

On Compiling Structured CNFs to OBDDs

Simone Bova¹ · Friedrich Slivovsky¹

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract We present new results on the size of OBDD representations of structurally characterized classes of CNF formulas. First, we prove that *variable convex* formulas (that is, formulas with incidence graphs that are convex with respect to the set of variables) have polynomial OBDD size. Second, we prove an exponential lower bound on the OBDD size of a family of CNF formulas with incidence graphs of bounded degree. We obtain the first result by identifying a simple sufficient condition—which we call the *few subterms* property—for a class of CNF formulas to have polynomial OBDD size, and show that variable convex formulas satisfy this condition. To prove the second result, we exploit the combinatorial properties of expander graphs; this approach allows us to establish an exponential lower bound on the OBDD size of formulas satisfying strong syntactic restrictions.

Keywords Knowledge compilation · Ordered binary decision diagram · Conjunctive normal form · Expander graphs

1 Introduction

The goal of *knowledge compilation* is to succinctly represent propositional knowledge bases in a format that supports a number of queries in polynomial time [8].

✉ Friedrich Slivovsky
fslivovsky@gmail.com

Simone Bova
bova@ac.tuwien.ac.at

¹ Institute of Computer Graphics and Algorithms, TU Wien, 1040 Vienna, Austria

Choosing a representation language generally involves a trade-off between succinctness and the range of queries that can be efficiently answered. In this paper, we study ordered binary decision diagram (OBDD) representations of propositional theories given as formulas in conjunctive normal form (CNF). Binary decision diagrams (also known as branching programs) and their variants are widely used and well-studied representation languages for Boolean functions [25]. OBDDs in particular enjoy properties, such as polynomial-time equivalence testing, that make them the data structure of choice for a range of applications.

The question of which classes of CNFs can be represented as (or *compiled into*, in the jargon of knowledge representation) OBDDs of polynomial size is largely unexplored [25, Chapter 4]. We approach this classification problem by considering *structurally* characterized CNF classes, more specifically, classes of CNF formulas defined in terms of properties of their *incidence graphs* (the incidence graph of a formula is the bipartite graph on clauses and variables where a variable is adjacent to the clauses it occurs in). Figure 1 depicts a hierarchy of well-studied bipartite graph classes as considered by Lozin and Rautenbach [20, Fig. 2]. This hierarchy is particularly well-suited for our classification project as it includes prominent cases such as beta acyclic CNFs [5] and bounded clique-width CNFs. When located within this hierarchy, the known bounds on the OBDD size of structural CNF classes leave a large gap (depicted on the left of Fig. 1):

- On the one hand, we have a polynomial upper bound on the OBDD size of bounded treewidth CNF classes proved recently by Razgon [23]. The corresponding graph classes are located at the bottom of the hierarchy.
- On the other hand, there is an exponential lower bound for the OBDD size of general CNFs, proved two decades ago by Devadas [9]. The corresponding graph




Fig. 1 The diagram depicts a hierarchy of classes of bipartite graphs under the inclusion relation (thin edges). \mathcal{B} , \mathcal{H} , \mathcal{D}_k , \mathcal{C} , \mathcal{C}_v , and \mathcal{C}_c denote, respectively, bipartite graphs, chordal bipartite graphs (corresponding to beta acyclic CNFs), bipartite graphs of degree at most k ($k \geq 3$), convex graphs, *left* (variable) convex graphs, and *right* (clause) convex graphs. The class $\mathcal{C}_v \cap \mathcal{C}_c$ of biconvex graphs and the class \mathcal{D}_k of bipartite graphs of degree at most k have unbounded clique-width. The class $\mathcal{H} \cap \mathcal{D}_k$ of chordal bipartite graphs of degree at most k has bounded treewidth. The *green* and *red* curved lines enclose, respectively, classes of incidence graphs whose CNFs have polynomial time OBDD compilation, and classes of incidence graphs whose CNFs have exponential size OBDD representations; the right hand picture shows the compilability frontier, updated in light of Results 1 and 2

class is not chordal bipartite, has unbounded degree and unbounded clique-width, and hence is located at the top of the hierarchy.

Contribution In this paper, we tighten this gap as illustrated *on the right* in Fig. 1. More specifically, we prove new bounds for two structural classes of CNFs.

Result 1 CNF formulas with *variable convex* incidence graphs have polynomial OBDD size (Theorem 7).

Convexity is a property of bipartite graphs that has been extensively studied in the area of combinatorial optimization [14, 15, 24], and that can be detected in linear time [4, 19].

To prove Result 1, we identify a property of CNF classes—called the *few subterms property*—that is sufficient for polynomial-size compilability (Theorem 4), and then prove that CNFs with variable convex incidence graphs have this property (Lemma 6). The few subterms property naturally arises as a sufficient condition for polynomial size compilability when considering OBDD representations of CNF formulas (cf. Oztok and Darwiche’s recent work on *CV-width* [22], which explores a similar idea). Aside from its role in proving polynomial-size compilation for variable convex CNFs, the few subterms property can also be used to explain the (known) fact that classes of CNFs with incidence graphs of *bounded treewidth* have OBDD representations of polynomial size (Lemma 10), and as such offers a unifying perspective on these results. Both the result on variable convex CNFs and the result on bounded treewidth CNFs can be improved to polynomial *time* compilation by appealing to a stronger version of the few subterms property (Theorems 7 and 11).

In an attempt to push the few subterms property further, we adopt the language of *parameterized complexity* to formally capture the idea that CNFs “close” to a class with few subterms have “small” OBDD representations. More precisely, defining the *deletion distance* of a CNF from a CNF class as the number of its variables or clauses that have to be deleted in order for the resulting formula to be in the class, we prove that CNFs have fixed-parameter tractable OBDD size parameterized by the deletion distance from a CNF class with few subterms (Theorem 13). This result can again be improved to fixed-parameter *time* compilation under additional assumptions (Theorem 14), yielding for instance fixed-parameter tractable time compilation of CNFs into OBDDs parameterized by the *feedback vertex set* size (Corollary 15).

On the negative side, we prove that some structurally characterized classes of CNF formulas do not have small OBDD representations:

Result 2 There is a class of CNF formulas with incidence graphs of bounded degree such that every formula F in this class has OBDD size at least $2^{\Omega(\text{size}(F))}$ (Theorem 19).¹

¹Here, $\text{size}(F)$ denotes the number of variable occurrences in the formula F .

This substantially improves a $2^{\Omega(\sqrt{\text{size}(F)})}$ lower bound for the OBDD size of a class of CNFs by Devadas [9]. Moreover, we establish this bound for a class that satisfies strong syntactic restrictions: every clause contains exactly two positive literals and each variable occurs at most 3 times.

The heavy lifting in our proof of this result is done by a family of *expander graphs*. Expander graphs have found applications in many areas of mathematics and computer science [16, 21], including circuit and proof complexity [17]. In this paper, we show how they can be used to derive lower bounds for OBDDs.

Organization The remainder of this paper is organized as follows. In Section 2, we introduce basic notation and terminology. In Section 3, we prove that formulas with few subterms have polynomial OBDD size and show that variable-convex CNFs (as well as bounded treewidth CNFs) enjoy the few subterms property. Fixed-parameter tractable size and time compilability results based on the few subterms property are presented in Section 3.4. In Section 4, we prove a strongly exponential lower bound on the OBDD size of CNF formulas based on expander graphs. We conclude in Section 5.

2 Preliminaries

Formulas Let X be a countable set of *variables*. A *literal* is a variable x or a negated variable $\neg x$. If x is a variable we let $\text{var}(x) = \text{var}(\neg x) = x$. A *clause* is a finite set of literals. For a clause c we define $\text{var}(c) = \{\text{var}(l) \mid l \in c\}$. If a clause contains a literal negated as well as unnegated it is *tautological*. A *conjunctive normal form (CNF)* is a finite set of non-tautological clauses. If F is a CNF formula we let $\text{var}(F) = \bigcup_{c \in F} \text{var}(c)$. The *size* of a clause c is $|c|$, and the *size* of a CNF F is $\text{size}(F) = \sum_{c \in F} |c|$. An *assignment* is a mapping $f: X' \rightarrow \{0, 1\}$, where $X' \subseteq X$. We identify f with the set $\{\neg x \mid x \in X', f(x) = 0\} \cup \{x \mid x \in X', f(x) = 1\}$. An assignment f *satisfies* a clause c if $f \cap c \neq \emptyset$; for a CNF F , we let $F[f]$ denote the CNF containing the clauses in F not satisfied by f , restricted to variables in $X \setminus \text{var}(f)$, that is, $F[f] = \{c \setminus \{x, \neg x \mid x \in \text{var}(f)\} \mid c \in F, f \cap c = \emptyset\}$; then, f *satisfies* F if $F[f] = \emptyset$, that is, if it satisfies all clauses in F . If F is a CNF with $\text{var}(F) = \{x_1, \dots, x_n\}$ we define the Boolean function $F(x_1, \dots, x_n)$ *computed by* F as $F(b_1, \dots, b_n) = 1$ if and only if the assignment $f_{(b_1, \dots, b_n)}: \text{var}(F) \rightarrow \{0, 1\}$ given by $f_{(b_1, \dots, b_n)}(x_i) = b_i$ satisfies the CNF F .

Binary Decision Diagrams A *binary decision diagram (BDD)* D on variables $\{x_1, \dots, x_n\}$ is a labelled directed acyclic graph satisfying the following conditions: D has at most two vertices without outgoing edges, called *sinks* of D . Sinks of D are labelled with 0 or 1; if there are exactly two sinks, one is labelled with 0 and the other is labelled with 1. Moreover, D has exactly one vertex without incoming edges, called the *source* of D . Each non-sink node of D is labelled by a variable x_i , and has exactly two outgoing edges, one labelled 0 and the other labelled 1. Each node v of D represents a Boolean function $F_v = F_v(x_1, \dots, x_n)$ in the following way. Let $(b_1, \dots, b_n) \in \{0, 1\}^n$ and let w be a node labelled with x_i . We say that

(b_1, \dots, b_n) activates an outgoing edge of w labelled with $b \in \{0, 1\}$ if $b_i = b$. Since (b_1, \dots, b_n) activates exactly one outgoing edge of each non-sink node, there is a unique sink that can be reached from v along edges activated by (b_1, \dots, b_n) . We let $F_v(b_1, \dots, b_n) = b$, where $b \in \{0, 1\}$ is the label of this sink. The function computed by D is F_s , where s denotes the (unique) source node of D . The size of a BDD is the number of its nodes.

An ordering σ of a set $\{x_1, \dots, x_n\}$ is a total order on $\{x_1, \dots, x_n\}$. If σ is an ordering of $\{x_1, \dots, x_n\}$ we let $\text{var}(\sigma) = \{x_1, \dots, x_n\}$. Let σ be the ordering of $\{x_1, \dots, x_n\}$ given by $x_{i_1} < x_{i_2} < \dots < x_{i_n}$. For every integer $0 < j \leq n$, the length j prefix of σ is the ordering of $\{x_{i_1}, \dots, x_{i_j}\}$ given by $x_{i_1} < \dots < x_{i_j}$. A prefix of σ is a length j prefix of σ for some integer $0 < j \leq n$. For orderings $\sigma = x_{i_1} < \dots < x_{i_n}$ of $\{x_1, \dots, x_n\}$ and $\rho = y_{i_1} < \dots < y_{i_m}$ of $\{y_1, \dots, y_m\}$, we let $\sigma\rho$ denote the ordering of $\{x_1, \dots, x_n, y_1, \dots, y_m\}$ given by $x_{i_1} < \dots < x_{i_n} < y_{i_1} < \dots < y_{i_m}$. Let D be a BDD on variables $\{x_1, \dots, x_n\}$ and let $\sigma = x_{i_1} < \dots < x_{i_n}$ be an ordering of $\{x_1, \dots, x_n\}$. The BDD D is a σ -ordered binary decision diagram (σ -OBDD) if $x_i < x_j$ (with respect to σ) whenever D contains an edge from a node labelled with x_i to a node labelled with x_j . A BDD D on variables $\{x_1, \dots, x_n\}$ is an ordered binary decision diagram (OBDD) if there is an ordering σ of $\{x_1, \dots, x_n\}$ such that D is a σ -OBDD. For a Boolean function $F = F(x_1, \dots, x_n)$, the OBDD size of F is the size of the smallest OBDD on $\{x_1, \dots, x_n\}$ computing F .

We say that a class \mathcal{F} of CNFs has polynomial-time compilation into OBDDs if there is a polynomial-time algorithm that, given a CNF $F \in \mathcal{F}$, returns an OBDD computing the same Boolean function as F . We say that a class \mathcal{F} of CNFs has polynomial size compilation into OBDDs if there exists a polynomial $p: \mathbb{N} \rightarrow \mathbb{N}$ such that, for all CNFs $F \in \mathcal{F}$, there exists an OBDD of size at most $p(\text{size}(F))$ that computes the same function as F .

Graphs For standard graph theoretic terminology, see [10]. Let $G = (V, E)$ be a graph. The (open) neighborhood of W in G , in symbols $\text{neigh}(W, G)$, is defined by

$$\text{neigh}(W, G) = \{v \in V \setminus W \mid \text{there exists } w \in W \text{ such that } vw \in E\}.$$

We freely use $\text{neigh}(v, G)$ as a shorthand for $\text{neigh}(\{v\}, G)$, and we write $\text{neigh}(W)$ instead of $\text{neigh}(W, G)$ if the graph G is clear from the context. A graph $G = (V, E)$ is bipartite if its vertex set V can be partitioned into two blocks V' and V'' such that, for every edge $vw \in E$, we either have $v \in V'$ and $w \in V''$, or $v \in V''$ and $w \in V'$. In this case we may write $G = (V', V'', E)$. The incidence graph of a CNF F , in symbols $\text{inc}(F)$, is the bipartite graph $(\text{var}(F), F, E)$ such that $vc \in E$ if and only if $v \in \text{var}(F)$, $c \in F$, and $v \in \text{var}(c)$; that is, the blocks are the variables and clauses of F , and a variable is adjacent to a clause if and only if the variable occurs in the clause.

A bipartite graph $G = (V, W, E)$ is left convex if there exists an ordering σ of V such that the following holds: if wv and wv' are edges of G and $v < v'' < v'$ (with respect to the ordering σ) then wv'' is an edge of G . The ordering σ is said to witness left convexity of G . A CNF F is variable convex if $\text{inc}(F) = (\text{var}(F), F, E)$ is left convex.

For an integer d , a CNF F has *degree d* if $\text{inc}(F)$ has degree at most d . A class \mathcal{F} of CNFs has *bounded degree* if there exists an integer d such that every CNF in \mathcal{F} has degree d .

3 Polynomial Time Compilability

In this section, we introduce the *few subterms* property, a sufficient condition for a class of CNFs to admit polynomial size compilation into OBDDs (Section 3.1). We prove that the classes of variable convex CNFs and bounded treewidth CNFs have the few subterms property (Sections 3.2 and 3.3). Finally, we establish fixed-parameter tractable size and time OBDD compilation results for CNFs, where the parameter is the deletion distance to a few subterms CNF class (Section 3.4).

3.1 The Few Subterms Property

Definition 1 (Subterm width) Let F be a CNF formula and let $V \subseteq \text{var}(F)$. The set of V -subterms of F is defined

$$\text{st}(F, V) = \{F[f] \mid f: V \rightarrow \{0, 1\}\}.$$

Given an ordering σ of $\text{var}(F)$, the *subterm width* of F with respect to σ is

$$\text{stw}(F, \sigma) = \max\{|\text{st}(F, \text{var}(\pi))| \mid \pi \text{ is a prefix of } \sigma\}.$$

The *subterm width* of F is the minimum subterm width of F with respect to an ordering of $\text{var}(F)$.

We now state a few simple properties of subterms that will be used throughout the paper. Let V be a set of variables and let F be a CNF.

1. Trivially, $|\text{st}(F, V)| \leq 2^{|V|}$.
2. Each V -subterm of F is the restriction of a subset of F , so $|\text{st}(F, V)| \leq 2^{|F|}$.
3. If $F = F' \cup F''$ then $F[f] = F'[f] \cup F''[f]$ for every variable assignment $f: V \rightarrow \{0, 1\}$, so $|\text{st}(F, V)| \leq |\text{st}(F', V)| + |\text{st}(F'', V)|$.

Definition 2 (Subterm Bound) Let \mathcal{F} be a class of CNF formulas. A function $b: \mathbb{N} \rightarrow \mathbb{N}$ is a *subterm bound* of \mathcal{F} if, for all $F \in \mathcal{F}$, the subterm width of F is bounded from above by $b(\text{size}(F))$. Let $b: \mathbb{N} \rightarrow \mathbb{N}$ be a subterm bound of \mathcal{F} , let $F \in \mathcal{F}$, and let σ be an ordering of $\text{var}(F)$. We call σ a *witness* of the subterm bound b with respect to F if $\text{stw}(F, \sigma) \leq b(\text{size}(F))$.

Definition 3 (Few Subterms) A class \mathcal{F} of CNF formulas has *few subterms* if it has a polynomial subterm bound $p: \mathbb{N} \rightarrow \mathbb{N}$; if, in addition, for all $F \in \mathcal{F}$, an ordering σ of $\text{var}(F)$ witnessing p with respect to F can be computed in polynomial time, \mathcal{F} is said to have *constructive few subterms*.

The few subterms property naturally presents itself as a sufficient condition for a polynomial size construction of OBDDs from CNFs.

Theorem 4 *There exists an algorithm that, given a CNF F and an ordering σ of $\text{var}(F)$, returns a σ -OBDD for F of size at most $|\text{var}(F)| \cdot \text{stw}(F, \sigma)$ in time polynomial in $|\text{var}(F)|$ and $\text{stw}(F, \sigma)$.*

Proof Let F be a CNF and $\sigma = x_1 < \dots < x_n$ be an ordering of $\text{var}(F)$. The algorithm computes a σ -OBDD D for F as follows.

At step $i = 1$, create the source of D , labelled by F , at level 0 of the diagram; if $\emptyset \in F$ (respectively, $F = \emptyset$), then identify the source with the 0-sink (respectively, 1-sink) of the diagram, otherwise make the source an x_1 -node.

At step $i + 1$ for $i = 1, \dots, n - 1$, let v_1, \dots, v_l be the x_i -nodes at level $i - 1$ of the diagram, respectively labelled F_1, \dots, F_l . For $j = 1, \dots, l$ and $b = 0, 1$, compute $F_j[x_i = b]$, where $x_i = b$ denotes the assignment $f: \{x_i\} \rightarrow \{0, 1\}$ mapping x_i to b . If $F_j[x_i = b]$ is equal to some label of an x_{i+1} -node v already created at level i , then direct the b -edge leaving the x_i -node labelled F_j to v ; otherwise, create a new x_{i+1} -node v at level i , labelled $F_j[x_i = b]$, and direct the b -edge leaving the x_i -node labelled F_j to v . If $\emptyset \in F_j[x_i = b]$, then identify v with the 0-sink of D , and if $\emptyset = F_j[x_i = b]$, then identify v with the 1-sink of D .

At termination, the diagram obtained computes F and respects σ . We analyze the runtime. At step $i + 1$ ($0 \leq i < n$), the nodes created at level i are labelled by CNFs of the form $F[f]$, where f ranges over all assignments of $\{x_1, \dots, x_i\}$ not falsifying F ; that is, these nodes correspond exactly to the $\{x_1, \dots, x_i\}$ -subterms $\text{st}(F, \{x_1, \dots, x_i\})$ of F not containing the empty clause, whose number is bounded above by $\text{stw}(F, \sigma)$. As level i is processed in time bounded above by its size times the size of level $i - 1$, and $|\text{var}(F)|$ levels are processed, the diagram D has size at most $|\text{var}(F)| \cdot \text{stw}(F, \sigma)$ and is constructed in time bounded above by a polynomial in $|\text{var}(F)|$ and $\text{stw}(F, \sigma)$. \square

Corollary 5 *A class of CNFs with the constructive few subterms property admits polynomial time compilation into OBDDs.*

Proof Let \mathcal{F} be a class of CNF formulas with constructive few subterms and let $p: \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial subterm bound of \mathcal{F} . The algorithm, given a CNF F , computes in polynomial time an ordering of $\text{var}(F)$ witnessing p with respect to F , and invokes the algorithm of Theorem 4, which runs in time polynomial in $|\text{var}(F)|$ and $\text{stw}(F, \sigma)$. Since $\text{stw}(F, \sigma) \leq p(\text{size}(F))$ the overall runtime is polynomial in $\text{size}(F)$. \square

3.2 Variable Convex CNF formulas

In this section, we prove that the class of variable convex CNFs has the constructive few subterms property (Lemma 6), and hence admits polynomial time compilation into OBDDs (Theorem 7); as a special case, CNFs whose incidence graphs are cographs admit polynomial time compilation into OBDDs (Example 8).

Lemma 6 *The class \mathcal{F} of variable convex CNFs has the constructive few subterms property.*

Proof Let $F \in \mathcal{F}$, so that $\text{inc}(F)$ is left convex, and let σ be an ordering of $\text{var}(F)$ witnessing the left convexity of $\text{inc}(F)$. Let π be any prefix of σ . Call a clause $c \in F$ π -active in F if $\text{var}(c) \cap \text{var}(\pi) \neq \emptyset$ and $\text{var}(c) \cap (\text{var}(F) \setminus \text{var}(\pi)) \neq \emptyset$. Let A denote the set of π -active clauses of F . For all $c \in A$, let $\text{var}_\pi(c) = \text{var}(c) \cap \text{var}(\pi)$.

Claim Let $c, c' \in A$. Then, $\text{var}_\pi(c) \subseteq \text{var}_\pi(c')$ or $\text{var}_\pi(c') \subseteq \text{var}_\pi(c)$.

Proof of Claim Let $c, c' \in A$. Assume for a contradiction that the statement does not hold, that is, there exist variables $v, v' \in \text{var}(\pi)$, $v \neq v'$, such that $v \in \text{var}_\pi(c) \setminus \text{var}_\pi(c')$ and $v' \in \text{var}_\pi(c') \setminus \text{var}_\pi(c)$. Assume that $\sigma(v) < \sigma(v')$; the other case is symmetric. Since c is π -active, by definition there exists a variable $w \in \text{var}(F) \setminus \text{var}(\pi)$ such that $w \in \text{var}(c)$. It follows that $\sigma(v') < \sigma(w)$. Therefore, we have $\sigma(v) < \sigma(v') < \sigma(w)$, where $v, w \in \text{var}(c)$ and $v' \notin \text{var}(c)$, contradicting the fact that σ witnesses the left convexity of $\text{inc}(F)$. \square

We now argue that there is a function g with domain A such that the image of A under g contains the set $\{A[f] \mid f \text{ does not satisfy } A\}$ of terms induced by assignments not satisfying A . Let $L = \{x, \neg x \mid x \in \text{var}(\pi)\}$ denote the set of literals associated with variables in $\text{var}(\pi)$. The function g is defined as follows. For $c \in A$, we let

$$g(c) = \{c' \setminus L \mid c' \in A, c' \cap L \subseteq c \cap L\}.$$

Let $f : \text{var}(\pi) \rightarrow \{0, 1\}$ be an assignment that does not satisfy A . Let $c \in A$ be a clause not satisfied by f such that $\text{var}_\pi(c)$ is maximal with respect to inclusion. We claim that $g(c) = A[f]$. To see this, let $c' \in A$ be an arbitrary clause. It follows from the claim proved above that either $\text{var}_\pi(c) \subsetneq \text{var}_\pi(c')$ or $\text{var}_\pi(c') \subseteq \text{var}_\pi(c)$. In the first case, c' is satisfied by choice of c . In the second case, c' is not satisfied by f if and only if $c' \cap L \subseteq c \cap L$. The formula $A[f]$ is precisely the set of clauses in A not satisfied by f , restricted to variables not in $\text{var}(\pi)$, so $g(c) = A[f]$ as claimed. Taking into account that an assignment might satisfy A , this implies

$$|\text{st}(A, \text{var}(\pi))| \leq |A| + 1 \leq \text{size}(F) + 1.$$

Let $A' = \{c \in F \mid \text{var}(c) \subseteq \text{var}(\pi)\}$ and $A'' = \{c \in F \mid \text{var}(c) \cap \text{var}(\pi) = \emptyset\}$, so that $F = A \cup A' \cup A''$. For every assignment $f : \text{var}(\pi) \rightarrow \{0, 1\}$ we have $A''[f] = A''$ and either $A'[f] = \emptyset$ or $A'[f] = \{\emptyset\}$, so the number of subterms of F under assignments to $\text{var}(\pi)$ is bounded as

$$|\text{st}(F, \text{var}(\pi))| \leq 2 \cdot (\text{size}(F) + 1).$$

This proves that the class of variable convex CNFs has few subterms. Moreover, an ordering witnessing the left convexity of $\text{inc}(F)$ can be computed in polynomial (even linear) time [4, 19], so the class of variable convex CNFs even has the constructive few subterms property. \square

Theorem 7 *The class of variable convex CNF formulas has polynomial time compilation into OBDDs.*

Proof Immediate from Corollary 5 and Lemma 6. \square

Example 8 (Bipartite Cographs) Let F be a CNF such that $\text{inc}(F)$ is a cograph. Note that $\text{inc}(F)$ is a complete bipartite graph. Indeed, cographs are characterized as graphs of clique-width at most 2 [7], and it is readily verified that if a bipartite graph has clique-width at most 2, then it is a complete bipartite graph. A complete bipartite graph is trivially left convex. Then Theorem 7 implies that CNFs whose incidence graphs are cographs have polynomial time compilation into OBDDs.

3.3 Bounded Treewidth CNF Formulas

In this section, we prove that if a class of CNFs has *bounded treewidth*, then it has the constructive few subterms property (Lemma 10), and hence admits polynomial time compilation into OBDDs (Theorem 11).

Let G be a graph. A *tree decomposition* of G is a triple $\mathcal{T} = (T, \chi, r)$, where $T = (V(T), E(T))$ is a tree rooted at r and $\chi : V(T) \rightarrow 2^{V(G)}$ is a labeling of the vertices of T by subsets of $V(G)$ (called *bags*) such that

1. $\bigcup_{t \in V(T)} \chi(t) = V(G)$,
2. for each edge $uv \in E(G)$, there is a node $t \in V(T)$ with $\{u, v\} \subseteq \chi(t)$, and
3. for each vertex $v \in V(G)$, the set of nodes t with $v \in \chi(t)$ forms a connected subtree of T .

The *width* of a tree decomposition (T, χ, r) is the size of a largest bag $\chi(t)$ minus 1. The *treewidth* of G is the minimum width of a tree decomposition of G . The *pathwidth* of G is the minimum width of a tree decomposition (T, χ, r) such that T is a path.

Let F be a CNF. We say that $\text{inc}(F) = (\text{var}(F), F, E)$ has *treewidth* (respectively, *pathwidth*) k if the graph $(\text{var}(F) \cup F, E)$ has *treewidth* (respectively, *pathwidth*) k . We identify the *pathwidth* (respectively, *treewidth*) of a CNF with the *pathwidth* (respectively, *treewidth*) of its incidence graph.

Let F be a CNF formula. If $\text{inc}(F)$ has *pathwidth* k , then an ordering σ of $\text{var}(F)$ is called a *forget ordering* for F if, with respect to an arbitrary linearization of some path decomposition of width k for $\text{inc}(F)$, if the first bag containing v is less than or equal to the first bag containing v' whenever $\sigma(v) < \sigma(v')$. A proof of the following lemma already appears, in essence, in works by Ferrara et al. [12, Theorem 2.1] and Razgon [23, Lemma 5].

Lemma 9 *Let F be a CNF of pathwidth $k - 1$, and let σ be a forget ordering for F . Then $\text{stw}(F, \sigma) \leq 2^{k+1}$.*

Proof Let F be a CNF such that $\text{inc}(F)$ has *pathwidth* $k - 1$, let σ be a forget ordering for F , and let π be any prefix of σ .

Let v be the last variable in $\text{var}(\pi)$ relative to the ordering σ , and let B be the first bag (in the linearization of P) that contains v . A clause $c \in F$ is called π -*active* in F if $\text{var}(c) \cap \text{var}(\pi) \neq \emptyset$ and $\text{var}(c) \cap (\text{var}(F) \setminus \text{var}(\pi)) \neq \emptyset$.

\emptyset . Let $\text{ac}(F, \text{var}(\pi))$ denote the CNF containing the π -active clauses of F . Let

$$\begin{aligned} C' &= \text{ac}(F, \text{var}(\pi)) \cap B, \\ C'' &= \{c \in \text{ac}(F, \text{var}(\pi)) \mid c \in B' \text{ only if } B' > B \text{ in } P\}; \end{aligned}$$

in words, C' contains π -active clauses in the bag B , and C'' contains π -active clauses occurring only in bags strictly larger than B in the total order of P . Clearly, $C' \cap C'' = \emptyset$.

Claim $\text{ac}(F, \text{var}(\pi)) = C' \cup C''$.

Proof of Claim First observe that a π -active clause c cannot occur only in bags strictly smaller than B in the total order of P . For otherwise, since $\text{var}(c) \cap (\text{var}(F) \setminus \text{var}(\pi)) \neq \emptyset$, let $v' \in \text{var}(c) \cap (\text{var}(F) \setminus \text{var}(\pi))$; if B' is the first bag that contains v' , then $B \leq B'$ (by the choice of v), hence v' is not contained in any bag strictly smaller than B , and the edge cv' is not witnessed in P , a contradiction.

Thus π -active clauses either occur in B (including the case where they occur in B and in bags smaller or larger than B in P), or occur only in bags strictly larger than B in P . Thus, $\text{ac}(F, \text{var}(\pi)) \subseteq C' \cup C''$; the other inclusion holds by definition. \square

From the claim, we get

$$|\text{st}(\text{ac}(F, \text{var}(\pi)), \text{var}(\pi))| \leq |\text{st}(C', \text{var}(\pi))| \cdot |\text{st}(C'', \text{var}(\pi))|;$$

thus, it suffices to bound above the size of the two sets on the right so that the product of the individual bounds is at most 2^k . Let $k' = |C'|$. Obviously, we have $|\text{st}(C', \text{var}(\pi))| \leq 2^{k'}$. Let $V' = \bigcup_{c \in C''} \text{var}(c) \cap \text{var}(\pi)$ and let $k'' = |V'|$.

Claim $V' \subseteq B$.

Proof of Claim Let c be a π -active clause occurring only in bags strictly larger than B in P . Let $v' \in \text{var}(c) \cap \text{var}(\pi)$. By the choice of v and the properties of the forget ordering σ , it holds that the first bag containing v' is less than or equal to B . Since B is the first bag that contains v , it holds that $v' \in B$ by the properties of P (the edge cv' is witnessed in a bag strictly larger than B in P). \square

Claim $|\text{st}(C'', \text{var}(\pi))| \leq 2^{k''}$.

Proof of Claim Define an equivalence relation on $\text{var}(\pi)$ -assignments as follows: For all $f, f': \text{var}(\pi) \rightarrow \{0, 1\}$, $f \equiv f'$ if and only if, for all $v \in V'$, $f(v) = f'(v)$. Since $|V'| = k''$, the equivalence relation has $2^{k''}$ many equivalence classes. Moreover, if $f \equiv f'$, then $C''[f] = C''[f']$, because $\text{var}(C'') \subseteq V'$. The claim follows. \square

Since $C', V' \subseteq B$ and $C' \cap V' = \emptyset$, it holds that $k' + k'' = |C'| + |V'| \leq |B| \leq k$. Hence,

$$|\text{st}(\text{ac}(F, \text{var}(\pi)), \text{var}(\pi))| \leq 2^{k'} \cdot 2^{k''} = 2^{k'+k''} \leq 2^k.$$

Note that F is the disjoint union of $\text{ac}(F, \text{var}(\pi))$, clauses $D' \subseteq F$ whose variables are all in $\text{var}(\pi)$, and clauses $D'' \subseteq F$ whose variables are all outside $\text{var}(\pi)$. Also, $\text{st}(D', \text{var}(\pi)) \subseteq \{\emptyset, \{\emptyset\}\}$, and $\text{st}(D'', \text{var}(\pi)) = \{D''\}$. It follows that

$$\begin{aligned} |\text{st}(F, \text{var}(\pi))| &\leq |\text{st}(\text{ac}(F, \text{var}(\pi)), \text{var}(\pi))| \cdot |\text{st}(D', \text{var}(\pi))| \cdot |\text{st}(D'', \text{var}(\pi))| \\ &\leq 2^{k+1}, \end{aligned}$$

and the statement is proved. \square

Lemma 10 *Let \mathcal{F} be a class of CNFs of bounded treewidth. Then \mathcal{F} has the constructive few subterms property.*

Proof Let $c - 1$ be the treewidth bound of \mathcal{F} and let $F \in \mathcal{F}$, so that the treewidth of $\text{inc}(F)$ is at most $c - 1$. We can compute a width $c - 1$ tree decomposition of $\text{inc}(F)$ in linear time $O(\text{size}(F))$ [3]. From this decomposition, we can compute a path decomposition of $\text{inc}(F)$ of width at most $(c - 1) \cdot \log |\text{var}(F) \cup F| \leq c \cdot \log |\text{var}(F) \cup F| - 1$ [2, Corollary 24] and a corresponding forget ordering of $\text{var}(F)$ in polynomial time. By Lemma 9, the subterm width of F with respect to σ is at most $2^{c \cdot \log |\text{var}(F) \cup F|} = |\text{var}(F) \cup F|^c \leq O(\text{size}(F)^c)$. Thus \mathcal{F} has a polynomial subterm bound, and a witnessing ordering σ can be computed for each $F \in \mathcal{F}$ in polynomial time. We conclude that \mathcal{F} has the constructive few subterms property. \square

Theorem 11 *Let \mathcal{F} be a class of CNFs of bounded treewidth. Then, \mathcal{F} has polynomial time compilation into OBDDs.*

Proof Immediate from Lemma 10 and Corollary 5. \square

3.4 Almost Few Subterms

In this section, we use the language of *parameterized complexity* to formalize the observation that CNF classes “close” to CNF classes with few subterms have “small” OBDD representations [11, 13].

Let F be a CNF and D a set of variables and clauses of F . Let E be the formula obtained by deleting D from F , that is,

$$E = \{c \setminus \{l \in c \mid \text{var}(l) \in D\} \mid c \in F \setminus D\};$$

we call D a *deletion set* of F with respect to E .

The following lemma shows that adding a few variables and clauses does not increase the subterm width of a formula too much.

Lemma 12 *Let F and E be CNFs such that D is a deletion set of F with respect to E . Let π be an ordering of $\text{var}(E)$ and let σ be an ordering of $\text{var}(F) \cap D$. Then $\text{stw}(F, \sigma\pi) \leq 2^k \cdot \text{stw}(E, \pi)$, where $k = |D|$.*

Proof Let $V = D \cap \text{var}(F)$ and $C = D \cap F$, and let $k' = |V|$ and $k'' = |C|$. Let ρ be a prefix of $\sigma\pi$ and $X = \text{var}(\rho)$. From $F = C \cup (F \setminus C)$ we get

$$|\text{st}(F, X)| \leq |\text{st}(C, X)| |\text{st}(F \setminus C, X)|. \quad (1)$$

By definition,

$$\text{st}(F \setminus C, X) = \{(F \setminus C)[f] \mid f \in \{0, 1\}^X\}.$$

Splitting the assignments f into two parts, we can write this as

$$\text{st}(F \setminus C, X) = \{(F \setminus C)[f'][f''] \mid f' \in \{0, 1\}^{V \cap X}, f'' \in \{0, 1\}^{X \setminus V}\}. \quad (2)$$

Let $f' \in \{0, 1\}^{V \cap X}$ be an assignment. The formula E is obtained from $F \setminus C$ by deleting variables in V . It follows that $(F \setminus C)[f'] \subseteq E$ and thus

$$(F \setminus C)[f'][f''] \subseteq E[f'']$$

for any assignment $f'' \in \{0, 1\}^{X \setminus V}$. This yields

$$|\{(F \setminus C)[f'][f''] \mid f'' \in \{0, 1\}^{X \setminus V}\}| \leq |E[f''] \mid f'' \in \{0, 1\}^{X \setminus V}|, \quad (3)$$

and the right hand side of this inequality corresponds to $|\text{st}(E, X \setminus V)|$. Combining this with (2), we obtain

$$\begin{aligned} |\text{st}(F \setminus C, X)| &= |\{(F \setminus C)[f'][f''] \mid f' \in \{0, 1\}^{V \cap X}, f'' \in \{0, 1\}^{X \setminus V}\}| \\ &\leq 2^{k'} |\{E[f''] \mid f'' \in \{0, 1\}^{X \setminus V}\}| = 2^{k'} \cdot |\text{st}(E, X \setminus V)| \\ &\leq 2^{k'} \cdot \text{stw}(E, \pi). \end{aligned}$$

Inserting into (1), we get

$$|\text{st}(F, X)| \leq |\text{st}(C, X)| \cdot |\text{st}(F \setminus C, X)| \leq 2^{k''} \cdot 2^{k'} \cdot \text{stw}(E, \pi) = 2^k \cdot \text{stw}(E, \pi),$$

and the lemma is proved. \square

In this section, the standard of efficiency we appeal to comes from the framework of *parameterized complexity* [11, 13]. The parameter we consider is defined as follows. Let \mathcal{F} be a class of CNF formulas. We say that \mathcal{F} is *closed under variable and clause deletion* if $E \in \mathcal{F}$ whenever E is obtained by deleting variables or clauses from $F \in \mathcal{F}$. Let \mathcal{F} be a CNF class closed under variable and clause deletion. The \mathcal{F} -deletion distance of F is the minimum size of a deletion set of F from any $E \in \mathcal{F}$. An \mathcal{F} -deletion set of F is a deletion set of F with respect to some $E \in \mathcal{F}$.

Let \mathcal{F} be a class of CNF formulas with few subterms closed under variable and clause deletion. We say that CNFs have *fixed-parameter tractable OBDD size*, parameterized by \mathcal{F} -deletion distance, if there is a computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ and a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that a CNF F with \mathcal{F} -deletion distance k has OBDD size at most $f(k) p(\text{size}(F))$.

Theorem 13 *Let \mathcal{F} be a class of CNF formulas with few subterms closed under variable and clause deletion. CNFs have fixed-parameter tractable OBDD size, parameterized by \mathcal{F} -deletion distance.*

Proof Let \mathcal{F} be a class of CNF formulas with few subterms closed under variable and clause deletion. Since \mathcal{F} has few subterms, it has a polynomial subterm bound $p : \mathbb{N} \rightarrow \mathbb{N}$. Let k be the \mathcal{F} -deletion distance of F . Let $E \in \mathcal{F}$ be a formula such that the deletion distance of F from E is k , and let D a deletion set of F with respect

to E . Let π be an ordering of $\text{var}(E)$ witnessing p for E , and let σ be an ordering of $\text{var}(F) \cap D$. By Lemma 12, the subterm width of F with respect to $\rho = \sigma\pi$ is at most $2^k p(\text{size}(E))$, so by Theorem 4 there is a ρ -OBDD for F of size at most $2^k p(\text{size}(E)) |\text{var}(F)|$. \square

The requirement that \mathcal{F} be closed under variable and clause deletion ensures that the deletion distance from \mathcal{F} is defined for every CNF. For our purposes, this can be assumed without loss of generality, as it is readily verified that if \mathcal{F} has few subterms with polynomial subterm bound $p : \mathbb{N} \rightarrow \mathbb{N}$, then the closure of \mathcal{F} under variable and clause deletion has few subterms with the same polynomial subterm bound.

Analogously, we say that CNFs have *fixed-parameter tractable time computable OBDDs* (respectively, \mathcal{F} -deletion sets), parameterized by \mathcal{F} -deletion distance, if an OBDD (respectively, a \mathcal{F} -deletion set) for a given CNF F of \mathcal{F} -deletion distance k is computable in time $f(k) p(\text{size}(F))$.

Theorem 14 *Let \mathcal{F} be a class of CNFs closed under variable and clause deletion satisfying the following:*

- *\mathcal{F} has the constructive few subterms property.*
- *CNFs have fixed-parameter tractable time computable \mathcal{F} -deletion sets, parameterized by \mathcal{F} -deletion distance.*

Then CNFs have fixed-parameter tractable time computable OBDDs, parameterized by \mathcal{F} -deletion distance.

Proof Given an input formula F , the algorithm first computes a smallest \mathcal{F} -deletion set D of F . Let E be the formula obtained from F by deleting D . The algorithm computes a variable ordering π of E witnessing a polynomial subterm bound $p : \mathbb{N} \rightarrow \mathbb{N}$ of \mathcal{F} . Since \mathcal{F} has the constructive few subterms property, this can be done in polynomial time. Next, the algorithm chooses an arbitrary ordering σ of $\text{var}(F) \cap D$. By Lemma 12 we have $\text{stw}(F, \sigma\pi) \leq 2^{|D|} \text{stw}(E, \pi) \leq 2^k p(\text{size}(E))$, where k is the \mathcal{F} -deletion distance of F . Invoking the algorithm of Theorem 4, our algorithm computes and returns an OBDD for F in time polynomial in $2^k p(\text{size}(E)) |\text{var}(F)|$. Since $\text{size}(E) \leq \text{size}(F)$ there is a polynomial $q : \mathbb{N} \rightarrow \mathbb{N}$ (independent of F) such that the last expression is bounded by $2^k q(\text{size}(F))$. \square

Corollary 15 (Feedback Vertex Set) *Let \mathcal{F} be the class of formulas whose incidence graphs are forests. CNFs have fixed-parameter tractable time computable OBDDs parameterized by \mathcal{F} -deletion distance.*

Proof Given a graph $G = (V, E)$, a set $D \subseteq V$ is called a *feedback vertex set* of G if the graph $G \setminus D$ is a forest; here, $G \setminus D$ is the graph $(V \setminus D, E')$ such that $vw \in E'$ if and only if $vw \in E$ and $v, w \in V \setminus D$. For any CNF F , a subset D of its variables and clauses is a feedback vertex set of the incidence graph $\text{inc}(F)$ if and only if it is a \mathcal{F} -deletion set, so a smallest feedback vertex set of $\text{inc}(F)$ is a smallest \mathcal{F} -deletion set. There is a fixed-parameter tractable algorithm that, given a graph G and a parameter k , computes a feedback vertex set D of G such that $|D| \leq k$ or reports that no such set

exists [6]. It follows that there is a fixed-parameter tractable algorithm, parameterized by the \mathcal{F} -deletion distance, for computing a smallest \mathcal{F} -deletion set of an input CNF. Moreover, the incidence graphs of formulas in \mathcal{F} have treewidth 1, so \mathcal{F} has the constructive few subterms property by Lemma 10. Clearly, \mathcal{F} is closed under variable and clause deletion. Hence, applying Theorem 14, we conclude that CNFs have fixed-parameter tractable time computable OBDDs parameterized by \mathcal{F} -deletion distance. \square

4 Polynomial Size Incompleteness

In this section, we introduce the *subfunction width* of a graph CNF, to which the OBDD size of the graph CNF is exponentially related (Section 4.1), and prove that *expander graphs* yield classes of graph CNFs of *bounded degree* with linear subfunction width, thus obtaining an exponential lower bound on the OBDD size for graph CNFs in such classes (Section 4.2).

4.1 Many Subfunctions

In this section, we introduce the *subfunction width* of a graph CNF (Definition 16), and prove that the OBDD size of a graph CNF is bounded below by an exponential function of its subfunction width (Theorem 17).

A *graph CNF* is a CNF F such that $F = \{\{u, v\} \mid uv \in E\}$ for some graph $G = (V, E)$ without isolated vertices.

Definition 16 (Subfunction Width) Let F be a graph CNF. Let σ be an ordering of $\text{var}(F)$ and let π be a prefix of σ . We say that a subset $\{c_1, \dots, c_e\}$ of clauses in F is *subfunction productive* relative to π if there exist $\{a_1, \dots, a_e\} \subseteq \text{var}(\pi)$ and $\{u_1, \dots, u_e\} \subseteq \text{var}(F) \setminus \text{var}(\pi)$ such that for all $i, j \in \{1, \dots, e\}$, $i \neq j$, and all $c \in F$,

- $c_i = \{a_i, u_i\}$;
- $c \neq \{a_i, a_j\}$ and $c \neq \{a_i, u_j\}$.

The *subfunction width* of F , in symbols $\text{sfw}(F)$, is defined by

$$\text{sfw}(F) = \min_{\sigma} \max_{\pi} \{ |M| \mid M \text{ is subfunction productive relative to } \pi \},$$

where σ ranges over all orderings of $\text{var}(F)$ and π ranges over all prefixes of σ .

Intuitively, in the graph G underlying the graph CNF F in Definition 16, there is a matching of the form $a_i u_i$ with $a_i \in \text{var}(\pi)$ and $u_i \in \text{var}(F) \setminus \text{var}(\pi)$, $i \in \{1, \dots, e\}$; such a matching is “almost” induced, in that G can contain edges of the form $u_i u_j$, but no edges of the form $a_i a_j$ or $a_i u_j$, $i, j \in \{1, \dots, e\}$, $i \neq j$.

Theorem 17 Let F be a graph CNF. The OBDD size of F is at least $2^{\text{sfw}(F)}$.

Proof Let F be a graph CNF. Let D be any OBDD computing F , let σ be the ordering of $\text{var}(F)$ respected by D , and let π be a prefix of σ such that $\{c_1, \dots, c_e\} \subseteq F$ is subfunction productive relative to σ and π and $e \geq \text{sfw}(F)$. Let $\{a_1, \dots, a_e\} \subseteq \text{var}(\pi)$ and $\{u_1, \dots, u_e\} \subseteq \text{var}(F) \setminus \text{var}(\pi)$ be as in Definition 16, so that in particular $c_i = \{a_i, u_i\}, i \in \{1, \dots, e\}$. Let

$$L = \{f: \text{var}(\pi) \rightarrow \{0, 1\} \mid f(v) = 1 \text{ for all } v \notin \{a_1, \dots, a_e\}\}; \quad (4)$$

in words, L is the set containing, for each assignment of $\{a_1, \dots, a_e\}$, its extension to $\text{var}(\pi)$ that sends all variables in $\text{var}(\pi) \setminus \{a_1, \dots, a_e\}$ to 1.

Claim Let $f \in L$ and let $c \in F$ be such that $c \subseteq \text{var}(\pi)$. Then, f satisfies c .

Proof of Claim Otherwise, since F is a graph CNF and by (4) the only variables sent to 0 by f are in $\{a_1, \dots, a_e\}$, it is the case that $c = \{a_i, a_j\}$ for some $i, j \in \{1, \dots, e\}$, $i \neq j$, which is impossible by the second item in Definition 16. \square

Claim Let f and g be distinct assignments in L . Then, f and g lead to different nodes in D .

Proof of Claim Let f and g be distinct assignments in L . By the previous claim, f and g satisfy each clause in F whose variables are contained in $\text{var}(\pi)$. Thus, the computation paths activated by f and g in D lead to some nodes in D distinct from the 0-sink of D .

Since f and g are distinct assignments in L , they differ on at least one variable in $\{a_1, \dots, a_e\}$; say without loss of generality that $f(a_1) = 0 \neq 1 = g(a_1)$. Let $h: \text{var}(F) \setminus \text{var}(\pi) \rightarrow \{0, 1\}$ be such that $h(v) = 0$ if and only if $v = u_1$. We show that that $f \cup h$ does not satisfy F , but $g \cup h$ satisfies F ; it follows that f and g lead to different nodes in D .

Clearly, $f \cup h$ does not satisfy F , because by Definition 16 the clause $c_1 = \{a_1, u_1\}$ is in F , and by construction $f(a_1) = h(u_1) = 0$. We show that $g \cup h$ satisfies F .

Let $c \in F$. If $c \subseteq \text{var}(\pi)$, then g satisfies c by the previous claim. If $c \subseteq \text{var}(F) \setminus \text{var}(\pi)$, then h satisfies c , because c contains two distinct variables, hence at least one of its variables differs from u_1 and is assigned to 1 by h . Otherwise, $c \cap \text{var}(\pi) \neq \emptyset$ and $c \cap (\text{var}(F) \setminus \text{var}(\pi)) \neq \emptyset$. If c contains a variable in $\text{var}(F) \setminus \text{var}(\pi)$ distinct from u_1 , then h satisfies c . Otherwise, $c = \{a, u_1\}$ for some $a \in \text{var}(\pi)$. In this case, if $a \in \{a_1, \dots, a_e\}$, then $a = a_1$ by Definition 16, and g satisfies c via $g(a_1) = 1$. Else, $a \in \text{var}(\pi) \setminus \{a_1, \dots, a_e\}$ and by definition of L we have $g(a) = 1$, so that again g satisfies c . \square

It is readily observed that $|L| = 2^e$. Then, by the above claims, the computation paths activated by the assignments in L lead to 2^e different nodes in D . We observed that $e \geq \text{sfw}(F)$. Then D has size at least $2^{\text{sfw}(F)}$. It follows that the OBDD size of F is at least $2^{\text{sfw}(F)}$. \square

4.2 Bounded Degree

In this section, we use the existence of a family of *expander graphs* to obtain a class of graph CNFs with linear subfunction width (Lemma 18), thus obtaining an exponential lower bound on the OBDD size of a class of CNFs of *bounded degree* (Theorem 19).

Let n and d be positive integers, $d \geq 3$, and let $\epsilon < 1$ be a positive real. A graph $G = (V, E)$ is a (n, d, ϵ) -*expander* if G has n vertices, degree at most d , and for all subsets $W \subseteq V$ such that $|W| \leq n/2$, the inequality

$$|\text{neigh}(W)| \geq \epsilon|W| \quad (5)$$

holds. It is known that for all integers $d \geq 3$, there exists a real $0 < \epsilon$, and a sequence

$$\{G_i \mid i \in \mathbb{N}\} \quad (6)$$

such that $G_i = (V_i, E_i)$ is an (n_i, d, ϵ) -expander ($i \in \mathbb{N}$), and n_i tends to infinity as i tends to infinity [1, Section 9.2].

Lemma 18 *Let F be a graph CNF whose underlying graph is an (n, d, ϵ) -expander, where $n \geq 2$, $\epsilon > 0$, and $d \geq 3$. Then*

$$\text{sfw}(F) \geq \frac{\epsilon}{16d} n.$$

Proof Let σ be any ordering of $\text{var}(F)$ and let π be the length $\lfloor n/2 \rfloor$ prefix of σ .

Claim There exists a subset $\{c_1, \dots, c_l\}$ of clauses of F , subfunction productive relative to π , such that $l \geq \frac{\epsilon}{16d} n$.

Proof of Claim We will construct a sequence $(a_1, b_1), \dots, (a_l, b_l)$ of pairs $(a_i, b_i) \in \text{var}(\pi) \times (\text{var}(F) \setminus \text{var}(\pi))$ of vertices such that $a_i \notin \text{neigh}(a_j)$, and such that $\{a_i, b_j\} \in F$ if and only if $i = j$, for $1 \leq i, j \leq l$. Letting $c_i = \{a_i, b_i\}$ for $1 \leq i \leq l$, this yields a set $\{c_1, \dots, c_l\}$ of clauses that are subfunction productive relative to π . Assume we have chosen a (possibly empty) sequence $(a_1, b_1), \dots, (a_j, b_j)$ of such pairs. For a vertex v in the underlying graph of F , let $N[v] = \{v\} \cup \text{neigh}(v)$ denote its solid neighborhood. Let $V = \bigcup_{i=1}^j (N[a_i] \cup N[b_i])$ and $A = \text{var}(\pi) \setminus V$. Then $|A| \leq n/2$ and we can use the expansion property (5) to conclude that $|\text{neigh}(A)| \geq \epsilon|A|$. Let $B = \text{neigh}(A) \setminus V$. If both A and B are nonempty we pick $(a_{j+1}, b_{j+1}) \in A \times B$ so that $a_{j+1}b_{j+1}$ is an edge. We have $A \subseteq \text{var}(\pi)$ as well as $B \subseteq \text{var}(F) \setminus (A \cup V) \subseteq \text{var}(F) \setminus \text{var}(\pi)$, so $(a_{j+1}, b_{j+1}) \in \text{var}(\pi) \times (\text{var}(F) \setminus \text{var}(\pi))$. By construction, $\{a_{j+1}, b_{j+1}\}$ is a clause in F ; moreover, $a_i \notin \text{neigh}(b_{j+1})$ as well as $b_i \notin \text{neigh}(a_{j+1})$, for $1 \leq i \leq j$. We conclude that the sequence $(a_1, b_1), \dots, (a_{j+1}, b_{j+1})$ has the desired properties. Otherwise, if either of the sets A or B is empty, we stop.

We now give a lower bound on the length l of a sequence constructed in this manner. Let $(a_1, b_1), \dots, (a_j, b_j)$ be such that one of the sets A and B as defined in the previous paragraph is empty, so that $j = l$. Since the degree of the underlying

graph is bounded by d , we have $|V| \leq 2dj$ and $|A| \geq \lfloor n/2 \rfloor - 2dj$. If A is empty, we must have $2dj \geq \lfloor n/2 \rfloor$ and thus

$$j \geq \left\lfloor \frac{n}{2} \right\rfloor \frac{1}{2d} \geq \frac{n-1}{4d} \geq \frac{n}{8d}, \quad (7)$$

where the last inequality follows from $n \geq 2$. Now suppose B is empty. We have $|B| \geq \epsilon|A| - |V|$, so

$$0 \geq \epsilon(\lfloor n/2 \rfloor - 2dj) - 2dj = \epsilon(\lfloor n/2 \rfloor) - 2dj(1 + \epsilon).$$

From this, we get

$$j \geq \frac{\epsilon(n-1)}{4d(1+\epsilon)} \geq \frac{\epsilon(n-1)}{8d} \geq \frac{\epsilon n}{16d}. \quad (8)$$

Here, the last inequality again follows from $n \geq 2$. Recalling that $\epsilon < 1$ and taking the minimum of the bounds in (7) and (8), we obtain the lower bound stated in the claim. \square

The lemma is an immediate consequence of the above claim. \square

Theorem 19 *There exist a class \mathcal{F} of CNF formulas and a constant $c > 0$ such that, for every $F \in \mathcal{F}$, the OBDD size of F is at least $2^{c \cdot \text{size}(F)}$. In fact, \mathcal{F} is a class of read 3 times, monotone, 2-CNF formulas.*

Proof Let $\mathcal{G} = \{G_i \mid i \in \mathbb{N}\}$ be a family of graphs as in (6), so that for all $i \in \mathbb{N}$ the graph $G_i = (V_i, E_i)$ is a (n_i, d, ϵ) -expander ($n_i \geq 2$, $d = 3$, $\epsilon > 0$) and $n_i \rightarrow \infty$ as $i \rightarrow \infty$. Using the expansion property, it is readily verified that each graph in \mathcal{G} is connected; in particular, it does not have isolated vertices. Therefore $\mathcal{F} = \{E_i : i \in \mathbb{N}\}$ is a class of graph CNFs; in fact, a class of read 3 times, monotone, 2-CNF formulas. By Lemma 18 each $F \in \mathcal{F}$ satisfies

$$\text{sffw}(F) \geq \frac{\epsilon}{16d} |\text{var}(F)|.$$

Since the underlying graph of F has degree at most d and $|\text{var}(F)|$ vertices, the formula F contains at most $d |\text{var}(F)|$ clauses (each variable occurs in at most d clauses), and each clause contains at most 2 literals. That is, $2d |\text{var}(F)| \geq \text{size} F$, and thus

$$\text{sffw}(F) \geq \frac{\epsilon}{16d} |\text{var}(F)| = \frac{\epsilon}{32d^2} 2d |\text{var}(F)| \geq \frac{\epsilon}{32d^2} \text{size} F.$$

Setting $c = \epsilon / 32d^2$, it follows from Theorem 17 that the OBDD size of F is at least $2^{c \cdot |\text{size} F|}$. \square

5 Conclusion

In closing, we briefly explain why completing the classification task laid out in this paper (and thus closing the gap depicted in Fig. 1) seems to require new ideas.

On the one hand, our upper bound for variable convex CNFs appears to push the few subterms property to its limits – natural variable orderings cannot be used to witness few subterms for (clause) convex CNFs and CNF classes of bounded clique-width. On the other hand, our lower bound technique based on expander graphs essentially requires bounded degree, but the candidate classes for improving lower bounds in our hierarchy, bounded clique-width CNFs and beta acyclic CNFs, have unbounded degree. In both cases, imposing a degree bound leads to classes of bounded treewidth [18] and thus polynomial bounds on the size of OBDD representations.

Acknowledgments Open access funding provided by TU Wien, Vienna, Austria. This research was supported by the European Research Council (Complex Reason, 239962) and the FWF Austrian Science Fund (Parameterized Compilation, P26200).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Alon, N., Spencer, J.H.: The Probabilistic Method. Wiley (2000)
2. Bodlaender, H.L.: A partial k -arboretum of graphs with bounded treewidth. *Theor. Comput. Sci.* **209**(1–2), 1–45 (1998)
3. Bodlaender, H.L., Kloks, T.: Efficient and constructive algorithms for the pathwidth and treewidth of graphs. *J. Algorithms* **21**(2), 358–402 (1996)
4. Booth, K.S., Lueker, G.S.: Testing for the consecutive ones property, interval graphs, and graph planarity using pq-tree algorithms. *J. Comput. Syst. Sci.* **13**(3), 335–379 (1976)
5. Brault-Baron, J., Capelli, F., Mengel, S.: Understanding model counting for β -acyclic CNF-formulas. In: *Proceedings of STACS* (2015)
6. Chen, J., Fomin, F.V., Liu, Y., Lu, S., Villanger, Y.: Improved algorithms for feedback vertex set problems. *J. Comput. Syst. Sci.* **74**(7), 1188–1198 (2008)
7. Courcelle, B., Olariu, S.: Upper bounds to the clique-width of graphs. *Discr. Appl. Math.* **101**(1–3), 77–114 (2000)
8. Darwiche, A., Marquis, P.: A Knowledge Compilation Map. *J. Artif. Intell. Res.* **17**, 229–264 (2002)
9. Devadas, S.: Comparing two-level and ordered binary decision diagram representations of logic functions. *IEEE Trans. Comput. Aided Design* **12**(5), 722–723 (1993)
10. Diestel, R.: *Graph Theory*. Springer (2000)
11. Downey, R.G., Fellows, M.R.: *Fundamentals of Parameterized Complexity*. Springer (2013)
12. Ferrara, A., Pan, G., Vardi, M.Y.: Treewidth in verification: Local vs. global. In: *Proceedings of LPAR* (2005)
13. Flum, J., Grohe, M.: *Parameterized Complexity Theory*. Springer (2006)
14. Gallo, G.: An $O(n \log n)$ algorithm for the convex bipartite matching problem. *Oper. Res. Lett.* **3**(1), 31–34 (1984)
15. Glover, F.: Maximum matching in a convex bipartite graph. *Naval Res. Logis. Q.* **14**(3), 313–316 (1967)
16. Hoory, S., Linial, N., Wigderson, A.: Expander graphs and their applications. *Bull. Am. Math. Soc.* **43**(4), 439–561 (2006)
17. Jukna, S.: *Boolean Function Complexity - Advances and Frontiers*. Springer (2012)
18. Kaminski, M., Lozin, V.V., Milanić, M.: Recent developments on graphs of bounded clique-width. *Discr. Appl. Math.* **157**(12), 2747–2761 (2009)

19. Köbler, J., Kuhnert, S., Laubner, B., Verbitsky, O.: Interval graphs: Canonical representations in logspace. *SIAM J. Comput.* **40**(5), 1292–1315 (2011)
20. Lozin, V., Rautenbach, D.: Chordal bipartite graphs of bounded tree- and clique-width. *Discrete Math.* **283**(1–3), 151–158 (2004)
21. Lubotzky, A.: Expander graphs in pure and applied mathematics. *Bull. Am. Math. Soc.* **49**(1), 113–162 (2012)
22. Oztok, U., Darwiche, A.: CV-width: A new complexity parameter for CNFs. In: *Proceedings of ECAI* (2014)
23. Razgon, I.: On OBDDs for CNFs of bounded treewidth. In: *Proceedings of KR* (2014)
24. Steiner, G., Yeomans, S.: Level schedules for mixed-model, just-in-time processes. *Manag. Sci.* **39**(6), 728–735 (1993)
25. Wegener, I.: *Branching Programs and Binary Decision Diagrams*. SIAM (2000)