



Counting Minimal Transversals of β -Acyclic Hypergraphs

Benjamin Bergougnoux, Florent Capelli, Mamadou Moustapha Kanté

► To cite this version:

Benjamin Bergougnoux, Florent Capelli, Mamadou Moustapha Kanté. Counting Minimal Transversals of β -Acyclic Hypergraphs . 2018. hal-01610959v2

HAL Id: hal-01610959

<https://hal.science/hal-01610959v2>

Preprint submitted on 15 Aug 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Counting Minimal Transversals of β -Acyclic Hypergraphs^{*}

Benjamin Bergougnoux^a, Florent Capelli^b, Mamadou Moustapha Kanté^a

^aUniversité Clermont Auvergne, LIMOS, CNRS, France

^bUniversité de Lille, CRISTAL, CNRS, INRIA, France

Abstract

We prove that one can count in polynomial time the number of minimal transversals of β -acyclic hypergraphs. In consequence, we can count in polynomial time the number of minimal dominating sets of strongly chordal graphs, continuing the line of research initiated in [M.M. Kanté and T. Uno, Counting Minimal Dominating Sets, TAMC'17].

Keywords: Counting problem, Minimal transversals, Dominating sets, β -Acyclic Hypergraphs, Strongly chordal graphs

Introduction

A *hypergraph* is a collection of subsets - called *hyperedges* - of a finite ground set, and a *transversal* is a subset of the ground set that intersects every hyperedge. In this paper, we consider the problem of counting the (inclusion-wise) minimal transversals of β -acyclic hypergraphs. Counting problems are usually harder than decision problems as for instance computing a (minimal) transversal of a hypergraph can be done in polynomial time while counting the number of (minimal) transversals is $\#P$ -complete [1]. Informally, $\#P$ denotes the set of functions corresponding to the number of accepting paths of a non-deterministic Turing machine, and $FP \subseteq \#P$, the set of functions computable in polynomial time. Under the assumption $\#P \neq FP$, one may be interested in classifying counting problems between those that are easy to compute, *i.e.*, belong to FP , and those that are hard, *i.e.*, are $\#P$ -hard [2], or even hard to approximate [3].

The HYPERGRAPH DUALISATION problem – a fifty years open problem – asks for the enumeration of all (inclusion-wise) minimal transversals of a given hypergraph. It is not yet known whether this problem can be solved by an algorithm that runs in time polynomial in the size of the output despite the fact

^{*}The authors are supported by the French Agency for Research under the GraphEN (ANR-15-CE40-0009) and AGGREG (ANR-14-CE25-0017) projects.

Email addresses: benjamin.bergougnoux@uca.fr (Benjamin Bergougnoux), florent.capelli@univ-lille.fr (Florent Capelli), mamadou.kante@uca.fr (Mamadou Moustapha Kanté)

that it has been extensively studied since it has several applications in many areas such as graph theory, artificial intelligence, datamining, model-checking, network modeling, or databases (see the survey [4]). For all these applications, the minimality of the transversals is needed. Thus, even if it is not hard to see that listing all transversals of a hypergraph can be done with a polynomial delay between two solutions, there may be an exponential gap between the number of minimal transversals and the number of transversals (see for instance the non-trivial class of affine formulas), which makes efficient enumeration of minimal transversals much harder.

The problem of computing the number of minimal transversals of a hypergraph is closely related to the HYPERGRAPH DUALISATION problem and has also many applications in several areas, see for instance the description given in [5] in the case of model checking. This problem has also applications in graph theory as it is closely related to the problem of counting the *minimal dominating sets* of a graph. It turns out that these counting problems are $\#P$ -complete in general, that is, it is very unlikely that it can be solved in polynomial time on every input. However, there is a rich literature on solving this problem in polynomial time by restricting the input hypergraph to a specific class, see for example [6, 7, 8].

Our contributions and approach. In this paper, we follow this line of research by exhibiting a new tractable class of hypergraphs for this problem, namely, the class of β -acyclic hypergraphs. The HYPERGRAPH DUALISATION problem was already known to be tractable for this class [9]. Moreover, it has been proved in [10] that the number of transversals of a β -acyclic hypergraph can be computed in polynomial time but the complexity of computing the number of minimal transversals was still unknown. More precisely, we prove the following theorem:

Main Theorem. *The number of minimal transversals of a β -acyclic hypergraph can be computed in polynomial time.*

A direct consequence of our result is the following corollary concerning the counting of *minimal dominating sets* in a subclass of chordal graphs, called *strongly chordal* graphs.

Main Corollary. *The number of minimal dominating sets of a strongly chordal graph is computable in polynomial time.*

Besides the polynomial time algorithm, the main contribution of this article is the modification of the framework considered in [11] in order to count minimal models. The techniques used in [7, 8] are based on structural restrictions and as shown in [11] cannot work for β -acyclic hypergraphs. Instead, Capelli showed in [11] how to construct, from the elimination ordering of a β -acyclic hypergraph associated to a boolean formula, a circuit whose satisfying assignments correspond to the models of the boolean formula. Such circuits are known as *decision Decomposable Negation Normal Form* in knowledge compilation. While the technique allows to count the models of non-monotone formulas, it cannot be used to count the minimal models. Indeed, the branchings of the constructed

circuit do not allow to control the minimality. We overcome this difficulty by introducing the notion of *blocked transversals*, which correspond roughly to the minimal transversals of a sub-hypergraph that are transversals of the whole hypergraph. We then show that blocked transversals can be used to control the minimality in the construction of the circuit. However, this control is only possible in the case of monotone Boolean formulas, corresponding to counting the minimal transversals of β -acyclic hypergraphs.

Because of technical definitions, we postpone the details of the algorithm in Section 3.2. The paper is organised as follows. Notations and some technical definitions are given in Section 1, while *blocked transversals* and intermediate lemmas are given in Section 2. The decomposition of β -acyclic hypergraphs proposed in [11] is refined in Section 3.1 to take into account *blocked transversals*. Finally, the algorithm is given in Section 3.2 and we talk about its consequences for the counting of minimal dominating sets in Section 4.

1. Definitions and notations

The power set of a set V is denoted by 2^V , and its cardinal is denoted by $\#V$. For two sets A and B , we let $A \setminus B$ denote the set $\{x \in A \mid x \notin B\}$. For a ground set V and subsets $\mathcal{A}_1, \dots, \mathcal{A}_k$ of 2^V , with $k \geq 2$, we let

$$\bigotimes_{1 \leq i \leq k} \mathcal{A}_i := \begin{cases} \emptyset & \text{if } \mathcal{A}_i = \emptyset \text{ for some } 1 \leq i \leq k, \\ \{T_1 \cup \dots \cup T_k \mid \forall i \leq k, T_i \in \mathcal{A}_i\} & \text{otherwise.} \end{cases}$$

If $k = 2$, we write $\mathcal{A}_1 \otimes \mathcal{A}_2$ as usual. For example, if $\mathcal{A}_1 = \{\{a, b\}, \{b, c\}\}$ and $\mathcal{A}_2 = \{\{d\}, \{e, f\}\}$, then $\mathcal{A}_1 \otimes \mathcal{A}_2 = \{\{a, b, d\}, \{a, b, e, f\}, \{b, c, d\}, \{b, c, e, f\}\}$.

1.1. Hypergraphs

A hypergraph \mathcal{H} is a collection of subsets of a finite ground set. The elements of \mathcal{H} are called the *hyperedges* of \mathcal{H} and the *vertex set* of \mathcal{H} is $V(\mathcal{H}) := \bigcup_{e \in \mathcal{H}} e$.

Given $S \subseteq V(\mathcal{H})$, we denote by $\mathcal{H}[S]$ the *hypergraph induced by S* , that is, $\mathcal{H}[S] := \{e \cap S \mid e \in \mathcal{H}\}$. Any subset \mathcal{H}' of \mathcal{H} is called a *sub-hypergraph* of \mathcal{H} . Observe that if there exists $e \in \mathcal{H}$ such that $e \subseteq V(\mathcal{H}) \setminus S$, then $\emptyset \in \mathcal{H}[S]$. We do not enforce hypergraphs to have non-empty edges or to be non-empty. However, a hypergraph with an empty edge may behave counter-intuitively. In the following definitions, we explicitly explain the extremal cases where $\emptyset \in \mathcal{H}$ or $\mathcal{H} = \emptyset$.

Given a hypergraph \mathcal{H} and a subset $S \subseteq V(\mathcal{H})$ of its vertex set, we denote by $\mathcal{H}(S)$ the set of hyperedges of \mathcal{H} containing at least one vertex in S , that is, $\mathcal{H}(S) := \{e \in \mathcal{H} \mid S \cap e \neq \emptyset\}$; to ease notations, we write $\mathcal{H}(x)$ instead of $\mathcal{H}(\{x\})$ for $x \in V(\mathcal{H})$. Observe that $\mathcal{H}(x)$ is the set of hyperedges containing x . For the hypergraph \mathcal{H} shown in Figure 1, $\mathcal{H}(\{d, x\})$ is the hypergraph $\{\{b, x\}, \{c, x\}, \{c, d\}\}$.

Given a hypergraph \mathcal{H} , a walk between two distinct edges e_1 and e_k is a sequence $(e_1, x_1, e_2, x_2, \dots, e_{k-1}, x_{k-1}, e_k)$ such that $x_i \in e_i \cap e_{i+1}$ for all

$1 \leq i \leq k-1$. Notice that, $(e_k, x_{k-1}, e_{k-1}, \dots, x_2, e_2, x_1, e_1)$ is also a walk between e_k and e_1 . A maximal set of edges of \mathcal{H} that are pairwise connected by a walk is called a *connected component* of \mathcal{H} . It is worth noticing that if C_1, \dots, C_k are the connected components of \mathcal{H} , then $V(C_i) \cap V(C_j) = \emptyset$ for distinct i, j in $\{1, \dots, k\}$.

A hypergraph \mathcal{H} is said *β -acyclic* if there exists an ordering x_1, \dots, x_n of $V(\mathcal{H})$ such that for each $1 \leq i \leq n$, the set $\{e \cap \{x_i, \dots, x_n\} \mid e \in \mathcal{H}, x_i \in e\}$ is linearly ordered by inclusion. Such an ordering is called a *β -elimination ordering*. It is well-known that every sub-hypergraph \mathcal{H}' of a β -acyclic hypergraph \mathcal{H} are β -acyclic as well (see for instance [12]).

1.2. Transversals

Let \mathcal{H} be a hypergraph. A transversal for \mathcal{H} is a subset $T \subseteq V(\mathcal{H})$ such that for every $e \in \mathcal{H}$, $T \cap e \neq \emptyset$. We denote by $\text{tr}(\mathcal{H})$ the set of transversals of \mathcal{H} . Observe that if $\emptyset \in \mathcal{H}$, then $\text{tr}(\mathcal{H}) = \emptyset$ as for every $T \subseteq V(\mathcal{H})$, $\emptyset \cap T = \emptyset$ so T cannot be a transversal of \mathcal{H} . Finally, observe that if $\mathcal{H} = \emptyset$, then $\text{tr}(\mathcal{H}) = \{\emptyset\}$.

A transversal T of \mathcal{H} is *minimal* if and only if for every $x \in T$, it holds that $T \setminus \{x\} \notin \text{tr}(\mathcal{H})$. A hyperedge e such that $e \cap T = \{x\}$ is said to be *private for x* w.r.t. T . When T is clear from the context, we may refer to such hyperedges as simply *privates for x* . The following fact follows directly from the definitions:

Fact 1. *T is a minimal transversal of a hypergraph \mathcal{H} if and only if T is a transversal of \mathcal{H} and each vertex $x \in T$ has a private.*

We denote by $\text{mtr}(\mathcal{H})$ the set of minimal transversals of \mathcal{H} . Again, observe that if $\mathcal{H} = \emptyset$ then $\text{mtr}(\mathcal{H}) = \{\emptyset\}$.

Figure 1 depicts a hypergraph together with its minimal transversals.




Figure 1: A hypergraph $\mathcal{H} = \{\{a, b\}, \{b, x\}, \{x, c\}, \{c, d\}\}$ having 4 minimal transversals. We have $\text{mtr}(\mathcal{H}) = \{\{a, x, d\}, \{a, x, c\}, \{b, x, d\}, \{b, c\}\}$.

It is worth noticing that since the sets $\text{tr}(\mathcal{H})$ and $\text{mtr}(\mathcal{H})$ are sets of subsets of $2^{V(\mathcal{H})}$, they may be seen as hypergraphs on $V(\mathcal{H})$. Thus, we will sometimes use the notations $\text{tr}(\mathcal{H})(S)$ (resp. $\text{mtr}(\mathcal{H})(S)$) which refer to the transversals (resp. minimal transversals) T of \mathcal{H} such that $S \cap T \neq \emptyset$.

2. Blocked transversals

Our algorithm uses a dynamic programming approach by finding a relation between the number of minimal transversals of a β -acyclic hypergraph \mathcal{H} with the number of minimal transversals of some specific subhypergraphs of \mathcal{H} . However, it is not possible to directly relate these quantities together. To illustrate this fact, let \mathcal{H} be the hypergraph depicted in Figure 1. The

minimal transversals of \mathcal{H} containing x are $\{a, x, c\}$, $\{a, x, d\}$ and $\{b, x, d\}$. Observe that removing x from these sets directly yields a minimal transversal of $\mathcal{H} \setminus \mathcal{H}(x) = \{\{a, b\}, \{c, d\}\}$. However, adding x to a minimal transversal of $\mathcal{H} \setminus \mathcal{H}(x)$ does not give necessarily a minimal transversal of \mathcal{H} . For example, we have $\{b, c\} \in \text{mtr}(\mathcal{H} \setminus \mathcal{H}(x))$, but $\{b, x, c\} \notin \text{mtr}(\mathcal{H})$.

In fact, we can show in general that $T \cup x \in \text{mtr}(\mathcal{H})$ if and only if T is a minimal transversal of $\mathcal{H} \setminus \mathcal{H}(x)$ and T is not a transversal of \mathcal{H} . Consequently, the number of minimal transversals of \mathcal{H} containing x is

$$\#\text{mtr}(\mathcal{H} \setminus \mathcal{H}(x)) - \#(\text{tr}(\mathcal{H}) \cap \text{mtr}(\mathcal{H} \setminus \mathcal{H}(x))). \quad (1)$$

We can infer from this fact a recursive formula to compute $\#\text{mtr}(\mathcal{H})$. In the general case, using this formula as it is will lead to the computation of an exponential number of terms. In this section, we will make this relation more precise by introducing the notion of *blocked transversal*. In the next section, we will show how we can use this relation to evaluate the number of minimal transversals of a β -acyclic hypergraphs with only a polynomial number of intermediate values.

Given a hypergraph \mathcal{H} , a sub-hypergraph $\mathcal{H}' \subseteq \mathcal{H}$ and $B \subseteq V(\mathcal{H})$, we define the *B-blocked transversals* of \mathcal{H}' to be the transversals T of \mathcal{H}' such that each vertex x of T has a private in $\mathcal{H}' \setminus \mathcal{H}'(B)$. In particular, if y is a vertex of B , then y cannot be in a B -blocked transversal of \mathcal{H}' . Observe also that if $y \in V(\mathcal{H}) \setminus V(\mathcal{H}')$, then y cannot be in a B -blocked transversal of \mathcal{H}' .

Fact 2. *Given a sub-hypergraph \mathcal{H}' of a hypergraph \mathcal{H} and $B \subseteq V(\mathcal{H})$, T is a B -blocked transversal of \mathcal{H}' if and only if T belongs to $\text{tr}(\mathcal{H}') \cap \text{mtr}(\mathcal{H}' \setminus \mathcal{H}'(B))$.*

Proof. If T is a transversal of \mathcal{H}' such that each vertex x of T has a private in $\mathcal{H}' \setminus \mathcal{H}'(B)$, then $T \subseteq V(\mathcal{H}' \setminus \mathcal{H}'(B))$ and is a transversal of $\mathcal{H}' \setminus \mathcal{H}'(B)$. Thus by Fact 1 T is a minimal transversal of $\mathcal{H}' \setminus \mathcal{H}'(B)$. Conversely, if T is a minimal transversal of $\mathcal{H}' \setminus \mathcal{H}'(B)$, then by Fact 1 each vertex x has a private in $\mathcal{H}' \setminus \mathcal{H}'(B)$. If in addition it belongs to the set of transversals of \mathcal{H}' , then it is a B -blocked transversal of \mathcal{H}' by definition. \square

As an example, let \mathcal{H} be the hypergraph depicted in Figure 1, $\mathcal{H}' = \mathcal{H}$ and $B = \{x\}$. The only B -blocked transversal of \mathcal{H} is $\{b, c\}$. While $\{b, d\}$ is a minimal transversal of $\mathcal{H} \setminus \mathcal{H}(x)$, it is not a B -blocked transversal as the hyperedge $\{x, c\}$ does not intersect $\{b, d\}$.

We call the set $\mathcal{H}'(B)$ the *blocked hyperedges*. Intuitively, $\mathcal{H}'(B)$ is the set of hyperedges that cannot be used as privates in a transversal of \mathcal{H}' . We denote by $\text{btr}_B(\mathcal{H}')$ the set of B -blocked transversals of \mathcal{H}' . By Fact 2, we have

$$\text{btr}_B(\mathcal{H}') := \text{tr}(\mathcal{H}') \cap \text{mtr}(\mathcal{H}' \setminus \mathcal{H}'(B)).$$

Observe that by definition, $\text{mtr}(\mathcal{H}) = \text{btr}_\emptyset(\mathcal{H})$. Moreover, if $\mathcal{H}(B) = \mathcal{H} \neq \emptyset$, then $\text{btr}_\mathcal{H}(\mathcal{H}) = \emptyset$ as $\text{mtr}(\emptyset) = \{\emptyset\}$ and $\emptyset \notin \text{tr}(\mathcal{H})$. When $B = \{x\}$, we denote $\text{btr}_B(\mathcal{H})$ by $\text{btr}_x(\mathcal{H})$.

Given $S \subseteq V(\mathcal{H})$, we denote by $\text{tr}(\mathcal{H}, S) := \{T \in \text{tr}(\mathcal{H}) \mid T \subseteq S\}$. We extend this notation to mtr and btr as well. The following summarises observations about blocked transversals.

Fact 3. *Let \mathcal{H}' be a sub-hypergraph of a hypergraph \mathcal{H} and $B, S \subseteq V(\mathcal{H})$. Then,*

1. $\text{btr}_B(\mathcal{H}') = \text{btr}_{B \cap V(\mathcal{H}')}(\mathcal{H}')$.
2. $\text{btr}_B(\mathcal{H}', S) = \text{btr}_B(\mathcal{H}', S \setminus B)$.
3. *If $x \notin V(\mathcal{H}')$, then $\text{btr}_B(\mathcal{H}', S) = \text{btr}_B(\mathcal{H}', S \setminus \{x\})$.*

Let us briefly explain how B -blocked transversals will be used in computing $\#\text{mtr}(\mathcal{H})$. One checks easily that $\#\text{mtr}(\mathcal{H}) = \#\text{mtr}(\mathcal{H}, V(\mathcal{H}) \setminus \{x\}) + \#\text{mtr}(\mathcal{H})(x)$. By Equation 1, $\#\text{mtr}(\mathcal{H})(x) = \#\text{mtr}(\mathcal{H} \setminus \mathcal{H}(x)) - \#\text{btr}_x(\mathcal{H})$. Therefore, $\#\text{mtr}(\mathcal{H}) = \#\text{btr}_\emptyset(\mathcal{H}) = \#\text{btr}_\emptyset(\mathcal{H}, V(\mathcal{H}) \setminus \{x\}) + \#\text{btr}_\emptyset(\mathcal{H} \setminus \mathcal{H}(x)) - \#\text{btr}_x(\mathcal{H})$. We will show in the next section that we can define, from the β -elimination ordering of a β -acyclic hypergraph \mathcal{H} , sub-hypergraphs $\mathcal{H}_1, \dots, \mathcal{H}_q \subseteq \mathcal{H}$ and vertices x_1, \dots, x_q such that, for all $1 \leq i \leq q$, $\#\text{btr}_{x_i}(\mathcal{H}_i)$ and $\#\text{btr}_\emptyset(\mathcal{H}_i)$ can be computed in polynomial time if $\#\text{btr}_{x_j}(\mathcal{H}_j)$ and $\#\text{btr}_\emptyset(\mathcal{H}_j)$ are known for all $j < i$. As a consequence, one can compute $\#\text{mtr}(\mathcal{H})$ by classical dynamic programming for any β -acyclic hypergraph.

The end of this section is dedicated to the proof of several crucial lemmas concerning recursive formulas for computing the number of blocked transversals, and that will be useful in our algorithm. We start by describing the blocked transversals of a hypergraph having more than one connected component.

Lemma 4. *Let \mathcal{H} be a hypergraph, $S, B \subseteq V(\mathcal{H})$ and C_1, \dots, C_k the connected components of $\mathcal{H}[S]$. For each $i \in [1, k]$, let $\mathcal{H}_i = \{e \in \mathcal{H} \mid e \cap S \in C_i\}$. We have:*

$$\text{btr}_B(\mathcal{H}, S) = \bigotimes_{i=1}^k \text{btr}_B(\mathcal{H}_i, S).$$

Proof. Assume that $\emptyset \in \mathcal{H}[S]$, it means that there exists a hyperedge $e \in \mathcal{H}$ such that $e \cap S = \emptyset$. In this case, $\text{btr}_B(\mathcal{H}, S) = \emptyset$. Moreover, there exists $i \in [1, k]$ such that $C_i = \{\emptyset\}$ and $e \in \mathcal{H}_i$. Thus, $\text{btr}_B(\mathcal{H}_i, S) = \emptyset$ and the equality holds in this case.

Assume from now that $\emptyset \notin \mathcal{H}[S]$. Let $T \in \text{btr}_B(\mathcal{H}, S)$. We show that for all $i \leq k$, $T_i = T \cap V(\mathcal{H}_i) \in \text{btr}_B(\mathcal{H}_i, S)$. Let $e \in \mathcal{H}_i$. Since $e \in \mathcal{H}$, we have $e \cap T \neq \emptyset$. As $e \in \mathcal{H}_i$, we have $e \subseteq V(\mathcal{H}_i)$. Thus $e \cap T_i \neq \emptyset$, that is, $T_i \in \text{tr}(\mathcal{H}_i, S)$. Moreover, let $y \in T_i$. By definition of T , there exists $e \in \mathcal{H} \setminus \mathcal{H}(B)$ such that e is private for y w.r.t. T . Observe that we have $T_i \subseteq S \cap V(\mathcal{H}_i) = V(C_i)$ since $T \subseteq S$. Thus, $y \in V(C_i)$ and $e \cap S \in C_i$, because C_i is a connected component. We can conclude that $e \in \mathcal{H}_i$ and e is private to y w.r.t. T_i and $\mathcal{H}_i \setminus \mathcal{H}(B) = \mathcal{H}_i \setminus \mathcal{H}_i(B)$. Thus T_i is a minimal transversal of $\mathcal{H}_i \setminus \mathcal{H}_i(B)$. That is $T_i \in \text{btr}_B(\mathcal{H}_i, S)$.

Now let $T_1 \in \text{btr}_B(\mathcal{H}_1, S), \dots, T_k \in \text{btr}_B(\mathcal{H}_k, S)$. We show that $T = \bigcup_{i=1}^k T_i \in \text{btr}_B(\mathcal{H}, S)$. Let $e \in \mathcal{H}$. As $\mathcal{H} = \bigcup_{i=1}^k \mathcal{H}_i$, there exists i such that $e \in \mathcal{H}_i$. Thus $e \cap T_i \neq \emptyset$ and thus $e \cap T \neq \emptyset$, that is, $T \in \text{tr}(\mathcal{H})$. It remains

to show that $T \in \text{mtr}(\mathcal{H} \setminus \mathcal{H}(B))$. Let $y \in T$. By definition of T , there exists i such that $y \in T_i$. Thus, there exists $e \in \mathcal{H}_i \setminus \mathcal{H}_i(B)$ that is private for y w.r.t. T_i . Moreover, since $\mathcal{H}_i(B) = \mathcal{H}_i \cap \mathcal{H}(B)$, we know that $e \notin \mathcal{H}(B)$. As C_1, \dots, C_k are the connected component of $\mathcal{H}[S]$, we have that, for every $j \neq i$, $V(C_i) \cap V(C_j) = \emptyset$. Moreover, for all $\ell \leq k$, we have $S \cap V(\mathcal{H}_\ell) = V(C_\ell)$. As $e \subseteq V(\mathcal{H}_i)$ and $T_j \subseteq S \cap V(\mathcal{H}_j)$, we have for every $j \neq i$:

$$e \cap T_j \subseteq V(\mathcal{H}_i) \cap S \cap V(\mathcal{H}_j) = V(C_i) \cap V(C_j) = \emptyset.$$

Thus, $e \cap T = e \cap T_i = \{y\}$. In other words, e is private for y w.r.t. T and $\mathcal{H} \setminus \mathcal{H}(B)$. That is $T \in \text{btr}_B(\mathcal{H})$. \square

We recall that $\text{btr}_B(\mathcal{H}, S)(x)$ is the set of B -blocked transversals T of \mathcal{H} such that $T \subseteq S$ and $x \in T$. The following lemma shows that for any B -blocked transversal $T \subseteq S$ of \mathcal{H} containing x , we have $T \setminus x$ is a B -blocked transversal of $\mathcal{H} \setminus \mathcal{H}(x)$.

Lemma 5. *Let \mathcal{H} be a hypergraph, $S, B \subseteq V(\mathcal{H})$ and $x \in S$. We have*

$$\text{btr}_B(\mathcal{H}, S)(x) \subseteq \{\{x\}\} \otimes \text{btr}_B(\mathcal{H} \setminus \mathcal{H}(x), S \setminus \{x\}).$$

Proof. Let $\mathcal{H}_1 := \mathcal{H} \setminus \mathcal{H}(x)$. Let $T \in \text{btr}_B(\mathcal{H}, S)(x)$. By definition, $x \in T$ and $T \subseteq S$, thus we only have to show that $T' = T \setminus \{x\} \in \text{btr}_B(\mathcal{H}_1)$. Let $e \in \mathcal{H} \setminus \mathcal{H}(x)$. Since T is a transversal of \mathcal{H} , there exists $y \in e \cap T$. Moreover, by definition, $x \notin e$, thus $y \in T'$, that is, T' is a transversal of $\mathcal{H} \setminus \mathcal{H}(x)$. It remains to show that T' is a minimal transversal of $\mathcal{H}_1 \setminus \mathcal{H}_1(B) = \mathcal{H} \setminus (\mathcal{H}(B) \cup \mathcal{H}(x))$. Let $y \in T'$. Since T is a minimal transversal of $\mathcal{H} \setminus \mathcal{H}(B)$, there exists $e \in \mathcal{H} \setminus \mathcal{H}(B)$ such that e is private for y w.r.t. T . Since $x \in T$, we have $x \notin e$, otherwise e would not be private for y w.r.t. T . Thus $e \in \mathcal{H} \setminus (\mathcal{H}(B) \cup \mathcal{H}(x))$, that is, e is private to y w.r.t. T' in $\mathcal{H}_1 \setminus \mathcal{H}_1(B)$. In other words, T' is a minimal transversal of $\mathcal{H}_1 \setminus \mathcal{H}_1(B)$ which concludes the proof. \square

To complete the previous lemma, we show that for each B -blocked transversal $T \subseteq S$ of $\mathcal{H} \setminus \mathcal{H}(x)$, we have $T \cup \{x\}$ is a B -blocked transversal of \mathcal{H} if and only if T is not a $(B \cup \{x\})$ -blocked transversal of $\mathcal{H} \setminus (\mathcal{H}(B) \cap \mathcal{H}(x))$.

Lemma 6. *Let \mathcal{H} be a hypergraph, $S, B \subseteq V(\mathcal{H})$ and $x \in S$. We have*

$$(\{\{x\}\} \otimes \text{btr}_B(\mathcal{H}_1, S \setminus \{x\})) \setminus \text{btr}_B(\mathcal{H}, S)(x) = \{\{x\}\} \otimes \text{btr}_{B \cup \{x\}}(\mathcal{H}_2, S \setminus \{x\})$$

where $\mathcal{H}_1 := \mathcal{H} \setminus \mathcal{H}(x)$ and $\mathcal{H}_2 := \mathcal{H} \setminus (\mathcal{H}(B) \cap \mathcal{H}(x))$.

Proof. We prove the lemma by proving first the left-to-right inclusion (Claim 1) and then the right-to-left inclusion (Claim 2). But first, notice that $\mathcal{H}_1 \setminus \mathcal{H}_1(B) = \mathcal{H}_2 \setminus \mathcal{H}_2(B \cup \{x\}) = \mathcal{H} \setminus (\mathcal{H}(B) \cup \mathcal{H}(x))$ since $\mathcal{H}(B) \cup \mathcal{H}(x) = \mathcal{H}(B \cup \{x\})$.

Claim 1. *For every $T \in (\{\{x\}\} \otimes \text{btr}_B(\mathcal{H}_1, S \setminus \{x\})) \setminus \text{btr}_B(\mathcal{H}, S)(x)$, we have $T \setminus \{x\} \in \text{btr}_{B \cup \{x\}}(\mathcal{H}_2, S \setminus \{x\})$.*

Proof. We start by proving that $T' = T \setminus \{x\}$ is in $\text{tr}(\mathcal{H}_2)$. Assume towards a contradiction that $T' \notin \text{tr}(\mathcal{H}_2)$, i.e., there exists $e \in \mathcal{H}_2$ such that $e \cap T' = \emptyset$. We prove that it implies $T \in \text{btr}_B(\mathcal{H}, S)(x)$. First, observe that $T \in \text{tr}(\mathcal{H})$, since $T' \in \text{tr}(\mathcal{H}_1) = \text{tr}(\mathcal{H} \setminus \mathcal{H}(x))$ and $T = T' \cup \{x\}$. Thus, we have $e \cap T = \{x\}$ and $e \in \mathcal{H}(x)$. As $e \in \mathcal{H}_2 = \mathcal{H} \setminus (\mathcal{H}(B) \cap \mathcal{H}(x))$, we have $e \in \mathcal{H} \setminus \mathcal{H}(B)$ and then e is a private hyperedge for x w.r.t. T and $\mathcal{H} \setminus \mathcal{H}(B)$. Furthermore, every vertex in T' has a private hyperedge w.r.t. T' and $\mathcal{H}_1 \setminus \mathcal{H}_1(B) = \mathcal{H} \setminus (\mathcal{H}(B) \cup \mathcal{H}(x))$ since $T' \in \text{mtr}(\mathcal{H}_1 \setminus \mathcal{H}_1(B))$. Thus, every vertex in T' has a private hyperedge w.r.t. T and $\mathcal{H} \setminus \mathcal{H}(B)$. As $T \in \text{tr}(\mathcal{H})$, we can conclude that $T \in \text{mtr}(\mathcal{H} \setminus \mathcal{H}(B))$. Finally, we have $T \subseteq S$ by assumption. Therefore $T \in \text{btr}_B(\mathcal{H}, S)(x)$ which is a contradiction. Thus, $T' \in \text{tr}(\mathcal{H}_2)$.

We now prove that $T' \in \text{mtr}(\mathcal{H}_2 \setminus \mathcal{H}_2(B \cup \{x\}))$, that is, we prove the minimality of T' in $\mathcal{H}_2 \setminus \mathcal{H}_2(B \cup \{x\})$. Let $y \in T'$. Since $T \in \text{btr}_B(\mathcal{H}_1)$, there exists $f \in \mathcal{H}_1 \setminus \mathcal{H}_1(B)$ such that $f \cap T = \{y\}$. Since $\mathcal{H}_1 \setminus \mathcal{H}_1(B) = \mathcal{H}_2 \setminus \mathcal{H}_2(B \cup \{x\})$, every $y \in T'$ have a private hyperedge in $\mathcal{H}_2 \setminus \mathcal{H}_2(B \cup \{x\})$, that is $T' \in \text{mtr}(\mathcal{H}_2 \setminus \mathcal{H}_2(B \cup \{x\}))$. As $T' \subseteq S \setminus \{x\}$, we can conclude that $T' \in \text{btr}_{B \cup \{x\}}(\mathcal{H}_2, S \setminus \{x\})$. \square

Claim 2. For every $T \in \{x\} \otimes \text{btr}_{B \cup \{x\}}(\mathcal{H}_2, S \setminus \{x\})$, we have $T \in \{x\} \otimes \text{btr}_B(\mathcal{H}_1, S \setminus \{x\}) \setminus \text{btr}_B(\mathcal{H}, S)(x)$.

Proof. We start by proving that $T' = T \setminus \{x\}$ is in $\text{btr}_B(\mathcal{H}_1, S \setminus \{x\})$. First, we show that T' is a transversal of \mathcal{H}_1 . Let $e \in \mathcal{H}_1$. By definition of \mathcal{H}_1 , $x \notin e$, thus $e \in \mathcal{H}_2$ as well. Therefore $e \cap T' \neq \emptyset$. We now prove that T' is minimal in $\mathcal{H}_1 \setminus \mathcal{H}_1(B)$. As $T' \in \text{mtr}(\mathcal{H}_2 \setminus \mathcal{H}_2(B \cup \{x\}))$, every vertex in T' has a private hyperedge in $\mathcal{H}_2 \setminus \mathcal{H}_2(B \cup \{x\})$. Moreover, recall that $\mathcal{H}_2 \setminus \mathcal{H}_2(B \cup \{x\}) = \mathcal{H}_1 \setminus \mathcal{H}_1(B)$. Thus, T' is minimal in $\mathcal{H}_1 \setminus \mathcal{H}_1(B)$. As $T' \subseteq S \setminus \{x\}$, we have $T' \in \text{btr}_B(\mathcal{H}_1, S \setminus \{x\})$.

We finish the proof by showing that $T \notin \text{btr}_B(\mathcal{H}, S)(x)$. In order to prove it, we show that there is no private hyperedge for x w.r.t. T and $\mathcal{H} \setminus \mathcal{H}(B)$. Indeed, since $T' \in \text{tr}(\mathcal{H}_2)$, every hyperedge in \mathcal{H}_2 contains a vertex in T' . By definition of \mathcal{H}_2 , we have $\mathcal{H} \setminus \mathcal{H}(B) \subseteq \mathcal{H}_2$, thus for every hyperedge e in $\mathcal{H} \setminus \mathcal{H}(B)$, we have $e \cap T \neq \{x\}$, i.e., e is not a private for x . \square

By Claim 1 and Claim 2 we can conclude the lemma. \square

Finally, we characterise the number of B -blocked transversals of a hypergraph that do not contain a given vertex. We use the symbol \uplus for the disjoint union of sets.

Lemma 7. Let \mathcal{H} be a hypergraph, $S, B \subseteq V(\mathcal{H})$ and $x \in S$. We have

$$\text{btr}_B(\mathcal{H}, S) = \text{btr}_B(\mathcal{H}, S)(x) \uplus \text{btr}_B(\mathcal{H}, S \setminus \{x\}).$$

Proof. Let $T \subseteq S$ be a B -blocked transversal of \mathcal{H} . Thus, either $x \in T$ and then $T \in \text{btr}_B(\mathcal{H}, S)(x)$, or $x \notin T$ and then $T \subseteq S \setminus \{x\}$, i.e., $T \in \text{btr}_B(\mathcal{H}, S \setminus \{x\})$. Since, the two cases are exclusive, we can conclude that $\text{btr}_B(\mathcal{H}, S)$ is the disjoint union of $\text{btr}_B(\mathcal{H}, S)(x)$ and of $\text{btr}_B(\mathcal{H}, S \setminus \{x\})$. \square

A direct consequence of Lemma 5, Lemma 6 and Lemma 7 is the following equality characterising the number of B -blocked transversals of \mathcal{H} . This will be a crucial step in our dynamic programming scheme.

Theorem 8. *Let \mathcal{H} be a hypergraph, $S, B \subseteq V(\mathcal{H})$ and $x \in S$. We have*

$$\#btr_B(\mathcal{H}, S) = \#btr_B(\mathcal{H}, S \setminus \{x\}) + \#btr_B(\mathcal{H}_1, S \setminus \{x\}) - \#btr_{B \cup \{x\}}(\mathcal{H}_2, S \setminus \{x\})$$

where $\mathcal{H}_1 := \mathcal{H} \setminus \mathcal{H}(x)$ and $\mathcal{H}_2 := \mathcal{H} \setminus (\mathcal{H}(B) \cap \mathcal{H}(x))$.

Proof. By Lemma 7, $\#btr_B(\mathcal{H}, S) = \#btr_B(\mathcal{H}, S \setminus \{x\}) + \#btr_B(\mathcal{H}, S)(x)$. By Lemma 5 and Lemma 6, $\{\{x\}\} \otimes btr_B(\mathcal{H}_1, S \setminus \{x\}) = btr_B(\mathcal{H}, S)(x) \uplus \{\{x\}\} \otimes btr_{B \cup \{x\}}(\mathcal{H}_2, S \setminus \{x\})$. Hence, $\#btr_B(\mathcal{H}, S)(x) = \#btr_B(\mathcal{H}_1, S \setminus \{x\}) - \#btr_{B \cup \{x\}}(\mathcal{H}_2, S \setminus \{x\})$. Therefore, the claimed equality holds. \square

3. Counting the minimal transversals of β -acyclic hypergraphs

In this section, we fix \mathcal{H} a β -acyclic hypergraph, \leq a β -elimination ordering of its vertices and we let $\leq_{\mathcal{H}}$ the induced lexicographic ordering on the hyperedges, i.e., $e \leq_{\mathcal{H}} f$ if $\min((e \setminus f) \cup (f \setminus e)) \in e$. We denote by \mathcal{H}_e^x the sub-hypergraph of \mathcal{H} formed by the hyperedges $f \in \mathcal{H}$ such that there exists a walk from f to e going only through hyperedges smaller than e and vertices smaller than x . For an example, take the β -elimination ordering a, b, x, c, d of the hypergraph in Figure 1 and the induced ordering $\{a, b\}, \{b, x\}, \{x, c\}, \{c, d\}$ on \mathcal{H} . For $e = \{x, c\}$, the hypergraph \mathcal{H}_e^x is composed of the hyperedges $\{b, x\}$ and $\{x, c\}$.

For a vertex x of $V(\mathcal{H})$, we write $[\leq x]$, $[< x]$ and $[\geq x]$ for, respectively, $\{y \in V(\mathcal{H}) \mid y \leq x\}$, $\{y \in V(\mathcal{H}) \mid y \leq x \wedge y \neq x\}$ and $\{y \in V(\mathcal{H}) \mid x \leq y\}$. Moreover, we write $\mathcal{H}[\leq x]$, $\mathcal{H}[< x]$ and $\mathcal{H}[\geq x]$ instead of, respectively, $\mathcal{H}[[\leq x]]$, $\mathcal{H}[[< x]]$ and $\mathcal{H}[[\geq x]]$.

3.1. Decomposition of β -acyclic hypergraphs

The following two lemmas have been proven in [11, Section III-A].

Lemma 9 (Theorem 3 in [11]). *For every hyperedge $e \in \mathcal{H}$, and $x \in V(\mathcal{H})$, we have $V(\mathcal{H}_e^x) \cap [\geq x] \subseteq e$.*

Lemma 10 (Lemma 2 in [11]). *Let e and f be two hyperedges of \mathcal{H} such that $e \leq_{\mathcal{H}} f$, and let x and y be vertices of \mathcal{H} such that $x \leq y$. If $V(\mathcal{H}_e^x) \cap V(\mathcal{H}_f^y) \cap [\leq x] \neq \emptyset$, then $\mathcal{H}_e^x \subseteq \mathcal{H}_f^y$.*

We prove a lemma on the decomposition of \mathcal{H}_e^x graphs that will be used with Lemma 4 to propagate the dynamic programming algorithm.

Lemma 11. *Let $x \in V(\mathcal{H})$, $e \in \mathcal{H}$ and $S \subseteq [\geq x]$. Let*

$$\mathcal{H}' := \begin{cases} \mathcal{H}_e^x & \text{if } S = \emptyset, \\ \mathcal{H}_e^x \setminus (\bigcap_{w \in S} \mathcal{H}_e^x(w)) & \text{otherwise.} \end{cases}$$

For every connected component C of $\mathcal{H}'[< x]$ different from $\{\emptyset\}$, there exists $y < x$ and $f \leq_{\mathcal{H}} e$ such that $C = \mathcal{H}_f^y[\leq y]$ and $\mathcal{H}_f^y = \{g \in \mathcal{H}' \mid g \cap [< x] \in C\}$.

Proof. Let $y = \max(V(C))$ and $f = \max\{g \in \mathcal{H}' \mid g \cap [< x] \in C\}$. We show that $\mathcal{H}_f^y = \{g \in \mathcal{H}' \mid g \cap [< x] \in C\}$.

First, we observe that $\mathcal{H}_f^y \subseteq \mathcal{H}'$. If $S = \emptyset$, it follows from Lemma 10 because, in this case, $\mathcal{H}' = \mathcal{H}_e^x$. Suppose that $S \neq \emptyset$, by definition of \mathcal{H}' , we have $S \not\subseteq f$. Moreover, by Lemma 9, we have that $V(\mathcal{H}_f^y) \cap [\geq y] \subseteq f$. As $S \subseteq [\geq x]$ and $x > y$, we have $S \subseteq [\geq y]$. Thus $S \not\subseteq V(\mathcal{H}_f^y)$ and for all $g \in \mathcal{H}_f^y$, we have $S \not\subseteq g$ since $g \subseteq V(\mathcal{H}_f^y)$. We can conclude that $\mathcal{H}_f^y \subseteq \mathcal{H}'$.

Now, we prove that every $g \in \mathcal{H}_f^y$, we have $g \cap [< x] \in C$. Let $g \in \mathcal{H}_f^y$. By definition of \mathcal{H}_f^y and because $\mathcal{H}_f^y \subseteq \mathcal{H}'$, there exists a path P from f to g going only through vertices smaller than y and hyperedges smaller than f in \mathcal{H}' . As $y < x$, we can conclude that $f \cap [< x]$ is connected to $g \cap [< x]$ in \mathcal{H}' , i.e., $g \cap [< x] \in C$. In other words, we have $\mathcal{H}_f^y \subseteq \{g \in \mathcal{H}' \mid g \cap [< x] \in C\}$.

It remains to prove the other inclusion. Let $g \in \mathcal{H}'$ with $g \cap [< x] \in C$. Since C is a connected component of $\mathcal{H}'[< x]$, there exists a path P from $f \cap [< x]$ to $g \cap [< x]$. By the maximality of y and f , P goes only through vertices smaller than y and hyperedges smaller than f . We can construct from P a path P' from f to g in \mathcal{H}' going through vertices smaller than y and hyperedges smaller than f . As $\mathcal{H}' \subseteq \mathcal{H}$, we can conclude that $g \in \mathcal{H}_f^y$ and thus, $\mathcal{H}_f^y = \{g \in \mathcal{H}' \mid g \cap [< x] \in C\}$. Finally, observe that $C = \mathcal{H}_f^y[< x] = \mathcal{H}_f^y[\leq y]$ since $y = \max(V(C))$. \square

3.2. The algorithm

In this subsection, we describe the dynamic programming algorithm we use to count the number of minimal transversals of a β -acyclic hypergraph. We denote by x_1 the smallest element of \leq .

Our goal is to compute $\#\text{btr}_\emptyset(\mathcal{H}_e^x, [\leq x])$ and $\#\text{btr}_w(\mathcal{H}_e^x, [\leq x])$ for every $e \in \mathcal{H}$, $x \in V(\mathcal{H})$ and $w \in V(\mathcal{H})$ such that $x < w$. Observe that it is enough for computing the number of minimal transversals of \mathcal{H} as $\#\text{mtr}(\mathcal{H}) = \#\text{btr}_\emptyset(\mathcal{H}_{e_m}^{x_n}, [\leq x_n])$ where e_m is the maximal hyperedge for $\leq_{\mathcal{H}}$ and x_n is the maximal vertex for \leq . Indeed, we have $\mathcal{H}_{e_m}^{x_n} = \mathcal{H}$ and $[\leq x_n] = V(\mathcal{H})$, thus $\text{btr}_\emptyset(\mathcal{H}_{e_m}^{x_n}, [\leq x_n]) = \text{btr}_\emptyset(\mathcal{H}) = \text{mtr}(\mathcal{H})$.

The propagation of the dynamic programming works as follows: we use Theorem 8 to reduce the computation of $\#\text{btr}_B(\mathcal{H}_e^x, [\leq x])$ to the computation of $\#\text{btr}$ for several hypergraphs that do not contain x . We then use Lemma 11 to show that these hypergraphs can be decomposed into disjoint hypergraphs of the form \mathcal{H}_f^y for $f \leq_{\mathcal{H}} e$ and $y < x$ which allows us to compute $\#\text{btr}_B(\mathcal{H}_e^x, [\leq x])$ from precomputed values of the form $\#\text{btr}_{B'}(\mathcal{H}_f^y, [\leq y])$, where $B' \in \{B, \{x\}\}$.

Before continuing, let us give a high-level description of the algorithm. For each $1 \leq i \leq n$ and each $1 \leq j \leq m$, let $\text{tab}[i, j, 0]$ be $\#\text{btr}_\emptyset(\mathcal{H}_{e_j}^{x_i}, [\leq x_i])$, and for each $\ell > i$, let $\text{tab}[i, j, \ell]$ be $\#\text{btr}_{x_\ell}(\mathcal{H}_{e_j}^{x_i}, [\leq x_i])$. Because the number of minimal transversals of \mathcal{H} is $\#\text{btr}_\emptyset(\mathcal{H}_{e_m}^{x_n}, [\leq x_n])$, it is enough to show how to compute $\text{tab}[n, m, 0]$ in polynomial time. The following is a high level description of the algorithm which computes $\text{tab}[i, j, \ell]$, for all $1 \leq i \leq n$, $1 \leq j \leq m$ and $\ell \in \{0, i+1, \dots, n\}$.

Algorithm CountMinTransversals(\mathcal{H})

\mathcal{H} : a β -acyclic hypergraph

1. Let x_1, \dots, x_n a β -elimination ordering of \mathcal{H} .
2. Let e_1, \dots, e_m the induced lexicographic ordering on \mathcal{H} .
3. Precompute $\mathcal{H}_{e_j}^{x_i}, [\leq x_i]$ for every $i \leq n, j \leq m$.
4. **for** $1 \leq i \leq n$ and $i < \ell \leq n$ **do**
5. **for** $1 \leq j \leq m$ **do**
6. Compute $tab[i, j, 0]$ from the recursive formula of $\#btr_{\emptyset}(\mathcal{H}_{e_j}^{x_i}, [\leq x_i])$.
7. Compute $tab[i, j, \ell]$ from the recursive formula of $\#btr_{x_\ell}(\mathcal{H}_{e_j}^{x_i}, [\leq x_i])$.
8. **end for**
9. **end for**
10. **return** $tab[n, m, 0]$

In order to ease the presentation, the computation of $\#btr_{\emptyset}(\mathcal{H}_e^x, [\leq x])$ and that of $\#btr_w(\mathcal{H}_e^x, [\leq x])$ are separated, even though many of the arguments are similar.

3.2.1. Base cases.

We observe that for every $e \in \mathcal{H}$, $\mathcal{H}_e^{x_1}[\leq x_1]$ is either equal to $\{x_1\}$ or $\{\emptyset\}$. Thus, for every $e \in \mathcal{H}$ and $w \in V(\mathcal{H})$ such that $w > x_1$, we can compute $\#btr_{\emptyset}(\mathcal{H}_e^x, [\leq x_1])$ and $\#btr_w(\mathcal{H}_e^x, [\leq x_1])$ in time $O(1)$.

3.2.2. Computing $\#btr_{\emptyset}(\mathcal{H}_e^x, [\leq x])$ by dynamic programming.

We start by explaining how we can compute $\#btr_{\emptyset}(\mathcal{H}_e^x, [\leq x])$ in polynomial time if the values $\#btr_{\emptyset}(\mathcal{H}_f^y, [\leq y])$ and $\#btr_w(\mathcal{H}_f^y, [\leq y])$ have been precomputed for $f \leq_{\mathcal{H}} e$ and $y < x, y < w$.

We start by applying Theorem 8.

$$\begin{aligned} \#btr_{\emptyset}(\mathcal{H}_e^x, [\leq x]) &= \#btr_{\emptyset}(\mathcal{H}_e^x, [< x]) \\ &\quad + \#btr_{\emptyset}(\mathcal{H}_e^x \setminus \mathcal{H}_e^x(x), [< x]) \\ &\quad - \#btr_x(\mathcal{H}_e^x, [< x]). \end{aligned}$$

Now, let C_1, \dots, C_k be the connected components of $\mathcal{H}_e^x[< x]$. If there exists i such that $C_i = \{\emptyset\}$, then $\#btr_{\emptyset}(\mathcal{H}_e^x, [< x]) = \#btr_x(\mathcal{H}_e^x, [< x]) = 0$. Otherwise, by applying Lemma 11 with $S = \emptyset$, there exists, for each $1 \leq i \leq k$, $y_i < x$ and $f_i \leq_{\mathcal{H}} e$ such that $\mathcal{H}_{f_i}^{y_i} = \{g \in \mathcal{H}_e^f \mid g \cap [< x] \in C_i\}$ and $C_i = \mathcal{H}_{f_i}^{y_i}[\leq y_i]$. By Lemma 4,

$$\begin{aligned} \#btr_{\emptyset}(\mathcal{H}_e^x, [< x]) &= \prod_{i=1}^k \#btr_{\emptyset}(\mathcal{H}_{f_i}^{y_i}, [\leq y_i]), \\ \#btr_x(\mathcal{H}_e^x, [< x]) &= \prod_{i=1}^k \#btr_x(\mathcal{H}_{f_i}^{y_i}, [\leq y_i]). \end{aligned}$$

We now show how to decompose $\#btr_{\emptyset}(\mathcal{H}_e^x \setminus \mathcal{H}_e^x(x), [< x])$ into a product of precomputed values. Let D_1, \dots, D_l be the connected components of $\mathcal{H}_e^x \setminus$

$\mathcal{H}_e^x(x)[< x]$. If there exists i such that $D_i = \{\emptyset\}$, then $\#btr_x(\mathcal{H}_e^x \setminus \mathcal{H}_e^x(x), [< x]) = 0$. Otherwise, if we apply Lemma 11 with $S = \{x\}$, then there exists, for each $1 \leq i \leq l$, $y_i < x$ and $f_i \leq_{\mathcal{H}} e$ such that $\mathcal{H}_{f_i}^{y_i} = \{g \in \mathcal{H}_e^x \setminus \mathcal{H}_e^x(x) \mid g \cap [< x] \in D_i\}$ and $D_i = \mathcal{H}_{f_i}^{y_i}[\leq y_i]$. We can thus conclude by Lemma 4 that

$$\#btr_{\emptyset}(\mathcal{H}_e^x \setminus \mathcal{H}_e^x(x), [< x]) = \prod_{i=1}^l \#btr_{\emptyset}(\mathcal{H}_{f_i}^{y_i}, [\leq y_i]).$$

Therefore, if $\#btr_{\emptyset}(\mathcal{H}_f^y, [\leq y])$ and $\#btr_x(\mathcal{H}_f^y, [\leq y])$ have already been computed for every $f <_{\mathcal{H}} e$ and $y \leq x$, we can compute $\#btr_{\emptyset}(\mathcal{H}_e^x, [\leq x])$ with at most $3 \times |\mathcal{H}_e^x|$ additional multiplications and 3 additions.

3.2.3. Computing $\#btr_w(\mathcal{H}_e^x, [\leq x])$ by dynamic programming.

Let $x \leq w$. By Theorem 8, we have:

$$\begin{aligned} \#btr_w(\mathcal{H}_e^x, [\leq x]) &= \#btr_w(\mathcal{H}_e^x, [< x]) \\ &\quad + \#btr_w(\mathcal{H}_e^x \setminus \mathcal{H}_e^x(x), [< x]) \\ &\quad - \#btr_{\{x,w\}}(\mathcal{H}_e^x \setminus (\mathcal{H}_e^x(x) \cap \mathcal{H}_e^x(w)), [< x]). \end{aligned}$$

We start by explaining how to compute $\#btr_w(\mathcal{H}_e^x, [< x])$. Let C_1, \dots, C_k be the connected components of $\mathcal{H}_e^x [< x]$. If there exists i such that $C_i = \{\emptyset\}$, then $\#btr_w(\mathcal{H}_e^x, [< x]) = 0$. Otherwise, by Lemma 11 with $S = \{w\}$, there exists, for each $1 \leq i \leq k$, $f_i \leq_{\mathcal{H}} e$ and $y_i < x$ such that $\mathcal{H}_{f_i}^{y_i} = \{g \in \mathcal{H}_e^x \mid g \cap [< x] \in C_i\}$ and $C_i = \mathcal{H}_{f_i}^{y_i}[\leq y_i]$. By Lemma 4, we can conclude that

$$\#btr_w(\mathcal{H}_e^x, [< x]) = \prod_{i=1}^k \#btr_w(\mathcal{H}_{f_i}^{y_i}, [\leq y_i]).$$

We now explain how to compute $\#btr_w(\mathcal{H}_e^x \setminus \mathcal{H}_e^x(x), [< x])$. Let D_1, \dots, D_l be the connected components of $(\mathcal{H}_e^x \setminus \mathcal{H}_e^x(x)) [< x]$. If there exists i such that $D_i = \{\emptyset\}$, then $\#btr_w(\mathcal{H}_e^x \setminus \mathcal{H}_e^x(x), [< x]) = 0$. Otherwise, by applying Lemma 11 with $S = \{x\}$, it follows that for every $1 \leq i \leq l$, there exists $y_i < x$ and $f_i \leq_{\mathcal{H}} e$ such that $\mathcal{H}_{f_i}^{y_i} = \{g \in \mathcal{H}_e^x \setminus \mathcal{H}_e^x(x) \mid g \cap [< x] \in D_i\}$ and $D_i = \mathcal{H}_{f_i}^{y_i}[\leq y_i]$. Thus, from Lemma 4,

$$\#btr_w(\mathcal{H}_e^x \setminus \mathcal{H}_e^x(x), [< x]) = \prod_{i=1}^l \#btr_w(\mathcal{H}_{f_i}^{y_i}, [\leq y_i]).$$

Finally, we explain how to decompose $\#btr_{\{x,w\}}(\mathcal{H}_e^x \setminus (\mathcal{H}_e^x(x) \cap \mathcal{H}_e^x(w)), [< x])$ into a product of pre-computed values. To ease the notation, we denote $\mathcal{H}_e^x \setminus (\mathcal{H}_e^x(x) \cap \mathcal{H}_e^x(w))$ by \mathcal{H}' . Let K_1, \dots, K_r be the connected components of $\mathcal{H}' [< x]$. If there exists i such that $K_i = \{\emptyset\}$, then $\#btr_{\{x,w\}}(\mathcal{H}', [< x]) = 0$. Otherwise, by Lemma 11 applied with $S = \{x, w\}$, we have that for every i ,

there exists $y_i < x$ and $f_i \leq_{\mathcal{H}} e$ such that $\mathcal{H}_{f_i}^{y_i} = \{g \in \mathcal{H}' \mid g \cap [< x] \in K_i\}$ and $K_i = \mathcal{H}_{f_i}^{y_i}[\leq y_i]$. By Lemma 4, we have:

$$\#\text{btr}_{\{x,w\}}(\mathcal{H}_e^x \setminus (\mathcal{H}_e^x(x) \cap \mathcal{H}_e^x(w)), [< x]) = \prod_{i=1}^r \#\text{btr}_{\{x,w\} \cap V(\mathcal{H}_{f_i}^{y_i})}(\mathcal{H}_{f_i}^{y_i}, [\leq y_i]).$$

Claim 3. *For every $i \leq p$, $\{x, w\} \cap V(\mathcal{H}_{f_i}^{y_i}) \neq \{x, w\}$.*

Proof. Assume towards a contradiction that $\{x, w\} \cap V(\mathcal{H}_{f_i}^{y_i}) = \{x, w\}$. Recall that $\mathcal{H}_{f_i}^{y_i}[\leq y_i]$. By Lemma 9, $\{x, w\} \subseteq V(\mathcal{H}_{f_i}^{y_i})$ implies $\{x, w\} \subseteq f$. Thus, we have $f \in \mathcal{H}_e^x(x) \cap \mathcal{H}_e^x(w)$. This is a contradiction, since $f \in \mathcal{H}_{f_i}^{y_i} \subseteq \mathcal{H}_e^x \setminus (\mathcal{H}_e^x(x) \cap \mathcal{H}_e^x(w))$. \square

Thus, $\{x, w\} \cap V(\mathcal{H}_{f_i}^{y_i})$ equals either $\{x\}$, or $\{w\}$ or \emptyset by Claim 3. That is, we can compute $\#\text{btr}_{\{x,w\}}(\mathcal{H}_e^x \setminus (\mathcal{H}_e^x(x) \cap \mathcal{H}_e^x(w)), [< x])$ from precomputed terms.

We can conclude that, if $\#\text{btr}_{\emptyset}(\mathcal{H}_f^y, [\leq y])$ and $\#\text{btr}_w(\mathcal{H}_f^y, [\leq y])$ have already been computed for every $f <_{\mathcal{H}} e$ and $y \leq w$, we can compute $\#\text{btr}_w(\mathcal{H}_e^x, [\leq x])$ with at most $3 \times |\mathcal{H}_e^x|$ additional multiplications and 3 additions.

It is easy to see that a straightforward greedy algorithm can be used to compute a β -elimination ordering in polynomial time (see [13] for a better algorithm due to Paige and Tarjan). Moreover, the dynamic programming algorithm describes above computes at most $O(n^2|\mathcal{H}|)$ terms and each of them can be computed from the others with a polynomial number of arithmetic operations. Finally, all these terms can be bounded by 2^n since they are all the cardinals of some collection of subsets of the vertices. Thus these arithmetic operations can be done in polynomial time in the size of the input. It follows.

Theorem 12 (Main Theorem). *Let \mathcal{H} be a β -acyclic hypergraph. One can compute in polynomial time the number of minimal transversals of \mathcal{H} .*

4. Applications to the counting of Dominating Sets

We refer to [14] for our graph terminology. For a graph G , let $V(G)$ be its set of vertices and $E(G)$ be its set of edges. For a vertex x of a graph G , let $N(x)$ be the set of neighbours of x and we let $N[x]$ be the set $\{x\} \cup N(x)$. The *closed neighbourhood hypergraph* of G , denoted by $\mathcal{N}[G]$, is the hypergraph $\{N[x] \mid x \in V(G)\}$. A *dominating set* in a graph G is a transversal of $\mathcal{N}[G]$. DOMINATING SET problems are classic and well-studied graph problems, and has applications in many areas such as networks and graph theory [15].

In [16] the authors reduce the HYPERGRAPH DUALISATION problem into the enumeration of minimal dominating sets, showing that the two problems are equivalent in the area of enumeration problems (a fact already established in the case of optimisation). The reduction indeed shows that the counting versions are equivalent (under Turing reductions), and such a reduction is of big interest because it allows to study counting and enumeration problems associated with

the HYPERGRAPH DUALISATION in the perspectives of graph theory, where tools had been developed to tackle combinatorial problems.

Despite the broad application of counting the minimal dominating sets in (hyper)graphs, the problem was not investigated until recently, except in [17] where it is proved that the models of any monadic second-order formula can be counted in polynomial time in graphs of bounded clique-width. Indeed, as far as we know the counting of minimal dominating sets is only considered in [7, 8, 18]. This problem is known to be polynomial on interval graphs and permutation graphs [7]. However, the systematic study of its computational complexity in graph classes is only considered in [18], where the authors proved the $\#P$ -completeness in several graph classes and asked whether the following dichotomy conjecture is true. A k -sun is a graph obtained from a cycle of length $2k$ ($k \geq 3$) by adding edges to make the even-indexed vertices pairwise adjacent. A graph is *chordal* if it does not contain cycles of length at least 4 as induced subgraphs.

Conjecture 13. *Let \mathcal{C} be a class of chordal graphs. If \mathcal{C} does not contain a k -sun as an induced subgraph, for $k \geq 4$, then one can count in polynomial time the number of minimal dominating sets of any graph in \mathcal{C} . Otherwise, the problem is $\#P$ -complete.*

This conjecture is motivated by the recursive structure of the chordal graphs without k -sun, for $k \leq 4$.

We make a first step towards a proof of the first statement of the conjecture and provide a polynomial time algorithm for computing the minimal dominating sets in *strongly chordal graphs*, which are exactly chordal graphs without k -suns, for $k \geq 3$.

Corollary 14 (Main Corollary). *Let G be a strongly chordal graph. One can count in polynomial time the number of minimal dominating sets of G .*

Proof. Let G be a strongly chordal graph. It is well-known that the hypergraph $\mathcal{N}[G]$ is β -acyclic [19]. By Theorem 12, one can count in polynomial time the minimal transversals of $\mathcal{N}[G]$, which are exactly the minimal dominating sets of G . \square

5. Conclusion

We proposed a polynomial time algorithm for counting the minimal transversals of any β -acyclic hypergraph, supporting Conjecture 13, and it seems that the technique can be easily adapted to consider (inclusion-wise) minimal d -dominating sets, which are dominating sets such that each vertex is dominated by at least d vertices [8].

Besides resolving Conjecture 13, there are two immediate questions that deserve to be considered. Firstly, can we count the minimal models of any non-monotone β -acyclic formula in polynomial time? Secondly, for which graph classes and counting graph problems the techniques of this paper apply?

References

- [1] S. P. Vadhan, The complexity of counting in sparse, regular, and planar graphs, *SIAM J. Comput.* 31 (2) (2001) 398–427 (electronic). doi:10.1137/S0097539797321602.
URL <http://dx.doi.org/10.1137/S0097539797321602>
- [2] L. G. Valiant, The complexity of computing the permanent, *Theoret. Comput. Sci.* 8 (2) (1979) 189–201. doi:10.1016/0304-3975(79)90044-6.
URL [http://dx.doi.org/10.1016/0304-3975\(79\)90044-6](http://dx.doi.org/10.1016/0304-3975(79)90044-6)
- [3] L. A. Goldberg, R. Gysel, J. Lapinskas, Approximately counting locally-optimal structures, *J. Comput. Syst. Sci.* 82 (6) (2016) 1144–1160. doi:10.1016/j.jcss.2016.04.001.
URL <https://doi.org/10.1016/j.jcss.2016.04.001>
- [4] T. Eiter, K. Makino, G. Gottlob, Computational aspects of monotone dualization: a brief survey, *Discrete Appl. Math.* 156 (11) (2008) 2035–2049. doi:10.1016/j.dam.2007.04.017.
URL <http://dx.doi.org/10.1016/j.dam.2007.04.017>
- [5] A. Durand, M. Hermann, On the counting complexity of propositional circumscription, *Inform. Process. Lett.* 106 (4) (2008) 164–170. doi:10.1016/j.ipl.2007.11.006.
URL <http://dx.doi.org/10.1016/j.ipl.2007.11.006>
- [6] F. Capelli, Structural restrictions of cnf formulas: application to model counting and knowledge compilation, Ph.D. thesis, Université Paris Diderot (2016).
- [7] M. M. Kanté, V. Limouzy, A. Mary, L. Nourine, T. Uno, On the enumeration and counting of minimal dominating sets in interval and permutation graphs, in: *Algorithms and computation*, Vol. 8283 of *Lecture Notes in Comput. Sci.*, Springer, Heidelberg, 2013, pp. 339–349. doi:10.1007/978-3-642-45030-3_32.
URL http://dx.doi.org/10.1007/978-3-642-45030-3_32
- [8] P. A. Golovach, P. Heggernes, M. M. Kanté, D. Kratsch, S. H. Sæther, Y. Villanger, Output-polynomial enumeration on graphs of bounded (local) linear mim-width, *Algorithmica* (2017) 1–28doi:10.1007/s00453-017-0289-1.
- [9] T. Eiter, G. Gottlob, Hypergraph transversal computation and related problems in logic and ai, in: *European Workshop on Logics in Artificial Intelligence*, Springer, 2002, pp. 549–564.
- [10] J. Brault-Baron, F. Capelli, S. Mengel, Understanding model counting for beta-acyclic cnf-formulas, in: *32nd International Symposium on Theoretical Aspects of Computer Science, STACS 2015, March 4-7, 2015, Garching*,

- Germany, 2015, pp. 143–156. doi:10.4230/LIPIcs.STACS.2015.143.
URL <https://doi.org/10.4230/LIPIcs.STACS.2015.143>
- [11] F. Capelli, Understanding the complexity of $\#sat$ using knowledge compilation, in: 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017, IEEE Computer Society, 2017, pp. 1–10. doi:10.1109/LICS.2017.8005121.
URL <https://doi.org/10.1109/LICS.2017.8005121>
 - [12] T. Eiter, G. Gottlob, Identifying the minimal transversals of a hypergraph and related problems, SIAM J. Comput. 24 (6) (1995) 1278–1304. doi:10.1137/S0097539793250299.
URL <http://dx.doi.org/10.1137/S0097539793250299>
 - [13] R. Paige, R. E. Tarjan, Three partition refinement algorithms, SIAM Journal on Computing 16 (6) (1987) 973–989.
 - [14] R. Diestel, Graph Theory, 3rd Edition, no. 173 in Graduate Texts in Mathematics, Springer, 2005.
 - [15] T. W. Haynes, S. T. Hedetniemi, P. J. Slater, Fundamentals of domination in graphs, Vol. 208 of Monographs and Textbooks in Pure and Applied Mathematics, Marcel Dekker, Inc., New York, 1998.
 - [16] M. M. Kanté, V. Limouzy, A. Mary, L. Nourine, On the enumeration of minimal dominating sets and related notions, SIAM J. Discrete Math. 28 (4) (2014) 1916–1929. doi:10.1137/120862612.
URL <http://dx.doi.org/10.1137/120862612>
 - [17] B. Courcelle, The monadic second-order logic of graphs XVI : Canonical graph decompositions, Logical Methods in Computer Science 2 (2).
 - [18] M. M. Kanté, T. Uno, Counting minimal dominating sets, in: Theory and Applications of Models of Computation - 14th Annual Conference, TAMC 2017, Bern, Switzerland, April 20-22, 2017, Proceedings, 2017, pp. 333–347.
 - [19] A. Brandstädt, V. B. Le, J. P. Spinrad, Graph classes: a survey, SIAM Monographs on Discrete Mathematics and Applications, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1999. doi:10.1137/1.9780898719796.
URL <http://dx.doi.org/10.1137/1.9780898719796>