

# Symmetry-Driven Decision Diagrams for Knowledge Compilation

Anicet Bart and Frédéric Koriche and Jean-Marie Lagniez and Pierre Marquis<sup>1</sup>

**Abstract.** In this paper, symmetries are exploited for achieving significant space savings in a knowledge compilation perspective. More precisely, the languages FBDD and DDG of decision diagrams are extended to the languages  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  of symmetry-driven decision diagrams, where  $X$  is a set of “symmetry-free” variables and  $Y$  is a set of “top” variables. Both the time efficiency and the space efficiency of  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  are analyzed, in order to put those languages in the knowledge compilation map for propositional representations. It turns out that each of  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  satisfies **CT** (the model counting query). We prove that no propositional language over a set  $X \cup Y$  of variables, satisfying both **CO** (the consistency query) and **CD** (the conditioning transformation), is at least as succinct as any of  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  unless the polynomial hierarchy collapses. The price to be paid is that only a restricted form of conditioning and a restricted form of forgetting are offered by  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$ . Nevertheless, this proves sufficient for a number of applications, including configuration and planning. We describe a compiler targeting  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  and give some experimental results on planning domains, highlighting the practical significance of these languages.

## 1 INTRODUCTION

It is well-known that many reasoning and optimization problems exhibit symmetries, and that recognizing and taking advantage of symmetries is a way to improve the computational time needed to solve those problems. Actually, much work has been devoted to this issue for decades. Among other highlights is the fact that the resolution system, equipped with a global symmetry rule, permits polynomial-length proofs of several combinatorial principles, including the pigeon/hole formulae [9], while such formulae require resolution proofs of exponential length [8, 14].

The main objective of this paper is to show that exploiting symmetries also proves valuable for *achieving space savings* in a knowledge compilation perspective, i.e., to derive more succinct compiled representations while preserving queries and transformations of interest. In order to reach this goal, we extend the language FBDD of free binary decision diagrams [7] to the language  $\text{Sym-FBDD}_{X,Y}$  of symmetry-driven free binary decision diagrams, containing free binary decision diagrams equipped with symmetries.  $X$  is a (possibly empty) set of “symmetry-free” variables, and  $Y$  is a (possibly full) set of “top” variables. We also extend the language DDG of decomposable decision diagrams [5] (a superset of FBDD where decomposable  $\wedge$ -nodes are allowed in the representations) to the language  $\text{Sym-DDG}$  of symmetry-driven decomposable decision dia-

grams, where the same conditions on  $X$  and  $Y$  are considered. We analyze  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  along the lines of the knowledge compilation map for propositional representations [4], by identifying the queries and transformations of interest for which some polynomial-time algorithms exist when the input is a representation from one of those languages; we also investigate the space efficiency of  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$ . Based on these investigations, it turns out that each of  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  satisfies the critical **CT** query (model counting) which is, for many languages, hard to satisfy (a **#P**-complete problem). We prove that no propositional language over a set  $X \cup Y$  of variables, satisfying both **CO** (the consistency query) and **CD** (the conditioning transformation), is at least as succinct as any of  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  unless the polynomial hierarchy collapses. The price to be paid is that only restricted forms of conditioning and of projection are offered by  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$ , namely conditioning over  $X$  and projection on  $Y$ . Nevertheless, this proves sufficient for a number of applications, including configuration and planning. We describe a compiler targeting  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  and give some experimental results on planning domains, highlighting the practical significance of these languages.

The paper is organized as follows. After introducing the formal background, the languages  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  are defined and analyzed. A CNF-to- $\text{Sym-DDG}_{X,Y}$  compiler is described in the next section. Before concluding, empirical results on some planning instances are presented, showing that the size of  $\text{Sym-DDG}_{X,Y}$  compilations can be significantly smaller than the size of the state-of-the-art d-DNNF compilations. Proofs are not provided in the paper due to space limitations, but can be found in an extended version, available at [www.cril.fr/~marquis/symddg.pdf](http://www.cril.fr/~marquis/symddg.pdf).

## 2 FORMAL PRELIMINARIES

Let  $PS$  be a finite set of propositional variables. A permutation  $\sigma$  over  $L_{PS}$ , the set of all literals over  $PS$ , is a bijective mapping from  $L_{PS} = PS \cup \{\neg x \mid x \in PS\}$  to  $L_{PS}$ . Any permutation  $\sigma$  can be extended easily to a morphism associating a propositional formula over  $PS$  with a propositional formula over  $PS$ , by stating that for every propositional connective  $c$  of arity  $k$ , we have  $\sigma(c(\alpha_1, \dots, \alpha_k)) = c(\sigma(\alpha_1), \dots, \sigma(\alpha_k))$ . We also note  $\sigma(X) = \{\sigma(x) \mid x \in X\}$  for any subset  $X$  of  $PS$ .

Every permutation  $\sigma$  under consideration in this paper is assumed to satisfy the following *stability condition*: for any pair of literals  $\ell_1, \ell_2$ ,  $\sigma(\ell_1) = \ell_2$  iff  $\sigma(\sim \ell_1) = \sim \ell_2$  where  $\sim \ell$  is the opposite of  $\ell$ , i.e.,  $\sim x = \neg x$  and  $\sim \neg x = x$ . Any permutation  $\sigma$  will be represented in a *simplified cycle notation*, i.e., as a product of cycles corresponding to its orbits (with at least two elements), where exactly one of

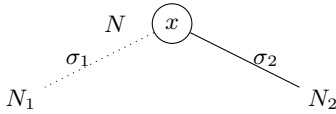
<sup>1</sup> CRIL-CNRS, Université d’Artois, France, email: name@cril.fr

the two orbits  $(l_1 \dots l_k)$  and  $(\sim l_1 \dots \sim l_k)$  are represented, whenever  $(l_1 \dots l_k)$  is an orbit of  $\sigma$ .

For instance, if  $PS = \{x_1, \dots, x_6\}$ ,  $(x_1 \neg x_3 x_4)(x_5 x_6)$  denotes the permutation  $\sigma$  associating  $x_1$  with  $\neg x_3$ ,  $\neg x_1$  with  $x_3$ ,  $x_3$  with  $\neg x_4$ ,  $\neg x_3$  with  $x_4$ ,  $x_4$  with  $x_1$ ,  $\neg x_4$  with  $\neg x_1$ ,  $x_5$  with  $x_6$ ,  $\neg x_5$  with  $\neg x_6$ ,  $x_6$  with  $x_5$ , and  $\neg x_6$  with  $\neg x_5$ , while  $x_2$  and  $\neg x_2$  are left unchanged by  $\sigma$ . The identity permutation is represented by the empty word using the simplified cycle notation.

By  $\Sigma$ , we denote the set of all bijective mappings from  $L_{PS}$  to  $L_{PS}$  satisfying the stability condition. Clearly enough,  $\Sigma$  is closed by composition: if  $\sigma_1, \sigma_2 \in \Sigma$  then  $\sigma_1 \circ \sigma_2 \in \Sigma$ . Since  $\Sigma$  is also closed for the inverse (if  $\sigma \in \Sigma$ , then  $\sigma^{-1} \in \Sigma$ ) and it contains the identity element (which is the neutral element for composition),  $\Sigma$  is a permutation group. Clearly enough, applying a permutation  $\sigma \in \Sigma$  to a propositional formula  $\alpha$  does not change the number of models of the latter; especially,  $\alpha$  is satisfiable (resp. valid) iff  $\sigma(\alpha)$  is satisfiable (resp. valid).

In the rest of the paper, we focus on subsets of Sym-EDD, the language of symmetry-driven extended decision diagrams, where permutations are defined over  $\Sigma$ . Basically, Sym-EDD generalizes the language of "extended" decision diagrams (i.e., binary decision diagrams in which  $\wedge$ -nodes are allowed) by associating some permutations to the arcs and to the root node. Diagrams from Sym-EDD are based on decision nodes, where a decision node  $N$  labeled with  $x \in PS$  is a node with two children, having the following form:



Such a node  $N$  is noted  $ite(x, N_1, N_2)$ , where "ite" stands for "if ... then ... else ...".

**Definition 1 (Sym-EDD).** Sym-EDD is the set of all finite, single-rooted multi-DAGs<sup>2</sup> (also referred to as "formulae")  $\alpha$  where:

- each leaf node of  $\alpha$  is either the  $\top$ -node (a node labeled by the Boolean constant  $\top$  – always true) or the  $\perp$ -node (a node labeled by the Boolean constant  $\perp$  – always false)
- each internal node of  $\alpha$  is labeled by  $\wedge$  and has a finite number of children ( $\geq 1$ ), or it is a decision node labeled with a variable from  $PS$ ;
- each arc of  $\alpha$  is labeled with a permutation from  $\Sigma$ ;
- the root of  $\alpha$  is labeled with a permutation from  $\Sigma$ .

The size  $|\alpha|$  of a Sym-EDD formula  $\alpha$  is the number of nodes, plus the number of arcs in the DAG, plus the sizes of the permutations labeling the arcs of  $\alpha$  and its root. The set  $Var(\alpha)$  of variables of a Sym-EDD formula  $\alpha$  rooted at node  $N$  is defined by  $\{\sigma_N(x) \mid x \in Var(N)\}$ , where  $\sigma_N$  is the permutation labeling  $N$ , and  $Var(N)$  is defined as follows:

- if  $N$  is a leaf node labeled by a Boolean constant, then  $Var(N) = \emptyset$ ;
- if  $N$  is a node labeled by  $\wedge$  and having  $k$  children  $N_1, \dots, N_k$  such that  $\forall i \in 1, \dots, k$ ,  $\sigma_i$  is the label of the arc  $(N, N_i)$ , then  $Var(N) = \bigcup_{i=1}^k \sigma_i(Var(N_i))$ ;
- if  $N = ite(x, N_1, N_2)$  is a decision node such that  $\sigma_1$  is the label of the arc  $(N, N_1)$  and  $\sigma_2$  is the label of the arc  $(N, N_2)$ , then  $Var(N) = \{x\} \cup \sigma_1(Var(N_1)) \cup \sigma_2(Var(N_2))$ .

Clearly enough,  $Var(\alpha)$  can be computed in time linear in  $|\alpha|$ . Note that  $Var(\alpha)$  may easily differ from the set of variables occurring in  $\alpha$ , when no permutation is taken into account (or equivalently, when each permutation is equal to the identity permutation).

Let us now define the semantics of Sym-EDD formulae. A simple way to do so consists in associating with every Sym-EDD formula  $\alpha$  a tree-shaped NNF formula  $T(\alpha)$  which is logically equivalent to  $\alpha$ . Formally,  $T(\alpha)$  is given by  $\sigma_N(T(N))$  where  $N$  is the root of  $\alpha$  and  $T(N)$  is defined inductively as follows:

- if  $N$  is a leaf node labeled by the Boolean constant  $\top$  (resp.  $\perp$ ), then  $T(N) = \top$  (resp.  $\perp$ );
- if  $N$  is a node labeled by  $\wedge$  and having  $k$  children  $N_1, \dots, N_k$  such that  $\forall i \in 1, \dots, k$ ,  $\sigma_i$  is the label of the arc  $(N, N_i)$ , then  $T(N) = \bigwedge_{i=1}^k \sigma_i(T(N_i))$ ;
- if  $N = ite(x, N_1, N_2)$  is a decision node such that  $\forall i \in 1, \dots, 2$ ,  $\sigma_i$  is the label of the arc  $(N, N_i)$ , then  $T(N) = (\neg x \wedge \sigma_1(T(N_1))) \vee (x \wedge \sigma_2(T(N_2)))$ .

Of course, the size of  $T(\alpha)$  is exponentially larger than the size of  $\alpha$  in the general case. Anyway, the models of  $\alpha$  are precisely those of  $T(\alpha)$ . We denote by  $\|\alpha\|$  the number of models of  $\alpha$  over  $Var(\alpha)$ .

For space reasons, we assume the reader is familiar with the languages FBDD, DDG, DNNF [7, 5, 3] which are considered in the following, and with the KC map [4]. The basic queries considered in the KC map include tests for consistency **CO**, validity **VA**, implicates (clausal entailment) **CE**, implicants **IM**, sentential entailment **SE**, model counting **CT**, and model enumeration **ME**. We add to them the *model checking query* **MC**, which is not obvious for the languages we introduce in the paper. The basic transformations include conditioning (**CD**), (possibly bounded) closures under the connectives ( $\wedge C$ ,  $\wedge BC$ ,  $\vee C$ ,  $\vee BC$ ,  $\neg C$ ), and forgetting (**FO**), or dually projection (**PR**). We add to them the *restricted conditioning transformation*, and the *restricted projection transformation*:

**Definition 2 (X-RCD).** Let  $\mathcal{L}$  be a subset of Sym-EDD, and  $X \subseteq PS$ .  $\mathcal{L}$  satisfies X-RCD iff there is a polynomial-time algorithm that maps every formula  $\alpha$  in  $\mathcal{L}$  and every consistent term  $\gamma$  over some variables in  $X$  to a formula  $\alpha \mid \gamma$  in  $\mathcal{L}$  which is logically equivalent to the most general logical consequence  $\beta$  of  $\alpha \wedge \gamma$ , where  $\beta$  is independent from the variables occurring in  $\gamma$ .

**Definition 3 (Y-RPR).** Let  $\mathcal{L}$  be a subset of Sym-EDD, and  $Y \subseteq PS$ .  $\mathcal{L}$  satisfies Y-RPR iff there is a polynomial-time algorithm that maps every formula  $\alpha$  in  $\mathcal{L}$  and every  $Z \subseteq Y$  to a formula in  $\mathcal{L}$  which is logically equivalent to the projection  $\exists PS \setminus Z. \alpha$  of  $\alpha$  on  $Z$ .<sup>3</sup>

Clearly enough, X-RCD (resp. Y-RPR) coincides with the usual conditioning transformation **CD** (resp. projection transformation **PR**) when  $X = PS$  (resp.  $Y = PS$ ).

The relative space efficiency of propositional languages is captured by a pre-order  $\leq_s$ , where  $\mathcal{L}_1 \leq_s \mathcal{L}_2$  means that  $\mathcal{L}_1$  is at least as succinct as  $\mathcal{L}_2$ , i.e., there exists a polynomial  $p$  such that for every formula  $\alpha \in \mathcal{L}_2$ , there exists an equivalent formula  $\beta \in \mathcal{L}_1$  where  $|\beta| \leq p(|\alpha|)$ .

$\sim_s$  denotes the symmetric part of  $\leq_s$  defined by  $\mathcal{L}_1 \sim_s \mathcal{L}_2$  iff  $\mathcal{L}_1 \leq_s \mathcal{L}_2$  and  $\mathcal{L}_2 \leq_s \mathcal{L}_1$ .  $<_s$  denotes the asymmetric part of  $\leq_s$  defined by  $\mathcal{L}_1 <_s \mathcal{L}_2$  iff  $\mathcal{L}_1 \leq_s \mathcal{L}_2$  and  $\mathcal{L}_2 \not\leq_s \mathcal{L}_1$ .  $\mathcal{L}_1 \not\leq_s^* \mathcal{L}_2$  means that  $\mathcal{L}_1 \not\leq_s \mathcal{L}_2$  unless the polynomial hierarchy PH collapses (which is considered very unlikely in complexity theory).  $\mathcal{L}_1 <_s^* \mathcal{L}_2$  is a short for  $\mathcal{L}_1 \leq_s \mathcal{L}_2$  and  $\mathcal{L}_2 \not\leq_s^* \mathcal{L}_1$ .

<sup>2</sup> More than one arc between two nodes is allowed.

<sup>3</sup> Or, in an equivalent way, to the forgetting of every variable in  $\alpha$  but those of  $Z$ .



	CO	VA	CE	IM
Sym-DDG <sub>X,Y</sub>	✓	✓	○	○
Sym-FBDD <sub>X,Y</sub>	✓	✓	○	○
DDG	✓	✓	✓	✓
FBDD	✓	✓	✓	✓

	SE	CT	ME	MC
Sym-DDG <sub>X,Y</sub>	○	✓	✓	✓
Sym-FBDD <sub>X,Y</sub>	○	✓	✓	✓
DDG	○	✓	✓	✓
FBDD	○	✓	✓	✓

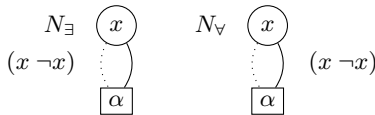
**Table 1.** Sym-DDG<sub>X,Y</sub>, Sym-FBDD<sub>X,Y</sub>, and the queries **CO**, **VA**, **CE**, **IM**, **SE**, **CT**, **ME**, **MC**. ✓ means “satisfies”, and ○ means “does not satisfy unless P = NP”.

	CD	X-RCD	PR	Y-RPR
Sym-DDG <sub>X,Y</sub>	○	✓	○	✓
Sym-FBDD <sub>X,Y</sub>	○	✓	○	✓
DDG	✓	✓	○	○
FBDD	✓	✓	•	•

	∧C	∧BC	∨C	∨BC	¬C
Sym-DDG <sub>X,Y</sub>	○	○	○	○	?
Sym-FBDD <sub>X,Y</sub>	○	○	○	○	✓
DDG	○	○	○	○	?
FBDD	•	○	•	○	✓

**Table 2.** Sym-DDG<sub>X,Y</sub>, Sym-FBDD<sub>X,Y</sub>, and the transformations **CD**, **X-RCD**, **PR**, **Y-RPR**, **∧C**, **∧BC**, **∨C**, **∨BC**, **¬C**. ✓ means “satisfies”, • means “does not satisfy”, and ○ means “does not satisfy unless P=NP”.

In a nutshell, it turns out that Sym-FBDD<sub>X,Y</sub> and Sym-DDG<sub>X,Y</sub> exhibit quite non-standard properties as target languages for knowledge compilation. Indeed, **CT** is typically hard to be satisfied (a #P-complete problem) while **CD** is usually obvious. In the same vein, model checking **MC** which is a straightforward query for usual knowledge compilation languages, is far from being easy for symmetry-driven graph-based languages, due to their ability of encoding quantifications in a succinct way. Indeed, assuming that  $\alpha$  is any Sym-EDD formula, the Sym-EDD formula rooted at node  $N_\exists$  on the figure below is equivalent to  $\exists x.\alpha$  while the formula rooted at node  $N_\forall$  is equivalent to  $\forall x.\alpha$ .<sup>5</sup> As a consequence, we get that Sym-EDD satisfies **CD** but also that **MC** for Sym-EDD formulae is PSPACE-hard!



The non-standard behavior of Sym-DDG<sub>X,Y</sub> (and its subset Sym-FBDD<sub>X,Y</sub>) w.r.t. unrestricted conditioning seems to be the price to be paid for an improved succinctness power. Indeed, the next proposition shows that the languages Sym-DDG<sub>X,Y</sub> and Sym-FBDD<sub>X,Y</sub>, are in some sense “very succinct”:

**Proposition 2.** *Let  $X, Y$  be two subsets of PS. No propositional language  $\mathcal{L}$  over  $X \cup Y$  satisfying **CD** and **CO** is at least as succinct as Sym-FBDD<sub>X,Y</sub> unless  $\Sigma_2^P = \Pi_2^P$ , i.e., we have  $\mathcal{L} \not\leq_s^* \text{Sym-FBDD}_{X,Y}$ .*

<sup>5</sup> Observe that, due to the read-once condition, none of these two formulae belongs to Sym-DDG, unless  $x \notin \text{Var}(\alpha)$ .

Consider the CNF formula  $\Delta_n$  containing every 3-clause  $\delta$  (i.e., a clause of size 3) over  $Y = \{y_1, \dots, y_n\}$  augmented by an additional literal  $x_\delta$  which identifies the clause.  $\Delta_n$  is thus a CNF over  $Y \cup X$  containing  $8 \times \binom{n}{3}$  4-clauses.  $X$  contains  $8 \times \binom{n}{3}$  variables  $x_\delta$ . The proof of Proposition 2 relies on the fact that  $\Delta_n$  can be represented by an equivalent Sym-FBDD formula of size linear in the size of  $\Delta_n$ . However, this statement does not hold for propositional languages satisfying **CO** and **CD** unless the polynomial hierarchy collapses. Indeed, to every CNF formula  $\alpha = \bigwedge_{i=1}^m \delta_i$  over  $Y = \{y_1, \dots, y_n\}$ , we can associate a consistent term  $\gamma_\alpha = \bigwedge_{i=1}^m \neg x_{\delta_i}$  such that  $\alpha$  is satisfiable iff  $\Delta_n \wedge \gamma_\alpha$  is satisfiable iff  $\Delta_n$  conditioned by  $\gamma_\alpha$  is satisfiable. Suppose now that  $\Delta_n$  has a polynomial-space representation  $\text{comp}(\Delta_n)$  in a propositional language  $\mathcal{L}$  over  $X \cup Y$ , where  $\mathcal{L}$  satisfies **CO** and **CD**. Then in order to check whether a CNF formula  $\alpha = \bigwedge_{i=1}^m \delta_i$  over  $\{y_1, \dots, y_n\}$  is satisfiable it would be enough to compute in polynomial time a  $\mathcal{L}$ -representation of  $\text{comp}(\Delta_n)$  conditioned by  $\gamma_\alpha$ , and to determine in polynomial time whether it is consistent or not. We would therefore get  $\text{NP} \subseteq \text{P/poly}$ , hence  $\Sigma_2^P = \Pi_2^P$  (see e.g. [13] for details).

As a consequence, state-of-the-art languages for knowledge compilation, like DNNEF, are not more succinct than any of the two languages Sym-FBDD, and Sym-DDG. Especially, due to the obvious inclusions  $\text{Sym-DDG} \supseteq \text{DDG}$ , and  $\text{Sym-FBDD} \supseteq \text{FBDD}$ , we have that  $\text{Sym-DDG} \leq_s \text{DDG}$ , and  $\text{Sym-FBDD} \leq_s \text{FBDD}$ . This implies that:

**Proposition 3.**  $\text{Sym-DDG} <_s^* \text{DDG}$  and  $\text{Sym-FBDD} <_s^* \text{FBDD}$ .

## 4 A CNF-TO-Sym-DDG<sub>X,Y</sub> COMPILER

Our compiler  $\text{symddg}_{X,Y}$  (see Algorithm 1) is essentially a CNF-to-DDG compiler, enriched with some modules for symmetry handling.

The two base cases are when the input CNF formula  $\Delta$  is the empty set of clauses (line 1) or contains the empty clause (line 2). In the first case,  $\Delta$  is equivalent to  $\top$  and a Sym-DDG<sub>X,Y</sub> reduced to the  $\top$ -node is returned. In the second case,  $\Delta$  is equivalent to  $\perp$  and a Sym-DDG<sub>X,Y</sub> reduced to the  $\perp$ -node is returned. Otherwise the connected components  $\Delta_1, \dots, \Delta_k$  of the constraint graph of  $\Delta$  are looked for (line 3); if  $\Delta$  has more than one connected component (line 4), then the root node of the resulting Sym-DDG<sub>X,Y</sub> formula is a  $\wedge$ -node (generating using function  $\text{anode}$ ) and its children are obtained by calling  $\text{symddg}_{X,Y}$  recursively on  $\Delta_1, \dots, \Delta_k$ .<sup>6</sup> Then, using a method  $\text{findSymmetry}$  for symmetry detection, we look whether there is an admissible permutation  $\sigma$  w.r.t.  $\Delta$  such that  $\sigma(\Delta)$  has already been encountered and is in the cache (line 5); if so, the Sym-DDG<sub>X,Y</sub> formula associated with  $\sigma(\Delta)$  is returned, and  $\sigma$  is associated with the arc which has been followed to reach the root node of the current formula  $\Delta$  or to the root node of the initial CNF formula at the first call.<sup>7</sup>  $\sigma$  is admissible w.r.t.  $\Delta$  when all the requirements imposed by Sym-DDG<sub>X,Y</sub> are met, i.e.,  $\sigma$  satisfies the stability condition,  $\sigma(\Delta)$  is read-once and satisfies the precedence condition on  $Y$ , and  $\forall x \in X, \sigma(x) = x$  (which is enough for ensuring that  $\sigma(\Delta)$  satisfies the symmetry-freeness condition on  $X$ ). Finally, in the remaining case (line 6), one chooses a variable from  $\Delta$ ; the root node of the resulting Sym-DDG<sub>X,Y</sub> formula is a decision node labeled with  $x$ , and its two children are obtained by calling  $\text{symddg}_{X,Y}$  recursively on  $\Delta$  conditioned by  $\neg x$ , and  $\Delta$  conditioned

<sup>6</sup> Removing lines 3 and 4 in the pseudo-code of  $\text{symddg}_{X,Y}$  leads to downsize it as a CNF-to-Sym-FBDD<sub>X,Y</sub> compiler.

<sup>7</sup> For the sake of clarity, this is not detailed in the pseudo-code. Also, though not explicitly indicated in the algorithm, each time a Sym-DDG<sub>X,Y</sub> formula is generated, it is added to the cache when it is not already in it.

---

**Algorithm 1:**  $\text{symddg}_{X,Y}(\Delta)$ 

---

input : a CNF formula  $\Delta$ , a set  $X$  of symmetry-free variables,  
and a set  $Y$  of top variables  
output: a  $\text{Sym-DDG}_{X,Y}$  formula equivalent to  $\Delta$

- 1 if  $\Delta$  is empty then return  $\text{leaf}(\top)$ ;
- 2 if  $\Delta$  contains an empty clause then return  $\text{leaf}(\perp)$ ;
- 3 let  $\Delta_1, \dots, \Delta_k$  be the connected components of  $\Delta$ ;
- 4 if  $k > 1$  then return  
     $\text{anode}(\text{symddg}_{X,Y}(\Delta_1), \dots, \text{symddg}_{X,Y}(\Delta_k))$ ;
- 5 if  $(\sigma \leftarrow \text{findSymmetry}(\Delta))$  such that  
     $(\text{key} \leftarrow \text{inCache}(\sigma(\Delta))) \neq \text{nil}$  then return  $\text{cache}(\text{key})$ ;
- 6 choose a variable  $x$  of  $\Delta$ ;  
    return  $\text{dnode}(x, \text{symddg}_{X,Y}(\Delta \mid_{x \leftarrow 0}), \text{symddg}_{X,Y}(\Delta \mid_{x \leftarrow 1}))$

---

by  $x$ . Specifically,  $x$  is chosen thanks to the VSADS heuristic function [12], adapted to ensure that the constraint on top variables  $Y$  is satisfied.

The key issue in the design of our  $\text{symddg}_{X,Y}$  compiler lies in an efficient implementation of the  $\text{findSymmetry}$  method for retrieving a formula  $\Delta'$  in the cache that is equivalent to the input formula  $\Delta$ , modulo an admissible symmetry  $\sigma$ . To this point, the problem of determining whether there exists a symmetry between two CNF formulae  $\Delta$  and  $\Delta'$  can be reduced to a graph isomorphism problem for which, unfortunately, no polynomial-time algorithm is known [6]. In the present study, this computational issue is circumvented using an *incomplete* method for detecting symmetries.

Specifically,  $\text{findSymmetry}$  is based on a two-stage filtering technique followed by a greedy search in the filtered space of permutations. In order to rapidly explore the cache, the first stage compares formulae according to their canonical *signature*. The signature of a CNF expression  $\Delta$  is given by two sorted vectors, which respectively encode the signature of the variables occurring in  $\Delta$  and the signature of the clauses occurring in  $\Delta$ . The signature of a variable  $x$  is given by a pair  $(p_x, n_x)$  where  $p_x$  (resp.  $n_x$ ) is the number of literals that occur positively (resp. negatively) in  $\Delta$ . The signature of a clause  $\delta$  is simply given by its size (i.e., its number of literals). Both vectors are sorted in increasing order of their entries. Based on this encoding, two CNF formulae  $\Delta$  and  $\Delta'$  with different signatures cannot be equivalent modulo an admissible symmetry.

During the second stage, the task of identifying an admissible symmetry between two comparable formulae  $\Delta$  and  $\Delta'$  is cast as a constraint satisfaction problem (CSP). The set of variables of the CSP is given by the collection of variables occurring in  $\Delta$ , which are renamed for convenience. The domain of each (renamed) variable  $x$  is formed by the set of all literals  $\ell$  occurring in  $\Delta'$ , such that  $x$  and  $\ell$  have the same signature.<sup>8</sup> Finally, a binary constraint  $x \neq y$  is associated with each pair of variables  $x, y$  occurring in the same clause  $\delta$  of  $\Delta$ . Based on this representation, the space of candidate permutations between  $\Delta$  and  $\Delta'$  is refined by enforcing arc-consistency in the CSP. An admissible permutation is searched in a greedy way by iteratively pruning values from the variable with largest domain, and propagating these unary constraints in the network.

## 5 EXPERIMENTAL RESULTS

We focus on the application of  $\text{Sym-DDG}_{X,Y}$  to planning, an area wherein symmetries naturally occur (see e.g. [11]). Given a time

horizon  $N$ , our objective is to compile a deterministic planning problem  $P = (F, O, I, G)$ , where the initial state  $I$  and the goal  $G$  vary. Here,  $F$  is a finite set of fluents,  $O$  is a set of deterministic actions with possibly conditional effects,  $I$  is a complete truth assignment of initial fluents in  $F$ , and  $G$  is a partial assignment of final fluents in  $F$  representing the goal situation. A plan  $\pi$  for  $P$  is a sequence  $\pi$  of sets of actions, one per time point between 0 and  $N - 1$ , which maps the initial state  $I$  to a goal state (i.e., a model of  $G$ ).

In order to compile  $P$ , we first encode a description of  $P$  into a corresponding CNF theory  $\Delta_P$  over the set of variables  $PS = (\bigcup_{i=0}^N \{f_i \mid f \in F\}) \cup \{a_i \mid a \in O, i = 0, \dots, N - 1\}$ .  $\Delta_P$  can be viewed as a compact representation of the transition model associated with  $O$ . In this encoding,  $f_i$  is true if and only if fluent  $f$  holds at time point  $i$ , and  $a_i$  is true if and only if action  $a$  holds at time point  $i$ . Since only deterministic actions are considered in  $O$ , the truth value of every fluent  $f_i$  ( $f \in F, i \in 1, \dots, N$ ) is fully determined in  $\Delta_P$  as soon as the truth values of the variables  $\{f_0 \mid f \in F\} \cup \bigcup_{i=0}^{N-1} \{a_i \mid a \in A\}$  are fixed (i.e., as soon as the initial state and the plan under consideration are specified).

Based on this encoding, the formula  $\Delta_P$  is compiled into a  $\text{Sym-DDG}_{X,Y}$  formula where  $X = \{f_0 \mid f \in F\} \cup \{f_N \mid f \in F\}$  and  $Y = PS$ . Thus, the permutation group for the target class is defined over all action variables and all fluent variables that exclude initial and goal descriptions. Once this compiled form has been computed, one can take advantage of the set of queries and transformations offered by  $\text{Sym-DDG}_{X,Y}$  to address in a computationally efficient way a number of issues which are NP-hard in the general case. Thus, since  $\text{Sym-DDG}_{X,Y}$  satisfies both **X-RCD** and **CO**, we can determine in polynomial time whether a plan  $\pi$  exists for any  $I$  and  $G$  given on-line; since  $\text{Sym-DDG}_{X,Y}$  satisfies **CT**, we can also count how many  $\pi$  exist in polynomial time.

The instances we selected cover a range of different planning benchmarks, with varying horizon length. “blocks- $n$ ” refers to the famous blocks-world domain with  $n$  blocks. “bomb- $m$ - $n$ ” is another popular domain involving  $m$  bombs,  $n$  toilets, and 2 actions. “comm- $m$ - $n$ ” is an IPC5 problem about communication signals with  $m$  stages,  $n$  packets, and 5 actions. “emptyroom- $n$ ” is about navigating a robot in an  $n \times n$  empty grid. Finally, “safe- $n$ ” is about opening a safe with  $n$  possible combinations.

All instances described in PDDL were translated into CNF theories using the DIMACS format, and then compiled according to three target languages: d-DNNF generated by the standard **c2d** compiler,<sup>9</sup> DDG generated by our  $\text{symddg}_{X,Y}$  compiler without symmetry detection, and  $\text{Sym-DDG}$  targeted by the full power of our compiler. Our experiments have been conducted on a Xeon E5-2643 (3.30GHz) with 7.6GB of memory. A time-out of one hour per instance has been considered for the compilation phase; for space reasons we do not provide detailed compilation times but it is worth noting that they remain reasonable: on the instances above, the mean (resp. maximum) compilation time of  $\text{symddg}_{X,Y}$  was 22.5 seconds (resp. 109.42 seconds, obtained for the bomb-20-05 instance).

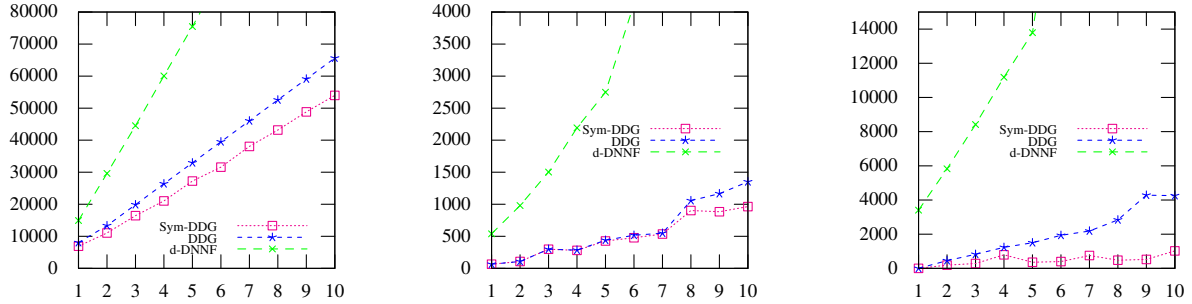
Table 3 presents the compilation results obtained from instances for which the horizon  $N$  was fixed to 5. “Nodes” (resp. “arcs”) refer to the numbers of nodes (resp. arcs) in the compiled representations. The last two columns give the percentage of size reduction achieved by  $\text{Sym-DDG}$  compared with DDG. Figure 1 plots the size (nodes + arcs) of the compiled formula versus the horizon length (from 1 to 10) for three of the test instances. Since the running time of on-line queries and transformations is governed by the size of the com-

<sup>8</sup> In an obvious way, the signature of  $\ell$  is  $(p_x, n_x)$  if  $\ell$  is the positive literal  $x$ , and  $(n_x, p_x)$  if  $\ell$  is the negative literal  $\neg x$ .

<sup>9</sup> **c2d**, available at [reasoning.cs.ucla.edu/c2d/](http://reasoning.cs.ucla.edu/c2d/), was run using the command `c2d -in file.cnf -dt.count 50 -smooth.all`

Instance	CNF		d-DNNF		DDG		Sym-DDG		% reduction	
	vars	clauses	nodes	arcs	nodes	arcs	nodes	arcs	nodes	arcs
blocks-2	406	1901	4679	37892	7332	15130	6792	14044	7.4	7.2
blocks-3	804	4343	157292	2357558	1208463	2598572	950142	2039622	<b>21.4</b>	<b>21.5</b>
bomb-5-1	564	1086	2745	4639	501	1194	352	896	<b>29.7</b>	<b>25.0</b>
bomb-5-5	1340	2610	6987	12233	1433	4290	984	3392	<b>31.3</b>	<b>20.9</b>
bomb-10-5	4680	9220	15324	28679	3273	12435	2224	10337	<b>32.1</b>	<b>16.9</b>
bomb-10-10	2760	5415	26730	48736	5863	27080	3964	23282	<b>32.4</b>	<b>14.0</b>
bomb-20-05	7100	14025	41469	76308	9203	50100	6204	44102	<b>32.6</b>	<b>12.0</b>
comm-5-2	781	2289	28292	143118	48978	108252	33346	73221	<b>31.9</b>	<b>32.4</b>
comm-6-3	1275	4032	89904	502804	132007	295113	69415	156382	<b>47.4</b>	<b>47.0</b>
emptyroom-4	188	584	731	2017	147	292	143	284	2.7	2.7
emptyroom-8	396	1292	11069	126398	8661	17432	8513	17136	1.7	1.7
safe-5	86	171	433	1061	73	163	35	87	<b>52.1</b>	<b>46.6</b>
safe-10	166	356	869	2927	191	420	40	99	<b>79.1</b>	<b>76.4</b>
safe-30	486	1346	2530	11262	451	1048	100	259	<b>77.8</b>	<b>75.3</b>

**Table 3.** Results for instances with horizon  $N = 5$ . Reduction gains above 10% are shown in boldface.



**Figure 1.** Results for Bomb-10-10 (left), Emptyroom-4 (middle), and Safe-30 (right) with varying horizon length.

pled representations, we can observe that both DDG and Sym-DDG, targeted by our compiler, are competitive with respect to d-DNNF, compiled using c2d. Furthermore, the experiments revealed that for many instances, a significant reduction of size in resulting diagrams is achieved by exploiting symmetries.

## 6 CONCLUSION

In this paper, we have shown how symmetries can be exploited for achieving significant space savings in a knowledge compilation perspective. We introduced two new languages  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  which generalize respectively the languages FBDD and DDG of decision diagrams. We have analyzed  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  both from a theoretical standpoint (following the lines of the knowledge compilation map) and from a practical standpoint (by compiling some planning benchmarks into them). The obtained results show  $\text{Sym-FBDD}_{X,Y}$  and  $\text{Sym-DDG}_{X,Y}$  as attractive; indeed, both languages offer sufficiently many queries and transformations for enabling efficient on-line reasoning for a number applications; furthermore, they achieve a high level of succinctness.

This work calls for many perspectives. One of them consists in taking advantage of complete methods for detecting symmetries, such as nauty [10, 1] and saucy [2]. While such methods are likely to lead to much longer off-line compilation times than our incomplete findSymmetry procedure, they are susceptible to explore the full symmetry group  $\Sigma$ , hence to provide smaller representations. Another perspective consists in studying the connections between  $\text{Sym-DDG}_{X,Y}$  and the language of first-order NNF circuits.

## REFERENCES

- [1] F. Aloul, K. Sakallah, and I. Markov, 'Efficient symmetry breaking for boolean satisfiability', in *Proc. of IJCAI'03*, pp. 271–276, (2003).
- [2] P. Darga, M. Liffiton, K. Sakallah, and I. Markov, 'Exploiting structure in symmetry detection for CNF', in *Proc. of DAC'04*, pp. 530–534, (2004).
- [3] A. Darwiche, 'Decomposable negation normal form', *Journal of the ACM*, **48**(4), 608–647, (2001).
- [4] A. Darwiche and P. Marquis, 'A knowledge compilation map', *Journal of Artificial Intelligence Research*, **17**, 229–264, (2002).
- [5] H. Fargier and P. Marquis, 'On the use of partially ordered decision graphs in knowledge compilation and quantified Boolean formulae', in *Proc. of AAAI'06*, pp. 42–47, (2006).
- [6] M.R. Garey and D.S. Johnson, *Computers and intractability: a guide to the theory of NP-completeness*, Freeman, 1979.
- [7] J. Gergov and C. Meinel, 'Efficient analysis and manipulation of OBDDs can be extended to FBDDs', *IEEE Transactions on Computers*, **43**(10), 1197–1209, (1994).
- [8] A. Haken, 'The intractability of resolution', *Theoretical Computer Science*, **39**, 297–308, (1985).
- [9] B. Krishnamurthy, 'Short proofs for tricky formulas', *Acta Informatica*, **22**, 253–275, (1985).
- [10] B. D. McKay, 'Practical graph isomorphism', *Congressus Numerantium*, **30**, 45–57, (1981).
- [11] H. Palacios, B. Bonet, A. Darwiche, and H. Geffner, 'Pruning conformant plans by counting models on compiled d-DNNF representations', in *Proc. of ICAPS'05*, pp. 141–150, (2005).
- [12] T. Sang, P. Beame, and H. Kautz, 'Heuristics for fast exact model counting', in *Proc. of the 8th International Conference on Theory and Applications of Satisfiability Testing (SAT'05)*, pp. 226–240, (2005).
- [13] B. Selman and H.A. Kautz, 'Knowledge compilation and theory approximation', *Journal of the ACM*, **43**, 193–224, (1996).
- [14] A. Urquhart, 'The symmetry rule in propositional logic', *Discrete Applied Mathematics*, **96/97**, 177–193, (1999).