any function which returns a set of $\mathcal{T}$ -lemmas $\{C_1[\alpha], \dots, C_K[\alpha]\}$ which rules out $P_{\alpha}(\varphi)$.

Theorem 6. Let $\varphi[\alpha]$ be a $\mathcal{T}$ -formula. Let $Cl_{\alpha}(\varphi) \stackrel{\text{def}}{=} \{C_1[\alpha], \dots, C_K[\alpha]\}$ be a set of $\mathcal{T}$ -lemmas which rules out $P_{\alpha}(\varphi)$. Then we have that:

$$\varphi^p[A] \wedge \bigwedge_{C_l[\alpha] \in Cl_{\alpha}(\varphi)} C_l^p[A] \equiv \bigvee_{\eta_i[\alpha] \in H_{\alpha}(\varphi)} \eta_i^p[A]. \quad (2)$$

Theorem 7. Let α denote generical sets of $\mathcal{T}$ -atoms with Boolean abstraction A . Consider some canonical form of DDs on some canonicity condition $\Lambda(A)$. Let $\mathcal{T}$ -DD with canonicity condition $\Lambda(\alpha)$ be such that, for all sets α and for all formulas $\varphi[\alpha]$:

$$\mathcal{T}\text{-DD}(\varphi[\alpha]) \stackrel{\text{def}}{=} \text{B2T}(\text{DD} \left(\varphi^p[A] \wedge \bigwedge_{C_l[\alpha] \in Cl_{\alpha}(\varphi)} C_l^p[A] \right)). \quad (3)$$

Then the $\mathcal{T}$ -DDs are $\mathcal{T}$ -canonical.

Theorem 7 suggests an easy way to implement $\mathcal{T}$ -canonical $\mathcal{T}$ -DDs by using as $Cl_{\alpha}(\varphi)$ the list of $\mathcal{T}$ -lemmas produced by an $\text{SMT}(\mathcal{T})$ solver during an AllSMT run over $\varphi[\alpha]$. (We will discuss this technique in §4.)

3.3 Dealing with extra $\mathcal{T}$ -atoms

Unfortunately, things are not so simple in practice. In order to cope with some theories, AllSMT solvers frequently need introducing extra $\mathcal{T}$ -atoms β on-the-fly, and need generating a set of some extra $\mathcal{T}$ -lemmas $\text{Defs}_{\alpha, \beta}(\varphi) \stackrel{\text{def}}{=} \{C_l[\alpha, \beta]\}_l$ relating the novel atoms β with those occurring in the original formula $\varphi[\alpha]$ [3]. Consequently, in this case the list of $\mathcal{T}$ -lemmas produced by an $\text{SMT}(\mathcal{T})$ solver during an AllSMT run over $\varphi[\alpha]$ may contain some of such $\mathcal{T}$ -lemmas, and thus they cannot be used as $Cl_{\alpha}(\varphi)$.

For instance, when the $\mathcal{LRA}$ -atom $(\sum_i a_i x_i = b)$ occurs in a $\mathcal{LRA}$ -formula, the SMT solver may need introducing also the $\mathcal{LRA}$ -atoms $(\sum_i a_i x_i \geq b)$ and $(\sum_i a_i x_i \leq b)$ and adding some or all the $\mathcal{LRA}$ -lemmas encoding $(\sum_i a_i x_i = b) \leftrightarrow ((\sum_i a_i x_i \geq b) \wedge (\sum_i a_i x_i \leq b))$.

Definition 4. We say that a set $\{C_1[\alpha, \beta], \dots, C_K[\alpha, \beta]\}$ of $\mathcal{T}$ -lemmas on α, β **rules out** a set $\{\rho_1[\alpha], \dots, \rho_M[\alpha]\}$ of $\mathcal{T}$ -inconsistent total truth assignments on α if and only if

$$\bigvee_{j=1}^M \rho_j[\alpha] \wedge \bigwedge_{l=1}^K C_l[\alpha, \beta] \equiv_{\mathbb{B}} \perp. \quad (4)$$




Figure 3: OBDD for $\varphi^p[A] \wedge \bigwedge_{i=1}^5 C_i^p[A, B]$ and its refinement for $\varphi[\alpha] \wedge \bigwedge_{i=1}^5 C_i[\alpha, \beta]$.

Given α, β and some $\mathcal{T}$ -formula $\varphi[\alpha]$, we denote as $Cl_{\alpha, \beta}(\varphi)$ any function which returns a set of $\mathcal{T}$ -lemmas $\{C_1[\alpha, \beta], \dots, C_K[\alpha, \beta]\}$ on α, β which rules out $P_{\alpha}(\varphi)$. (If $\beta = \emptyset$, then Definition 4 reduces to Definition 3 and $Cl_{\alpha, \beta}(\varphi)$ reduces to $Cl_{\alpha}(\varphi)$.) Notice that $Cl_{\alpha, \beta}(\varphi)$ is not unique, and it is not necessarily minimal: if $C[\alpha, \beta] \notin Cl_{\alpha, \beta}(\varphi)$ is a $\mathcal{T}$ -lemma, then $\{C[\alpha, \beta]\} \cup Cl_{\alpha, \beta}(\varphi)$ rules out $P_{\alpha}(\varphi)$ as well. The same fact applies to $Cl_{\alpha}(\varphi)$ as well.

Theorem 8. Let α and β be sets of $\mathcal{T}$ -atoms and let A and B denote their Boolean abstraction. Let $\varphi[\alpha]$ be a $\mathcal{T}$ -formula. Let $Cl_{\alpha, \beta}(\varphi) \stackrel{\text{def}}{=} \{C_1[\alpha, \beta], \dots, C_K[\alpha, \beta]\}$ be a set of $\mathcal{T}$ -lemmas on α, β which rules out $P_{\alpha}(\varphi)$. Then we have that:

$$\varphi^p[A] \wedge \exists B. \bigwedge_{C_l[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B] \equiv \bigvee_{\eta_i[\alpha] \in H_{\alpha}(\varphi)} \eta_i^p[A]. \quad (5)$$

Example 9. Consider the formula $\varphi_1[\alpha] \stackrel{\text{def}}{=} (x \leq 0) \vee (x = 1)$ and its Boolean abstraction $\varphi_1^p[A] \stackrel{\text{def}}{=} A_1 \vee A_2$, as in Example 2. If we run an AllSMT solver (e.g. MATHSAT) over it we obtain the set $\{\eta_1, \eta_2\} \stackrel{\text{def}}{=} \{(x \leq 0) \wedge \neg(x = 1), \neg(x \leq 0) \wedge (x = 1)\}$ but, instead of the $\mathcal{T}$ -lemma $\neg(x \leq 0) \vee \neg(x = 1)$, we might obtain instead five $\mathcal{T}$ -lemmas:

$$\begin{aligned} C_1 &: \neg(x \leq 0) \vee \neg(x \geq 1) \\ C_2 &: \neg(x \leq 0) \vee (x \leq 1) \\ C_3 &: \neg(x = 1) \vee (x \leq 1) \\ C_4 &: \neg(x = 1) \vee (x \geq 1) \\ C_5 &: (x = 1) \vee \neg(x \leq 1) \vee \neg(x \geq 1), \end{aligned}$$


because the SMT solver has introduced the extra atoms $\{\beta_1, \beta_2\} \stackrel{\text{def}}{=} \{(x \leq 1), (x \geq 1)\}$ and added the axiom $(x = 1) \leftrightarrow ((x \leq 1) \wedge (x \geq 1))$, which is returned in the list of the $\mathcal{T}$ -lemmas as C_3, C_4, C_5 .

Now, if we applied the OBDD construction simply to $\varphi^p[A] \wedge \bigwedge_{i=1}^5 C_i^p[A, B]$ and $\varphi[\alpha] \wedge \bigwedge_{i=1}^5 C_i[\alpha, \beta]$ respectively, we would obtain the OBDDs in Figure 3, which are much bigger than necessary. Instead, computing the Boolean abstraction and applying the existential quantification on B by Shannon's expansion, we obtain:

$$\varphi^p \wedge \exists B_1 B_2. \bigwedge_{l=1}^5 C_l^p \equiv \overbrace{A_1 \wedge \neg A_2}^{\eta_1^p[A]} \vee \overbrace{(\neg A_1 \wedge A_2)}^{\eta_2^p[A]}, \quad (6)$$

in line with (5) in Theorem 8. The resulting OBDD is that of Figure 2, bottom left.

Theorem 10. Let α, β and β' be sets of $\mathcal{T}$ -atoms and let A, B and B' denote their Boolean abstraction. Let $\varphi[\alpha]$ and $\varphi'[\alpha]$ be $\mathcal{T}$ -formulas. Let $Cl_{\alpha, \beta}(\varphi)$ be a set of $\mathcal{T}$ -lemmas on α, β which rules out $P_{\alpha}(\varphi)$ and $Cl_{\alpha, \beta'}(\varphi')$ be a set of $\mathcal{T}$ -lemmas on α, β' which



```

 $\mathcal{T}$ -DD  $\mathcal{DD\_Compiler}(\mathcal{T}\text{-formula } \varphi[\alpha]) \{$ 
  if (AllSMT_Solver( $\varphi[\alpha]$ ) == UNSAT)
    then return DD-False;
  //  $\{C_1, \dots, C_K\}$  are the  $\mathcal{T}$ -lemmas stored by AllSMT_Solver
   $\mathcal{DD}^p = \mathcal{DD\_Compiler}(\exists \mathbf{B}. \mathcal{T}2\mathcal{B}(\{\varphi \wedge \bigwedge_{i=1}^K C_i\}));$ 
   $\mathcal{DD} = \mathcal{B}2\mathcal{T}(\mathcal{DD}^p);$ 
  return  $\mathcal{DD}$ ;
}

```

Figure 4: Schema of the $\mathcal{T}$ -knowledge compiler: architecture (above) and algorithm (below).

rules out $P_\alpha(\varphi')$. Then $\varphi[\alpha] \equiv_{\mathcal{T}} \varphi'[\alpha]$ if and only if

$$\varphi^p[A] \wedge \exists \mathbf{B}. \bigwedge_{C_i[\alpha, \beta] \in \mathcal{C}l_{\alpha, \beta}(\varphi)} C_i^p[A, B] \equiv \varphi'^p[A] \wedge \exists \mathbf{B}'. \bigwedge_{C_i'[\alpha, \beta'] \in \mathcal{C}l_{\alpha, \beta'}(\varphi')} C_i'^p[A, B']. \quad (7)$$

As a direct consequence of Theorem 10, we have the following fact.

Theorem 11. Let α, β denote generical sets of $\mathcal{T}$ -atoms with Boolean abstraction $\mathbf{A}, \mathbf{B}$ respectively. Consider some canonical form of DDs on some canonicity condition $\Lambda(\mathbf{A})$. Let $\mathcal{T}$ -DD with canonicity condition $\Lambda(\alpha)$ be such that, for all sets α, β and for all formulas $\varphi[\alpha]$:

$$\mathcal{T}\text{-DD}(\varphi[\alpha]) \stackrel{\text{def}}{=} \mathcal{B}2\mathcal{T}(\mathcal{DD}(\varphi^p[A] \wedge \exists \mathbf{B}. \bigwedge_{C_i[\alpha, \beta] \in \mathcal{C}l_{\alpha, \beta}(\varphi)} C_i^p[A, B])). \quad (8)$$

Then the $\mathcal{T}$ -DDs are $\mathcal{T}$ -canonical.

4 Building Canonical $\mathcal{T}$ -DDs

Given some background theory $\mathcal{T}$ and given some form of knowledge compiler for Boolean formulas into some form of Boolean decision diagrams (e.g., OBDDs, SDDs, ...), Theorem 11, suggests us an easy way to implement a compiler of an SMT formula into a $\mathcal{T}$ -canonical $\mathcal{T}$ -DD.

4.1 General ideas.

The procedure is reported in Figure 4. The input $\mathcal{T}$ -formula $\varphi[\alpha]$ is first fed to an AllSMT solver which, overall, enumerates the set of $\mathcal{T}$ -satisfiable total assignments $H_\alpha(\varphi[\alpha]) \stackrel{\text{def}}{=} \{\eta_1[\alpha], \dots, \eta_N[\alpha]\}$ propositionally satisfying $\varphi[\alpha]$. To do that, the AllSMT solver has to produce a set of $\mathcal{T}$ -lemmas $\mathcal{C}l_{\alpha, \beta}(\varphi[\alpha]) \stackrel{\text{def}}{=} \{C_1[\alpha, \beta], \dots, C_K[\alpha, \beta]\}$ to rule out all the $\mathcal{T}$ -unsatisfiable assignments in $P_\alpha(\varphi[\alpha])$ s. SMT solvers like MATHSAT can produce these $\mathcal{T}$ -lemmas as output.

We ignore the η_i s and we conjoin the $\mathcal{T}$ -lemmas to $\varphi[\alpha]$.² We apply the Boolean abstraction and feed $\exists \mathbf{B}. \mathcal{T}2\mathcal{B}(\{\varphi \wedge \bigwedge_{i=1}^K C_i\})$ to a

² In principle, we could feed the DD package directly the disjunction of all assignments η in $H_\alpha(\varphi[\alpha])$. In practice, this would be extremely inefficient, since the η s are all total assignments and there are a large amount of them. Rather, the $\mathcal{T}$ -lemmas typically involve only a very small subset of $\mathcal{T}$ -atoms, and thus each $\mathcal{T}$ -lemma of length k rules out up to $2^{|\alpha|-k}$ $\mathcal{T}$ -inconsistent total assignments in $P_\alpha(\varphi[\alpha])$.




Figure 5: Top left: OBDD of $\varphi^p \stackrel{\text{def}}{=} A_1 \wedge (A_2 \vee A_3)$; Top right: refinement of the OBDD of $\varphi \stackrel{\text{def}}{=} (x \leq 0) \wedge ((x \geq 1) \vee (x \leq 2))$; Bottom left: OBDD of $\varphi^p \wedge C_1^p \stackrel{\text{def}}{=} (A_1 \wedge \neg A_2 \wedge A_3)$; Bottom right: refinement of the OBDD of $\varphi^p \wedge C_1^p \stackrel{\text{def}}{=} \tau \wedge ((x \leq 0) \wedge \neg(x \geq 1) \wedge (x \leq 2))$.

Boolean DD-Compiler, which returns one decision diagram $\mathcal{DD}^p$, which is equivalent to $\exists \mathbf{B}. \mathcal{T}2\mathcal{B}(\{\varphi \wedge \bigwedge_{i=1}^K C_i\})$ in the Boolean space. $\mathcal{DD}^p$ is then mapped back into a $\mathcal{T}$ -DD $\mathcal{DD}$ via $\mathcal{B}2\mathcal{T}$.

Example 12. Let $\alpha \stackrel{\text{def}}{=} \{(x \leq 0), (x \geq 1), (x \leq 2)\}$. Consider the $\mathcal{LRA}$ -formula $\varphi \stackrel{\text{def}}{=} (x \leq 0) \wedge ((x \geq 1) \vee (x \leq 2))$, and its Boolean abstraction $\varphi^p \stackrel{\text{def}}{=} A_1 \wedge (A_2 \vee A_3)$. Assume the order $\{(x \leq 0), (x \geq 1), (x \leq 2)\}$. The OBDD of φ^p and its refinement are reported in Figure 5, top left and right. Notice that the latter has one $\mathcal{T}$ -inconsistent branch $\{(x \leq 0), (x \geq 1)\}$ (in red).

The AllSMT solver can enumerate the satisfying assignments:

$$\begin{aligned} \rho_1 &\stackrel{\text{def}}{=} (x \leq 0) \wedge (x \geq 1) \wedge (x \leq 2) \\ \rho_2 &\stackrel{\text{def}}{=} (x \leq 0) \wedge (x \geq 1) \wedge \neg(x \leq 2) \\ \eta_1 &\stackrel{\text{def}}{=} (x \leq 0) \wedge \neg(x \geq 1) \wedge (x \leq 2) \end{aligned}$$

causing the generation of the following $\mathcal{T}$ -lemma to rule out ρ_1, ρ_2 : $C_1 \stackrel{\text{def}}{=} \neg(x \leq 0) \vee \neg(x \geq 1)$, whose Boolean abstraction is $C_1^p \stackrel{\text{def}}{=} \neg A_1 \vee \neg A_2$. (Since $\beta = \emptyset$ here, there is no need to existentially quantify $\mathbf{B}$.)

Passing $\varphi^p \wedge C_1^p \stackrel{\text{def}}{=} (A_1 \wedge \neg A_2 \wedge A_3)$ to an OBDD compiler, the OBDD returned is the one in Figure 5, bottom left, corresponding to the $\mathcal{T}$ -OBDD on bottom right. Notice that the $\mathcal{T}$ -inconsistent branch has been removed, and that there is no $\mathcal{T}$ -inconsistent branch left. $\diamond$

Remark 1. We stress the fact that our approach is not a form of eager SMT encoding for DD construction. The latter consists in enumerating a priori all possible $\mathcal{T}$ -lemmas which can be constructed on top of the $\mathcal{T}$ -atom set α , regardless the formula $\varphi[\alpha]$ [3]. With the exception of very simple theories like $\mathcal{E}$ or $\mathcal{EUF}$, this causes a huge amount of $\mathcal{T}$ -lemmas. With our approach, which is inspired instead to the “lemma-lifting” approach for SMT unsat-core extraction [14] and MaxSMT [15], only the $\mathcal{T}$ -lemmas which are needed to rule out the $\mathcal{T}$ -inconsistent truth assignments in $P_\alpha(\varphi)$, which are generated on demand by the AllSMT solver.

4.2 About canonicity

A big advantage of our approach is that Theorem 10 guarantees that, given an ordered set of atoms α , we produce canonical $\mathcal{T}$ -OBDDs. The importance of $\mathcal{T}$ -canonicity is shown in the following example.




Figure 6: OBDDs for $(\neg A_0 \vee (A_1 \vee A_2)) \wedge (A_0 \vee (A_1 \oplus A_2))$ and $\mathcal{T}$ -OBDD for $(\neg(y \leq 0) \vee ((x \leq 0) \vee (x = 1))) \wedge ((y \leq 0) \vee ((x \leq 0) \oplus (x = 1)))$.

Example 13. Consider $\Lambda(\alpha) \stackrel{\text{def}}{=} \{(y \leq 0), (x \leq 0), (x = 1)\}$ and let $\varphi \stackrel{\text{def}}{=} (\neg(y \leq 0) \vee \varphi_1) \wedge ((y \leq 0) \vee \varphi_2)$, which contains $\varphi_1 \stackrel{\text{def}}{=} ((x \leq 0) \vee (x = 1))$ and $\varphi_2 \stackrel{\text{def}}{=} ((x \leq 0) \oplus (x = 1))$ of Example 2, which are $\mathcal{T}$ -equivalent, but which may not be recognized as such by previous forms of $\mathcal{T}$ -OBDDs. If this is the case, the final $\mathcal{T}$ -OBDD(φ) contains $\mathcal{T}$ -OBDD(φ_1) and $\mathcal{T}$ -OBDD(φ_2), without realizing they are $\mathcal{T}$ -equivalent. For example, if $\mathcal{T}$ -OBDD(φ_1) and $\mathcal{T}$ -OBDD(φ_2) are these in Figure 2 top right and bottom right respectively, then $\mathcal{T}$ -OBDD(φ) is represented in Figure 6 right.

With our approach, instead, since $\varphi \equiv_{\mathcal{T}} \varphi_1 \equiv_{\mathcal{T}} \varphi_2$ and thanks to Theorem 10, we produce the OBDD and $\mathcal{T}$ -OBDD of Figure 2, bottom left and right. In fact, φ is $\mathcal{T}$ -equivalent to φ_1 and to φ_2 , because it is in the form $(\neg\beta_1 \vee \varphi_1) \wedge (\beta_1 \vee \varphi_2)$ where $\varphi_1 \equiv_{\mathcal{T}} \varphi_2$.

Remark 2. The notion of $\mathcal{T}$ -canonicity, like that of Boolean canonicity, assumes that the formulas $\varphi[\alpha]$ and $\varphi'[\alpha]$ are compared on the same (super)set of $\mathcal{T}$ -atoms α . Therefore, in order to compare two formulas on different atom sets, $\varphi[\alpha, \beta]$ and $\varphi'[\alpha, \beta']$, we need to consider them as formulas on the union of atoms sets: $\varphi[\alpha, \beta, \beta']$ and $\varphi'[\alpha, \beta, \beta']$. If so, our technique produces also the necessary $\mathcal{T}$ -lemmas which relate α, β, β' .

Example 14. In order to compare $\psi_1 \stackrel{\text{def}}{=} (x = 0) \wedge (y = 1)$ and $\psi_2 \stackrel{\text{def}}{=} (x = 0) \wedge (y = x + 1)$, we need to consider them on the $\mathcal{T}$ -atoms $\{(x = 0), (y = 1), (y = x + 1)\}$. If so, with our procedure the AllSMT solver produces the $\mathcal{T}$ -lemmas $C_1 \stackrel{\text{def}}{=} \neg(x = 0) \vee \neg(y = 1) \vee (y = x + 1)$ for ψ_1 and $C_2 \stackrel{\text{def}}{=} \neg(x = 0) \vee \neg(y = x + 1) \vee (y = 1)$ for ψ_2 , and builds the DDs of $\psi_1^p \wedge C_1^p$ and $\psi_2^p \wedge C_2^p$ which are equivalent, so that the two $\mathcal{T}$ -DDs are identical, and are identical to that of $(x = 0) \wedge (y = 1) \wedge (y = x + 1)$.

Ensuring $\mathcal{T}$ -canonicity is not straightforward. E.g., DDDs [30] and LDDs [10] are not $\mathcal{T}$ -canonical: e.g., in the example in Figure 7 both produce different $\mathcal{T}$ -DDs for two $\mathcal{T}$ -equivalent formulas. Additionally, LDDs are not even $\mathcal{T}$ -semicanonical: in Figure 8 we show an example where LDDs' theory-specific simplifications fail to reduce a $\mathcal{T}$ -inconsistent formula to the node $\perp$.

5 A Preliminary Empirical Evaluation

To test the feasibility of our approach, we developed a prototype of the $\mathcal{T}$ -DD generator described in the algorithm in §4. Our tool, coded in Python using PySMT [24] for parsing and manipulating formulas, leverages: (i) CUDD [32] for OBDD generation; (ii) SDD [17] for SDD generation; (iii) MATHSAT5 [16] for AllSMT enumeration and theory lemma generation. The source code of the algorithms is available at <https://github.com/MaxMicheluttiUnitn/TheoryConsistentDecisionDiagrams>, and that of the experiments at <https://github.com/MaxMicheluttiUnitn/DecisionDiagrams>.




Figure 7: Let $\Lambda(\alpha) \stackrel{\text{def}}{=} \{(x \leq y), (x \leq z), (y \leq z)\}$. The DDDs/LDDs for the $\mathcal{T}$ -formulas $\phi_1 = (x \leq y) \leftrightarrow (y \leq z)$ and $\phi_2 = \phi_1 \wedge (\neg(x \leq y) \vee (x \leq z) \vee \neg(y \leq z))$ are reported above. (On these formulas, the output of DDD and LDD is the same). Notice that $(\neg(x \leq y) \vee (x \leq z) \vee \neg(y \leq z))$ is $\mathcal{T}$ -valid so that $\phi_1 \equiv_{\mathcal{T}} \phi_2$, but the diagrams are different.




Figure 8: LDD for $\phi_3 \stackrel{\text{def}}{=} (x - z \leq -3) \vee (y - x \leq 2) \vee (z - y \leq -1)$. Notice that ϕ_3 is $\mathcal{T}$ -unsatisfiable, yet the LDD is not the $\perp$ node.

5.1 Comparison with other tools

As reported in §1, the existing toolsets in the field are very limited. Indeed, for $\mathcal{T}$ -OBDDs only LDD [10] have a public and directly usable implementation. (See the analysis of tools in the appendix.) Moreover, LDD's implementation is confined to $\mathcal{TVPI}$ over real or integer variables. For $\mathcal{T}$ -SDDs, the implementation of XSDD is intricately tailored for Weighted Model Integration problems, making the extraction of $\mathcal{T}$ -SDD from its code non-trivial.

Thus, we conducted a comparative analysis of our tool against the following tools: (i) *Abstract OBDD*, baseline $\mathcal{T}$ -OBDD obtained from the refinement of the OBDD of the Boolean abstraction built with CUDD; (ii) *Abstract SDD*, baseline $\mathcal{T}$ -SDD obtained from the refinement of the SDD of the Boolean abstraction built with SDD. We remark that, as indicated in [28], the XSDD construction aligns with *Abstract SDD*; (iii) *LDD* from [10]. To ensure fairness in comparing the algorithms, uniformity in variable ordering (for OBDDs) and v-tree (for SDDs) across the tools is assumed. Notice that none of our “competitors” here is $\mathcal{T}$ -canonical, nor even $\mathcal{T}$ -semicanonical.

5.2 Benchmark

Due to the limited literature on $\mathcal{T}$ -DDs, there is a scarcity of benchmarks available. As a first step, we tested our tool on a subset of SMT-LIB benchmark problems. The main issue is that SMT-LIB problems are thought for SMT solving, and not for knowledge compilation. As a result, most of the problems are UNSAT or too difficult to compile into a $\mathcal{T}$ -DD in a feasible amount of time by any tool. Hence, we generated problems inspired by the Weighted Model Integration application, drawing inspiration from [28]. In this context, we consulted recent papers on the topic [33] and crafted a set of synthetic benchmarks accordingly. We set the weight function to 1 to prioritize the generation of the $\mathcal{T}$ -DD of the support formula, and adjusted the generation code to align with theories supported by the competitors (i.e., $\mathcal{TVPI}$ for LDD and $\mathcal{LRA}$ for XSDD).




Figure 9: Results obtained on synthetic $\mathcal{LRA}$ benchmarks (250 problems), comparing number of nodes (left) and computational times (right). Timeouts on the horizontal and vertical lines. $\mathcal{T}$ -OBDD timeouts: 79. $\mathcal{T}$ -SDD timeouts: 111.

5.3 Results

Figures 9 and 10 show the comparison of our algorithm (x -axis) against all baseline solvers (y -axis). The results are shown through scatter plots, comparing the size of the generated $\mathcal{T}$ -DDs and the taken computational time. We set the timeout to 3600s for AllSMT computation, and additional 3600s for $\mathcal{T}$ -DD generation. Notice that both axes are log-scaled. On the one hand, the plots show that our algorithms have longer computational times compared to the other tools. This outcome is not surprising, given the additional overhead associated with enumerating the lemmas via AllSMT and performing Boolean existential quantification. On the other hand, our tools generate smaller $\mathcal{T}$ -DDs, which is particularly noticeable for $\mathcal{T}$ -SDDs.

Our tools offer several distinctive advantages that set them apart within the field. Notably, these advantages may not be readily discernible from scatter plots or other visualization methods.

The $\mathcal{T}$ -DDs built with our approach ensure that every extension of a partial assignment leading to the $\top$ node represents a $\mathcal{T}$ -consistent total assignment. Consequently, our algorithm stands as the sole contender capable of performing $\#SMT$, aligning with the definition of $\#SMT$ proposed in [31]. This characteristic holds substantial implications for various applications, particularly in fields like Quantitative Information Flow, where precise enumeration is crucial.

Furthermore, benchmarks from the SMT-LIB, predominantly comprising UNSAT instances, proved our capability to identify $\mathcal{T}$ -inconsistent formulas and condense them into a single $\perp$ node. In contrast, LDD do not generate a $\perp$ $\mathcal{T}$ -DD for these formulas, highlighting once again their lack in achieving $\mathcal{T}$ -semicanonicity.

Finally, our algorithm supports the combination of theories and addresses theories not supported by other available implementations. In the tool repository, we provide a collection of problems spanning various theories, all of which are compatible with our implementation. Notably, our tool is the only one capable of generating theory decision diagrams for these problem domains.




Figure 10: Results obtained on synthetic $\mathcal{TVPI}$ benchmarks (200 problems), comparing number of nodes (left) and computational times (right). Timeouts on the horizontal and vertical lines. LDD timeouts: 0. $\mathcal{T}$ -OBDD timeouts: 22. $\mathcal{T}$ -SDD timeouts: 81.

6 Conclusions and Future Work

In this paper, we have investigated the problem of leveraging Boolean decision diagrams (DDs) to SMT level ($\mathcal{T}$ -DDs). We have presented a general theory-agnostic and DD-agnostic formal framework for $\mathcal{T}$ -DDs. We have shown a straightforward way to leverage DDs to $\mathcal{T}$ -DDs by simply combining an AllSMT solver and a DD package, both used as black boxes. This approach requires little effort to implement, since it does not require to modify the code of the AllSMT solver and of the DD package, and is very general, since it can be applied to any theory supported by the AllSMT solver and combinations thereof, and to any DD with a compiler admitting Boolean existential quantification. Importantly, this technique has a fundamental feature: it allows leveraging canonical DDs into $\mathcal{T}$ -canonical $\mathcal{T}$ -DDs. To the best of our knowledge, this is the first case of provably canonical $\mathcal{T}$ -DDs in the literature. We have implemented our approach on top of the MathSAT AllSMT solver and of both OBDD and SDD packages, and shown empirically its effectiveness.

This work opens several research avenues, and will progress along several directions. From a theoretical viewpoint, we are going to investigate how $\mathcal{T}$ -DDs can be effectively composed and how querying can be performed; also, we plan to extend our analysis to other forms of DDs, and on NNF formulas in general (in particular d-DNNF), investigating how their properties can be preserved by leveraging to SMT level. From a practical viewpoint, our approach currently suf-

fers from two main bottlenecks: (a) the need to perform AllSMT upfront and (b) the need to perform Boolean existential quantification to remove the extra $\mathcal{T}$ -atoms. For the former, we plan to investigate alternative and less-expensive ways to enumerate $\mathcal{T}$ -lemmas ruling out $\mathcal{T}$ -inconsistent assignments. For the latter, we plan to investigate alternative SMT techniques which reduce or even eliminate the presence of novel $\mathcal{T}$ -atoms in the $\mathcal{T}$ -lemmas. From an application viewpoint, we plan to use our $\mathcal{T}$ -SDDs package for Weighted Model Integration (WMI), with the idea of merging the best features of AllSMT-based WMI [33] and those of KC-based WMI [20, 28].

References

- [1] B. Badban and J. van de Pol. An Algorithm to Verify Formulas by means of (O,S,=)-BDDs. In *Proceedings of the 9th Annual Computer Society of Iran Computer Conference*, Tehran, Iran, 2004.
- [2] B. Badban and J. van de Pol. Zero, successor and equality in BDDs. *Annals of Pure and Applied Logic*, 133(1):101–123, May 2005. ISSN 0168-0072. doi: 10.1016/j.apal.2004.10.005.
- [3] C. W. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli. Satisfiability Modulo Theories. In *Handbook of Satisfiability*, volume 336, pages 1267–1329. IOS Press, 2 edition, 2021.
- [4] S. Bova. SDDs Are Exponentially More Succinct than OBDDs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), Feb. 2016. ISSN 2374-3468. doi: 10.1609/aaai.v30i1.10107.
- [5] G. Broeck. On the completeness of first-order knowledge compilation for lifted probabilistic inference. *Advances in Neural Information Processing Systems*, 24, 2011.
- [6] R. E. Bryant. Graph-Based Algorithms for Boolean Function Manipulation. *IEEE Transactions on Computers*, C-35(8):677–691, Aug. 1986. ISSN 1557-9956. doi: 10.1109/TC.1986.1676819.
- [7] R. E. Bryant and M. N. Velev. Boolean satisfiability with transitivity constraints. *ACM Transactions on Computational Logic*, 3(4):604–627, Oct. 2002. ISSN 1529-3785. doi: 10.1145/566385.566390.
- [8] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic model checking: 10^{20} States and beyond. *Information and Computation*, 98(2):142–170, June 1992. ISSN 0890-5401. doi: 10.1016/0890-5401(92)90017-A.
- [9] R. Cavada, A. Cimatti, A. Franzen, K. Kalyanasundaram, M. Roveri, and R. Shyamasundar. Computing Predicate Abstractions by Integrating BDDs and SMT Solvers. In *Formal Methods in Computer Aided Design*, pages 69–76, Nov. 2007. doi: 10.1109/FAMCAD.2007.35.
- [10] S. Chaki, A. Gurfinkel, and O. Strichman. Decision diagrams for linear arithmetic. In *Formal Methods in Computer Aided Design*, pages 53–60, Nov. 2009. doi: 10.1109/FMCAD.2009.5351143.
- [11] W. Chan, R. Anderson, P. Beame, and D. Notkin. Combining constraint solving and symbolic model checking for a class of systems with non-linear constraints. In O. Grumberg, editor, *Computer Aided Verification*, Lecture Notes in Computer Science, pages 316–327, Berlin, Heidelberg, 1997. Springer. ISBN 978-3-540-69195-2. doi: 10.1007/3-540-63166-6_32.
- [12] M. Chavira and A. Darwiche. On probabilistic inference by weighted model counting. *Artificial Intelligence*, 172(6-7):772–799, 2008.
- [13] A. Cimatti, A. Franzen, A. Griggio, K. Kalyanasundaram, and M. Roveri. Tighter integration of BDDs and SMT for Predicate Abstraction. In *2010 Design, Automation & Test in Europe Conference & Exhibition (DATE 2010)*, pages 1707–1712, Mar. 2010. doi: 10.1109/DATE.2010.5457090.
- [14] A. Cimatti, A. Griggio, and R. Sebastiani. Computing Small Unsatisfiable Cores in Satisfiability Modulo Theories. *Journal of Artificial Intelligence Research*, 40:701–728, Apr. 2011. ISSN 1076-9757. doi: 10.1613/jair.3196.
- [15] A. Cimatti, A. Griggio, B. J. Schaafsma, and R. Sebastiani. A Modular Approach to MaxSAT Modulo Theories. In *16th International Conference on Theory and Applications of Satisfiability Testing*, Lecture Notes in Computer Science, pages 150–165, Berlin, Heidelberg, July 2013. Springer-Verlag. ISBN 978-3-642-39070-8. doi: 10.1007/978-3-642-39071-5_12.
- [16] A. Cimatti, A. Griggio, B. J. Schaafsma, and R. Sebastiani. The MathSAT 5 SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems, TACAS’13.*, volume 7795 of *LNCS*, pages 95–109. Springer, 2013.
- [17] A. Darwiche. SDD: A new canonical representation of propositional knowledge bases. In *22nd International Joint Conference on Artificial Intelligence*, volume 2 of *IJCAI’11*, pages 819–826, Barcelona, Catalonia, Spain, July 2011. AAAI Press. ISBN 978-1-57735-514-4.
- [18] A. Darwiche and P. Marquis. A knowledge compilation map. *Journal of Artificial Intelligence Research*, 17(1):229–264, Sept. 2002. ISSN 1076-9757.
- [19] D. Deharbe and S. Ranise. Light-weight theorem proving for debugging and verifying units of code. In *First International Conference on Software Engineering and Formal Methods, 2003.Proceedings.*, pages 220–228, Sept. 2003. doi: 10.1109/SEFM.2003.1236224.
- [20] P. Z. Dos Martires, A. Dries, and L. De Raedt. Exact and Approximate Weighted Model Integration with Probability Density Functions Using Knowledge Compilation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):7825–7833, July 2019. ISSN 2374-3468, 2159-5399. doi: 10.1609/aaai.v33i01.33017825.
- [21] D. Fierens, G. V. d. Broeck, I. Thon, B. Gutmann, and L. De Raedt. Inference in probabilistic logic programs using weighted cnf’s. *arXiv preprint arXiv:1202.3719*, 2012.
- [22] P. Fontaine and E. P. Gribomont. Using BDDs with Combinations of Theories. In M. Baaz and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning*, Lecture Notes in Computer Science, pages 190–201, Berlin, Heidelberg, 2002. Springer. ISBN 978-3-540-36078-0. doi: 10.1007/3-540-36078-6_13.
- [23] J. Friso Groote and J. van de Pol. Equational Binary Decision Diagrams. In M. Parigot and A. Voronkov, editors, *Logic for Programming and Automated Reasoning*, Lecture Notes in Artificial Intelligence, pages 161–178, Berlin, Heidelberg, 2000. Springer. ISBN 978-3-540-44404-6. doi: 10.1007/3-540-44404-1_11.
- [24] M. Gario and A. Micheli. PySMT: A solver-agnostic library for fast prototyping of SMT-based algorithms. In *SMT Workshop 2015*, 2015.
- [25] A. Goel, K. Sajid, H. Zhou, A. Aziz, and V. Singhal. BDD based procedures for a theory of equality with uninterpreted functions. In A. J. Hu and M. Y. Vardi, editors, *Computer Aided Verification*, Lecture Notes in Computer Science, pages 244–255, Berlin, Heidelberg, 1998. Springer. ISBN 978-3-540-69339-0. doi: 10.1007/BFb0028749.
- [26] A. Goel, K. Sajid, H. Zhou, A. Aziz, and V. Singhal. BDD Based Procedures for a Theory of Equality with Uninterpreted Functions. *Formal Methods in System Design*, 22(3):205–224, May 2003. ISSN 1572-8102. doi: 10.1023/A:1022988809947.
- [27] J. Huang et al. Combining knowledge compilation and search for conformant probabilistic planning. In *ICAPS*, pages 253–262, 2006.
- [28] S. Kolb, P. Z. D. Martires, and L. D. Raedt. How to Exploit Structure while Solving Weighted Model Integration Problems. In *35th Conference on Uncertainty in Artificial Intelligence*, pages 744–754. PMLR, Aug. 2020.
- [29] S. K. Lahiri, R. Nieuwenhuis, and A. Oliveras. SMT Techniques for Fast Predicate Abstraction. In *Computer Aided Verification*, pages 424–437, 2006. ISBN 978-3-540-37411-4.
- [30] J. Möller, J. Lichtenberg, H. R. Andersen, and H. Hulgaard. Difference Decision Diagrams. In J. Flum and M. Rodríguez-Artalejo, editors, *Computer Science Logic*, Lecture Notes in Computer Science, pages 111–125, Berlin, Heidelberg, 1999. Springer. ISBN 978-3-540-48168-3. doi: 10.1007/3-540-48168-0_9.
- [31] Q.-S. Phan. Model counting modulo theories. *arXiv preprint arXiv:1504.02796*, 2015.
- [32] F. Somenzi. Cudd: Cu decision diagram package-release 2.4. 0. *University of Colorado at Boulder*, 21, 2009.
- [33] G. Spallitta, G. Masina, P. Morettin, A. Passerini, and R. Sebastiani. Enhancing smt-based weighted model integration by structure awareness. *Artificial Intelligence*, 328:104067, 2024.
- [34] J. van de Pol and O. Tveretina. A BDD-Representation for the Logic of Equality and Uninterpreted Functions. In J. Jędrzejowicz and A. Szepietowski, editors, *Mathematical Foundations of Computer Science 2005*, Lecture Notes in Computer Science, pages 769–780, Berlin, Heidelberg, 2005. Springer. ISBN 978-3-540-31867-5. doi: 10.1007/11549345_66.
- [35] G. Van den Broeck. *Lifted inference and learning in statistical relational models*. PhD thesis, PhD thesis, KU Leuven, 2013.

A Appendix: Proofs of the theorems

Proof of Theorem 3

Proof. Let $\Psi[\alpha] \stackrel{\text{def}}{=} \mathcal{T}\text{-DD}(\varphi[\alpha])$ and $\Psi'[\alpha] \stackrel{\text{def}}{=} \mathcal{T}\text{-DD}(\varphi'[\alpha])$. Then $\Psi[\alpha] = \Psi'[\alpha]$ if and only if $\Psi^p[\alpha] = \Psi'^p[\alpha]$. By Definition 1, $\Psi^p[\alpha]$ and $\Psi'^p[\alpha]$ are DDs, which are canonical by hypothesis. Thus $\Psi^p[\alpha] = \Psi'^p[\alpha]$ if and only if $\Psi^p[\alpha] \equiv \Psi'^p[\alpha]$, that is, if and only if $\Psi[\alpha] \equiv \Psi'[\alpha]$. $\square$

Proof of Theorem 5

Proof. Consider two SMT($\mathcal{T}$) formulas $\varphi[\alpha]$ and $\varphi'[\alpha]$. By Proposition 1(d), $\varphi[\alpha] \equiv_{\mathcal{T}} \varphi'[\alpha]$ if and only if $H_{\alpha}(\varphi) = H_{\alpha}(\varphi')$. By the definition of $H_{\alpha}(\dots)$ all the η s in $H_{\alpha}(\dots)$ are total on α and pairwise disjoint, so that $H_{\alpha}(\varphi) = H_{\alpha}(\varphi')$ if and only if $\bigvee_{\eta_i \in H_{\alpha}(\varphi)} \eta_i \equiv \bigvee_{\eta'_i \in H_{\alpha}(\varphi')} \eta'_i$. Since $\mathcal{T}\text{-DD}(\varphi[\alpha]) \equiv \bigvee_{\eta_i \in H_{\alpha}(\varphi)} \eta_i$ and $\mathcal{T}\text{-DD}(\varphi'[\alpha]) \equiv \bigvee_{\eta'_i \in H_{\alpha}(\varphi')} \eta'_i$, then we have that $H_{\alpha}(\varphi) = H_{\alpha}(\varphi')$ if and only if $\mathcal{T}\text{-DD}(\varphi[\alpha]) \equiv \mathcal{T}\text{-DD}(\varphi'[\alpha])$. By Theorem 3, $\mathcal{T}\text{-DD}(\varphi[\alpha]) \equiv \mathcal{T}\text{-DD}(\varphi'[\alpha])$ if and only if $\mathcal{T}\text{-DD}(\varphi[\alpha]) = \mathcal{T}\text{-DD}(\varphi'[\alpha])$. $\square$

Proof of Theorem 6

Proof. Theorem 6 is a corollary of Theorem 8 (which we prove below) by setting $\beta \stackrel{\text{def}}{=} \emptyset$. $\square$

Proof of Theorem 7

Proof. Theorem 7 is a corollary of Theorem 11 (which we prove below) by setting $\beta \stackrel{\text{def}}{=} \emptyset$. $\square$

Proof of Theorem 8

Proof. Since $Cl_{\alpha, \beta}(\varphi)$ rules out $P_{\alpha}(\varphi)$, we have:

$$\bigvee_{\rho_j[\alpha] \in P_{\alpha}(\varphi)} \rho_j[\alpha] \wedge \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} Cl[\alpha, \beta] \equiv \perp, \quad (9)$$

$$\text{i.e.:} \quad \bigvee_{\rho_j[\alpha] \in P_{\alpha}(\varphi)} \rho_j^p[A] \wedge \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B] \equiv \perp, \quad (10)$$

$$\text{equiv.:} \quad \bigvee_{\rho_j[\alpha] \in P_{\alpha}(\varphi)} \rho_j^p[A] \wedge \exists B. \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B] \equiv \perp \quad (11)$$

Let $\varphi^{*p}[A] \stackrel{\text{def}}{=} \exists B. \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B]$. Since the $\eta_i[\alpha]$ s in $H_{\alpha}(\varphi)$ are all total on α , then for each $\eta_i[\alpha]$, either $\eta_i^p[A] \models \varphi^{*p}[A]$ or $\eta_i^p[A] \models \neg \varphi^{*p}[A]$. The latter is not possible, because it would mean that $\eta_i^p[A] \wedge \varphi^{*p}[A] \models \perp$, and hence, by (11), $\eta_i \in P_{\alpha}(\varphi)$, which would contradict the fact that $H_{\alpha}(\varphi)$ and $P_{\alpha}(\varphi)$ are disjoint. Thus we have:

$$\bigvee_{\eta_i[\alpha] \in H_{\alpha}(\varphi)} \eta_i^p[A] \models \exists B. \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B]. \quad (12)$$

Hence, we have that:

$$\begin{aligned} & \varphi^p[A] \wedge \exists B. \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B] \\ \text{by (1):} & \equiv \left(\bigvee_{\eta_i[\alpha] \in H_{\alpha}(\varphi)} \eta_i^p[A] \vee \bigvee_{\rho_j[\alpha] \in P_{\alpha}(\varphi)} \rho_j^p[A] \right) \wedge \\ & \exists B. \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B] \end{aligned} \quad (13)$$

$$\begin{aligned} \wedge/\vee: & \equiv \left(\bigvee_{\eta_i[\alpha] \in H_{\alpha}(\varphi)} \eta_i^p[A] \wedge \exists B. \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B] \right) \vee \\ & \left(\bigvee_{\rho_j[\alpha] \in P_{\alpha}(\varphi)} \rho_j^p[A] \wedge \exists B. \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B] \right) \end{aligned} \quad (14)$$

$$\text{by (11):} \equiv \left(\bigvee_{\eta_i[\alpha] \in H_{\alpha}(\varphi)} \eta_i^p[A] \wedge \exists B. \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B] \right)$$

$$\text{by (12):} \equiv \bigvee_{\eta_i[\alpha] \in H_{\alpha}(\varphi)} \eta_i^p[A].$$

$\square$

Proof of Theorem 10

Proof. By applying Theorem 8 to both $\varphi[\alpha]$ and $\varphi'[\alpha]$ we obtain:

$$\varphi^p[A] \wedge \exists B. \bigwedge_{Cl[\alpha, \beta] \in Cl_{\alpha, \beta}(\varphi)} C_l^p[A, B] \equiv \bigvee_{\eta_i[\alpha] \in H_{\alpha}(\varphi)} \eta_i^p[A] \quad (15)$$

$$\varphi'^p[A] \wedge \exists B'. \bigwedge_{Cl'[\alpha, \beta'] \in Cl_{\alpha, \beta'}(\varphi')} C_l'^p[A, B'] \equiv \bigvee_{\eta'_i[\alpha] \in H_{\alpha}(\varphi')} \eta'^p_i[A]. \quad (16)$$

By Property 1(d), $\varphi[\alpha] \equiv_{\mathcal{T}} \varphi'[\alpha]$ if and only if $H_{\alpha}(\varphi) = H_{\alpha}(\varphi')$, that is, if and only if:

$$\bigvee_{\eta_i[\alpha] \in H_{\alpha}(\varphi)} \eta_i^p[A] \equiv \bigvee_{\eta'_i[\alpha] \in H_{\alpha}(\varphi')} \eta'^p_i[A]. \quad (17)$$

Thus, by combining (17) with (15) and (16), we have that $\varphi[\alpha] \equiv_{\mathcal{T}} \varphi'[\alpha]$ if and only if (7) holds. $\square$

Proof of Theorem 11

Proof. Let $\varphi[\alpha]$ and $\varphi'[\alpha]$ be $\mathcal{T}$ -formulas. By Theorem 10, $\varphi[\alpha] \equiv_{\mathcal{T}} \varphi'[\alpha]$ if and only if (7) holds. Since the DDs are canonical and $\mathcal{B}2\mathcal{T}$ is injective, (7) holds if and only if $\mathcal{T}\text{-DD}(\varphi[\alpha]) = \mathcal{T}\text{-DD}(\varphi'[\alpha])$. $\square$

B Appendix: Extended Related Work

Several works have tried to leverage Decision Diagrams from the propositional to the SMT level. Most of them are theory-specific, in particular focusing on $\mathcal{E}$, $\mathcal{EUF}$ and (fragments of) arithmetic. $\mathcal{EUF}$ -DDs are of particular interest in hardware verification [25, 26], while DDs for arithmetic have been mainly studied for the verification of infinite-state systems [11].

In the following, we present an analysis of the most relevant works that leverage DDs from the propositional to the SMT level. We focus on generalization of OBDDs and SDDs, as they are the most used DDs in the literature. In Table 1, we summarize the main properties of the analyzed works. From the table, we can see that most

	Solver	Theory	Avail.	Prune inconsistent paths	Semi-canonical	Canonical
BDD	EQ-BDDs [23]	$\mathcal{E}$	*	✓	✓	✗
	GOEL-FM [25, 26]	$\mathcal{E}$	✗	✓	✓	?
	GOEL- e_{ij} [25, 26]	$\mathcal{E}$	✗	✗	✗	✗
	BRYANT [7]	$\mathcal{E}$	✗	✓	✓	?
	EUF-BDDs [34]	$\mathcal{EUF}$	*	✓	✓	✗
	(0,S,=)-BDDs [1, 2]	$\mathcal{E} \cup \{0, S\}$	✗	✓	✓	✗
	DDD [30]	$\mathcal{DL}$	✗	✓	✓	✗
	LDD [10]	$\mathcal{LA}$	✓	✗	✗	✗
	CHAN [11]	$\mathcal{NLA}$	✗	✓	✗	✗
	HARVEY [19]	any	✗	✓	✗	✗
	FONTAINE [22]	any	✗	✓	✗	✗
	BDD+SMT [9, 13]	any	*	✓	✗	✗
SDD	XSDD [20, 28]	$\mathcal{LRA}$	*	✗	✗	✗

Table 1: Theory-aware decision diagram solvers. The column *Available* indicates whether the solver is available for the given theory. The symbol “✓” means that the solver is publicly available, * means that the solver is implemented within another tool and not directly usable, and ✗ means that the solver is not publicly available. In the column *Canonical*, the symbol “?” indicates that the DDs may be canonical, but the authors do not provide a proof.

of the works are theory-specific, and while several $\mathcal{T}$ -semicanonical DDs have been proposed, $\mathcal{T}$ -canonical representations have been achieved only in some very-specific cases. With the only exception of LDDs [10], all the analyzed works do not have a public implementation or are implemented within other tools, making them not directly usable.

$\mathcal{T}$ -DDs for $\mathcal{EUF}$. In Goel et al. [25, 26], the authors describe two techniques to build OBDDs for the theory of equality ($\mathcal{E}$). The first consists in encoding each of the n variables with $\lceil \log(n) \rceil$ bits, reducing to a Boolean formula. The resulting OBDD is, therefore, canonical, but its size is unmanageable even for small instances. The second approach consists in introducing a Boolean atom e_{ij} for each equality $x_i = x_j$, and building a OBDD over these atoms. This essentially builds the OBDD of the Boolean abstraction of the formula, which allows for theory-inconsistent paths. This problem has been addressed in Bryant and Velev [7], where transitivity lemmas are instantiated in advance and conjoined with the OBDD. This approach is similar in flavour to our approach, and produces OBDDs whose refinement is $\mathcal{EUF}$ -canonical; the main difference is that the procedure used to generate the lemmas is specific to the theory of equality, whereas our approach is general and can be applied to any theory supported by the SMT solver.

EQ-BDDs [23] extend OBDDs to allow for nodes with atoms representing equation between variables. EUF-BDDs [34] extend EQ-BDDs to atoms involving also uninterpreted functions. (0,S,=)-BDDs [1, 2] extend EQ-BDDs to atoms involving also the zero constant and the successor function. In all three cases, rewriting rules are applied to prune inconsistent paths. The resulting OBDDs are $\mathcal{T}$ -semicanonical, but not $\mathcal{T}$ -canonical.

$\mathcal{T}$ -DDs for arithmetic. Difference Decision Diagrams (DDDs) [30] are a generalization of OBDDs to the theory of difference logic ($\mathcal{DL}$). The building procedure consists in first building the refinement of the OBDD of the Boolean abstraction of the formula, and then pruning inconsistent paths by applying local and path reductions. Local reductions are based on rewriting rules, leveraging implications between predicates to reduce redundant splitting. Path reductions prune inconsistent paths, both those going to the $\top$ and $\perp$ terminals. The resulting DDD is $\mathcal{T}$ -semicanonical, as $\mathcal{T}$ -valid and $\mathcal{T}$ -inconsistent formulas are represented by the $\top$ and $\perp$ DDDs, respectively. In general, however, they are not $\mathcal{T}$ -canonical, even for formulas on the same atoms. Some desirable properties are

discussed, and they conjecture that DDDs with these properties are canonical.

LDDs [10] generalize DDDs to $\mathcal{LA}$ formulas. However, the implementation restricts to the theory of Two Variables Per Inequality ($\mathcal{TVPI}$) over real or integer variables. Moreover, only local reductions are applied, making them not even $\mathcal{T}$ -semicanonical. Most importantly, not even contradictions are recognized.

In [11], a procedure is described to build $\mathcal{T}$ -OBDDs for nonlinear arithmetic ($\mathcal{NLA}$). The procedure consists in building the refinement of the OBDD of the Boolean abstraction of the formula, and then using an (incomplete) quadratic constraint solver to prune inconsistent paths. As a result, the $\mathcal{T}$ -OBDD is not $\mathcal{T}$ -semi-canonical, since $\mathcal{T}$ -valid formulas may have different representations.

To the best of our knowledge, XSDDs [20, 28] are the only tentative to extend SDDs to support first-order theories. XSDDs have been proposed in the context of Weighted Model Integration (WMI), and extend SDDs by allowing for atoms representing linear inequalities on real variables in decision nodes. However, they only propose to refine the SDD of the Boolean abstraction of the formula, without any pruning of inconsistent paths. Simplifications are only done at later stages during the WMI computation.

$\mathcal{T}$ -DDs for arbitrary theories. In [19], a general way has been proposed to build $\mathcal{T}$ -OBDDs. The tool named HARVEY first builds the refinement of the OBDD of the Boolean abstraction of the formula. Then, it looks for a $\mathcal{T}$ -inconsistent path, from which it extracts a subset of $\mathcal{T}$ -inconsistent constraints. The negation of this subset, which is a $\mathcal{T}$ -lemma, is conjoined to the $\mathcal{T}$ -OBDD to prune this and possibly other $\mathcal{T}$ -inconsistent paths. The procedure is iterated until no inconsistent paths are found. Here, only the lemmas necessary to prune $\mathcal{T}$ -inconsistent partial assignments satisfying the formula are generated, making the resulting $\mathcal{T}$ -OBDD not $\mathcal{T}$ -semicanonical, as $\mathcal{T}$ -valid formulas may have different representations.

The technique described in [22] is similar, but it generalizes to combination of theories.

In [9], the authors propose a general method to build $\mathcal{T}$ -OBDD by integrating an OBDD compiler with an SMT($\mathcal{T}$) solver, which is invoked to check the consistency of a path during its construction. The approach was refined in [13], where the authors propose many optimizations to get a tighter integration of the SMT solver within the $\mathcal{T}$ -OBDD construction. In both cases, all inconsistent paths are pruned, but the resulting $\mathcal{T}$ -OBDD is not $\mathcal{T}$ -semicanonical.