

On CNF Encodings of Decision Diagrams

Ignasi Abío¹, Graeme Gange³, Valentin Mayer-Eichberger^{1,2(✉)},
and Peter J. Stuckey^{1,3}

¹ NICTA, West Melbourne, Australia

{ignasi.abio,valentin.mayer-eichberger,peter.stuckey}@nicta.com.au

² University of New South Wales, Sydney, Australia

³ University of Melbourne, Melbourne, Australia

gkgange@unimelb.edu.au

Abstract. Decisions diagrams such as Binary Decision Diagrams (BDDs), Multi-valued Decision Diagrams (MDDs) and Negation Normal Forms (NNFs) provide succinct ways of representing Boolean and other finite functions. Hence they provide a powerful tool for modelling complex constraints in discrete satisfaction and optimization problems. Generic propagators for these global constraints exist, but they are complex and hard to implement. An alternative approach to making use of them for solving is to encode them to CNF, using SAT style solving technology to implement them efficiently. This may also have advantages since it is naturally incremental and exposes intermediate literals which may well be useful as search decisions for solving the problem.

In this paper we explore different ways that we can map these constraints to CNF, and the different properties these mappings maintain. Surprisingly the most used encoding of BDDs does not maintain domain consistency in arbitrary BDDs. We also consider the strength of propagation with respect to the intermediate literals. We give experiments which compare the performance of the different encodings.

1 Introduction

Decisions diagrams such as Binary Decision Diagrams (BDDs), Multi Decision Diagrams (MDDs) and Negation Normal Forms (NNFs) provide succinct ways of representing Boolean and other finite functions. Hence they provide a powerful tool for modelling complex constraints in discrete satisfaction and optimization problems.

Constraint programming solvers include generic propagators for propagating constraints represented by BDDs [16], MDDs [8] and NNFs [15], since they are highly flexible, and hence useful in many different models. But these propagators are complex and hard to implement.

An alternative approach to making use of them for solving is to encode them to CNF, using SAT style solving technology to implement them efficiently. If the remainder of the problem is naturally modelled in CNF then this allows a SAT solver to tackle the problem.

A SAT encoding may also be preferable within a CP solver, as it avoids the need for implementing complex propagators, is naturally incremental, and exposes intermediate literals as candidates for search and learning. A good encoding is critical in lazy decomposition approaches [1], where a propagator that participates in many conflicts is replaced by a CNF decomposition during runtime.

In this paper we explore different approaches for encoding decision diagrams to CNF.¹ The contributions of this paper are:

- An investigation of a large design space for encoding decision diagrams.
- We clarify the picture of BDD/MDD/NNF encodings, analyse their propagation strength and correct some misunderstandings in the literature.
- We introduce an encoding of BDDs and MDDs where unit propagation implements propagation completeness.
- Experiments which compare the performance of the different encodings.

2 Preliminaries

2.1 SAT Solving

We denote the Boolean value true by \top and false by \perp .

Let $\mathcal{Y} = \{y_1, y_2, \dots\}$ be a fixed set of propositional *variables*. If $y \in \mathcal{Y}$ then y and $\neg y$ are *positive* and *negative literals*, respectively. The *negation* of a literal l , written $\neg l$, denotes $\neg y$ if l is y , and y if l is $\neg y$. A *clause* is a disjunction of literals $\neg y_1 \vee \dots \vee \neg y_p \vee y_{p+1} \vee \dots \vee y_n$, sometimes written as $y_1 \wedge \dots \wedge y_p \rightarrow y_{p+1} \vee \dots \vee y_n$. A *CNF formula* F is a conjunction of clauses.

A set of literals A is *contradictory* if $\exists y. \{y, \neg y\} \subset A$. A (partial) *assignment* A is a set of literals which is not contradictory. A literal l is *true* in A if $l \in A$, is *false* in A if $\neg l \in A$, and is *undefined* in A otherwise. An *extension* of an assignment A is an assignment A' where $A' \supset A$. A *complete assignment* is an assignment with no undefined literals. Given a partial assignment A , a *completion* of A is an extension of A which is a complete assignment.

A complete assignment A satisfies formula ϕ if replacing each y in ϕ which is true in A with \top and replacing each y in ϕ which is false in A with \perp gives an expression which evaluates to \top . A partial assignment A satisfies formula ϕ , written $A \models \phi$ if every completion of A satisfies ϕ .

Systems that decide whether a CNF formula F has any model are called SAT solvers, and the main inference rule they implement is *unit propagation*: given a CNF F and an assignment A , find a clause in F such that all its literals are false in A except at most one, say l , which is undefined, add l to A and repeat the process until reaching a fix-point. See e.g. [21] for more details.

For some set of clauses C , we shall use $UP_C(A)$ to denote the set of literals inferred by unit propagation on C starting from assignment A . We will omit the C subscript when clear from the context. Note that $UP_C(A)$ may be contradictory, in which case unit propagation has detected unsatisfiability.

¹ A longer version of this paper including proofs of all Theorems can be found at <http://people.eng.unimelb.edu.au/pstuckey/mddenc.pdf>.

2.2 Propositional Encodings

Problems of interest rarely (if ever) begin in CNF form. Boolean formulae ϕ must be first converted into some equisatisfiable conjunction of clauses F_ϕ . The seminal work here is the Tseytin transformation [25], later refined by Plaisted and Greenbaum [22], which introduces a variable for each sub-formula and adds clauses to enforce the semantics of each connective.

While equisatisfiability is sufficient for correctness, the choice of decomposition can have a great impact on solver performance. A major consideration here is *propagation strength* – that is, given some partial assignment A and formula ϕ , what can be said of $UP_{F_\phi}(A)$.

There are a number of properties we may wish of F_ϕ .

- An encoding F_ϕ for a formula ϕ is *correct* if any complete assignment A on $vars(\phi)$ where $A \models \phi$, then A has an extension satisfying F_ϕ , and any complete assignment $A \models \neg\phi$ has no extension satisfying F_ϕ .
- An encoding F_ϕ for a formula ϕ *implements consistency* if for every assignment A over $vars(\phi)$ where $A \models \neg\phi$, then $UP_{F_\phi}(A)$ is contradictory.
- An encoding F_ϕ for a formula ϕ *implements domain consistency* when for each literal l over $vars(\phi)$, if $A \models \phi \rightarrow l$ then $l \in UP_{F_\phi}(A)$.
- An encoding F_ϕ for a formula ϕ *implements unit refutation completeness* [26] (also called *SLUR* [19]) when for assignment B over $vars(F_\phi)$ where $B \models \neg F_\phi$, then $UP_{F_\phi}(B)$ is contradictory.
- An encoding F_ϕ for a formula ϕ *implements propagation completeness* [6, 19] when for each literal l over $vars(F_\phi)$, $B \models F_\phi \rightarrow l$ then $l \in UP_{F_\phi}(B)$.

Another important consideration is the encoding size. In general, smaller encodings are more efficient than larger ones, if both have the same propagation strength.

2.3 At-most-one and Exactly-one Constraints

Given a set of literals l_1, \dots, l_n , the *At-most-one* (AMO) constraint over these literals is defined as $l_1 + l_2 + \dots + l_n \leq 1$.

There are several ways to encode AMO into SAT [3, 7, 14]. Here, we consider the ladder encoding. It introduces variables $\{a_i := l_1 \vee \dots \vee l_i \mid 1 \leq i < n\}$ and clauses $\{a_i \rightarrow a_{i+1}, l_i \rightarrow a_i, l_{i+1} \rightarrow \neg a_i\}$. It is easy to see that this encoding is propagation complete.

Given a set of literals l_1, \dots, l_n , the *Exactly-one* (EO) constraint over these literals is defined as $l_1 + l_2 + \dots + l_n = 1$. Notice that

$$EO(\{l_1, \dots, l_n\}) = AMO(\{l_1, \dots, l_n\}) \wedge (l_1 \vee \dots \vee l_n)$$

This defines a propagation complete encoding for EO given a propagation complete encoding of AMO.

2.4 Direct Encoding for Integer Variables

There are different methods for encoding integer variables into SAT (see for instance [18, 27]). In this paper we use the direct encoding.

Let x be an integer variable with domain $[a, b]$. The *direct encoding* introduces Boolean variables $\llbracket x = i \rrbracket$ for $a \leq i \leq b$. A variable $\llbracket x = i \rrbracket$ is true iff $x = i$. The encoding also introduces the constraint $\text{EO}(\{\llbracket x = i \rrbracket \mid a \leq i \leq b\})$.

We will sometimes treat Boolean variables b as integers with domain $[0, 1]$.

We will implicitly assume that the direct encoding clauses $\text{EO}(\{\llbracket x = i \rrbracket \mid a \leq i \leq b\})$ are part of any encoding of formula using integers x . We also assume all assignments A are closed under unit propagation of these clauses.

We extend the notion of satisfaction to formulae involving integer variables, as follows. A complete assignment A satisfies ϕ if replacing each Boolean variable as before, and each integer variable x_i by j if $\llbracket x_i = j \rrbracket \in A$ (since $A \models \text{EO}(\{\llbracket x_i = j \rrbracket \mid a \leq j \leq b\})$ there must be exactly one) and evaluating the resulting ground expression gives \top . We extend the notation $A \models \phi$ as before.

2.5 Multi-valued Decision Diagrams

A directed acyclic graph \mathcal{M} is called an *ordered Multi-valued Decision Diagram (MDD)* if it satisfies the following properties:

- It has two terminal nodes, namely \mathcal{T} (true) and \mathcal{F} (false).
- Each non-terminal node is labeled by an integer variable $\{x_1, x_2, \dots, x_n\}$. This variable is called *selector variable*.
- Every node labeled by x_i has the same number of outgoing edges, namely $b_i - a_i + 1$, where $[a_i, b_i]$ is the domain of x_i .
- If an edge connects a node with a selector variable x_i and a node with a selector variable x_j , then $j > i$.

The MDD is *quasi-reduced* if no isomorphic subgraphs exist. It is *reduced* if, moreover, no nodes with only one child exist. A *long edge* is an edge connecting two nodes with selector variables x_i and x_j such that $j > i + 1$. In the following we only consider quasi-reduced ordered MDDs without long edges, and we just refer to them as MDDs for simplicity.² We refer to [24] for further details about MDDs.

Given an MDD \mathcal{M} we use ρ to refer to its *root node*. Given a node $\nu \in \mathcal{M}$, we write $\text{var}(\nu) = x_j$ when node ν is labelled by x_j . Given an edge $\varepsilon \in \mathcal{M}$, we write $\varepsilon = \text{edge}(\nu, \mu, \llbracket x_i = j \rrbracket)$ if ε joins the node ν and μ when $x_i = j$.

An MDD represents a formula over integer variables: a MDD node ν with selector x with domain $[a, b]$ and children $\nu_a, \nu_{a+1}, \dots, \nu_b$ represents the formula ϕ_ν where

$$\phi_\nu \equiv \bigvee_{i \in [a, b]} x = i \wedge \phi_{\nu_i}$$

where ϕ_{ν_i} is the formula represented by node ν_i , and $\phi_{\mathcal{T}} = \top$ and $\phi_{\mathcal{F}} = \perp$.

² Notice, however, that every result in this paper holds for non-reduced MDDs without long edges, and with some modifications of the rules the results also extend to non-reduced MDDs with long edges.

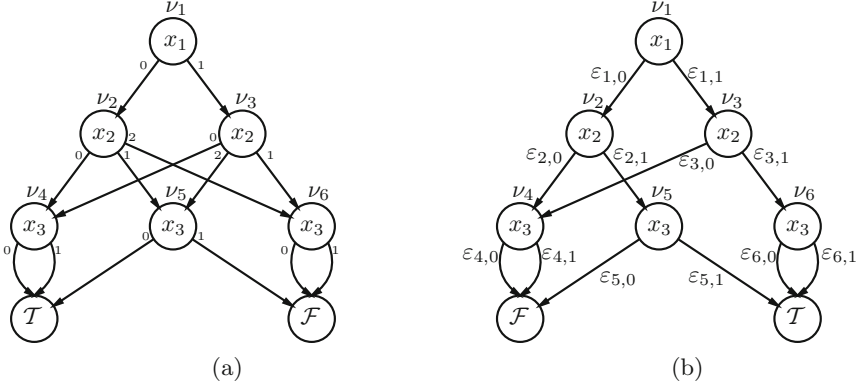


Fig. 1. (a) MDD of $x_2 = 0 \vee (x_3 = 0 \wedge x_2 - x_1 = 1)$ and (b) BDD of $x_2 \wedge (x_1 \vee x_3)$

Example 1. Let us consider the MDD encoding of $x_2 = 0 \vee (x_3 = 0 \wedge x_2 - x_1 = 1)$, with $x_1, x_3 \in \{0, 1\}$ and $x_2 \in \{0, 1, 2\}$, shown in Fig. 1(a). In this case $\rho = \nu_1$, $\text{var}(\nu_3) = x_2$, and the rightmost edge from ν_3 is $\text{edge}(\nu_3, \nu_6, x_2 = 1)$. $\phi_{\nu_4} \leftrightarrow \top$, $\phi_{\nu_5} \leftrightarrow x_3 = 0$, $\phi_{\nu_6} \leftrightarrow \perp$, and hence $\phi_{\nu_2} \leftrightarrow (x_2 = 0 \wedge \top) \vee (x_2 = 1 \wedge x_3 = 0) \vee (x_2 = 2 \wedge \perp)$ or equivalently $\phi_{\nu_2} \leftrightarrow x_2 = 0 \vee (x_2 = 1 \wedge x_3 = 0)$. \square

A *binary decision diagram (BDD)* is an MDD with only Boolean variables. For a BDD \mathcal{M} we can consider a non-terminal node ν as a triple (x, t, f) where there are two outgoing edges $\text{edge}(\nu, t, x)$ and $\text{edge}(\nu, f, \neg x)$. The BDD node ν represents the formula $\phi_\nu \equiv \text{ITE}(x, \phi_t, \phi_f)$ or equivalently $(x \wedge \phi_t) \vee (\neg x \wedge \phi_f)$.

2.6 Negation Normal Form Formulae

A *negation normal form* formula (NNF) is a rooted, directed acyclic graph (DAG) where each leaf node is labeled with x or $\neg x$ and each internal node is labeled with \wedge or \vee and can have arbitrarily many children.

NNFs are a more general form of decision diagram than BDDs, and can be exponentially more compact to represent the same formula [11]. We can use NNFs to express formulae over finite domain integer variables using the direct encoding.

But NNFs in general are too expressive, so usually we require some additional properties, such as:

Decomposable. An NNF \mathcal{N} is *decomposable* if for each conjunction ϕ in \mathcal{N} , the conjuncts of ϕ do not share variables. That is, if ϕ_1, \dots, ϕ_n are the children of and-node ϕ , then $\text{vars}(\phi_i) \cap \text{vars}(\phi_j) = \emptyset$ for $i \neq j$.

Deterministic. An NNF \mathcal{N} is *deterministic* if for each disjunction ϕ in \mathcal{N} , each two disjuncts of ϕ are logically contradictory. That is, if ϕ_1, \dots, ϕ_n are the children of or-node ϕ , then $\phi_i \wedge \phi_j \models \perp$ for $i \neq j$.

Smooth. An NNF \mathcal{N} is *smooth* if for each disjunction ϕ in \mathcal{N} , each disjunct of ϕ mentions the same variables. That is, if ϕ_1, \dots, ϕ_n are the children of or-node ϕ , then $\text{vars}(\phi_i) = \text{vars}(\phi_j)$ for $i \neq j$.

3 Encoding MDDs

3.1 Encoding BDDs

The BDD encoding of MiniSat+ [13] is defined as follows: For each non-terminal BDD node $\nu = (x, t, f)$ we generate a Boolean variable ν which represents the truth value of the BDD rooted at ν .

For each non-terminal node $\nu = (x, t, f)$, we generate the following clauses:

$$\begin{array}{ll} \text{B1 } t \wedge x \rightarrow \nu. & \text{B4 } \neg f \wedge \neg x \rightarrow \neg \nu. \\ \text{B2 } \neg t \wedge x \rightarrow \neg \nu. & \text{B5 } t \wedge f \rightarrow \nu. \\ \text{B3 } f \wedge \neg x \rightarrow \nu. & \text{B6 } \neg t \wedge \neg f \rightarrow \neg \nu. \end{array}$$

Define encoding MiniSAT as B1–B6, together with the terminal and root clauses: \mathcal{T} (the true terminal is true), $\neg\mathcal{F}$ (the false terminal is false) and ρ (the root of the tree must be true).

Note while Een and Sorensen [13] refer to this as a Tseytin encoding, it is not since Tseytin [25] does not include an ITE constructor, so in the Tseytin encoding $\text{ITE}(x, t, f)$ needs to be encoded as $(x \wedge t) \vee (\neg x \wedge f)$.

The encoding contains $O(s)$ variables and clauses, where s is the size of the BDD.

Een and Sorensen [13] show that this encoding maintains domain consistency when used to encode (sorted) pseudo-Boolean constraints.

Theorem 1 ([13]). *Unit propagation on the MiniSAT encoding for a BDD for pseudo-Boolean constraint $\sum_{i=1}^n c_i x_i \geq d$ maintains domain consistency, assuming the coefficients c_i are in non-increasing order.* \square

This theorem does not hold without the ordering criterion. Consider the BDD encoding $x_1 + 2x_2 + x_3 \geq 3$ (or equivalently $x_2 \wedge (x_1 \vee x_3)$) shown in Fig. 1(b). Any solution of the BDD requires x_2 is \top . Unit propagation on the MiniSAT encoding generates $\neg\mathcal{F}, \mathcal{T}, \nu_1, \neg\nu_4, \nu_6$ and nothing else.

Theorem 2. *Unit propagation on the clauses (B2), (B4), (B6), $\neg\mathcal{F}$, ρ for a BDD maintains consistency.* \square

All in all, the encoding is compact (especially if only clauses (B2), (B4), (B6), $\neg\mathcal{F}$ and ρ are used), but the propagation strength is low.

3.2 Encodings MDDs with One Variable Per Node

The first set of encodings for MDDs, used for example in [2], are generalizations of the MiniSat+ encoding. This is natural since they are also used to encode pseudo-Boolean and linear constraints.

For each node ν at level i , with children $\nu_{a_i}, \nu_{a_i+1}, \dots, \nu_{b_i}$, where the domain of x_i is $[a_i, b_i]$.

- M1 $\neg\nu_j \wedge \llbracket x_i = j \rrbracket \rightarrow \neg\nu$ (generalizes B2 and B4).
M2 $\nu_j \wedge \llbracket x_i = j \rrbracket \rightarrow \nu$ (generalizes B1 and B3).
M3 $\nu_{a_i} \wedge \nu_{a_i+1} \wedge \cdots \wedge \nu_{b_i} \rightarrow \nu$ (weakly generalizes B5).
M4 $\neg\nu_{a_i} \wedge \neg\nu_{a_i+1} \wedge \cdots \wedge \neg\nu_{b_i} \rightarrow \neg\nu$ (weakly generalizes B6).

With these clauses, we can define different encodings:

Minimal: Clauses M1, $\neg\mathcal{F}$, ρ .

GenMiniSAT: Clauses M1–M4, \mathcal{T} , $\neg\mathcal{F}$, ρ .

Minimal is very compact, but its propagation strength is low, moreover when the original variables are fixed it does not necessarily fix all the node variables, and hence does not preserve solution counts. GenMiniSAT is the natural generalization of the BDD encoding from [13] to MDDs. Again, it is not the Tseytin encoding [25] of the MDD. Both encodings use $O(s)$ variables and $O(sd)$ clauses, where s is the MDD size and d is the maximum domain size of variables x .

Proposition 1. *Let $A = \{\llbracket x_i = v_i \rrbracket \mid 1 \leq i \leq n\}$ be a complete assignment on variables x satisfying the MDD \mathcal{M} . Then, there exists a complete assignment $B \supset A$ over the variables x, ν satisfying clauses GenMiniSAT.* \square

Proposition 2. *Let $A = \{\llbracket x_i = v_i \rrbracket \mid 1 \leq i \leq n\}$ be a complete assignment on variables x not satisfying the MDD \mathcal{M} , then clauses ρ and M1 propagate \mathcal{F} .* \square

Corollary 1. *Minimal and GenMiniSAT are correct.* \square

These two encodings, however, do not detect inconsistencies:

Example 2. Consider again the MDD of $x_2 = 0 \vee (x_3 = 0 \wedge x_2 - x_1 = 1)$, with $x_1, x_3 \in \{0, 1\}$ and $x_2 \in \{0, 1, 2\}$ shown in Fig. 1(a).

After simplification, GenMiniSAT consists of the following clauses:

$$\begin{aligned} &\neg\llbracket x_1 = 0 \rrbracket \vee \nu_2, & \neg\llbracket x_1 = 1 \rrbracket \vee \nu_3, & \nu_2 \vee \nu_3, & \neg\llbracket x_2 = 0 \rrbracket \vee \nu_2, \\ &\neg\nu_4 \vee \neg\llbracket x_2 = 1 \rrbracket \vee \nu_2, & \nu_4 \vee \neg\llbracket x_2 = 1 \rrbracket \vee \neg\nu_2, & \neg\llbracket x_2 = 2 \rrbracket \vee \neg\nu_2 & \neg\llbracket x_2 = 0 \rrbracket \vee \nu_3, \\ &\neg\nu_4 \vee \neg\llbracket x_2 = 2 \rrbracket \vee \nu_3, & \nu_4 \vee \neg\llbracket x_2 = 2 \rrbracket \vee \neg\nu_3, & \neg\llbracket x_2 = 1 \rrbracket \vee \neg\nu_3 & \neg\llbracket x_3 = 0 \rrbracket \vee \nu_4, \\ &\neg\llbracket x_3 = 1 \rrbracket \vee \neg\nu_4. \end{aligned}$$

Consider the partial assignment $A = \{\neg\llbracket x_2 = 0 \rrbracket, \neg\llbracket x_3 = 0 \rrbracket, \llbracket x_3 = 1 \rrbracket\}$. It cannot be extended to a complete assignment satisfying the MDD. However, unit propagation does not fail.

The same happens with Minimal, since it is a subset of GenMiniSAT. \square

3.3 Tseytin Encoding of an MDD

In this section we describe an alternative encodings for an MDD, the Tseytin encoding [25]. It detects inconsistencies with respect to the original variables but does not enforce domain consistency.

The Tseytin encoding introduce Boolean variables representing the formula of each edge. Let ν be a node at level i , with outgoing edges $\{\varepsilon_j \mid j \in J\}$. Let $\varepsilon = \text{edge}(\nu, \mu, \llbracket x_i = j \rrbracket)$ be an edge of \mathcal{M} , then the Boolean variable ε encoding the edge represents the formula $\llbracket x_i = j \rrbracket \wedge \phi_\mu$.

The clauses of the Tseytin encoding are, for each node ν and edge ε

- T1 $\nu \rightarrow \bigvee_j \varepsilon_j.$
 T2 $\varepsilon \rightarrow \nu.$
 T3 $\varepsilon \rightarrow \mu.$
 T4 $\varepsilon \rightarrow \llbracket x_i = j \rrbracket.$
 T5 $\mu \wedge \llbracket x_i = j \rrbracket \rightarrow \varepsilon.$

The Tseytin encoding, $\mathcal{T}_{\text{seytin}}$, consists of clauses T1–T5, \mathcal{T} , $\neg\mathcal{F}$ and ρ . Therefore, it consists in $O(sd)$ variables and clauses, where s is the MDD size and d the maximum domain size of variables x .

Proposition 3. *Let $A = \{\llbracket x_i = v_i \rrbracket \mid 1 \leq i \leq n\}$ be a complete assignment on variables x satisfying the MDD \mathcal{M} . Then, there exists a complete assignment $B \supset A$ over the variables x, ν, ε satisfying clauses $\mathcal{T}_{\text{seytin}}$. \square*

Proposition 4. *Let A be a partial assignment on variables $\{x_i, x_{i+1}, \dots, x_n\}$, and let ν be a node of \mathcal{M} at level i . Assume that there is no completion A' of A satisfying the MDD rooted at ν . Then, unit propagation on clauses $\mathcal{T}_{\text{seytin}}$ and A enforces $\neg\nu$. \square*

As a corollary, we can prove:

Theorem 3. *Tseytin is correct; i.e., given a complete assignment of the input variables, this encoding finds an inconsistency if and only if the assignment does not satisfy \mathcal{M} . Moreover, it implements consistency. \square*

However, $\mathcal{T}_{\text{seytin}}$ does not preserve domain consistency.

Example 3. Let us consider the BDD of $x_2 \wedge (x_1 \vee x_3)$, shown in Fig. 1(b). $\mathcal{T}_{\text{seytin}}$, once simplified, generates the following clauses:

$$\begin{array}{lll}
 \varepsilon_{1,0} \vee \varepsilon_{1,1}, & \neg\nu_2 \vee x_1 \vee \varepsilon_{1,0}, & \neg\varepsilon_{1,0} \vee \neg x_1, \quad \neg\varepsilon_{1,0} \vee \nu_2, \\
 \neg\nu_3 \vee \neg x_1 \vee \varepsilon_{1,1}, & \neg\varepsilon_{1,1} \vee x_1, & \neg\varepsilon_{1,0} \vee \nu_3, \quad \neg\nu_2 \vee \varepsilon_{2,1}, \\
 \neg\nu_5 \vee \neg x_2 \vee \varepsilon_{2,1}, & \neg\varepsilon_{2,1} \vee \nu_2, & \neg\varepsilon_{2,1} \vee x_2, \quad \neg\varepsilon_{2,1} \vee \nu_5, \\
 \neg\nu_3 \vee \varepsilon_{3,1}, & \neg x_2 \vee \varepsilon_{3,1}, & \neg\varepsilon_{3,1} \vee \nu_3, \quad \neg\varepsilon_{3,1} \vee x_2, \\
 \neg\nu_5 \vee \varepsilon_{5,1}, & \neg x_3 \vee \varepsilon_{5,1}, & \neg\varepsilon_{5,1} \vee \nu_5, \quad \neg\varepsilon_{5,1} \vee x_3.
 \end{array}$$

Consider the partial assignment $A = \emptyset$. Notice that x_2 is not propagated even though that there is no solution of \mathcal{M} with $\neg x_2$. Clause $x_2 \vee \varepsilon_{2,0} \vee \varepsilon_{3,0}$ would propagate x_2 . \square

Also, $\mathcal{T}_{\text{seytin}}$ does not implement unit refutation completeness:

Example 4. Consider the BDD of the constraint $\text{XOR}(x_1, x_2, x_3, x_4)$ shown in Fig. 2. Node ν_2 represents the constraint $\text{XOR}(x_2, x_3, x_4)$, and node ν_3 represents $\neg\text{XOR}(x_2, x_3, x_4)$. It is clear, therefore, that the partial assignment $B = \{\nu_2, \nu_3\}$ cannot be extended to a complete assignment satisfying \mathcal{M} . However, $\mathcal{T}_{\text{seytin}}$ does not find any conflict. \square

3.4 Path-Based Encodings

Under the encodings described in Sects. 3.2 and 3.3, the semantics of variables match the Boolean formula they represent – a node/edge variable is true (in a complete assignment) iff the corresponding formula is true.

In this section, we describe a set of *path-based* encodings. Like the Tseytin encoding these introduce one variable per node and per edge, but the interpretation of these variables is different. Under a path-based encoding, ν (or ε) is true iff the path from the root r to \mathcal{T} defined by the selector variables passes through ν (resp. ε).

Unlike the previous encodings, the variables introduced here cannot be re-used if a sub-formula occurs in multiple constraints. However, we shall see that this interpretation allows us to make much stronger inferences.

A related treatment of path-based encodings of the **regular** constraint to CNF can be found in Bacchus work in [4] and by Quimper and Walsh in [23] in context of the **grammar** constraint. Our study provides a complete analysis of such encodings for decision diagrams and introduces a novel encoding with stronger propagation properties.

We generate clauses for each node ν and connecting it to each of its outgoing edge ε_j and each of its incoming edges δ_j , as well as clauses for each edge $\varepsilon = \text{edge}(\nu, \mu, \llbracket x_i = j \rrbracket)$.

P1 $\nu \wedge \llbracket x_i = j \rrbracket \rightarrow \varepsilon_j$.

P2 $\nu \rightarrow \bigvee_j \delta_j$ where $\nu \neq \rho$.

P3 $\llbracket x_i = j \rrbracket \rightarrow \bigvee \{\varepsilon' \mid \varepsilon' = \text{edge}(\nu, \mu, \llbracket x_i = j \rrbracket) \text{ for some } \nu, \mu \in \mathcal{M}\}$.

P4 $\text{EO}(\{\nu' \in \mathcal{M} \mid \text{Level}(\nu') = i\})$.

Clauses P1 enforce that a node on the path puts its outgoing edge on the path. Clauses P2 require each node on the path (except the root) has an incoming edge. Clauses P3 require that each integer value has an edge that supports it. Clauses P4 require that exactly one node on each level is \top .

With these clauses, we can define different encodings:

BasicPath: Clauses P1–P2, T1–T4, \mathcal{T} , $\neg\mathcal{F}$, ρ .

NNFPath: BasicPath and clauses P3.

LevelPath: BasicPath and clauses P4.

CompletePath: BasicPath and clauses P3–P4.

All the encodings require $O(sd)$ variables and clauses, where s is the MDD size and d the maximum domain size of variables x .

A complete assignment A over the variables x_i defines a path in \mathcal{M} in the obvious way. This path is denoted by $\nu_1 = \rho, \varepsilon_1, \nu_2, \varepsilon_2, \dots$. By definition of the MDD, the assignment is compatible with \mathcal{M} if and only if $\nu_{n+1} = \mathcal{T}$.

A complete assignment B over variables x_i, ν, ε is compatible with \mathcal{M} if

- $A := B \cap (\{\llbracket x_i = j \rrbracket \mid 1 \leq i \leq n, j \in [a_i, b_i]\} \setminus \{\neg\llbracket x_i = j \rrbracket \mid 1 \leq i \leq n, j \in [a_i, b_i]\})$ is compatible with \mathcal{M} .

- $\nu \in B$ iff $\nu = \nu_i$ for some i on the path defined by A .
- $\varepsilon \in B$ iff $\varepsilon = \varepsilon_i$ for some i on the path defined by A .

Proposition 5. *Given a complete assignment A on the variables x compatible with \mathcal{M} , there exists a complete assignment $B \supset A$ over the variables x, ν, ε satisfying clauses *CompletePath*. \square*

Proposition 6. *Let A be a partial assignment on variables x . Let $UP(A)$ be the set of propagated literals with *BasicPath*. Let ν be a node of \mathcal{M} , and ε be an edge of \mathcal{M} . Then:*

- $\neg\nu \in UP(A)$ if $A \wedge \nu \models \neg\mathcal{M}$.
- $\neg\varepsilon \in UP(A)$ if $A \wedge \varepsilon \models \neg\mathcal{M}$. \square

Let us explain the idea behind the proof. If ν has not been propagated to false, we can create a path from ρ to \mathcal{T} passing through ν , where all the nodes of this path have not been propagated to false. This path will define a completion B satisfying \mathcal{M} with $\nu \in B$.

To build this path, we start from ν . Since $\neg\nu \notin UP(A)$, ν must have a parent that has also not been propagated to false. This node, again, has a parent that has not been propagated to false, etc. That gives a path from ρ to ν . In the same way, ν has a child that has not been propagated to false, and this child has a child that has not been propagated to false, etc. That gives a path from ν to \mathcal{T} . Concatenating both paths, we obtain the desired path from ρ to \mathcal{T} .

Theorem 4. *BasicPath maintains consistency by unit propagation. \square*

BasicPath, however, does not maintain domain consistency. For that we need clauses P3.

Example 5. Let us consider the BDD of $x_2 \wedge (x_1 \vee x_3)$, shown at Fig. 1(b). *BasicPath*, once simplified, generates the following clauses:

$$\begin{aligned}
& x_1 \vee \varepsilon_{1,0}, \quad \neg x_1 \vee \varepsilon_{1,1}, \quad \neg\nu_2 \vee x_2, \quad \neg\nu_3 \vee x_2, \\
& \neg\nu_5 \vee x_3, \quad \varepsilon_{1,0} \vee \varepsilon_{1,1}, \quad \neg\nu_2 \vee \varepsilon_{2,1}, \quad \neg\nu_3 \vee \varepsilon_{3,1}, \\
& \neg\nu_5 \vee \varepsilon_{5,1}, \quad \neg\nu_2 \vee \varepsilon_{1,0}, \quad \neg\nu_3 \vee \varepsilon_{1,1}, \quad \neg\nu_5 \vee \varepsilon_{2,1}, \\
& \varepsilon_{3,1} \vee \varepsilon_{5,1} \quad \neg\varepsilon_{2,1} \vee \nu_2, \quad \neg\varepsilon_{3,1} \vee \nu_3, \quad \neg\varepsilon_{5,1} \vee \nu_5, \\
& \neg\varepsilon_{1,0} \vee \nu_2, \quad \neg\varepsilon_{1,1} \vee \nu_3, \quad \neg\varepsilon_{2,1} \vee \nu_5, \quad \neg\varepsilon_{1,0} \vee \neg x_1, \\
& \neg\varepsilon_{1,1} \vee x_1, \quad \neg\varepsilon_{2,1} \vee x_2, \quad \neg\varepsilon_{3,1} \vee x_2, \quad \neg\varepsilon_{5,1} \vee x_3.
\end{aligned}$$

Consider the partial assignment $A = \emptyset$. Then, unit propagation does not propagate x_2 even though that there is no solution of \mathcal{M} with $\neg x_2$. Clause $x_2 \vee \varepsilon_{2,0} \vee \varepsilon_{3,0}$, from P3, would propagate x_2 . \square

As Corollary of Proposition 5 and Theorem 4, it follows that

Theorem 5. *Encodings *BasicPath*, *NNFPath*, *LevelPath* and *CompletePath* are correct; i.e., given a complete assignment of the input variables, these encodings find an inconsistency if and only if the assignment does not satisfy \mathcal{M} . \square*

Theorem 6. *NNFPath maintains domain consistency by unit propagation.* \square

NNFPath maintains domain consistency with respect to the original variables. However, since a SAT solver will not differentiate between original variables and auxiliary ones, partial assignments, in general, contain both type of variables. And, without clauses P4, the encodings are not propagation complete:




Fig. 2. BDD of $\text{XOR}(x_1, x_2, x_3, x_4)$

Example 6. Consider the MDD shown in Fig. 2, representing the constraint $\text{XOR}(x_1, x_2, x_3, x_4)$. Consider the partial assignment $B = \{\nu_4, \nu_5\}$. It is clear that B cannot be extended to a complete assignment satisfying \mathcal{M} , since no path can contain two nodes on the same level. However, NNFPath does not find any conflict. \square

To maintain consistency with respect to all variables, clauses P4 are needed. In that case, we can generalize the previous results to assignments containing auxiliary variables:

Proposition 7. *Let B be a partial assignment on all the variables. Let $UP(B)$ be the set of propagated literals with LevelPath. Let ν be a node of \mathcal{M} , and ϵ be an edge of \mathcal{M} . Then:*

1. $\neg\nu \in UP(B)$ if $B \wedge \nu \models \neg\mathcal{M}$.
2. $\neg\epsilon \in UP(B)$ if $B \wedge \epsilon \models \neg\mathcal{M}$.
3. $\nu \in UP(B)$ if $B \wedge \neg\nu \models \neg\mathcal{M}$.
4. $\epsilon \in UP(B)$ if $B \wedge \neg\epsilon \models \neg\mathcal{M}$.

Theorem 7. *LevelPath is unit refutation complete.* \square

LevelPath does not maintain domain consistency on all variables, though. Example 5 shows a counterexample. To obtain domain consistency we once more need the clauses P3.

Theorem 8. *CompletePath is propagation complete.* \square

The path based encoding do have one weakness compared to the Tseytin encoding. Since they require only a single path throught the MDD, we cannot allow different MDD constraints that share a sub-MDD to reuse the same encoding, we need a different copy of the encoding for each constraint. This is not the case for Tseytin encodings where the node variable ν just represents the truth value of the sub-formula encoded by the MDD rooted at ν . To our knowledge this restriction is not very significant in the CP context. No such sharing exists in any of our benchmarks. The bulk of nodes in an MDD are in the middle and unlikely to be shared. Moreover, separating MDDs per constraint for translation allows us to use different variable orderings for each MDD and thus reduce the number of nodes required. On the other hand, if substantial sharing of nodes among the different MDDs happens then a Tseytin encoding could be beneficial, since it translates this sharing to the CNF level.

The table below shows the sizes and propagation strength of the different encodings. As before, s is the size of the MDD, d is the maximum domain size of variables x and n is the number of variables x . Notice that usually $n \ll s$.

	Minimal	GMinisat	Tseytin	BasicP	NNFP	LevelP	ComplP
Variables	s	s	$s(d + 1)$	$s(d + 1)$	$s(d + 1)$	$s(d + 2)$	$s(d + 2)$
Clauses	sd	$s(2d + 2)$	$s(4d + 1)$	$s(4d + 2)$	$s(4d + 2)$ $+nd$	$s(4d + 5)$	$s(4d + 5)$ $+ nd$
Consistent	✗	✗	✓	✓	✓	✓	✓
Dom. Consis.	✗	✗	✗	✗	✓	✗	✓
Ref. Compl.	✗	✗	✗	✗	✗	✓	✓
Prop. Compl	✗	✗	✗	✗	✗	✗	✓

4 Encoding NNFs

BDDs are a special case of NNFs and hence NNF encodings provide an alternate approach to encoding BDDs. There is an existing encoding for NNFs given by [20]. When applied correctly to MDDs it results in the NNFPPath (hence the name). But care has to be taken in NNF encodings, without the right restrictions on the form of the NNF the encodings are incorrect!

An encoding of an NNF \mathcal{N} to clauses is given by [20]. Each node ν is associated with a literal, also called ν . For leaf nodes the literal is just the label of the node. For non-leaf nodes the literal is a new Boolean variable. The clauses we make use of are

N1 $\nu \rightarrow \nu_1 \vee \dots \vee \nu_k$ for each \vee -node ν with children ν_1, \dots, ν_k

N2 $\nu \rightarrow \nu_i, 1 \leq i \leq k$ for each \wedge -node ν with children ν_1, \dots, ν_k

N3 $\nu \rightarrow p_1 \vee \dots \vee p_m$ for each node ν with incoming edges from nodes p_1, \dots, p_m .




Fig. 3. NNF for formula (a) $(x \wedge (p \vee q)) \vee (\neg x \wedge (\neg p \vee \neg q))$ and (b) $(\neg q \wedge p) \vee (p \wedge q)$

We consider two encodings: **BaseNNF** Clauses N1–N2 and ρ , and **ExtNNF** Clauses N1–N3 and ρ as defined in [20].

Theorem 9. *Given an NNF \mathcal{N} then **BaseNNF** is a correct encoding.* \square

Note that this *correctness* result *does not apply* to **ExtNNF** unless the NNF is smooth and decomposable. Jung [20] also claim that **ExtNNF** enforces domain consistency for decomposable NNFs, but this too is incorrect.

Example 7. The NNF shown in Fig. 3(a) is decomposable, deterministic but not smooth (e.g. the two children of node ν_4 do not mention the same variables). The **ExtNNF** encoding is

$$\begin{array}{llll}
 N1 : \nu_1 \rightarrow \nu_2 \vee \nu_3 & \nu_4 \rightarrow p \vee q & \nu_5 \rightarrow \neg p \vee \neg q & \\
 N2 : \nu_2 \rightarrow x & \nu_2 \rightarrow \nu_4 & \nu_3 \rightarrow \neg x & \nu_3 \rightarrow \nu_5 \\
 N3 : \nu_2 \rightarrow \nu_1 & \nu_3 \rightarrow \nu_1 & x \rightarrow \nu_2 & \nu_4 \rightarrow \nu_2 \quad \neg x \rightarrow \nu_3 \\
 & \nu_5 \rightarrow \nu_3 & p \rightarrow \nu_4 & q \rightarrow \nu_4 \quad \neg p \rightarrow \nu_5 \quad \neg q \rightarrow \nu_5 \\
 \rho : \nu_1 & & &
 \end{array}$$

Consider the assignment $A = \{x, \neg q\}$ unit propagation determines $\nu_1, \nu_2, \nu_4, p, \nu_5, \nu_3, \neg x$ and hence a contradiction. This is wrong since there is a model of the NNF $\{x, \neg q, p\}$. \square

Example 8. Consider the smooth, decomposable and deterministic NNF for $(\neg q \wedge p) \vee (p \wedge q)$ shown in Fig. 3(b). Then the clauses of **ExtNNF** are

$$\begin{array}{llll}
 \rho : \nu_1 & N1 : & \nu_1 \rightarrow \nu_2 \vee \nu_3 & \\
 N2 : \nu_2 \rightarrow \neg q & \nu_2 \rightarrow p & \nu_3 \rightarrow p & \nu_3 \rightarrow q \\
 N3 : \nu_2 \rightarrow \nu_1 & \nu_3 \rightarrow \nu_1 & \neg q \rightarrow \nu_2 & p \rightarrow \nu_2 \vee \nu_3 \quad q \rightarrow \nu_3
 \end{array}$$

Any model of the formula must make p true, but unit propagation on these clauses derives only ν_1 . What is missing is information that $\neg p$ does not appear in the NNF. This means p *must hold!* \square

Table 1. Results on nurse rostering, shift scheduling and pentominoes.

Bench	Type	Search	#Inst	Prop	Minimal	GMinisat	Tseytin	BasicP	NNFP	LevelP	CompLP
Nurse	SAT	VSIDS	286 #sol	282	88	<u>195</u>	184	150	185	157	187
			78 com	1.97	–	<u>5.33</u>	27.81	58.73	14.09	42.51	24.86
			286 all	23.82	903.64	<u>395.26</u>	473.05	617.42	457.79	607.16	457.85
		prog	179 #sol	132	143	<u>151</u>	156	156	108	156	104
			80 com	3.42	–	<u>6.19</u>	6.61	18.39	54.63	29.96	50.86
			179 all	329.63	284.73	212.5	181.65	171.95	516.36	<u>177.19</u>	526.63
	UNSAT	VSIDS	46 #sol	32	29	46	27	31	<u>33</u>	32	32
			26 com	42.73	–	8.09	229.45	98.31	<u>26.87</u>	71.55	69.26
			46 all	402.57	626.02	231.34	631.03	450.35	<u>380.4</u>	413.69	437.38
	Shift	OPT	VSIDS	120 #sol	<u>114</u>	85	96	97	116	115	110
78 com				109.8	–	166.51	161.91	51.54	88.07	<u>68.91</u>	117.44
120 all				<u>213.84</u>	535.59	457.94	444.8	174.65	252.11	224.41	276.17
prog			56 #sol	49	48	56	48	<u>55</u>	50	52	48
			48 com	100.11	–	<u>28.64</u>	113.06	24.52	74.09	34.02	79.97
			56 all	257.02	240.44	60.42	268.34	<u>161.76</u>	232.17	<u>176.28</u>	239.97
Pent	ALL	prog	14 #sol	14	<u>12</u>	<u>12</u>	6	<u>12</u>	9	<u>12</u>	6
			6 com	6.67	–	<u>8.21</u>	18.27	14.57	16.02	8.8	15.36
			14 all	279.43	<u>352.82</u>	505.92	693.54	626.07	653.08	387.67	692.3

To fix Jung’s encoding we add the following clauses

N4 $\neg l$ for each literal l for $vars(\mathcal{N})$ which does not appear in \mathcal{N} .

We denote by FullNNF Clauses N1–N4 and ρ .

Theorem 10. *Given a smooth decomposable NNF \mathcal{N} then FullNNF is a correct encoding.*

Theorem 11. *Given a smooth decomposable NNF \mathcal{N} , then unit propagation on FullNNF enforces domain consistency.* \square

It follows that FullNNF is equivalent to NNFPPath if applied to MDDs rewritten as NNF. To summarise the results in this section we provide the following table.

	BaseNNF	ExtNNF	FullNNF
Clauses	N1–N2	N1–N3	N1–N4
Correctness	Always	Smooth and Decomposable	Smooth and Decomposable
Domain consistent	✗	✗	✓

5 Experiments

We show results on three benchmarks: nurse rostering, shift scheduling and pentominoes (Nurse, Shift and Pent).³ The MDD encodings are implemented as eager translations of MDDs within the LCG solver Chuffed [9, 10] and compared with a native MDD propagator with learning [17]. We use SAT branching heuristics

³ Benchmarks are available from <http://people.eng.unimelb.edu.au/pstuckey/mddenc.tar.gz>.

(VSIDS) and the programmed search as specified in the models (prog). We omit instances not solved by any solver using that search. For each model we show: (#sol) the number of instances solved (SAT and UNSAT for *Nurse*, to optimality for *Shift*, all solutions for *Pent*); (com) the mean solving time in seconds for all benchmarks solved by all solvers (except *Minimal*); and (all) the mean solving time of all benchmarks using timeout (1200 s) for unsolved instances. The results on the encoding *Minimal* are omitted for *com* and for *Pent* since it does not preserve solution counting. Best results are in bold, and second best are underlined (Table 1).

In case of satisfiable instances of *Nurse* the results show that encodings do not compete with the native propagator. This is not surprising as the search quickly finds the solutions without being disturbed by the complete CNF model generated by the eager encodings. For the UNSAT instances decompositions and their intermediate literals show their strength and beat the propagator. *GenMiniSAT* shows best performance for these UNSAT instances with VSIDS. The encodings also have an advantage over the propagator when programmed search is used, but it is unclear which one dominates.

For *Shift* the results show that when using activity based search and branching takes place on auxiliary variables, the path based approaches are generally superior.

The main advantage of the native propagator is that its explanations are built in a more deterministic fashion and hence tend to be more reusable. Furthermore, since the propagator only generates a fraction of the variables of the eager encoding, the search is less likely get trapped in an unfruitful search space using VSIDS. The difference in results on SAT and UNSAT instances of *Nurse* clearly indicate that a combination of the propagator and a lazy encoding as in [1] would be a strong approach.

6 Conclusion and Future Work

This paper resulted from discussions that uncovered our own misunderstanding of the strength of decision diagram encodings. We were surprised to discover that the usual BDD encoding is not domain consistent. In this paper we seek to remove this confusion, and demonstrate a wealth of different encoding possibilities, with different properties.

The experimental results show that there is unlikely to be one single best encoding for MDDs, and hence an important direction of future work is to determine when each encoding is best. Possibly a portfolio approach varying over encodings of each constraint is a fruitful and pragmatic technique to solve hard problems in practice.

Another interesting direction of future work is to determine a propagation complete encoding for NNFs. It appears the result may require restricting to Sentential Decision Diagrams [12] a form of NNF with a uniform V-tree.

The literature on CNF encodings focuses on consistencies wrt. primary variables of the constraint, whereas we have shown that consistency on auxiliary variables are worthwhile to look at. Our work concentrated on translations of decision diagrams and we would like to extend this research to other constraints like

linear and **sequence**. State-of-the-art CNF encodings of **cardinality** are the next candidate for this investigation.

In case of theoretical results, an interesting direction is to establish lower bounds on the size of encodings implementing certain consistencies for concrete constraints. The strong relationship between CNF encodings and monotone circuits established in [5, 19] demonstrates a powerful tool for this purpose.

Acknowledgement. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

References

1. Abío, I., Nieuwenhuis, R., Oliveras, A., Rodríguez-Carbonell, E., Stuckey, P.J.: To encode or to propagate? The best choice for each constraint in SAT. In: Schulte, C. (ed.) CP 2013. LNCS, vol. 8124, pp. 97–106. Springer, Heidelberg (2013)
2. Abío, I., Stuckey, P.J.: Encoding linear constraints into SAT. In: O’Sullivan, B. (ed.) CP 2014. LNCS, vol. 8656, pp. 75–91. Springer, Heidelberg (2014)
3. Ansótegui, C., Manyà, F.: Mapping problems with finite-domain variables into problems with boolean variables. In: SAT 2004 - The Seventh International Conference on Theory and Applications of Satisfiability Testing, Vancouver, BC, Canada, Online Proceedings, 10–13 May 2004
4. Bacchus, F.: GAC via unit propagation. In: Bessière, C. (ed.) CP 2007. LNCS, vol. 4741, pp. 133–147. Springer, Heidelberg (2007)
5. Bessiere, C., Katsirelos, G., Narodytska, N., Walsh, T.: Circuit complexity and decompositions of global constraints. In: Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009, Pasadena, California, USA, 11–17 July 2009, pp. 412–418 (2009)
6. Bordeaux, L., Marques-Silva, J.: Knowledge compilation with empowerment. In: Bieliková, M., Friedrich, G., Gottlob, G., Katzenbeisser, S., Turán, G. (eds.) SOFSEM 2012. LNCS, vol. 7147, pp. 612–624. Springer, Heidelberg (2012)
7. Chen, J.: A new SAT encoding of the at-most-one constraint. In: Proceedings of the Tenth International Workshop of Constraint Modelling and Reformulation (2010)
8. Cheng, K.C.K., Yap, R.H.C.: Maintaining generalized arc consistency on Ad Hoc r -Ary constraints. In: Stuckey, P.J. (ed.) CP 2008. LNCS, vol. 5202, pp. 509–523. Springer, Heidelberg (2008)
9. Chu, G.: Chuffed. <https://github.com/geoffchu/chuffed>
10. Chu, G.G.: Improving combinatorial optimization. Ph.D. thesis, The University of Melbourne (2011)
11. Darwiche, A.: On the tractable counting of theory models and its application to truth maintenance and belief revision. *J. Appl. Non-class. Logics* **11**(1–2), 11–34 (2001)
12. Darwiche, A.: SDD: a new canonical representation of propositional knowledge bases. In: IJCAI, pp. 819–826 (2011)
13. Eén, N., Sörensson, N.: Translating pseudo-boolean constraints into SAT. *JSAT* **2**(1–4), 1–26 (2006)
14. Frisch, A.M., Peugniez, T.J., Doggett, A.J., Nightingale, P.: Solving non-boolean satisfiability problems with stochastic local search: a comparison of encodings. *J. Autom. Reason.* **35**(1–3), 143–179 (2005)

15. Gange, G., Stuckey, P.J., Van Hentenryck, P.: Explaining propagators for edge-valued decision diagrams. In: Schulte, C. (ed.) CP 2013. LNCS, vol. 8124, pp. 340–355. Springer, Heidelberg (2013)
16. Gange, G., Stuckey, P.J., Lagoon, V.: Fast set bounds propagation using a BDD-SAT hybrid. *J. Artif. Intell. Res. (JAIR)* **38**, 307–338 (2010)
17. Gange, G., Stuckey, P.J., Szymanek, R.: MDD propagators with explanation. *Constraints* **16**(4), 407–429 (2011)
18. Gent, I.P.: Arc consistency in SAT. In: Proceedings of the 15th European Conference on Artificial Intelligence, ECAI 2002, Lyon, France, pp. 121–125, July 2002
19. Gwynne, M., Kullmann, O.: A framework for good SAT translations, with applications to CNF representations of XOR constraints. CoRR abs/1406.7398 (2014). <http://arxiv.org/abs/1406.7398>
20. Jung, J.C., Barahona, P., Katsirelos, G., Walsh, T.: Two encodings of DNNF theories. In: ECAI 2008 Workshop on Inference Methods Based on Graphical Structures of Knowledge (2008)
21. Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Solving SAT and SAT modulo theories: from an abstract davis-putnam-logemann-loveland procedure to DPLL(T). *J. ACM* **53**(6), 937–977 (2006)
22. Plaisted, D.A., Greenbaum, S.: A structure-preserving clause form translation. *J. Symb. Comput.* **2**(3), 293–304 (1986)
23. Quimper, C.-G., Walsh, T.: Decomposing global grammar constraints. In: Bessière, C. (ed.) CP 2007. LNCS, vol. 4741, pp. 590–604. Springer, Heidelberg (2007)
24. Srinivasan, A., Ham, T., Malik, S., Brayton, R.: Algorithms for discrete function manipulation. In: 1990 IEEE International Conference on Computer-Aided Design, ICCAD-90, Digest of Technical Papers, pp. 92–95 (1990)
25. Tseytin, G.: On the complexity of derivation in propositional calculus. *Stud. Constr. Math. Math. Logic Part 2*, 115–125 (1968)
26. del Val, A.: Tractable databases: how to make propositional unit resolution complete through compilation. In: Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning (KR 1994), Bonn, Germany, 24–27 May 1994, pp. 551–561 (1994)
27. Walsh, T.: SAT v CSP. In: Dechter, R. (ed.) CP 2000. LNCS, vol. 1894, pp. 441–456. Springer, Heidelberg (2000)