

Connecting Width and Structure in Knowledge Compilation

Antoine Amarilli¹, Mikaël Monet^{1,3}, and Pierre Senellart^{2,3}

- 1 LTCI, Télécom ParisTech, Université Paris-Saclay
 2 DI ENS, ENS, CNRS, PSL Research University; Paris, France
 3 Inria Paris; Paris, France

Abstract

Several query evaluation tasks can be done via *knowledge compilation*: the query result is compiled as a *lineage circuit* from which the answer can be determined. For such tasks, it is important to leverage some width parameters of the circuit, such as bounded treewidth or pathwidth, to convert the circuit to structured classes, e.g., deterministic structured NNFs (d-SDNNFs) or OBDDs. In this work, we show how to connect the width of circuits to the size of their structured representation, through upper and lower bounds. For the upper bound, we show how bounded-treewidth circuits can be converted to a d-SDNNF, in time linear in the circuit size. Our bound, unlike existing results, is constructive and only singly exponential in the treewidth. We show a related lower bound on monotone CNF or DNF formulas, assuming a constant bound on the arity (size of clauses) and degree (number of occurrences of each variable). Specifically, any d-SDNNF for such a formula must be of exponential size in its treewidth; and the same holds for pathwidth when compiling to OBDDs. Our lower bounds, in contrast with most previous work, apply to *any* formula of this class, not just a well-chosen family. Hence, for our language of CNF and DNF, pathwidth and treewidth respectively characterize the efficiency of compiling to OBDDs and d-SDNNFs, that is, compilation is singly exponential in the corresponding width parameter. We conclude by applying our lower bound results to the task of query evaluation.

1 Introduction

Uncertainty and errors in data can be modeled using *probabilistic databases* [38], annotating every tuple with a probability of existence. Query evaluation on probabilistic databases must then handle the uncertainty by computing the probability that each query result holds. A common technique to evaluate queries on probabilistic databases is the *intensional approach*: first compute a representation of the *lineage* of the query on the database, which intuitively describes how the query depends on the possible database facts; then use this lineage to compute probabilities efficiently. Specifically, the lineage can be computed as a *circuit* [32], and efficient probability computation can be achieved by restricting to tractable circuit classes via *knowledge compilation*. Thus, to evaluate queries on probabilistic databases, we can use knowledge compilation algorithms to translate circuits to tractable classes; conversely, lower bounds in knowledge compilation can identify the limits of the intensional approach.

In this paper, we study the relationship between two kinds of tractable circuit classes in knowledge compilation: *width-based* classes, specifically, bounded-treewidth and bounded-pathwidth circuits; and *structure-based* classes, specifically, OBDDs (ordered binary decision diagrams [16], following a *variable order*) and d-SDNNFs (structured deterministic decomposable negation normal forms [35], following a *v-tree*). Circuits of bounded treewidth can be obtained as a result of practical query evaluation [3, 6, 30], whereas OBDDs and d-DNNFs have been studied to show theoretical characterizations of the query lineages they can represent [31]. Both classes enjoy tractable probabilistic computation: for width-based classes, using *message passing* [34], in time linear in the circuit and exponential in the treewidth; for



licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

OBDDs and d-SDNNFs, in linear time by definition of the class [21]. Hence the question that we study: can we compile width-based classes efficiently into structure-based classes?

We first study how to perform this transformation, and show corresponding *upper bounds*. Existing work has already studied the compilation of bounded-pathwidth circuits to OBDDs [32], which can be made constructive [7, Lemma 6.9]. Accordingly, we focus on compiling *bounded-treewidth circuits* to *d-SDNNF circuits*. Our first contribution, stated in Section 3 and proved in Section 4, is to show the following:

► **Result 1 (Theorem 3.3 and subsequent remark).** *Given as input a Boolean circuit C of treewidth k of C , we can compute a d-SDNNF equivalent to C in time $O(|C| \times f(k))$ where f is singly exponential.*

The algorithm transforms the input circuit bottom-up, considering all possible valuations of the gates in each bag of the tree decomposition, and keeping track of additional information to remember which guessed values have been substantiated by a corresponding input. Our result recaptures a recent theorem of Bova and Szeider in [15], except that our bound depends on $|C|$ (the circuit size) whereas their bound depends on the number of variables of C . In exchange for this, we improve on their result in two ways. First, our result is constructive, whereas [15] only shows a bound on the size of the d-SDNNF, without bounding the complexity of effectively computing it. Second, our bound is singly exponential in k , whereas [15] is doubly exponential; this allows us to be competitive with message passing (also singly exponential in k), and we believe it can be useful for practical applications. Indeed, beyond probabilistic query evaluation, our result implies that all tractable tasks on d-SDNNFs (e.g., enumeration [2] and MAP inference [27]) are also tractable on bounded-treewidth circuits.

Second, we study *lower bounds* on how efficiently we can convert from width-based to structure-based classes. Our bounds already apply to a weaker formalism of width-based circuits, namely monotone CNFs and DNFs of bounded width, so we focus on them. Our second contribution (in Section 5) concerns pathwidth and OBDD representations: we show that, up to factors in the formula arity (maximal size of clauses) and degree (maximal number of variable occurrences), any OBDD for a monotone CNF or DNF must be of width exponential in the pathwidth of the formula. Formally:

► **Result 2 (Theorem 5.3).** *Let φ be a monotone DNF (or monotone CNF), let $a := \text{arity}(\varphi)$ and $d := \text{degree}(\varphi)$. Then any OBDD for φ has width $2^{\Omega\left(\frac{\text{pw}(\varphi)}{a^3 \times d^2}\right)}$.*

This result generalizes several existing lower bounds in knowledge compilation that exponentially separate CNFs from OBDDs, such as [24] and [14, Theorem 19].

Our third contribution (Section 6) is to show an analogue for treewidth and d-SDNNFs:

► **Result 3 (Theorem 6.1).** *Let φ be a monotone DNF (or monotone CNF), let $a := \text{arity}(\varphi)$ and $d := \text{degree}(\varphi)$. Then any d-SDNNF for φ has size $2^{\Omega\left(\frac{\text{tw}(\varphi)}{a^3 \times d^2}\right)}$.*

Our two lower bounds contribute to a vast landscape of knowledge compilation results giving lower bounds on compiling specific Boolean functions to restricted circuits classes, e.g., [14, 24, 37] to OBDDs, [17] to *decision* structured DNNF, [8] to *sentential decision diagrams* (SDDs), [12, 36] to d-SDNNF, [12, 18, 19] to d-DNNFs and DNNFs. However, all those lower bounds (with the exception of some results in [18, 19] discussed in Section 6) apply to well-chosen families of Boolean functions (usually CNF), whereas Result 2 and 3 apply to *any* monotone CNF and DNF. Together with Result 1, these generic lower bounds point to a strong relationship between width parameters and structure representations, on monotone CNFs and DNFs of constant arity and degree. Specifically, the smallest width of

OBDD representations of any such formula φ is in $2^{\Theta(\text{pw}(\varphi))}$, i.e., precisely singly exponential in the pathwidth; and an analogous bound applies to d-SDNNF size and treewidth.

To prove our lower bounds, we rephrase pathwidth and treewidth to new notions of *pathsplitwidth* and *treewidth*, which intuitively measure the performance of a variable ordering or v-tree. We also use the *disjoint non-covering prime implicant sets* (dnpci-sets), a tool introduced in [1, 7] by some of the present authors, and generalizing *subfunction width* [14]. These dnpci-sets allow us to derive lower bounds on OBDD width directly using [1]. We show how they can also imply lower bounds on d-SDNNF size, using the recent communication complexity approach of Bova, Capelli, Mengel and Slivovsky [12].

Our fourth contribution (Section 7) applies our lower bounds to intensional query evaluation on relational databases. We reuse the notion of *intricate* queries of [7], and show that d-SDNNF representations of the lineage of these queries have size exponential in the treewidth of *any* input instance. This extends the result of [7] from OBDDs to d-SDNNFs:

► **Result 4 (Theorem 7.1).** *There is a constant $d \in \mathbb{N}$ such that the following is true. Let σ be an arity-2 signature, and Q be a connected UCQ $^\neq$ which is intricate on σ . For any instance I on σ , any d-SDNNF representing the lineage of Q on I has size $\geq 2^{\Omega(\text{tw}(I)^{1/d})}$.*

As in [7], this result shows that, on arity-2 signatures and under constructibility assumptions, treewidth is the right parameter on instance families to ensure that all queries (in monadic second-order) have tractable d-SDNNF lineage representations.

We start in Section 2 with preliminaries. Full proofs of all results are in the appendix.

2 Preliminaries

Hypergraphs, treewidth, pathwidth. A *hypergraph* $H = (V, E)$ consists of a finite set of *nodes* (or *vertices*) V and of a set E of *hyperedges* (or simply *edges*) which are non-empty subsets of V . We always assume that hypergraphs have at least one edge. For a node v of H , we write $E(v)$ for the set of edges of H that contain v . The *arity* of H , written $\text{arity}(H)$, is the maximal size of an edge of H . The *degree* of H , written $\text{degree}(H)$, is the maximal number of edges to which a vertex belongs, i.e., $\max_{v \in V} |E(v)|$.

A *tree decomposition* of a hypergraph $H = (V, E)$ is a finite, rooted tree T , whose nodes b (called *bags*) are labeled by a subset $\lambda(b)$ of V , and which satisfies:

1. for every fact $e \in E$, there is a bag $b \in T$ with $e \subseteq \lambda(b)$;
2. for all $v \in V$, the set of bags $\{b \in T \mid v \in \lambda(b)\}$ is a connected subtree of T .

For brevity, we identify a bag b with its domain $\lambda(b)$. The *width* of T is $\max_{b \in T} |\lambda(b)| - 1$. The *treewidth* of H is the minimal width of a tree decomposition of H . Pathwidth is defined similarly but with *path decompositions*, where T is a path rather than a tree.

It is NP-hard to determine the treewidth of a hypergraph, but we can compute a tree decomposition in linear time when parametrizing by the treewidth:

► **Theorem 2.1 [9].** *Given a hypergraph H and an integer $k \in \mathbb{N}$ we can check in time $O(|C| \times g(k))$ whether H has treewidth $\leq k$, and if yes output a tree decomposition of H of width $\leq k$, where g is a fixed function in $O(2^{(32+\varepsilon)k^3})$ for any $\varepsilon > 0$.*

For simplicity, we will often assume that a tree decomposition is *nice*, meaning that:

1. it is a full binary tree, i.e., each node has exactly zero or two children;
2. for every internal bag b with children b_l, b_r we have $b \subseteq b_l \cup b_r$;
3. for every leaf bag b we have $|b| \leq 1$;
4. for every non-root bag b with parent b' , we have $|b \setminus b'| \leq 1$;
5. for the root bag b we have $|b| \leq 1$.

► **Lemma 2.2.** *Given a tree decomposition T of width k having n nodes, we can compute in time $O(k \times n)$ a nice tree decomposition T' of width k and $O(k \times n)$ nodes.*

Boolean circuits and functions. A (Boolean) *valuation* of a set V is a function $\nu : V \rightarrow \{0, 1\}$. A *Boolean function* φ on variables V is a mapping that associates to each valuation ν of V a Boolean value in $\{0, 1\}$ called the *evaluation* of φ according to ν .

A (Boolean) *circuit* $C = (G, W, g_{\text{output}}, \mu)$ is a directed acyclic graph (G, W) whose vertices G are called *gates*, whose edges W are called *wires*, where $g_{\text{output}} \in G$ is the *output gate*, and where each gate $g \in G$ has a *type* $\mu(g)$ among **var** (a *variable gate*), **not**, **or**, and. The *inputs* of a gate $g \in G$ are the gates $g' \in G$ such that $(g', g) \in W$; the *fan-in* of g is its number of inputs. We require **not**-gates to have fan-in 1 and **var**-gates to have fan-in 0. The *treewidth* of C , and its *size*, are those of the graph (G, W) . The set C_{var} of *variable gates* of C are those of type **var**. Given a valuation ν of C_{var} , we extend it to an *evaluation* of C by mapping each variable $g \in C_{\text{var}}$ to $\nu(g)$, and evaluating the other gates according to their type. The Boolean function on C_{var} *captured* by the circuit is the one that maps ν to the evaluation of g_{output} under ν . Two circuits are *equivalent* if they capture the same function.

We recall restricted circuit classes from knowledge compilation. We say that C is in *negation normal form* (NNF) if the inputs of **not**-gates are always variable gates. For a gate g in a Boolean circuit C , we write $\text{Vars}(g)$ for the set of variable gates of C_{var} that have a directed path to g in C . An **and**-gate g of C is *decomposable* if for every two input gates $g_1 \neq g_2$ of g we have $\text{Vars}(g_1) \cap \text{Vars}(g_2) = \emptyset$. We call C *decomposable* if each **and**-gate is.

A stronger requirement than decomposability is *structuredness*. A *v-tree* [35] over a set V is a rooted ordered binary tree T whose leaves are in bijection with V ; we identify each leaf with the associated element of V . For $n \in T$, we denote by T_n the subtree of T rooted at n , and for a subset $U \subseteq T$ of nodes of T , we denote by $\text{Leaves}(U)$ the leaves that are in U , i.e., $U \cap V$. We say that T *structures* a Boolean circuit C (and call it a *v-tree for C*) if T is over the set C_{var} and if, for every **and**-gate g of C with inputs g_1, \dots, g_m and $m > 0$, there is a node $n \in T$ that *structures* g , i.e., n has m children n_1, \dots, n_m and we have $\text{Vars}(g_i) \subseteq \text{Leaves}(T_{n_i})$ for all $1 \leq i \leq m$. We call C *structured* if some v-tree structures it. Note that structured Boolean circuits are always decomposable, and their **and**-gates have at most two inputs.

A last requirement on circuits is *determinism*. An **or**-gate g of C is *deterministic* if there is no pair $g_1 \neq g_2$ of input gates of g and valuation ν of C_{var} such that g_1 and g_2 both evaluate to 1 under ν . A Boolean circuit is *deterministic* if each **or**-gate is.

The main structured class of circuits that we study in this work are *deterministic structured decomposable NNFs*, which we denote d-SDNNF for brevity as in [35].

DNFs and CNFs. We also study other representations of Boolean functions, namely, Boolean formulas in *conjunctive normal form* (CNFs) and in *disjunctive normal form* (DNFs). A DNF (resp., CNF) φ on a set of variables V is a disjunction of *clauses*, each of which is a conjunction (resp., disjunction) of *literals* on V , i.e., variables of V (a *positive literal*) or their negation (a *negative literal*). A *monotone DNF* (resp., monotone CNF) is one where all literals are positive, in which case we often identify a clause to the set of variables that it contains. We always assume that monotone DNFs and monotone CNFs are *minimized*, i.e., no clause is a subset of another. This ensures that every monotone Boolean function has a unique representation as a monotone DNF (the disjunction of its prime implicants), and likewise for CNF. We assume that CNFs and DNFs always contain at least one non-empty clause (in particular, they cannot represent constant functions). Monotone DNFs and CNFs φ are isomorphic to *hypergraphs*: the vertices are the variables of φ , and the hyperedges are

the clauses of φ . We often identify φ to its hypergraph. In particular, the *pathwidth* and *treewidth* of φ , and its *arity* and *degree*, are defined as that of its hypergraph.

3 Upper Bounds

Our upper bound result studies how to compile a Boolean circuit to a d-SDNNF, parametrized by the treewidth of the input circuit. To present it, we first review the independent result that was recently shown by Bova and Szeider [15] about these circuit classes:

► **Theorem 3.1 [15, Theorem 3 and Equation (22)].** *Given a Boolean circuit C with n variables and of treewidth $\leq k$, there exists an equivalent d-SDNNF of size $O(f(k) \times n)$, where f is doubly exponential.*

An advantage of their result is that it depends only on the *number of variables* of the circuit (and on the width parameter), not on the *size* of the circuit. None of our results will have this advantage, and we will always measure complexity as a function of the size of the input circuit. In exchange for this advantage, their result has two drawbacks: (i) the doubly exponential dependency on the width; and (ii) its nonconstructive aspect, because [15] gives no time bound on the computation, leaving open the question of effectively compiling bounded-treewidth circuits to d-SDNNFs.

Naive constructive bound. We first address the second drawback by showing an easy constructive result. The argument is very simple and appeals to techniques from our earlier works on provenance circuits [6, 7]; it is independent from the techniques of [15].

► **Theorem 3.2.** *Given any circuit C of treewidth k , we can compute an equivalent d-SDNNF in linear time parametrized by k , i.e., in time $O(|C| \times f(k))$ for some computable function f .*

Proof sketch. We encode in linear time the input circuit C to a relational instance I with same treewidth. We use [7, Theorem 6.11] to construct in linear time a provenance representation C' on I of a fixed MSO formula that describes Boolean circuit evaluation. This allows us to obtain in linear time from C' the desired equivalent d-SDNNF representation. ◀

This result shows that we can effectively compile in linear time parametrized by the treewidth k , but does not address the first drawback, namely, the dependency in k .

Improved bound. Our main upper bound result subsumes the naive bound above, with a more elaborate proof, again independent of the techniques of [15]. It addresses both drawbacks and shows that we can effectively compile in time singly exponential in k ; formally:

► **Theorem 3.3.** *Given as input a Boolean circuit C and tree decomposition T of width k , we can compute a d-SDNNF equivalent to C with its v -tree, in $O(|T| \times 2^{(4+\varepsilon)k})$ for any $\varepsilon > 0$.*

We prove Theorem 3.3 in the next section. Observe how we assume the tree decomposition to be given as part of the input. If it is not, we can compute one with Theorem 2.1, but this becomes the bottleneck: the complexity becomes $O(|C| \times 2^{(32+\varepsilon)k^3})$ for any $\varepsilon > 0$.

Applications. Theorem 3.3 implies several consequences for bounded-treewidth circuits. The first one deals with *probability computation*: we are given a *probability valuation* π mapping each variable $g \in C_{\text{var}}$ to a probability that g is true (independently from other variables), and we wish to compute the probability $\pi(C)$ that C evaluates to true under π , assuming that arithmetic operations (sum and product) take unit time. This problem is #P-hard for arbitrary circuits, but it is tractable for d-SDNNF [21]. Hence, our result implies the following, where $|\pi|$ denotes the size of writing the probability valuation π :

► **Corollary 3.4.** *Given a Boolean circuit C , a tree decomposition T of width k of C , and a probability valuation π of C , we can compute $\pi(C)$ in $O(|\pi| + |T| \times 2^{(4+\varepsilon)k})$ for any $\varepsilon > 0$.*

This improves on the bound obtained when applying message passing techniques [34] directly on the bounded-treewidth input circuit (as presented, e.g., in [5, Theorem D.2]). Indeed, message passing applies to *moralized* representations of the input: for each gate, the tree decomposition must contain a bag containing all inputs of this gate *simultaneously*, which is problematic for circuits of large fan-in. Indeed, if the original circuit has a tree decomposition of width k , rewriting it to make it moralized results in a tree decomposition of width $3k^2$ (see [4, Lemmas 53 and 55]), and the bound of [5, Theorem D.2] then yields an overall complexity of $O(|\pi| + |T| \times 2^{3k^2})$ for message passing. Our Corollary 3.4 achieves a more favorable bound because Theorem 3.3 uses directly the associativity of **and** and **or**. We note that the connection between message-passing techniques and structured circuits has also been investigated by Darwiche, but his result [22, Theorem 6] produces arithmetic circuits rather than d-DNNFs, and it also needs the input to be moralized.

A second consequence concerns the task of *enumerating* the accepting valuations of circuits, i.e., producing them one after the other, with small *delay* between each accepting valuation. The valuations are concisely represented as *assignments*, i.e., as a set of variables that are set to true, omitting those that are set to false. This task is of course NP-hard on arbitrary circuits (as it implies that we can check whether an accepting valuation exists), but was recently shown in [2] to be feasible on d-SDNNFs with linear-time preprocessing and delay linear in the Hamming weight of each produced assignment. Hence, we have:

► **Corollary 3.5.** *Given a Boolean circuit C and a tree decomposition T of width k of C , we can enumerate the accepting assignments of C with preprocessing in $O(|T| \times 2^{(4+\varepsilon)k})$ and delay linear in the size of each produced assignment.*

Other applications of Theorem 3.3 include counting the number of satisfying valuations of the circuit (a special case of probability computation), MAP inference [27] or random sampling of possible worlds (which can be done on the d-DNNF in a straightforward manner).

4

Proof of the Main Upper Bound Result

In this section, we present the construction used to prove Theorem 3.3. We start with prerequisites, and then describe how to build the d-SDNNF equivalent to the input bounded-treewidth circuit. Last, we sketch the correctness proof.

Prerequisites. Let C be the input circuit, and T the input tree decomposition. By Lemma 2.2, we assume that T is nice. Further, up to adding a constant number of bags and re-rooting T , we can assume that the root bag of T contains only the output gate g_{output} . For any bag b of T , we define $\text{VarT}(b)$ to be the set of variable gates such that b is the topmost bag in which they appear; as T is nice, $\text{VarT}(b)$ is either empty or is a singleton $\{g\}$, in which case we call b *responsible for the variable gate g* . We can explicitly compute the function VarT in $O(|T|)$, i.e., compute $\text{VarT}(b)$ for each $b \in T$; see for instance [28, Lemma 3.1].

To abstract away the type of gates and their values in the construction, we will talk of *strong* and *weak* values. Intuitively, a value is *strong* for a gate g if any input g' of g which carries this value determines the value of g ; and *weak* otherwise. Formally:

- **Definition 4.1.** Let g be a gate and $c \in \{0, 1\}$:
- If g is an **and**-gate, we say that $c = 0$ is *strong* for g and $c = 1$ is *weak* for g ;
 - If g is an **or**-gate, we say that $c = 1$ is *strong* for g and $c = 0$ is *weak* for g ;

- If g is a not-gate, $c = 0$ and $c = 1$ are both *strong* for g ;
- If g is a var-gate, $c = 0$ and $c = 1$ are both *weak* for g .

If we take any valuation $\nu : C_{\text{var}} \rightarrow \{0, 1\}$ of the circuit $C = (G, W, g_{\text{output}}, \mu)$, and extend it to an evaluation $\nu : G \rightarrow \{0, 1\}$, then ν will respect the semantics of gates. In particular, it will *respect strong values*: for any gate g of C , if g has an input g' for which $\nu(g')$ is a strong value, then $\nu(g)$ is determined by $\nu(g')$, specifically, it is $\nu(g)$ if g is an **or**- or an **and**-gate, and $1 - \nu(g')$ if g is a **not**-gate. In our construction, we will need to guess how gates of the circuit are evaluated, focusing on a subset of the gates (as given by a bag of T); we will then call *almost-evaluation* an assignment of these gates that respects strong values. Formally:

► **Definition 4.2.** Let U be a set of gates of C . We call $\nu : U \rightarrow \{0, 1\}$ a (C, U) -*almost-evaluation* if it *respects strong values*, i.e., for every gate $g \in U$, if there is an input g' of g in U and $\nu(g')$ is a strong value for g , then $\nu(g)$ is determined from $\nu(g')$ in the sense above.

Respecting strong values is a necessary condition for such an assignment to be extensible to a valuation of the entire circuit. However, it is not sufficient: an almost-evaluation ν may map a gate g to a strong value even though g has no input that can justify this value. This is hard to avoid: when we focus on the set U , we do not know about other inputs of g . For now, let us call *unjustified* the gates of ν that carry a strong value that is not justified by ν :

► **Definition 4.3.** Let U be a set of gates of a circuit C and ν a (C, U) -almost-evaluation. We call $g \in U$ *unjustified* if $\nu(g)$ is a strong value for g , but, for every input g' of g in U , the value $\nu(g')$ is weak for g ; otherwise, g is *justified*. The set of unjustified gates is written $\text{Unj}(\nu)$.

Let us start to explain how to construct the d-SDNNF circuit D equivalent to the input circuit C . We do so by traversing T bottom-up, and for each bag b of T we create gates $G_b^{\nu, S}$ in D , where ν is a (C, b) -almost-evaluation and S is a subset of $\text{Unj}(\nu)$ which we call the *suspicious gates* of $G_b^{\nu, S}$. We will connect the gates of D created for each internal bag b with the gates created for its children in T , in a way that we will explain later. Intuitively, for a gate $G_b^{\nu, S}$ of D , the *suspicious gates* g in the set S are gates of b whose strong value is not justified by ν (i.e., $g \in \text{Unj}(\nu)$), and is not justified either by any of the almost-evaluations at descendant bags of b to which $G_b^{\nu, S}$ is connected. We call *innocent* the other gates of b ; they are the gates that are justified in ν (in particular, those who carry weak values), and the gates that are unjustified in ν but have been justified by an almost-evaluation at a descendant bag b' of b . Crucially, in the latter case, the gate g' justifying the strong value in b' may no longer appear in b , making g unjustified for ν ; this is why we remember the set S .

We still have to explain how we connect the gates $G_b^{\nu, S}$ of D to the gates $G_{b_l}^{\nu_l, S_l}$ and $G_{b_r}^{\nu_r, S_r}$ created for the children b_l and b_r of b in T . The first condition is that ν_l and ν_r must *mutually agree*, i.e., $\nu_l(g) = \nu_r(g)$ for all $g \in b_l \cap b_r$, and ν must then be the union of ν_l and ν_r , restricted to b . Remember that T is nice, so b is a subset of $b_l \cup b_r$, and it is easy to verify that ν is then a (C, b) -almost-evaluation. We impose a second condition to prohibit suspicious gates from escaping before they have been justified, which we formalize as *connectibility* of a pair (ν, S) at bag b to the parent bag of b .

► **Definition 4.4.** Let b be a non-root bag, b' its parent bag, and ν a (C, b) -almost-evaluation. For any set $S \subseteq \text{Unj}(\nu)$, we say that (ν, S) is *connectible* to b' if $S \subseteq b'$, i.e., the suspicious gates of ν must still appear in b' .

If a gate $G_b^{\nu, S}$ is such that (ν, S) is not connectible to the parent bag b' , then this gate will not be used as input to any other gate (but we do not try to preemptively remove these useless gates in the construction). We are now ready to give the formal definition that will be used to explain how gates are connected:

► **Definition 4.5.** Let b be an internal bag with children b_l and b_r , let ν_l and ν_r be respectively (C, b_l) and (C, b_r) -almost-evaluations that mutually agree, and $S_l \subseteq \text{Unj}(\nu_l)$ and $S_r \subseteq \text{Unj}(\nu_r)$ be sets of suspicious gates such that both (ν_l, S_l) and (ν_r, S_r) are connectible to b . The *result* of (ν_l, S_l) and (ν_r, S_r) is then defined as the pair (ν, S) where:

- ν is a (C, b) -almost-evaluation defined as the restriction of $\nu_l \cup \nu_r$ to b .
- $S \subseteq \text{Unj}(\nu)$ is the new set of suspicious gates, defined as follows. A gate $g \in b$ is innocent (i.e., $g \in b \setminus S$) if it is justified for ν or if it is innocent for some child. Formally, $b \setminus S := (b \setminus \text{Unj}(\nu)) \cup [b \cap [(b_l \setminus S_l) \cup (b_r \setminus S_r)]]$.

Construction. We now use these definitions to present the construction formally. For every variable gate g of C , we create a corresponding variable gate $G^{g,1}$ of D , and we create $G^{g,0} := \text{not}(G^{g,1})$. For every internal bag b of T , for each (C, b) -almost-evaluation ν and set $S \subseteq \text{Unj}(\nu)$ of suspicious gates of ν , we create one **and**-gate $G_b^{\nu,S}$ and one **or**-gate $G_{b,\text{children}}^{\nu,S}$ which is an input of $G_b^{\nu,S}$. For every leaf bag b of T , we create one gate $G_b^{\nu,S}$ for every (C, b) -almost-evaluation ν , where we set $S := \text{Unj}(\nu)$; intuitively, in a leaf bag, an unjustified gate is always suspicious (it cannot have been justified at a descendant bag).

Now, for each internal bag b of T with children b_l, b_r , for each pair of gates $G_{b_l}^{\nu_l, S_l}$ and $G_{b_r}^{\nu_r, S_r}$ that are both connectible to b and where ν_l and ν_r mutually agree, letting (ν, S) be the result of (ν_l, S_l) and (ν_r, S_r) , we create a gate $G_b^{\nu_l, S_l, \nu_r, S_r} = \text{and}(G_{b_l}^{\nu_l, S_l}, G_{b_r}^{\nu_r, S_r})$ and make it an input of $G_{b,\text{children}}^{\nu,S}$. Last, for each bag b which is responsible for a variable gate g , for each (C, b) -almost-evaluation ν and set of suspicious gates $S \subseteq \text{Unj}(\nu)$, set the gate $G^{g,\nu(g)}$ to be an input of $G_b^{\nu,S}$. The output gate of D is the gate $G_{b_{\text{root}}}^{\nu,\emptyset}$ where b_{root} is the root of T (remember that b_{root} contains only g_{output}).

Correctness. We have formally described the construction of our d-SDNNF D . The construction clearly works in linear time, and we can prove that the dependency on k of the running time is as stated. Further, we easily see that D is structured by a v-tree constructed from the tree decomposition T . To show that D is equivalent to C , one direction is easier: any valuation χ that satisfies C also satisfies D , because we can construct an *accepting trace* in D using the gates $G_b^{\nu,S}$ for ν the restriction of the evaluation χ to b , and for $S := \text{Unj}(\chi|_{T_b})$ where T_b denotes the gates of C occurring in the bags of the subtree of T rooted at b . The converse is trickier: we show that any accepting trace of D describes an evaluation of C that respects strong values by definition of almost-evaluations, and eventually justifies every gate which is given a strong value thanks to our bookkeeping of suspicious gates. Last, we show that D is deterministic: this is unexpected because we freely guess the values of gates of C at leaf bags, but it holds because, when we know the valuation of the variable gates, knowing the valuation of all gates of a bag b uniquely fixes the valuation at the subtree rooted at b . This concludes the proof of Theorem 3.3; see Appendix C for the full correctness proof.

5 Lower Bounds on OBDDs

We now move on to lower bounds on the size of structured representations of Boolean functions, in terms of the width of a circuit for that function. Our end goal is to obtain a lower bound for d-SDNNFs, that will form a counterpart to the upper bound of Theorem 3.3. We will do so in Section 6. For now, in this section, we consider a weaker class of lineage representations than d-SDNNFs, namely, *OBDDs*.

► **Definition 5.1.** An *ordered binary decision diagram* (or *OBDD*) on a set of variables $V = \{v_1, \dots, v_n\}$ is a rooted DAG O whose leaves are labeled by 0 or 1, and whose internal nodes are labeled with a variable of V and have two outgoing edges labeled 0 and 1. We

require that there exists a total order $\mathbf{v} = v_{i_1}, \dots, v_{i_n}$ on the variables such that, for every path from the root to a leaf, the sequence of variables which labels the internal nodes of the path is a subsequence of \mathbf{v} and does not contain duplicate variables. The OBDD O captures a Boolean function on V defined by mapping each valuation ν to the value of the leaf reached from the root by following the path given by ν . The *size* $|O|$ of O is its number of nodes, and the *width* w of O is the maximum number of nodes at every *level*, where a level is defined for a prefix of \mathbf{v} as the set of nodes reached by enumerating all possible valuations of this prefix. Note that we clearly have $|O| \leq |V| \times w$.

Our upper bound in the previous section applied to arbitrary Boolean circuits; however, our lower bounds in this section and the next one will already apply to much weaker formalisms for Boolean functions, namely, monotone DNFs and monotone CNFs (recall their definition from Section 2). Some lower bounds are already known for the compilation of monotone CNFs into OBDDs: Bova and Slivovsky have constructed a family of CNFs of bounded degree whose OBDD width is exponential in their number of variable occurrences [14, Theorem 19], following an earlier result of this type by Razgon [37, Corollary 1]. The result is as follows:

► **Theorem 5.2 [14, Theorem 19].** *There is a class of monotone CNF formulas of bounded degree and arity such that every formula φ in this class has OBDD size at least $2^{\Omega(|\varphi|)}$.*

We adapt some of these techniques to show a more general result: our lower bound applies to *any* monotone DNF or monotone CNF, not to one specific family. Specifically, we show:

► **Theorem 5.3.** *Let φ be a monotone DNF (or monotone CNF), let $a := \text{arity}(\varphi)$ and $d := \text{degree}(\varphi)$. Then any OBDD for φ has width $\geq 2^{\lfloor \frac{\text{pw}(\varphi)}{a^3 \times d^2} \rfloor}$.*

From our Theorem 5.3, we can easily derive Theorem 5.2 using the fact (also used in the proof of [14, Theorem 19]) that there exists a family of monotone CNFs of bounded degree and arity whose treewidth (hence pathwidth) is linear in their size, namely, the CNFs built from *expander graphs* (see [29, Theorem 5 and Proposition 1]). Note that expander graphs can also be used to show lower bounds for (*non-deterministic*) *DNNFs* for a CNF formula [11]; our lower bound on d-SDNNFs of Section 6 does not capture this result.

We observe that, for a family of formulas with bounded arity and degree, the bound of Theorem 5.3 is optimal, up to constant factors in the exponent. Indeed, following earlier work [26, 37], Bova and Slivovsky have shown that any CNF φ can be compiled to OBDDs of width $2^{\text{pw}(\varphi)+2}$ [14, Theorem 4 and Lemma 9]. (Their upper bound result also applies to DNFs, and does not assume monotonicity nor a bound on the arity or degree.) In other words, for any monotone DNF or monotone CNF of bounded arity and degree, pathwidth characterizes the width of an OBDD for the formula, in the following sense:

► **Corollary 5.4.** *For any constant c , for any monotone DNF (or monotone CNF) φ with arity and degree bounded by c , the smallest width of an OBDD for φ is $2^{\Theta(\text{pw}(\varphi))}$.*

This corollary talks about the pathwidth of φ measured as that of its hypergraph, but note that the same result would hold when measuring the pathwidth of the incidence graph or dual hypergraph of φ . Indeed, all these pathwidths are within a constant factor of one another when the degree and arity are bounded by a constant.

We prove Theorem 5.3 in the rest of this section. We present the proof in the case of monotone DNFs to reuse existing lower bound techniques from [1, 7], but explain at the end of this section how the proof adapts to monotone CNFs. We first present *pathsplittedwidth*, a new notion which intuitively measures the performance of a variable ordering for an OBDD on the monotone DNF φ , and connect it to the pathwidth of φ . Second, we recall the definition of *dnepi-sets* introduced in [1, 7] to show lower bounds from the structure of Boolean functions. Last, we conclude the proof by connecting pathsplittedwidth to the size of dnepi-sets.

Pathsplitwidth. The first step of the proof is to rephrase the bound on pathwidth, arity, and degree, in terms of a bound on the performance of variable orderings. Intuitively, a good variable ordering is one which does not *split* too many clauses. Formally:

► **Definition 5.5.** Let $H = (V, E)$ be a hypergraph, and $\mathbf{v} = v_1, \dots, v_{|V|}$ be an ordering on the variables of V . For $1 \leq i \leq |V|$, we define $\text{Split}_i(\mathbf{v}, H)$ as the set of hyperedges e of H that contain both a variable at or before v_i , and a variable strictly after v_i , formally: $\text{Split}_i(\mathbf{v}, H) := \{e \in E \mid \exists l \in \{1, \dots, i\} \text{ and } \exists r \in \{i + 1, \dots, |V|\} \text{ such that } \{v_l, v_r\} \subseteq e\}$. Note that $\text{Split}_{|V|}(\mathbf{v}, H)$ is always empty. The *pathsplitwidth* of \mathbf{v} relative to H is the maximum size of the split, formally, $\text{psw}(\mathbf{v}, H) := \max_{1 \leq i \leq |V|} |\text{Split}_i(\mathbf{v}, H)|$. The *pathsplitwidth* $\text{psw}(H)$ of H is then the minimum of $\text{psw}(\mathbf{v}, H)$ over all variable orderings \mathbf{v} of V .

In other words, $\text{psw}(H)$ is the smallest integer $n \in \mathbb{N}$ such that, for any variable ordering \mathbf{v} of the nodes of H , there is a moment at which n hyperedges of H are *split*, i.e., for n hyperedges e , we have begun enumerating the nodes of e but we have not enumerated all of them yet. For a monotone DNF φ with hypergraph H , the quantity $\text{psw}(H)$ is intuitively connected to the quantity of information that an OBDD will have to remember when evaluating φ following any variable ordering, which we will formalize via dncpi-sets. This being said, the definition of pathsplitwidth is reminiscent of that of pathwidth, and we can indeed connect the two (up to a factor in the arity):

► **Lemma 5.6.** *For any hypergraph $H = (V, E)$, we have $\text{pw}(H) \leq \text{arity}(H) \times \text{psw}(H)$.*

Proof sketch. From a variable ordering \mathbf{v} , we construct a path decomposition of H by creating $|V|$ bags in sequence, each of which containing v_i plus $\bigcup \text{Split}_i(\mathbf{v}, H)$. This is a path decomposition of width $\leq \text{arity}(H) \times \text{psw}(H)$. First, each hyperedge of H is contained in a bag where it is split. Second, each vertex v_i occurs in the corresponding bag b_i and at the union of where the edges containing v are split, which forms a segment of \mathbf{v} : thus, the connectedness condition of tree decompositions is respected. ◀

Thanks to Lemma 5.6, it suffices to show that an OBDD for φ has width $\geq 2^{\lfloor \frac{\text{psw}(\varphi)}{(a \times d)^2} \rfloor}$, which we will do in the rest of this section.

dncpi-sets. To show this lower bound, we use the technical tool of *dncpi-sets* [1, 7]. We recall the definitions here, adapting the notation slightly: in particular, we do not use the term “prime implicant” because our monotone DNFs are assumed to be minimized so prime implicants correspond to clauses. Note that dncpi-sets are reminiscent of *subfunction width* in [14] (see Theorem 17 in this work), but the latter notion is only defined for graph CNFs.

► **Definition 5.7 [1, Definition 6.4.6].** Given a monotone DNF φ on variables V , a *disjoint non-covering prime implicant set* (dncpi-set) of φ is a set S of clauses of φ which:

- are pairwise disjoint: for any $D_1 \neq D_2$ in S , we have $D_1 \cap D_2 = \emptyset$.
 - are *non-covering* in the following sense: for any clause D of φ , if $D \subseteq \bigcup S$, then $D \in S$.
- The *size* of S is the number of clauses that it contains.

Given a variable ordering \mathbf{v} of V , we say that \mathbf{v} *shatters* a dncpi-set S if there exists $1 \leq i \leq |V|$ such that $S \subseteq \text{Split}_i(\mathbf{v}, H)$, where H is the hypergraph of φ .

Observe the analogy between shattering and splitting, which we will substantiate below. We recall the main result on dncpi-sets:

► **Lemma 5.8 [1, Lemma 6.4.7].** *Let φ be a monotone DNF on variables V and $n \in \mathbb{N}$. Assume that, for every variable ordering \mathbf{v} of V , there is some dncpi-set S of φ with $|S| \geq n$, such that \mathbf{v} shatters S . Then any OBDD for φ has width $\geq 2^n$.*

Proof sketch. Considering the point at which the dncpi-set is shattered, the OBDD must remember exactly the status of each clause of the set: any valuation that satisfies a subset of these clauses gives rise to a different continuation function. This is where we use the fact that the DNF is monotone: it ensures that we can freely choose a valuation of the variables that do not occur in the dncpi-set without making the formula true. ◀

Concluding the proof. We conclude the proof of Theorem 5.3 by showing that any variable ordering of the variables of a monotone DNF φ shatters a dncpi-set of the right size. The formal statement is as follows, and it is the last result to prove:

► **Lemma 5.9.** *Let φ be a monotone DNF, H its hypergraph, and \mathbf{v} an enumeration of its variables. Then there is a dncpi-set S of φ shattered by \mathbf{v} such that $|S| \geq \left\lfloor \frac{\text{psw}(H)}{(\text{arity}(H) \times \text{degree}(H))^2} \right\rfloor$.*

We prove this result in the rest of the section. Our goal is to construct a dncpi-set, which intuitively consists of clauses that are disjoint and which do not cover another clause. We can do so by picking clauses sufficiently “far apart”. Let the *exclusion graph* of $H = (V, E)$ be the graph on E where two edges $e \neq e'$ are adjacent if there is an edge e'' of E with which they both share a variable: this is in particular the case when e and e' intersect as we can take $e'' := e$. Formally, the exclusion graph is $G_H = (E, \{\{e, e'\} \in E^2 \mid e \neq e' \wedge \exists e'' \in E, (e \cap e'') \neq \emptyset \wedge (e' \cap e'') \neq \emptyset\})$. In other words, two hyperedges are adjacent in G_H iff they are different and are at distance at most 4 in the incidence graph of H .

Remember that an *independent set* in the graph G_H is a subset S of E such that no two elements of S are adjacent in G_H . The definition of G_H then ensures:

► **Lemma 5.10.** *For any monotone DNF φ , letting H be its hypergraph, any independent set of the exclusion graph G_H is a dncpi-set of φ .*

In other words, our goal is to compute a large independent set of the exclusion graph. To do this, we will use the following straightforward lemma about independent sets:

► **Lemma 5.11.** *Let $G = (V, E)$ be a graph and let $V' \subseteq V$. Then G has an independent set $S \subseteq V'$ of size at least $\left\lfloor \frac{|V'|}{\text{degree}(G)+1} \right\rfloor$.*

Moreover, we can bound the degree of G_H using the degree and arity of H :

► **Lemma 5.12.** *Let H be a hypergraph. Then $\text{degree}(G_H) \leq (\text{arity}(H) \times \text{degree}(H))^2 - 1$.*

Proof sketch. The bound on the arity and degree of H implies a bound on the number of edges that can be at distance ≤ 4 of another edge in the incidence graph of H , hence bounding the degree of the exclusion graph. ◀

We are now ready to conclude the proof of Lemma 5.9:

Proof of Lemma 5.9. Let φ be a monotone DNF, $H = (V, E)$ its hypergraph, and \mathbf{v} an enumeration of its variables. By definition of pathsplitwidth, there is $v_i \in V$ such that, for $E' := \text{Split}_{v_i}(\mathbf{v}, H)$, we have $|E'| \geq \text{psw}(H)$. Now, by Lemma 5.11, G_H has an independent set $S \subseteq E'$ of size at least $\left\lfloor \frac{|E'|}{\text{degree}(G_H)+1} \right\rfloor$ which is $\geq \left\lfloor \frac{\text{psw}(H)}{(\text{arity}(H) \times \text{degree}(H))^2} \right\rfloor$ by Lemma 5.12. Hence, S is a dncpi-set by Lemma 5.10, has the desired size, and is shattered since $S \subseteq E'$. ◀

Combining this result with Lemma 5.6 and Lemma 5.8 concludes the proof of Theorem 5.3.

From DNFs to CNFs. We now argue that Theorem 5.3 also holds for monotone CNFs. Let φ be a monotone CNF, $a := \text{arity}(\varphi)$ and $d := \text{degree}(\varphi)$, and suppose for a contradiction that there is an OBDD O for φ of width $< 2^{\lfloor \frac{\text{pw}(\varphi)}{a^3 \times d^2} \rfloor}$. Consider the DNF φ' build from φ by replacing each \wedge by a \vee and each \vee by a \wedge . Now, let O' be the OBDD built from O by replacing the label $b \in \{0, 1\}$ of each edge by $1 - b$, and replacing the label b of each leaf by $1 - b$. It is clear, by De Morgan's laws, that O' is an OBDD for φ' (which is a monotone DNF) of size $< 2^{\lfloor \frac{\text{pw}(\varphi)}{a^3 \times d^2} \rfloor}$, which contradicts Theorem 5.3 applied to monotone DNFs.

6 Lower Bounds on d-SDNNFs

In the previous section, we have shown that *pathwidth* measures how concisely an OBDD can represent a monotone DNF or CNF formula with bounded degree and arity. In this section, we move from OBDDs to d-SDNNFs, and show that *treewidth* plays a similar role than pathwidth in this setting. Formally, we show the following analogue of Theorem 5.3:

► **Theorem 6.1.** *Let φ be a monotone DNF (or monotone CNF), let $a := \text{arity}(\varphi)$ and $d := \text{degree}(\varphi)$. Then any d-SDNNF for φ has size $\geq 2^{\lfloor \frac{\text{tw}(\varphi)}{3 \times a^3 \times d^2} \rfloor} - 1$.*

Combined with Theorem 3.3 (or with existing results specific to CNF formulas such as [13, Corollary 1]), this yields an analogue of Corollary 5.4. However, its statement is less neat: unlike OBDDs, d-SDNNFs have no obvious notion of width, so the lower bound above refers to size rather than width, and it does not exactly match our upper bound. We obtain:

► **Corollary 6.2.** *For any constant c , for any monotone DNF (or monotone CNF) φ with arity and degree bounded by c , there is a d-SDNNF for φ having size $|\varphi| \times 2^{O(\text{tw}(\varphi))}$, and any d-SDNNF for φ has size $2^{\Omega(\text{tw}(\varphi))}$.*

Our proof of Theorem 6.1 will follow the same overall structure as in the previous section. We present the proof for monotone DNFs (see Appendix E.2 for the extension to monotone CNFs). Recall that d-SDNNFs are structured by *v-trees*, which generalize variable orders. We first introduce *treewidth*, a width notion that measures the performance of a v-tree by counting how many clauses it splits; and we connect *treewidth* to *treewidth*. We use again dncpi-sets, and argue that a d-SDNNF structured by a v-tree must shatter a dncpi-set whose size follows the *treewidth* of the v-tree. We then show that shattering a dncpi-set forces d-SDNNFs to be large: instead of the easy OBDD result of the previous section (Lemma 5.8), we will need a much deeper result of Pipatsrisawat and Darwiche [36, Theorem 3], rephrased in the setting of communication complexity by Bova, Capelli, Mengel, and Slivovsky [12].

Note that [12], by a similar approach, shows an exponential lower bound on the size of d-SDNNF which is reminiscent of ours. However, their bound again applies to one well-chosen family of Boolean functions; our contribution is to show a general lower bound. In essence, our result is shown by observing that the family of functions used in their lower bound occurs “within” any bounded-degree, bounded-arity monotone DNF or CNF. Also note that a result similar to the lower bound of Corollary 6.2 is proven by Capelli [18, Corollary 6.35] as an auxiliary statement to separate structured DNNFs and FBDDs. The result uses *MIM-width*, but Theorem 4.2.5 of [39], as degree and arity are bounded, implies that we could rephrase it to *treewidth*; further, the result assumes arity-2 formulas, but it could be extended to arbitrary arity as in [19, Theorem 12]. More importantly, the result applies only to monotone CNFs and not to DNFs (thanks to which the bound is shown to hold even for non-deterministic SDNNF).

Treesplitwidth. Informally, treesplitwidth is to v-trees what pathsplitwidth is to variable orders: it bounds the “best performance” of any v-tree.

► **Definition 6.3.** Let $H = (V, E)$ be a hypergraph, and T be a v-tree over V . For any node n of T , we define $\text{Split}_n(T, H)$ as the set of hyperedges e of H that contain both a variable in T_n and one outside T_n (recall that T_n denotes the subtree of T rooted at n). Formally: $\text{Split}_n(T, H) := \{e \in E \mid \exists v_i \in \text{Leaves}(T_n) \text{ and } \exists v_o \in \text{Leaves}(T \setminus T_n) \text{ such that } \{v_i, v_o\} \subseteq e\}$.

The *treesplitwidth* of T relative to H is $\text{tsw}(T, H) := \max_{n \in T} |\text{Split}_n(T, H)|$. The *treesplitwidth* $\text{tsw}(H)$ of H is then the minimum of $\text{tsw}(T, H)$ over all v-trees T of V .

As with pathsplitwidth and pathwidth (Lemma 5.6), we can bound the treewidth of a hypergraph by its treesplitwidth:

► **Lemma 6.4.** *For any hypergraph $H = (V, E)$, we have $\text{tw}(H) \leq 3 \times \text{arity}(H) \times \text{tsw}(H)$.*

Proof sketch. We construct a tree decomposition from a v-tree T that has same skeleton as T , where leaf bags contain the corresponding variable in the v-tree, and where internal bags contain the split at this v-tree node unioned with the split at the child nodes. This is indeed a tree decomposition because each non-singleton edge is split, and the nodes of the v-tree where a vertex of H occurs always form a connected subtree. ◀

Moreover, using the same techniques that we used in the last section, we can show the analogue of Lemma 5.9. Specifically, given a monotone DNF φ on variables V , a v-tree T over V , and a dncpi-set S of φ , we say that T *shatters* S if there is a node n in T such that $S \subseteq \text{Split}_n(T, \varphi)$. We now show that any v-tree over V must shatter a large dncpi-set (depending on the treewidth, degree, and arity):

► **Lemma 6.5.** *Let φ be a monotone DNF, H its hypergraph, and T be a v-tree over its variables. Then there is a dncpi-set S of φ shattered by T such that $|S| \geq \left\lfloor \frac{\text{tsw}(H)}{(\text{arity}(H) \times \text{degree}(H))^2} \right\rfloor$.*

Proof sketch. The proof is just like that of Lemma 5.9, except with the new definition of split on v-trees. In particular, we use Lemmas 5.10, 5.11, and 5.12. ◀

Hence, to conclude the proof of Theorem 6.1, the only missing ingredient is a lower bound on the size of d-SDNNFs that shatter large dncpi-sets. Specifically, we need the following analogue of Lemma 5.8:

► **Lemma 6.6.** *Let φ be a monotone DNF on variables V and $n \in \mathbb{N}$. Assume that, for every v-tree T over V , there is some dncpi-set S of φ with $|S| \geq n$, such that T shatters S . Then any d-SDNNF for φ has size $\geq 2^n - 1$.*

We will prove Lemma 6.6 in the rest of this section using a recent lower bound by Bova, Capelli, Mengel, and Slivovsky [12]. They bound the size of any d-SDNNF for the *set intersection* function, defined as $\text{SINT}_n := (x_1 \wedge y_1) \vee \dots \vee (x_n \wedge y_n)$. This bound is useful for us: a dncpi-set intuitively isolates some variables on which φ computes exactly SINT_n :

► **Lemma 6.7.** *Let φ be a DNF with variables V , and let $S = \{D_1, \dots, D_n\}$ be a dncpi-set of φ where every clause has size ≥ 2 . Pick two variables $x_i \neq y_i$ in D_i for each $1 \leq i \leq n$, and let $V' := \{x_1, y_1, \dots, x_n, y_n\}$. Then there is a partial valuation ν of V with domain $V \setminus V'$ such that $\nu(\varphi) = \text{SINT}_n$.*

Proof sketch. The valuation ν sets to 1 the variables V'' which are in $\bigcup S$ but not in V' , and sets to 0 all remaining variables. This amounts to discarding the clauses not in the dncpi-set, and discarding the variables of V'' in the dncpi-set: what remains of the DNF is then precisely SINT_n . Note that this result relies on monotonicity, and on the fact that φ is a DNF. (However, as we show in Appendix E.2, a dual result holds for monotone CNF.) ◀

This observation allows us to leverage the bound of [12] on the size of d-SDNNFs that compute SINT_n , assuming that they are structured by an “inconvenient” v-tree:

► **Proposition 6.8 [12, Proposition 14].** *Let $X_n = \{x_1, \dots, x_n\}$ and $Y_n = \{y_1, \dots, y_n\}$ for $n \in \mathbb{N}$, and let T be a v-tree over $X_n \sqcup Y_n$ such that there exists a node $n \in T$ with $X_n \subseteq \text{Leaves}(T_n)$ and $Y_n \subseteq \text{Leaves}(T \setminus T_n)$. Then any d-SDNNF structured by T computing SINT_n has size $\geq 2^n - 1$.*

In our setting, an “inconvenient” v-tree for a dncpi-set is one that shatters it: each clause of the dncpi-set is then partitioned in two non-empty subsets where we can pick x_i and y_i for Lemma 6.7. Hence, when every v-tree shatters a large dncpi-set of φ , Proposition 6.8 allows us to deduce the lower bound on the size of every d-SDNNF for φ . We have thus shown Lemma 6.6, and this concludes the proof of Theorem 6.1 (in the DNF case).

7 Application to Query Lineages

In this section, we adapt the lower bound of the previous section to the computation of query lineages on relational instances. Like in [7], for technical reasons, we must assume a graph signature. We first recall some preliminaries and then state our result.

Preliminaries. We fix a *graph signature* σ of relation names and arities in $\{1, 2\}$, with at least one relation of arity 2. An *instance* I on σ is a finite set of *facts* of the form $R(a_1, \dots, a_n)$ for n the arity of R ; we call a_1, \dots, a_n *elements* of I . An instance I' is a *subinstance* of I if the facts of I' are a subset of those of I . The *Gaifman graph* of I has the elements of I as vertices and has one edge between each pair of elements that co-occur in some fact of I . The *treewidth* $\text{tw}(I)$ of I is that of its Gaifman graph.

A *Boolean conjunctive query* (CQ) is an existentially quantified conjunction of *atoms* of the form $R(x_1, \dots, x_n)$ where the x_i are *variables*. A *UCQ* is a disjunction of CQs, and a UCQ^\neq also allows atoms of the form $x \neq y$. A UCQ^\neq is *connected* if the Gaifman graph of each disjunct (seen as an instance, and ignoring \neq -atoms) is connected. For instance, letting σ_R consist of one arity-2 relation R , the following connected UCQ^\neq tests if there are two facts that share one element: $Q_p : \exists xyz (R(x, y) \vee R(y, x)) \wedge (R(y, z) \vee R(z, y)) \wedge x \neq z$. (While Q_p is not given as a disjunction of CQs, it can be rewritten to one using distributivity.)

The *lineage* of a UCQ^\neq Q over I is a Boolean formula $\varphi(Q, I)$ on the facts of I that maps each Boolean valuation $\nu : I \rightarrow \{0, 1\}$ to 1 or 0 depending on whether I_ν satisfies Q or not, where $I_\nu := \{F \in I \mid \nu(F) = 1\}$. The lineage intuitively represents which facts of I suffice to satisfy Q . Lineages are useful to evaluate queries on *probabilistic databases* [38]: we can obtain the probability of the query from an OBDD or d-DNNF representing its lineage.

Problem statement. We study when query lineages can be computed efficiently in *data complexity*, i.e., as a function of the input instance, with the query being fixed. A first question asks which *queries* have *tractable lineages* on all instances: Jha and Suciu [32, Theorem 3.9] showed that *inversion-free* UCQ^\neq admit OBDD representations in this sense, and Bova and Szeider [15, Theorem 5] have recently shown that UCQ^\neq queries with inversions do not even have tractable d-SDNNF lineages. A second question asks which *instance classes* ensure that *all queries* have tractable lineages on them. This was studied for OBDD representations in [7]: bounded-treewidth instances have tractable OBDD lineage representations for any MSO query ([7, Theorem 6.5], using [32]); conversely there are *intricate* queries (a class of connected UCQ^\neq queries) whose lineages never have tractable OBDD representations in the instance treewidth [7, Theorem 8.7]. The query Q_p above is an example of an intricate query on the signature σ_R (refer to [7, Definition 8.5] for the formal definition of intricate queries).

This result shows that we must bound instance treewidth for all queries to have tractable OBDDs, but leaves the question open for more expressive lineage representations.

Result. Our bound in the previous section allows us to extend Theorem 8.7 of [7] from OBDDs to d-SDNNFs, yielding the following:

► **Theorem 7.1.** *There is a constant $d \in \mathbb{N}$ such that the following is true. Let σ be an arity-2 signature, and Q a connected UCQ $^\neq$ which is intricate on σ . For any instance I on σ , any d-SDNNF representing the lineage of Q on I has size $2^{\Omega(\text{tw}(I)^{1/d})}$.*

Proof sketch. As in [7], we use a result of Chekuri and Chuzhoy [20] to show that the Gaifman graph of I has a degree-3 topological minor S of treewidth $\Omega(\text{tw}(I)^{1/d})$ for some constant $d \in \mathbb{N}$; we also ensure that S has sufficiently high *girth* relative to Q . We focus on a subinstance I' of I that corresponds to S : this suffices to show our lower bound, because we can always compute a tractable representation of $\varphi(Q, I')$ from one of $\varphi(Q, I)$. Now, we can represent $\varphi(Q, I')$ as a minimized DNF ψ by enumerating its minimal matches: ψ has constant arity because the number of atoms of Q is fixed, and it has constant degree because S has constant degree and Q is connected. Further, as Q is intricate and I' has high girth relative to Q , we can ensure that this DNF has treewidth $\Omega(\text{tw}(I'))$. We conclude by Theorem 6.1 that d-SDNNF representations of $\varphi(Q, I')$, hence $\varphi(Q, I)$, have size $2^{\Omega(\text{tw}(I)^{1/d})}$. ◀

To summarize, given an instance family \mathcal{I} satisfying the constructibility requirement of Theorem 8.1 of [7], there are two regimes: (i.) \mathcal{I} has bounded treewidth and then all MSO queries have d-SDNNF lineages on instances of \mathcal{I} that are computable in linear time; or (ii.) the treewidth is unbounded and then there are UCQ $^\neq$ queries (the intricate ones) whose lineages on instances of \mathcal{I} have no d-SDNNF representations polynomial in the instance size.

8 Conclusion

We have shown tight connections between structured circuit classes and width measures on circuits. We constructively rewrite bounded-treewidth circuits to d-SDNNFs in time linear in the circuit and singly exponential in the treewidth, and show matching lower bounds for arbitrary monotone CNFs or DNFs under degree and arity assumptions; we also show a lower bound for pathwidth and OBDDs. Our results have applications to rich query evaluation: probabilistic query evaluation, computation of lineages, enumeration, etc.

Our work also raises a number of open questions. First, the d-SDNNF obtained in the proof of Theorem 3.3 does *not* respect the definition of a *sentential decision diagram* (SDD) [23]. Can this be fixed, and Theorem 3.3 extended to SDDs? Or is it impossible, which could solve the open question [10] of separating SDDs and d-SDNNFs? Second, can we weaken the hypotheses of bounded degree and arity in Corollaries 5.4 and 6.2, and can we rephrase the latter to a notion of d-SDNNF width to match more closely the statement of the former? Last, Section 7 shows that d-SDNNF representations of the lineages of intricate queries are exponential in the treewidth; we conjecture a similar result for pathwidth and OBDDs, but this would require a pathwidth analogue of the minor extraction results of [20].

Acknowledgments. We acknowledge Chandra Chekuri for his helpful comments at <https://csttheory.stackexchange.com/a/38943/>, as well as Florent Capelli for pointing out the connection to [18, Corollary 6.35] and [39].

References

- 1 Antoine Amarilli. *Leveraging the Structure of Uncertain Data*. PhD thesis, Télécom Paris-Tech, 2016.
- 2 Antoine Amarilli, Pierre Bourhis, Louis Jachiet, and Stefan Mengel. A circuit-based approach to efficient enumeration. In *ICALP*, 2017.
- 3 Antoine Amarilli, Pierre Bourhis, Mikaël Monet, and Pierre Senellart. Combined tractability of query evaluation via tree automata and cycluits. In *ICDT*, 2017.
- 4 Antoine Amarilli, Pierre Bourhis, Mikaël Monet, and Pierre Senellart. Combined tractability of query evaluation via tree automata and cycluits (Extended Version). *CoRR*, abs/1612.04203, 2017. Extended version of [3].
- 5 Antoine Amarilli, Pierre Bourhis, and Pierre Senellart. Provenance circuits for trees and treelike instances (extended version). *CoRR*, abs/1511.08723, 2015. Extended version of [6].
- 6 Antoine Amarilli, Pierre Bourhis, and Pierre Senellart. Provenance circuits for trees and treelike instances. In *ICALP*, July 2015.
- 7 Antoine Amarilli, Pierre Bourhis, and Pierre Senellart. Tractable lineages on treelike instances: limits and extensions. In *PODS*, 2016.
- 8 Paul Beame and Vincent Liew. New limits for knowledge compilation and applications to exact model counting. In *UAI*, 2015.
- 9 Hans L Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6), 1996.
- 10 Simone Bova. SDDs are exponentially more succinct than OBDDs. In *AAAI*, 2016.
- 11 Simone Bova, Florent Capelli, Stefan Mengel, and Friedrich Slivovsky. A strongly exponential separation of DNNFs from CNF formulas. *CoRR*, abs/1411.1995, 2015.
- 12 Simone Bova, Florent Capelli, Stefan Mengel, and Friedrich Slivovsky. Knowledge compilation meets communication complexity. In *IJCAI*, 2016.
- 13 Simone Bova and Friedrich Slivovsky. On compiling structured CNFs to OBDDs. In *CSR*, 2015.
- 14 Simone Bova and Friedrich Slivovsky. On compiling structured CNFs to OBDDs. *TCS*, 61(2), 2017.
- 15 Simone Bova and Stefan Szeider. Circuit treewidth, sentential decision, and query compilation. In *PODS*, 2017.
- 16 Randal E. Bryant. Symbolic Boolean manipulation with ordered binary-decision diagrams. *ACM Comput. Surv.*, 24(3), 1992.
- 17 Andrea Cali, Florent Capelli, and Igor Razgon. Non-FPT lower bounds for structural restrictions of decision DNNF. *CoRR*, abs/1708.07767, 2017.
- 18 Florent Capelli. *Structural restrictions of CNF-formulas: applications to model counting and knowledge compilation*. PhD thesis, Université Paris-Diderot, 2016.
- 19 Florent Capelli. Understanding the complexity of #SAT using knowledge compilation. In *LICS*, 2017.
- 20 Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. In *STOC*, 2014.
- 21 Adnan Darwiche. On the tractable counting of theory models and its application to truth maintenance and belief revision. *J. Applied Non-Classical Logics*, 11(1-2), 2001.
- 22 Adnan Darwiche. A differential approach to inference in Bayesian networks. *JACM*, 50(3), 2003.
- 23 Adnan Darwiche. SDD: A new canonical representation of propositional knowledge bases. In *IJCAI*, 2011.
- 24 Srinivas Devadas. Comparing two-level and ordered binary decision diagram representations of logic functions. *IEEE TCAD*, 12(5), 1993.

- 25 Feodor F. Dragan, Fedor V. Fomin, and Petr A. Golovach. Spanners in sparse graphs. *JCSS*, 77(6), 2011.
- 26 Andrea Ferrara, Guoqiang Pan, and Moshe Y Vardi. Treewidth in verification: Local vs. global. In *LPAR*, 2005.
- 27 Daan Fierens, Guy Van den Broeck, Joris Renkens, Dimitar Shterionov, Bernd Gutmann, Ingo Thon, Gerda Janssens, and Luc De Raedt. Inference and learning in probabilistic logic programs using weighted Boolean formulas. *TPLP*, 15(3), 2015.
- 28 Jörg Flum, Markus Frick, and Martin Grohe. Query evaluation via tree-decompositions. *J. ACM*, 49(6), 2002.
- 29 Martin Grohe and Dániel Marx. On tree width, bramble size, and expansion. *J. Combinatorial Theory, Series B*, 99(1), 2009.
- 30 Abhay Kumar Jha, Dan Olteanu, and Dan Suciu. Bridging the gap between intensional and extensional query evaluation in probabilistic databases. In *EDBT*, 2010.
- 31 Abhay Kumar Jha and Dan Suciu. Knowledge compilation meets database theory: compiling queries to decision diagrams. In *ICDT*, 2011.
- 32 Abhay Kumar Jha and Dan Suciu. On the tractability of query compilation and bounded treewidth. In *ICDT*, 2012.
- 33 Stasys Jukna. *Boolean Function Complexity*. Springer, 2012.
- 34 Steffen L. Lauritzen and David J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *J. Royal Statistical Society. Series B*, 1988.
- 35 Knot Pipatsrisawat and Adnan Darwiche. New compilation languages based on structured decomposability. In *AAAI*, 2008.
- 36 Knot Pipatsrisawat and Adnan Darwiche. A lower bound on the size of decomposable negation normal form. In *AAAI*, 2010.
- 37 Igor Razgon. On OBDDs for CNFs of bounded treewidth. In *KR*, 2014.
- 38 Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch. *Probabilistic Databases*. Morgan & Claypool, 2011.
- 39 Martin Vatshelle. *New width parameters of graphs*. PhD thesis, University of Bergen, 2012.

A Proofs for Section 2 (Preliminaries)

► **Lemma 2.2.** *Given a tree decomposition T of width k having n nodes, we can compute in time $O(k \times n)$ a nice tree decomposition T' of width k and $O(k \times n)$ nodes.*

Proof. We first make the tree decomposition binary (but not necessarily full) by replacing each bag b with children b_1, \dots, b_n with $n > 2$ by a chain of bags with the same label as b to which we attach the children b_1, \dots, b_n . This process is in time $O(n)$ and does not change the width.

We then ensure the second and third conditions, by applying a transformation to leaf bags and to internal bags. We modify every leaf bag b containing more than one vertex by a chain of at most k internal bags with leaves where the vertices are added one after the other. At the same time, we modify every internal bag b that contains elements v_1, \dots, v_n not present in the union D of its children: we replace b by a chain of at most k internal bags b'_1, \dots, b'_n containing respectively $b, b \setminus \{v_n\}, b \setminus \{v_n, v_{n-1}\}, \dots, D$, each bag having a child introducing the corresponding gate v_i . This is in time $O(kn)$, and again it does not change the width; further, the result of the process satisfies the second and third conditions and obviously it is still a binary tree.

We next ensure the fourth and fifth condition. To do this, whenever a non-root bag b and its parent b' violate the condition, we create a chain of intermediate nodes where the gates in $b \setminus b'$ are removed one after the other; and we replace the root bag by a chain of bags where the elements of the root bag are removed one after the other. This does not affect the fact that the tree decomposition is binary, or the second and third conditions, it does not change the width, and it runs in time $O(kn)$. Observe that the bound is $O(kn)$ in the original tree decomposition (not in the output of the previous step), because the transformation never needs to be performed within the chains of nodes that we have introduced in the previous step; it only needs to be performed on interfaces between bags that correspond to interfaces between original bags in T .

The only missing part is to ensure that the tree decomposition is full, which we can simply ensure in linear time by adding bags with an empty label as a second children for internal nodes that have only one child. This is obviously in linear time, does not change the width, and does not affect the other conditions, concluding the proof. ◀

B Proofs for Section 3 (Upper Bounds)

► **Theorem 3.2.** *Given any circuit C of treewidth k , we can compute an equivalent d -SDNNF in linear time parametrized by k , i.e., in time $O(|C| \times f(k))$ for some computable function f .*

Proof. We first define a signature σ to encode circuits into relational instances: the elements of such an instance are gates, and σ features unary relations Or, And, Not, Variable to describe the type of each gate, a unary relation Output to identify the output gate, and a binary relation Wire to describe the wires. We also add a unary relation True to describe a valuation: the relation applies to variable gates to indicate whether they are true or not.

Now, we write an MSO formula φ on σ which checks whether a circuit evaluates to true under the indicated valuation. The formula φ existentially guesses a set S of true gates: it asserts (i) that the variable gates of S are exactly the ones indicated as True in the input instance, (ii) that the output gate of the circuit is in S , and (iii) that S satisfies the semantics

of internal gate, i.e., each or-gate is in S iff it has an input in S , each and-gate is in S iff it has all its inputs in S , and each not-gate is in S iff its input is not in S .

Now, given an input circuit C with variable gates V , we encode it in linear time into a relational structure I on σ in the expected way. For any subset V' of V , we let $I_{V'}$ be I where we add a fact $\text{True}(g)$ for each $g \in V$. In particular, the instance I_V can be constructed in linear time, and its treewidth is the same as that of C . We now use [7, Theorem 6.11] to compute in linear time in I_V (hence in C) a d-SDNNF¹ C' capturing the *provenance* of φ on I_V , i.e., for any Boolean valuation ν of I_V , letting $I' := \{F \in I_V \mid \nu(F) = 1\}$, we have $\nu(C') = 1$ iff $I' \models \varphi$. We fix to 1 the value of the inputs of C' that stands for facts other than True-facts (i.e., those that describe the structure of the circuit). The result C'' is computed in linear time from C , it is still a d-SDNNF, and it ensures that for any Boolean valuation ν of V , letting $V' := \{g \in V \mid \nu(g) = 1\}$, we have $\nu(C') = 1$ iff $I_{V'} \models \varphi$. By construction, the latter holds iff $\nu(C) = 1$. In other words, C'' is equivalent to C , which concludes the proof. ◀

► **Corollary 3.4.** *Given a Boolean circuit C , a tree decomposition T of width k of C , and a probability valuation π of C , we can compute $\pi(C)$ in $O(|\pi| + |T| \times 2^{(4+\varepsilon)k})$ for any $\varepsilon > 0$.*

Proof. Use Theorem 3.3 to compute an equivalent d-SDNNF C' ; as C and C' are equivalent, it is clear that $\pi(C) = \pi(C')$. Now, compute the probability $\pi(C')$ in linear time in C' and $|\pi|$ by a simple bottom-up pass, using the fact that C' is a d-DNNF [21]. ◀

► **Corollary 3.5.** *Given a Boolean circuit C and a tree decomposition T of width k of C , we can enumerate the accepting assignments of C with preprocessing in $O(|T| \times 2^{(4+\varepsilon)k})$ and delay linear in the size of each produced assignment.*

Proof. Use Theorem 3.3 to compute an equivalent d-SDNNF C' , which has the same accepting valuations, along with a v-tree T' of C' . We now conclude using [2, Theorem 2.1]. ◀

C Proofs for Section 4 (Proof of the Main Upper Bound Result)

We now prove that the circuit D constructed in the main text is indeed a d-SDNNF equivalent to the initial circuit C , and that it can be constructed together with its v-tree in $O(|T| \times 2^{(4+\varepsilon)k})$ for any $\varepsilon > 0$.

C.1 D is a Structured DNNF

Negations only apply to the input gates, so D is an NNF. To justify that D is structured, we will define an appropriate v-tree T' . Consider the nice tree decomposition T of C that was used to construct D . For each bag $b \in T$, the v-tree T' has a node b' . For each internal bag $b \in T$ with children b_1, b_2 , T' has a node b'_{children} , whose children are b'_1 and b'_2 , and whose parent is b' . For every bag $b \in T$ that is responsible of some variable gate g , T' has a node $G^{g,1}$ whose parent is b' . Hence T' is a tree, and one can check that any any $G_b^{\nu_l, I_l, \nu_r, I_r}$ is structured by b'_{children} and any $G_b^{\nu, S}$ is structured by b' , so that T' structures D .

¹ The result only states that it computes a d-DNNF, not a d-SDNNF. However, it is immediate from the construction that the circuit actually follows a v-tree, which is given by the tree encoding of the input instance. This can be checked from the proof (given as that of [1, Theorem 3.5.8]).

C.2 D is Equivalent to C

In order to prove that D is equivalent to C , we introduce the standard notion of a *trace* in an NNF:

► **Definition C.1.** Let D be an NNF, χ a valuation of its variable gates, and g a gate that evaluates to 1 under χ . A *trace* of D starting at g according to χ is a set Ξ of gates of D that is minimal by inclusion and such that:

- $g \in \Xi$;
- If $g' \in \Xi$ and g' is an and gate, then $W(g') \subseteq \Xi$, where $W(g)$ denotes the set of gates that are input to g ;
- If $g' \in \Xi$ and g' is an or gate, then exactly one input of g' that evaluates to 1 is in Ξ .

The first step is then to prove that traces have exactly one almost-evaluation corresponding to each descendant bag, and that these almost-evaluations mutually agree.

► **Lemma C.2.** Let χ be a valuation of the variable gates, $G_b^{\nu,S}$ a gate in D that evaluates to 1 under χ and Ξ be a trace of D starting at $G_b^{\nu,S}$ according to χ . Then for any bag $b' \leq b$ (meaning that b' is b or a descendant of b), Ξ contains exactly one gate of the form $G_{b'}^{\nu',S'}$. Moreover, over all $b' \leq b$, all the almost-evaluations of the gates $G_{b'}^{\nu',S'}$ that are in Ξ mutually agree.

Proof. The fact that Ξ contains exactly one gate $G_{b'}^{\nu',S'}$ for any bag $b' \leq b$ is obvious by construction of D , as or-gates are assumed to have exactly one input evaluated to 1 in Ξ . For the second claim, suppose by contradiction that not all the almost-evaluations of the gates $G_{b'}^{\nu',S'}$ that are in Ξ mutually agree. We would then have $G_{b_1}^{\nu_1,S_1}$ and $G_{b_2}^{\nu_2,S_2}$ in Ξ and $g \in b_1 \cap b_2$ such that $\nu_1(g) \neq \nu_2(g)$. But because T is a tree decomposition, g appears in all the bags on the path from b_1 and b_2 , and by construction the almost-evaluations of the $G_{b'}^{\nu',S'}$ on this path that are in Ξ mutually agree, hence a contradiction. ◀

Therefore, Lemma C.2 allows us to define the union of the almost-evaluations in such a trace:

► **Definition C.3.** Let χ be a valuation of the variable gates, G_b^ν a gate in D that evaluates to 1 under χ and Ξ be a trace of D starting at G_b^ν according to χ . Then $\gamma(\Xi) := \bigcup_{G_{b'}^{\nu',S'} \in \Xi} \nu'$ (the union of the almost-evaluations in Ξ , which is a valuation from $\bigcup_{G_{b'}^{\nu',S'} \in \Xi} b'$ to $\{0, 1\}$) is properly defined.

We now need to prove a few lemmas about the behavior of gates that are innocent (i.e., not suspicious).

► **Lemma C.4.** Let χ be a valuation of the variable gates, $G_b^{\nu,S}$ a gate in D that evaluates to 1 under χ and Ξ be a trace of D starting at $G_b^{\nu,S}$ according to χ . Let $g \in b$ be a gate that is innocent ($g \notin S$). Then the following holds:

- If $\nu(g)$ is a weak value of g , then for every input g' of g that is in the domain of $\gamma(\Xi)$ (i.e., g' appears in a bag $b' \leq b$), we have that $\gamma(\Xi)$ maps g' to a weak value of g ;
- If $\nu(g)$ is a strong value of g , then there exists an input g' of g that is in the domain of $\gamma(\Xi)$ such that $\gamma(\Xi)(g')$ is $\nu(g)$ if g is an and or or gate, and $\gamma(\Xi)(g')$ is $1 - \nu(g)$ if g is a not gate.

Proof. We prove the claim by bottom-up induction on $b \in T$. One can easily check that the claim is true when b is a leaf bag, remembering that in this case we must (crucially) have $S = \text{Unj}(\nu)$ by construction (that is, all the gates that are unjustified are suspicious). For the induction case, let b_l, b_r be the children of b . Suppose first that $\nu(g)$ is the weak value of g , and suppose for a contradiction that there is an input g' of g in the domain of $\gamma(\Xi)$ such that $\gamma(\Xi)(g')$ is a strong value of g . By the occurrence and connectedness properties of tree decompositions, there exists a bag $b' \leq b$ in which both g and g' occur. Consider the gate $G_{b'}^{\nu', S'}$ that is in Ξ : by Lemma C.2, this gate exists and is unique. By definition of $\gamma(\Xi)$ we have $\nu'(g') = \gamma(\Xi)(g')$. Because ν' is a (C, b') -almost-evaluation that maps g' to a strong value of g , we must have that $\nu'(g)$ is also a strong value of g , thus contradicting our hypothesis that $\nu(g) = \gamma(\Xi)(g) = \nu'(g)$ is a weak value for g .

Suppose now that $\nu(g)$ is a strong value of g . We only treat the case when g is an or or an and gate, as the case of a not gate is similar. We distinguish two sub-cases:

- g is justified. Then clearly, because ν is a (C, b) -almost-evaluation, there must exist an input g' of g that is also in b such that $\nu(g')$ is a strong value of g , which proves the claim.
- g is unjustified but innocent ($g \notin S$). By construction (precisely, by the second item of Definition 4.5), g must then be innocent for a child of b , and the claim clearly holds by induction hypothesis. ◀

Lemma C.4 allows us to show that for a gate g , letting b be the topmost bag in which g appears (hence, each input of g must occur in some bag $b' \leq b$), if g is innocent then for any trace Ξ starting at a gate for bag b , $\gamma(\Xi)$ respects the semantics of g . Formally:

► **Lemma C.5.** *Let χ be a valuation of the variable gates, $G_b^{\nu, S}$ a gate in D that evaluates to 1 under χ and Ξ be a trace of D starting at $G_b^{\nu, S}$ according to χ . Let $g \in b$ be a gate such that b is the topmost bag in which g appears (hence $W(g) \subseteq \text{domain}(\gamma(\Xi))$). If g is innocent ($g \notin S$) then $\gamma(\Xi)$ respects the semantics of g , that is $\gamma(\Xi)(g) = \odot \gamma(\Xi)(W(g))$ where \odot is the type of g .*

Proof. Clearly implied by Lemma C.4. ◀

We need one last lemma about the behavior of suspicious gates, which intuitively tells us that if we have already seen all the input gates of a gate g and g is still suspicious, then g can never escape:

► **Lemma C.6.** *Let χ be a valuation of the variable gates, $G_b^{\nu, S}$ a gate in D that evaluates to 1 under χ and Ξ be a trace of D starting at $G_b^{\nu, S}$ according to χ . Let g be a gate such that the topmost bag b' in which g appears is $\leq b$, and consider the unique gate of the form $G_{b'}^{\nu', S'}$ that is in Ξ . If $g \in S'$ then $b' = b$ (hence $G_b^{\nu, S} = G_{b'}^{\nu', S'}$ by uniqueness).*

Proof. Let $g \in S'$. Suppose by contradiction that $b' \neq b$. Let p be the parent of b' (which exists because $b' < b$). It is clear that by construction (ν', S') is connectible to p (recall Definition 4.4), hence g must be in p , contradicting the fact that b' should have been the topmost bag in which g occurs. Hence $b' = b$. ◀

We now have all the results that we need to show that $D \implies C$, i.e., for every valuation χ of the variables of C , if $\chi(D) = 1$ then $\chi(C) = 1$. We prove a stronger result:

► **Lemma C.7.** *Let χ be a valuation of the variable gates, $G_{\text{root}(T)}^{\nu, \emptyset} \in D$ a gate that evaluates to 1 under χ , and Ξ a trace of D starting at $G_{\text{root}(T)}^{\nu, \emptyset}$ according to χ . Then $\gamma(\Xi)$ corresponds to the evaluation χ of C .*

Proof. We prove by induction on C (as its graph is a DAG) that for all $g \in C$, $\gamma(\Xi)(g) = \chi(g)$. When g is a variable gate, consider the bag b' that is responsible of g , and consider the gate $G_{b'}^{\nu', S'}$ that is in Ξ : this gate exists and is unique according to Lemma C.2. This gate evaluates to 1 under χ (because it is in the trace), which is only possible if $G^{g, \nu'(g)}$ evaluates to 1 under χ , hence by construction we must have $\nu'(g) = \chi(g)$ and then $\gamma(\Xi)(g) = \chi(g)$. Now suppose that g is an internal gate, and consider the topmost bag b' in which g appears. Consider again the unique $G_{b'}^{\nu', S'}$ that is in Ξ . By induction hypothesis we have that $\gamma(\Xi)(g') = \chi(g')$ for every input g' of g . We now distinguish two cases:

- $b' = \text{root}(T)$. Therefore by Lemma C.5 we know that $\gamma(\Xi)$ respects the semantics of g , which means that $\gamma(\Xi)(g) = \odot \gamma(\Xi)(W(g)) = \odot \chi(W(g)) = \chi(g)$ (where the second equality comes from the induction hypothesis and the third equality is just the definition of the evaluation χ of C), which proves the claim.
- $b' < \text{root}(T)$. But then by Lemma C.6 we must have $g \notin S'$ (because otherwise we should have $b' = \text{root}(T)$ and then $S' = \emptyset$), that is g is innocent for $G_{b'}^{\nu', S'}$. Therefore, again by Lemma C.5, it must be the case that $\gamma(\Xi)$ respects the semantics of g , and we can again show that $\gamma(\Xi)(g) = \chi(g)$, concluding the proof. ◀

This indeed implies that $D \implies C$: let χ be a valuation of the variable gates and suppose $\chi(D) = 1$. Then by definition of the output of D , it means that the gate $G_{\text{root}(T)}^{\nu, \emptyset}$ such that $\nu(g_{\text{output}}) = 1$ evaluates to 1 under χ . But then, considering a trace Ξ of D starting at $G_{\text{root}(T)}^{\nu, \emptyset}$ according to χ , we have that $\chi(g_{\text{output}}) = \gamma(\Xi) = \nu(g_{\text{output}}) = 1$. To show the converse ($C \implies D$), once can simply observe the following phenomenon:

► **Lemma C.8.** *Let χ be a valuation of the variable gates. Then for every bag $b \in T$, the gate $G_b^{\chi|_b, S}$ evaluates to 1 under χ , where S is the set of gates $g \in \text{Unj}(\nu)$ such that for all g' input of g that appears in some bag $b' \leq b$, then $\chi(g')$ is a weak value of g .*

Proof. Easily proved by bottom-up induction. ◀

Now suppose $\chi(C) = 1$. By Lemma C.8 we have that $G_{\text{root}(T)}^{\chi|_{\text{root}(T)}, \emptyset}$ evaluates to 1 under χ , and because $\chi(g_{\text{output}}) = 1$ we have that $\chi(D) = 1$. Hence, we have proved that D is equivalent to C .

C.3 D is Deterministic

We now prove that D is deterministic, i.e., that every or gate in D is deterministic. Recall that the only or gates in D are the gates of the form $G_{b, \text{children}}^{\nu, S}$. We will in fact prove that traces are unique, which clearly implies that all the or gates are deterministic.

We just need to prove the following lemma to reach our goal:

► **Lemma C.9.** *Let χ be a valuation of the variable gates, $G_b^{\nu, S}$ a gate in D that evaluates to 1 under χ and Ξ be a trace of D starting at $G_b^{\nu, S}$ according to χ . Let $g \in b$. Then the following is true:*

- if g is innocent ($g \notin S$) and $\nu(g)$ is a strong value of g , then there exists an input g' of g such that $\gamma(\Xi)(g)$ is a strong value for g .

- if $g \in S$, then for every input g' of g that is in the domain of $\gamma(\Xi)$, we have that $\gamma(\Xi)(g')$ is a weak value for g .

Proof. We prove the two claims independently:

- Let $g \in b$ such that $g \notin S$ and $\nu(g)$ is a strong value for g . Then the claim directly follows from the second item of Lemma C.4.
- We prove the second claim via a bottom-up induction on T . When b is a leaf then it is trivially true because g has no input g' in b because $|b| \leq 1$ because T is nice. For the induction case, let $G_{b_l}^{\nu_l, S_l}$ and $G_{b_r}^{\nu_r, S_r}$ be the (unique) gates in Ξ corresponding to the children b_l, b_r of b . By hypothesis we have $g \in S$. By definition of a gate being suspicious, we know that $\nu(g)$ is a strong value for g . To reach a contradiction, assume that there is an input g' of g in the domain of $\gamma(\Xi)$ such that $\gamma(\Xi)(g')$ is a strong value for g . Clearly this g' is not in b , because g is unjustified by ν (because $S \subseteq \text{Unj}(\nu)$). Either g' occurs in a bag $b'_l \leq b_l$, or it occurs in a bag $b'_r \leq b_r$. The two cases are symmetric, so we assume that we are in the former. As $g \in b$ and $g' \in b'_l$, by the properties of tree decompositions and because $g' \notin b$, we must have $g \in b_l$. Hence, by the contrapositive of the induction hypothesis on b_l applied to g , we deduce that $g \notin S_l$. But then by the second item of Definition 4.5, g should be innocent for $G_b^{\nu, S}$, that is $g \notin S$, which is a contradiction. ◀

We are ready to prove that traces are unique. Let us first introduce some useful notations: Let U, U' be sets, ν, ν' be valuations. We write $(\nu, U) = (\nu', U')$ to mean $\nu = \nu'$ and $U = U'$, and we write $(\nu, U)(g) = (\nu', U')(g)$ to mean that $\nu(g) = \nu'(g)$ and that we have $g \in U$ iff $g \in U'$. We show the following:

► **Lemma C.10.** *Let χ be a valuation of the variable gates such that $G_{b, \text{children}}^{\nu, S}$ evaluates to 1 under χ . Then there is a unique trace of D starting at $G_{b, \text{children}}^{\nu, S}$ according to χ .*

Proof. We will prove the claim by bottom-up induction on T . The case when b is a leaf is vacuous because there are no gates of the form $G_{b, \text{children}}^{\nu, S}$ for a leaf. For the inductive case, let b be an internal bag with children b_l and b_r . By induction hypothesis for every $G_{b_l, \text{children}}^{\nu_l, S_l}$ (resp., $G_{b_r, \text{children}}^{\nu_r, S_r}$) that evaluates to 1 under χ there exists a unique trace Ξ_l (resp., Ξ_r) of D starting at $G_{b_l, \text{children}}^{\nu_l, S_l}$ (resp., $G_{b_r, \text{children}}^{\nu_r, S_r}$). Hence, if by contradiction there are more than two traces of D starting at $G_{b, \text{children}}^{\nu, S}$, it can only be because $G_{b, \text{children}}^{\nu, S}$ is not deterministic, i.e., because at least two different inputs of $G_{b, \text{children}}^{\nu, S}$ evaluate to 1 under χ , say $G_b^{\nu_l, S_l, \nu_r, S_r}$ and $G_b^{\nu'_l, S'_l, \nu'_r, S'_r}$ with $(\nu_l, S_l) \neq (\nu'_l, S'_l)$ or $(\nu_r, S_r) \neq (\nu'_r, S'_r)$. W.l.o.g. we can suppose that it is $(\nu_l, S_l) \neq (\nu'_l, S'_l)$. Hence there exists $g_0 \in b_l$ such that $(\nu_l, S_l)(g_0) \neq (\nu'_l, S'_l)(g_0)$. Let Ξ_l be the trace of D starting at $G_{b_l}^{\nu_l, S_l}$ and Ξ'_l be the trace of D starting at $G_{b_l}^{\nu'_l, S'_l}$. We observe the following simple fact about Ξ_l and Ξ'_l :

- (*) for any g , if $\gamma(\Xi_l)(g) \neq \gamma(\Xi'_l)(g)$ then $g \notin b$. Indeed otherwise we should have $\nu(g) = \nu_l(g) = \gamma(\Xi_l)(g)$ and $\nu(g) = \nu'_l(g) = \gamma(\Xi'_l)(g)$, which is impossible.

Now we will define an operator θ that takes as input a gate g such that $(\gamma(\Xi_l), S_l)(g) \neq (\gamma(\Xi'_l), S'_l)(g)$, and outputs another gate $\theta(g)$ which is an input of g and such that again $(\gamma(\Xi_l), S_l)(\theta(g)) \neq (\gamma(\Xi'_l), S'_l)(\theta(g))$. This will lead to a contradiction because for any $n \in \mathbb{N}$, starting with g_0 and applying θ n times consecutively we would obtain n mutually distinct gates (because C is acyclic), but C has a finite number of gates.

Let us now prove that θ exists: let g such that $(\gamma(\Xi_l), S_l)(g) \neq (\gamma(\Xi'_l), S'_l)(g)$. We distinguish two cases:

- We have $(\gamma(\Xi_l), S_l)(g) \neq (\gamma(\Xi'_l), S'_l)(g)$ because $\gamma(\Xi_l)(g) \neq \gamma(\Xi'_l)(g)$. Then by (*), we know for sure that $g \notin b$. Therefore the topmost bag b' in which g occurs is $\leq b_l$. Let $G_{b',S'}^{\nu',S'}$ be the gate in Ξ_l and $G_{b',S''}^{\nu'',S''}$ the gate in Ξ'_l (they exist and are unique by Lemma C.2). Then by Lemma C.6 we must have $g \notin S'$ and $g \notin S''$, because otherwise we should have $b' = b$, which is not true. Hence, by Lemma C.5 we know that both $\gamma(\Xi_l)$ and $\gamma(\Xi'_l)$ respect the semantics of g . But we have $\gamma(\Xi_l)(g) \neq \gamma(\Xi'_l)(g)$, so there must exist an input g' of g such that $\gamma(\Xi_l)(g') \neq \gamma(\Xi'_l)(g')$! We can thus take $\theta(g)$ to be g' .
- We have $(\gamma(\Xi_l), S_l)(g) \neq (\gamma(\Xi'_l), S'_l)(g)$ because (w.l.o.g.) $g \notin S_l$ and $g \in S'_l$. Observe that this implies that $g \in b_l$, and that $\nu'_l(g)$ is a strong value for g . We can assume that $\nu_l(g) = \nu'_l(g)$, as otherwise we would have $\gamma(\Xi_l)(g) \neq \gamma(\Xi'_l)(g)$, which is a case already covered by the last item. Hence $\nu_l(g)$ is also a strong value for g , but we have $g \notin S_l$, so by the first item of Lemma C.9 we know that there exists an input g' of g that occurs in some bag $\leq b_l$ and such that $\gamma(\Xi_l)(g')$ is a strong value for g . We show that $\gamma(\Xi'_l)(g')$ must in contrast be a weak value for g , so that we can take $\theta(g)$ to be g' and conclude the proof. Indeed suppose by way of contradiction that $\gamma(\Xi'_l)(g')$ is a strong value for g . By the contrapositive of the second item of Lemma C.9, we get that $g \notin S'_l$, which contradicts our assumption.

Hence we proved that θ exists, which shows a contradiction, which means that in fact we must have $G_b^{\nu_l, S_l, \nu_r, S_r} = G_b^{\nu'_l, S'_l, \nu'_r, S'_r}$, so that $G_{b, \text{children}}^{\nu, S}$ is deterministic, which proves that there is a unique trace of D starting at $G_{b, \text{children}}^{\nu, S}$ according to χ , which was our goal. ◀

This concludes the proof that D is deterministic, and thus that D is a d-SDNNF equivalent to C .

C.4 Analysis of the Running Time

We last check that the construction can be performed in time $O(|T| \times f(k))$, where f is in $O(2^{(4+\varepsilon)k})$ for any $\varepsilon > 0$:

- From the initial tree decomposition T of C , we paid $O(k|T|)$ to compute the nice tree decomposition T_{nice} of size $O(k|T|)$;
- We computed the mapping $\text{Var}T_{\text{nice}}$ in linear time in T_{nice} ;
- We can clearly compute the v-tree in linear time from T_{nice} ;
- For each bag b of T_{nice} we have $2^{2|b|} \leq 2^{2k+2}$ different pairs of a valuation ν of b and of a subset S of b , and checking if ν is a (C, b) -almost-evaluation and if S is a subset of the unjustified gates of ν can be done in polynomial time in $|b| \leq k+1$ (we access the inputs and the type of each gate in RAM, i.e., in constant time, from C), hence we pay $O(|T_{\text{nice}}| \times p(k) \times 2^{2k})$ to create the gates of the form $G_b^{\nu, S}$, for p some polynomial;
- We pay an additional $O(|T_{\text{nice}}| \times 2^{4k})$ to create the gates of the form $G_b^{\nu_l, S_l, \nu_r, S_r}$;
- We pay an additional $O(|T_{\text{nice}}| \times p'(k) \times 2^{4k})$ to connect the gates of the form $G_{b, \text{children}}^{\nu, S}$ to their inputs (p' being again some polynomial).

Hence the total cost is indeed in $O(|T| \times f(k))$, where f is in $O(2^{(4+\varepsilon)k})$ for any $\varepsilon > 0$.

D Proofs for Section 5 (Lower Bounds on OBDDs)

► **Lemma 5.6.** *For any hypergraph $H = (V, E)$, we have $\text{pw}(H) \leq \text{arity}(H) \times \text{psw}(H)$.*

Proof. Let $H = (V, E)$ be a hypergraph, and let \mathbf{v} be an enumeration of the nodes of H witnessing that H has pathsplitwidth $\text{psw}(H)$. We will construct a path decomposition of H of width $\leq \text{arity}(H) \times \text{psw}(H)$. Consider the path $P = b_1, \dots, b_{|V|}$ and the labeling

function λ where $\lambda(b_i) := \{v_i\} \cup \bigcup \text{Split}_i(\mathbf{v}, H)$ for $1 \leq i \leq |V|$. Let us show that (P, λ) is a path decomposition of H : once this is established, it is clear that its width will be $\leq \text{arity}(H) \times \text{psw}(H)$.

First, we verify the occurrence condition. Let $e \in E$. If e is a singleton $\{v_i\}$ then e is included in b_i . Now, if $|e| \geq 2$, then let v_i be the first element of e enumerated by \mathbf{v} . We have $e \in \text{Split}_i(\mathbf{v}, H)$, and therefore e is included in b_i .

Second, we verify the connectedness condition. Let v be a vertex of H , then by definition $v \in b_i$ iff $v = v_i$ or there exists $e \in \text{Split}_i(\mathbf{v}, H)$ with $v \in e$. We must show that the set T_v of the bags that contain v forms a connected subpath in P . To show this, first observe that for every $e \in E$, letting $\text{Split}(e) = \{v_i \mid 1 \leq i < |V| \wedge e \in \text{Split}_i(\mathbf{v}, H)\}$, then $\text{Split}(e)$ is clearly a connected segment of \mathbf{v} . Second, note that for every e with $v \in e$, then either $v \in \text{Split}(e)$ or v and the connected subpath $\text{Split}(e)$ are adjacent (in the case where v is the last vertex of e in the enumeration). Now, by definition T_v is the union of the $b_{v'}$ for $v' \in \text{Split}(e)$ with $v \in e$ and of b_i , so it is a union of connected subpaths which all contain b_i or are adjacent to it: this establishes that T_v is a connected subpath, which shows in turn that (T, λ) is a path decomposition, concluding the proof. \blacktriangleleft

For completeness with the preceding result, we note that the following also holds, although we do not use it in the proof of Theorem 5.3:

► **Lemma D.1.** *For any hypergraph $H = (V, E)$, we have $\text{psw}(H) \leq \text{degree}(H) \times (\text{pw}(H) + 1)$.*

Proof. Let $P = b_1 - \dots - b_m$ be a path decomposition of H of width $\text{pw}(H)$. For $1 \leq i \leq m$ we define $\text{first}(b_i)$ to be the set of all $v \in b_i$ such that b_i is the first bag containing v (this set can be empty). Let \mathbf{v}_i be any ordering on $\text{first}(b_i)$. Consider the ordering $\mathbf{v} := \mathbf{v}_1 \dots \mathbf{v}_m \mathbf{v}'$, where \mathbf{v}' is any ordering of the remaining vertices of H (i.e., those that do not appear in P because they are not present in any hyperedge). Let $n = |\mathbf{v}| = |H|$. We claim that for any $1 \leq i \leq m$, we have $|\text{Split}_i(\mathbf{v}, H)| \leq \text{degree}(H) \times (\text{pw}(H) + 1)$, which clearly implies that $\text{psw}(H) \leq \text{degree}(H) \times (\text{pw}(H) + 1)$. This is clear for v_i in \mathbf{v}' , since then $|\text{Split}_i(\mathbf{v}, H)| = 0$. Now suppose $v_i \in \mathbf{v}_j$ for some $1 \leq i \leq m$. Let $e \in \text{Split}_i(\mathbf{v}, H)$. We will show that $\exists v' \in b_j$ such that $v' \in e$, which will imply that $|\text{Split}_i(\mathbf{v}, H)| \leq \text{degree}(H) \times (\text{pw}(H) + 1)$. Assume by way of contradiction that there is no such v' . We know that $e \in \text{Split}_i(\mathbf{v}, H)$, hence there exist v^-, v^+ with $\{v^-, v^+\} \subseteq e$ and $v^- \in \mathbf{v}_{j^-}$ for some $j^- < j$ and $v^+ \in \mathbf{v}_{j^+}$ for some $j^+ > j$. But, as P is a path decomposition of H and $\{v^-, v^+\} \subseteq e$, v^- and v^+ must appear together in a bag! Now, as b_{j^+} is the first bag in which v^+ appears, it must be the case that $v^- \in b_{j^+}$, and therefore $v^- \in b_j$ (otherwise the connectedness property would be violated), which leads to a contradiction and concludes the proof. \blacktriangleleft

► **Lemma 5.10.** *For any monotone DNF φ , letting H be its hypergraph, any independent set of the exclusion graph G_H is a dncpi-set of φ .*

Proof. The vertices of G_H are clauses of φ by construction. Now, the elements of an independent set S are pairwise disjoint clauses, because whenever two clauses e and e' intersect, then taking $e'' := e$, we have that e'' intersects both e and e' , so there is an edge between e and e' in the exclusion graph, so e and e' cannot both occur in an independent set. Now, to show why S is non-covering, assume by contradiction that there exists a clause e'' of φ which is not in S and such that $e'' \subseteq \bigsqcup S$. Remember that φ has been minimized, so e'' cannot be a strict subset of a single clause of S , and it cannot be a clause of S by hypothesis. Hence, there must be two clauses $e \neq e'$ in S such that e'' intersects both e and e' . Thus, e''

witnesses that there is an edge between e and e' in the exclusion graph, so they cannot be both part of S , a contradiction. This concludes the proof. ◀

► **Lemma 5.11.** *Let $G = (V, E)$ be a graph and let $V' \subseteq V$. Then G has an independent set $S \subseteq V'$ of size at least $\left\lfloor \frac{|V'|}{\text{degree}(G)+1} \right\rfloor$.*

Proof. We construct the independent S set with the following trivial algorithm: start with $S := \emptyset$ and, while V' is non-empty, pick an arbitrary vertex v in V' , add it to S , and remove v and all its neighbors from G . It is clear that this algorithm terminates and add the prescribed number of vertices to S , so all that remains is to show that S is an independent set at the end of the algorithm. This is initially true for $S = \emptyset$; let us show that it is preserved throughout the algorithm. Assume by way of contradiction that, at a stage of the algorithm, we add a vertex v to S and that it stops being an independent set. This means that S contains a neighbor v' of v which must have been added earlier; but when we added v' to S we have removed all its neighbors from G , so we have removed v and we cannot add it later, a contradiction. Hence, the algorithm is correct and the claim is shown. ◀

► **Lemma 5.12.** *Let H be a hypergraph. Then $\text{degree}(G_H) \leq (\text{arity}(H) \times \text{degree}(H))^2 - 1$.*

Proof. Any edge e of H contains $\leq \text{arity}(H)$ vertices, each of which occurs in $\leq \text{degree}(H) - 1$ edges that are different from e , so any edge e of H intersects at most $n := \text{arity}(H) \times (\text{degree}(H) - 1)$ edges different from e . Hence, the degree of G_H is at most $n + n^2$ (counting the edges that intersect e or those at distance 2 from e). Now, we have $n + n^2 = n(n + 1)$, and as $\text{degree}(H) \geq 1$ and $\text{arity}(H) \geq 1$ (because we assume that hypergraphs contain at least one non-empty edge), the degree of G_H is $< \text{arity}(H) \times \text{degree}(H) \times (1 + \text{arity}(H) \times (\text{degree}(H) - 1))$, i.e., it is indeed $< (\text{arity}(H) \times \text{degree}(H))^2$, which concludes. ◀

E Proofs for Section 6 (Lower Bounds on d-SDNNFs)

E.1 Proof of Auxiliary Lemmas

► **Lemma 6.4.** *For any hypergraph $H = (V, E)$, we have $\text{tw}(H) \leq 3 \times \text{arity}(H) \times \text{tsw}(H)$.*

Proof. Let $H = (V, E)$ be a hypergraph, and T a v-tree over V witnessing that H has treesplitwidth $\text{tsw}(H)$. We will construct a tree decomposition T' of H of width $\leq 3 \times \text{arity}(H) \times \text{tsw}(H)$. The skeleton of T' is the same as that of T . Now, for each node $n \in T$, we call b_n the corresponding bag of T' , and we define the labeling $\lambda(b_n)$ of b_n .

If n is an internal node of T with children n_l, n_r (recall that v-trees are assumed to be binary), then we define $\lambda(b_n) := \bigcup \text{Split}_n(T, H) \cup \bigcup \text{Split}_{n_l}(T, H) \cup \bigcup \text{Split}_{n_r}(T, H)$, and if n is a variable $v \in V$ (i.e., n is a leaf of T) then $\lambda(b_n) := \{v\}$. It is clear that the width of P is $\leq \max(3 \times \text{arity}(H) \times \text{tsw}(H), 1) - 1 \leq 3 \times \text{arity}(H) \times \text{tsw}(H)$.

The occurrence condition is verified: let e be an edge of H . If e is a singleton edge $\{v\}$ then it is included in b_v . If $|e| \geq 2$ then there must exist a node $n \in T$ such that $e \in \text{Split}_n(T, H)$. If n is an internal node of T then $e \subseteq \bigcup \text{Split}_n(T, H) \subseteq b_n$, and if n is a leaf node of T then it must have a parent p (since e is split), and $e \subseteq \bigcup \text{Split}_n(T, H) \subseteq b_p$.

Connectedness is proved in the same way as in the proof of Lemma 5.6: for a given vertex $v \in V$, the nodes of T where each edge e containing v is split is a connected subtree

of T without its root node: more precisely, they are all the ancestors of a leaf in e strictly lower than their the least common ancestor. Adding the missing root to each such subtree and unioning them all results in the subtree of all ancestors of a vertex adjacent to v (v itself included) up to their least common ancestor a . Consequently, the set of nodes of T' containing v is a connected subtree of T' , rooted in b_a . ◀

► **Lemma 6.5.** *Let φ be a monotone DNF, H its hypergraph, and T be a v -tree over its variables. Then there is a dncpi-set S of φ shattered by T such that $|S| \geq \left\lfloor \frac{\text{tsw}(H)}{(\text{arity}(H) \times \text{degree}(H))^2} \right\rfloor$.*

Proof. Let φ be a monotone DNF, $H = (V, E)$ its hypergraph, and T a v -tree over V . By definition of treesplitwidth, there exists $n \in V$ such that, letting $E' := \text{Split}_n(T, H)$, we have $|E'| \geq \text{tsw}(H)$. Now, by Lemma 5.11, G_H has an independent set $S \subseteq E'$ of size at least $\left\lfloor \frac{|E'|}{\text{degree}(G_H)+1} \right\rfloor$ which is $\geq \left\lfloor \frac{\text{tsw}(H)}{(\text{arity}(H) \times \text{degree}(H))^2} \right\rfloor$ by Lemma 5.12. Hence, S is a dncpi-set by Lemma 5.10, has the desired size, and is shattered since $S \subseteq E'$. ◀

► **Lemma 6.6.** *Let φ be a monotone DNF on variables V and $n \in \mathbb{N}$. Assume that, for every v -tree T over V , there is some dncpi-set S of φ with $|S| \geq n$, such that T shatters S . Then any d -SDNNF for φ has size $\geq 2^n - 1$.*

► **Lemma 6.7.** *Let φ be a DNF with variables V , and let $S = \{D_1, \dots, D_n\}$ be a dncpi-set of φ where every clause has size ≥ 2 . Pick two variables $x_i \neq y_i$ in D_i for each $1 \leq i \leq n$, and let $V' := \{x_1, y_1, \dots, x_n, y_n\}$. Then there is a partial valuation ν of V with domain $V \setminus V'$ such that $\nu(\varphi) = \text{SINT}_n$.*

Proof. Define the following partial valuation $\nu : V \setminus V' \rightarrow \{0, 1\}$ that maps all the variables of $\bigcup_{i=1, \dots, n} (D_i \setminus \{x_i, y_i\})$ to 1 and all the other variables of $V \setminus \bigcup S$ to 0. Let us show that for a clause $D \in \varphi \setminus S$ we have $\nu(D) = 0$. Otherwise, as all the variables that ν maps to 1 are in $\bigcup_{i=1, \dots, n} X_i \cup Y_i$, we should have $D \subseteq \bigcup S$, but because S is a dncpi-set we should have $D \in S$ which is a contradiction. Now, ν maps all the variables of $D_i \setminus \{x_i, y_i\}$ to 1, hence $\nu(\varphi)$ indeed captures SINT_n . ◀

Proof of Lemma 6.6. Let C be a d -SDNNF structured by a v -tree T that captures φ . Consider the dncpi-set $S = \{D_1, \dots, D_m\}$ of size $\geq n$ of φ that is shattered by T (note that this implies in particular that every clause contains at least two variables). Consider the node u of T which witnesses this. We can write each clause D_i of S as $X_i \sqcup Y_i$, where X_i is $D_i \cap \text{Leaves}(T \setminus T_u)$ and Y_i is $D_i \cap \text{Leaves}(T_u)$. Then according to Lemma 6.7, there exists a valuation ν of the variables of φ with domain $V \setminus \{x_1, y_1, \dots, x_m, y_m\}$, where $x_i \in X_i$ and $y_i \in Y_i$ for $1 \leq i \leq m$, such that $\nu(\varphi)$ captures the Boolean function SINT_m , hence we know that $\nu(C)$ also captures SINT_m . But by Proposition 6.8, we have $|\nu(C)| \geq 2^m - 1 \geq 2^n - 1$, hence $|C| \geq 2^n - 1$. ◀

E.2 From DNFs to CNFs

We now argue that Theorem 6.1 also holds for monotone CNFs. Note that we cannot use a dualization argument as we did in the previous section, as we cannot easily construct the negation of a d-SDNNF as a similar d-SDNNF. Observe that Definition 5.7 and Lemma 6.5 can also apply to monotone CNFs as these only use the hypergraph corresponding to the formula, not the semantics of the formula. Hence, in order to apply the same arguments as in the DNF case and prove an analogue of Lemma 6.6 in the CNF case, the only difference is that we would need to consider the function $f_n := (x_1 \vee y_1) \wedge \dots \wedge (x_n \vee y_n)$ and obtain analogues of Lemma 6.7 and Proposition 6.8 for that function. This is clear for Lemma 6.7, so we only need to check that the analogue of Proposition 6.8 holds. To understand why, we need to go deeper into the proof from [12]. They paraphrase a result of Pipatsrisawat and Darwiche [36] in the following way:

► **Theorem E.1 ([12, Theorem 13] and [36, Theorem 3]).** *Let C be a d -SDNNF on variables V capturing function f structured by a v -tree T . For every node $n \in T$, f has a disjoint rectangle cover of size $\leq |C|$ with partition $(V \cap \text{Leaves}(T_n), V \cap \text{Leaves}(T \setminus T_n))$.*

where a *disjoint rectangle cover* of a Boolean function $f : X \sqcup Y \rightarrow \{0, 1\}$ with partition (X, Y) is a deterministic disjunction $\bigvee_{i_1, \dots, i_m} (g_i(X) \wedge h_i(Y))$ equivalent to f such that g_i (resp., h_i) is a Boolean function from X (resp., Y) to $\{0, 1\}$, and m is its *size*. This notion is a standard tool for showing lower bounds in communication complexity. Therefore, we are interested in the smallest size of a disjoint rectangle cover for the function $f_n : X \sqcup Y \mapsto (x_1 \vee y_1) \wedge \dots \wedge (x_n \vee y_n)$ under partition (X, Y) . But it is known from communication complexity that any disjoint rectangle cover for the function *set disjunction* $\text{SDISJ}_n : X \sqcup Y \mapsto (\neg x_1 \vee \neg y_1) \wedge \dots \wedge (\neg x_n \vee \neg y_n)$ has size $\geq 2^n$ (see [33] and [12, Theorem 15 and Proposition 1]). Moreover, it is easy to see that we can turn any disjoint rectangle cover of size m for f_n with partition (X, Y) into a disjoint rectangle cover for SDISJ_n of the same size and under the same partition, which implies that any such cover for f_n must be of size at least $\geq 2^n$ and concludes the proof. Indeed, let $\bigvee_{i_1, \dots, i_m} (g_i(X) \wedge h_i(Y))$ be a disjoint rectangle cover for f_n with partition (X, Y) . When ν is a Boolean valuation from S to $\{0, 1\}$, let us write $\bar{\nu}$ for the Boolean valuation $S \rightarrow \{0, 1\}$ defined by $\bar{\nu}(s) = 1 - \nu(s)$ for $s \in S$. We then define \bar{g}_i for $1 \leq i \leq n$ (resp., \bar{h}_i) to be the Boolean function from X (resp., Y) to $\{0, 1\}$ defined by $\bar{g}_i(\nu) = g_i(\bar{\nu})$ for all valuations $\nu : X \rightarrow \{0, 1\}$ (resp., $\bar{h}_i(\nu) = h_i(\bar{\nu})$). One can then check that $\bigvee_{i_1, \dots, i_m} (\bar{g}_i(X) \wedge \bar{h}_i(Y))$ is a disjoint rectangle cover for SDISJ_n of size m with partition (X, Y) .

F Proofs for Section 7 (Application to Query Lineages)

► **Theorem 7.1.** *There is a constant $d \in \mathbb{N}$ such that the following is true. Let σ be an arity-2 signature, and Q a connected UCQ^\neq which is intricate on σ . For any instance I on σ , any d -SDNNF representing the lineage of Q on I has size $2^{\Omega(\text{tw}(I)^{1/d})}$.*

In this section, we prove Theorem 7.1. We will use the restatement of the main result of [20] given in [7] (where a *degree-3* graph is one where the maximal degree is 3):

► **Lemma F.1 ([20], rephrased as [7, Lemma 4.4]).** *There is $c \in \mathbb{N}$ such that, for any degree-3 planar graph H , for any graph G of treewidth $\geq |V(H)|^c$, H is a topological minor of G .*

We set $d := 2c$, for the c of this lemma. Fix the arity-2 signature σ and the intricate query Q . As σ is nonempty, the tautological and vacuous UCQ^\neq queries are not intricate, so

we can assume that Q is not trivial in this sense. We denote by $|Q|$ the number of atoms of Q .

We now define the class of subgraphs that we wish to extract. Recall that the *girth* of an undirected graph is defined as the length of the shortest simple cycle in the graph (or ∞ if the graph is acyclic). Let us define an infinite family $\mathcal{S} = S_2, \dots, S_n, \dots$ of graphs to extract, such that, for each $i \in \mathbb{N}$, the graph S_i has:

1. maximal degree 3;
2. treewidth i ;
3. $\leq \alpha \times i^2$ vertices for some constant $\alpha \geq 1$ depending only on Q ;
4. no vertex of degree 1;
5. $\text{girth} > 2|Q| + 2$;

We can define each S_i by starting, for instance, with a wall graph [25], to satisfy the first three conditions (for some fixed α). We then iteratively remove all vertices of degree 1, which clearly does not impact treewidth. Indeed, treewidth cannot increase when we do this, the graph cannot become empty (because its initial treewidth is ≥ 2 , so it has a cycle, which will never be removed), and treewidth cannot decrease either. Specifically, if we consider a graph G and the result G' of removing one vertex v of degree 1 in G , given a tree decomposition T' of G' , we can construct a tree decomposition T of G by adding one bag with v and its one incident vertex w , and connecting it to a bag containing w in T' (if one exists; we connect it arbitrarily otherwise); the result is clearly a tree decomposition of G , and the width is unchanged because G' is non-empty so the maximal bag size in T' is ≥ 2 . This satisfies requirement 4 and does not break requirements 1–3. Last, we subdivide each edge into a path of length $2|Q| + 3$ to ensure that the girth condition is respected: this satisfies requirement 5, does not affect requirements 1–2 or 4, and requirement 3 is still satisfied up to multiplying α by $3 \times (2|Q| + 2) + 1$ (each path replacing an edge introduces $(2|Q| + 2)$ new vertices, and since the graph is degree-3, an upper bound on the number of edges is three times the number of vertices).

We now make explicit the function hidden in the Ω -notation in the exponent of the bound that we wish to show. This function will only depend on Q . Define the increasing function $f : k \mapsto \frac{1}{\alpha} k^{1/d}$, and let $k_0 \in \mathbb{N}$ be the smallest value of k such that $f(k) \geq 2$. We will show that the size of a d-SDNNF for an input instance I is $\geq 2^{\beta f(\text{tw}(I))}$ when $\text{tw}(I)$ is large enough, for some constant $\beta > 0$ to be defined later, depending only on Q . This means indeed that it is a $\Omega(\text{tw}(I)^{1/d})$. We assume $\text{tw}(I) \geq k_0$ (and thus $f(k) \geq 2$) in what follows.

Let I be the input instance on σ , let G be the Gaifman graph of I , and let $k := \text{tw}(I) = \text{tw}(G)$. Let $k' := f(k)$, and consider $S_{k'}$, which is well-defined because k' is an integer which is ≥ 2 . We know that the number $n_{k'}$ of vertices of $S_{k'}$ is such that $n_{k'} \leq \alpha k'^2$, i.e., $n_{k'} \leq k^{1/c}$, so the treewidth k of G is $\geq n_{k'}^c$. Hence, we know by Lemma F.1 that $S_{k'}$ is a topological minor of G . Let G' be the subgraph of G corresponding to this topological minor: it is a subgraph of G , and a subdivision of $S_{k'}$.

We will extract a corresponding subinstance I' of I whose Gaifman graph is G' . For simplicity, we will ensure that I' is *Gaifman-tight*. An instance I_0 is *Gaifman-tight* if two conditions hold: first, letting G_0 be the Gaifman graph of I_0 , for each edge $\{a, b\}$ of G_0 , there is exactly one fact of I_0 containing a and b (hence, of the form $R(a, b)$ or $R(b, a)$, with $a \neq b$); second, every fact of I_0 is a binary fact with two distinct elements (of the form $R(c, d)$ with $c \neq d$). Intuitively, an instance is Gaifman-tight if it is exactly obtained from its Gaifman graph by choosing one relation name and orientation for each edge of the Gaifman graph.

We define a Gaifman-tight subinstance I' of I with Gaifman graph G' by keeping, for every edge $\{a, b\}$ of G' , exactly one binary fact of I containing the two elements a and b

(which must exist by definition of the Gaifman graph). By construction, the Gaifman graph of I' is then G' . Hence, we know the following about the subinstance I' of I and its Gaifman graph G' (the numbering of this list follows the list of conditions on \mathcal{S}):

1. For every element a of I' , there are at most 3 facts where a occurs (because G' has maximal degree 3).
2. The treewidth of I' is k' .
3. (N/A: There is no analogue of the requirement 3 imposed on \mathcal{S})
4. There are no vertices of degree 1 in G' .
5. The girth of G' is $> 2|Q| + 2$ (because as a subdivision of $S_{k'}$ its girth is at least that of $S_{k'}$).
6. I' is Gaifman-tight.

We will now construct a DNF representation of $\varphi(Q, I')$. Remember that Q is a UCQ $^\neq$, so it is monotone, hence we can construct $\varphi(Q, I')$ to be a monotone DNF. As Q is not trivial, $\varphi(Q, I')$ will contain at least one nonempty clause. Further, the DNF can be computed as a minimized DNF by taking the disjunction of conjunctions that stand for each *minimal match* of Q in I' . Specifically, a *minimal match* of Q in I' is a subinstance M of I' such that $M \models Q$ and there is no $M' \subsetneq M$ such that $M' \models Q$. The following is then easy to see (and this monotone DNF representation is clearly unique):

$$\varphi(Q, I') := \bigvee_{\substack{M \text{ minimal} \\ \text{match of } Q \\ \text{in } I'}} \bigwedge_{F \in M} F$$

Let H be the hypergraph of this DNF. To be able to usefully apply Theorem 6.1, we must show that the arity and degree of H are constant, and that $\text{tw}(H)$ is $\Omega(\text{tw}(I'))$. We first show the first claim. The arity of H is clearly bounded from above by the size of a minimal match of Q in I' , whose size is clearly bounded from above by $|Q|$, which is constant. As for the degree of H , as Q is a connected query, any minimal match of Q on I' involving some fact F must be contained in the subinstance of I' induced by the ball of radius $|Q|$ centered around the elements of F in G' : as the degree of G' is at most 3, this ball has constant size, so, as σ is fixed, F can only occur in constantly many different matches, and the degree is constant. We now show that $\text{tw}(H)$ is $\Omega(\text{tw}(I'))$: we show this in the following lemma, which captures the essence of intricate queries (namely: under some conditions, their lineage never has lower treewidth than the input instance):

► **Lemma F.2.** *Let σ be an arity-2 signature, let Q be a connected UCQ $^\neq$ which is intricate for σ , and let I' be a Gaifman-tight instance on σ whose Gaifman graph has no degree-1 vertex and has girth $> 2|Q| + 2$. Then, letting H be the hypergraph of the monotone DNF representing $\varphi(Q, I')$, we have $\text{tw}(H) \geq \lfloor \frac{\text{tw}(I')}{2} \rfloor$.*

Let us conclude the proof of Theorem 7.1 using Lemma F.2, and show Lemma F.2 afterwards. As the arity and degree of H are bounded by constants, by Theorem 6.1, we know that any d-SDNNF for $\varphi(Q, I')$ has size $\geq 2^{\beta' \text{tw}(H)}$ for some constant $\beta' > 0$ (depending only on the arity and degree bounds on I' given above, which depend only on Q), which by Lemma F.2 is $\geq 2^{\beta \text{tw}(I')}$ for a different constant $\beta > 0$ and $\text{tw}(I')$ large enough. By definition of $S_{k'}$, we obtain the lower bound of $2^{\beta f(k)}$ for k large enough. Now, to conclude, we must show that this lower bound also applies to any d-SDNNF for $\varphi(Q, I)$. But it is clear that, from any d-SDNNF C for $\varphi(Q, I)$, we can obtain a d-SDNNF C' for $\varphi(Q, I')$ which is no larger than C (structured by a v-tree obtained from that of C), simply by evaluating to 0 all

inputs corresponding to facts of $I \setminus I'$. Hence, the lower bound also applies to a d-SDNNF for $\varphi(Q, I)$, establishing the result of Theorem 7.1.

All that remains is to show Lemma F.2. Let us fix the graph signature σ , the connected UCQ $^\neq$ Q which is intricate for σ , and the instance I' on σ satisfying the conditions. We say that two different facts $R(a, b)$ and $S(c, d)$ of I' *touch* if they share an element, formally, $|\{a, b\} \cap \{c, d\}| = 1$: as I' is Gaifman-tight, remember that we must have $a \neq b$, $c \neq d$, and $\{a, b\} \neq \{c, d\}$. The key for Lemma F.2 is then captured in the following auxiliary claim:

► **Claim F.3.** *Let F and F' be two facts of I' that touch. Then there is a minimal match M of Q such that $\{F, F'\} \subseteq M$.*

Proof. Let G be the Gaifman graph of I' . Consider the two edges e and e' standing for F and F' in G : these edges are incident in G , so we write without loss of generality $e = \{u, v\}$ and $e' = \{v, w\}$. Fix $n := |Q|$. Define a path $\pi = uu_1 \dots u_n$ in G of $|Q|$ edges by exploring G from u : initially we are at u and call v the *predecessor vertex*, and whenever we reach some vertex x , we visit a neighbor of x which is different from the predecessor of x , and set x to be the new predecessor. Such a path exists, because this exploration can only get stuck on a vertex of degree 1 (i.e., a vertex that we cannot exit except by going back on its predecessor), and this cannot happen by our assumption that G has no vertex of degree 1. We define a path $\pi' = ww_1 \dots w_n$ in G of $|Q|$ edges by exploring from w with predecessor v in the same way. Now, we consider the path ρ obtained by concatenating the reverse of π , e , e' , and π' , namely: $\rho: u_n, \dots, u_1, u, v, w, w_1, \dots, w_n$. We claim that this path is a simple path, i.e., no two vertices in the path are the same. Indeed, by definition, no two consecutive vertices can be the same in π , in π' , or in u, v, w . Further, two vertices separated by one single vertex cannot be the same: this is the case in π and π' because we do not go back to the predecessor vertex in the exploration, and initially we do not go back on v : and for u and w we know that they are different because F and F' touch and I' is Gaifman-tight. Last, two vertices further apart in ρ cannot be equal, because otherwise the path ρ would contain a simple cycle of G , which would contradict the hypothesis on the girth of G .

Hence, ρ is a simple path of the Gaifman graph G of I' . Consider the sequence of facts L of I' that witness the existence of each edge of ρ , which is unique because I' is Gaifman-tight; in particular we choose F and F' as witnesses for e and e' . Recall now the definition of a line instance, and of a UCQ $^\neq$ Q being intricate (Definitions 8.4 and 8.5 of [7]). The sequence of facts L is a line instance, with $|L| = 2|Q| + 2$, and the two facts incident to the middle element are F and F' . Hence, the definition of intricate queries ensures that there is a minimal match M of Q on L that includes both F and F' . As L is a subinstance of I' , the match M is still a match of Q on I' , and it is still minimal, because any match $M' \subseteq M$ would also satisfy $M' \subseteq L$ and contradict the minimality of M on L . Hence, M is the desired minimal match, which concludes the proof. ◀

We are now ready to prove Lemma F.2 from Claim F.3, which is the only missing part of the proof of Theorem 7.1:

Proof of Lemma F.2. Fix σ , Q , and I' , consider the monotone DNF representation of $\varphi(Q, I')$ and its hypergraph H . To show the desired inequality, it suffices to show that, from a tree decomposition T of H where the maximal bag size is k , we can construct a tree decomposition T' of I' whose maximal bag size is no greater than $2k$. Let T be a tree decomposition of H , and construct T' to have same skeleton as T . We define the labeling $\lambda(b')$ of every bag b' of T' to be the set of vertices occurring in the label $\lambda(b)$ of the corresponding

bag b of T (which consists of variables of $\varphi(Q, I')$, hence of facts of I'): this clearly satisfies the size requirements. We must now show that T' is a tree decomposition of I' .

To show the occurrence requirement, we must show that for every fact F of I' , there is a bag of T' containing its two elements. To show this, it suffices to show that there is a bag of T that contains F (as a vertex of H). As the Gaifman graph of I' has no vertex of degree 1, there must be a fact F' of I' that touches F , and we can conclude using a consequence of Claim F.3: F must occur in a minimal match of Q on I' (together with F' , but we do not use this), hence it occurs in a clause of $\varphi(Q, I')$, and the occurrence requirement on T ensures that F occurs in a bag of T .

To show the connectedness requirement, pick an element a of I' . Its occurrences in T' are the union of the occurrences in T of the facts that contain a , which are connected subtrees of T by the connectedness requirement of T . Hence, it suffices to show that their union is connected. To do this, let us show that for any two facts F and F' that contain a , then the subtrees T_F and $T_{F'}$ of their occurrences in T necessarily intersect. This is trivial if $F = F'$; now, if $F \neq F'$, since I' is Gaifman-tight, the facts F and F' must touch in I' (they cannot share exactly the same elements). Now, we use Claim F.3 to conclude that F and F' occur together in a minimal match M of Q on I' . Hence, there is a clause of $\varphi(Q, I')$ which contains both F and F' , which ensures that F and F' occur together in a bag of T , so T_F and $T_{F'}$ intersect. This shows that T' is indeed a tree decomposition of I' , which concludes the proof. ◀

This concludes the proof of Theorem 7.1.