

Efficient enumeration of solutions produced by closure operations

Arnaud Mary¹Yann Strozecki²¹ Université Lyon 1 ; CNRS, UMR5558, LBBE / INRIA - ERABLE² Université de Versailles Saint-Quentin-en-Yvelines, DAVID laboratoryreceived 12th Dec. 2017, revised 20th Apr. 2019, 14th Sep. 2018, accepted 5th June 2019.

In this paper we address the problem of generating all elements obtained by the saturation of an initial set by some operations. More precisely, we prove that we can generate the closure of a boolean relation (a set of boolean vectors) by polymorphisms with a polynomial delay. Therefore we can compute with polynomial delay the closure of a family of sets by any set of “set operations”: union, intersection, symmetric difference, subsets, supersets . . .). To do so, we study the MEMBERSHIP \mathcal{F} problem: for a set of operations \mathcal{F} , decide whether an element belongs to the closure by \mathcal{F} of a family of elements. In the boolean case, we prove that MEMBERSHIP \mathcal{F} is in P for any set of boolean operations \mathcal{F} . When the input vectors are over a domain larger than two elements, we prove that the generic enumeration method fails, since MEMBERSHIP \mathcal{F} is NP-hard for some \mathcal{F} . We also study the problem of generating minimal or maximal elements of closures and prove that some of them are related to well known enumeration problems such as the enumeration of the circuits of a matroid or the enumeration of maximal independent sets of a hypergraph.

Keywords: enumeration, set saturation, incremental polynomial time, polynomial delay, Post’s lattice, maximal independent sets

1 Introduction

An enumeration problem is the task of listing all elements of a set without redundancies. Since the set to generate may be of exponential cardinality in the size of the input, the complexity of enumeration problems generally are measured in term of the input size *and* output size. Enumeration algorithms whose complexity depends both on the input and the output are called *output sensitive* and when the dependency is polynomial in the sum of both measures, they are called *output polynomial*. Another more precise notion of complexity, is the *delay* which measures the time between the production of two consecutive solutions. We are especially interested in problems solvable with a delay polynomial in the input size, which are considered as the *tractable problems* in enumeration complexity. For instance, the spanning trees, the simple cycles [33] or the maximal independent sets [22] of a graph can be enumerated with polynomial delay.

If we allow the delay to grow during the algorithm, we obtain *polynomial incremental time* algorithms: the first k solutions can be enumerated in a time polynomial in k and in the size of the input. Many problems which can be solved in polynomial incremental time have the following form: given a set of elements

and a polynomial time function acting on tuples of elements, produce the closure of the set by the function. For instance the following problems can be solved in polynomial incremental time: the enumeration of the circuits of a matroid [24] and the enumeration of the vertices of restricted polyhedra [17].

In this article, we try to understand when saturation problems, which by definition can be solved in polynomial incremental time, can be in fact solved by a polynomial delay algorithm. We would also like to get rid of the exponential space which is necessary in the enumeration algorithm by saturation. To tackle this question we need to restrict the set of saturation operations we consider. An element will be a vector over some finite set and in most of this article, we require the saturation operations to act *coefficient-wise* and in the same way on each coefficient. We prove that, when the vector is over the boolean domain, every possible saturation can be computed with polynomial delay. To do that we study a decision version of our problem, denoted by $\text{MEMBERSHIP}_{\mathcal{F}}$: given a vector v and a set of vectors \mathcal{S} decide whether v belongs to the closure of \mathcal{S} by the operations of \mathcal{F} . We prove that $\text{MEMBERSHIP}_{\mathcal{F}} \in \text{P}$ for all set of operations \mathcal{F} over the boolean domain.

When the domain is boolean, the problem can be reformulated in terms of set systems or hypergraphs. It is equivalent to the generation of the smallest hypergraph which contains a given hypergraph and which is closed under some operations. We show how to efficiently compute the closure of a hypergraph by any family of set operations (any operation that is the composition of unions, intersections and complementations) on the hyperedges. Some particular cases were already known previously. For instance, it is well known that if a family of subsets ordered by inclusion forms a lattice, then the set of so called *meet irreducible elements* is a generator with respect to the intersection operation (all other sets can be expressed as the intersections of some meet irreducible elements). In general, knowing how to compute a closure may serve as a good tool to design other enumeration algorithms. One only has to express an enumeration problem as the closure of some sufficiently small and easy to compute set of elements (called a generator) and then to apply the algorithms presented in this article.

The closure computation is also related to constraint satisfaction problems (CSP). Indeed, the set of vectors can be seen as a relation R and the problem of generating its closure by some operations \mathcal{F} is equivalent to the computation of the smallest relation R' containing R and for which all functions of \mathcal{F} are polymorphisms of R' . There are several works related to the enumeration in the context of CSP. They deal with the enumeration of solutions of a CSP with polynomial delay [13, 7]. The simplest such result [13] states that in the boolean case, there is a polynomial delay algorithm if and only if the constraint language is Horn, anti-Horn, bijunctive or affine. While those works deal with the enumeration of solutions of CSPs, this paper focuses on finding the closure of relations. However, we use tools from CSPs such as Post's lattice [32], used by Schaefer in his seminal paper [35], and the Baker-Pixley theorem [2].

The main theorem of this article settles the complexity of a whole family of decision problems and implies, quite surprisingly, that the backtrack search is enough to obtain a polynomial delay algorithm to enumerate any closure of boolean vectors. For all these enumeration problems, compared to the naive saturation algorithm, our method has a better time complexity (even from a practical point of view) and a better space complexity (polynomial rather than exponential). Moreover, besides the generic enumeration algorithm, we give for each closure rule an algorithm with the best possible complexity. In doing so, we illustrate several classical methods used to enumerate objects such as amortized backtrack search, hill climbing, Gray code . . . It is interesting to note that most algorithms we provide have a delay polynomial in the maximum size of a solution we generate. However some are polynomial in the instance which can be much larger than the size of a solution, for instance the algorithm computing the closure of sets by intersection as explained in Section 3.1. In that case, we provide a reduction from the problem of

generating the assignments of a monotone *DNF* formula which suggests that the delay must depend on the instance size.

In a second part of the article, we generalize the set of operators used to compute a closure. The aim is to capture more interesting problems and to better understand the difference between polynomial incremental time and polynomial delay. The first generalization is to consider larger domains for the input vectors. In that setting, the problem $\text{MEMBERSHIP}_{\mathcal{F}}$ is NP-complete for some \mathcal{F} and we are not able to settle the question in general. The second generalization is to allow the operators to act differently on each coefficient. We prove that $\text{MEMBERSHIP}_{\mathcal{F}} \in \text{P}$ when the operators are extremely simple: they change only a single coefficient of the vector and leave the rest unchanged. However, allowing the operators to act on three coefficients is already enough to make $\text{MEMBERSHIP}_{\mathcal{F}}$ NP complete.

It is classical in enumeration to try to reduce the number of generated objects, which can be very large, by requiring additional properties. The most classical properties are the maximality or minimality for inclusion, since this is often compatible with other notion of optimality of solutions. Hence algorithms are known for maximal matchings [41], maximal cliques [22], minimal transversals [16] . . . Therefore, as a third generalization we propose to enumerate only the minimal or maximal elements of the closures. In these settings, the problems are not automatically in polynomial incremental time, since no saturation algorithm generate the maximal or minimal elements only. We prove that either these problems have a polynomial delay algorithm or that they are equivalent to well known problems such as the generation of the circuits of a binary matroid or of the maximal independent sets of a hypergraph which can be solved in polynomial incremental time.

This article is a long version of our previous work see [28]. The proofs have been improved and are more detailed, the complexity of several enumeration algorithms have been improved, sometimes using new techniques and more lower bounds are provided. The subject of the last two sections (non uniform operators and maximal sets) have been proposed in [29] and are new for the most part.

1.1 Organization of the paper

In Section 2, we define enumeration complexity, the closure problem ENUMCLOSURE and the backtrack search. In Section 3, we use Post's lattice, restricted through suitable reductions between clones, to prove that $\text{MEMBERSHIP}_{\mathcal{F}}$ is polynomial for all set of binary operations \mathcal{F} . It turns out that there are only a few types of closure operations to study: the monotone operations (Section 3.1), the addition over \mathbb{F}_2 (Section 3.2), the set of all operations (Section 3.3), two infinite hierarchies related to the majority function (Section 3.5) and the limit cases of the previous hierarchies (Section 3.4). For all these closure operations we give for ENUMCLOSURE both a generic backtrack algorithm in polynomial delay and also an ad hoc efficient enumeration algorithm, sometimes with a conditional lower bound matching its complexity. In Section 4, we give polynomial delay algorithms for three classes of closure operations over any domain and prove that the backtrack search we use in the boolean case fails. In Section 5, we enrich the set of possible operations by operators acting on a single coefficient, and prove that the problem ENUMCLOSURE is still in polynomial delay for any set of such operators. Finally, in Section 6 we try to enumerate only the minimal or maximal elements of a closure and we show it is either trivially in polynomial delay or equivalent to famous enumeration problems for which no polynomial delay algorithm are known such as enumerating the circuits of a binary matroid or the maximal independent sets of a k uniform hypergraph.

2 Preliminary

Given $n \in \mathbb{N}$, $[n]$ denotes the set $\{1, \dots, n\}$. For a set D and a vector $v \in D^n$, we denote by v_i the i^{th} coordinate of v . Let $i, j \in [n]$, we denote by $v_{i,j}$ the vector (v_i, v_j) . More generally, for a subset $I = \{i_1, \dots, i_k\}$ of $[n]$ with $i_1 < \dots < i_k$ we denote by v_I the vector $(v_{i_1}, \dots, v_{i_k})$. Let \mathcal{S} be a set of vectors. We denote by \mathcal{S}_I the set $\{v_I \mid v \in \mathcal{S}\}$, we say that v_i is a *projection* of v and \mathcal{S}_I is a projection of \mathcal{S} . The characteristic vector v of a subset E of $[n]$, is the vector in $\{0, 1\}^n$ such that $v_i = 1$ if and only if $i \in E$.

2.1 Complexity

In this section, we recall basic definitions about enumeration problems and their complexity, for further details and examples see [38].

Let Σ be some finite alphabet. An *enumeration problem* is a function A from Σ^* to $\mathcal{P}(\Sigma^*)$. That is to each input word, A associates a set of words. An algorithm which solves the enumeration problem A takes any input word w and produces the set $A(w)$ word by word and *without redundancies*. We always require the sets $A(w)$ to be finite. We may also ask $A(w)$ to contain only words of polynomial size in the size of w and that one can test whether an element belongs to $A(w)$ in polynomial time. If those two conditions hold, the problem is in the class EnumP which is the counterpart of NP for enumeration. Because of this relationship to NP, we often call solutions the elements we enumerate.

The computational model is the random access machine model (RAM) with addition, subtraction and multiplication as its basic arithmetic operations. We have additional output registers, and when a special OUTPUT instruction is executed, the contents of the output registers is outputted and considered as an element of the outputted set. The RAM model is chosen over Turing machines, because it allows to use data structures more efficiently which impacts the complexity measure used in enumeration.

The *delay* is the time between the productions of two consecutive solutions. Usually we want to bound the delay of an algorithm for all pairs of consecutive solutions and for all inputs of the same size. If this delay is polynomial in the size of the input, then we say that the algorithm is in *polynomial delay* and the problem is in the class DelayP. If the delay is polynomial in the input and the number of already generated solutions, we say that the algorithm is in *incremental delay* and the problem is in the class IncP. By definition we have DelayP \subset IncP. Moreover, if we assume the exponential time hypothesis, $(\text{DelayP} \cap \text{EnumP}) \neq (\text{IncP} \cap \text{EnumP})$ [11]. In practice problems in DelayP are considered to be tractable, because the total time is linear in the number of solutions generated and because the solutions can be regularly generated. Moreover, these problems often enjoy algorithms with polynomial space in the size of the input, which is not the case for most polynomial incremental algorithms. Note that in a polynomial delay algorithm we allow a *polynomial precomputation step*, usually to set up data structures, which is not taken into account in the delay. This is why we can have a delay *smaller than the size of the input*.

For S a set of boolean formulas, we define the enumeration problem ENUMS as the function which associates to a formula of S the set of its models. One classical example is when S is the set of formulas in conjunctive normal form denoted by *CNF*, and the associated problem is ENUMCNF or equivalently ENUMSAT. In this article, we consider other families of boolean formulas: the Horn formulas *HORN*, the bijnunctive formulas *2SAT*, the formulas in disjunctive normal form *DNF* and the monotone formulas in disjunctive normal form *MONDNF*.

To compare two enumeration problems, we use often the notion of *parsimonious reduction* from count-

ing complexity as it is adapted to enumeration complexity contrarily to Cook or Karp reductions. Given two enumeration problems A and B , we say that there is a parsimonious reduction from A to B if there are two polynomial time functions f and g such that for all instances x of A , g maps bijectively $B(f(x))$ to $A(x)$. The problem ENUMCNF is EnumP complete for parsimonious reductions by adapting the classic proof of Cook [12]. In this article, we need a slightly more general notion of reduction, such that a solution of B may yield several solutions of A as long as they can be efficiently enumerated, as introduced in [27].

Definition 1. Let A and B be two enumeration problems. We say that there is a *polynomial delay reduction* from A to B if there is a polynomial time function f and C an enumeration problem in DelayP such that for all instances x of A , we have:

- if $y \neq z$ and $y, z \in B(f(x))$ then $C(xy) \cap C(xz) = \emptyset$
- $A(x) = \bigcup_{y \in B(f(x))} C(xy)$

where xy is the concatenation of x and y .

When a problem A reduces to B and B reduces to A , we say that A and B are *equivalent*. Remark that the class DelayP is closed under polynomial delay reduction. We will use tight reduction, where the size of $f(x)$ is closely related to x and C has a linear delay to obtain lower and upper bounds on the complexity of several enumeration problems.

We now explain a very classical and natural enumeration method called the *Backtrack Search* (sometimes also called the *flashlight method*) used in many previous articles [33, 39]. It can be used to solve all auto-reducible problems, in particular ENUMC when C is a set of formulas solvable in polynomial time [13] and closed by partial assignment of variables. We represent the solutions we want to enumerate as vectors of size n and coefficients in D . In practice solutions are often subsets of $[n]$ which means that $D = \{0, 1\}$ and the vector is the characteristic vector of the subset.

The enumeration algorithm is a depth first traversal of a tree whose nodes are partial solutions. The nodes of the tree are all vectors v of size l , for all $l \leq n$, such that $v = w_{[l]}$ and w is a solution. The children of the node v are the vectors of size $l + 1$, which restricted to $[l]$ are equal to v . The leaves of this tree are the solutions of our problem, therefore a depth first traversal visits all leaves and thus outputs all solutions. Since a branch of the tree is of size n , it is enough to find the children of a node in a time polynomial in n to obtain a polynomial delay. The delay also depends linearly on $|D|$, but in the rest of the paper $|D|$ will be constant. Therefore the backtrack search is in polynomial delay if and only if the following *decision* problem is in P: given v of size l is there w a solution such that $v = w_{[l]}$? This problem is called the *extension problem* associated to the enumeration problem.

Proposition 2. *Given an enumeration problem A , such that for all w , $A(w)$ is a set of vectors of size n and coefficients in D , with n and $|D|$ polynomially related to $|w|$. If the extension problem associated to A is in P, then A is in DelayP.*

We will see in the next part, that the complexity of solving the extension problem can be amortized over a whole branch of the tree, since we solve it many times, using well chosen data structures.

There is a second classical enumeration method to design a polynomial delay algorithm, named the *supergraph* method. The idea is to organize the solutions (and not the partial solutions) as a strongly connected digraph instead of a tree, and to traverse this supergraph. For that we should be able to visit

all the successors of a node in polynomial time. To avoid the storage of the nodes of the supergraph, the *reverse search* method is often used [1]. It consists on defining a canonical parent computable in polynomial time for each node and thus defining a spanning arborescence of the supergraph. This method may have a better delay than the backtrack search, because the traversal goes over solutions only.

2.2 Closure of families by set operations

We fix a finite domain D . Given a t -ary operation f (a function from D^t to D), f can be naturally extended to a t -ary operation over vectors of the same size. Let (v^1, \dots, v^t) be a t -tuple of vectors of size n , f acts coefficient-wise on it, that is for all $i \leq n$, $f(v^1, \dots, v^t)_i = f(v_i^1, \dots, v_i^t)$.

Definition 3. Let \mathcal{F} be a finite set of operations over D . Let \mathcal{S} be a set of vectors over D . Let $\mathcal{F}^i(\mathcal{S}) = \{f(v_1, \dots, v_t) \mid v_1, \dots, v_t \in \mathcal{F}^{i-1}(\mathcal{S}) \text{ and } f \in \mathcal{F}\} \cup \mathcal{F}^{i-1}(\mathcal{S})$ and $\mathcal{F}^0(\mathcal{S}) = \mathcal{S}$. The closure of \mathcal{S} by \mathcal{F} is $Cl_{\mathcal{F}}(\mathcal{S}) = \cup_i \mathcal{F}^i(\mathcal{S})$.

Notice that $Cl_{\mathcal{F}}(\mathcal{S})$ is also the smallest set which contains \mathcal{S} and which is closed by the operations of \mathcal{F} . The set $Cl_{\mathcal{F}}(\mathcal{S})$ is invariant under the operations of \mathcal{F} : these operations are called *polymorphisms* of the set $Cl_{\mathcal{F}}(\mathcal{S})$, a notion which comes from universal algebra.

As an illustration, assume that $D = \{0, 1\}$ and that $\mathcal{F} = \{\vee\}$. Then the elements of \mathcal{S} can be seen as subsets of $[n]$ (each vector of size n is the characteristic vector of a subset of $[n]$) and $Cl_{\{\vee\}}(\mathcal{S})$ is the closure by union of all sets in \mathcal{S} . Let $\mathcal{S} = \{\{1, 2, 4\}, \{2, 3\}, \{1, 3\}\}$ then

$$Cl_{\{\vee\}}(\mathcal{S}) = \{\{1, 2, 4\}, \{1, 2, 3, 4\}, \{2, 3\}, \{1, 3\}, \{1, 2, 3\}\}.$$

Note that $Cl_{\{\vee\}}(\mathcal{S})$ is indeed closed by union, that is \vee is a polymorphism of $Cl_{\{\vee\}}(\mathcal{S})$.

The problem we try to solve in this article, for all set of operations \mathcal{F} over D , is `ENUMCLOSURE $_{\mathcal{F}}$` : given a set of vectors \mathcal{S} list the elements of $Cl_{\mathcal{F}}(\mathcal{S})$. We always denote the size of the vectors of \mathcal{S} by n and the cardinality of \mathcal{S} by m . A naive saturation algorithm solves `ENUMCLOSURE $_{\mathcal{F}}$` : the elements of \mathcal{F} are applied to all tuples of \mathcal{S} and any new element is added to \mathcal{S} , the algorithm stops when no new element can be produced. If the largest arity of an operation in \mathcal{F} is t then this algorithm adds a new element to a set of k elements (or stops) in time $O(k^t n)$, therefore it is in `IncP`.

The aim of this article is to find better algorithms to solve `ENUMCLOSURE $_{\mathcal{F}}$` , more specifically polynomial delay algorithms with polynomial memory. To do that, let us introduce the extension problem associated to a set of operations \mathcal{F} , denoted by `EXTCLOSURE $_{\mathcal{F}}$` : given \mathcal{S} a set of vectors of size n , and a vector v of size $l \leq n$, is there a vector $v' \in Cl_{\mathcal{F}}(\mathcal{S})$ such that $v'_{[l]} = v$. Solving this problem in polynomial time yields a polynomial delay algorithm using backtrack search by Proposition 2. In fact, we need only to solve the simpler membership problem, denoted by `MEMBERSHIP $_{\mathcal{F}}$` : given \mathcal{S} a set of vectors of size n , and a vector v of size n does v belongs to $Cl_{\mathcal{F}}(\mathcal{S})$?

Proposition 4. *If `MEMBERSHIP $_{\mathcal{F}}$` \in P then `ENUMCLOSURE $_{\mathcal{F}}$` \in DelayP.*

Proof: `EXTCLOSURE $_{\mathcal{F}}$` can be reduced to `MEMBERSHIP $_{\mathcal{F}}$` . Indeed, given a vector v of size l , because the operations of \mathcal{F} act coordinate-wise, the two following predicates are equivalent:

- $\exists v' \in Cl_{\mathcal{F}}(\mathcal{S})$ such that $v'_{[l]} = v$
- $v \in Cl_{\mathcal{F}}(\mathcal{S}_{[l]})$

Therefore if $\text{MEMBERSHIP}_{\mathcal{F}} \in \text{P}$ then we have also $\text{EXTCLOSURE}_{\mathcal{F}} \in \text{P}$. We use Proposition 2 to conclude. \square

We have introduced an infinite family of problems, whose complexity we want to determine. Several families of operations may always produce the same closure. To deal with that, we need to introduce the notion of functional clone. We write clone instead of functional clone in the rest of the paper since there is no ambiguity in our context.

Definition 5. Let \mathcal{F} be a finite set of operations over D , the functional clone generated by \mathcal{F} , denoted by $\langle \mathcal{F} \rangle$, is the set of operations obtained by any composition of the operations of \mathcal{F} and of the projections $\pi_k^n : D^n \rightarrow D$ defined by $\pi_k^n(x_1, \dots, x_n) = x_k$.

This notion is useful, because two sets of functions which generate the same clone applied to the same set produce the same closure.

Lemma 6. For all set of operations \mathcal{F} and all set of vectors \mathcal{S} , $Cl_{\mathcal{F}}(\mathcal{S}) = Cl_{\langle \mathcal{F} \rangle}(\mathcal{S})$.

The number of clones over D is infinite even when D is the boolean domain (of size 2). However, in this case the clones form a countable lattice, called Post's lattice [32]. Moreover there is a *finite* number of well described clones plus a few number of infinite families of very regular clones.

3 The Boolean Domain

In this part, we prove the main theorem on the complexity of $\text{MEMBERSHIP}_{\mathcal{F}}$, when the domain is boolean. An instance of one such problem, denoted by \mathcal{S} , will be indifferently seen as a set of vectors of size n or a set of subsets of $[n]$.

Theorem 7. Let \mathcal{F} be any fixed finite set of operations over the boolean domain, then $\text{MEMBERSHIP}_{\mathcal{F}} \in \text{P}$ and $\text{ENUMCLOSURE}_{\mathcal{F}} \in \text{DelayP}$.

There is also a uniform version of $\text{MEMBERSHIP}_{\mathcal{F}}$, where \mathcal{F} is given as input. It turns out that this problem is NP-hard as proven in Section 3.5.

To prove Theorem 7, we prove that $\text{MEMBERSHIP}_{\mathcal{F}} \in \text{P}$, for each clone \mathcal{F} of the Post's lattice. Since the Post's lattice contains many classes, we first show that many of them are equivalent with regard to the complexity of $\text{MEMBERSHIP}_{\mathcal{F}}$: for some \mathcal{F} the problem $\text{MEMBERSHIP}_{\mathcal{F}}$ can be reduced to $\text{MEMBERSHIP}_{\mathcal{G}}$ where \mathcal{G} is another clone obtained from a simple transformation of \mathcal{F} . This reduces the number of cases we need to consider.

To an operation f we can associate its dual \bar{f} defined by $\bar{f}(s_1, \dots, s_t) = \neg f(\neg s_1, \dots, \neg s_t)$. If \mathcal{F} is a set of operations, $\bar{\mathcal{F}}$ is the set of duals of operation in \mathcal{F} . We denote by $\mathbf{0}$ and $\mathbf{1}$ the constant functions which always return 0 and 1. By a slight abuse of notation, we also denote by $\mathbf{0}$ the all zero vector and by $\mathbf{1}$ the all one vector.

Proposition 8. The following problems can be reduced to $\text{MEMBERSHIP}_{\mathcal{F}}$ by a polynomial time parsimonious reduction:

1. $\text{MEMBERSHIP}_{\mathcal{F} \cup \{\mathbf{0}\}}$, $\text{MEMBERSHIP}_{\mathcal{F} \cup \{\mathbf{1}\}}$, $\text{MEMBERSHIP}_{\mathcal{F} \cup \{\mathbf{0}, \mathbf{1}\}}$
2. $\text{MEMBERSHIP}_{\bar{\mathcal{F}}}$
3. $\text{MEMBERSHIP}_{\mathcal{F} \cup \{\neg\}}$ when $\mathcal{F} = \bar{\mathcal{F}}$

Proof: The reductions follow easily from these observations:

1. $Cl_{\mathcal{F} \cup \{f\}}(\mathcal{S}) = Cl_{\mathcal{F}}(\mathcal{S} \cup \{f\})$ for $f = \mathbf{0}$ or $f = \mathbf{1}$ and $\mathcal{S} \neq \emptyset$.
2. $Cl_{\overline{\mathcal{F}}}(\mathcal{S}) = \overline{Cl_{\mathcal{F}}(\overline{\mathcal{S}})}$ where $\overline{\mathcal{S}}$ denotes the set of negation of vectors in \mathcal{S} .
3. $Cl_{\mathcal{F} \cup \{\neg\}}(\mathcal{S}) = Cl_{\mathcal{F}}(\mathcal{S} \cup \overline{\mathcal{S}})$ since for every $f \in \mathcal{F}$, there exists $g \in \mathcal{F}$ such that $\neg f(v_1, \dots, v_t) = \overline{f}(\neg v_1, \dots, \neg v_t) = g(\neg v_1, \dots, \neg v_t)$.

□

In Figure 1, we represent the clones which cannot be reduced to another one using Proposition 8 and their bases. We also represent BF the clone of all function since it is useful in several reductions presented in the rest of the section. For a modern presentation of all boolean clones, their bases and the Post's lattice see [4, 34]. Each clone correspond to a case to settle in our proof of Theorem 7. We can further group clones by the algorithm used to solve their associated enumeration problems and we represent these by ellipses enclosing similar clones in Figure 1. Each ellipse in Figure 1 corresponds to a subsection of this section.

To further simplify the enumeration problems we solve, we remark the following. When there is an index i such that the value of v_i for all elements v to be enumerated is determined by the value of another index or is not constrained, then we can project it out. The formal statement of this fact is given in the next proposition, and the proof is by obvious polynomial delay reductions. Note that these polynomial delay reductions decrease the size of the instance and incurs an overhead for each solution which is constant. Since these reductions can be applied at most once per index, the general overhead for each solution is bounded by the size of the solution.

Proposition 9. *The problem $\text{ENUMCLOSURE}_{\mathcal{F}}$ can be reduced by a polynomial delay reduction to $\text{ENUMCLOSURE}_{\mathcal{F}}$ where the instances \mathcal{S} do not satisfy these two properties:*

- $\exists i, j \in [n], \exists f : \{0, 1\} \rightarrow \{0, 1\}, \forall v \in Cl_{\mathcal{F}}(\mathcal{S}), v_i = f(v_j)$
- $\exists i \in [n], \forall v \in Cl_{\mathcal{F}}(\mathcal{S}), \exists v' \in Cl_{\mathcal{F}}(\mathcal{S}), v_{[n] \setminus i} = v'_{[n] \setminus i}$ and $v_i = \neg v'_i$

Proof: In the two cases, the reduction is to $\text{ENUMCLOSURE}_{\mathcal{F}}$ on the instance $\mathcal{S}_{[n] \setminus i}$. Then each solution v is extended by $f(v_j)$ in the first case or by 0 and 1 in the second case. □

3.1 Conjunction

We first study one of the simplest clones: $E_2 = \langle \wedge \rangle$. We give an elementary proof that $\text{MEMBERSHIP}_{E_2} \in \text{P}$, then we explain how to obtain a good delay for ENUMCLOSURE_{E_2} . For a binary vector v , let us denote by $\mathbb{0}(v)$ (resp. $\mathbb{1}(v)$) the set of indices i for which $v_i = 0$ (resp. $v_i = 1$).

Proposition 10. $\text{MEMBERSHIP}_{E_2} \in \text{P}$.

Proof: Let \mathcal{S} be a set of boolean vectors, if we apply \wedge to a couple of vectors in \mathcal{S} it produces the intersection of two vectors when seen as sets. Since the intersection operation is associative and commutative, $Cl_{E_2}(\mathcal{S})$ is the set of arbitrary intersections of elements of \mathcal{S} . Let v be a vector and let \mathcal{S}_1 be the set $\{w \in \mathcal{S} \mid w_{\mathbb{1}(v)} = \mathbf{1}\}$. Assume now that v can be obtained as an intersection of elements v_1, \dots, v_t , those

Clone	Base
I_2	\emptyset
L_2	$x + y + z$
L_0	$x + y$
E_2	\wedge
S_{10}	$x \wedge (y \vee z)$
S_{10}^2	$maj, x \wedge (y \vee z)$
S_{10}^k	$Th_k^{k+1}, k \geq 3$
S_{12}	$x \wedge (y \rightarrow z)$
S_{12}^k	$Th_k^{k+1}, x \wedge (y \rightarrow z)$
D_2	maj
D_1	$maj, x + y + z$
M_2	\vee, \wedge
R	$x ? y : z$
R_0	$\vee, +$
BF	\vee, \neg




Fig. 1: The reduced Post's lattice, upward edges represent inclusions of clones

elements must be in \mathcal{S}_1 because of the monotonicity of the intersection for the inclusion. On the other hand, by definition of \mathcal{S}_1 , v is contained in $\bigcap_{w \in \mathcal{S}_1} w$. Therefore, $v \in Cl_{E_2}(\mathcal{S})$ if and only if $v = \bigcap_{w \in \mathcal{S}_1} w$.

This intersection can be computed in time $O(mn)$ which concludes the proof. \square

By Proposition 2, we can turn the algorithm for $MEMBERSHIP_{E_2}$ into an enumeration algorithm for $ENUMCLOSURE_{E_2}$ with delay $O(mn^2)$. We explain in the next proposition how to reduce this delay to $O(mn)$.

Proposition 11. *There is an algorithm solving $ENUMCLOSURE_{E_2}$ with delay $O(mn)$.*

Proof: We use the backtrack search described in Proposition 2 but we maintain data structures which allow to decide $MEMBERSHIP_{E_2}$ quickly. Let \mathcal{S} be the input set of m vectors of size n . During the traversal of the tree we update the partial solution p , represented by an array of size n which stores whether $p_i = 1$, $p_i = 0$ or is yet undefined.

A vector v of \mathcal{S} is compatible with the partial solution if $\mathbb{1}(p) \subseteq \mathbb{1}(v)$. We maintain an array $COMP$ indexed by the sets of \mathcal{S} , which stores whether each vector of \mathcal{S} is compatible or not with the current

partial solution. Finally we update an array $COUNT$, such that $COUNT[i]$ is the number of compatible vectors $v \in \mathcal{S}$ such that $v_i = 0$. Notice that a partial solution p can be extended into a vector of $Cl_{E_2}(\mathcal{S})$ if and only if for all $i \in \mathbb{0}(p)$ $COUNT[i] > 0$, the solution is then the intersection of all compatible vectors.

At each step of the traversal, we select an index i such that p_i is undefined and we set first $p_i = 0$ then $p_i = 1$. When we set $p_i = 0$, there is no change to do in $COUNT$ and $COMP$ and we can check whether this extended partial solution is correct by checking if $COUNT[i] > 0$ in constant time. When we set $p_i = 1$, we need to update $COMP$ by removing from it every vector v such that $v_i = 0$. Each time we remove such a vector v , we decrement $COUNT[j]$ for all j such that $v_j = 0$. If there is a j such that $COUNT[j]$ is decremented to 0 then the extension of p by $p_i = 1$ is not possible.

When we traverse a whole branch of the tree of partial solutions during the backtrack search, we set $p_i = 1$ for each i at most once and then we need to remove each vector from $COMP$ at most once. Therefore the total number of operations we do to maintain $COMP$ and $COUNT$ is $O(mn)$ and so is the delay. \square

We now relate the problem $ENUMCLOSURE_{E_2}$ to $ENUMDNF$ which is the problem of listing all satisfying assignments of a formula in disjunctive normal form. The problem $ENUMDNF$ appears in several contexts such as the enumeration of satisfying assignments of an existential formula with second order variables [15] or knowledge compilation [10]. The best algorithms to solve $ENUMDNF$ have a delay of $O(mn)$ where n is the number of variables and m the number of clauses. The following question is open and seems hard to solve:

Is there an algorithm for $ENUMDNF$ with delay polynomial in n only or which has a sublinear dependency on m ? In this paper we use $ENUMMONDNF$, the variant with only positive variables, as a *hard* problem, that is we conjecture it cannot be solved with a delay better than $O(m)$. Then, in the spirit of fine grained complexity [42], it can be used to prove conditional lower bounds on the delay of other enumeration problems.

Note that if we ask for all subsets of the sets in \mathcal{S} instead of all intersections, we exactly get the problem of enumerating the solutions of a monotone DNF formula. Moreover, the algorithm of Proposition 11 is similar to one used to generate the solutions of a DNF formula or of a monotone CNF formula [31]. We now show that we can reduce the problem $ENUMMONDNF$ to $ENUMCLOSURE_{E_2}$. Moreover, an instance of $ENUMMONDNF$ with n variables and m clauses is transformed into an instance of $ENUMCLOSURE_{E_2}$ with mn vectors of size n . Therefore any improvement in the delay for $ENUMCLOSURE_{E_2}$ yields a better delay for $ENUMMONDNF$.

Proposition 12. *There is a polynomial delay reduction from $ENUMMONDNF$ to $ENUMCLOSURE_{E_2}$.*

Proof: First recall that by the second point of Proposition 8, $ENUMCLOSURE_{\{\vee\}}$ can be reduced to $ENUMCLOSURE_{E_2}$ with an instance of the exact same size. Let $\phi \equiv \bigvee_{i=1}^m C_i$ be a monotone DNF formula over the variables $X = \{x_0, \dots, x_{n-1}\}$. The C_i are clauses over positive variables and can be seen as the set of their variables. To C_i we associate the vectors $v^{i,j}$ which are equal to the characteristic vector of C_i except on the coordinate j where it is equal to 1. Let $\mathcal{S} = \{v^{i,j} \mid i \in [m], j \in [n]\}$. Notice that a solution of ϕ is any assignment which contains all variables of some clause C_i . The set $Cl_{\{\vee\}}(\mathcal{S})$ is also the set of all sets containing some clause which proves the reduction from $ENUMMONDNF$ to $ENUMCLOSURE_{\{\vee\}}$ and thus to $ENUMCLOSURE_{D_2}$. \square

3.2 Algebraic operations

We first deal with the clone $L_0 = \langle + \rangle$ where $+$ is the boolean addition (or equivalently the boolean operation *XOR*). Note that $Cl_{L_0}(\mathcal{S})$ is the vector space generated by the vectors in \mathcal{S} . Seen as an operation on sets, $+$ is the symmetric difference of the two sets.

Proposition 13. $MEMBERSHIP_{L_0} \in P$.

Proof: Let \mathcal{S} be the set of input vectors, let v be a vector and let A be the matrix whose rows are the elements of \mathcal{S} . The vector v is in $Cl_{L_0}(\mathcal{S})$ if and only if there is a solution to $Ax = v$. Solving a linear system over \mathbb{F}_2 can be done in polynomial time which proves the proposition. \square

The previous proposition yields a polynomial delay algorithm by applying Proposition 2. One can get a better delay, by computing in polynomial time a maximal free family M of \mathcal{S} , which is a basis of $Cl_{L_0}(\mathcal{S})$. The basis M is a succinct representation of $Cl_{L_0}(\mathcal{S})$. One can generate all elements of $Cl_{L_0}(\mathcal{S})$ by going over all possible subsets of elements of M and summing them. The subsets can be enumerated in constant time by using Gray code enumeration (see [25]). The sum can be done in time n by adding a single vector since two consecutive sets differ by a single element in the Gray code order. Therefore we have, after the polynomial time computation of M , an enumeration with delay $O(n)$. If one allows to output the elements represented in the basis M , the algorithm even has constant delay.

The closure by the clone L_2 (generated by the sum modulo two of three elements) corresponds to an affine space rather than a vector space. Hence, we can easily extend the previous result to L_2 .

Proposition 14. $MEMBERSHIP_{L_2} \in P$.

Proof: Since the sum of three elements is associative and commutative, the vectors in $Cl_{L_2}(\mathcal{S})$ are the sum of an odd number of vectors in \mathcal{S} . In other words $v \in Cl_{L_2}(\mathcal{S})$ if and only if there is a x such that $Ax = v$ and the Hamming weight of x is odd. One can compute a basis B of the vector space of the solutions to the equation $Ax = v$. If all elements of B have even Hamming weight, then their sums also have even Hamming weight. Therefore $v \in Cl_{L_2}(\mathcal{S})$ if and only if there is an element in B with odd Hamming weight, which can be decided in polynomial time. \square

3.3 Boolean algebras

In this subsection, we deal with the largest clones of our reduced Post lattice: $M_2 = \langle \wedge, \vee \rangle$, $BF = \langle \vee, \neg \rangle$, $R_0 = \langle \vee, + \rangle$ and $R = \langle x ? y : z \rangle$, where $x ? y : z$ is the "if then else" operator which is equal to $(\neg x \vee y) \wedge (x \vee z)$.

Proposition 15. $MEMBERSHIP_{M_2} \in P$.

Proof: Let \mathcal{S} be a vector set and for all $i \in [n]$, let $x^i = \bigwedge_{v \in \mathcal{S} \text{ s.t. } v_i=1} v$, we call x^i an atom. We show that a vector u belongs to $Cl_{M_2}(\mathcal{S})$ if and only if $u = \bigvee_{i \in \mathbb{1}(u)} x^i$. By definition of $Cl_{M_2}(\mathcal{S})$, it contains the atoms since they are intersections of elements in \mathcal{S} and it contains their unions $u = \bigvee_{i \in \mathbb{1}(u)} x^i$. Since the intersection distributes with the union, we can write any element v in $Cl_{M_2}(\mathcal{S})$ as $v = \bigvee_i \left(\bigwedge_j s^{i,j} \right)$ where $s^{i,j} \in \mathcal{S}$. Hence it is enough to show that $v^i = \bigwedge_j s^{i,j}$ is an union of atoms. Let k be an index

such that $v_k^i = 1$, then for all vectors $s^{i,j}$, $s_k^{i,j} = 1$. It implies that seen as a set, x^k is included in v^i and that $v^i = \bigvee_{k \in \mathbb{1}(v^i)} x^k$ which proves the characterization of $Cl_{M_2}(\mathcal{S})$.

We can compute the atoms in time $O(mn^2)$, and then to decide whether $v \in Cl_{M_2}(\mathcal{S})$, one must check whether v contains all atoms x^i such that $v_i = 1$ in time $O(n^2)$ which proves the proposition. \square

Applying Proposition 2, we get an enumeration algorithm with delay $O(mn^3)$. Moreover, $Cl_{M_2}(\mathcal{S}) = Cl_{<v>}(\{x^i\}_{i \in [n]})$. Recall that $ENUMCLOSURE_{<v>}$ can be reduced to $ENUMCLOSURE_{E_2}$ by Proposition 8 and that $ENUMCLOSURE_{E_2}$ has an algorithm with delay $O(mn)$ thanks to Proposition 11. Hence the n atoms can be precomputed and their union generated with delay $O(n^2)$ since here $m = n$. We can do better by using the inclusion structure of the x^i 's to obtain a $O(n)$ delay.

Proposition 16. $ENUMCLOSURE_{M_2}$ can be solved with delay $O(n)$.

Proof: Let \mathcal{S} be the input. We first build for all i , $x^i = \bigwedge_{v \in \mathcal{S}, v_i = 1} v$. Inclusion is a partial order between the elements x^i seen as sets. We consider a linear extension T of that partial order, i.e. a total order obtained by ordering all the incomparable pairs. We then generate all elements of $Cl_{M_2}(\mathcal{S})$ by a Hill climbing algorithm: we go from one solution to another by adding a single x^i . Let v be the current solution, we maintain a list L ordered according to T of the indices i of v such that $v_i = 0$. At each step we select i the smallest element of L and we set $v_j = 1$ and remove j from L for all $j \in x^i$. This produces a new solution in time $O(n)$. We then recursively call the algorithm on this new solution and list. When the recursive call is finished, we call the algorithm on v and $L \setminus \{i\}$.

This algorithm is correct, because the solutions generated in the two recursive calls are disjoint. Indeed, in the second call v_i will always be 0, because all indices in L are bigger than i in T . It means that x^j for $j \in L$ is either smaller or incomparable. Since x^i is the smallest element with $x_i^i = 1$ it implies that $x_i^j = 0$. \square

The problem $ENUMCLOSURE_{BF}$, which can be reduced to $ENUMCLOSURE_{M_2}$ by Proposition 17 is in fact even easier to enumerate. Let \mathcal{S} be a set of vectors, let $X^i = \{v \mid v \in \mathcal{S}, v_i = 1\} \cup \{-v \mid v \in \mathcal{S}, v_i = 0\}$ and let $x^i = \bigwedge_{v \in X^i} v$. The set $Cl_{BF}(\mathcal{S})$ is in fact a boolean algebra, whose atoms are the disjoint elements x^i . Indeed, either $x_{i,j}^i = x_{i,j}^j$ and they are equal or $\mathbb{1}_{x^i} \cap \mathbb{1}_{x^j} = \emptyset$. Let $A = \{x^i \mid i \in [n]\}$, two distinct unions of elements in A produce distinct elements. Hence by enumerating all possible subsets of A with a Gray code, we can generate $Cl_{BF}(\mathcal{S})$ with delay $O(n)$ or $O(1)$ if groups of always equal coefficients are represented by a single coefficient.

We now prove that the closures by the clones R and R_0 are equal to the closure by BF up to some coefficients which are fixed to 0 or 1, thus they are as easy to enumerate as stated in the next proposition.

Proposition 17. The problems $ENUMCLOSURE_R$ and $ENUMCLOSURE_{R_0}$ can be reduced to $ENUMCLOSURE_{BF}$ by a polynomial delay reduction.

Proof: First notice that both R and R_0 contains M_2 which means that they contain \vee and \wedge . By Proposition 9 we can assume that for all i , there is $u \in \mathcal{S}$ such that $u_i = 1$ and v such that $v_i = 0$. Since \vee is in both clones the vector $\mathbf{1}$ is in their closure. Since the clones also contain \wedge , the vector $\mathbf{0}$ is also in their closures. Now notice that $x \cdot \mathbf{0} : \mathbf{1} = \bar{x}$ and that $x + \mathbf{1} = \bar{x}$. As a consequence we have $Cl_R(\mathcal{S}) = Cl_{R_0}(\mathcal{S}) = Cl_{BF}(\mathcal{S})$ which proves the proposition. \square

3.4 Limits of the infinite parts

We deal with the two infinite hierarchies of clones in Subsection 3.5. However, the method for the hierarchy does not directly apply to the clones which are the limits of these hierarchies: $S_{10} = \langle x \wedge (y \vee z) \rangle$ and $S_{12} = \langle x \wedge (y \rightarrow z) \rangle$. We show here how we can see them as a union of closures by M_2 or BF , which complexity wise makes them similar to ENUMCLOSURE_{E_2} . As a consequence, they have a polynomial delay algorithm but the delay depends on m the size of \mathcal{S} and a reduction from ENUMMONDNF suggests it is hard to get rid of this dependency.

Lemma 18. *Let \mathcal{S} be a set of vectors then:*

- $Cl_{S_{10}}(\mathcal{S}) = \bigcup_{s \in \mathcal{S}} Cl_{M_2}(\mathcal{S}_{\mathbb{1}(s)})$
- *there is a set of indices I such that if $v \in Cl_{S_{12}}(\mathcal{S}_I)$, $v_I = \mathbf{1}$ and*

$$Cl_{S_{12}}(\mathcal{S}_I) = \bigcup_{s \in \mathcal{S}} Cl_{BF}(\mathcal{S}_{I \cap \mathbb{1}(s)})$$

where the elements of $Cl_{M_2}(\mathcal{S}_{\mathbb{1}(s)})$ and of $Cl_{BF}(\mathcal{S}_{\mathbb{1}(s)})$ are extended by zeros outside of $\mathbb{1}(s)$.

Proof: By induction on the elements of $Cl_{S_{10}}(\mathcal{S})$ (respectively $Cl_{S_{12}}(\mathcal{S})$), one proves that their support are always contained in the support of some element of \mathcal{S} .

Hence, for $v \in Cl_{S_{10}}(\mathcal{S})$, we assume that $v \subseteq s \in \mathcal{S}$. Notice that $x \wedge (y \vee y)$ is equal to $x \wedge y$, therefore we have the operation \wedge in S_{10} . If $v_{\mathbb{1}(s)} \in Cl_{S_{10}}(\mathcal{S}_{\mathbb{1}(s)})$, it implies that there is a $v' \in Cl_{S_{10}}(\mathcal{S})$ such that $v'_{\mathbb{1}(s)} = v_{\mathbb{1}(s)}$. Therefore $v' \wedge s = v$ and $v \in Cl_{S_{10}}(\mathcal{S})$. We have proved $v \in Cl_{S_{10}}(\mathcal{S})$ if and only if $v_{\mathbb{1}(s)} \in Cl_{S_{10}}(\mathcal{S}_{\mathbb{1}(s)})$ since the other direction is just the projection on $\mathbb{1}(s)$.

As a consequence, if we consider $\mathcal{S}_{\mathbb{1}(s)}$, it contains the vector $\mathbf{1}$. We can simulate \vee , since $\mathbf{1} \wedge (x \vee y) = x \vee y$. Therefore $Cl_{M_2}(\mathcal{S}_{\mathbb{1}(s)}) \subseteq Cl_{S_{10}}(\mathcal{S}_{\mathbb{1}(s)})$ and the reverse inclusion is clear since $M_2 = \langle \vee, \wedge \rangle$ contains S_{10} which proves the first item of the lemma.

Assume that there is a coefficient i such that for all $s \in \mathcal{S}$, $s_i = 1$, then for all $v \in Cl_{S_{12}}(\mathcal{S})$, $v_i = 1$. We denote by I the set of such coefficients and by the previous remark, for all v in $Cl_{S_{12}}(\mathcal{S})$, we have $v_I = \mathbf{1}$. By definition, for each $i \in \bar{I}$, there is a $s \in \mathcal{S}$ with $s_i = 0$. Since $x \wedge (x \rightarrow y) = x \wedge y$, the vector $\mathbf{0}$ is in $Cl_{S_{12}}(\mathcal{S}_{\bar{I}})$.

As in the first part of the proof, $v \in Cl_{S_{12}}(\mathcal{S}_{\bar{I}})$ if and only if there is $v \subseteq s \in \mathcal{S}_{\bar{I}}$, such that $v_{\mathbb{1}(s)} \in Cl_{S_{12}}(\mathcal{S}_{\bar{I} \cap \mathbb{1}(s)})$. By definition, $s_{\mathbb{1}(s)} = \mathbf{1}$. Thus we can simulate the negation in $Cl_{S_{12}}(\mathcal{S}_{\bar{I} \cap \mathbb{1}(s)})$ by $\mathbf{1} \wedge (x \rightarrow \mathbf{0}) = \neg x$. We also simulate the disjunction by $\mathbf{1} \wedge (\neg x \rightarrow y) = x \vee y$. Therefore $Cl_{S_{12}}(\mathcal{S}_{\bar{I} \cap \mathbb{1}(s)})$ contains $Cl_{BF}(\mathcal{S}_{\bar{I} \cap \mathbb{1}(s)})$. The reverse inclusion follows from the fact that BF contains S_{12} which proves the second item of the lemma. \square

As a consequence of this structural characterization, is it easy to decide the problems $\text{MEMBERSHIP}_{S_{10}}$ and $\text{MEMBERSHIP}_{S_{12}}$ in time $O(mn)$. Therefore using a backtrack search, we obtain enumeration algorithms with delay $O(mn^2)$. However, using the structures of $Cl_{S_{10}}(\mathcal{S})$ and $Cl_{S_{12}}(\mathcal{S})$ described in Lemma 18 we obtain better algorithms described in the following proposition.

Proposition 19. *The problems $\text{ENUMCLOSURE}_{S_{10}}$ and $\text{ENUMCLOSURE}_{S_{12}}$ can be solved with delay $O(mn)$.*

Proof: We first prove the result for $\text{ENUMCLOSURE}_{S_{12}}$. By Lemma 18, there is a set I , such that $Cl_{S_{12}}(\mathcal{S}_{\bar{I}}) = \cup_{s \in \mathcal{S}} Cl_{BF}(\mathcal{S}_{\bar{I} \cap \mathbb{1}(s)})$ where the elements of $\mathcal{S}_{\mathbb{1}(s)}$ are extended by zeros. The coefficients in I are always one and do not matter for an enumeration algorithm. Moreover, ENUMCLOSURE_{BF} can be solved with delay $O(n)$. In particular, we can output the solutions in lexicographic order, since they are union of disjoint atoms. If the elements of $Cl_{BF}(\mathcal{S}_{\mathbb{1}(s)})$ are generated for each s we obtain all the elements we want to output but with *redundancies*. We can obtain an algorithm without redundancies by simulating an enumeration of the elements of $Cl_{BF}(\mathcal{S}_{\mathbb{1}(s)})$ for all s and always outputting the smallest one. This method is described in [38] (Lemma 3, Chapter 2) and the delay is the sum of delays of the enumeration procedure for each s , that is $O(mn)$.

The proof for $\text{ENUMCLOSURE}_{S_{10}}$ is similar and follows from Proposition 16 which shows that the elements of $Cl_{M_2}(\mathcal{S})$ can be enumerated with delay $O(n)$ in lexicographic order. \square

Proposition 20. *There is a polynomial delay reduction from ENUMMONDNF to $\text{ENUMCLOSURE}_{S_{10}}$ and $\text{ENUMCLOSURE}_{S_{12}}$.*

Proof: Let $\phi = \bigvee_{i \in [m]} C_i$ where each C_i is a conjunction of negative variables in $X = \{x_0, \dots, x_{n-1}\}$. We assume that no variable appear in all clauses otherwise we could do a trivial polynomial delay reduction to get rid of it. To a clause C_i we associate the vector v^i such that $v_j^i = 0$ if and only if $x_j \in C_i$. Let e^i denote the vector which has the coefficient $e_i^i = 1$ and its other coefficients are zero. We define $\mathcal{S} = \{v^i \mid i \in [m]\} \cup \{e^i \mid i \in [n]\}$. Note that there is natural bijection between assignments of ϕ and vectors of size n . Let $\phi(X)$ stand for the set of satisfying assignments of ϕ . Note that $\phi(X) = \cup_{i \in [m]} C_i(X)$. And by definition $C_i(X)$ is equal to the set of all assignments where the variables in C_i are equal to 0. Since both S_{10} and S_{12} contain \wedge , $C_i(X) = Cl_{S_{10}}(\{v^i\} \cup \{e^j \mid j \in [n]\})_{\mathbb{1}_{v^i}} = Cl_{S_{12}}(\{v^i\} \cup \{e^j \mid j \in [n]\})_{\mathbb{1}_{v^i}}$. Since $C_i(X)$ contains $Cl_{S_{10}}(\mathcal{S}_{\mathbb{1}_{v^i}})$ and $Cl_{S_{12}}(\mathcal{S}_{\mathbb{1}_{v^i}})$ it is equal to these sets. Since no variable appear in all clauses, $Cl_{S_{10}}(\mathcal{S}_{\mathbb{1}_{e^i}})$ and $Cl_{S_{12}}(\mathcal{S}_{\mathbb{1}_{e^i}})$ are in $\phi(X)$ for all i . Hence, using the structure of $Cl_{S_{10}}(\mathcal{S})$ and $Cl_{S_{12}}(\mathcal{S})$ given in Lemma 18, we obtain $Cl_{S_{10}}(\mathcal{S}) = Cl_{S_{12}}(\mathcal{S}) = \phi(X)$. \square

Since the reduction transforms an instance of ENUMMONDNF with m clauses and n variables into an instance of $\text{ENUMCLOSURE}_{S_{10}}$ or $\text{ENUMCLOSURE}_{S_{12}}$ with $m+n$ vectors of size n any algorithm improving the delay of $\text{ENUMCLOSURE}_{S_{10}}$ or $\text{ENUMCLOSURE}_{S_{12}}$ would improve the one of ENUMMONDNF .

3.5 Majority and threshold

An operation f is a *near unanimity* of arity k if it satisfies $f(x_1, x_2, \dots, x_k) = x$ for each k -tuple with at most one element different from x . The *threshold* function of arity k , denoted by Th_{k-1}^k , is defined by $Th_{k-1}^k(x_1, \dots, x_k)$ is equal to 1 if and only if at least $k-1$ of the elements x_1, \dots, x_k are equal to one. It is the smallest near unanimity operation over the booleans. We use the Baker-Pixley theorem from universal algebra to characterize the closure by any clone which contains a threshold function by its projections of fixed size.

Majority

The threshold function Th_2^3 is the majority operation over three booleans that we denote by maj and the clone it generates is D_2 . We could directly apply the Baker-Pixley theorem in this case, but we felt that for a computer science readership, it would be useful to present a proof of a special case: the characterization of $Cl_{D_2}(\mathcal{S})$. We also show how to obtain the best possible enumeration algorithm for deciding ENUMCLOSURE_{D_2} .

Lemma 21. *Let \mathcal{S} be set of vectors. Then v belongs to $Cl_{D_2}(\mathcal{S})$ if and only if for all $i, j \in [n]$, $i \neq j$, there exists $x \in \mathcal{S}$ such that $x_{i,j} = v_{i,j}$.*

Proof:

(\implies) Given $a, b \in \{0, 1\}$ and $i, j \in [n]$, $i \neq j$, we first show that if for all $v \in \mathcal{S}$, $v_i \neq a$ or $v_j \neq b$ then for all $u \in Cl_{D_2}(\mathcal{S})$, $u_i \neq a$ or $u_j \neq b$. It is sufficient to prove that this property is preserved by applying maj to a vector set i.e. that if \mathcal{S} has this property, then $maj(\mathcal{S})$ has also this property. Let $x, y, z \in \mathcal{S}$, $v := maj(x, y, z)$, and assume for contradiction that $v_{i,j} = (a, b)$. Since $v_i = a$, there is at least two vectors among $\{x, y, z\}$ that are equal to a at index i . Without loss of generality, let x and y be these two vectors. Since for all $u \in \mathcal{S}$, $u_i \neq a$ or $u_j \neq b$, we have $x_j \neq b$ and $y_j \neq b$ and then $v_j \neq b$ which contradicts the assumption. We conclude that if $v \in Cl_{D_2}(\mathcal{S})$, then for all $i, j \in [n]$, there exists $u \in \mathcal{S}$ with $v_{i,j} = u_{i,j}$.

(\impliedby) Let $k \leq n$ and let $a_1, \dots, a_k \in \{0, 1\}$. We show by induction on k , that if for all $i, j \leq k$ there exists $v \in \mathcal{S}$ with $v_i = a_i$ and $v_j = a_j$, then there exists $u \in Cl_{D_2}(\mathcal{S})$ with $u_1 = a_1, u_2 = a_2, \dots, u_k = a_k$. The assertion is true for $k = 2$. Assume it is true for $k - 1$, and let $a_1, \dots, a_k \in \{0, 1\}$. By induction hypothesis there exists a vector $w \in Cl_{D_2}(\mathcal{S})$ with $w_1 = a_1, \dots, w_{k-1} = a_{k-1}$. By hypothesis, for all $i \leq k$ there exists $v^i \in \mathcal{S}$ with $v^i_i = a_i$ and $v^i_k = a_k$. We then construct a sequence of vectors $(u^i)_{i \leq k}$ as follow. We let $u^1 = v^1$ and for all $1 < i < k$, $u^i = maj(w, u^{i-1}, v^i)$. We claim that $u := u^{k-1}$ has the property sought i.e. for all $i \leq k$, $u_i = a_i$. First let prove that for all $i < k$ and for all $j \leq i$, $u^i_j = a_j$. It is true for u^1 by definition. Assume now that the property holds for u^{i-1} , $i < k$. Then, by construction, for all $j \leq i - 1$, we have $u^i_j = a_j$ since $w_j = a_j$ and $u^{i-1}_j = a_j$. Furthermore, we have $u^i_i = maj(w_i, u^{i-1}_i, v^i_i) = a_i$ since $w_i = a_i$ and $v^i_i = a_i$. We conclude that for all $i \leq k - 1$, $u_i = u^{k-1}_i = a_i$.

We claim now that for all $i < k$, $u^i_k = a_k$. It is true for u^1 . Assume it is true for u^{i-1} , $i < k$. Then we have $u^i_k = maj(w_k, u^{i-1}_k, v^i_k)$ which is equal to a_k since $u^{i-1}_k = a_k$ by induction and $v^i_k = a_k$ by definition. We then have $u_i = a_i$ for all $i \leq k$ which concludes the proof. \square

As an immediate consequence we get the following corollary and proposition.

Corollary 22. $MEMBERSHIP_{D_2} \in P$.

Proof: Using Lemma 21, one decides whether a vector v is in $Cl_{D_2}(\mathcal{S})$, by checking for every pair of indices i, j whether there is a vector $w \in \mathcal{S}$ such that $v_{i,j} = w_{i,j}$. \square

The complexity of the algorithm of the previous corollary is $O(mn^2)$, hence applying Prop 2, we get an enumeration algorithm with delay $O(mn^3)$. We explain how to improve this delay in the next proposition.

Proposition 23. $ENUMCLOSURE_{D_2}$ can be solved with delay $O(n^2)$.

Proof: We do a backtrack search and we explain how to efficiently decide $MEMBERSHIP_{D_2}$ during the enumeration. We first precompute for each pair (i, j) all values (a, b) such that there exists $v \in \mathcal{S}$, $v_{i,j} = (a, b)$. When we want to decide whether the vector v of size l can be extended into a solution, it is enough that it satisfies the condition of Lemma 21. Moreover, we already know that $v_{[l-1]}$ satisfies the condition of Lemma 21. Hence we only have to check that the values of $v_{i,l}$ for all $i < l$ can be found in $\mathcal{S}_{i,l}$ which can be done in time $O(l)$. The delay is the sum of the complexity of deciding $MEMBERSHIP_{D_2}$ for each partial solution in a branch: $O(n^2)$. \square

We can yet improve the delay by using an algorithm which does not use the backtrack search of Prop 2. The closure of \mathcal{S} can be represented by a boolean formula $\phi_{\mathcal{S}}$ which describes the obstructions to being in $Cl_{D_2}(\mathcal{S})$. Its variables are $X = \{x_0, \dots, x_{n-1}\}$, therefore its assignments are naturally in bijection with the boolean vectors of size n and we do not distinguish both. Recall that ENUM2SAT is the problem of generating all satisfying assignments of a $2CNF$ formula, that is a formula in conjunctive normal form where each clause is of size at most two.

Proposition 24. $ENUMCLOSURE_{D_2}$ can be solved with delay $O(n)$.

Proof: We prove that $ENUMCLOSURE_{D_2}$ can be reduced to ENUM2SAT by a polynomial delay reduction. To an instance \mathcal{S} , we associate a formula $\phi_{\mathcal{S}}$ in conjunctive normal form. Given $i \neq j$ and $v_{i,j}$ which is not in $\mathcal{S}_{i,j}$, we define the clause $l_i \vee l_j$ where l_k , for $k = i, j$, is x_k if $v_k = 0$ and $\neg x_k$ otherwise. The formula $\phi_{\mathcal{S}}$ is the conjunction of all such clauses. By construction, the formula $\phi_{Cl_{D_2}(\mathcal{S})}$ is a $2CNF$.

Notice that each clause forbids the assignment to take a pair of values which is not present in \mathcal{S} . Therefore a boolean vector v seen as an assignment satisfies $\phi_{\mathcal{S}}$ if and only if every projection of size 2 of v is in \mathcal{S} . By Lemma 21, it means that $v \in Cl_{D_2}(\mathcal{S})$ if and only if v satisfies $\phi_{\mathcal{S}}$.

The proposition is proved since the problem ENUM2SAT can be solved with delay $O(n)$, where n is the number of variables using an algorithm of Feder [18]. \square

We remark that we can apply the idea of Prop 24 to the clones of Subsection 3.3 and to D_1 which all contain the *maj* function. Indeed, closures by these clones are all characterized by their projections of size two, which allows to do a reduction to a $2CNF$ formula. The algorithms we obtain have delay $O(n)$. However, we felt it was relevant to deal with clones of Subsection 3.3 separately, since the algorithms presented in Section 3.3 are different and simpler. In particular, for the clones R_0 and R , the delay is only $O(1)$ if we allow a compact representation of the solutions.

The Baker-Pixley theorem

We now state the Baker-Pixley theorem which generalizes Lemma 21 to all clones which contain near unanimity terms. For more context on this theorem and universal algebra see [9].

Theorem 25 (Baker-Pixley, adapted from [2]). *Let \mathcal{F} be a clone which contains a near unanimity term of arity k , then $v \in Cl_{\mathcal{F}}(\mathcal{S})$ if and only if for each set of indices I of size $k - 1$, $v_I \in Cl_{\mathcal{F}}(\mathcal{S})_I$.*

This allows to settle the case of the two infinite families of clones of our restricted lattice $S_{10}^k = \langle Th_k^{k+1} \rangle$ and $S_{12}^k = \langle Th_k^{k+1}, x \wedge (y \rightarrow z) \rangle$.

Corollary 26. *If a clone \mathcal{F} contains Th_k^{k+1} then $MEMBERSHIP_{\mathcal{F}}$ is solvable in time $O(mn^k)$.*

Proof: Let \mathcal{S} be a set of vectors and let v be a vector. By Theorem 25, $v \in Cl_{\mathcal{F}}(\mathcal{S})$ if and only if for all I of size k , $v_I \in Cl_{\mathcal{F}}(\mathcal{S})_I$. First notice that $Cl_{\mathcal{F}}(\mathcal{S})_I = Cl_{\mathcal{F}}(\mathcal{S}_I)$ because the functions of \mathcal{F} act coefficient-wise on \mathcal{S} . The algorithm generates for each I of size k the set $Cl_{\mathcal{F}}(\mathcal{S}_I)$. For a given I , we first need to build the set \mathcal{S}_I in time m and then the generation of $Cl_{\mathcal{F}}(\mathcal{S}_I)$ can be done in constant time. Indeed, we can apply the classical incremental algorithm to generate the elements in $Cl_{\mathcal{F}}(\mathcal{S}_I)$, and the cardinal of $Cl_{\mathcal{F}}(\mathcal{S}_I)$ only depends on k which is a constant. The time to generate all $Cl_{\mathcal{F}}(\mathcal{S}_I)$ is $O(mn^k)$ and then all the tests can be done in $O(n^k)$. \square

We have proved that the complexity of any closure problem in one of our infinite families is polynomial. We can use the method of Proposition 23 to obtain an algorithm with delay $O(n^k)$ to enumerate the

elements of a set closed by a clone containing a near unanimity function of arity $k + 1$. When the arity is 3, we can use the method of Proposition 24 to obtain an algorithm with the better delay $O(n)$.

The complexity of $\text{MEMBERSHIP}_{\mathcal{F}}$ is increasing with the smallest arity of a near unanimity function in \mathcal{F} . Hence we now investigate the complexity of the *uniform problem* when the clone is given as input. Let CLOSURETRESHOLD be the following problem: given a set \mathcal{S} of vectors and an integer k decide whether the vector $\mathbf{1} \in \text{Cl}_{S_{10}^k}(\mathcal{S})$. It is a restricted version of the uniform problem, but it is already hard to solve because we can reduce the problem HITTINGSET to its complement.

Theorem 27. CLOSURETRESHOLD is coNP-complete.

Proof: First notice that the problem is in coNP since by Theorem 25, the answer to the problem is negative if and only if one can exhibit a subset of indices I of size k such that no element of $\text{Cl}_{S_{10}^k}(\mathcal{S}_I)$ is equal to $\mathbf{1}$. Let us show that the later holds if and only if no element of \mathcal{S}_I is equal to $\mathbf{1}$. We assume that $k \geq 3$ hence $S_{10}^k = \langle \text{Th}_k^{k+1} \rangle$. If no element in \mathcal{S}_I is equal to $\mathbf{1}$, then the application of Th_k^{k+1} to \mathcal{S}_I preserves this property. Indeed, consider $\text{Th}_k^{k+1}(v^1, \dots, v^{k+1})$. Each v^i has a zero coefficient and since there are $k + 1$ such vectors and the vectors are of size k , by the pigeonhole principle, there are i, j, l such that $v_l^i = v_l^j = 0$. This implies that $\text{Th}_k^{k+1}(v^1, \dots, v^{k+1}) \neq \mathbf{1}$.

Let us show that HITTINGSET can be reduced to CLOSURETRESHOLD . Given a hypergraph $\mathcal{H} = (V, \mathcal{E})$ and an integer k , HITTINGSET asks whether there exists a subset $X \subseteq V$ of size k that intersects all the hyperedges of \mathcal{H} . This problem is a classical NP-complete problem [21]. Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph and k be an integer. Let $\bar{\mathcal{H}}$ be the hypergraph on V whose hyperedges are the complement of the hyperedges of \mathcal{H} , and let \mathcal{S} be the set of characteristic vectors of the hyperedges of $\bar{\mathcal{H}}$. Then \mathcal{H} has a transversal of size k if and only if there is a set I of indices of size k such for all $v \in \mathcal{S}_I, v \neq \mathbf{1}$. Indeed, I is a hitting set of \mathcal{H} if for all $E \in \mathcal{E}$, there exists $i \in I$ such that $i \in E$ which implies that $i \notin \bar{E}$ and then the characteristic vector v of \bar{E} is such that $v_i = 0$.

Since the other direction is straightforward, we have proved that there is a set I of indices of size k such that for all $v \in \mathcal{S}_I, v \neq \mathbf{1}$ if and only if there is a set I of indices of size k such that for all $v \in \text{Cl}_{S_{10}^k}(\mathcal{S}_I), v \neq \mathbf{1}$. By Theorem 25, the later property is equivalent to $\mathbf{1} \notin \text{Cl}_{S_{10}^k}(\mathcal{S})$. Therefore we have given a polynomial time reduction from HITTINGSET to the complement of CLOSURETRESHOLD which proves the proposition. \square

In fact, the result is even stronger. We cannot hope to get an FPT algorithm for CLOSURETRESHOLD parametrized by k since HITTINGSET parametrized by the size of the hitting set is $\text{W}[2]$ -complete [19]. It implies that if we want to significantly improve the delay of our enumeration algorithm for the clone S_{10}^k , we should drop the backtrack search since it relies on solving $\text{MEMBERSHIP}_{S_{10}^k}$.

4 Larger Domains

In this section, we try to extend some results of the boolean domain to larger domains, which is the simplest generalization possible. The main problem is that, for domain of size 3 or more, the lattice of clones is uncountable and not as well understood as Post's lattice. We show how we can do a backtrack search for threshold functions and algebraic functions on larger domains. We also prove that for domain of size 3, there is a clone generated by a single binary function f such that $\text{MEMBERSHIP}_{\langle f \rangle}$ is NP-hard.

4.1 Tractable cases

We show here how the cases of near unanimity term (subsection 3.5 and 3.3) and of the addition (subsection 3.2) can be generalized. Using the Baker-Pixley theorem (Theorem 25), we can get an equivalent to Corollary 26 and to Proposition 23 in any domain size.

Corollary 28. *If \mathcal{F} contains a near unanimity operation, then $\text{MEMBERSHIP}_{\mathcal{F}} \in P$.*

Proposition 29. *If \mathcal{F} contains a near unanimity term of arity $k+1$, then $\text{ENUMCLOSURE}_{\mathcal{F}}$ can be solved with delay $O(n^k)$.*

The second tractable case is a generalization of Subsection 3.2. The simplest generalization of the sum over \mathbb{F}_2 is to consider any finite group operation. Using Sim's stabilizer chain it is possible to decide membership in groups given by generators [20, 43].

Proposition 30. *Let f be a group operation over D . Then $\text{MEMBERSHIP}_{\langle f \rangle} \in P$.*

The problem $\text{MEMBERSHIP}_{\mathcal{F}}$ is known in the universal algebra community under the name *subpower membership problem* or SMP. The method for groups also works for extensions of groups by multilinear operations such as ring or modules. For semigroups, some cases are known to be PSPACE-complete and other are in P [8, 37]. There are also works on SMP for clones which satisfy conditions from universal algebras: for expansions of finite nilpotent Maltsev algebras of prime power size or for residually finite algebras with cube terms, SMP is in P [30, 40], while it is hard for some algebras satisfying any strong Maltsev conditions [36].

4.2 A limit to the backtrack search

The last case we would like to extend is the clone generated by the conjunction. A natural generalization is to fix an order on D and to study the complexity of $\text{MEMBERSHIP}_{\langle f \rangle}$ with f monotone. Let f be the function over $D = \{0, 1, 2\}$ defined by $f(x, y) = \min(x + y, 2)$. This function is monotone in each of its arguments. We prove that EXACT3COVER reduces to $\text{MEMBERSHIP}_{\langle f \rangle}$, where EXACT3COVER is the problem of finding an exact cover of a set by subsets of size 3. The next proposition is from [28] and it has been independently proven using the same reduction for semigroups with idempotent elements in Lemma 5.2 of [8].

Proposition 31. *$\text{MEMBERSHIP}_{\langle f \rangle}$ is NP-complete.*

Proof: We reduce EXACT3COVER, which is NP-hard [21], to $\text{MEMBERSHIP}_{\langle f \rangle}$. Let \mathcal{S} be an instance of EXACT3COVER, that is a set of subsets of $[n]$ of size 3. If we consider that the sets in \mathcal{S} are represented by their characteristic vectors, then \mathcal{S} is an instance of $\text{MEMBERSHIP}_{\langle f \rangle}$ and we now prove that $\mathbf{1} \in Cl_{\langle f \rangle}(\mathcal{S})$ if and only if there is an exact cover of \mathcal{S} . First note that f is associative, therefore any element of $Cl_{\langle f \rangle}(\mathcal{S})$ can be written $f(s^1, f(s^2, f(s^3, \dots)))$ with $s^i \in \mathcal{S}$. Since f is also commutative, the order of the operations is not relevant either. Hence we can represent any element $v \in Cl_{\langle f \rangle}(\mathcal{S})$ by a multiset of elements of \mathcal{S} , which combined by f yields v . Notice that all elements in $Cl_{\langle f \rangle}(\mathcal{S})$ can be seen as a multiset containing at most twice an element since $f(s, s) = f(s, f(s, s))$. If $s_i > 0$ then $f(s_i, s_i) = 2$, therefore a multiset which yields $\mathbf{1}$ by f should in fact be a set. Moreover a set which yields $\mathbf{1}$ satisfies that for all $i \leq n$ there is one and only one of its elements with a coefficient 1 at the index i . Such a set is an exact cover of \mathcal{S} , which proves the reduction.

The problem is in NP because an element v is in $Cl_{\langle f \rangle}(\mathcal{S})$, if and only if there is a multiset of elements of \mathcal{S} such that applying f to its elements yields v . We impose that the witnesses are multiset containing

element at most twice. We have seen these witnesses always exist and it guarantees that they are of polynomial size. \square

This hardness result implies that we cannot use the backtrack search to solve the associated enumeration algorithm. However, if we allow a space proportional to the number of solutions, we can still get a polynomial delay algorithm for associative functions, a property satisfied by the function f of the last proposition. Note that the space used is proportional to the *number of solutions* while the backtrack search only requires a space polynomial in the input size.

Proposition 32. *If f is an associative function of arity 2, then $\text{ENUMCLOSURE}_{\langle f \rangle} \in \text{DelayP}$.*

Proof: Let \mathcal{S} be an instance of $\text{ENUMCLOSURE}_{\langle f \rangle}$. Let G be the directed graph with vertices $Cl_{\langle f \rangle}(\mathcal{S})$ and from each $v \in Cl_{\langle f \rangle}(\mathcal{S})$, there is an arc to $f(v, s)$ for all $s \in \mathcal{S}$. Since f is associative, by definition of G , every vertex of $Cl_{\langle f \rangle}(\mathcal{S})$ is accessible from a vertex in \mathcal{S} . Therefore we can do a depth-first traversal of the graph G to enumerate all solutions. A step of the traversal is in polynomial time: from an element v we generate its neighborhood: $f(v, s)$ for $s \in \mathcal{S}$. The computation of $f(v, s)$ is in time $O(n)$ and $|\mathcal{S}| = m$. We must also test whether the solution $f(v, s)$ has already been generated. This can be done in time $O(n)$ by maintaining a trie containing the generated solutions. In conclusion the delay of the enumeration algorithm is in $O(mn)$ thus polynomial. \square

To obtain a polynomial space algorithm, we could try to use the *reverse search* method [1]. To do that, we want the graph G to be a directed acyclic graph, which is the case if we require the function to be monotone. The monotonicity also ensures that the depth of G is at most $n(|D| - 1)$. However we also need to be able to compute for each element of G a canonical ancestor in polynomial time and it does not seem to be easy even when f is monotone.

Open problem 1. *Find a property of f which ensures the existence of an easy to compute ancestor so that we obtain a polynomial delay and polynomial space enumeration algorithm.*

Open problem 2. *Generalize the previous algorithm to clones generated by several functions or functions of arity larger than three.*

5 Non uniform operators

In this section we generalize the class of operators acting on vectors of booleans. We relax the condition that an operator acts in the same way on each coefficient. On the other hand, we consider operators acting on k solutions at most and in polynomial time, so that the saturation algorithm still works in incremental polynomial time.

We study the simplest case of *unary* operators acting on a *single* coefficient. There are four functions from the boolean domain to the boolean domain. When computing the closure, the identity has no effect, while the negation allows to have any value for a coefficient which is also trivial to deal with. Therefore we are only interested in the two constant functions.

Let \uparrow_i be the function which to a vector v associates $\uparrow_i(v)$ such that for all $j \neq i$, $\uparrow_i(v)_j = v_j$ and $\uparrow_i(v)_i = 1$. We call this function an *upward closure*. The *downward closures* are defined similarly: for each vector v and index $j \neq i$, $\downarrow_i(v)_j = v_j$ and $\downarrow_i(v)_i = 0$. We prove that we can add any sets of such functions to a set of operators over the boolean domain and enumerate the elements of a closure with polynomial delay.

We generalize the problem $\text{MEMBERSHIP}_{\mathcal{F}}$ and $\text{ENUMCLOSURE}_{\mathcal{F}}$, by adding to the input the set of upward and downward closures that can be used to compute the closure. Indeed, these operators work on a single coefficient and we should allow to grow their numbers when the size of the vectors increases to make our problem meaningfully different. In particular we want to allow the set of all downward closures as input to represent solutions closed by taking subsets.

For each set of boolean operators \mathcal{F} , we define $\text{ENUMUDCLOSURE}_{\mathcal{F}}$: given a set of vectors \mathcal{S} , a set of upward closures \mathcal{U} and a set of downward closures \mathcal{D} , enumerate $Cl_{\mathcal{F} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S})$. We present proofs that all problems $\text{ENUMUDCLOSURE}_{\mathcal{F}}$ can be reduced to $\text{ENUMCLOSURE}_{\mathcal{F}}$ in a tight way, which provides good enumeration algorithms except for $\text{ENUMUDCLOSURE}_{\emptyset}$. Note that when both $\uparrow_i \in \mathcal{U}$ and $\downarrow_i \in \mathcal{D}$, the value of the i th coefficient does not matter. Therefore we always assume that there is no i such that $\uparrow_i \in \mathcal{U}$ and $\downarrow_i \in \mathcal{D}$.

Proposition 33. $\text{ENUMUDCLOSURE}_{\emptyset}$ and $\text{ENUMUDCLOSURE}_{E_2}$ can be solved with delay $O(mn)$.

Proof: The problem $\text{ENUMUDCLOSURE}_{\emptyset}$ can be reduced to ENUMDNF by mapping each element of $v \in \mathcal{S}$ to a clause C_v of a formula $\phi_{\mathcal{S}}$. The variable x_i is in C_v if $v_i = 1$, $\downarrow_i \notin \mathcal{D}$ and $\neg x_i$ is in C_v if $v_i = 0$, $\uparrow_i \notin \mathcal{U}$. By construction, the satisfying assignments of $\phi_{\mathcal{S}}$ are in bijection with $Cl_{\mathcal{U}, \mathcal{D}}(\mathcal{S})$. Since the formula $\phi_{\mathcal{S}}$ has m clauses of size at most n , one can enumerate its models with delay $O(mn)$.

Let \mathcal{S} be a set of vectors, \mathcal{U} be a set of upward closures and \mathcal{D} be a set of downward closures. We want to decide whether $v \in Cl_{E_2 \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S})$. If $v_i = 1$ and $\uparrow_i \in \mathcal{U}$ (resp. if $v_i = 0$ and $\downarrow_i \in \mathcal{D}$), then there is no constraint on the coefficient i which can always be turned to the value 1 (resp. 0). Consider i such that $v_i = 0$ and $\uparrow_i \in \mathcal{U}$. Assume that $v \in Cl_{E_2 \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S})$, therefore there is a sequence of operations in $E_2 \cup \mathcal{U} \cup \mathcal{D}$ which allows to build v from \mathcal{S} . There are only two operators which can change the value of the coefficient i : \uparrow_i and \wedge . Since \wedge is decreasing and acts independently on each coefficient, we can remove all the uses of \uparrow_i in the sequence of operations yielding v without changing the final result. Similarly, if $v_i = 1$ it cannot be useful to use $\downarrow_i \in \mathcal{D}$. Therefore we can reduce the problem of deciding whether $v \in Cl_{E_2 \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S})$ to the problem MEMBERSHIP_{E_2} on some projection of \mathcal{S} . A simple adaptation of Proposition 11 yields an algorithm with delay $O(mn)$. \square

Note that it seems hard to improve the complexity of $\text{ENUMUDCLOSURE}_{\emptyset}$ since we can reduce ENUMMONDNF to it while preserving the size of the instance. To each clause C of a monotone DNF formula ϕ we map the vector v_C , such that $(v_C)_i = 1$ if and only if x_i is in the clause C . Let \mathcal{U} be the set of all upward closures and \mathcal{S}_{ϕ} be the set of vectors v_C . The satisfying assignments of ϕ are in bijection with $Cl_{\mathcal{U}}(\mathcal{S}_{\phi})$ and \mathcal{S}_{ϕ} has as many vectors as ϕ has clauses.

Proposition 34. $\text{ENUMUDCLOSURE}_{L_0}$ and $\text{ENUMUDCLOSURE}_{L_2}$ can be solved with delay $O(n)$.

Proof: Consider an index i , such that $\uparrow_i \in \mathcal{U}$. Then if we consider $\uparrow_i (v + v)$, we obtain the base vector e^i with 0 everywhere except at position i . Therefore the value of coefficient i is not constrained in $Cl_{L_0 \cup \mathcal{D} \cup \mathcal{U}}(\mathcal{S})$ and we can solve the problem on $\mathcal{S}_{[n] \setminus \{i\}}$.

If $\downarrow_i \in \mathcal{D}$ and there is some vector v with $v_i = 1$ then $v + \downarrow_i (v)$ is the vector e^i . If all vectors $v \in \mathcal{S}$ are such that $v_i = 0$ then all vectors $v' \in Cl_{L_0 \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S})$ have $v'_i = 0$. In both cases we can also solve the problem on $\mathcal{S}_{[n] \setminus \{i\}}$. Therefore to solve the problem $\text{ENUMUDCLOSURE}_{L_0}$ it is enough to solve ENUMCLOSURE_{L_0} on a projection of \mathcal{S} .

Since $Cl_{L_2}(\mathcal{S})$ is an affine space, it is equal to $v + Cl_{L_0}(\mathcal{S})$, therefore we can use the same argument to solve $\text{ENUMUDCLOSURE}_{L_2}$ by reduction to ENUMCLOSURE_{L_2} . \square

Proposition 35. *Let \mathcal{F} be a clone with $Th_k^{k+1} \in \mathcal{F}$ then $\text{ENUMUDCLOSURE}_{\mathcal{F}}$ can be solved with delay $O(n^k)$.*

Proof: First notice that by definition of the Cl operations, $Cl_{\mathcal{F} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S}) = Cl_{\mathcal{F}}(Cl_{\mathcal{F} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S}))$. Therefore we can apply Theorem 25 to the right hand side of the equation, since \mathcal{F} contains Th_k^{k+1} . As a consequence we have $Cl_{\mathcal{F} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S}) = \{v \mid \forall |I| = k, v_I \in Cl_{\mathcal{F}}(Cl_{\mathcal{F} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S}_I))\}$. Applying the first equality again, we have $Cl_{\mathcal{F} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S}) = \{v \mid \forall |I| = k, v_I \in Cl_{\mathcal{F} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S}_I)\}$. The characterization we obtain extends the one of Theorem 25. Since $Cl_{\mathcal{F} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S}_I)$ can be generated in polynomial time for all I of size k , we can use the same algorithm as in Corollary 26. \square

Proposition 36. $\text{ENUMUDCLOSURE}_{S_{10}}$ and $\text{ENUMUDCLOSURE}_{S_{12}}$ can be solved with delay $O(mn)$.

Proof: We consider a set of upward closures \mathcal{U} and we let s' be the vector s to which each element of \mathcal{U} has been applied. Let e^i be the vector such that $e^i_i = 1$ and $e^i_j = 0$ for $i \neq j$. We let $\mathcal{S}' = \mathcal{S} \cup \{s' \mid s \in \mathcal{S}\} \cup \{e^i \mid \uparrow_i \in \mathcal{U}\}$. Let us prove that $Cl_{S_{10} \cup \mathcal{U}}(\mathcal{S}) = Cl_{S_{10}}(\mathcal{S}')$. Note that the vector $\mathbf{0}$ is in $Cl_{S_{10}}(\mathcal{S})$ since \wedge is in the clone S_{10} and for every index i there are two vectors v, w such that $v_i = 0$ and $w_i = 1$. The inclusion from right to left is clear: it is easy to build s' by definition and $e^i = \uparrow_i(\mathbf{0})$. The other inclusion is true since $\uparrow_i(x) = x' \wedge (x \vee e^i)$.

We can prove by induction, using the monotonicity of \vee and \wedge that $Cl_{S_{10} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S}) = Cl_{\mathcal{D}}(Cl_{\mathcal{U}}(Cl_{S_{10}}(\mathcal{S})))$. Therefore we only need to characterize $Cl_{\mathcal{D}}(Cl_{S_{10}}(\mathcal{S}'))$, since we have shown that $Cl_{\mathcal{U}}(Cl_{S_{10}}(\mathcal{S})) = Cl_{S_{10}}(\mathcal{S}')$. Let us define y^i as the intersection of all elements $s \in \mathcal{S}'$ with $s_i = 1$ to which are applied all $\downarrow_j \in \mathcal{D}$ for $j \neq i$. Let us define $\mathcal{S}'' = \mathcal{S}' \cup \{y^i \mid i \in [n]\}$, we now prove that $Cl_{\mathcal{D}}(Cl_{S_{10}}(\mathcal{S}')) = Cl_{S_{10}}(\mathcal{S}'')$. The inclusion from right to left is clear: it is easy to build y^i since \wedge is in the clone S_{10} . Now consider $v \in Cl_{S_{10}}(\mathcal{S}')$ and $\downarrow_i \in \mathcal{D}$. By Lemma 18, v can be seen as the union of some atoms of \mathcal{S} intersected with v . Therefore, if we consider the union of y^j with $j \in \mathbb{1}(v)$ and $j \neq i$, it is equal to $\downarrow_i(v)$ once intersected with v since $y^j_i = 0$ by construction. It proves that $\downarrow_i(v) \in Cl_{S_{10}}(\mathcal{S}'')$ which proves the equality of the two closures.

As a conclusion, we have reduced an instance of $\text{ENUMUDCLOSURE}_{S_{10}}$ with m vectors of size n to an instance of $\text{ENUMCLOSURE}_{S_{10}}$ with at most $O(m+n)$ vectors of size n which proves the result, using the algorithm of Proposition 19.

We now consider $Cl_{S_{12} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S})$. Let $\uparrow_i \in \mathcal{U}$, we have assumed that there is $s \in \mathcal{S}$ with $s_i = 0$. Consider $x \in Cl_{S_{12} \cup \mathcal{U} \cup \mathcal{D}}(\mathcal{S})$, we have $x \wedge (\uparrow_i(s) \rightarrow s) = \downarrow_i(x)$. Therefore when we have an upper closure \uparrow_i , we can simulate the corresponding downward closure $\downarrow_i(x)$ and thus forget about the coefficient which is not constrained. The proof is the same for downward closures, therefore we can reduce $\text{ENUMUDCLOSURE}_{S_{12}}$ to $\text{ENUMCLOSURE}_{S_{12}}$ on a projection of the instance, which by Proposition 19 yields the desired algorithm. \square

If we allow more general operators as inputs, that is any k -ary function from k -tuples of vectors to vectors, it is possible to use the incremental delay saturation algorithm. However, the problem of deciding membership in the closure becomes hard. For instance, we consider unary operators $O_{i,j,k}$, which act on the coefficients i, j, k of a vector: if $v_i = v_j = v_k = 0$ then $O_{i,j,k}(v)_i = O_{i,j,k}(v)_j = O_{i,j,k}(v)_k = 1$ otherwise the coefficients are left unchanged. It is easy to reduce the problem EXACT3COVER to deciding membership in the closure by these operators (a subset of them are given as input). To each subset $\{i, j, k\}$ is associated the operator $O_{i,j,k}$, and the vector $\mathbf{1}$ belongs to the closure of $\mathcal{S} = \{\mathbf{0}\}$ by these operators if and only if the set of subsets has an exact cover.

Open problem 3. *Is the membership problem hard for unary operators acting on 2 coefficients ?*

There is another natural view on downward and upward closures. The question is to enumerate the subsets or supersets of $Cl_{\mathcal{F}}(\mathcal{S})$ for any fixed \mathcal{F} . In that way, upward and downward closures do not mix with the operators in \mathcal{F} . We let the reader check that these problems are still solvable with polynomial delay by inspecting every case of the reduced Post's lattice. When generating a set of elements closed by subsets or supersets, we are often only interested in the minimal or maximal elements for inclusion. The next section deal with this natural extension.

6 Minimal and maximal elements

In this section, we are interested in enumerating minimal (resp. maximal) elements of $Cl_{\mathcal{F}}(\mathcal{S})$ for a given family \mathcal{F} . Given two vectors $v, u \in \{0, 1\}^n$, we say that v is smaller than u if $\mathbb{1}(v) \subseteq \mathbb{1}(u)$, that is the vectors are ordered by the inclusion of the corresponding sets. A vector $v \in Cl_{\mathcal{F}}(\mathcal{S})$ is said to be minimal (resp. maximal), if it is smaller (resp. larger) or incomparable with all vectors of $Cl_{\mathcal{F}}(\mathcal{S})$. We denote by $\text{ENUMCLOSUREMAX}_{\mathcal{F}}$ and $\text{ENUMCLOSUREMIN}_{\mathcal{F}}$ the problem of, given \mathcal{S} , enumerating respectively the maximal and the minimal elements of $Cl_{\mathcal{F}}(\mathcal{S}) \setminus \{0, 1\}$. We exclude 0 and 1 to obtain non trivial problems when $Cl_{\mathcal{F}}(\mathcal{S})$ contains them. Note that $\text{ENUMCLOSUREMAX}_{\mathcal{F}}$ and $\text{ENUMCLOSUREMIN}_{\mathcal{F}}$ cannot be solved by a simple saturation algorithm, therefore they are not naturally in IncP contrarily to all the problems studied in the previous sections.

6.1 Cases with few solutions

There are many clones which have only a polynomial number of minimal or maximal elements which can be easily found, making their enumeration trivial.

The maximal elements of a closure by E_2 are the original vectors since the operation is decreasing. The minimal elements are the intersection of all base vectors containing i for each $i \in [n]$. The closures by the clones M_2 , R_0 and R are characterized as union of atoms or the intersection with a base vector of a union of atoms for S_{10} and S_{12} . As a consequence the minimal elements are these atoms which can be computed in polynomial time. For the maximal elements, if we assume that the union of all atoms is equal to 1 then it is the unions of all atoms but one.

The hierarchies S_{10}^k and S_{12}^k contain respectively S_{10} and S_{12} and thus a closure by S_{10}^k and S_{12}^k contains the atoms defined for S_{10} and S_{12} . They are their minimal elements, which are thus easy to produce. However since the structure of S_{10}^k and S_{12}^k is slightly different from S_{10} and S_{12} , they have many more maximal elements and we will see in Section 6.4 that they are much harder to enumerate.

We are left with the cases of minimal and maximal elements when the clone contains the majority function, maximal elements when the clone contains a threshold function of arity larger than 3 and minimal and maximal elements for the algebraic clones L_0 and L_2 .

6.2 Majority

For the five clones which contain the majority operator, we can use the reduction to a $2CNF$ formula of Proposition 24. Then we can generate the maximal (resp. minimal) elements of their closure by generating maximal (resp. minimal) models of a $2CNF$ formula. Note that the minimal models of a $2CNF$ formula are the negations of the maximal models of the same formula where each variable is negated. Therefore the problem of enumerating maximal models is equivalent to the problem of generating minimal models.

The problem of enumerating the maximal models of a $2CNF$ formula has already been studied and can be solved with polynomial delay as explained in [23]. The algorithm works by eliminating by resolution all positive variables in the formula. It means that for a given variable x , and for all clauses $A \vee x$ and $B \vee \neg x$ all clauses $A \vee B$ are created and $A \vee x$, $B \vee \neg x$ are removed. This operation does not change the set of maximal models. This step can be successively applied in time $O(n^3)$ to obtain m' clauses where n is the number of variables. Then the formula is a $2CNF$, such that each variable is always positive or always negative. When enumerating maximal models of such a formula all positive variables are set to 1 and the formula is thus an antimonotone $2CNF$. Therefore the problem can be reduced to the enumeration of maximal independent sets of the graph whose vertices are variables and the edges are clauses. The later problem can be solved with delay $O(n^3)$ and space $O(n^2)$ since $m' \leq n^2$, see [26]. Note that any graph can be encoded into an antimonotone $2CNF$, therefore improving the enumeration of maximal solutions of antimonotone $2CNF$ means finding a better enumeration algorithm for maximal independent sets in general graphs which is a hard open problem.

6.3 Linear algebra

The clone L_0 is generated by $+$, hence the closure by L_0 is a vector space. Therefore we want to understand the problem of finding the minimal and maximal vectors of a vector space over \mathbb{F}_2 . Since the minimum of a vector space is always $\mathbf{0}$, we consider the minimal *non zero vectors*, a concept which is related to circuits of a matroid. A binary matroid \mathcal{M} is represented by a set of vectors over \mathbb{F}_2 . A circuit of \mathcal{M} is a minimal –for the inclusion– dependent set, that is a subset S of vectors in \mathcal{M} such that there is a linear combination of elements in S equal to $\mathbf{0}$.

Theorem 37. *The problem $\text{ENUMCLOSUREMIN}_{L_0}$ is equivalent to the enumeration of the circuits of a binary matroid.*

Proof: Let \mathcal{M} be a binary matroid and let M be the matrix whose columns are the vectors of \mathcal{M} . The dependent sets of \mathcal{M} are the solutions to $Mx = 0$ or supersets of these elements. We can find a basis of the kernel of M in polynomial time and we call it B . Notice that by definition, $Cl_{L_0}(B)$ contains only dependent sets of \mathcal{M} and amongst them all its minimal dependent sets. Hence solving $\text{ENUMCLOSUREMIN}_{L_0}$ on B is equivalent to the enumeration of the circuits of \mathcal{M} .

Now assume that we are given a set S of binary vectors. We can compute a basis B of the vector space $Cl_{L_0}(S)$ and then we compute in polynomial time a matrix M such that B is a basis of the kernel of M . The binary matroid defined by the columns vectors of M has for circuits the minimal elements of $Cl_{L_0}(B) = Cl_{L_0}(S)$, which proves the reduction. \square

The enumeration of the circuits of a matroid can be done in incremental polynomial time [24], but it is an open question to find a polynomial delay for this problem even when the matroid is binary. In particular, it is not possible to enumerate the circuits by increasing Hamming weight, since finding the one with smallest Hamming weight is NP-hard [3]. Moreover, we cannot use the backtrack search, since deciding whether a set of vectors can be extended into a circuit is NP-complete [14]. Therefore we have an incremental polynomial time algorithm for $\text{ENUMCLOSUREMIN}_{L_0}$ but improving it would solve an important open question.

The reduction done for $\text{ENUMCLOSUREMIN}_{L_0}$ also works for $\text{ENUMCLOSUREMIN}_{L_2}$, but since it is an affine space, the problem is equivalent to finding the minimal solutions of a non homogeneous system. As a consequence solving $\text{ENUMCLOSUREMIN}_{L_2}$ is equivalent to enumerating the circuits of a matroid containing a *given element*. This later problem is also in IncP [24].

Open problem 4. *Can we relate $\text{ENUMCLOSUREMAX}_{L_0}$ to a known enumeration problem as we have done for $\text{ENUMCLOSUREMIN}_{L_0}$? In particular, is $\text{ENUMCLOSUREMAX}_{L_0} \in \text{IncP}$.*

6.4 The infinite part

We now study $\text{ENUMCLOSUREMAX}_{S_{10}^k}$ and $\text{ENUMCLOSUREMAX}_{S_{12}^k}$ and we show that they are both equivalent to the enumeration of maximal independent sets of a hypergraph of dimension k , a problem denoted by ENUMKMAXINDSET . The dimension of a hypergraph is the size of its largest hyperedge. An independent set of a hypergraph is a subset of its vertices that does not contain any hyperedge. It is said to be maximal, if it is maximal by inclusion. The problem of enumerating the maximal independent sets or equivalently the minimal transversal of a hypergraph of dimension bounded by k belongs to IncP (c.f. [16, 6, 5]) but no polynomial delay algorithm is known. Therefore the equivalence of ENUMKMAXINDSET with $\text{ENUMCLOSUREMAX}_{S_{10}^k}$ and $\text{ENUMCLOSUREMAX}_{S_{12}^k}$ proves that they are in IncP since this class is closed under polynomial delay reduction.

We first give a characterization of the vectors of $Cl_{S_{10}^k}(\mathcal{S})$ and $Cl_{S_{12}^k}(\mathcal{S})$. By Theorem 25, a vector v belongs to $Cl_{S_{10}^k}(\mathcal{S})$ (resp. $Cl_{S_{12}^k}(\mathcal{S})$) if and only if for every subset of indices I of size k , $v_I \in Cl_{S_{10}^k}(\mathcal{S})_I$ (resp. $v_I \in Cl_{S_{12}^k}(\mathcal{S})_I$). So the vectors of the closure are completely determined by the closures of the projections on the k -subsets of indices of \mathcal{S} . Since the projection of the closure is equal to the closure of the projection, the following two lemmas give a characterization of the vectors in $Cl_{S_{10}^k}(\mathcal{S})_I$ and $Cl_{S_{12}^k}(\mathcal{S})_I$.

Lemma 38. *Let $k > 3$, \mathcal{S} be a set of vectors each of size k and let v be a vector of size k . Then $v \in Cl_{S_{10}^k}(\mathcal{S})$ if and only if:*

- *There exists $u \in \mathcal{S}$ such that $\mathbb{1}(v) \subseteq \mathbb{1}(u)$*
- *For all $i, j \leq k$ such that $v_i = 1$ and $v_j = 0$, there exists $u \in \mathcal{S}$ with $u_i = 1$ and $u_j = 0$*

Proof: (\Rightarrow) Let us prove that if for a given set of indices I , $|I| \leq k$, \mathcal{S} has the property that no projection of its vectors on I is equal to $\mathbf{1}$, then this property is preserved by the application of Th_k^{k+1} and thus $Cl_{S_{10}^k}(\mathcal{S})$ has the same property. This implies that if no vector u of \mathcal{S} is such that $\mathbb{1}(v) \subseteq \mathbb{1}(u)$, then $v \notin Cl_{S_{10}^k}(\mathcal{S})$. Let us consider $Th_k^{k+1}(u^1, \dots, u^{k+1})$, each u^i has at least one zero on indices in I . Since there are $k+1$ such vectors and I is of size at most k , by the pigeonhole principle, there are $i, j \leq k$ and $l \in I$ such that $u_l^i = u_l^j = 0$. This implies that $Th_k^{k+1}(u^1, \dots, u^{k+1}) \neq \mathbf{1}$. Thus, if $v \in Cl_{S_{10}^k}(\mathcal{S})$, there exists $u \in \mathcal{S}$ with $\mathbb{1}(v) \subseteq \mathbb{1}(u)$.

Let us prove now that if, for all vectors $u \in \mathcal{S}$, we have $u_{i,j} \neq (1, 0)$, then the same is true in $Cl_{S_{10}^k}(\mathcal{S})$, since the property is preserved by the application of Th_k^{k+1} . Assume that $v = Th_k^{k+1}(u^1, \dots, u^{k+1})$ with $v_i = 1$. Then at least k vectors from u^1, \dots, u^{k+1} are equal to 1 on the i^{th} coordinate. By hypothesis, all these k vectors are equal to 1 on the j^{th} coordinate which implies that $v_j = 1$.

(\Leftarrow) Let v be a vector and assume that there exists $u \in \mathcal{S}$ such that $\mathbb{1}(v) \subseteq \mathbb{1}(u)$ and assume that for all $i, j \leq k$ such that $v_i = 1$ and $v_j = 0$, there exists $u \in \mathcal{S}$ with $u_i = 1$ and $u_j = 0$. We show that $v \in Cl_{S_{10}^k}(\mathcal{S})$. By rearranging the coordinates, assume without loss of generality that $v_i = 1$ for all $i \leq \ell$ and $v_i = 0$ for all $\ell < i \leq k$. We show by induction that for all $\ell \leq r \leq k$, $v_{[r]} \in Cl_{S_{10}^k}(\mathcal{S})_{[r]}$. For $r = \ell$, we have $v_{[r]} = \mathbf{1}$. By assumption there exists $u \in \mathcal{S}$ with $\mathbb{1}(v) \subseteq \mathbb{1}(u)$, so we have $v_{[r]} = u_{[r]} \in \mathcal{S}_{[r]} \subseteq Cl_{S_{10}^k}(\mathcal{S})_{[r]}$. Assume now that $v_{[r-1]} \in Cl_{S_{10}^k}(\mathcal{S})_{[r-1]}$ for a given r with $\ell < r \leq k$

and let us show that $v_{[r]} \in Cl_{S_{10}^k}(\mathcal{S})_{[r]}$. Since $v_{[r-1]} \in Cl_{S_{10}^k}(\mathcal{S})_{[r-1]}$ there exists $v' \in Cl_{S_{10}^k}(\mathcal{S})$ such that $v'_{[r-1]} = v_{[r-1]}$. Observe first that if $v'_r = 0$ then the result directly holds since $v_r = 0$ and then $v'_{[r]} = v_{[r]}$. So assume that $v'_r = 1$. Since by assumption, for each $i, j \leq k$ such that $v_i = 1$ and $v_j = 0$ there exists a vector $u \in \mathcal{S}$ with $u_i = 1$ and $u_j = 0$, we have that for each $i \leq \ell$, there exists $u^i \in \mathcal{S}$ with $u^i_i = 1$ and $u^i_r = 0$. We now construct a sequence of vectors $(z^i)_{i \leq \ell}$ that belong to $Cl_{S_{10}^k}(\mathcal{S})$ and we prove by induction that each z^i is such that :

- $z^i_j = 1$ for all $j \leq i$
- $z^i_j = 0$ for all $\ell < j \leq r$

Then we have that $z^\ell_{[r]} = v_{[r]}$ which will prove that $v_{[r]} \in Cl_{S_{10}^k}(\mathcal{S})_{[r]}$. We first let $z^1 = Th_k^{k+1}(v', v', \dots, v', u^1, u^1)$. Clearly $z^1_1 = 1$ since $u^1_1 = 1$ and $v'_1 = 1$. Furthermore, since $u^1_r = 0$, $z^1_r = 0$ and since $v'_j = 0$ for all $\ell < j < r$, $z^1_j = 0$ for all $\ell < j < r$. Now for $1 < i \leq \ell$ we let $z^i = Th_k^{k+1}(v', v', \dots, v', u^i, z^{i-1})$. Let us prove that z^i has the sought properties. Let $j \leq i$. Notice that since $v'_j = 1$, $z^i_j = 1$ if either u^i_j or z^{i-1}_j is equal to one. If $j = i$ then $u^i_j = 1$ and then $z^i_j = 1$. Now if $j < i$ then by induction, $z^{i-1}_j = 1$ and then $z^i_j = 1$. Now let $\ell < j \leq r$. If $j < r$, we have that $v'_j = 0$ and then $z^i_j = 0$. Now since by induction $z^{i-1}_r = 0$ and since by definition $u^i_r = 0$, we have $z^i_r = 0$. So for every $i \leq \ell$, z^i has the sought property which implies in particular that $z^\ell_{[r]} = v_{[r]}$. Since at each step, z^i is constructed by applying the threshold operator on vectors of $Cl_{S_{10}^k}(\mathcal{S})$, z^ℓ belongs to $Cl_{S_{10}^k}(\mathcal{S})$ and thus $v_{[r]} \in Cl_{S_{10}^k}(\mathcal{S})_{[r]}$. \square

Lemma 39. *Let $k > 3$, \mathcal{S} be a set of vectors each of size k and let v be a vector of size k . Assume that for all $i < k$, there is $s \in \mathcal{S}$ such that $s_i = 0$. Then $v \in Cl_{S_{12}^k}(\mathcal{S})$ if and only if:*

- *There exists $u \in \mathcal{S}$ such that $\mathbb{1}(v) \subseteq \mathbb{1}(u)$*
- *For all $i, j \leq k$ such that $v_i \neq v_j$, there exists $u \in \mathcal{S}$ with $u_i \neq u_j$*

Proof: (\Rightarrow) Assume that $v \in Cl_{S_{12}^k}(\mathcal{S})$. Observe first that if for two coordinates i, j we have $u_i = u_j$ for all vectors in \mathcal{S} , then applying any boolean operator produce a vector for which both coordinates are equal. So if $v_i \neq v_j$ there must be a vector in \mathcal{S} with $u_i \neq u_j$. Now let us show that there exists a vector $u \in \mathcal{S}$ with $\mathbb{1}(v) \subseteq \mathbb{1}(u)$. By Lemma 38 if no vector u of \mathcal{S} is such that $\mathbb{1}(v) \subseteq \mathbb{1}(u)$, then $v \notin Cl_{S_{10}^k}(\mathcal{S})$. Furthermore the extra operator $x \wedge (y \rightarrow z)$ which is in S_{12}^k is decreasing because of the \wedge . So if no vector u of \mathcal{S} is such that $\mathbb{1}(v) \subseteq \mathbb{1}(u)$, then $v \notin Cl_{S_{12}^k}(\mathcal{S})$.

(\Leftarrow) Assume that v has the following properties

- For all $i, j \leq k$ such that $v_i \neq v_j$, there exists $u \in \mathcal{S}$ with $u_i \neq u_j$
- There exists $s \in \mathcal{S}$ such that $\mathbb{1}(v) \subseteq \mathbb{1}(s)$

and let us show that $v \in Cl_{S_{12}^k}(\mathcal{S})$. The first property shows it is enough to prove it projected on $\mathbb{1}(s)$. We can apply Lemma 18 to $Cl_{S_{12}^k}(\mathcal{S}_{\mathbb{1}(s)})$, which shows it is equal to $Cl_{BF}(\mathcal{S}_{\mathbb{1}(s)})$, since there is no index on which all elements of \mathcal{S} are 1. Since S_{12} is included in S_{12}^k , $Cl_{S_{12}^k}(\mathcal{S})$ contains $Cl_{BF}(\mathcal{S}_{\mathbb{1}(s)})$. Using the second property of v , we have that $v \in Cl_{BF}(\mathcal{S}_{\mathbb{1}(s)})$ which implies that $v \in Cl_{S_{12}^k}(\mathcal{S})$. \square

Notice that if two coordinates i and j are equal for every vector of \mathcal{S} , they will remain equal in the closure. So one can simplify the input vector set by deleting all but one equivalent coordinates. So assuming that no coordinates are equivalent, the two previous lemmas together with Theorem 25 give the two following corollaries.

Corollary 40. *Let $k > 3$, $n \geq k$, \mathcal{S} be a set of vectors each of size n and let v be a vector of size n . Then $v \in Cl_{S_{10}^k}(\mathcal{S})$ if and only if:*

- *For all $i, j \leq n$ such that $v_i = 1$ and $v_j = 0$, there exists $u \in \mathcal{S}$ with $u_i = 1$ and $u_j = 0$*
- *For all $I \subseteq [n]$, $|I| = k$, there exists $u \in \mathcal{S}$ such that $\mathbb{1}(v_I) \subseteq \mathbb{1}(u_I)$*

Corollary 41. *Let $k > 3$, $n \geq k$, \mathcal{S} be a set of vectors each of size n and let v be a vector of size n . Then $v \in Cl_{S_{12}^k}(\mathcal{S})$ if and only if for all $I \subseteq [n]$, $|I| = k$, there exists $u \in \mathcal{S}$ such that $\mathbb{1}(v_I) \subseteq \mathbb{1}(u_I)$.*

We are now ready to reduce ENUMKMAXINDSET to CLOSUREMAX $_{S_{10}^k}$ and CLOSUREMAX $_{S_{12}^k}$. Let $\mathcal{H} = (C, \mathcal{E})$ be a hypergraph of dimension k on n vertices and assume for simplicity that \mathcal{H} is k -regular (i.e. all its hyperedges are of size exactly k). Let $\mathcal{T} = \{\chi_T \mid T \in \binom{V}{k} \setminus \mathcal{E}\}$ where $\binom{V}{k}$ denotes the subsets of V of size k and χ_T denotes the characteristic vector of T . Now for $i \leq n$, let e^i be the vector of size n such that $e^i_i = 1$ and $e^i_j = 0$ for all $j \neq i$. If we denote by \mathcal{U} the set of vectors $\mathcal{T} \cup \{e^i : i \leq n\}$ we have the following property.

Lemma 42. *$Cl_{S_{10}^k}(\mathcal{U})$ (resp. $Cl_{S_{12}^k}(\mathcal{U})$) is the set of the characteristic vectors of the independent sets of \mathcal{H} .*

Proof: Since for every $i, j \leq n$ there is $e^i \in \mathcal{U}$ such that $e^i_i = 1$ and $e^i_j = 0$, by Corollary 40 and 41, a vector belongs to $Cl_{S_{10}^k}(\mathcal{U})$ (resp. $Cl_{S_{12}^k}(\mathcal{U})$) if and only if for every subset of indices $I \subseteq [n]$ of size k , there exists $u \in \mathcal{U}$ such that $\mathbb{1}(v_I) \subseteq \mathbb{1}(u_I)$. So $v \in Cl_{S_{10}^k}(\mathcal{U})$ (resp. $v \in Cl_{S_{12}^k}(\mathcal{U})$) if and only if for every subset I of $\mathbb{1}(v)$ of size k there exists $u \in \mathcal{U}$ such that $v_I = u_I = \mathbf{1}$. Since every vectors of \mathcal{U} has at most k ones, $v \in Cl_{S_{10}^k}(\mathcal{U})$ (resp. $v \in Cl_{S_{12}^k}(\mathcal{U})$) if and only if for all subsets I of $\mathbb{1}(v)$ of size k there exists $u \in \mathcal{T}$ such that u is the characteristic vector of $\mathbb{1}(v_I)$. By construction of \mathcal{T} , the later is equivalent to the fact that I is not a hyperedge of \mathcal{H} . So, since \mathcal{H} is k -regular, $v \in Cl_{S_{10}^k}(\mathcal{U})$ (resp. $v \in Cl_{S_{12}^k}(\mathcal{U})$) if and only if $\mathbb{1}(v)$ contains no hyperedge of \mathcal{H} , i.e. v is the characteristic vector of an independent set of \mathcal{H} . \square

The bijections between the closures of \mathcal{U} and the independent sets of \mathcal{H} are also bijections between the maximal elements of the closures and the maximal independent sets. Since \mathcal{U} can be built in time polynomial in \mathcal{H} , we have proved the existence of polynomial delay reductions as stated in the next corollary.

Corollary 43. *There are polynomial delay reductions from ENUMKMAXINDSET to ENUMCLOSUREMAX $_{S_{10}^k}$ and ENUMCLOSUREMAX $_{S_{12}^k}$.*

Now we show that ENUMCLOSUREMAX $_{S_{10}^k}$ and ENUMCLOSUREMAX $_{S_{12}^k}$ can be reduced to ENUMKMAXINDSET, proving that these three problems are equivalent. Let \mathcal{S} be a set of binary vectors of size n , and $k < n$ be an integer. We denote by $\mathcal{H}_k(\mathcal{S})$ the hypergraph on n vertices such that the hyperedges are the subsets $I \subseteq [n]$, $|I| \leq k$ such that $u_I \neq \mathbf{1}$ for all $u \in \mathcal{S}$. A hyperedge of $\mathcal{H}_k(\mathcal{S})$ is a subset of coordinates that are never set to one at the same time in \mathcal{S} . We denote by $I(\mathcal{H}_k(\mathcal{S}))$ the set of independent sets of \mathcal{H}_k and

we denote by $\text{Max}(I(\mathcal{H}_k(\mathcal{S})))$ its maximal ones. For a family of sets \mathcal{T} , let us denote by $\chi(\mathcal{T})$ the set of characteristic vectors of the sets in \mathcal{T} .

We first show that $Cl_{S_{12}^k}(\mathcal{S})$ is the set of characteristic vectors of $I(\mathcal{H}_k(\mathcal{S}))$. Since the inclusion ordering between sets correspond to the classical ordering between their characteristic vectors, this implies that $\text{Max}(Cl_{S_{12}^k}(\mathcal{S})) = \chi(\text{Max}(I(\mathcal{H}_k(\mathcal{S}))))$. On the other hand $Cl_{S_{10}^k}(\mathcal{S}) \neq \chi(I(\mathcal{H}_k(\mathcal{S})))$, but we will see that $\text{Max}(Cl_{S_{10}^k}(\mathcal{S})) = \text{Max}(Cl_{S_{12}^k}(\mathcal{S})) = \chi(\text{Max}(I(\mathcal{H}_k(\mathcal{S}))))$.

Theorem 44. $\text{ENUMCLOSUREMAX}_{S_{12}^k}$ and ENUMKMAXINDSET are equivalent.

Proof: We first prove that $Cl_{S_{12}^k}(\mathcal{S}) = \chi(I(\mathcal{H}_k(\mathcal{S})))$.

We have $T \in \chi(I(\mathcal{H}_k(\mathcal{S})))$

\Leftrightarrow for each subset $I \subseteq \mathbb{1}(v)$, $|I| \leq k$, I is not a hyperedge of $\mathcal{H}_k(\mathcal{S})$

\Leftrightarrow for each subset $I \subseteq \mathbb{1}(v)$, $|I| \leq k$, there exists $u \in \mathcal{S}$ such that $u_I = 1$

$\Leftrightarrow v \in Cl_{S_{12}^k}(\mathcal{S})$ (By Corollary 41)

Given \mathcal{S} , $\mathcal{H}_k(\mathcal{S})$ can be constructed in time $O(mn^k)$ where m is the number of vectors in \mathcal{S} . Moreover we have proved that $Cl_{S_{12}^k}(\mathcal{S}) = \chi(I(\mathcal{H}_k(\mathcal{S})))$ and χ can be computed in linear time, which proves that this is a polynomial delay reduction from $\text{ENUMCLOSUREMAX}_{S_{12}^k}$ to ENUMKMAXINDSET . Since Corollary 43 proves the reduction in the other direction, the two problems are equivalent. \square

We now prove that $\text{ENUMCLOSUREMAX}_{S_{10}^k}$ and ENUMKMAXINDSET are equivalent by proving that the maximal elements of $Cl_{S_{12}^k}(\mathcal{S})$ and $Cl_{S_{10}^k}(\mathcal{S})$ are the same.

Theorem 45. $\text{ENUMCLOSUREMAX}_{S_{10}^k}$ and ENUMKMAXINDSET are equivalent.

Proof: Since $S_{10}^k \subseteq S_{12}^k$ we have that $Cl_{S_{10}^k}(\mathcal{S}) \subseteq Cl_{S_{12}^k}(\mathcal{S})$.

Claim 1: $\text{Max}(Cl_{S_{12}^k}(\mathcal{S})) \subseteq \text{Max}(Cl_{S_{10}^k}(\mathcal{S}))$

Let $v \in \text{Max}(Cl_{S_{12}^k}(\mathcal{S}))$ and let us show that $v \in Cl_{S_{10}^k}(\mathcal{S})$. Assume it is not, by corollaries 40 and 41 and since $v \in Cl_{S_{12}^k}(\mathcal{S})$, we know that there exists $i, j \leq n$ such that $v_{i,j} = (1, 0)$ while the vector $(1, 0)$ does not belong to $\mathcal{S}_{i,j}$. By maximality of v in $Cl_{S_{12}^k}(\mathcal{S})$, the vector v' obtained from v by setting j to 1 does not belong to $Cl_{S_{12}^k}(\mathcal{S})$, so by corollary 41, there exists $I \subseteq \mathbb{1}(v')$, $|I| \leq k$ such that $v'_I \notin \mathcal{S}_I$. Notice that j must belong to I since otherwise v would not belong to $Cl_{S_{12}^k}(\mathcal{S})$. Now let $I' := I \setminus \{j\} \cup \{i\}$. We have that $I' \subseteq \mathbb{1}(v)$ and $|I'| \leq k$. So again, by corollary 41 there exists $u \in \mathcal{S}$ such that $u_{I'} = v_{I'}$. Since $u_I \neq v'_I$ and $v'_j = 1$ we have that $u_j = 0$. But since $u_i = v_i = 1$, $u_{i,j} = (1, 0)$ contradicting the fact that the vector $(1, 0)$ does not belong to $\mathcal{S}_{i,j}$. So $v \in Cl_{S_{10}^k}(\mathcal{S})$, and since v is maximal in $Cl_{S_{12}^k}(\mathcal{S})$ it is also maximal in $Cl_{S_{10}^k}(\mathcal{S})$.

Assume that $v \in \text{Max}(Cl_{S_{10}^k}(\mathcal{S}))$. Since $Cl_{S_{10}^k}(\mathcal{S}) \subseteq Cl_{S_{12}^k}(\mathcal{S})$, $v \in Cl_{S_{12}^k}(\mathcal{S})$. Now assume that v is not maximal in $Cl_{S_{12}^k}(\mathcal{S})$, then there exists $v' \in \text{Max}(Cl_{S_{12}^k}(\mathcal{S}))$ such that $v < v'$. But then, by Claim 1, $v' \in Cl_{S_{10}^k}(\mathcal{S})$ which contradicts the maximality of v in $Cl_{S_{10}^k}(\mathcal{S})$.

We have thus proved that $\text{Max}(Cl_{S_{12}^k}(\mathcal{S})) = \text{Max}(Cl_{S_{10}^k}(\mathcal{S}))$ which proves that $\text{ENUMCLOSUREMAX}_{S_{10}^k}$ and $\text{ENUMCLOSUREMAX}_{S_{12}^k}$ are equivalent. Then by Theorem 44 and the transitivity of equivalence, we have that $\text{ENUMCLOSUREMAX}_{S_{10}^k}$ and ENUMKMAXINDSET are equivalent. \square

Acknowledgements

The authors have been partially supported by the French Agence Nationale de la Recherche, AGGREG project reference ANR-14-CE25-0017-01 and we thank the members of the project and Mamadou Kanté for interesting discussions about enumeration. We also thank Florent Madelaine for his help with CSP and universal algebra. We thank Bruno Zanuttini for his remarks about representing closures as boolean formulas.

References

- [1] David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65(1-3):21–46, 1996.
- [2] Kirby A. Baker and Alden F. Pixley. Polynomial interpolation and the chinese remainder theorem for algebraic systems. *Mathematische Zeitschrift*, 143(2):165–174, 1975.
- [3] E.R. Berlekamp, R.J. McEliece, and H.C.A. Van Tilborg. On the inherent intractability of certain coding problems. *IEEE Transactions on Information Theory*, 24(3):384–386, 1978.
- [4] Elmar Böhler and Heribert Vollmer. Boolean functions and posets lattice with applications to complexity theory. *Lecture Notes for Logic and Interaction*, 2002.
- [5] Endre Boros, Vladimir Gurvich, Khaled Elbassioni, and Leonid Khachiyan. An efficient incremental algorithm for generating all maximal independent sets in hypergraphs of bounded dimension. *Parallel Processing Letters*, 10(04):253–266, 2000.
- [6] Endre Boros, Vladimir Gurvich, and Peter L. Hammer. Dual subimplicants of positive boolean functions. *Optimization Methods and Software*, 10:147–156, 1998.
- [7] Andrei Bulatov, Víctor Dalmau, Martin Grohe, and Dániel Marx. Enumerating homomorphisms. *Journal of Computer and System Sciences*, 78(2):638–650, 2012.
- [8] Andrei Bulatov, Marcin Kozik, Peter Mayr, and Markus Steindl. The subpower membership problem for semigroups. *International Journal of Algebra and Computation*, 26(07):1435–1451, 2016.
- [9] Stanley Burris and Hanamantagida Pandappa Sankappanavar. *A Course in Universal Algebra*, volume 78 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1981 (2012 online version).
- [10] Florent Capelli. *Structural restrictions of CNF-formulas: applications to model counting and knowledge compilation*. PhD thesis, Université Paris Diderot, 2016.
- [11] Florent Capelli and Yann Strozecki. Incremental delay enumeration: Space and time. *Discrete Applied Mathematics*, 2018.
- [12] Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, pages 151–158. ACM, 1971.

- [13] Nadia Creignou and J-J. Hébrard. On generating all solutions of generalized satisfiability problems. *RAIRO-Theoretical Informatics and Applications*, 31(6):499–511, 1997.
- [14] Arnaud Durand and Miki Hermann. The inference problem for propositional circumscription of affine formulas is conp-complete. In *STACS*, pages 451–462. Springer, 2003.
- [15] Arnaud Durand and Yann Strozecki. Enumeration complexity of logical query problems with second-order variables. In *Proceedings of the 20th Conference on Computer Science Logic*, pages 189–202, 2011.
- [16] Thomas Eiter and Georg Gottlob. Identifying the minimal transversals of a hypergraph and related problems. *SIAM Journal on Computing*, 24(6):1278–1304, 1995.
- [17] Khaled Elbassioni and Kazuhisa Makino. Enumerating vertices of 0/1-polyhedra associated with 0/1-totally unimodular matrices. In *16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [18] Tomás Feder. Network flow and 2-satisfiability. *Algorithmica*, 11(3):291–319, 1994.
- [19] Jörg Flum and Martin Grohe. *Parameterized complexity theory*. Springer Science & Business Media, 2006.
- [20] Merrick Furst, John Hopcroft, and Eugene Luks. Polynomial-time algorithms for permutation groups. In *Foundations of Computer Science, 1980., 21st Annual Symposium on*, pages 36–41. IEEE, 1980.
- [21] Michael R. Garey and David S. Johnson. *Computers and intractability*, volume 29. W Freeman New York, 2002.
- [22] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. On generating all maximal independent sets. *Information Processing Letters*, 27(3):119–123, 1988.
- [23] Dimitris J. Kavvadias, Martha Sideri, and Elias C. Stavropoulos. Generating all maximal models of a boolean expression. *Information Processing Letters*, 74(3-4):157–162, 2000.
- [24] Leonid Khachiyan, Endre Boros, Khaled Elbassioni, Vladimir Gurvich, and Kazuhisa Makino. On the complexity of some enumeration problems for matroids. *SIAM Journal on Discrete Mathematics*, 19(4):966–984, 2005.
- [25] Donald E. Knuth. Combinatorial Algorithms, part 1, volume 4a of *The Art of Computer Programming*, 2011.
- [26] Eugene L. Lawler, Jan Karel Lenstra, and A. H. G. Rinnooy Kan. Generating all maximal independent sets: Np-hardness and polynomial-time algorithms. *SIAM J. Comput.*, 9(3):558–565, 1980.
- [27] Arnaud Mary. *Énumération des Dominants Minimaux dun graphe*. PhD thesis, Université Blaise Pascal, 2013.

- [28] Arnaud Mary and Yann Strozecki. Efficient enumeration of solutions produced by closure operations. In *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016*, volume 47, 2016.
- [29] Arnaud Mary and Yann Strozecki. Generating maximal solutions given by closure operations. Workshop on Enumeration Problems and Applications, 2016.
- [30] Peter Mayr. The subpower membership problem for Mal'cev algebras. *International Journal of Algebra and Computation*, 22(07):1250075, 2012.
- [31] Keisuke Murakami and Takeaki Uno. Efficient algorithms for dualizing large-scale hypergraphs. *Discrete Applied Mathematics*, 170:83–94, 2014.
- [32] Emil Leon Post. *The two-valued iterative systems of mathematical logic*. Princeton University Press, 1941.
- [33] Robert C. Read. Bounds on backtrack algorithms for listing cycles, paths, and spanning trees. *Networks*, 5:237–252, 1975.
- [34] Steffen Reith and Heribert Vollmer. Optimal satisfiability for propositional calculi and constraint satisfaction problems. *Information and Computation*, 186(1):1–19, 2003.
- [35] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, pages 216–226. ACM, 1978.
- [36] Jeff Shriver. Hardness results for the subpower membership problem. *International Journal of Algebra and Computation*, 28(05):719–732, 2018.
- [37] Markus Steindl. The subpower membership problem for bands. *Journal of Algebra*, 489:529–551, 2017.
- [38] Yann Strozecki. *Enumeration complexity and matroid decomposition*. PhD thesis, Université Paris Diderot - Paris 7, 2010.
- [39] Yann Strozecki. On enumerating monomials and other combinatorial structures by polynomial interpolation. *Theory of Computing Systems*, 53(4):532–568, 2013.
- [40] Ágnes Szendrei, Peter Mayr, and Andrei Bulatov. The subpower membership problem for finite algebras with cube terms. *Logical Methods in Computer Science*, 15, 2019.
- [41] Takeaki Uno. Algorithms for enumerating all perfect, maximum and maximal matchings in bipartite graphs. *Algorithms and Computation*, pages 92–101, 1997.
- [42] Virginia Vassilevska Williams and Ryan Williams. Subcubic equivalences between path, matrix and triangle problems. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 645–654. IEEE, 2010.
- [43] Stephan Zwick. Computing in direct powers of expanded groups. Master's thesis, 2013.