Understanding the complexity of #SAT using knowledge compilation[*]

Florent Capelli

Birkbeck College, University of London

florent@dcs.bbk.ac.uk

August 14, 2018

Abstract

Two main techniques have been used so far to solve the #P-hard problem #SAT. The first one, used in practice, is based on an extension of DPLL for model counting called exhaustive DPLL. The second approach, more theoretical, exploits the structure of the input to compute the number of satisfying assignments by usually using a dynamic programming scheme on a decomposition of the formula. In this paper, we make a first step toward the separation of these two techniques by exhibiting a family of formulas that can be solved in polynomial time with the first technique but needs an exponential time with the second one. We show this by observing that both techniques implicitly construct a very specific boolean circuit equivalent to the input formula. We then show that every β-acyclic formula can be represented by a polynomial size circuit corresponding to the first method and exhibit a family of β-acyclic formulas which cannot be represented by polynomial size circuits corresponding to the second method. This result shed a new light on the complexity of #SAT and related problems on β-acyclic formulas. As a byproduct, we give new handy tools to design algorithms on β-acyclic hypergraphs.

1 Introduction

The problem #SAT of counting the satisfying assignments of a given CNF-formula is a central problem to several areas such as probabilistic reasoning [Rot96, BDP03] and probabilistic databases [BLRS14, BLRS13, JS13]. This problem is much harder than SAT, its associated decision problem. For example, the problem 2-SAT of deciding if a formula having at most two literals per clause if satisfiable is easy where counting those satisfying assignments is as hard as #SAT. Even computing a $2^{n^{1-\epsilon}}$-approximation in the restricted case of monotone 2-SAT is hard for any $\epsilon > 0$ [Rot96].

In order to tackle this problem, two main approaches have been used so far. The first approach – applied in practical tools for solving #SAT – follows the successful road paved by SAT-solvers: it is based on a variation of DPLL [DP60] called *exhaustive DPLL* [HD05] and the approach is mainly focused on improving the heuristics used for eliminating variables and choosing which subformulas should be cached during the computation. The performance of such tools – though impressive for such a hard problem [HD05, SBB+04, Thu06, BDP03] – lag far behind the state-of-the-art SAT-solvers. This gap is mainly explained by the differences between the hardness of both problems,

---

but also by the fact that optimizations for exhaustive DPLL are inspired by those used in SAT-solvers and not always relevant for model counting [SBK05]. The second – more theoretical – approach focuses on *structural restrictions* of the input formula. The main idea of this approach is to solve #SAT more quickly on formulas where interaction between the clauses and the variables is restricted. This interaction is usually represented by a graph derived from the input CNF-formula. The complexity of #SAT is then studied on inputs where the associated graph belongs to a restricted class of graphs. Samer and Szeider [SS10] were the first to formalize this idea for #SAT by showing that if this graph is of bounded tree width, then #SAT can be solved in polynomial time. This result has then been improved and completed by different work showing the tractability of #SAT for more general or incomparable classes of formulas [PSS16, SS13, HSTV14, CDM14], the intended goal being to understand the frontier of tractibility for #SAT.

Contributions. The main contribution of this paper is to propose a formal framework, using tools from *knowledge compilation*, to study both algorithmic techniques and to compare their respective power. We then make a first step toward the separation of both techniques by exhibiting a class of formulas having the following property: for every formula F of this class, there exists an elimination order of the variables for which exhaustive DPLL returns the number of satisfying assignments of F in linear time while algorithms based on structural restrictions needs exponential time.

The class of formulas we are using to separate both technique are β-acyclic formulas, a class already known to be tractable [BCM15]. The algorithm used to solve this class was however very different from the one that are usually used by structure-based algorithms. Our result gives a formal explanation of why the usual techniques fail on this class, a question that has puzzled the community since SAT has been shown tractable on this class of formulas without generalizing to counting [OPS13].

Moreover, in Section 3, we give tools that are useful for designing algorithms on β-acyclic hypergraphs and are of independent interest.

Methodology. It has been observed that the trace of every implementation of exhaustive DPLL actually constructs a very specific Boolean circuit equivalent to the input formula [HD05]. Such circuits are known in knowledge compilation under the name of *decision Decomposable Negation Normal Form* (dec-DNNF). We first show in Section 3 that β-acyclic formulas can be represented by linear size dec-DNNF, which can be interpreted as the fact that exhaustive DPLL may solve this class of formula in polynomial time, if it chooses the right order to eliminate variables and the right caching methods.

Similarly, all structure-based algorithms for #SAT use the same kind of dynamic programming scheme and it has been shown that they all implicitly construct a very specific Boolean circuit equivalent to the input formula [BCMS15]. Such circuits are known under the name of *structured deterministic* DNNF. We start by arguing in Section 2 that every algorithm using techniques similar to the one used by structure-based algorithms will implicitly construct a circuit having a special property called *determinism*. In Section 4, We exhibit a class of β-acyclic formulas having no polynomial size equivalent *structured* DNNF, thus separating both methods.

Related work. The class of β-acyclic formulas we use to prove the separation have already been shown to be tractable for #SAT and not tractable to the state-of-the-art structure-based algorithms [BCM15] but this result does not rule out the existence of a more general algorithm

2

based on the same technique and solving every known tractable class. Our result is sufficiently strong to rule out the existence of such an algorithm.

New lower bounds have been recently shown for circuits used in knowledge compilation [BCMS16, BLRS14, BLRS13, BL15, PD10]. Moreover, knowledge compilation has already been used to prove limits of algorithmic techniques in the context of model counting. Beame et al. [BLRS13] for example have exhibited a very interesting class of queries on probabilistic databases that can be answered in polynomial time by using specific techniques but that cannot be represented by circuits corresponding to exhaustive DPLL. They conclude that solving the query by using well-known reduction to #SAT and then calling a #SAT-solver is weaker than using their technique. Our result uses somehow the same philosophy but on a different algorithmic technique.

Organization of the paper. The paper is organized as follows: Section 2 contains the needed definitions and concepts used through the paper. Section 3 describes the algorithm to transform β-acyclic formulas into circuits corresponding to the execution of an exhaustive DPLL algorithm. Finally, Section 4 contains the formalization of the framework for studying algorithms based on dynamic programming along a branch decomposition and a proof that the β-acyclic case is not covered by this framework.

2   Preliminaries

2.1   CNF-formulas.

A *literal* is a variable x or a negated variable $\neg x$. A *clause* is a finite set of literals. A clause is *tautological* if it contains the same variable negated as well as unnegated. A *(CNF) formula* (or *CNF*, for short) is a finite set of non-tautological clauses. If x is a variable, we let $\mathsf{var}(x) = \mathsf{var}(\neg x) = x$. Given a clause C, we denote by $\mathsf{var}(C) = \bigcup_{\ell \in C} \mathsf{var}(\ell)$ and given a CNF-formula, we denote by $\mathsf{var}(F) = \bigcup_{C \in F} \mathsf{var}(C)$. The *size* of a CNF-formula F, denoted by $\mathsf{size}(F)$, is defined to be $\sum_{C \in F} |\mathsf{var}(C)|$. A CNF-formula is *monotone* if it does not contain negative literals.

Let X be a set of variables. An *assignment* τ of X is a mapping from X to $\{0, 1\}$. The set of assignments of X is denoted by $\{0, 1\}^X$. Given an assignment τ of X and $X' \subseteq X$, we denote by $\tau|_{X'}$ the restriction of τ on X'. Given two sets X, X', $\tau \in \{0, 1\}^X$ and $\tau' \in \{0, 1\}^{X'}$, we denote by $\tau \simeq \tau'$ if $\tau|_{X \cap X'} = \tau'|_{X \cap X'}$. If $\tau \simeq \tau'$, we denote by $\tau \cup \tau'$ the assignment of $X \cup X'$ such that for all $x \in X$, $(\tau \cup \tau')(x) = \tau(x)$ and for all $x \in X'$, $(\tau \cup \tau')(x) = \tau'(x)$.

A *boolean function* f on variables X is a mapping from $\{0, 1\}^X$ to $\{0, 1\}$. We denote by $\tau \models f$ if $\tau \in \{0, 1\}^X$ is such that $f(\tau) = 1$ and by $\mathsf{sat}(f) = \{\tau \in \{0, 1\}^X \mid \tau \models f\}$. Given $Y \subseteq X$ and $\tau \in \{0, 1\}^Y$, we denote by $f[\tau]$ the boolean function on variables $X \setminus Y$ defined by for every $\tau' \in \{0, 1\}^{X \setminus Y}$, $f[\tau](\tau') = f(\tau \cup \tau')$.

A CNF-formula F naturally induces a boolean function. Extending assignments to literals in the usual way, we say that an assignment τ *satisfies* a clause C if there is a literal $\ell \in C$ such that $\tau(\ell) = 1$. An assignment *satisfies* a formula F if it satisfies every clause $C \in F$. In this paper, we often identify the CNF-formula and its underlying boolean function. Thus, given a CNF-formula F on variables X and an assignment τ of $Y \subseteq X$, we will use the notation $F[\tau]$ in the same way as for any other boolean function. Observe that $F[\tau]$ is still represented by a CNF-formula of size less than $\mathsf{size}(F)$: it is the CNF-formula where we have removed satisfied clauses from F and removed the variables of Y in each remaining clause.
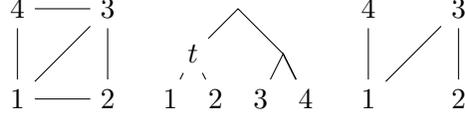
Figure 1: From left to right: a graph $G = (V, E)$, a branch decomposition of G and $G[V_t, V \setminus V_t]$

2.2 Graphs and branch decompositions.

We assume the reader is familiar with the basics of graph theory. An introduction to the topic can be found in [Die12]. Given a graph $G = (V, E)$, we often denote by $V(G)$ the set of vertices of G and by $E(G)$ the set of edges of G if they have not been named explicitly before. G is said *bipartite* if there exists a partition $V_1 \uplus V_2$ of V such that for every $e \in E$, one end-point of e is in V_1 and the other is in V_2. Given a graph $G = (V, E)$ and $X, Y \subseteq V$, we denote by $G[X, Y] = (V', E')$ the bipartite graph such that $V' = X \cup Y$ and $E' = \{\{u, v\} \in E \mid u \in X, v \in Y\}$. An *induced matching* M is a matching of G such that for every $e, f \in M$, if $e = \{u, v\}$ and $f = \{u', v'\}$, we have $\{u, u'\} \notin E$, $\{u, v'\} \notin E$, $\{v, u'\} \notin E$ and $\{v, v'\} \notin E$.

A *branch decomposition* of G is a binary rooted tree T whose leaves are in one-to-one correspondence with V. Given t a vertex of T, we denote by T_t the subtree of T rooted in t. We denote by V_t the set of leaves of T_t.

The *maximal induced matching width* [Vat12], MIM-width for short, of a vertex t of T is the size of the largest induced matching M of $G[V \setminus V_t, V_t]$. The MIM-width of T, denoted as $\mathbf{mimw}(T)$, is the maximal MIM-width of its vertices. The *maximal induced matching width* of a graph G, denoted as $\mathbf{mimw}(G)$, is the minimal MIM-width of all branch decomposition of G. Figure 1 depicts a graph G together with a branch decomposition of G. The distinguished node t of this branch decomposition has MIM-width 1 as the biggest induced matching of $G[V_t, V \setminus V_t]$ is of size one because the matching $\{\{1, 4\}, \{2, 3\}\}$ is not induced.

2.3 Hypergraphs and β-acyclicity.

A *hypergraph* $\mathcal{H}$ is a finite set of finite sets, called *edges*. We denote by $V(\mathcal{H}) = \bigcup_{e \in \mathcal{H}} e$ the set of vertices of hypergraph $\mathcal{H}$.

Most notions on graphs may be naturally generalized to hypergraph. A hypergraph $\mathcal{H}'$ is a subhypergraph of $\mathcal{H}$ if $\mathcal{H}' \subseteq \mathcal{H}$. Given $S \subseteq V(\mathcal{H})$, we denote by $\mathcal{H} \setminus S = \{e \setminus S \mid e \in \mathcal{H}\}$. A *walk* of length n from edge $e \in \mathcal{H}$ to $f \in \mathcal{H}$ is a sequence $(e_1, x_1, \ldots, x_n, e_{n+1})$ of vertices and edges such that: $e = e_0$, $f = e_{n+1}$ and for every $i \leq n$, $x_i \in e_i \cap e_{i+1}$. A *path* is a walk that never goes twice through the same vertex nor the same edge. It is easy to check that if there is a walk from e to f in $\mathcal{H}$, then there is also a path from e to f.

There exist several generalizations of acyclicity to hypergraph introduced by Fagin [Fag83] in the context of database query answer. An extensive presentation of hypergraph acyclicity notions may be found in [BB14]. In this paper, we focus on the β-acyclicity, which is the most general of such notions for which #SAT is still tractable. A hypergraph $\mathcal{H}$ is β-acyclic if there exists an order $(x_1, \ldots, x_n)$ of $V(\mathcal{H})$ such that for all $i \leq n$, for all $e, f \in \mathcal{H}$ such that $x_i \in e \cap f$, then either $e \setminus \{x_1, \ldots, x_i\} \subseteq f$ or $f \setminus \{x_1, \ldots, x_i\} \subseteq e$. Such an order is called a β-elimination order of $\mathcal{H}$. A β-acyclic hypergraph can be found on Figure 2. The order $\{1, 2, 3, 4, 5\}$ is an β-elimination order.

Given a hypergraph $\mathcal{H}$, the incidence graph of $\mathcal{H}$ is defined as the bipartite graph whose vertices
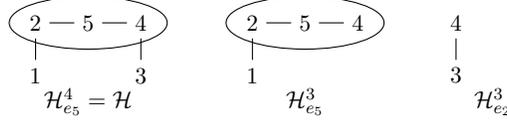
Figure 2: An example of $\mathcal{H}_e^x$

are $V(\mathcal{H}) \cup \mathcal{H}$ and there is an edge between $e \in \mathcal{H}$ and $x \in V(\mathcal{H})$ if and only if $x \in e$. The *incidence* MIM-*width* of a hypergraph is the MIM-width of its incidence graph. The incidence MIM-width of β-acyclic hypergraphs can be very large:

Theorem 1 ([BCM15]). *There exists an infinite family of β-acyclic hypergraphs of incidence* MIM-*width* $\Omega(n)$ *where n is the number of vertices of the hypergraph.*

2.4 Structure of formulas.

Let F be a CNF-formula. The *incidence graph* of F, denoted by $\mathcal{G}_{\mathsf{inc}}(F)$, is the bipartite graph whose vertices are the variables and the clauses of F and there is an edge between a variable x and a clause C if and only if $x \in \mathsf{var}(C)$. The *incidence* MIM-*width* of a formula F is the MIM-width of $\mathcal{G}_{\mathsf{inc}}(F)$. The hypergraph of F, denoted by $\mathcal{H}(F)$, is defined as $\mathcal{H}(F) = \{\mathsf{var}(C) \mid C \in F\}$. A CNF-formula is said to be β-*acyclic* if and only if its hypergraph is β-acyclic.

2.5 Knowledge compilation

DNNF. In this paper we focus on so-called DNNF introduced by Darwiche [Dar01]. An extensive presentation of different target languages with their properties may be found in [DM02]. A Boolean circuit C on variables X is in *Negation Normal Form*, NNF for short, if its input are labeled by literals on X and its internal gates are labeled with either a $\wedge$-gate or a $\vee$-gate. We assume that such circuit has a distinguished gate called the *output*. An NNF circuit D computes the boolean function computed by its output gate and we will often identify the circuits and its computed Boolean function. We denote by $\mathsf{size}(D)$ the number of gates of D and by $\mathsf{var}(D)$ the set of variables labeling its input. If v is a gate of D, we denote by D_v the circuit given by the maximal the sub-circuits of D rooted in v and whose output is v. If v is an $\wedge$-gate, it is said *decomposable* if for every v_1, v_2 that are distinct inputs of v, it holds that $\mathsf{var}(D_{v_1}) \cap \mathsf{var}(D_{v_2}) = \emptyset$. An NNF circuit is in *Decomposable Normal Form* if all its $\wedge$-gates are decomposable. We will refer to such circuits as DNNF. It is easy to see that one can find a satisfying assignment of a DNNF D in time $O(\mathsf{size}(D))$. Moreover, if D is a DNNF on variables X, $Y \subseteq X$ and $\tau \in \{0, 1\}^Y$, then $D[\tau]$ is computed by a DNNF smaller than D since we can plug the values of literals in Y in the circuit D.

Deterministic and Decision DNNF. Let D be a DNNF. An $\vee$-gate in D is called *deterministic* if for every v_1, v_2 that are distinct inputs of v, it holds that $D_{v_1} \wedge D_{v_2} \equiv 0$. D is said deterministic if all its $\vee$-gates are deterministic. Observe that determinism is a semantic condition and is hard to decide from the DNNF only. In this paper, we will be mostly interested in *decision* gates that are a special case of deterministic gates. An $\vee$-gate v of D is a *decision gate* if it is binary and if there exists a variable x and two gates v_1, v_2 of D such that v is of the form $(x \wedge v_1) \vee (\neg x \wedge v_2)$. A *decision* DNNF, dec-DNNF for short, is a DNNF for which every $\vee$-gate is a decision gate. It is

5

Figure 3: A DNNF and a vtree

easy to see that a dec-DNNF is deterministic. Figure 3 depicts a DNNF. The output is represented by a square and the DNNF computes the boolean function $(\neg x \wedge z) \vee (x \wedge (y \vee z))$. It is easy to check that both $\wedge$-gates are decomposable. The output gate is also a decision gate on variable x. The other $\vee$-gate is not deterministic since the boolean function $y \wedge z$ is satisfiable.

Structuredness. Structuredness is a constraint on the way variables can be partitioned by a DNNF. It may be seen as a generalization to trees of the variable ordering that is sometimes required in data structures such as OBDD [Weg00] and was introduced in [PD08]. Let D be a DNNF on variables X. A *vtree* T on X is a binary tree whose leaves are in one-to-one correspondence with X. An $\wedge$-gate v of D *respects* a vertex t of T if it has exactly two inputs v_1, v_2 and if $\mathsf{var}(D_{v_1}) \subseteq X_{t_1}$ and $\mathsf{var}(D_{v_2}) \subseteq X_{t_2}$ where t_1, t_2 are the children of t in T and X_{t_1} (resp. X_{t_2}) is the set of variables that appears in the leaves of T_{t_1} (resp. T_{t_2}). A DNNF D respects a vtree T if for every $\wedge$-gate v of D, there exists a vertex t of T such that v respects t. A DNNF D is *structured* if there exists a vtree T such that D respects T. It can be checked that the DNNF depicted in Figure 3 respects the vtree given on the same figure.

2.6 Structuredness and Branch Decomposition

In this section, we explain how most of the structure-based algorithms for #SAT work and how we can relate this to the fact that they are implicitly constructing a structured DNNF equivalent to the input formula.

The current techniques for solving #SAT by exploiting the structure of the input are all based on the same technique: they start by computing a "good" branch decomposition T of the incidence graph of the formula F. Each vertex t of the branch decomposition is then used to define a sub-formula F_t and partial assignments $a_1, \ldots, a_k$ of its variables. The number of solutions of $F_t[a_i]$ is then computed by dynamic programming along the branch decomposition in a bottom-up fashion. In all algorithms, the variables of F_t are the variables of F that label the leaves of T_t. The number of solutions of F_t on some partial assignment a_i is computed by multiplying and summing the number of solutions of F_{t_1} and of F_{t_2}, where t_1, t_2 are the children of t on restrictions of a_i to the variables of F_{t_1} and F_{t_2} respectively. Those multiplications can be seen as a decomposable $\wedge$-gate and the sums can be seen as deterministic $\vee$-gates. Thus, the underlying DNNF constructed by those algorithms is naturally structured along the vtree obtained from the branch decomposition T by forgetting the leaves that are labeled by clauses of the formula.

In this paper, we will thus say that a class of formula can be solved by using the *standard (dynamic programming) technique* if it can be compiled into deterministic structured DNNF. The most general known algorithm exploiting the structure of the input, that we will call, from the author names, the *STV-algorithm* [HSTV14], uses exactly this technique. It has been observed in [BCMS15] that this algorithm is actually implicitly constructing a deterministic structured DNNF

equivalent to the input CNF-formula, which reinforces the idea that the notion of structuredness captures the essence of the standard technique for solving #SAT.

3 Compilation of β-acyclic formulas into dec-DNNF

We show how to construct a linear size dec-DNNF equivalent to a given β-acyclic formula F (Theorem 8). We use a dynamic programming approach by iteratively constructing dec-DNNF for subformulas of F. These subformulas are defined using general remarks on the structure of β-acyclic hypergraphs.

3.1 Structure of β-acyclic hypergraphs.

In this section, we fix a β-acyclic hypergraph $\mathcal{H}$ with n vertices and a β-elimination order $(x_1, \ldots, x_n)$ of its vertices denoted by $<$. We denote by $<_\mathcal{H}$ the order on $\mathcal{H}$ defined as the lexicographical order on $\mathcal{H}$ where $e \in \mathcal{H}$ is seen as the $\{0,1\}^n$-vector $\vec{e}$ such that $\vec{e}_i = 1$ if $x_{n-i} \in e$ and $\vec{e}_i = 0$ otherwise. In other words, $e <_\mathcal{H} f$ if and only if $\max(e\Delta f) \in f$.

From these orders, we construct a family of subhypergraphs of $\mathcal{H}$ which will be interesting for us later. Let $x \in V$ and $e \in \mathcal{H}$. We denote by $V_{\leq x} = \{y \in V \mid y \leq x\}$. $V_{<x}$, $V_{\geq x}$ and $V_{>x}$ are defined similarly. We denote by $\mathcal{H}_e^x$ the subhypergraph of $\mathcal{H}$ that contains the edges $f \in \mathcal{H}$ such that there is a walk from f to e that goes only through edges smaller than e and vertices smaller than x.

Observe that, by definition, $\mathcal{H}_e^x$ is a connected subhypergraph of $\mathcal{H}$, with $e \in \mathcal{H}_e^x$ and for all $f \in \mathcal{H}_e^x$, $f \leq_\mathcal{H} e$. Observe also that even if there is a walk from $f \in \mathcal{H}_e^x$ to e that goes only through vertices smaller than x, f may hold vertices that are bigger than x. We insist on the fact that the whole edge f is in $\mathcal{H}_e^x$ and not only its restriction to $V_{\leq x}$.

We start by giving an example. Let $\mathcal{H} = \{\{1,2\}, \{3,4\}, \{2,5\}, \{4,5\}, \{2,4,5\}\}$ be the hypergraph depicted on Figure 2. One can easily check that $1 < 2 < 3 < 4 < 5$ is a β-elimination order and that the order $<_\mathcal{H}$ is the following: $e_1 = \{1,2\} <_\mathcal{H} e_2 = \{3,4\} <_\mathcal{H} e_3 = \{2,5\} <_\mathcal{H} e_4 = \{4,5\} <_\mathcal{H} e_5 = \{2,4,5\}$. $\mathcal{H}_{e_5}^4$ is the whole hypergraph since one can reach any edge from e_5 by going through vertices smaller than 4. $\mathcal{H}_{e_5}^3$ however is lacking the edge $e_2 = \{3,4\}$ since the only way of reaching e_2 from e_5 is to go through the vertex 4 which is not allowed.

Observe that these subhypergraphs are naturally ordered by inclusion:

Lemma 2. *Let* $x, y \in V(\mathcal{H})$ *such that* $x \leq y$ *and* $e, f \in \mathcal{H}$ *such that* $e \leq_\mathcal{H} f$ *and* $V(\mathcal{H}_e^x) \cap V(\mathcal{H}_f^y) \cap V_{\leq x} \neq \emptyset$. *Then* $\mathcal{H}_e^x \subseteq \mathcal{H}_f^y$. *In particular, for all* y, *if* $e \in \mathcal{H}_f^y$ *then* $\mathcal{H}_e^y \subseteq \mathcal{H}_f^y$.

Proof. Let $z \in V(\mathcal{H}_e^x) \cap V(\mathcal{H}_f^y) \cap V_{\leq x}$ and let $g_1 \in \mathcal{H}_e^x$ and $g_2 \in \mathcal{H}_f^y$ be such that $z \in g_1 \cap g_2$. By definition, there exists a walk $\mathcal{P}_1$ from f to g_2 going through vertices smaller than y and edges smaller than f and a walk $\mathcal{P}_2$ from g_1 to e going through vertices smaller than x and edges smaller than e. Since $z \leq x \leq y$ and $e \leq_\mathcal{H} f$, $\mathcal{P} = (\mathcal{P}_1, z, \mathcal{P}_2)$ is a walk from f to e going through edges smaller than f and vertices smaller than y, that is $e \in \mathcal{H}_f^y$. Now let $h \in \mathcal{H}_e^x$ and let $\mathcal{P}_3$ be a path from e to h going through vertices smaller than x and edges smaller than e. Then $(\mathcal{P}, \mathcal{P}_3)$ is a walk from f to h going through vertices smaller than y and edges smaller than f. That is $h \in \mathcal{H}_f^y$, so $\mathcal{H}_e^x \subseteq \mathcal{H}_f^y$. □

We now state the main result of this section. Theorem 3 relates the variables of $\mathcal{H}_e^x$ to those of $V_{\geq x}$ and e. This is crucial for the dynamic programming scheme of our algorithm:

Theorem 3. *For every $x \in V$ and $e \in \mathcal{H}$, $V(\mathcal{H}_e^x) \cap V_{\geq x} \subseteq e$.*

In order to prove Theorem 3, we need two easy intermediate lemmas:

Lemma 4. *Let $e, f \in \mathcal{H}$ such that there exists $x \in e \cap f$. If $e <_{\mathcal{H}} f$ then $e \cap V_{\geq x} \subseteq f$.*

Proof. By definition of β-acyclic elimination order, we must have either $e \cap V_{\geq x} \subseteq f \cap V_{\geq x}$ or $f \cap V_{\geq x} \subseteq e \cap V_{\geq x}$. Now since $e <_{\mathcal{H}} f$, we have $m = \max(e \Delta f) \in f$. If $m \leq x$, we have $e \cap V_{\geq x} = f \cap V_{\geq x}$. Otherwise, we have $e \cap V_{\geq x} \subseteq f \cap V_{\geq x}$ since $m \in (f \setminus e) \cap V_{\geq x}$. $\qquad\square$

A path $\mathcal{P} = (e_0, x_0, \ldots, x_{n-1}, e_n)$ in $\mathcal{H}$ is called *decreasing* if for all i, $e_i >_{\mathcal{H}} e_{i+1}$ and $x_i > x_{i+1}$.

Lemma 5. *For every $x \in V$, $e \in \mathcal{H}$ and $f \in \mathcal{H}_e^x$, there exists a decreasing path from e to f going through vertices smaller than x.*

Proof. By definition of $\mathcal{H}_e^x$, there exists a path $\mathcal{P} = (e_0, x_0, \ldots, x_{n-1}, e_n)$ with $e_0 = e$ and $e_n = f$ such that for all $i \leq n$, $e_i \leq_{\mathcal{H}} e$ and $x_i \leq x$. We show that if $\mathcal{P}$ is a shortest path among those going through vertices smaller than x, then it is also decreasing. Assume toward a contradiction that $\mathcal{P}$ is a non-decreasing such shortest path. Remember that by definition of paths, the edges (e_i) are pairwise distinct. The same is true for the vertices (x_i). Moreover, observe that since $\mathcal{P}$ is a shortest path, then it holds that:

$$\forall i < n \forall j \notin \{i, i+1\}, x_i \notin e_j. \tag{$\star$}$$

Indeed, if there exists i and $j \notin \{i, i+1\}$ such that $x_i \in e_j$, $\mathcal{P}$ could be shortened by going directly from e_i to e_j if $j > i+1$ or from e_j to e_{i+1} if $j < i$.

Let $i = \min\{j \mid x_{j+1} > x_j \text{ or } e_{j+1} >_{\mathcal{H}} e_j\}$ be the first indices where $\mathcal{P}$ does not respect the decreasing condition, which exists if $\mathcal{P}$ is not decreasing by assumption.

First assume $i = 0$. By definition of $\mathcal{P}$, $e_0 = e >_{\mathcal{H}} e_1$. Thus it must be that $x_0 < x_1$. By definition, $x_0 \in e_0 \cap e_1$ and by Lemma 4, $e_1 \cap V_{\geq x_0} \subseteq e_0$. Since $x_1 > x_0$, $x_1 \in e_1 \cap V_{\geq x_0}$, thus $x_1 \in e_0$ which contradicts $(\star)$.

Now assume $i > 0$. First, assume that $e_{i+1} >_{\mathcal{H}} e_i$. By definition of $\mathcal{P}$, it holds that $x_i \in e_i \cap e_{i+1}$ and then by Lemma 4, $e_i \cap V_{\geq x_i} \subseteq e_{i+1}$. Now observe that by minimality of i, $x_{i-1} > x_i$. Since $x_{i-1} \in e_i$, $x_{i-1} \in e_i \cap V_{\geq x_i} \subseteq e_{i+1}$, which contradicts $(\star)$.

Otherwise, $e_i >_{\mathcal{H}} e_{i+1}$ and then $x_{i+1} > x_i$. By Lemma 4 again, $e_{i+1} \cap V_{\geq x_i} \subseteq e_i$. Since $x_{i+1} \in e_{i+1}$, this implies that $x_{i+1} \in e_{i+1} \cap V_{\geq x_i} \subseteq e_i$, which contradicts $(\star)$. It follows that such i does not exist, that is, $\mathcal{P}$ is decreasing. $\qquad\square$

Proof (of Theorem 3). We show by induction on n that for any decreasing path $\mathcal{P} = (e_0, x_0, \ldots, e_n)$ from e_0 to e_n, we have $e_0 \supseteq e_n \cap V_{\geq x_0}$. If $n = 0$, then $e_n = e_0$ and the inclusion is obvious. Now, let $\mathcal{P} = (e_0, x_0, \ldots, e_n, x_n, e_{n+1})$. By induction, $e_0 \supseteq e_n \cap V_{\geq x_0}$ since $(e_0, x_0, \ldots, e_n)$ is a decreasing path from e_0 to e_n. Now by Lemma 4, since $x_n \in e_{n+1} \cap e_n$ and $e_{n+1} <_{\mathcal{H}} e_n$, we have $e_{n+1} \cap V_{\geq x_n} \subseteq e_n$. Since $x_0 > x_n$, $e_{n+1} \cap V_{\geq x_0} \subseteq e_{n+1} \cap V_{\geq x_n} \subseteq e_n$. Thus $e_{n+1} \cap V_{\geq x_0} \subseteq e_n \cap V_{\geq x_0} \subseteq e_0$ which concludes the induction.

Now let $e \in \mathcal{H}$, $x \in V(\mathcal{H})$ and $f \in \mathcal{H}_e^x$. By Lemma 5, there exists a decreasing path from e to f going through vertices smaller than x. From what precedes, $f \cap V_{\geq x} \subseteq e$. Therefore $V(\mathcal{H}_e^x) \cap V_{\geq x} \subseteq e$. $\qquad\square$

8

3.2 Constructing the dec-DNNF.

Given a CNF-formula F with hypergraph $\mathcal{H}$, we can naturally define a family of subformulas F_e^x from $\mathcal{H}_e^x$ as the conjunction of clauses corresponding to the edges in $\mathcal{H}_e^x$, that is $F_e^x = \{C \in F \mid \mathsf{var}(C) \in \mathcal{H}_e^x\}$. Theorem 3 implies in particular that $\mathsf{var}(F_e^x) \subseteq (e \cup V_{<x})$. Thus, if τ is an assignment of variables $(e \cap V_{>x})$, then $F_e^x[\tau]$ has all its variables in $V_{\leq x}$. We will be particularly interested in such assignments: for a clause $C \in F$, denote by τ_C the only assignment of $\mathsf{var}(C)$ such that $\tau_C \not\models C$ and by $\tau_C^x := \tau_C|_{V_{>x}}$. We construct a dec-DNNF D by dynamic programming such that for each clause C with $\mathsf{var}(C) = e$ and variable $x \in V$, there exists a gate in D computing $F_e^x[\tau_C^x]$, which is a formula with variables in $V_{\leq x}$. Lemma 6 and Corollary 7 describe everything needed for the dynamic programming algorithm by expressing F_e^x as a decomposable conjunction of precomputed values.

Lemma 6. *Let $x \in \mathsf{var}(F)$ such that $x \neq x_1$ and let $y \in \mathsf{var}(F)$ be the predecessor of x for $<$. Let $e \in \mathcal{H}(F)$ and $\tau : (e \cap V_{\geq x}) \to \{0, 1\}$. Then either $F_e^x[\tau] \equiv 1$ or there exists $U \subseteq \mathcal{H}_e^x$ and for all $g \in U$ a clause $C(g) \in F_e^x$ with $\mathsf{var}(C(g)) = g$ such that*

$$F_e^x[\tau] \equiv \bigwedge_{g \in U} F_g^y[\tau_{C(g)}^y].$$

Moreover, this conjunction is decomposable and U can be found in polynomial time in $\mathsf{size}(F)$.

Proof. Assume first that for all $C \in F_e^x$, $\tau \models C$. Thus $F_e^x[\tau] \equiv 1$ since every clause of F_e^x is satisfied by τ.

Now assume that there exists $C \in F_e^x$ is such that $\tau \not\models C$. This means that $\tau \simeq \tau_C$. We let $A = \{\mathsf{var}(C) \mid C \in F_e^x \text{ and } \tau \not\models C\} \neq \emptyset$ by assumption. Observe that

$$F_e^x[\tau] \equiv \bigwedge_{\substack{C \in F_e^x \\ \mathsf{var}(C) \in A}} C[\tau]$$

since for every $C \in F_e^x$, if $\mathsf{var}(C) \notin A$, $\tau \models C$ by construction of A.

Let $U = \{g \in A \mid \forall f \in A \setminus \{g\}, g \notin \mathcal{H}_f^y\}$. For each $g \in U$, we choose an arbitrary clause $C(g)$ such that $\mathsf{var}(C(g)) = g$ and $\tau \not\models C(g)$. Such a clause exists since $U \subseteq A$. We claim that U meets the conditions given in the statement of the lemma.

We start by observing that U can be computed in polynomial time in $\mathsf{size}(F)$. Indeed, computing F_e^x for all e, x can be done in polynomial time as it boils down to a computation of connected component in a hypergraph. Now to compute A, it is enough to test for every $C \in F_e^x$ that $\tau \not\models C$ which can be done in polynomial time in $\mathsf{size}(F)$. Finally, extracting U from A can also be done in polynomial time by testing for every $g \in A$ if g respects the given condition: it is enough to test for every $f \in A \setminus \{g\}$ if $g \notin \mathcal{H}_f^y$, which is possible since we can compute $\mathcal{H}_f^y$ easily.

Now let $f \in A$. We show that there exists $g \in U$ such that $f \in \mathcal{H}_g^y$. If $f \in U$, then we are done since $f \in \mathcal{H}_f^y$. Now assume that $f \notin U$. By definition of U, $B = \{g \in A \setminus \{f\} \mid f \in \mathcal{H}_g^y\} \neq \emptyset$. We choose g to be the maximum of B for $\leq_{\mathcal{H}}$. We claim that $g \in U$. Indeed, assume there exists $g' \in A$ such that $g \in \mathcal{H}_{g'}^y$ and $g < g'$. By Lemma 2, $\mathcal{H}_g^y \subseteq \mathcal{H}_{g'}^y$ and since $f \in \mathcal{H}_g^y$, we also have $f \in \mathcal{H}_{g'}^y$, that is, $g' \in B$. Yet, $g = \max(B)$ and $g \leq g'$, that is, $g = g'$. Thus $g \in U$.

We thus have proved that for all $f \in A$, there exists $g \in U$ such that $f \in \mathcal{H}_g^y$. Thus if C is a clause of F_e^x, either $\mathsf{var}(C) \notin A$ and then $\tau \models C$ by definition of A, or $\mathsf{var}(C) \in A$, then there exists

9

$g \in U$ such that $\mathsf{var}(C) \in \mathcal{H}_g^y$, that is, $C \in F_g^y$. Now, if $C \in F_g^y$ for some $g \in U$, then $C \in F_e^x$ too since by Lemma 2, $F_g^y \subseteq F_e^x$. Thus

$$F_e^x[\tau] \equiv \bigwedge_{g \in U} F_g^y[\tau].$$

Let $g \in U$. We show that $\tau|_{\mathsf{var}(F_g^y)} = \tau_{C(g)}^y$. Observe that by Theorem 3, $\mathsf{var}(F_g^y) \cap V_{\geq x} = V(\mathcal{H}_g^y) \cap V_{\geq x} \subseteq g \cap V_{\geq x}$. Since τ assigns variables from $e \cap V_{\geq x}$:

$$\tau|_{\mathsf{var}(F_g^y)} = \tau|_{\mathsf{var}(F_g^y) \cap V_{\geq x} \cap e}$$
$$= \tau|_{g \cap V_{\geq x} \cap e}$$

Moreover, since $g \in \mathcal{H}_e^x$, we have $g \cap V_{\geq x} \subseteq e \cap V_{\geq x}$ by Theorem 3 again. Thus $g \cap V_{\geq x} \cap e = g \cap V_{\geq x}$. In other words, $\tau|_{\mathsf{var}(F_g^y)} = \tau|_{g \cap V_{\geq x}}$.

Since τ assigns all variables of $e \cap V_{\geq x}$ by assumption, $\tau|_{g \cap V_{\geq x}}$ assigns all variables of $g \cap V_{\geq x}$. Finally, since $\tau \not\models C(g)$ by construction of $C(g)$, we have $\tau \simeq \tau_{C(g)}^y$. Since by definition $\mathsf{var}(C(g)) = g$, it follows that $\tau|_{\mathsf{var}(F_g^y)} = \tau_{C(g)}^y$. So far, we have proven that

$$F_e^x[\tau] \equiv \bigwedge_{g \in U} F_g^y[\tau_{C(g)}^y].$$

It remains to show that this conjunction is decomposable, that is, for all $g_1, g_2 \in U$, $\mathsf{var}(F_{g_1}^y[\tau_{C(g_1)}^y]) \cap \mathsf{var}(F_{g_2}^y[\tau_{C(g_2)}^y]) = \emptyset$. Let $g_1, g_2 \in U$ with $g_1 <_{\mathcal{H}} g_2$ and assume there exists $z \in \mathsf{var}(F_{g_1}^y[\tau_{C(g_1)}^y]) \cap \mathsf{var}(F_{g_2}^y[\tau_{C(g_2)}^y])$, that is, $z \in \mathsf{var}(F_{g_1}^y) \cap \mathsf{var}(F_{g_2}^y) \cap V_{\leq y}$. From what precedes, τ assigns every variable of $F_{g_1}^y$ greater than x. By Lemma 2, we have $F_{g_1}^y \subseteq F_{g_2}^y$, which contradicts the fact that $g_1 \in U$. $\square$

Corollary 7. *Let $x \in \mathsf{var}(F)$ such that $x \neq x_1$ and let $y \in \mathsf{var}(F)$ be the predecessor of x for $<$. For every $C \in \mathcal{H}(F)$, there exist $U_0, U_1 \subseteq \mathcal{H}_{\mathsf{var}(C)}^x$ and for all $g \in U_0 \cup U_1$ a clause $C(g) \in F_{\mathsf{var}(C)}^x$ with $\mathsf{var}(C(g)) = g$ such that*

$$F_{\mathsf{var}(C)}^x[\tau_C^x] \equiv \left(x \wedge \bigwedge_{g \in U_1} F_g^y[\tau_{C(g)}^y]\right) \vee \left(\neg x \wedge \bigwedge_{g \in U_0} F_g^y[\tau_{C(g)}^y]\right).$$

Moreover, all conjunctions are decomposable and U_0, U_1 can be found in polynomial time in $\mathsf{size}(F)$.

Proof. Let $\tau_1 = \tau_C^x \cup \{x \mapsto 1\}$ and $\tau_0 = \tau_C^x \cup \{x \mapsto 0\}$. We observe that

$$F_{\mathsf{var}(C)}^x[\tau_C^x] = (x \wedge F_{\mathsf{var}(C)}^x[\tau_1]) \vee (\neg x \wedge F_{\mathsf{var}(C)}^x[\tau_0]).$$

Clearly, $x \notin \mathsf{var}(F_{\mathsf{var}(C)}^x[\tau_1])$ and $x \notin \mathsf{var}(F_{\mathsf{var}(C)}^x[\tau_0])$, thus, both conjunctions are decomposable. Now, applying Lemma 6 on $F_{\mathsf{var}(C)}^x[\tau_0]$ and on $F_{\mathsf{var}(C)}^x[\tau_1]$ yields the desired decomposition. $\square$

Theorem 8. *Let F be a β-acyclic CNF-formula. One can construct in polynomial time in $\mathsf{size}(F)$ a dec-DNNF D of size $O(\mathsf{size}(F))$ and fanin at most $|\mathcal{H}|$ computing F.*

Proof. Let $\mathcal{H}$ be the hypergraph of F and $<$ a β-elimination order. Let $\mathsf{var}(F) = \{x_1, \ldots, x_n\}$ where $x_i < x_j$ if and only if $i < j$. We construct by induction on i a dec-DNNF D_i of fanin $|\mathcal{H}|$ at most such that for each $e \in \mathcal{H}$, $C \in F$ such that $\mathsf{var}(C) = e$ and $j \leq i$, there exists a gate in D_i

computing $F_e^{x_j}[\tau_C^{x_j}]$ and $|D_i| \leq 7 \cdot (\sum_{j=1}^i c(x_j))$ where $c(x_j)$ is the number of clauses in F holding x_j.

We start by explaining how D_1 is constructed. Let $e \in \mathcal{H}$. If $x_1 \notin e$, then $F_e^{x_1}$ contains only the clauses C such that $e = \mathsf{var}(C)$. For such a C, $\tau_C^{x_1} = \tau_C$, thus $F_e^{x_1}[\tau_C] = 0$. Now, if $x_1 \in e$, $F_e^{x_1}$ contains only clauses D such that $x_1 \in \mathsf{var}(D) \subseteq e$ since x_1 is the first element of the elimination order. Let C be a clause such that $\mathsf{var}(C) = e$. For every $D \in F_e^{x_1}$, $\mathsf{var}(D) \subseteq \mathsf{var}(C)$, thus $F_e^{x_1}[\tau_C^{x_1}]$ has only one variable: x_1. Thus $F_e^{x_1}[\tau_C^{x_1}]$ is equivalent to either $x_1, \neg x_1$ or 0. We thus define D_1 to be the dec-DNNF with at most three gates $x_1, \neg x_1$ and 0, which are input gates. We have $|D_1| \leq 3 \leq 7 \cdot c(x_1)$.

Now lets assume D_i is constructed. To ease notations, let $x = x_{i+1}$. Let $e \in \mathcal{H}$ and C be a clause such that $\mathsf{var}(C) = e$. We want to add a gate in D_i that will compute $F_e^x[\tau_C^x]$. If $x \notin e$, then $\mathcal{H}_e^x = \mathcal{H}_e^{x_i}$ since by Theorem 3, $\mathsf{var}(\mathcal{H}_e^{x_i}) \subseteq (e \cup V_{<x_i})$. Thus $F_e^x = F_e^{x_i}$ and $\tau_C^x = \tau_C^{x_i}$. Therefore, there is already a gate computing $F_e^x[\tau_C^x]$ in D_i.

Assume now that $x \in e$. By Corollary 7, we can compute $F_{\mathsf{var}(C)}^x[\tau_C^x]$ for every C with $\mathsf{var}(C) = e$ by adding at most one decision-gate and a fanin $|\mathcal{H}|$ decomposable and-gate to D_i since for all values appearing in the statement of Corollary 7 there exists a gate in D_i computing it. Moreover such gate can be found in polynomial time. That is, we add to D_i at most 7 gates to compute $F_{\mathsf{var}(C)}^x[\tau_C^x]$. We have to do this for each $C \in F$ such that $x \in \mathsf{var}(C)$. We thus add at most $7c(x)$ gates in D_i. Thus $|D_{i+1}| \leq 7 \cdot \sum_{j \leq i+1} c(x_j)$.

To conclude, assume that $\mathcal{H}$ is connected and let $e = \max(\mathcal{H})$. We have $\mathcal{H}_e^{x_n} = \mathcal{H}$ since there is a path from e to every other edge in $\mathcal{H}$. Thus $F_e^{x_n} = F$. Let C be a clause with $\mathsf{var}(C) = e$. The assignment $\tau_C^{x_n}$ is empty, thus $F_e^{x_n}[\tau_C^{x_n}] \equiv F$. Hence, there is a gate in D_n that computes F and D_n is of size at most $7 \cdot \mathsf{size}(F)$ and fanin $|\mathcal{H}|$ at most. Each step can be done in polynomial time in $\mathsf{size}(F)$.

If $\mathcal{H}$ is not connected, then each connected component of $\mathcal{H}$ is β-acyclic, thus we can compile them independently and take the decomposable conjunction of these dec-DNNF. $\square$

We conclude this section by giving insights on the significance of Theorem 8 from a practical point of view. Most practical tools for #SAT are based on an algorithm called exhaustive DPLL with caching [HD05, SBB$^+$04, Thu06, BDP03] which works as follows: given F, the algorithm starts by trying to write F as $F_1 \wedge F_2$ with F_1 and F_2 having no common variables. If it succeeds, it computes recursively $\#F_1$, $\#F_2$ and returns $\#F_1 \cdot \#F_2$. Otherwise, it chooses a variable x and returns $\#F[x \mapsto 0] + \#F[x \mapsto 1]$. In addition, these tools use caching techniques to avoid redoing the same computation twice. It was observed in [HD05] that the trace of such algorithms is exactly a dec-DNNF. It is not hard to see that the construction given in Theorem 8 is the trace of a run of an exhaustive DPLL algorithm where the variables are chosen in a reverse β-elimination order. This shows that if the right elimination order of the variables is chosen (and this order can be computed greedily in polynomial time), then practical tools for solving #SAT can in theory solve β-acyclic formulas in polynomial time.

4   Deviation from the technique based on branch decompositions

In this section, we finally prove that standard techniques based on branch decompositions fail on β-acyclic formulas. Recall that we have defined in Section 2.6 the standard technique to be the

implicit construction of a polynomial size structured DNNF equivalent to the input formula. We formally prove the following:

Theorem 9. *There exists an infinite family $\mathcal{F}$ of β-acyclic CNF-formulas such that for every $F \in \mathcal{F}$ having n variables, there is no structured DNNF of size less than $2^{\Omega(\sqrt{n})}$ computing F.*

We use techniques based on communication complexity tools developed in [BCMS16] to prove lower bounds on the size of structured DNNF.

Definition 10. *Let r be a boolean function on variables X and let (Y, Z) be a partition of X. The function r is a (Y, Z)-rectangle if and only if for every $\tau, \tau' \in \{0, 1\}^X$ such that $\tau \models r$ and $\tau' \models r$, we have $(\tau|_Y \cup \tau'|_Z) \models r$. A (Y, Z)-rectangle cover of a boolean function f is a set $R = \{r_1, \dots, r_q\}$ of (Y, Z)-rectangles such that $\mathsf{sat}(f) = \bigcup_{i=1}^{q} \mathsf{sat}(r_i)$.*

Theorem 11 ([BCMS16],[PD10])**.** *Let D be a DNNF on variables X respecting the vtree T. For every vertex t of T, there exists a $(X_t, X \setminus X_t)$-rectangle cover of D of size at most $|D|$, where $X_t = \mathsf{var}(T_t)$.*

Given a CNF-formula F, we define $\hat{F}$ to be the formula $\{K \cup \{c_K\} \mid K \in F\}$ on variables $\{c_K \mid K \in F\} \cup \mathsf{var}(F)$. Intuitively, $\hat{F}$ is the formula obtained by adding one fresh variable c_K in each clause K of F. Our main lower bound relates the incidence MIM-width of a monotone CNF-formula to the size of structured DNNF computing $\hat{F}$.

Theorem 12. *Let F be a monotone formula of incidence MIM-width k. Any structured DNNF computing $\hat{F}$ is of size at least $2^{k/2}$.*

The proof of Theorem 12 heavily relies on the following lower bound and on Theorem 11:

Lemma 13. *Let $X = \{x_1, \dots, x_k\}$ and $Y = \{y_1, \dots, y_k\}$ be two disjoint sets of k variables. The number of (X, Y)-rectangles needed to cover the CNF-formula $F = \bigwedge_{i=1}^{k}(x_i \vee y_i)$ is at least 2^k.*

Proof. Let $\{R_1, \dots, R_q\}$ be a (X, Y)-rectangle cover of F. For $K \subseteq \{1, \dots, k\}$, we denote by τ_K the assignment such that $\tau_K(x_i) = 1$ if $i \in K$ and 0 otherwise and $\tau_K(y_i) = 1 - \tau_K(x_i)$. Observe that by definition, for every $K \subseteq \{1, \dots, k\}$, $\tau_K \models F$. We claim that if $\tau_K \models R_i$ then for any $K' \neq K$, we have $\tau_{K'} \not\models R_i$. For the sake of contradiction, assume there exist K, K' such that $K \neq K'$, $\tau_K \models R_i$ and $\tau_{K'} \models R_i$. Without loss of generality, we can assume that there exists $i \in K \setminus K'$. By definition of rectangles, $\tau' = \tau_{K'}|_X \cup \tau_K|_Y \models R_i$. But $\tau'(x_i) = \tau'(y_i) = 0$ and then $\tau' \not\models F$ which contradicts the definition of R_i. Since there are 2^k different subsets of $\{1, \dots, k\}$ and each τ_K satisfies disjoint rectangles, we have that $q \geq 2^k$. $\square$

Proof (of Theorem 12). Let $G = \mathcal{G}_{inc}(F)$ and D be a structured DNNF computing $\hat{F}$. We claim that $|D| \geq 2^{k/2}$.

Let T be the vtree respected by D. Observe that the variables of $\hat{F}$ are in one to one correspondence with $V(G)$ thus we can see T as a branch decomposition of G. Since G is of MIM-width k, there exists a vertex t of T such that there is an induced matching $M = \{(x_1, y_1), \dots, (x_q, y_q)\}$ with $q \geq k$ in $G[V_t, V(G) \setminus V_t]$ where V_t denotes the labels of the leaves of T_t. Let $e = (x, y)$ be an edge of M. Since it is an edge of G, too, one end point of e corresponds to a variable x_e of F and the other to a clause $c_e \in F$. Let M' be the set of edges e of M such that $x_e \in V_t$ and $c_e \notin V_t$ and let M'' be the set of edges e of M such that $x_e \notin V_t$ and $c_e \in V_t$. It is readily verified that $M = M' \uplus M''$.

Let N be the largest of these two sets. N is thus an induced matching of $G[V_t, V(G) \setminus V_t]$ of size at least $k/2$. Moreover, if $e, e' \in N$ are distinct, we have $x_{e'} \notin c_e$. Indeed, if $x_{e'} \in c_e$ then they are connected by an edge of G and this edge is across V_t and $V(G) \setminus V_t$ by construction of N. Thus, if such an edge exists, it violates the assumption that N is an induced matching of $G[V_t, V(G) \setminus V_t]$.

Now let τ be the following partial assignment of $\mathsf{var}(\hat{F})$: if C is a clause that does not appear in N, we let $\tau(C) = 1$. If x is a variable of F that does not appear in N, we let $\tau(x) = 0$. We claim that $\hat{F}[\tau] \equiv \bigwedge_{e \in N}(x_e \vee c_e)$. Indeed, each clause C that does not appear in N is already satisfied in $\hat{F}[\tau]$ since $\tau(C) = 1$ and for the remaining clauses, the variables that do not appear in N disappear as they are set to 0 (remember that F is monotone). Moreover, if $e, e' \in N$ are distinct edges of N, we have that $x_e \notin c_{e'}$ thus the only variables remaining in the clause c_e is x_e for each $e \in N$.

Now since $\hat{F}$ is computed by D, $\hat{F}[\tau]$ is computed by $D' = D[\tau]$ which is a structured DNNF smaller than D. By Theorem 11, there is a $(V_t, V(G) \setminus V_t)$-rectangle cover of D' of size at most $\mathsf{size}(D')$ and by Lemma 13, we need at least $2^{|N|} \geq 2^{k/2}$ rectangles to cover $F[\tau]$. Thus, $\mathsf{size}(D) \geq \mathsf{size}(D') \geq 2^{k/2}$. □

Theorem 9 is a corollary of Theorem 12 and Theorem 1:

Proof of Theorem 9. Let F be a β-acyclic formula. We claim that $\hat{F}$ is also β-acyclic. Indeed, let $(x_1, \ldots, x_n)$ be a β-elimination order for $\mathcal{H}(F)$. We claim that $(c_1, \ldots, c_m, x_1, \ldots, x_n)$ is a β-elimination order of $\mathcal{H}(\hat{F})$ where $c_1, \ldots, c_m$ are the variables of $\hat{F}$ corresponding to the clauses of F. Indeed, for all i, c_i is in exactly one edge of $\mathcal{H}(\hat{F})$ and can thus be eliminated from the start. Finally, $\mathcal{H}(\hat{F}) \setminus \{c_1, \ldots, c_m\} = \mathcal{H}(F)$, thus $(x_1, \ldots, x_n)$ is a β-elimination order of $\mathcal{H}(\hat{F}) \setminus \{c_1, \ldots, c_m\}$.

To every hypergraph $\mathcal{H}$, we can associate a monotone formula $\mathsf{CNF}(\mathcal{H})$ whose variables are the vertices of $\mathcal{H}$ and clauses are the edges of $\mathcal{H}$ without negations. It is readily verified that the hypergraph of $\mathsf{CNF}(\mathcal{H})$ is $\mathcal{H}$. Let $\mathcal{G}$ be the family of β-acyclic hypergraphs with MIM-width of $\Omega(n)$ from Theorem 1 and let $\mathcal{F} = \{\widehat{\mathsf{CNF}(\mathcal{H})} \mid \mathcal{H} \in \mathcal{G}\}$. From what precedes, $\mathcal{F}$ is a family of β-acyclic hypergraphs and by Theorem 12, if $F \in \mathcal{F}$ has m clauses and $N = n + m$ variables then any structured DNNF computing F is of size at least $2^{\Omega(n)}$. The statement of Theorem 9 follows since the number of edges in a β-acyclic hypergraph with n vertices is at most $n(n+1)/2$ (Remark 13 in [BB14]). Thus, $N = O(n^2)$, *i.e.* $n = \Omega(\sqrt{N})$. □

5 Discussion

We discuss here further directions that can be studied from the results presented in this paper. In Section 4, we have shown that β-acyclic formulas cannot be compiled into structured DNNF contrary to other known tractable classes. It would be interesting to study the opposite question, that is, to understand if classes tractable with respect to the STV-algorithm can be compiled into dec-DNNF. A positive answer to this question would open interesting perspectives as it would imply that all known tractable structural restrictions for #SAT can be processed using exhaustive DPLL with caching, which could lead to a practical use of such theoretical result and to the design of interesting heuristic for the order in which variables are eliminated in DPLL based on structural restrictions. A negative answer would show that some "easy" cases are missed by practical tools and that it would be worth investing time to develop practical tools taking the formula structure into account.

Another direction is suggested by Theorem 12 which says that the MIM-width of the formula is closely related to the size of the smallest structured DNNF for $\hat{F}$. The most general graph parameter

that is known to lead to polynomial time execution with the STV-algorithm is the MIM-width: #SAT can be solved in time $m^{\Omega(k)}\mathsf{poly}(n+m)$ on a formula with m clauses, n variables and of MIM-width k. Theorem 12 almost proves the optimality of such running time for compilation into structured DNNF.

References

[BB14]   Johann Brault-Baron. Hypergraph Acyclicity Revisited. *ArXiv e-prints*, March 2014.

[BCM15]   Johann Brault-Baron, Florent Capelli, and Stefan Mengel. Understanding Model Counting for beta-acyclic CNF-formulas. In *32nd International Symposium on Theoretical Aspects of Computer Science*, volume 30 of *LIPIcs*, pages 143–156. Schloss Dagstuhl, 2015.

[BCMS15]   Simone Bova, Florent Capelli, Stefan Mengel, and Friedrich Slivovsky. On Compiling CNFs into Structured Deterministic DNNFs. In *Theory and Applications of Satisfiability Testing*, Lecture Notes in Computer Science, pages 199–214. Springer International Publishing, September 2015.

[BCMS16]   Simone Bova, Florent Capelli, Stefan Mengel, and Friedrich Slivovsky. Knowledge Compilation Meets Communication Complexity. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 1008–1014, 2016.

[BDP03]   Fahiem Bacchus, Shannon Dalmao, and Toniann Pitassi. Algorithms and Complexity Results for #SAT and Bayesian Inference. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '03, pages 340–, Washington, DC, USA, 2003. IEEE Computer Society.

[BL15]   Paul Beame and Vincent Liew. New limits for knowledge compilation and applications to exact model counting. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 131–140, 2015.

[BLRS13]   Paul Beame, Jerry Li, Sudeepa Roy, and Dan Suciu. Lower bounds for exact model counting and applications in probabilistic databases. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, 2013.

[BLRS14]   Paul Beame, Jerry Li, Sudeepa Roy, and Dan Suciu. Counting of query expressions: Limitations of propositional methods. In *Proc. 17th International Conference on Database Theory (ICDT)*, pages 177–188, 2014.

[CDM14]   Florent Capelli, Arnaud Durand, and Stefan Mengel. Hypergraph Acyclicity and Propositional Model Counting. In *Theory and Applications of Satisfiability Testing - SAT 2014 - 17th International Conference*, pages 399–414, 2014.

[Dar01]   Adnan Darwiche. Decomposable Negation Normal Form. *J. ACM*, 48(4):608–647, 2001.

[Die12]   Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.

[DM02]   Adnan Darwiche and Pierre Marquis. A Knowledge Compilation Map. *Journal of Artificial Intelligence Research*, 17:229–264, 2002.

[DP60]   Martin Davis and Hilary Putnam. A Computing Procedure for Quantification Theory. *J. ACM*, 7(3):201–215, July 1960.

[Fag83]   Ronald Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *Journal of the ACM*, 30(3):514–550, 1983.

[HD05]   Jinbo Huang and Adnan Darwiche. DPLL with a trace: From SAT to knowledge compilation. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 156–162, 2005.

[HSTV14]   Sigve Hortemo Sæther, Jan Arne Telle, and Martin Vatshelle. Solving MaxSAT and #SAT on Structured CNF Formulas. In *Theory and Applications of Satisfiability Testing*, pages 16–31, 2014.

[JS13]   Abhay Kumar Jha and Dan Suciu. Knowledge compilation meets database theory: Compiling queries to decision diagrams. *Theory Comput. Syst.*, 52(3):403–440, 2013.

[OPS13]   S. Ordyniak, D. Paulusma, and S. Szeider. Satisfiability of acyclic and almost acyclic CNF formulas. *Theoretical Computer Science*, 481:85–99, 2013.

[PD08]   Knot Pipatsrisawat and Adnan Darwiche. New Compilation Languages Based on Structured Decomposability. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI*, pages 517–522, 2008.

[PD10]   Thammanit Pipatsrisawat and Adnan Darwiche. A Lower Bound on the Size of Decomposable Negation Normal Form. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*, July 2010.

[PSS16]   Daniël Paulusma, Friedrich Slivovsky, and Stefan Szeider. Model Counting for CNF Formulas of Bounded Modular Treewidth. *Algorithmica*, 76(1):168–194, 2016.

[Rot96]   Dan Roth. On the Hardness of Approximate Reasoning. *Artificial Intelligence*, 82(12):273 – 302, 1996.

[SBB+04]   Tian Sang, Fahiem Bacchus, Paul Beame, Henry A Kautz, and Toniann Pitassi. Combining Component Caching and Clause Learning for Effective Model Counting. *Theory and Applications of Satisfiability Testing*, 4:7th, 2004.

[SBK05]   Tian Sang, Paul Beame, and Henry A. Kautz. Heuristics for fast exact model counting. In *Theory and Applications of Satisfiability Testing, 8th International Conference, SAT 2005, St. Andrews, UK, June 19-23, 2005, Proceedings*, pages 226–240, 2005.

[SS10]   Marko Samer and Stefan Szeider. Algorithms for propositional model counting. *Journal of Discrete Algorithms*, 8(1):50–64, 2010.

[SS13]   Friedrich Slivovsky and Stefan Szeider. Model Counting for Formulas of Bounded Clique-Width. In *Algorithms and Computation - 24th International Symposium, ISAAC*, pages 677–687, 2013.

[Thu06]   Marc Thurley. SharpSAT – Counting Models with Advanced Component Caching and Implicit BCP. In *Theory and Applications of Satisfiability Testing*, pages 424–429. Springer, 2006.

[Vat12]   Martin Vatshelle. *New Width Parameters of Graphs.* PhD thesis, University of Bergen, 2012.

[Weg00]   Ingo Wegener. *Branching Programs and Binary Decision Diagrams.* SIAM, 2000.