

# On the Role of Canonicity in Bottom-up Knowledge Compilation

Guy Van den Broeck and Adnan Darwiche

Computer Science Department  
University of California, Los Angeles  
{guyvdb, darwiche}@cs.ucla.edu

## Abstract

We consider the problem of bottom-up compilation of knowledge bases, which is usually predicated on the existence of a polytime function for combining compilations using Boolean operators (usually called an `Apply` function). While such a polytime `Apply` function is known to exist for certain languages (e.g., OBDDs) and not exist for others (e.g., DNNF), its existence for certain languages remains unknown. Among the latter is the recently introduced language of Sentential Decision Diagrams (SDDs), for which a polytime `Apply` function exists for unreduced SDDs, but remains unknown for reduced ones (i.e. canonical SDDs). We resolve this open question in this paper and consider some of its theoretical and practical implications. Some of the findings we report question the common wisdom on the relationship between bottom-up compilation, language canonicity and the complexity of the `Apply` function.

## Introduction

The Sentential Decision Diagram (SDD) is a recently proposed circuit representation of propositional knowledge bases (Darwiche 2011). The SDD is a target language for knowledge compilation (Selman and Kautz 1996; Darwiche and Marquis 2002), meaning that once a propositional knowledge base is compiled into an SDD, the SDD can be reused to answer multiple hard queries efficiently (e.g., clausal entailment or model counting).

SDDs subsume Ordered Binary Decision Diagrams (OBDD) (Bryant 1986) and come with a tighter size bound (Darwiche 2011), while still being equally powerful as far as their polytime support for classical queries (e.g., the ones in Darwiche and Marquis (2002)). Moreover, SDDs are subsumed by d-DNNFs (Darwiche 2001), which received much attention over the last decade, for fault diagnosis (Elliott and Williams 2006), planning (Palacios et al. 2005), databases (Suciu et al. 2011), but most importantly for probabilistic inference (Chavira, Darwiche, and Jaeger 2006; Chavira and Darwiche 2008; Fierens et al. 2011). Even though SDDs are less succinct than d-DNNFs, they can be compiled *bottom-up*, just like OBDDs. For example, a clause can be compiled by disjoining the SDDs corresponding to its literals, and a CNF can be compiled by conjoining the SDDs corresponding to its clauses. This bottom-up compilation is implemented using the `Apply` function, which

combines two SDDs using Boolean operators.<sup>1</sup> Bottom-up compilation makes SDDs attractive for certain applications (e.g., probabilistic inference (Choi, Kisa, and Darwiche 2013)) and can be critical when the knowledge base to be compiled is constructed incrementally (see the discussion in Pipatsrisawat and Darwiche (2008)).

According to common wisdom, a language supports bottom-up compilation only if it supports a polytime `Apply` function. For example, OBDDs are known to support bottom-up compilation and have traditionally been compiled this way. In fact, the discovery of SDDs was mostly driven by the need for bottom-up compilation, which was preceded by the discovery of *structured decomposability* (Pipatsrisawat and Darwiche 2008): a property that enables some Boolean operations to be applied in polytime. SDDs satisfy this property and stronger ones, leading to a polytime `Apply` function (Darwiche 2011). This function, however, assumes that the SDDs are unreduced (i.e., not canonical). For reduced SDDs, the existence of a polytime `Apply` function has been an open question since SDDs were first introduced (note, however, that both reduced and unreduced OBDDs are supported by a polytime `Apply` function).

We resolve this open question in this paper, showing that such an `Apply` function does not exist in general. We also pursue some theoretical and practical implications of this result, on bottom-up compilation in particular. On the practical side, we reveal an empirical finding that seems quite surprising: bottom-up compilation with reduced SDDs is much more feasible practically than with unreduced ones, even though the latter supports a polytime `Apply` function while the former does not. This finding questions common convictions on the relative importance of a polytime `Apply` in contrast to canonicity as desirable properties for a language that supports efficient bottom-up compilation. On the theoretical side, we show that certain transformations (e.g., conditioning) can lead to blowing up the size of reduced SDDs, while they don't for unreduced SDDs. Finally, we identify a subset of SDDs for which a polytime `Apply` exists even under reduction.

---

<sup>1</sup>The `Apply` function (and its name) originated in the OBDD literature (Bryant 1986)

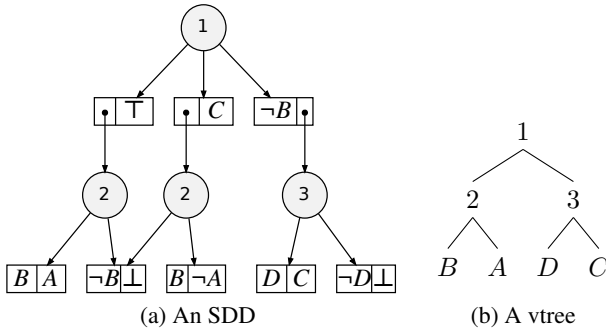


Figure 1: An SDD and vtree for  $(A \wedge B) \vee (B \wedge C) \vee (C \wedge D)$ .

## Background

We will use the following notation for propositional logic. Upper-case letters (e.g.,  $X$ ) denote propositional *variables* and lower-case letters denote their *instantiations* (e.g.,  $x$ ). Bold letters represent sets of variables (e.g.,  $\mathbf{X}$ ) and their instantiations (e.g.,  $\mathbf{x}$ ). A *literal* is a variable or its negation. A *Boolean function*  $f(\mathbf{X})$  maps each instantiation  $\mathbf{x}$  to  $\top$  (true) or  $\perp$  (false).

### The SDD Representation

The Sentential Decision Diagram (SDD) is a newly introduced representation language for propositional knowledge bases (Darwiche 2011). Figure 1a depicts an SDD: paired-boxes  $\boxed{p | s}$  are called *elements* and represent conjunctions  $(p \wedge s)$ , where  $p$  is called a *prime* and  $s$  is called a *sub*. Circles are called *decision nodes* and represent disjunctions of their child elements.

An SDD is constructed for a given *vtree*, which is a full binary tree whose leaves are variables; see for example Figure 1b. Every node in an SDD respects some vtree node (except for  $\top$  and  $\perp$ ). SDD literals respect the vtree leaf labeled with their variable. In Figure 1a, decision nodes are labeled with the vtree node they respect. Consider a decision node with elements  $\boxed{p_1 | s_1}, \dots, \boxed{p_n | s_n}$ , and suppose that it respects a vtree node  $v$  which has variables  $\mathbf{X}$  in its left subtree and variables  $\mathbf{Y}$  in its right subtree. We are then guaranteed that each prime  $p_i$  will only mention variables in  $\mathbf{X}$  and that each sub  $s_i$  will only mention variables in  $\mathbf{Y}$ . Moreover, the primes are guaranteed to represent propositional sentences that are consistent, mutually exclusive, and exhaustive. This type of decomposition is called an  $(\mathbf{X}, \mathbf{Y})$ -partition (Darwiche 2011). For example, the top decision node in Figure 1a has the following elements

$$\left\{ \underbrace{(A \wedge B)}_{\text{prime}}, \underbrace{\top}_{\text{sub}}, \underbrace{(\neg A \wedge B)}_{\text{prime}}, \underbrace{C}_{\text{sub}}, \underbrace{(\neg B)}_{\text{prime}}, \underbrace{(D \wedge C)}_{\text{sub}} \right\}, \quad (1)$$

which correspond to an  $(AB, CD)$ -partition of the function  $(A \wedge B) \vee (B \wedge C) \vee (C \wedge D)$ . One can verify that the primes and subs of this partition satisfy the properties mentioned above.

An  $(\mathbf{X}, \mathbf{Y})$ -partition is *compressed* if it has distinct subs, and an SDD is compressed if its  $(\mathbf{X}, \mathbf{Y})$ -partitions are compressed. A Boolean function may have multiple  $(\mathbf{X}, \mathbf{Y})$ -

Query	Description	OBDD	SDD	d-DNNF
<b>CO</b>	consistency	✓	✓	✓
<b>VA</b>	validity	✓	✓	✓
<b>CE</b>	clausal entailment	✓	✓	✓
<b>IM</b>	implicant check	✓	✓	✓
<b>EQ</b>	equivalence check	✓	✓	?
<b>CT</b>	model counting	✓	✓	✓
<b>SE</b>	sentential entailment	✓	✓	◦
<b>ME</b>	model enumeration	✓	✓	✓

Table 1: ✓ means that a polytime algorithm exists for the corresponding language/query, while ◦ means that no such algorithm exists unless  $P = NP$ .

partitions, but the compressed partition is unique. Our example function has another  $(AB, CD)$ -partition, which is not compressed:

$$\left\{ (A \wedge B, \top), (\neg A \wedge B, C), (A \wedge \neg B, D \wedge C), (\neg A \wedge \neg B, D \wedge C) \right\}. \quad (2)$$

An uncompressed  $(\mathbf{X}, \mathbf{Y})$ -partition can be compressed by merging all elements  $(p_1, s), \dots, (p_n, s)$  that share the same sub into one element  $(p_1 \vee \dots \vee p_n, s)$ . Compressing (2) combines the two last elements into  $([A \wedge \neg B] \vee [\neg A \wedge \neg B], D \wedge C) = (\neg B, D \wedge C)$ , resulting in (1). This is the unique compressed  $(AB, CD)$ -partition.

Given a vtree, each Boolean function also has a unique *compressed* SDD, when this property is combined with either *trimming* or *normalization* properties (Darwiche 2011). These are weaker properties that mildly affect the size of an SDD. For example, a trimmed SDD contains no decision nodes of the form  $\{(\top, \alpha)\}$  or  $\{(\alpha, \top), (\neg \alpha, \perp)\}$  (we can trim an SDD by replacing these nodes with  $\alpha$ ). Compressed and trimmed SDDs are canonical, and so are compressed and normalized SDDs (Darwiche 2011).

OBDDs correspond precisely to SDDs that are constructed using a special type of vtree, called a *right-linear* vtree (Darwiche 2011). The left child of each inner node in these vtrees is a variable. With right-linear vtrees, compressed/trimmed SDDs correspond to reduced OBDDs, while compressed/normalized SDDs correspond to oblivious OBDDs (Xue, Choi, and Darwiche 2012) (reduced and oblivious OBDDs are also canonical). We will refer to compressed and trimmed SDDs as *reduced* SDDs and restrict our attention to them in the rest of the paper.

The size of an OBDD depends critically on the underlying variable order. Similarly, the size of an SDD depends critically on the vtree used (right-linear vtrees correspond to variable orders). Vtree search algorithms can sometimes find SDDs that are orders-of-magnitude more succinct than OBDDs found by searching for variable orders (Choi and Darwiche 2013). Such algorithms assume canonical SDDs, allowing one to search the space of SDDs by searching the space of vtrees.

### Queries and Transformations

SDDs are a strict subset of deterministic, decomposable negation normal form (d-DNNF). They are actually a strict

---

**Algorithm 1**  $\text{Apply}(\alpha, \beta, \circ)$ 

---

```
1: if  $\alpha$  and  $\beta$  are constants or literals then
2:   return  $\alpha \circ \beta$  // result is a constant or literal
3: else if  $\text{Cache}(\alpha, \beta, \circ) \neq \text{nil}$  then
4:   return  $\text{Cache}(\alpha, \beta, \circ)$  // has been computed before
5: else
6:    $\gamma \leftarrow \{\}$ 
7:   for all elements  $(p_i, s_i)$  in  $\alpha$  do
8:     for all elements  $(q_j, r_j)$  in  $\beta$  do
9:        $p \leftarrow \text{Apply}(p_i, q_j, \wedge)$ 
10:      if  $p$  is consistent then
11:         $s \leftarrow \text{Apply}(s_i, r_j, \circ)$ 
12:        add element  $(p, s)$  to  $\gamma$ 
13:      (optionally)  $\gamma \leftarrow \text{Compress}(\gamma)$  // compression
// get unique decision node and return it
14:   return  $\text{Cache}(\alpha, \beta, \circ) \leftarrow \text{UniqueD}(\gamma)$ 
```

---

subset of structured d-DNNF and, hence, support the same polytime queries supported by structured d-DNNF (Pipatsrisawat and Darwiche 2008); see Table 1. This makes SDDs as powerful as OBDDs in terms of their support for these queries.

SDDs satisfy stronger properties than structured d-DNNFs, allowing one, for example, to conjoin or disjoin two SDDs in polytime. Algorithm 1 shows the outline of an  $\text{Apply}$  function (Darwiche 2011) that takes two SDDs  $\alpha$  and  $\beta$ , and a binary Boolean operator  $\circ$  (e.g.,  $\wedge, \vee, \text{xor}$ ), and returns the SDD for  $\alpha \circ \beta$ . This code assumes that the SDD is normalized instead of trimmed. The code for trimmed SDDs is similar, although a bit more detailed. Line 13 optionally compresses each partition, in order to return a compressed SDD. Without compression, this algorithm has a time and space complexity of  $O(nm)$ , where  $n$  and  $m$  are the sizes of input SDDs (Darwiche 2011). This comes at the expense of losing canonicity. Whether a polytime complexity can be attained under compression is an open question.

There are several implications of this question. For example, depending on the answer, one would know whether certain *transformations*, such as conditioning and existential quantification, can be supported in polytime on reduced SDDs. Moreover, according to common wisdom, a negative answer may preclude bottom-up compilation from being feasible on reduced SDDs. We explore this question and its implications next.


### Complexity of $\text{Apply}$ on Reduced SDDs

The size of a decision node is the number of its elements, and the size of an SDD is the sum of sizes attained by its decision nodes. We now show that reduction, given a fixed vtree, may blow up the size of an SDD.

**Theorem 1.** *There exists a class of Boolean functions  $f_m(X_1, \dots, X_m)$  and corresponding vtrees  $T_m$  such that  $f_m$  has an SDD of size  $O(m^2)$  wrt vtree  $T_m$ , yet the reduced SDD of function  $f_m$  wrt vtree  $T_m$  has size  $\Omega(2^m)$ .*

*Proof.* Consider the function  $f_n^a(\mathbf{X}, \mathbf{Y}, Z) =$

$\bigvee_{i=1}^n \left( \bigwedge_{j=1}^{i-1} \neg Y_j \right) \wedge Y_i \wedge X_i$  which has  $2n + 1$  variables. Of these,  $Z$  is non-essential. Consider a vtree  $T_n$  of the form



where the sub-vtrees over variables  $\mathbf{X}$  and  $\mathbf{Y}$  are arbitrary. We will now construct an uncompressed SDD for this function using vtree  $T_n$  and whose size is  $O(n^2)$ . We will then show that the compressed SDD for this function and vtree has a size  $\Omega(2^n)$ .

The first step is to construct a partition of FUNCTION  $f_n^a$  that respects the root vtree node, that is, an  $(\mathbf{XY}, Z)$ -partition. Consider

$$\left\{ \begin{array}{l} (Y_1 \wedge X_1, \top), \\ (\neg Y_1 \wedge Y_2 \wedge X_2, \top), \\ \dots, \\ (\neg Y_1 \wedge \dots \wedge \neg Y_{n-1} \wedge Y_n \wedge X_n, \top), \\ (Y_1 \wedge \neg X_1, \perp), \\ (\neg Y_1 \wedge Y_2 \wedge \neg X_2, \perp), \\ \dots, \\ (\neg Y_1 \wedge \dots \wedge \neg Y_{n-1} \wedge Y_n \wedge \neg X_n, \perp), \\ (\neg Y_1 \wedge \dots \wedge \neg Y_n, \perp) \end{array} \right\},$$

which is equivalently written as

$$\bigcup_{i=1}^n \left\{ \left( \bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \wedge X_i, \top \right), \left( \bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \wedge \neg X_i, \perp \right) \right\} \cup \left\{ \left( \bigwedge_{j=1}^n \neg Y_j, \perp \right) \right\}.$$

The size of this partition is  $2n + 1$ , and hence linear in  $n$ . It is *uncompressed*, because there are  $n$  elements that share sub  $\top$  and  $n + 1$  elements that share sub  $\perp$ . The subs already respect the leaf vtree node labelled with variable  $Z$ .

In a second step, each of prime above is written as a compressed  $(\mathbf{X}, \mathbf{Y})$ -partition that respects the left child of the vtree root. Prime  $\bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \wedge X_i$  becomes

$$\left\{ \left( X_i, \bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \right), (\neg X_i, \perp) \right\},$$

prime  $\bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \wedge \neg X_i$  becomes

$$\left\{ \left( \neg X_i, \bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \right), (X_i, \perp) \right\}$$

and prime  $\bigwedge_{j=1}^n \neg Y_j$  becomes

$$\left\{ \left( \top, \bigwedge_{j=1}^n \neg Y_j \right) \right\}.$$

The sizes of these partitions are bounded by 2.

Finally, we need to represent the above primes as SDDs over variables  $\mathbf{X}$  and the subs as SDDs over variables  $\mathbf{Y}$ . Since these primes and subs correspond to terms (i.e. conjunctions of literals), each has a compact SDD representation, independent of the chosen sub-vtree over variables  $\mathbf{X}$  and  $\mathbf{Y}$ . For example, we can choose a right-linear vtree over variables  $\mathbf{X}$ , and similarly for variables  $\mathbf{Y}$ , leading to an OBDD representation of each prime and sub, with a size linear in  $n$  for each OBDD. The full SDD for function  $f_n^a$  will then have a size which is  $O(n^2)$ . Recall that this SDD is uncompressed as some of its decision nodes have elements with equal subs.

The compressed SDD for this function and vtree is unique. We now show that its size must be  $\Omega(2^n)$ . We first observe that the unique, compressed  $(\mathbf{X}, \mathbf{Y})$ -partition of function  $f_n^a$  is

$$\left\{ \left( \bigvee_{i=1}^n \left( \bigwedge_{j=1}^{i-1} \neg Y_j \right) \wedge Y_i \wedge X_i, \top \right), \left( \left[ \bigvee_{i=1}^n \left( \bigwedge_{j=1}^{i-1} \neg Y_j \right) \wedge Y_i \wedge \neg X_i \right] \vee \left[ \bigwedge_{j=1}^n \neg Y_j \right], \perp \right) \right\}.$$

Its first sub is the function

$$f_n^b(\mathbf{X}, \mathbf{Y}) = \bigvee_{i=1}^n \left( \bigwedge_{j=1}^{i-1} \neg Y_j \right) \wedge Y_i \wedge X_i,$$

which we need to represent as an  $(\mathbf{X}, \mathbf{Y})$ -partition to respect left child of the vtree root. However, Xue, Choi, and Darwiche (2012) proved the following.

**Lemma 2.** *The compressed  $(\mathbf{X}, \mathbf{Y})$ -partition of  $f_n^b(\mathbf{X}, \mathbf{Y})$  has  $2^n$  elements.*

This becomes clear when looking at the function  $f_n^b$  after instantiating the  $\mathbf{X}$ -variables. Each distinct  $\mathbf{x}$  results in a unique subfunction  $f_n^b(\mathbf{x}, \mathbf{Y})$ , and all states  $\mathbf{x}$  are mutually exclusive and exhaustive. Therefore,

$$\{(\mathbf{x}, f_n^b(\mathbf{x}, \mathbf{Y})) \mid \mathbf{x} \text{ instantiates } \mathbf{X}\}$$

is the unique, compressed  $(\mathbf{X}, \mathbf{Y})$ -partition of function  $f_n^b(\mathbf{X}, \mathbf{Y})$ , which has  $2^n$  elements. Hence, the compressed SDD must have size  $\Omega(2^n)$ .  $\square$

Theorem 1 has a number of implications, which are summarized in Table 2; see also Darwiche and Marquis (2002).

**Theorem 3.** *The results in Table 2 hold.*

The proof of this theorem is in the supplementary material. First, combining two reduced SDDs (e.g., using the conjoin or disjoin operator) may lead to a reduced SDD whose size is exponential in the size of inputs. Hence, if we activate compression in Algorithm 1, the algorithm may take exponential time in the worst-case. Second, conditioning a reduced SDD on a literal may exponentially increase its size (assuming the result is also reduced). Third, forgetting a variable (i.e., existentially quantifying it) from a reduced

Notation	Transformation	SDD	Reduced SDD
CD	polytime conditioning	✓	•
FO	polytime forgetting	•	•
SFO	polytime singleton forgetting	✓	•
$\wedge C$	polytime conjunction	•	•
$\wedge BC$	polytime bounded conjunction	✓	•
$\vee C$	polytime disjunction	•	•
$\vee BC$	polytime bounded disjunction	✓	•
$\neg C$	polytime negation	✓	✓

Table 2: ✓ means “satisfies”; • means “does not satisfy”. Satisfaction means the existence of a polytime algorithm that implements the corresponding transformation.

SDD may exponentially increase its size (again, assuming that the result is also reduced). These results may seem discouraging. However, we argue next that, in practice, working with reduced SDDs is actually favorable despite the lack of polytime guarantees on these transformations. Moreover, we identify a class of vtrees, called *bounded vtrees*, which include right-linear ones, and which induce reduced SDDs with a polytime Apply function.

### Canonicity or a Polytime Apply?

One has two options when working with SDDs. The first option is to work with unreduced SDDs, which are not canonical, but are supported by a polytime Apply function. The second option is to work with reduced SDDs, which are canonical but lose the advantage of a polytime Apply function. The classical reason for seeking canonicity is that it leads to a very efficient equivalence test, which takes constant time (both reduced and unreduced SDDs support a polytime equivalence test, but the one known for unreduced SDDs is not a constant time test). The classical reason for seeking a polytime Apply function is to enable bottom-up compilation, that is, compiling a knowledge base (e.g., CNF or DNF) into an SDD by repeated application of the Apply function to components of the knowledge base (e.g., clauses or terms). If our goal is efficient bottom-up compilation, one may expect that unreduced SDDs provide a better alternative. Empirically, however, this turns out to be false. Our goal in this section is to shed some light on this phenomena through some empirical evidence and then an explanation.

We used the SDD package provided by the Automated Reasoning Group at UCLA<sup>2</sup> in our experiments. The package works with reduced SDDs, but can be adjusted to work with unreduced SDDs as long as dynamic vtree search is not invoked.<sup>3</sup> In our first experiment, we compiled CNFs from the LGSynth89 benchmarks into the following:

- Reduced SDDs respecting an arbitrary vtree. Dynamic vtree search is used to minimize the size of the SDD during compilation, starting from a balanced vtree.

<sup>2</sup>Available at <http://reasoning.cs.ucla.edu/sdd/>

<sup>3</sup>Dynamic vtree search requires reduced SDDs as this reduces the search space over SDDs into one over vtrees.

Name	Reduced SDDs+s	Reduced SDDs	Unreduced SDDs
C17	99	171	286
majority	123	193	384
b1	166	250	514
cm152a	149	3,139	18,400
cm82a	225	363	683
cm151a	614	1,319	24,360
cm42a	394	823	276,437
cm138a	463	890	9,201,336
decod	471	810	1,212,302
tcon	596	1,327	618,947
parity	549	978	2,793
cmb	980	2,311	81,980
cm163a	886	1,793	21,202
pcl	785	1,366	n/a
x2	785	1,757	12,150,626
cm85a	1,015	2,098	19,657
cm162a	907	2,050	153,228
cm150a	1,603	5,805	17,265,164
pcler8	1,518	4,335	15,532,667
cu	1,466	5,789	n/a
pm1	1,810	3,699	n/a
mux	1,825	6,517	n/a
cc	1,451	6,938	n/a
unreg	3,056	668,531	n/a
ldd	1,610	2,349	n/a
count	4,168	51,639	n/a
comp	2,212	4,500	205,105
f51m	3,290	6,049	n/a
my_adder	2,793	4,408	35,754
cht	4,832	13,311	n/a

Table 3: LGSynth89 SDD sizes.

Name	Reduced SDDs+s	Reduced SDDs	Unreduced SDDs
C17	0.00	0.00	0.00
majority	0.00	0.00	0.00
b1	0.00	0.00	0.00
cm152a	0.01	0.01	0.02
cm82a	0.01	0.00	0.00
cm151a	0.04	0.00	0.04
cm42a	0.03	0.00	0.10
cm138a	0.02	0.01	109.05
decod	0.04	0.01	1.40
tcon	0.05	0.00	0.33
parity	0.02	0.00	0.00
cmb	0.12	0.02	0.06
cm163a	0.06	0.00	0.02
pcl	0.07	0.01	n/a
x2	0.08	0.02	19.87
cm85a	0.08	0.01	0.03
cm162a	0.08	0.01	0.16
cm150a	0.16	0.06	60.37
pcler8	0.18	0.05	33.32
cu	0.19	0.10	n/a
pm1	0.27	0.05	n/a
mux	0.19	0.09	n/a
cc	0.22	0.04	n/a
unreg	0.66	263.06	n/a
ldd	0.23	0.10	n/a
count	1.05	0.24	n/a
comp	0.24	0.01	0.22
f51m	0.52	0.32	n/a
my_adder	0.24	0.02	0.04
cht	1.24	0.36	n/a

Table 4: LGSynth89 compilation times, in seconds.

- Reduced SDDs respecting a fixed balanced vtree.
- Unreduced SDDs respecting a fixed balanced vtree.

Tables 3 and 4 show the corresponding sizes and compilation times. According to these results, unreduced SDDs end up several orders of magnitude larger than the reduced ones, with or without dynamic vtree search. For the harder problems, this translates to orders-of-magnitude increase in compilation times. Often, we cannot even compile the input without reduction (due to running out of 4GB of memory), even on relatively easy benchmarks. For the easiest benchmarks, dynamic vtree search is slower due to the overhead, but yields smaller compilations. The benefit of vtree search shows only in harder problems (e.g., “unreg”).

This experiment clearly shows the advantage of reduced SDDs over unreduced ones, even though the latter supports a polytime `Apply` function while the former does not. This begs an explanation and we provide one next that we back up by additional experimental results.

The benefit of reduced SDDs is canonicity, which plays a critical role in the performance of the `Apply` function. Consider in particular Line 4 of Algorithm 1. The test `Cache( $\alpha, \beta, \circ$ ) \neq \text{nil}` checks whether SDDs  $\alpha$  and  $\beta$  have been previously combined using the Boolean operator  $\circ$ . Without canonicity, it is possible that we would have combined some  $\alpha'$  and  $\beta'$  using  $\circ$ , where SDD  $\alpha'$  is equivalent

to, but distinct from SDD  $\alpha$  (and similarly for  $\beta'$  and  $\beta$ ). In this case, the cache test would fail, causing `Apply` to recompute the same result again. Worse, the SDD returned by `Apply( $\alpha, \beta, \circ$ )` may be distinct from the SDD returned by `Apply( $\alpha', \beta', \circ$ )`, even though the two SDDs are equivalent. This redundancy also happens when  $\alpha$  is not equivalent to  $\alpha'$  (and similarly for  $\beta$  and  $\beta'$ ),  $\alpha \circ \beta$  is equivalent to  $\alpha' \circ \beta'$ , but the result returned by `Apply( $\alpha, \beta, \circ$ )` is distinct from the one returned by `Apply( $\alpha', \beta', \circ$ )`.

Two observations are due here. First, this redundancy is still under control when calling `Apply` only once: `Apply` runs in  $O(nm)$  time, where  $n$  and  $m$  are the sizes of input SDDs. However, this redundancy becomes problematic when calling `Apply` multiple times (as in bottom-up compilation), in which case quadratic performance is no longer as attractive. For example, if we use `Apply` to combine  $m$  SDDs of size  $n$  each, all we can say is that the output will be of size  $O(n^m)$ . The second observation is that the previous redundancy will not occur when working with reduced SDDs due to their canonicity: Two SDDs are equivalent iff they are represented by the same structure in memory.<sup>4</sup>

This analysis points to the following conclusion: While `Apply` has a quadratic complexity on unreduced SDDs, it

<sup>4</sup>This is due to the technique of *unique nodes* from OBDDs; see `UniqueD` in Algorithm 1.




Figure 2: Relative SDD size.




Figure 3: Relative number of recursive Apply calls.

may have a worse average complexity than `Apply` on reduced SDDs. Our next experiment is indeed directed towards this hypothesis.

For all benchmarks in Table 3 that can be compiled without vtree search, we intercept all non-trivial calls to `Apply` (when  $|\alpha| \cdot |\beta| > 500$ ) and report the size of the output  $|\alpha \circ \beta|$  divided by  $|\alpha| \cdot |\beta|$ . For unreduced SDDs, we know that  $|\alpha \circ \beta| = O(|\alpha| \cdot |\beta|)$  and that these ratios are therefore bounded above by some constant. For reduced SDDs, however, Theorem 3 states that there exists no constant bound.

Figure 2 shows the distribution of these ratios for the two methods (note the log scale). The number of function calls is 67,809 for reduced SDDs, vs. 1,626,591 for unreduced ones. The average ratio is 0.027 for reduced, vs. 0.101 for unreduced. Contrasting the theoretical bounds, reduced `Apply` incurs much smaller blowups than unreduced `Apply`. This is most clear for ratios in the range  $[0.48, 0.56]$ , covering 30% of the unreduced, but only 2% of the reduced calls.

The results are similar when looking at runtime for individual `Apply` calls, which we measure by the number of recursive `Apply` calls  $r$ . Figure 3 reports these, again relative to  $|\alpha| \cdot |\beta|$ . The ratio  $r / (|\alpha| \cdot |\beta|)$  is on average 0.013 for reduced SDDs, vs. 0.034 for unreduced ones. These results corroborate our earlier analysis, suggesting that canonicity is quite important for the performance of bottom-up compilers as they make repeated calls to the `Apply` function. In fact, this can be more important than a polytime `Apply`, perhaps contrary to common wisdom which seems to emphasize the importance of polytime `Apply` in effective bottom-up compilation (e.g., Pipatsrisawat and Darwiche (2008)).

## Bounded Vtrees

A *bounded vtree* is one for which the number of variables in any left subtree is bounded. This includes right-linear vtrees

which give rise to OBDDs, since each left subtree contains a single variable in this case. We now have the following.

**Theorem 4.** *The time and space complexity of Algorithm 1, with compression, is in  $O(nm)$ , where  $n$  and  $m$  are the sizes of its inputs, assuming that the input SDDs are reduced and respect a bounded vtree.*

The compression step of Algorithm 1 identifies elements  $(p_i, s)$  and  $(p_j, s)$  that share sub  $s$ , and merges these elements into the element  $(p_i \vee p_j, s)$  by calling `Apply` recursively to disjoin primes  $p_i$  and  $p_j$ . Since the vtree is bounded, primes  $p_i$  and  $p_j$  must be over a bounded number of variables. Hence, the complexity of compression is bounded, leading `Apply` to have the same complexity with or without compression.

For example, in right-linear vtrees (i.e., OBDDs), primes are literals over a single variable. Hence, all decision nodes are of the form  $\{(X, \alpha), (\neg X, \beta)\}$ . On these, compression occurs when  $\alpha = \beta$ , resulting in the partition  $\{(X \vee \neg X, \alpha)\} = \{(\top, \alpha)\}$ , which trimming replaces by  $\alpha$ . This corresponds to the OBDD reduction rule that eliminates decision nodes with isomorphic children (Bryant 1986).

Xue, Choi, and Darwiche (2012) showed a class of Boolean functions whose OBDDs have exponential size with respect to certain orders (right-linear vtrees), but which have SDDs of linear size when the vtrees are not right-linear (but have the same left-to-right variable order). The used vtrees, however, were not bounded. It would be interesting to see if a similar result can be obtained for bounded vtrees.

## Conclusions

We have shown that the `Apply` function on reduced SDDs can take exponential time in the worst case, resolving a question that has been open since SDDs were first introduced. We have also pursued some of the theoretical and practical implications of this result. On the theoretical side, we showed that it implies an exponential complexity for various transformations, such as conditioning and existential quantification. We also identified the class of bounded vtrees, for which the `Apply` function can be implemented in polynomial time even for reduced SDDs. On the practical side, we argued empirically that working with reduced SDDs remains favorable, despite the polytime complexity of the `Apply` function on unreduced SDDs. The canonicity of reduced SDDs, we argued, is more valuable for bottom-up compilation than a polytime `Apply` due to its role in facilitating caching and dynamic vtree search. Our findings appear contrary to some of the common wisdom on the relationship between bottom-up compilation, canonicity and the complexity of the `Apply` function.

## Acknowledgments

We thank Arthur Choi, Doga Kisa, and Umut Oztok for helpful suggestions. This work was supported by ONR grant #N00014-12-1-0423, NSF grant #IIS-1118122, NSF grant #IIS-0916161, and the Research Foundation-Flanders (FWO-Vlaanderen). GVdB is also at KU Leuven, Belgium.

## Complexity of Transformations

We now prove Theorem 3, stating that the results in Table 2 hold. We will first show the results for unreduced SDDs, and then prove the results for reduced SDDs.

For *unreduced SDDs*, Darwiche (2011) showed support for  $\wedge\mathbf{BC}$ ,  $\vee\mathbf{BC}$ , and  $\neg\mathbf{C}$  (see Algorithm 1). We show support for unreduced  $\mathbf{CD}$  next.

**Theorem 5.** *We can condition an unreduced SDD on a literal  $\ell$  in polynomial time by replacing  $\ell$  by  $\top$  and  $\neg\ell$  by  $\perp$ . Moreover, when removing all elements whose prime is equivalent to  $\perp$ , the resulting sentence is an unreduced SDD.*

*Proof.* It is clear that the procedure transforms  $\alpha$  into a sentence that is logically equivalent to  $\alpha|\ell$ : the first step directly follows the definition of conditioning, and the second step maintains logical equivalence. We need to show next that the result is syntactically an SDD, by showing that the primes in its partitions are consistent, exhaustive, and mutually exclusive. The second step enforces consistency of the primes. Moreover, if the primes are exhaustive, that is,  $p_1 \vee \dots \vee p_n \equiv \top$ , then  $p_1|\ell \vee \dots \vee p_n|\ell \equiv (p_1 \vee \dots \vee p_n)|\ell \equiv \top|\ell \equiv \top$ , and the result of conditioning is also exhaustive. Finally, when  $p_i$  and  $p_j$  are mutually exclusive, that is,  $p_i \wedge p_j \equiv \perp$ , then  $p_i|\ell \wedge p_j|\ell \equiv (p_i \wedge p_j)|\ell \equiv \perp|\ell \equiv \perp$ , and the conditioned primes are also mutually exclusive.  $\square$

Support for  $\mathbf{SFO}$  follows from the support for  $\mathbf{CD}$  and  $\vee\mathbf{BC}$ . The negative results for  $\mathbf{FO}$ ,  $\wedge\mathbf{C}$  and  $\vee\mathbf{C}$  follow from identical OBDD results in Darwiche and Marquis (2002), and the fact that OBDDs are a special case of SDDs.

For *reduced SDDs*, the negative  $\mathbf{FO}$ ,  $\wedge\mathbf{C}$  and  $\vee\mathbf{C}$  results also follow from OBDD results. It is also clear from Algorithm 1 that negating a reduced SDD  $\alpha$  by computing  $\text{Apply}(\alpha, \top, \text{xor})$  does not cause any subs to become equivalent. Therefore, negating a reduced SDD leads to a reduced result, and reduced SDDs support  $\neg\mathbf{C}$ . The remaining results in Table 2, on  $\mathbf{CD}$ ,  $\mathbf{SFO}$ ,  $\wedge\mathbf{BC}$  and  $\vee\mathbf{BC}$  are discussed next.

**Theorem 6.** *There exists a class of Boolean functions  $f(X_1, \dots, X_n)$  and vtrees  $T_n$  for which the reduced SDD has size  $O(n)$ , yet the reduced SDD for the function  $f(X_1, \dots, X_n)|\ell$  has size  $\Omega(2^n)$  for some literal  $\ell$ .*

*Proof.* Consider the function

$$f_n^c(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, W) = \bigvee_{i=1}^n \bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \wedge [(X_i \wedge (W \vee Z_i)) \vee (\neg X_i \wedge Z_i)]$$

and the vtree depicted in Figure 4a.

The root of the reduce SDD for  $f_n^c$  is an  $(\mathbf{XY}, \mathbf{ZW})$ -partition that respects vtree node 1, consisting of elements

$$\bigcup_{i=1}^n \left\{ \left( \bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \wedge X_i, W \vee Z_i \right), \left( \bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \wedge \neg X_i, Z_i \right) \right\},$$




Figure 4: Vtree structures used in proofs

together with  $(\bigwedge_{j=1}^n \neg Y_j, \perp)$  to make it exhaustive. The size of this partition is linear in  $n$ . It has the same primes as the unreduced SDD for  $f_n^a$  used in the proof of Theorem 1, only now the partition is compressed, as all subs are distinct.

The primes of this partition can be represented as compressed  $(\mathbf{X}, \mathbf{Y})$ -partitions, exactly as in the second step for Theorem 1. The remaining primes and subs (over  $\mathbf{X}$ , over  $\mathbf{Y}$ , and over  $\mathbf{Z} \cup \{W\}$ ) are all simple conjunctions or disjunctions of literals that have a linear reduced SDD representation for any vtree.

We have now obtained a polysize SDD. However, when we condition this SDD on the literal  $W$ , all  $n$  subs of the form  $W \vee Z_i$  become equivalent to  $\top$ . Their elements need to be compressed into the single element  $(\bigvee_{i=1}^n \bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \wedge X_i, \top)$ . Its prime is again the function  $f_n^b(\mathbf{X}, \mathbf{Y})$  for which Lemma 2 states that the reduced SDD wrt vtree node 2 has exponential size.  $\square$

**Theorem 7.** *There is a class of Boolean functions  $f(X_1, \dots, X_n)$  and vtrees  $T_n$  for which the reduced SDD has size  $O(n)$ , yet the reduced SDD for the Boolean function  $f(X_1, \dots, X_n) \wedge \ell$  has size  $\Omega(2^n)$  for some literal  $\ell$ .*

*Proof.* Consider again the reduced SDD for  $f_n^c$  that was constructed in the proof of Theorem 6 for the vtree in Figure 4a. Conjoining this SDD with the SDD for literal  $W$  makes the  $n$  subs of the form  $W \vee Z_i$  equivalent to  $W \wedge (W \vee Z_i) = W$ . Compressing these creates the element  $(\bigvee_{i=1}^n \bigwedge_{j=1}^{i-1} \neg Y_j \wedge Y_i \wedge X_i, W)$ , whose prime is again  $f_n^b(\mathbf{X}, \mathbf{Y})$ , which has no polysize reduced SDD for vtree node 2.  $\square$

This already proves that  $\text{Apply}$  is worst-case exponential when performing conjunctions on reduced SDDs. Given that reduced SDDs support polytime negation, this result generalizes to any binary Boolean operator  $\circ$  that is functionally complete together with negation (Wernick 1942). Support for these operators would allow us to do polytime conjunction by combining  $\circ$  and negation. One such operator is disjunction, which is therefore also worst-case exponential.

Suppose now that we can perform singleton forgetting in polytime, which is defined as  $\exists L. \alpha = (\alpha|L) \vee (\alpha|\neg L)$ . Then given any two reduced SDDs  $\beta$  and  $\gamma$  respecting the same vtree  $T$ , we can obtain  $\beta \vee \gamma$  in polytime as follows. Add a new variable  $L$  to vtree  $T$ , as depicted in Figure 4b. The reduced SDD  $\alpha$  for the function  $(L \wedge \beta) \vee (\neg L \wedge \gamma)$  has the

root partition  $\{(L, \beta), (\neg L, \gamma)\}$ . Forgetting  $L$  from  $\alpha$  results in the reduced SDD for  $\beta \vee \gamma$ . Hence, if single forgetting can be done in polytime, then bounded disjunction can also be done in polytime. Since the latter is impossible, the former is also impossible.

## References

- Bryant, R. E. 1986. Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers* C-35:677–691.
- Chavira, M., and Darwiche, A. 2008. On probabilistic inference by weighted model counting. *Artificial Intelligence Journal* 172(6–7):772–799.
- Chavira, M.; Darwiche, A.; and Jaeger, M. 2006. Compiling relational bayesian networks for exact inference. *International Journal of Approximate Reasoning* 42(1):4–20.
- Choi, A., and Darwiche, A. 2013. Dynamic minimization of sentential decision diagrams. In *Proceedings of the 27th Conference on Artificial Intelligence (AAAI)*.
- Choi, A.; Kisa, D.; and Darwiche, A. 2013. Compiling probabilistic graphical models using sentential decision diagrams. In *Proceedings of the 12th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*.
- Darwiche, A., and Marquis, P. 2002. A knowledge compilation map. *Journal of Artificial Intelligence Research* 17:229–264.
- Darwiche, A. 2001. On the tractability of counting theory models and its application to belief revision and truth maintenance. *Journal of Applied Non-Classical Logics* 11(1-2):11–34.
- Darwiche, A. 2011. SDD: A new canonical representation of propositional knowledge bases. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, 819–826.
- Elliott, P., and Williams, B. 2006. DNNF-based belief state estimation. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*.
- Fierens, D.; Van den Broeck, G.; Thon, I.; Gutmann, B.; and Raedt, L. D. 2011. Inference in probabilistic logic programs using weighted CNF's. In *UAI*, 211–220.
- Palacios, H.; Bonet, B.; Darwiche, A.; and Geffner, H. 2005. Pruning conformant plans by counting models on compiled d-DNNF representations. In *Proceedings of the 15th International Conference on Automated Planning and Scheduling*, 141–150.
- Pipatsrisawat, K., and Darwiche, A. 2008. New compilation languages based on structured decomposability. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI)*, 517–522.
- Selman, B., and Kautz, H. 1996. Knowledge compilation and theory approximation. *Journal of the ACM (JACM)* 43(2):193–224.
- Suciu, D.; Olteanu, D.; Ré, C.; and Koch, C. 2011. Probabilistic databases. *Synthesis Lectures on Data Management* 3(2):1–180.
- Wernick, W. 1942. Complete sets of logical functions. *Transactions of the American Mathematical Society* 51(1):117–132.
- Xue, Y.; Choi, A.; and Darwiche, A. 2012. Basing decisions on sentences in decision diagrams. In *In Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI)*.