



HAL
open science

The enumeration of meet-irreducible elements based on hierarchical decompositions of implicational bases

Lhouari Nourine, Simon Vilmin

► **To cite this version:**

Lhouari Nourine, Simon Vilmin. The enumeration of meet-irreducible elements based on hierarchical decompositions of implicational bases. 2022. hal-03565763

HAL Id: hal-03565763

<https://hal.science/hal-03565763v1>

Preprint submitted on 11 Feb 2022

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

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

The enumeration of meet-irreducible elements based on hierarchical decompositions of implicational bases

Lhouari Nourine* and Simon Vilmin*

¹Université Clermont-Auvergne, CNRS, Mines de Saint-Étienne, Clermont-Auvergne-INP, LIMOS, 63000 Clermont-Ferrand, France.

February 11, 2022

Abstract

We study the well-known problem of translating between two representations of closure systems, namely implicational bases and meet-irreducible elements. Albeit its importance, the problem is open. In this paper, we introduce *splits* of an implicational base. It is a partitioning operation of the implications which we recursively apply to obtain a binary tree representing a decomposition of the implicational base. We show that this decomposition can be conducted in polynomial time and space in the size of the input implicational base. Focusing on the case of *acyclic* splits, we obtain a recursive characterization of the meet-irreducible elements of the associated closure system. We use this characterization and hypergraph dualization to derive new results for the translation problem in acyclic convex geometries.

Keywords: meet-irreducible enumeration, implicational bases, closure systems, hierarchical decomposition.

1 Introduction

Finite closure systems over a (finite) ground set are set systems containing the ground set and closed under set intersections. When ordered by inclusion, they are also known as (closure) lattices [11, 24]. These structures are well-known in mathematics and computer science. They show up in Knowledge Space Theory (KST) [16], database theory [14, 36], propositional logic [29, 31], Formal Concept Analysis (FCA) [23], or argumentation frameworks [17, 20] for example.

Albeit ubiquitous, closure systems suffer from their size, which can be exponential in the size of their ground set. For this reason, numerous research works have been conducted over the last decades to construct space efficient representations of lattices [4, 9, 23, 25, 26, 31, 36, 38, 43]. The surveys [8, 46] are also recent witnesses of the importance and the relevance of compactly representing closure systems.

Among all possible representations, there are two prominent candidates: *implications* and *meet-irreducible elements*. An implication is a mathematical expression $A \rightarrow B$, where A and B are subsets of the ground set, modeling a causality relation between A and B in the closure system: “If a set includes A , it must also include B ”. Every closure system \mathcal{C} over some ground set V can be represented by a set Σ of implications called an *implicational base*. Dually, every set of implications gives birth to a closure system [43]. As several implicational bases can represent the same closure system, numerous bases with “good” properties have been studied. Among them, the Duquenne-Guigues base [25] being minimum (the least number of implications) or the canonical direct base [9] have attracted much attention. More recently, Adaricheva *et al.*

[1, 2, 3] have proposed refinements of the canonical direct base such as the D -base and the E -base. Because of their simple nature, implications have been used under different shapes and names such as functional dependencies in databases [14, 36], Horn functions in propositional logic [29, 31], queries in KST [16] or attribute implications in FCA [23, 25] for instance. A second way to compactly represent a closure system \mathcal{C} is its family of meet-irreducible elements \mathcal{M} . It is the unique minimal collection of sets from which the whole closure system can be recovered by taking set-intersection. In Horn logic, meet-irreducible elements are called *characteristic models* [31, 29] for they completely identify a given Horn function. Moreover, they appear in the poset of irreducibles in [26, 38], in the Armstrong relations in databases [36], in the base of knowledge spaces [16] or in the reduced context of FCA [23].

In this paper, we study the problem of translating between these two representations. This problem is twofold. Either it asks to list the meet-irreducible elements of a closure system given by an implicational base, or vice-versa, to construct an implicational base from a set of meet-irreducible elements.

The choice of the representation impacts the complexity of several problems, thus making the translation a crucial task. For example, it is **NP**-complete to decide whether an element belongs to a minimal generator of a closure system if the latter is given by an implicational base [35]. When the closure system is represented by its meet-irreducible elements, we can answer the question in polynomial time [8]. The complexity of recognizing a class of closure systems also depends on the representation. For instance, recognizing convex geometries and join-semidistributivity can be achieved in polynomial time from a family of meet-irreducible elements [18, 39]. Whether we can recognize convex geometries and join-semidistributive lattices from an implicational base is open. Another example where the representation matters comes from propositional logic [29], where abductive reasoning can be conducted in polynomial time from meet-irreducible elements, while it is **NP**-complete with implications.

Translating is also important to enjoy the most compact representation for a given closure system. Indeed, implicational bases and meet-irreducible elements are generally much shorter than the closure systems they represent. However, when we compare the two representations, there are cases where an implicational base has size exponential in the number of meet-irreducible elements, or dually, where the number of meet-irreducible elements can be exponential in the size of an implicational base [32, 37].

Known results We now review the principal results on the translation task. It has attracted much attention during the last decades [2, 5, 7, 31, 36, 44]. The surveys [8, 46] provide a detailed account of all the progresses made on this question. Since the size of the output can be exponential in the size of the input, we express the complexity results in terms of the combined size of the input and the output. This is *output-sensitive* complexity [28].

For completeness we discuss four representations for a closure system: implications, meet-irreducible elements, the closure system itself or the closure operator. The closure operator is seen as a black-box oracle returning the smallest closed set including a given set. We explain each direction of Figure 1, which summarizes hardness results about the translation task. Numbers in the figure refer to the following explanations.

(1). *From any representation to the closure operator.* The closure operation can be simulated in polynomial-time from any other representation of the closure system, using intersections and the closure algorithm (or the forward chaining) [23].

(2). *From any representation to the closure system.* The whole closure system can be constructed in output-polynomial time from any other representation, with the help of well-known algorithms such as `NextClosure` [23].

(3). *From the closure operator to meet-irreducible elements and implications.* Lawler *et al.* prove in [33] that meet-irreducible elements or implications cannot be enumerated in output-




Figure 1: The complexity of translating between the representations of a closure system.

polynomial time unless $\mathbf{P} = \mathbf{NP}$ from a closure oracle.

(4). *From the closure system to its meet-irreducible elements.* It is sufficient to perform a traversal of the closed sets, and check for the meet-irreducible property. This is done in (output)-polynomial time.

(5). *From the closure system to an implicational base.* To find a (minimum) implicational base, it is for instance possible to use the attribute-incremental approach of Duquenne and Obiedkov [40] in output-polynomial time.

(6). *From an implicational base to meet-irreducible elements and vice-versa.* Remark that undertaking the construction of the whole closure system as an intermediate will necessarily produce output-exponential time algorithms in the worst case. In the landmark paper [31], written in the framework of Horn logic, these problems are called CCM for *Computing Characteristic Models* and SID for *Structure Identification*. We keep these names for historicity.

MEET-IRREDUCIBLE ELEMENTS ENUMERATION (CCM)

Input: An implicational base Σ of a closure system \mathcal{C} over V .

Output: The meet-irreducible elements \mathcal{M} of \mathcal{C} .

MINIMUM IMPLICATIONAL BASE IDENTIFICATION (SID)

Input: The family \mathcal{M} of meet-irreducible elements of a closure system \mathcal{C} over V .

Output: A minimum implicational base Σ corresponding to \mathcal{C} .

In [31] the author consider right-optimum implicational bases (minimizing the right-hand sides of implications) and shows that both directions of the translation are equivalent. Whether this equivalence also holds for minimum implicational bases is not clear as going from right-optimum to minimum is much easier than the other way around [4, 41]. In any case, the task is already harder than enumerating the maximal independent sets of a hypergraph [31]. This latter problem, also known as *hypergraph dualization*, is a famous open problem [19, 21]. The best known algorithm for this task is the one of Fredman and Khachiyan [21], running in output quasi-polynomial time. Babin and Kuznetsov prove in [5] that it is \mathbf{coNP} -complete to decide whether an implication belongs to a minimum implicational base from the meet-irreducible elements. In [30], the authors state that co-atoms of a closure system cannot be enumerated in output-polynomial time unless $\mathbf{P} = \mathbf{NP}$. In [15], it is shown that the minimal pseudo-closed sets of the Duquenne-Guigues basis cannot be enumerated in output-polynomial time unless $\mathbf{P} = \mathbf{NP}$ either. More recently [13], it has been shown that CCM and SID are harder than hypergraph dualization, even in acyclic convex geometries. In spite of these hardness results, the complexity of translating between meet-irreducible elements and implications remains unsettled.

On the positive side, finding the canonical direct base from the meet-irreducible elements (and vice-versa) is equivalent to hypergraph dualization [8, 9, 31]. Adaricheva *et al.* [2] obtain similar results for the D -base. More generally, exponential time algorithms have been designed, see *e.g.* [23, 40, 36, 44]. In [45], Wild shows that SID can be solved in polynomial time in modular lattices. The authors in [7] devise output-polynomial time algorithms for both CCM and SID in k -meet-semidistributive lattices. Finally, it has been proved [13] that CCM and SID are polynomially equivalent to hypergraph dualization in the class of *ranked convex geometries*.

Contributions and outline. We are mostly interested in the problem CCM in the class of acyclic convex geometries. It is a well-studied class of closure systems [3, 27, 43, 47], lying in the intersection of convex geometries and lower-bounded closure systems [1, 22]. They also contain distributive closure systems, in which the translation can be solved efficiently.

Let Σ be an implicational base for some closure system \mathcal{C} over V . We start with some preliminary definitions in Section 2. Then, we give the following results:

1. We introduce a partitioning operation of an implicational base called a *split*, inspired by [10, 34]. We use this operation to hierarchically decompose Σ and its associated closure system \mathcal{C} . This part is detailed in Section 3.
2. Section 4 is devoted to *acyclic splits*:
 - (1) We characterize \mathcal{C} with respect to this partitioning operation, see Subsection 4.1.
 - (2) We derive a recursive characterization of the set of meet-irreducible elements \mathcal{M} associated to \mathcal{C} , see Subsection 4.2.
 - (3) We devise an algorithm solving CCM in the presence of acyclic splits. We highlight cases where this procedure performs in output-quasipolynomial time using the algorithm of Fredman and Khachiyan [21] for hypergraph dualization. This result includes ranked convex geometries as a particular case. This is Subsection 4.3.

The paper gathers results communicated at the 21st conference ICTCS (for Section 3) and the 8th workshop FCA4AI (for Section 4), without published proceedings.

2 Preliminaries

All the object considered in this paper are *finite*. For more definitions about closure systems and implications, we refer the reader to [8]. If V is a set, we refer to $\mathbf{2}^V$ as the family of all subsets of V . Sometimes, and mostly in examples, we shall write the subset $\{u_1, \dots, u_k\}$ of V as the concatenation of its elements, that is $u_1 \dots u_k$. Let \mathcal{S} be a family of subsets of V . We say that \mathcal{S} is *simple* or an *antichain* if for every $S_1, S_2 \in \mathcal{S}$, $S_1 \not\subseteq S_2$. Let $V' \subseteq V$. The *trace* of \mathcal{S} on V' , denoted by $\mathcal{S}: V'$ is obtained by intersecting each element of \mathcal{S} with V' , that is $\mathcal{S}: V' = \{S \cap V' \mid S \in \mathcal{S}\}$.

Closure systems, closure operators Let V be a set. A *closure system* over V is a family \mathcal{C} of subsets of V such that $V \in \mathcal{C}$ and $C_1 \cap C_2 \in \mathcal{C}$ for every $C_1, C_2 \in \mathcal{C}$. The sets in \mathcal{C} are called *closed (sets)*. When ordered by inclusion, the pair (\mathcal{C}, \subseteq) is a (*closure*) *lattice*. In this paper, we always assume that a closure system is equipped with this order. Hence, we write \mathcal{C} to denote the lattice (\mathcal{C}, \subseteq) . Let $C_1, C_2 \in \mathcal{C}$. We say that C_1 and C_2 are *comparable* if $C_1 \subseteq C_2$ or $C_2 \subseteq C_1$. We write $C_1 \subset C_2$ if $C_1 \subseteq C_2$ but $C_1 \neq C_2$. We say that C_2 *covers* C_1 , denoted by $C_1 \prec C_2$, if $C_1 \subset C_2$ and there is no closed set $C \in \mathcal{C}$ such that $C_1 \subset C \subset C_2$. In this case, C_2 is a *successor* of C_1 and C_1 a *predecessor* of C_2 . Let $C \in \mathcal{C}$. The *ideal* of C in \mathcal{C} , denoted $\downarrow C$ contains all the closed subsets of C , *i.e.* $\downarrow C = \{C' \in \mathcal{C} \mid C' \subseteq C\}$. The *filter* $\uparrow C$ of C in \mathcal{C} is defined dually

with the closed supersets of C . If \mathcal{C}' is a subset of \mathcal{C} , the ideal of \mathcal{C}' is $\downarrow \mathcal{C}' = \bigcup_{C' \in \mathcal{C}'} \downarrow C'$ and its filter is $\uparrow \mathcal{C}' = \bigcup_{C' \in \mathcal{C}'} \uparrow C'$. Let $V' \subseteq V$. A closed set M is *meet-irreducible* if $M = C_1 \cap C_2$ with $C_1, C_2 \in \mathcal{C}$ implies $M = C_1$ or $M = C_2$. The set of meet-irreducible elements of \mathcal{C} is denoted $\mathcal{M}(\mathcal{C})$ or simply \mathcal{M} when clear from the context. The whole closure system can be recovered by taking the intersections of every combinations of meet-irreducible elements. For a given closed set C , we put $\mathcal{M}(C) = \{M \in \mathcal{M} \mid C \subseteq M\}$. We have $C = \bigcap \mathcal{M}(C)$.

Closure systems are closely related to closure operators. A mapping $\phi: \mathbf{2}^V \rightarrow \mathbf{2}^V$ is a *closure operator* if for every $X, Y \subseteq V$, $X \subseteq \phi(X)$ (ϕ is extensive), $X \subseteq Y$ implies that $\phi(X) \subseteq \phi(Y)$ (ϕ is monotone) and $\phi(\phi(X)) = \phi(X)$ (ϕ is idempotent). The family $\mathcal{C} = \{\phi(X) \mid X \subseteq V\} = \{X \subseteq V \mid \phi(X) = X\}$ is a closure system. Similarly, every closure system \mathcal{C} induces a closure operator ϕ defined by $\phi(X) = \bigcap \{C \in \mathcal{C} \mid X \subseteq C\}$ for every $X \subseteq V$. Note that since \mathcal{C} is closed by intersection, we also have that $\phi(X) = \min_{\subseteq}(\{C \in \mathcal{C} \mid X \subseteq C\})$. Thus, the correspondence between closure operators and closure systems is one-to-one.

Let \mathcal{C} be a closure system over V with associated closure operator ϕ . We say that \mathcal{C} is *standard* if for every $u \in V$, $\phi(u) \setminus \{u\}$ is closed. In particular, \emptyset is closed. In this paper, all the closure systems are considered standard, a common assumption [3, 46].

A standard closure system \mathcal{C} over V is *Boolean* if $\mathcal{C} = \mathbf{2}^V$. It is *distributive* if $C_1 \cup C_2 \in \mathcal{C}$ for every pair of closed sets C_1, C_2 . Let $\mathcal{C}_1, \mathcal{C}_2$ be two closure systems over disjoint V_1, V_2 (resp.). The *direct product* of \mathcal{C}_1 and \mathcal{C}_2 is defined by $\mathcal{C}_1 \times \mathcal{C}_2 = \{C_1 \cup C_2 \mid C_1 \in \mathcal{C}_1, C_2 \in \mathcal{C}_2\}$.

Implicational bases An *implication* over V is an expression $A \rightarrow B$ where A and B are subsets of V . An *implicational base* Σ over V is a family of implications (over V). A subset C of V *satisfies* or *models* an implicational base Σ if for every $A \rightarrow B \in \Sigma$, $A \subseteq C$ implies that $B \subseteq C$. It is known [8, 46] that the family $\mathcal{C} = \{C \subseteq V \mid C \text{ satisfies } \Sigma\}$ is a closure system. Its associated closure operator ϕ can be computed with the closure procedure (or the forward chaining) [23]. For a given $X \subseteq V$, this procedure starts from X and constructs a sequence $X = X_0 \subseteq \dots \subseteq X_k = \phi(X)$ of subsets of V such that for every $1 \leq i \leq k$, $X_i = X_{i-1} \cup \bigcup \{B \mid \exists A \rightarrow B \in \Sigma \text{ such that } A \subseteq X_{i-1}\}$. The routine stops when $X_{i-1} = X_i$.

On the other hand, every closure system \mathcal{C} can be represented by at least one implicational base Σ [46]. An implication $A \rightarrow B$ holds in a closure system \mathcal{C} if all the closed sets of \mathcal{C} are models of $A \rightarrow B$. Equivalently, $A \rightarrow B$ holds in \mathcal{C} if $B \subseteq \phi(A)$. Two implicational bases are *equivalent* if they represent the same closure system. In particular, an implicational base Σ is equivalent to its *unit-expansion* $\Sigma_u = \{A \rightarrow b \mid A \rightarrow B \in \Sigma, b \in B\}$. We will interchangeably use an implicational base or its unit-expansion.

Remark 1. As we restrict our attention to standard closure systems, we consider that an implicational base Σ has no implications of the form $\emptyset \rightarrow B$ for some $B \subseteq V$.

Let Σ be an implicational base over V . The *restriction* of Σ to a subset V' of V is the implicational base $\Sigma[V'] = \{A \rightarrow b \in \Sigma \mid A \cup \{b\} \subseteq V'\}$. Then, $\Sigma[V']$ is a *sub-base* of Σ . Let V_1, V_2 be a non-trivial (full) bipartition of V , that is $V_1 \cup V_2 = V$, $V_1 \cap V_2 = \emptyset$ and $V_1 \neq \emptyset, V_2 \neq \emptyset$. An implicational base is *bipartite* (w.r.t. V_1, V_2) if every implication $A \rightarrow B$ satisfies $A \subseteq V_1$ and $B \subseteq V_2$ or vice-versa. We write $\Sigma[V_1, V_2]$ to denote a bipartite implicational base. A *path* in Σ is a sequence v_1, \dots, v_k of elements of V such that for every $1 \leq i < k$ there exists an implication $A_i \rightarrow B_i$ with $v_i \in A_i$ and $v_{i+1} \in B_i$. The path is a *cycle* when $v_1 = v_k$. An implicational base without cycles is called *acyclic*. A closure system which admits an acyclic implicational base is an *acyclic convex geometry* [18]. Acyclic convex geometries are also known as G -geometries [43] or poset type convex geometries [1]. The term acyclic comes from Horn logic and acyclic Horn formulas [27, 47].

Directed hypergraphs [4] are a convenient graphical representation for (unit-expansions of) implicational bases. A *directed hypergraph* \mathcal{D} (over V) is a pair (V, \mathcal{A}) where \mathcal{A} is a set of hyperarcs. A hyperarc is a pair (A, b) where $A \cup \{b\} \subseteq V$, A is the *body* and b the *head* of the

hyperarc. A hyperarc can be used to model an implication $A \rightarrow b$ in the unit-expansion of an implicational base Σ .

Example 1. Let $V = \{1, 2, 3, 4\}$ and $\Sigma = \{12 \rightarrow 3, 23 \rightarrow 4, 4 \rightarrow 1\}$. The sequence 1, 3, 4 is a cycle in Σ . We represent Σ and its associated closure system \mathcal{C} in Figure 2. The meet-irreducible elements of \mathcal{C} are 2, 14, 13 and 134.




Figure 2: On the left, the (associated directed hypergraph of the) implicational base Σ in Example 1. On the right, its associated closure system \mathcal{C} .

Enumeration complexity We conclude with a brief reminder on enumeration algorithms [28]. Let A be an algorithm with input x of size n and output a set of solutions $R(x)$ with m elements. In our case, each solution in $R(x)$ has size $\text{poly}(n)$. We say that A is running in *output-polynomial* time if its execution time is bounded by $\text{poly}(m + n)$. If the execution time of A is instead bounded by $(n + m)^{\log(n+m)}$, A is said to run in *output-quasipolynomial* time.

3 Splits and hierarchical decomposition of implicational bases

Inspired by [10, 34], we define the *split* operation for an implicational base Σ over V . A split is a bipartition (V_1, V_2) of the groundset V which *completely* partitions the implications of Σ in three sub-bases:

$\Sigma[V_1]$: the implications of Σ fully contained in V_1 ,

$\Sigma[V_2]$: the implications of Σ fully contained in V_2 ,

$\Sigma[V_1, V_2]$: the implications of Σ whose premises are included in V_1 and their conclusions in V_2 , or vice-versa.

This partitioning operation can be conducted recursively and leads to a *hierarchical decomposition* (*H-decomposition*) of Σ , represented by a full rooted binary tree. The root of the tree is labelled by $\Sigma[V_1, V_2]$, its left-child corresponds to a decomposition of $\Sigma[V_1]$, its right-child to a decomposition of $\Sigma[V_2]$. This tree is called a Σ -tree. We illustrate the structure of a Σ -tree in Figure 3.




Figure 3: A bipartition of Σ by a split.

We characterize the implicational bases having a hierarchical decomposition into trivial bases, and give a polynomial time and space algorithm, `BuildTree`, which takes an implicational base Σ as an input, and outputs a Σ -tree if it exists. Afterwards, we relax the requirement of the H-decomposition into trivial bases to H-factors, which are indecomposable sub-bases of Σ .

Finally, we consider the decomposition of \mathcal{C} , when a split (V_1, V_2) of Σ is given. We show that \mathcal{C} is obtained by combining closed sets of \mathcal{C}_1 , the closure system of $\Sigma[V_1]$, with closed sets of \mathcal{C}_2 , the closure system of $\Sigma[V_2]$. The way \mathcal{C}_1 and \mathcal{C}_2 are combined depends on the implications in $\Sigma[V_1, V_2]$.

3.1 Split operation

Our first step is to define the split operation.

Definition 1. Let Σ be an implicational base over V . A split of Σ is a non-trivial bipartition (V_1, V_2) of V such that for every $A \rightarrow b \in \Sigma$, $A \subseteq V_1$ or $A \subseteq V_2$.

A split (V_1, V_2) induces three sub-bases $\Sigma[V_1]$, $\Sigma[V_2]$ and a bipartite base $\Sigma[V_1, V_2]$. Moreover, every implication of Σ belongs to exactly one of $\Sigma[V_1]$, $\Sigma[V_2]$ or $\Sigma[V_1, V_2]$ (recall that Σ has no implications $\emptyset \rightarrow b$). Intuitively, the split shows that Σ is fully described by two smaller distinct bases $\Sigma[V_1]$ and $\Sigma[V_2]$ acting on each other through the bipartite implicational base $\Sigma[V_1, V_2]$.

Example 2. Let $V = \{1, 2, 3, 4, 5, 6, 7\}$ and consider the implicational base Σ with implications $12 \rightarrow 3$, $3 \rightarrow 1$, $56 \rightarrow 2$, $23 \rightarrow 7$, $45 \rightarrow 6$ and $5 \rightarrow 7$. Figure 4 represents Σ .




Figure 4: The implicational base of Example 2.

In Figure 5 we consider two possible bipartitions of V . The bipartition illustrated on the left separates V in two sets $V_1 = \{1, 3\}$ and $V_2 = \{2, 4, 5, 6, 7\}$. It is not a split since the premises of $12 \rightarrow 3$ and $23 \rightarrow 7$ intersect both V_1 and V_2 . The bipartition on the right puts $V_1 = \{1, 2, 3\}$ and $V_2 = \{4, 5, 6, 7\}$. It is a split with $\Sigma[V_1] = \{12 \rightarrow 3, 3 \rightarrow 1\}$, $\Sigma[V_2] = \{45 \rightarrow 6, 5 \rightarrow 7\}$, and $\Sigma[V_1, V_2] = \{56 \rightarrow 2, 23 \rightarrow 7\}$.




Figure 5: Two bipartitions of V , the left one is not a split of Σ , the right one is.

Before giving a characterization of implicational bases having a split, we make two observations. First, if Σ is empty or contains only implications of the form $a \rightarrow b$. In this case, every non-trivial bipartition of V —every cut of the associated directed (hyper)graph—is a split. In fact, an implication of the form $a \rightarrow b$ always satisfies the condition of Definition 1. Thus, these

implications have no impact on the existence of a split. Second, there may be implicational bases where no bipartition corresponds to a split, as shown by the next example.

Example 3. Consider $V = \{1, 2, 3\}$ and the implicational base $\Sigma = \{12 \rightarrow 3, 13 \rightarrow 2\}$. Here, none of the three possible bipartitions is a split:

- $V_1 = \{1, 2\}$ and $V_2 = \{3\}$ fails to separate the implication $13 \rightarrow 2$;
- $V_1 = \{1, 3\}$, $V_2 = \{2\}$ omits the implication $12 \rightarrow 3$; and
- $V_1 = \{2, 3\}$, $V_2 = \{1\}$ breaks the two implications of Σ .

In the following, we show that the implicational base's connectivity is important for the notion of a split. Let Σ be an implicational base over V . A *premise-path* in Σ is a sequence v_1, \dots, v_k of (distinct) elements of V such that for every $1 \leq i < k$ there exists an implication $A_i \rightarrow b_i$ in Σ such that $\{v_i, v_{i+1}\} \subseteq A_i$. Two vertices $u, v \in V$ are said to be *premise-connected* in Σ if there exists a premise-path from u to v . We say that Σ is *premise-connected* when every pair of vertices in V is premise-connected. A subset C of V is a *premise-connected component* of Σ if there exists a premise-path between each pair of vertices of C , and if C is inclusion-wise maximal for this property. A singleton premise-connected component of Σ is *trivial*.

Example 4. Consider the implicational base Σ given in Example 2. For instance, $6, 5, 4$ is a premise-path and hence 4 and 6 are premise-connected. Here Σ is not premise-connected as there is no premise-path between 2 and 6. The premise-connected components of Σ are $\{1, 2, 3\}$, $\{4, 5, 6\}$ and $\{7\}$ being trivial.

Using premise-connectivity, we are now in position to identify whether a given implicational base admits a split or not.

Proposition 1. *An implicational base Σ over V has a split if and only if it is not premise-connected.*

Proof. We begin with the only if part. Suppose that Σ has a split (V_1, V_2) , and let $u \in V_1$ and $v \in V_2$. Since a split is a non-trivial bipartition of V , such u and v must exist. Now let us assume for contradiction there exists a premise-path $u = v_1, \dots, v_k = v$ for some $k \in \mathbb{N}$. Such a premise-path exists if there is some j with $1 \leq j \leq k$ such that $A_j \rightarrow b_j$ is an implication of Σ , $A_j \cap V_1 \neq \emptyset$ and $A_j \cap V_2 \neq \emptyset$. However, the implication $A_j \rightarrow b_j$ does not satisfy Definition 1. This contradicts the assumption that (V_1, V_2) is a split of Σ . Hence, u, v cannot be premise-connected and Σ is not premise-connected either.

We move to the if part. Suppose that Σ is not premise-connected and let C be a premise-connected component of Σ . We show that $(C, V \setminus C)$ is a split of Σ . Let $A \rightarrow b$ be an implication in Σ . If $A \subseteq C$ or A is a singleton element, it is clear that it satisfies Definition 1. Assume that $A \not\subseteq C$ and that A is not a singleton element. Recall that no implication of the form $\emptyset \rightarrow b$ lies in Σ . Let u, v be distinct elements in A and assume for contradiction $u \in C$ and $v \notin C$. Clearly, u, v is a premise path between u and v . Let w be any element of C . Since $u \in C$, u and w are premise connected. Consider any premise-path from w to u and append v to its end. The new path is a premise-path connecting w and v . Hence, $C \cup \{v\}$ is premise-connected, a contradiction with the fact that C is maximal. We deduce that $A \not\subseteq C$ implies that $A \cap C = \emptyset$. So $(C, V \setminus C)$ is indeed a split of Σ . \square

It is important to note that premise-connectivity is not inherited. That is, a sub-base induced by a premise-connected component need not be premise-connected in general.

Example 5. Consider the implicational base of Example 2 with the split $V_1 = \{1, 2, 3\}$, $V_2 = \{4, 5, 6, 7\}$. The elements 5 and 6 are premise-connected in Σ but not in $\Sigma[V_2] = \{5 \rightarrow 7, 45 \rightarrow 6\}$. This happens because the implication $56 \rightarrow 2$ is in $\Sigma[V_1, V_2]$.

Henceforth, premise-connected components of an implicational base may be further decomposed. Consequently, the split operation can be conducted in a recursive manner, leading to a hierarchical decomposition of implicational bases, up to trivial cases.

3.2 The decomposition tree of an implicational base

Based on the split operation, we define a hierarchical decomposition of an implicational base Σ . We call it a *H-decomposition* of Σ . The strategy is to recursively split Σ into smaller implicational bases until we reach trivial cases. This recursive decomposition can be conveniently represented by a full rooted binary tree T (full means that each node has precisely two children). An interior node of the tree corresponds to a split (V_1, V_2) of Σ whose children are H-decompositions of $\Sigma[V_1]$ and $\Sigma[V_2]$. The leaves of the tree represent the ground set V . Since the splits (V_1, V_2) and (V_2, V_1) are equivalent, the children of a node are unordered.

Definition 2 (Σ -tree and H-decomposition). *Let Σ be an implicational base over V and T be a full rooted binary tree. Then (T, λ) is a Σ -tree of Σ if there exists a labelling map $\lambda: T \rightarrow V \cup 2^\Sigma$ satisfying the following conditions:*

1. $\lambda(t)$ equals v for some $v \in V$ if t is a leaf of T ;
2. $\lambda(t) \subseteq \Sigma$ if t is an interior node (possibly $\lambda(t) = \emptyset$);
3. for every $A \rightarrow b \in \lambda(t)$, elements of A are labels of leaves in the subtree of one child of t and b is the label of a leaf in the subtree of the other child.
4. the set $\{\lambda(t) \mid t \in T\}$ is a full partition of $V \cup \Sigma$ and may contain the empty set.

If such labelling exists, we say that Σ is hierarchically decomposable (H-decomposable for short), and H-indecomposable otherwise.

In the particular case where $V = \emptyset$, we must have that $\Sigma = \emptyset$. If it happens, we say for convenience that Σ is trivially H-decomposable and that its Σ -tree is empty.

Example 6. The implicational base Σ from Example 2 is H-decomposable. In Figure 6, we represent a possible Σ -tree for Σ .




Figure 6: An Σ -tree for the implicational base of Example 2.

There are cases where a H-decomposition can be computed easily. For instance, if Σ is empty, every full rooted binary tree whose leaves are labelled by a permutation of V and every interior node by \emptyset is a Σ -tree. The case where Σ only contains implications of the form $a \rightarrow b$ for some $a, b \in V$ behaves similarly, except that the interior nodes of the tree contain the implications of Σ .

However, there are also some implicational bases that cannot be H-decomposed, for example when they admit no split at all. Next, our objective is to characterize H-decomposable implicational bases and devise a polynomial-time algorithm to build decomposition trees whenever possible. We first need two preparatory propositions.

Proposition 2. *An H-decomposable implicational base Σ is not premise-connected.*

Proof. Suppose that Σ is H-decomposable, and let (T, λ) be a Σ -tree with root r . Let (V_1, V_2) be the split of V corresponding to r , i.e. V_1 corresponds to the leaves of the left subtree of r and V_2 to those of the right subtree. Then, according to Proposition 1, Σ is not premise-connected. \square

Remark that the converse of Proposition 2 does not hold in general. We exhibit a counter-example. The main idea is to hide a premise-connected implicational base into a sub-base of a non premise-connected one.

Example 7. Let $V = \{1, 2, 3, 4\}$ and $\Sigma = \{12 \rightarrow 3, 13 \rightarrow 2, 23 \rightarrow 4\}$. The implicational base Σ has a unique split, $V_1 = \{1, 2, 3\}$ and $V_2 = \{4\}$. Thus it is not premise-connected and any possible Σ -tree must have the split (V_1, V_2) in the label of its root. After splitting, we are left with the sub-bases $\Sigma[V_2] = \emptyset$, $\Sigma[V_1, V_2] = \{23 \rightarrow 4\}$ and $\Sigma[V_1] = \{12 \rightarrow 3, 13 \rightarrow 2\}$. Observe that $\Sigma[V_1]$ is exactly the implicational base of Example 3. Hence, it is premise-connected and using Proposition 2, it cannot be H-decomposed. It follows that Σ admits no H-decomposition either.

Inspired by the previous example, we show that H-decomposability is hereditary, i.e. if an implicational base Σ has a Σ -tree then each of its sub-bases has a H-decomposition too.

Proposition 3. *Let Σ be an implicational base over V and let $X \subseteq V$. Then Σ has a H-decomposition only if $\Sigma[X]$ is H-decomposable.*

Proof. Let Σ be an implicational base over V , $X \subseteq V$, and let (T, λ) be a Σ -tree. If $X = \emptyset$, then the result trivially holds. We construct a subtree not necessarily induced by T which corresponds to a $\Sigma[X]$ -tree. We start from the root r of T and apply the following operation for each interior node t : if the sets of leaves of the left child and those of the right one both intersect X , keep t with label $\lambda(t) = \lambda(t) \cap \Sigma[X]$. Otherwise, there is a child of t whose set of leaves do not intersect X . In this case replace t by the child whose set of leaves intersects X . The obtained subtree has X as the set of label of its leaves, and the set of labels of the internal nodes are exactly $\Sigma[X]$. \square

The following theorem characterizes H-decomposability and gives the strategy of an algorithm computing a H-decomposition.

Theorem 1. *Let Σ be a non premise-connected implicational base and let C be a premise-connected component of Σ . Then Σ is H-decomposable if and only if $\Sigma[C]$ and $\Sigma[V \setminus C]$ are H-decomposable.*

Proof. The only if part directly follows from Proposition 3. Let us show the if part. Let C be a premise-connected component of Σ , (T_1, λ_1) be a $\Sigma[C]$ -tree and (T_2, λ_2) be a $\Sigma[V \setminus C]$ -tree. We consider a new tree (T, λ) such that T has root r with left subtree T_1 and right subtree T_2 . As for λ , we put $\lambda(t_1) = \lambda_1(t_1)$ if $t_1 \in T_1$, $\lambda(t_2) = \lambda_2(t_2)$ if $t_2 \in T_2$ and $\lambda(r) = \Sigma \setminus (\Sigma[C] \cup \Sigma[V \setminus C])$. In words, $\lambda(r)$ contains each implication whose premise is not fully contained in C or $V \setminus C$. It is clear that conditions (i), (ii), (iv) of Definition 2 are fulfilled for (T, λ) as they are for (T_1, λ_1) , (T_2, λ_2) and $C \cup V \setminus C = V$. Hence, we have to check (iii). Let $A \rightarrow b$ be an implication in $\lambda(v)$. If $A \cap C \neq \emptyset$, then $A \subseteq C$ since C is a premise-connected component of Σ . As $A \rightarrow b$ is not an implication of $\Sigma[C]$, it follows that $b \in V \setminus C$. Dually, if $A \cap C = \emptyset$, then $b \in C$ since $A \rightarrow b$ is not in $\Sigma[V \setminus C]$. Consequently, condition (iii) is satisfied and (T, λ) is a Σ -tree as required. \square

Theorem 1 suggests a recursive algorithm which returns a Σ -tree for an implicational base Σ if it is H-decomposable. If $V = \emptyset$, we simply output \emptyset . If V is a singleton element v , we output a leaf with label v . Otherwise, we compute a premise-connected component C of Σ if Σ is not premise-connected. We label the corresponding node by the implications of $\Sigma[C, V \setminus C]$, and we recursively call the algorithm on $\Sigma[C]$ and $\Sigma[V \setminus C]$. This strategy is formalized in Algorithm 1, whose correctness and complexity are studied in Theorem 2.

Algorithm 1: BuildTree.

Input: An implicational base Σ over V

Output: A Σ -tree, if it exists, **FAIL** otherwise

```
1 if  $V = \emptyset$  then
2   | return  $\emptyset$  ;
3 if  $V$  has one vertex  $v$  then
4   | create a new leaf  $r$  with appropriated  $\lambda(r)$ ;
5   | return  $r$  ;
6 else
7   | compute a premise-connected component  $C$  of  $\Sigma$  ;
8   | if  $|C| = |V|$  then
9     | stop and return FAIL ;
10  | else
11    | let  $r$  be a new node with  $\lambda(r) = \Sigma \setminus (\Sigma[C] \cup \Sigma[V \setminus C])$  ;
12    |  $\text{left}(r) = \text{BuildTree}(\Sigma[C])$  ;
13    |  $\text{right}(r) = \text{BuildTree}(\Sigma[V \setminus C])$  ;
14    | return  $r$  ;
```

Theorem 2. *Given an implicational base Σ over V , the Algorithm BuildTree computes a Σ -tree if it exists, in polynomial time and space in the size of Σ and V .*

Proof. First, we show by induction on $|V|$ that the algorithm returns a Σ -tree if and only if Σ is H-decomposable. Clearly if $V = \emptyset$, the algorithm returns \emptyset . In the case where V is reduced to a vertex v , the algorithm returns a Σ -tree corresponding to a leaf with label v .

Now, assume that the algorithm is correct for implicational bases with $|V| < n$, $n \in \mathbb{N}$, and consider a base Σ over V with $|V| = n$. Suppose Σ is H-decomposable. By Proposition 1, Σ is not premise-connected. Let C be a premise-connected component of Σ . Inductively, the algorithm is correct for $\Sigma[C]$ and $\Sigma[V \setminus C]$ since $1 \leq |C| < n$. From Theorem 1, we have that both $\Sigma[C]$ and $\Sigma[V \setminus C]$ are H-decomposable. By induction, the algorithm computes a $\Sigma[C]$ -tree (T_1, λ_1) and a $\Sigma[V \setminus C]$ -tree (T_2, λ_2) . Hence, the algorithm returns a labelled tree (T, λ) with root r whose label is $\lambda(r) = \Sigma \setminus (\Sigma[C] \cup \Sigma[V \setminus C])$ and children T_1 and T_2 . This tree satisfies all conditions to be a Σ -tree. Thus, the algorithm computes a Σ -tree for every H-decomposable implicational base.

Now suppose Σ is not H-decomposable. We have two cases:

1. Σ is premise-connected and the algorithm returns **FAIL** in Line 7.
2. Σ is not premise-connected. The algorithm chooses a premise-connected component C with $1 \leq |C| < n$. By Theorem 1, either $\Sigma[C]$ or $\Sigma[V \setminus C]$ is H-indecomposable. Thus, by induction, the algorithm will return **FAIL** for the input $\Sigma[C]$ or $\Sigma[V \setminus C]$ in lines 11-12. Since the algorithm stops, the output of the algorithm is **FAIL**.

Hence, the algorithm fails if the input Σ is H-indecomposable. We conclude that the algorithm returns a Σ -tree if and only if the input Σ is H-decomposable.

Finally, we show that the total time and space complexity of the algorithm are polynomial. The space required for the algorithm is bounded by the size of the implicational base Σ , the ground set V and the size of the Σ -tree. As the size of the Σ -tree is bounded by $O(|\Sigma| \times |V|)$, the overall space is bounded by $O(|\Sigma| \times |V|)$.

The time complexity is bounded by the sum of the costs of all nodes (or calls) of the search tree. The number of calls is bounded by $O(|V|)$, the size of the search tree. The cost of a call is

dominated by the computation of a premise-connected component of the input Σ . For this, we use union-find data structure of [42], which runs in almost linear time, *i.e.* $O(|\Sigma| \times |V| \times \alpha(|\Sigma| \times |V|, |V|))$ where $\alpha(\cdot, \cdot)$ is the inverse Ackermann function. The almost linear comes from the fact that $\alpha(|V|) \leq 4$ for every practical implicational base (see [42]). Thus, the total time complexity is $O(|V| \times (|\Sigma| \times |V| \times \alpha(|\Sigma| \times |V|, |V|)))$. \square

It is worth noticing, that the Σ -tree we obtain by the end of Algorithm 1 depends on the choice of a premise-connected component in line 5. As shown by the following example, the structure of the resulting Σ -tree is impacted by this choice.

Example 8. Let $V = \{1, 2, 3, 4, 5, 6, 7, 8\}$ and let Σ be the implicational base $\{12 \rightarrow 3, 3 \rightarrow 1, 23 \rightarrow 4, 34 \rightarrow 5, 56 \rightarrow 7, 67 \rightarrow 8\}$. For convenience, we represent Σ in Figure 7.




Figure 7: The implicational base of Example 8.

The premise-connected components of Σ are $\{1, 2, 3, 4\}$, $\{5, 6, 7\}$ and $\{8\}$. Thus, we can devise at least three distinct Σ -trees for Σ . In Figure 8, we give two of them. Observe that the first one (on the left) balances the size of labels of its interior nodes. On the other hand, the second one is a balanced tree.




Figure 8: Two Σ -trees for the implicational base of Example 8.

Following the previous example, a natural question arises: are all Σ -trees *equivalently interesting*? In particular, a balanced Σ -tree is a good candidate as the balancing is a common desirable property for decomposition trees to obtain efficient algorithms. This question, which uniquely depends on the syntax of the implicational base, is left open for further research.

3.3 Extension of the H-decomposition

As seen before, there are implicational bases that cannot have a split and thus a H-decomposition into trivial sub-bases. Such implicational bases are premise-connected, and will be called *irreducible H-factors* (H-factors for short). Now we describe a slight modification of Algorithm 1 to obtain a H-decomposition of implicational bases into H-factors. Instead of returning FAIL at line 7 in Algorithm BuildTree, we replace it by the following:

7' create a new leaf r with $\lambda(r) = \Sigma$ and return r ;

Example 9. Consider $V = \{1, 2, 3, 4, 5, 6\}$ and let $\Sigma = \{45 \rightarrow 1, 12 \rightarrow 3, 23 \rightarrow 1, 13 \rightarrow 2, 3 \rightarrow 6\}$. We represent Σ on the left of Figure 9. Clearly, Σ is not premise-connected and its premise-connected components are $\{4, 5\}$, $\{1, 2, 3\}$ and $\{3\}$. On the right of Figure 9, we present a H-decomposition of Σ into H-factors.




Figure 9: H-decomposition into H-factors.

With this modification, each possible implicational base has now a H-decomposition where leaves can be H-factors. To conclude this subsection, we show that H-factors are independent of the choice of the Σ -tree.

Proposition 4. *Let Σ be an implicational base over V and let (T_1, λ_1) and (T_2, λ_2) be two Σ -trees. Then, T_1 and T_2 have the same number of leaves and $\{\lambda_1(t_1) \mid t_1 \text{ is a leaf of } T_1\} = \{\lambda_2(t_2) \mid t_2 \text{ is a leaf of } T_2\}$.*

Proof. If Σ is H-decomposable or $(T_1, \lambda_1) = (T_2, \lambda_2)$, the result is clear due to Theorem 2. Assume that Σ is not H-decomposable and that the trees are different. Let t_1 be a leaf of T_1 such that $\lambda(t_1) = \Sigma_H$ is a H-factor of Σ . Let V_H be the set of elements spanned by Σ_H and let t_2 be the lowest node of T_2 such that $\Sigma_H \subseteq \bigcup\{\lambda(t'_2) \mid t'_2 \text{ is an ancestor of } t_2 \text{ in } T_2\}$. In other words, t_2 is the ancestor of all the elements in V_H . If t_2 is not a leaf, there exists a split in the sub-base induced by t_2 which separates the elements of V_H , a contradiction with Σ_H being a H-factor of Σ in (T_1, λ_1) . Hence, t_2 is also a leaf, and $\lambda_2(t_2) = \Sigma_H$ follows by applying the same reasoning in T_1 , which concludes the proof. \square

3.4 Splits and decomposition of a closure system

Naturally, the H-decomposition of an implicational base Σ induces a decomposition of the closure system \mathcal{C} defined by Σ . We also call the decomposition of \mathcal{C} a H-decomposition. The H-decomposition of \mathcal{C} is obtained from the H-decomposition of Σ , where the label of a node of its Σ -tree is replaced by the closure system associated to the implicational base induced by its subtree. The closure systems in leaves are the irreducible H-factors of the input closure system. Figure 10 illustrates the H-decomposition of the closure system associated to the H-decomposition of Example 9.

Theorem 3. *Let Σ be an implicational base over V with closure system \mathcal{C} , and let (V_1, V_2) be a split of Σ . Let \mathcal{C}_1 and \mathcal{C}_2 be closure systems associated to $\Sigma[V_1]$ and $\Sigma[V_2]$ (resp.). Then:*

1. $C \in \mathcal{C}$ implies that $C \cap V_1 \in \mathcal{C}_1$ and $C \cap V_2 \in \mathcal{C}_2$. Hence, $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$;
2. $\mathcal{C} = \mathcal{C}_1 \times \mathcal{C}_2$ holds whenever $\Sigma[V_1, V_2] = \emptyset$ (i.e. \mathcal{C} is the direct product of \mathcal{C}_1 and \mathcal{C}_2);
3. if for every implication $A \rightarrow b$ in $\Sigma[V_1, V_2]$, we have $A \subseteq V_1$, then $\mathcal{C} : V_1 = \mathcal{C}_1$ and $\mathcal{C} : V_2 = \mathcal{C}_2$; and
4. dually, if $A \subseteq V_2$ for every $A \rightarrow b$ in $\Sigma[V_1, V_2]$, we have $\mathcal{C} : V_1 = \mathcal{C}_1$ and $\mathcal{C} : V_2 = \mathcal{C}_2$.




Figure 10: H-decomposition of the closure system corresponding to Example 9.

Proof. Consider a split (V_1, V_2) of Σ , \mathcal{C}_1 and \mathcal{C}_2 the closure systems corresponding to $\Sigma[V_1]$ and $\Sigma[V_2]$. Their respective closure operators are ϕ_1, ϕ_2 . We prove items (i), (ii) and (iii). Items (iii) and (iv) are similar.

Item (i). Let $C \in \mathcal{C}$, $C_1 = C \cap V_1$ and let $A \rightarrow b$ be an implication of $\Sigma[V_1]$. Suppose $A \subseteq C_1$ and $b \notin C_1$. Then we also have $A \subseteq C$ and $b \notin C$ which contradicts $C \in \mathcal{C}$ as $A \rightarrow b \in \Sigma$. Thus $C_1 \in \mathcal{C}_1$. A similar reasoning applies to \mathcal{C}_2 , and $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$ holds.

Item (ii). We readily have that $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$ by item (i). For the other inclusion, let $C_1 \in \mathcal{C}_1$ and $C_2 \in \mathcal{C}_2$. We show that $C_1 \cup C_2 \in \mathcal{C}$. Let $A \rightarrow b$ be an implication of Σ with $A \subseteq C_1 \cup C_2$. As $\Sigma[V_1, V_2]$ is empty, $A \rightarrow b$ is either an implication of $\Sigma[V_1]$ or $\Sigma[V_2]$. As C_1, C_2 are closed for $\Sigma[V_1], \Sigma[V_2]$ (resp.), it follows that $C_1 \cup C_2 \in \mathcal{C}$.

Item (iii). Let $C_1 \in \mathcal{C}_1$. We show that $\phi(C_1)$ satisfies $\phi(C_1) \cap V_1 = C_1$. We readily have that $C_1 \subseteq \phi(C_1) \cap V_1$. Let $C_1 = X_0 \subset X_1 \subset \dots \subset X_k = \phi(C_1)$ be the sequence of sets obtained by applying the forward chaining algorithm on C_1 with Σ . We show by induction on $0 \leq i \leq k$ that $X_i \cap V_1 = C_1$. For the initial case $X_0 = C_1$, the result is clear. Now assume that the results holds true for any $0 \leq i < k$ and consider X_{i+1} . Let $A \rightarrow b$ be an implication such that $A \subseteq X_i$. Since (V_1, V_2) is a split of Σ , either $A \subseteq V_1$ or $A \subseteq V_2$. We have three cases

- (1) $A \subseteq V_2$. Then $A \rightarrow b \in \Sigma[V_2]$ and $b \in V_2$ so that $b \notin X_{i+1} \cap V_1$.
- (2) $A \rightarrow b$ is in $\Sigma[V_1]$. Then, $A \subseteq X_i \cap V_1$ which equals C_1 by inductive hypothesis. Since C_1 models $\Sigma[V_1]$ we have that $b \in X_i \cap V_1 = C_1$.
- (3) $A \rightarrow b$ is an implication of $\Sigma[V_1, V_2]$. Then $A \subseteq V_1$ and $b \in V_2$ since we assumed that every implication of $\Sigma[V_1, V_2]$ has its premise in V_1 and its conclusion in V_2 . Therefore, $b \notin X_{i+1} \cap V_1$.

Consequently $X_{i+1} \setminus X_i \subseteq V_2$, from which we deduce that $X_{i+1} \cap V_1 = C_1$, finishing the induction. Applying the result on $X_k = \phi(C_1)$, $\phi(C_1) \cap V_1 = C_1$ follows. So $C_1 \in \mathcal{C} : V_1$ and $C_1 \subseteq \mathcal{C} : V_1$. The reverse inclusion holds by item (i). As for \mathcal{C}_2 , we have $\mathcal{C}_2 \subseteq \mathcal{C}$ as $A \subseteq V_1$ for every implication $A \rightarrow b$ of $\Sigma[V_1, V_2]$. \square

According to Theorem 3 item (i), every closure system is a subset of the product of its H-factors closure systems. So the idea is to compute in parallel \mathcal{C}_1 and \mathcal{C}_2 for every split (V_1, V_2)

in the Σ -tree, and then use the bipartite implicational base $\Sigma[V_1, V_2]$ to compute \mathcal{C} . But this strategy is expensive, since the size of \mathcal{C}_1 and \mathcal{C}_2 may be exponential in the size of \mathcal{C} .

Example 10. Let $V = \{u_1, \dots, u_k, x, y\}$ for some $k \in \mathbb{N}$ and let $\Sigma = \bigcup\{u_i u_j \rightarrow x, u_i u_j \rightarrow y \mid 1 \leq i, j, \leq k, i \neq j\} \cup \{xy \rightarrow u_i \mid 1 \leq i \leq k\}$. Clearly, the unique possible split is $(V \setminus \{x, y\}, \{x, y\})$. Since $\Sigma[V \setminus \{x, y\}]$ is empty, its associated closure system is Boolean and has 2^k elements. However, $\mathcal{C} = \{v \mid v \in V\} \cup \{\{u, v\} \mid \{u, v\} \in (V \setminus \{x, y\}) \times \{x, y\}\} \cup \{\emptyset, V\}$ so that $|\mathcal{C}| = 3k + 4$.

However, this exponential reduction cannot occur when the sub-closure systems \mathcal{C}_1 and \mathcal{C}_2 appear as traces of \mathcal{C} .




Figure 11: Possible H-indecomposable factors.

To conclude this section, we relate H-decomposition to the subdirect product decomposition [23, 24]. Consider the closure system \mathcal{C} over $V = \{1, 2, 3\}$ in Figure 11(a) encoded by the implicational base $\{2 \rightarrow 1, 13 \rightarrow 2\}$. It is known that it cannot be decomposed using the subdirect product. Clearly Σ is not premise-connected and $V_1 = \{1, 3\}$ et $V_2 = \{2\}$ is the unique split where $\mathcal{C}_1 = \{\emptyset, 1, 3, 13\}$ and $\mathcal{C}_2 = \{\emptyset, 2\}$ are traces. Yet, \mathcal{C} is not a sublattice of $\mathcal{C}_1 \times \mathcal{C}_2$, since $\{1, 3\}$, the upper bound of 1 and 3 in $\mathcal{C}_1 \times \mathcal{C}_2$ is not preserved in \mathcal{C} . However, systems of Figure 11(b), (c) and (d) are both subdirectly irreducible and irreducible H-factors. Hence, we end the section with the following.

Corollary 1. *The closure system associated to an implicational base Σ is included in the direct product of its H-factors.*

Proof. This follows from Theorem 3, item (i) and the fact that a closure system is closed under intersection. \square

In the next section, we pay more attention to particular splits called *acyclic*. We show how they can be applied to the problem of translating between the representations of a closure system.

4 Closure systems with acyclic splits

In this section, we give a characterization of closure systems with acyclic splits. Then, we derive a recursive expression of their meet-irreducible elements. Finally, we devise an algorithm solving CCM in the case of acyclic splits. To illustrate our results, we will use the following running example all along the section.

Example 11 (Running example). Let $V = \{1, 2, 3, 4, 5, 6\}$ and $\Sigma = \{12 \rightarrow 3, 13 \rightarrow 4, 23 \rightarrow 5, 2 \rightarrow 4, 1 \rightarrow 5, 5 \rightarrow 6, 4 \rightarrow 6\}$. We represent Σ and its associated closure system \mathcal{C} in Figure 12.

The bipartition $V_1 = \{1, 2, 3\}$ and $V_2 = \{4, 5, 6\}$ is an *acyclic* split of Σ and \mathcal{C} : every implications has its premise included in V_1 and its conclusion in V_2 . We have $\Sigma[V_1] = \{12 \rightarrow 3\}$, $\Sigma[V_2] = \{4 \rightarrow 6, 5 \rightarrow 6\}$ and $\Sigma[V_1, V_2] = \{13 \rightarrow 4, 2 \rightarrow 4, 23 \rightarrow 5, 1 \rightarrow 5\}$.

We formally introduce *acyclic split* of an implicational base Σ . They are a restriction of a split (V_1, V_2) where all implications of $\Sigma[V_1, V_2]$ have to go from V_1 to V_2 , *i.e.* they satisfy




Figure 12: An implicational base and its associated closure system.

condition (iii) or (iv) of Theorem 3. The definition of acyclic split for implicational bases extend to closure systems.

Definition 3 (Acyclic split). *Let Σ be an implicational base over V and (V_1, V_2) a split of Σ . The split (V_1, V_2) is acyclic if for every $A \rightarrow b \in \Sigma[V_1, V_2]$, $A \subseteq V_1$.*

Definition 4 (Acyclic split of a closure system). *Let \mathcal{C} be a closure system over V and let (V_1, V_2) be a non-trivial bipartition of V such that $V_2 \in \mathcal{C}$. Then, (V_1, V_2) is an acyclic split of \mathcal{C} if there exists an implicational base Σ for \mathcal{C} with acyclic split (V_1, V_2) .*

4.1 Acyclic split of a closure system

Let Σ be an implicational base over V with acyclic split (V_1, V_2) . Let \mathcal{C} be its corresponding closure system. We first show how to construct \mathcal{C} from \mathcal{C}_1 , the closure system associated to $\Sigma[V_1, V_2]$, \mathcal{C}_2 , the closure system of $\Sigma[V_2]$ and the implications $\Sigma[V_1, V_2]$.

We draw intuition from the particular case where $\Sigma[V_1, V_2] = \emptyset$. According to Theorem 3, \mathcal{C} is the direct product of \mathcal{C}_1 and \mathcal{C}_2 , that is $\mathcal{C} = \{C_1 \cup C_2 \mid C_1 \in \mathcal{C}_1, C_2 \in \mathcal{C}_2\}$. Intuitively, \mathcal{C} is obtained by “extending” each closed set of \mathcal{C}_2 with a copy of \mathcal{C}_1 (see the left part of Figure 13). This point of view will be particularly well-suited for us, and naturally leads to the following definition.

Definition 5. *Let \mathcal{C} be a closure system over V , (V_1, V_2) be a non-trivial bipartition of V such that $V_2 \in \mathcal{C}$. Let $C_2 \in \mathcal{C}$, $C_2 \subseteq V_2$ and $C \in \mathcal{C}$. We say that C is an extension of C_2 with respect to V_2 if $C \cap V_2 = C_2$. We denote by $\text{Ext}(C_2)$ the extensions of C_2 in \mathcal{C} . The trace $\text{Ext}(C_2)$ on V_1 is written $\text{Ext}(C_2): V_1$.*

In our definition, V_2 is closed. Therefore, for every $C \in \mathcal{C}$, $C \cap V_2$ is also closed. We deduce that C belongs to the extension of a unique closed set C_2 included in V_2 . As a consequence, we can write \mathcal{C} as the (disjoint) union of its extensions with respect to V_2 , *i.e.*

$$\mathcal{C} = \bigcup_{C_2 \in \mathcal{C}, C_2 \subseteq V_2} \text{Ext}(C_2)$$

This definition of extensions allows to formally express the intuition that the direct product of \mathcal{C}_1 and \mathcal{C}_2 (when $\Sigma[V_1, V_2] = \emptyset$) is obtained by extending each closed set of \mathcal{C}_2 with a copy of \mathcal{C}_1 .




Figure 13: Building of \mathcal{C} with an acyclic split: on the left, the case where $\Sigma[V_1, V_2] = \emptyset$ (direct product). On the right, the more general case where $\Sigma[V_1, V_2] \neq \emptyset$ (increasing extensions).

Indeed, we have $\mathcal{C} = \bigcup_{C_2 \in \mathcal{C}_2} \text{Ext}(C_2)$ with the particularity that the trace of $\text{Ext}(C_2)$ on V_1 is exactly C_1 for every $C_2 \in \mathcal{C}_2$. This construction is illustrated on the left of Figure 13.

In the more general case where $\Sigma[V_1, V_2]$ is not-empty, we show that the extensions of C_2 are no longer full copies of C_1 , but increasing copies of ideals of C_1 , as illustrated on the right side of Figure 13. We begin with the following proposition, which characterizes extensions with the bipartite set of implications $\Sigma[V_1, V_2]$.

Proposition 5. *Let $C_2 \in \mathcal{C}_2$ and $C_1 \subseteq V_1$. Then, $C = C_1 \cup C_2$ is an extension of C_2 if and only if $C_1 \in \mathcal{C}_1$ and for each implication $A \rightarrow b$ in $\Sigma[V_1, V_2]$, $A \subseteq C_1$ implies $b \in C_2$.*

Proof. We begin with the only if part. Let C_1 be a subset of V_1 such that let C_1 be a closed set of \mathcal{C}_1 such that $C_1 \cup C_2$ is an extension of C_2 . By Theorem 3, $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$ so that for every $C_1 \subseteq V_1$ such that $C_1 \cup C_2 \in \mathcal{C}$, $C_1 \in \mathcal{C}_1$ holds. Now let $A \rightarrow b \in \Sigma[V_1, V_2]$. If $A \subseteq C_1$, it must be that $b \in C_2$ since we would contradict $C_1 \cup C_2 \in \mathcal{C}$ otherwise.

We move to the if part. Let C_1 be a closed set of \mathcal{C}_1 and C_2 a closed set of \mathcal{C}_2 such that for each implication $A \rightarrow b$ in $\Sigma[V_1, V_2]$, $A \subseteq C_1$ implies $b \in C_2$. We have to show that $C_1 \cup C_2$ is closed. Let $A \rightarrow b$ be an implication of Σ with $A \subseteq C_1 \cup C_2$. As (V_1, V_2) is an acyclic split of V , we have two cases: either $A \rightarrow b$ is in $\Sigma[V_1, V_2]$ or it is not. In the second case, assume $A \rightarrow b$ belongs to $\Sigma[V_1]$. As $A \subseteq C_1 \cup C_2$, we have $A \subseteq C_1$. Furthermore, C_1 is closed for $\Sigma[V_1]$. Hence, $b \in C_1 \subseteq C_1 \cup C_2$. The same reasoning can be applied if $A \rightarrow b$ is in $\Sigma[V_2]$. Now assume $A \rightarrow b$ is in $\Sigma[V_1, V_2]$. We have that $A \subseteq V_1$ by definition of an acyclic split. In particular, we have $A \subseteq C_1$ which entails $b \in C_2$ by assumption. In any case, $C_1 \cup C_2$ already contains b for every implication $A \rightarrow b$ in Σ such that $A \subseteq C_1 \cup C_2$. Hence, $C_1 \cup C_2$ is closed. \square

We readily deduce from Proposition 5 that $\text{Ext}(V_2): V_1$ is equal to C_1 . Proposition 5 is also a step towards the next proposition. It settles the fact that in an acyclic split, extensions coincide with ideals of \mathcal{C}_1 .

Proposition 6. *Let $C_1 \in \mathcal{C}_1$, $C_2 \in \mathcal{C}_2$. If $C_1 \cup C_2$ is an extension of C_2 , then for every $C'_1 \in \mathcal{C}_1$ such that $C'_1 \subseteq C_1$, $C'_1 \cup C_2$ is also an extension of C_2 .*

Proof. Let $C_1 \in \mathcal{C}_1$, $C_2 \in \mathcal{C}_2$ such that $C_1 \cup C_2 \in \mathcal{C}$. Let $C'_1 \in \mathcal{C}_1$ such that $C'_1 \subseteq C_1$. As $C_1 \cup C_2$ is an extension of C_2 , for each $A \rightarrow b$ in $\Sigma[V_1, V_2]$ such that $A \subseteq C_1$, we have $b \in C_2$ by Proposition

5. Since $C'_1 \subseteq C_1$, this condition holds in particular if $A \subseteq C'_1$. Applying Proposition 5, we deduce that $C'_1 \cup C_2$ is closed. \square

In fact, the preceding proposition can be further strengthened. Not only extensions of \mathcal{C}_2 correspond to ideals of \mathcal{C}_1 , but they are increasing. That is, if C_1 contributes to an extension of C_2 , it will also contribute to an extension of any closed set $C'_2 \in \mathcal{C}_2$ including C_2 .

Lemma 1. *Let $C_2, C'_2 \in \mathcal{C}_2$ such that $C_2 \subseteq C'_2$. Then $\text{Ext}(C_2): V_1 \subseteq \text{Ext}(C'_2): V_1$.*

Proof. We need to show that for every $C_2, C'_2 \in \mathcal{C}_2$ such that $C_2 \subseteq C'_2$, if $C_1 \cup C_2 \in \mathcal{C}$ for some $C_1 \subseteq V_1$, we also have $C_1 \cup C'_2 \in \mathcal{C}$. Observe that due to Proposition 5, $C_1 \in \mathcal{C}_1$. As $C_1 \cup C_2$ is an extension of C_2 , for every implication $A \rightarrow b$ of $\Sigma[V_1, V_2]$ such that $A \subseteq C_1$, we have $b \in C_2 \subseteq C'_2$ by Proposition 5. Therefore, $C_1 \cup C'_2$ is indeed an extension of C'_2 . \square

Corollary 2. *Let $C_2, C'_2 \in \mathcal{C}_2$ such that $C_2 \prec C'_2$ and let $C_1 \in \mathcal{C}_1$ such that $C_1 \cup C_2 \in \mathcal{C}$. Then $C_1 \cup C'_2 \in \mathcal{C}$ and $C_1 \cup C_2 \prec C_1 \cup C'_2$.*

Proof. The fact that $C_1 \cup C'_2$ is closed follows from Lemma 3. By Theorem 3, $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$ so that any closed set C such that $C_1 \cup C_2 \subset C \subseteq C_1 \cup C'_2$ satisfies $C \cap V_2 \in \mathcal{C}_2$. Since $C_2 \prec C'_2$ in \mathcal{C}_2 , $C = C_1 \cup C'_2$ follows. \square

Thus, we have shown that if (V_1, V_2) is an acyclic split of Σ , \mathcal{C} can be constructed by extending each closed set C_2 of \mathcal{C}_2 , with an ideal of \mathcal{C}_1 , in an increasing fashion. This construction is illustrated in Figure 13 and in Figure 15 on an example. In the next theorem, we demonstrate that this construction by increasing extensions is in fact a characterization of acyclic splits.

Theorem 4. *Let \mathcal{C} be a closure system over V and (V_1, V_2) be a non-trivial bipartition of V such that $V_2 \in \mathcal{C}$. Let $\mathcal{C}_1 = \uparrow V_2: V_1$ and $\mathcal{C}_2 = \downarrow V_2$. Then, (V_1, V_2) is an acyclic split for \mathcal{C} if and only if for every $C_2, C'_2 \in \mathcal{C}_2$ such that $C_2 \subseteq C'_2$, we have $\text{Ext}(C_2): V_1 \subseteq \text{Ext}(C'_2): V_1$.*

Proof. The only if part follows from Lemma 1. To show the if part, we build an implicational base Σ with the acyclic split (V_1, V_2) . Beforehand, we outline the main ideas:

- Σ should contain an implicational base for \mathcal{C}_2 as it is an ideal of \mathcal{C} ;
- Σ should also include an implicational base for \mathcal{C}_1 since it is a filter of \mathcal{C} and Σ must respect the split (V_1, V_2) ;
- Σ must describe, for each $C_2 \in \mathcal{C}_2$, which closed sets of \mathcal{C}_1 contribute to extensions of C_2 or not. The most direct way to express this relationship is to explicitly write it in Σ by putting implications $C_1 \rightarrow \phi(C_1) \cap V_2$, if C_1 does not participate in an extension of C_2 .

Actually, we can readily optimize the last item. Indeed, since the property of not contributing to an extension is monotone, it is sufficient to put an implication $C_1 \rightarrow \phi(C_1) \cap V_2$ if C_1 is a minimal closed set of \mathcal{C}_1 which does not yield an extension of C_2 .

With these ideas in mind, we proceed now to the proof. Let $\mathcal{C}_1 = \uparrow V_2: V_1$ and $\mathcal{C}_2 = \downarrow V_2$. Observe that both \mathcal{C}_1 and \mathcal{C}_2 are closure systems. We aim to construct an implicational base Σ representing \mathcal{C} with acyclic split (V_1, V_2) .

First, we prove that $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$. Let $C \in \mathcal{C}$ and let $C_1 = C \cap V_1$ and $C_2 = C \cap V_2$. As C and V_2 are closed in \mathcal{C} we deduce that $C_2 \in \mathcal{C}_2$ and hence that $C \in \text{Ext}(C_2)$. As $C_2 \subseteq V_2$, we have $\text{Ext}(C_2): V_1 \subseteq \text{Ext}(V_2): V_1$ with $\text{Ext}(V_2): V_1 = \mathcal{C}_1$ by assumption. Hence $C_1 \in \mathcal{C}_1$. We deduce that $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$.

Now, let $\Sigma[V_1]$ be an implicational base for \mathcal{C}_1 , $\Sigma[V_2]$ an implicational base for \mathcal{C}_2 and let

$$\Sigma[V_1, V_2] = \{C_1 \rightarrow \phi(C_1) \cap V_2 \mid C_1 \in \min_{\subseteq}(\mathcal{C}_1 \setminus \text{Ext}(C_2): V_1) \text{ for some } C_2 \in \mathcal{C}_2\}$$

Finally we put $\Sigma = \Sigma[V_1, V_2] \cup \Sigma[V_1] \cup \Sigma[V_2]$. Clearly (V_1, V_2) is an acyclic split for Σ . We prove that Σ is an implicational base for \mathcal{C} . Let \mathcal{C}_Σ be the closure system associated to Σ .

To show that $\mathcal{C}_\Sigma \subseteq \mathcal{C}$, we prove that $C \notin \mathcal{C}$ entails $C \notin \mathcal{C}_\Sigma$, for every $C \subseteq V$. Let $C \subseteq V$ such that $C \notin \mathcal{C}$ and put $C_1 = C \cap V_1$ and $C_2 = C \cap V_2$. First, assume that $C \notin \mathcal{C}_1 \times \mathcal{C}_2$. Since $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$, $C \notin \mathcal{C}$ readily holds. Then, $C_1 \notin \mathcal{C}_1$ or $C_2 \notin \mathcal{C}_2$ so that C fails $\Sigma[V_1]$ or $\Sigma[V_2]$ and $C \notin \mathcal{C}_\Sigma$ holds. Now assume that $C \in \mathcal{C}_1 \times \mathcal{C}_2$ but $C \notin \mathcal{C}$. By construction of \mathcal{C} , we have that $C \notin \text{Ext}(C_2)$, or equivalently, $C_1 \notin \text{Ext}(C_2): V_1$. Let $C'_1 \in \mathcal{C}_1$ with $C'_1 \subseteq C_1$ and $C'_1 \in \min_{\subseteq}(\mathcal{C}_1 \setminus \text{Ext}(C_2): V_1)$. We show that C fails the implication $C'_1 \rightarrow \phi(C'_1) \cap V_2$ of $\Sigma[V_1, V_2]$. We have $\phi(C'_1) \in \mathcal{C}$ so that $\phi(C'_1) \cap V_2 \in \mathcal{C}_2$ and $C'_1 \in \text{Ext}(\phi(C'_1) \cap V_2): V_1$. By assumption, for every closed set $C''_2 \in \mathcal{C}_2$ such that $\phi(C'_1) \cap V_2 \subseteq C''_2$, $\text{Ext}(\phi(C'_1) \cap V_2): V_1 \subseteq \text{Ext}(C''_2): V_1$. Therefore, $C'_1 \notin \text{Ext}(C_2): V_1$ implies that $\phi(C'_1) \cap V_2 \not\subseteq C_2$. Consequently, $C'_1 \subseteq C_1 \subseteq C$ but $\phi(C'_1) \cap V_2 \not\subseteq C \cap V_2 = C_2$. We deduce that $C \notin \mathcal{C}_\Sigma$, and hence that $\mathcal{C}_\Sigma \subseteq \mathcal{C}$.

Now we demonstrate that $\mathcal{C} \subseteq \mathcal{C}_\Sigma$. Let $C \in \mathcal{C}$ and put $C_1 = C \cap V_1$, $C_2 = C \cap V_2$. Recall that $\mathcal{C}_2 = \downarrow V_2$ and that $\Sigma[V_2]$ is an implicational base for \mathcal{C}_2 . Therefore, $C_2 \in \mathcal{C}_2$ and C is a model of $\Sigma[V_2]$ since $C_2 \subseteq C$. Now, because $C_2 \subseteq V_2$, we have $\text{Ext}(C_2): V_1 \subseteq \text{Ext}(V_2): V_1 = \mathcal{C}_1$ by assumption. Moreover, $\Sigma[V_1]$ is an implicational base for \mathcal{C}_1 . Consequently, we obtain that $C_1 \in \mathcal{C}_1$ and hence that C is a model for $\Sigma[V_1]$. It remains to show that C also models $\Sigma[V_1, V_2]$. But this is clear as $C = \phi(C)$ and each implication $C_1 \rightarrow \phi(C'_1) \cap V_2$ of $\Sigma[V_1, V_2]$ satisfies $\phi(C'_1) \cap V_2 \subseteq \phi(C'_1)$. Hence, $C'_1 \subseteq C$ implies that $\phi(C'_1) \subseteq C$. Consequently, $\mathcal{C} \subseteq \mathcal{C}_\Sigma$ and $\mathcal{C} = \mathcal{C}_\Sigma$ holds, concluding the proof. \square

Example 12 (Running example). The closure system \mathcal{C}_1 associated to $\Sigma[V_1] = \{12 \rightarrow 3\}$ is given on the left of Figure 14. On the right, we give \mathcal{C}_2 , the closure system of $\Sigma[V_2] = \{4 \rightarrow 6, 5 \rightarrow 6\}$.




Figure 14: The closure systems \mathcal{C}_1 and \mathcal{C}_2 .

The construction of \mathcal{C} using extensions with respect to \mathcal{C}_1 and \mathcal{C}_2 suggested by Theorem 4 is highlighted in Figure 15. For instance, the extensions of 6 are \emptyset and 36. Remark that \emptyset and 3 also contribute to the extensions 46, 346 of 46. Moreover, 346 is a maximal extension of 46, along with 246. Finally, the extensions of 456 (that is, V_2) coincide with \mathcal{C}_1 .

In the particular case where \mathcal{C} is a direct product of $\mathcal{C}_1, \mathcal{C}_2$, the pair (V_1, V_2) becomes a strong decomposition pair of [34]. It is worth noticing that Theorem 4 hints a strategy to recursively compute the meet-irreducible elements of \mathcal{C} . This is the aim of the next subsection.

4.2 The meet-irreducible elements of a closure system with acyclic split

Now we use Theorem 4 to obtain a recursive expression of \mathcal{M} , the meet-irreducible elements of \mathcal{C} in terms of \mathcal{M}_1 and \mathcal{M}_2 , the meet-irreducible elements of \mathcal{C}_1 and \mathcal{C}_2 respectively. We prove that the decomposition of \mathcal{C} with extensions captures the structure of \mathcal{M} . Again, we start from the case of the direct product. This result has already been formulated in lattice theory, for instance in [11]. We give a proof in our framework for self-containment.

Proposition 7. *Let \mathcal{C}_1 and \mathcal{C}_2 be two closure systems over V_1 and V_2 (resp.) where V_1 and V_2 are disjoint. Let $\mathcal{C} = \mathcal{C}_1 \times \mathcal{C}_2$. Then $\mathcal{M} = \{M_1 \cup V_2 \mid M_1 \in \mathcal{M}_1\} \cup \{M_2 \cup V_1 \mid M_2 \in \mathcal{M}_2\}$.*




Figure 15: The closure \mathcal{C} constructed from \mathcal{C}_1 and \mathcal{C}_2 (black dots are closed set of \mathcal{C}_2).

Proof. Let $M \in \mathcal{M}$, $M_1 = M \cap V_1$ and $M_2 = M \cap V_2$. Since $\mathcal{C} = \mathcal{C}_1 \times \mathcal{C}_2$, we have $M_1 \in \mathcal{C}_1$ and $M_2 \in \mathcal{C}_2$. As $M \neq V_1 \cup V_2$, either $V_1 \not\subseteq M$ or $V_2 \not\subseteq M$. Suppose both statements hold. Then, there exists $C_1 \in \mathcal{C}_1$ such that $M_1 \prec C_1$ in \mathcal{C}_1 . Similarly, there exists $C_2 \in \mathcal{C}_2$ such that $M_2 \prec C_2$ in \mathcal{C}_2 . However $\mathcal{C} = \mathcal{C}_1 \times \mathcal{C}_2$. Hence, $M_1 \cup C_2$ and $C_1 \cup M_2$ belong to \mathcal{C} . Furthermore, they are incomparable and we have $M \prec M_1 \cup C_2$ and $M \prec C_1 \cup M_2$ which contradicts $M \in \mathcal{M}$. Therefore, either $V_1 \subseteq M$ or $V_2 \subseteq M$. Assume without loss of generality that $V_1 \subseteq M$. Let M' be the unique cover of M in \mathcal{C} . Then, $V_1 \subseteq M'$ and it follows that $M_2 \prec M' \cap V_2$ in \mathcal{C}_2 . As M' is the unique cover of M in \mathcal{C} , we conclude that $M' \cap V_2$ is the unique cover of M_2 in \mathcal{C}_2 and $M_2 \in \mathcal{C}_2$.

Let $M_1 \in \mathcal{M}_1$ and consider $M_1 \cup V_2 \in \mathcal{C}_2$. Let M'_1 be the unique cover of M_1 in \mathcal{C}_1 . As $\mathcal{C} = \mathcal{C}_1 \times \mathcal{C}_2$, we have that $M_1 \cup V_2 \prec M'_1 \cup V_2$ is in \mathcal{C} . Let C be any closed set such that $M_1 \cup V_2 \subset C$. We have $C \cap V_2 = V_2$ and hence $M_1 \subset C \cap V_1$. Since $\mathcal{C} = \mathcal{C}_1 \times \mathcal{C}_2$, we get $C \cap V_1 \in \mathcal{C}_1$. As $M_1 \prec M'_1$ in \mathcal{C}_1 and $M_1 \in \mathcal{M}_1$, we conclude that $M'_1 \subseteq C \cap V_1$ and hence that $M'_1 \cap V_2 \subseteq C$. Therefore, $M_1 \cup V_2 \in \mathcal{M}$. Similarly, we obtain $M_2 \cup V_1 \in \mathcal{M}$, for $M_2 \in \mathcal{M}_2$. \square




Figure 16: Meet-irreducible elements of \mathcal{C} with an acyclic split: on the left, the direct product. On the right, the case of acyclic splits in general.

If we adopt the point of view of extensions with respect to \mathcal{C}_2 , as in the previous subsection, the meet-irreducible elements of $\mathcal{C}_1 \times \mathcal{C}_2$ can be partitioned into two classes:

- (1) those belonging to extensions of V_2 , that is $\{M_1 \cup V_2 \mid M_1 \in \mathcal{M}_1\}$;
- (2) meet-irreducible elements of \mathcal{M}_2 which we extended with V_1 , that is $\{M_2 \cup V_1 \mid M_2 \in \mathcal{M}_2\}$.
Observe that $M_2 \cup V_1$ is the unique inclusion-wise maximal extension of M_2 , for each $M_2 \in \mathcal{M}_2$.

This construction is illustrated on the left part of Figure 16.

We show next that when \mathcal{C} has an acyclic split (V_1, V_2) but it is not the direct product of \mathcal{C}_1 and \mathcal{C}_2 , the structure of \mathcal{M} preserves this partitioning:

- (1) $\{M_1 \cup V_2 \mid M_1 \in \mathcal{M}_1\}$ remains unchanged;
- (2) $\{M_2 \cup V_1 \mid M_2 \in \mathcal{M}_2\}$ is adapted to replace $M_2 \cup V_1$ by the possible maximal extensions of elements of \mathcal{M}_2 .

This construction is represented on the right of Figure 16. Let \mathcal{C} be a closure system with acyclic split (V_1, V_2) . Again, let $\mathcal{C}_1 = \uparrow V_2: V_1$ and $\mathcal{C}_2 = \downarrow V_2$. We begin with the following two lemmas.

Lemma 2. *Let $C_2 \in \mathcal{C}_2, C_2 \neq V_2$ and $C_1 \in \mathcal{C}_1$ such that $C_1 \cup C_2$ is a non-maximal extension of C_2 . Then $C_1 \cup C_2 \notin \mathcal{M}$.*

Proof. Let $C_2 \in \mathcal{C}_2, C_2 \neq V_2$ and $C_1 \in \mathcal{C}_1$ such that $C_1 \cup C_2$ is a non-maximal extension of C_2 . As $C_2 \neq V_2$, there exists at least one closed set $C'_2 \in \mathcal{C}_2$ such that $C_2 \prec C'_2$. By Corollary 2 we have that $C_1 \cup C_2 \prec C_1 \cup C'_2$ in \mathcal{C} . Furthermore, $C_1 \cup C_2$ is not a maximal extension of C_2 . Therefore, there exists a closed set C'_1 in \mathcal{C}_1 such that $C_1 \prec C'_1$ and $C'_1 \cup C_2 \in \mathcal{C}$. As $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$ by Theorem 4 and extensions are increasing by Lemma 1, it follows that $C_1 \cup C_2 \prec C'_1 \cup C_2$ in \mathcal{C} with $C_1 \cup C'_2 \neq C'_1 \cup C_2$. Therefore, $C_1 \cup C_2$ is not a meet-irreducible element of \mathcal{C} . \square

Lemma 3. *Let $C_2 \in \mathcal{C}_2$ such that $C_2 \neq V_2$ and $C_2 \notin \mathcal{M}_2$. Then $C \notin \mathcal{M}$ for every $C \in \text{Ext}(C_2)$.*

Proof. Let $C_2 \in \mathcal{C}_2$ such that $C_2 \neq V_2$ and $C_2 \notin \mathcal{M}_2$. Let $C \in \text{Ext}(C_2)$ and $C_1 = C \cap V_1$. As $C_2 \notin \mathcal{M}_2$, it has at least two covers C'_2, C''_2 in \mathcal{C}_2 . By Corollary 2, it follows that both $C'_2 \cup C_1$ and $C''_2 \cup C_1$ are covers of C in \mathcal{C} . Hence $C \notin \mathcal{M}$. \square

These lemmas suggest that meet-irreducible elements of \mathcal{C} arise from maximal extensions of meet-irreducible elements of \mathcal{C}_2 . They might also come from meet-irreducible extensions of V_2 since $\text{Ext}(V_2): V_1 = \mathcal{C}_1$. These ideas are proved in the following theorem, which characterizes the meet-irreducible elements \mathcal{M} of \mathcal{C} according to the two types we described.

Theorem 5. *Let \mathcal{C} be a closure system over V with acyclic split (V_1, V_2) . Let $\mathcal{C}_1 = \uparrow V_2: V_1$ and $\mathcal{C}_2 = \downarrow V_2$. Meet-irreducible elements \mathcal{M} of \mathcal{C} satisfy $|\mathcal{M}| \geq |\mathcal{M}_1| + |\mathcal{M}_2|$ and are subject to the following equality:*

$$\mathcal{M} = \{M_1 \cup V_2 \mid M_1 \in \mathcal{M}_1\} \cup \{C \in \max_{\subseteq}(\text{Ext}(M_2)) \mid M_2 \in \mathcal{M}_2\}$$

Proof. First, $\{M_1 \cup V_2 \mid M_1 \in \mathcal{M}_1\} \subseteq \mathcal{M}$ follows from the fact that $\mathcal{C}_1 = \uparrow V_2: V_1$. We prove that $\max_{\subseteq}(\text{Ext}(M_2)) \subseteq \mathcal{M}$ for every $M_2 \in \mathcal{M}_2$. Let $M_2 \in \mathcal{M}_2$ and let C be a maximal extension of M_2 with $C = C_1 \cup M_2$. Since $M_2 \in \mathcal{C}_2$, it has a unique cover M'_2 in \mathcal{C}_2 . By Corollary 2, we get $C \prec M'_2 \cup C_1$ in \mathcal{C} . Let $C' \in \mathcal{C}$ such that $C \subset C'$. Recall that $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$ follows from Theorem 4, so that $C' \cap V_1 \in \mathcal{C}_1$ and $C' \cap V_2 \in \mathcal{C}_2$. Furthermore, $C \in \max_{\subseteq}(\text{Ext}(M_2))$, therefore $C \subset C'$ implies that $M_2 \subset C' \cap V_2$ and hence that $M'_2 \subseteq C' \cap V_2$ as $M_2 \in \mathcal{C}_2$. Since $C_1 \subseteq C' \cap V_1$, we get $C \prec M'_2 \cup C_1 \subseteq C'$ and $C \in \mathcal{M}$ as it has a unique cover.

Now we prove the other side of the equation. Let $M \in \mathcal{M}$. As $\mathcal{C} \subseteq \mathcal{C}_1 \times \mathcal{C}_2$ since (V_1, V_2) is an acyclic split of \mathcal{C} , $M \cap V_2 \in \mathcal{C}_2$ and we can distinguish two cases. Either $M \cap V_2 = V_2$ or $M \cap V_2 \subset V_2$. If $M \cap V_2 = V_2$ then M is a meet-irreducible element of the closure system $\uparrow V_2$. Since $\uparrow V_2: V_1 = \mathcal{C}_1$, we obtain that $M \cap V_1 = M_1 \in \mathcal{M}_1$. Now assume that $M \cap V_2 \subset V_2$. Let $M_1 = M \cap V_1$ and $M_2 = M \cap V_2$. Then by contrapositive of Lemma 2 we have that $M \in \max_{\subseteq}(\text{Ext}(M_2))$ as $M_2 \neq V_2$. Similarly, we get $M_2 \in \mathcal{M}_2$ by Lemma 3. The inequality $|\mathcal{M}| \geq |\mathcal{M}_1| + |\mathcal{M}_2|$ follows from the description of \mathcal{M} . \square

Example 13 (Running example). The meet-irreducible elements \mathcal{M}_1 of \mathcal{C}_1 are 1, 13, 2 and 23. Similarly, the meet-irreducible elements of \mathcal{C}_2 are \emptyset , 46 and 56. In Figure 17 we highlight the two types of meet-irreducible elements of \mathcal{C} , based on Theorem 5. For instance 23456 is of type (1) as it is obtained from the meet-irreducible element 23 of \mathcal{C}_1 and V_2 . Dually, 356 is of type (2) because it is a maximal extension of the meet-irreducible element 56 of \mathcal{C}_2 .




Figure 17: The two types of meet-irreducible elements in \mathcal{C} (black dots are closed sets of \mathcal{C}_2).

To conclude this section, we briefly discuss another characterization of acyclic splits based on Theorem 4 and Theorem 5. Because extensions are hereditary, the extensions of \mathcal{M}_2 completely capture extensions of \mathcal{C}_2 . In other words, if $C_2 \in \mathcal{C}_2$ and C_1 contributes to an extension of C_2 , then $C_1 \cup M_2$ is also an extension of M_2 , for every $M_2 \in \mathcal{M}_2(C_2)$. Therefore, $C_1 \cup C_2$ results from the intersection of the closed sets $M_2 \cup C_1$, $M_2 \in \mathcal{M}_2(C_2)$. We illustrate this idea in Figure 18.

Corollary 3. *Let \mathcal{C} be a closure system over V and (V_1, V_2) a non-trivial bipartition of V with $V_2 \in \mathcal{C}$. Let $\mathcal{C}_1 = \uparrow V_2: V_1$ and $\mathcal{C}_2 = \downarrow V_2$. The pair (V_1, V_2) is an acyclic split for \mathcal{C} if and only if for every $C_2 \in \mathcal{C}_2$ and $C'_2 \in \mathcal{M}_2(C_2) \cup \{V_2\}$, $\text{Ext}(C_2): V_1 \subseteq \text{Ext}(C'_2): V_1$.*

Proof. The only if part follows from Theorem 4. Let $C_2, C'_2 \in \mathcal{C}_2$ with $C_2 \subseteq C'_2$. If $C_2 = V_2$ or $C'_2 = V_2$, the fact that $\text{Ext}(C_2): V_1 \subseteq \text{Ext}(C'_2): V_1$ is clear. Assume that $C_2 \subseteq C'_2 \subset V_2$ so that $\mathcal{M}_2(C_2)$ and $\mathcal{M}_2(C'_2)$ are not empty. From $C_2 \subseteq C'_2$, we deduce $\mathcal{M}_2(C'_2) \subseteq \mathcal{M}_2(C_2)$. Let $C \in \text{Ext}(C_2)$ with $C_1 = C \cap V_1$. Remark that $C_1 \in \mathcal{C}_1$ holds by assumption. Moreover, for every $M_2 \in \mathcal{M}_2(C_2)$, we have $C_1 \cup M_2 \in \text{Ext}(M_2)$. This holds in particular for every $M_2 \in \mathcal{M}_2(C'_2)$ so that $\bigcap_{M_2 \in \mathcal{M}_2(C'_2)} (M_2 \cup C_1) = (\bigcap_{M_2 \in \mathcal{M}_2(C'_2)} M_2) \cup C_1 = C'_2 \cup C_1 \in \mathcal{C}$. Consequently, $C_1 \cup C'_2 \in \text{Ext}(C'_2)$ holds, concluding the proof. \square

4.3 Acyclic splits and CCM

We apply Theorem 5 to the problem CCM. Let \mathcal{C} be a closure system over V and Σ be an implicational base for \mathcal{C} . We assume that Σ has an acyclic split (V_1, V_2) . According to Theorem




Figure 18: Computing extensions of a closed set using extensions of meet-irreducible elements of \mathcal{C}_2 .

5, computing \mathcal{M} from \mathcal{M}_1 and \mathcal{M}_2 requires finding maximal extensions of every meet-irreducible element $M_2 \in \mathcal{M}_2$.

FIND MAXIMAL EXTENSIONS (MAXEXT)

Input: A triple $\Sigma[V_1], \Sigma[V_2], \Sigma[V_1, V_2]$ given by an acyclic split of an implicational base Σ , meet-irreducible elements $\mathcal{M}_1, \mathcal{M}_2$, and a closed set C_2 of $\Sigma[V_2]$.

Output: The maximal extensions of C_2 in \mathcal{C} , i.e. $\max_{\subseteq}(\text{Ext}(C_2))$.

This problem relates to the *dualization in closure systems*. Let \mathcal{C} be a closure system over V and $\mathcal{B}^-, \mathcal{B}^+$ two antichains of \mathcal{C} . We say that \mathcal{B}^- and \mathcal{B}^+ are *dual* in \mathcal{C} if $\downarrow \mathcal{B}^+ \cup \uparrow \mathcal{B}^- = \mathcal{C}$ and $\downarrow \mathcal{B}^+ \cap \uparrow \mathcal{B}^- = \emptyset$. The antichain \mathcal{B}^+ is referred as the *positive border*, while \mathcal{B}^- is the *negative border*. Observe that $\mathcal{B}^+ = \max_{\subseteq}(\{C \in \mathcal{C} \mid C \not\subseteq \uparrow \mathcal{B}^-\})$ and similarly $\mathcal{B}^- = \min_{\subseteq}(\{C \in \mathcal{C} \mid C \not\subseteq \downarrow \mathcal{B}^+\})$ so that \mathcal{B}^- is the unique negative border associated to \mathcal{B}^+ , and vice-versa for \mathcal{B}^+ .

We connect maximal extensions of a closed set with dualization. Consider a closure system \mathcal{C} with acyclic split (V_1, V_2) . Let $C_2 \in \mathcal{C}_2$. Since $\text{Ext}(C_2): V_1$ is an ideal of \mathcal{C}_1 , the antichain $\max_{\subseteq}(\text{Ext}(C_2): V_1)$, we call it \mathcal{B}^+ , has a dual antichain \mathcal{B}^- in \mathcal{C}_1 . We have $\mathcal{B}^- = \min_{\subseteq}(\mathcal{C}_1 \setminus \text{Ext}(C_2): V_1)$. In words, \mathcal{B}^- is the family of minimal closed sets of \mathcal{C}_1 that are not participating in extensions of C_2 .

Proposition 8. *Let $C_2 \in \mathcal{C}_2$, and $C_1 \in \mathcal{C}_1$. Then, $C_1 \in \mathcal{B}^-$ if and only if $C_1 \in \min_{\subseteq}\{\phi_1(A) \mid A \rightarrow b \in \Sigma[V_1, V_2], b \notin C_2\}$.*

Proof. We show the if part. We denote by ϕ_1 the closure operator associated to $\Sigma[V_1]$. Let $C_1 \in \min_{\subseteq}\{\phi_1(A) \mid A \rightarrow b \in \Sigma[V_1, V_2], b \notin C_2\}$. We show that for any closed set $C'_1 \subseteq C_1$ in \mathcal{C}_1 , C'_1 contributes to an extension of C_2 . It is sufficient to show this property to the case where $C'_1 \prec C_1$ as $\text{Ext}(C_2): V_1$ is an ideal of \mathcal{C}_1 by Proposition 6. Hence, consider a closed set C'_1 in \mathcal{C}_1 such that $C'_1 \prec C_1$. Note that such C'_1 exists since $\emptyset \in \mathcal{C}_1$ and no implication $A \rightarrow b$ in Σ has $A = \emptyset$ so that $\emptyset \subset \phi_1(A)$ for any implication $A \rightarrow b$ of $\Sigma[V_1, V_2]$ such that $b \notin C_2$. Then, by construction of C'_1 , for any $A \rightarrow b$ in $\Sigma[V_1, V_2]$ such that $b \notin C_2$, we have $\phi_1(A) \not\subseteq C'_1$. As

ϕ_1 is a closure operator, it is monotone and $\phi_1(A) \not\subseteq \phi_1(C'_1) = C'_1$ entails $A \not\subseteq C'_1$ for any such implication $A \rightarrow b$. Therefore $C'_1 \in \text{Ext}(C_2): V_1$ and $C_1 \in \mathcal{B}^-$.

We prove the only if part using contrapositive. Assume $C_1 \notin \min_{\subseteq} \{\phi_1(A) \mid A \rightarrow b \in \Sigma[V_1, V_2], v \notin C_2\}$. We have two cases. First, for any implication $A \rightarrow b$ in $\Sigma[V_1, V_2]$ such that $b \notin C_2$, $\phi_1(A) \not\subseteq C_1$. Since ϕ_1 is monotone and C_1 is closed in \mathcal{C}_1 , we have $A \not\subseteq C_1$ and $C_1 \in \text{Ext}(C_2): V_1$ by Lemma 1. Hence $C_1 \notin \mathcal{B}^-(C_2)$. In the second case, there is an implication $A \rightarrow b$ with $b \notin C_2$ in $\Sigma[V_1, V_2]$ such that $\phi_1(A) \subseteq C_1$ which implies $C_1 \notin \text{Ext}(C_2): V_1$. If $\phi_1(A) \subset C_1$, then clearly $C_1 \notin \mathcal{B}^-$ as $\phi_1(A) \in \mathcal{C}_1$ and $\phi_1(A) \notin \text{Ext}(C_2): V_1$. Hence, assume that $C = \phi_1(A)$. Since $C_1 \notin \min_{\subseteq} \{\phi_1(A) \mid A \rightarrow b \in \Sigma[V_1, V_2], b \notin C_2\}$ by hypothesis, there exists another implication $A' \rightarrow b' \in \Sigma[V_1, V_2]$ such that $b' \notin C_2$ and $\phi_1(A') \subset C_1$. Hence $\phi_1(A') \notin \text{Ext}(C_2): V_1$ and $C_1 \notin \mathcal{B}^-$ as it is not an inclusion-wise minimum closed set which does not belong to $\text{Ext}(C_2): V_1$. \square

We can build \mathcal{B}^- in polynomial time from Σ using Proposition 5 and $\Sigma[V_1, V_2]$: we compute $\phi_1(A)$ for every implication $A \rightarrow b$ in $\Sigma[V_1, V_2]$ and we keep the closed sets (in \mathcal{C}_1) that are inclusion-wise minimal. Therefore, the problem MAXEXT relates to the following generation version of dualization.

LOWER DUALIZATION IN CLOSURE SYSTEMS (LDUAL(α))

Input: A representation α for a closure system \mathcal{C} over V , an antichain \mathcal{B}^- of \mathcal{C}
Output: The antichain \mathcal{B}^+ dual to \mathcal{B}^- .

When α is an implicational base Σ or the set of meet-irreducible elements \mathcal{M} , the problem LDUAL(α) is impossible to solve in output-polynomial time unless $\mathbf{P} = \mathbf{NP}$ [6, 12]. However, in MAXEXT we have access to both Σ_1 and \mathcal{M}_1 so that the version of LDUAL we have to consider is the one where α is both an implicational base and a set of meet-irreducible elements, that is LDUAL(Σ, \mathcal{M}). This version of LDUAL is open, even if not harder than SID [6]. When $\Sigma = \emptyset$, *i.e.* when the closure system is Boolean, the problem reduces to hypergraph dualization.

Now, we describe an algorithm for solving CCM in the presence of acyclic splits. First, we have $|\mathcal{M}| \geq |\mathcal{M}_1| + |\mathcal{M}_2|$ due to Theorem 5. Furthermore, each $M \in \mathcal{M}$ arise from a unique element of $M' \in \mathcal{M}_1 \cup \mathcal{M}_2$, and each $M' \in \mathcal{M}_1 \cup \mathcal{M}_2$ is used to construct at least one new meet-irreducible element $M \in \mathcal{M}$. Therefore, the algorithm will output every meet-irreducible element only once. Furthermore, the space needed to store intermediate solutions is bounded by the size of the output \mathcal{M} which prevents an exponential blow up during the execution.

The algorithm proceeds as follows. If Σ has no acyclic split, we use routines such as in [36, 7] to compute \mathcal{M} . When V is a singleton, the unique meet-irreducible to find is \emptyset and hence no call to other algorithm is required. Otherwise, we find an acyclic split (V_1, V_2) of Σ and we recursively call the algorithm on $\Sigma[V_1]$ and $\Sigma[V_2]$. Then, we compute \mathcal{M} using $\Sigma, \mathcal{M}_1, \mathcal{M}_2$ and by solving MAXEXT. Observe that it takes polynomial time in the size of Σ and V to compute an acyclic split, if it exists:

- compute the premise-connected components of Σ ;
- construct a directed graph on these components, with an arc from a component C_1 to C_2 if there is an implication $A \rightarrow b$ in Σ such that $A \subseteq C_1$ and $b \in C_2$;
- then, an acyclic split exists if and only if there are at least two strongly connected components, and each non-trivial bipartition of the strongly connected components will represent an acyclic split.

Thus, the algorithm BuildTree can be adapted to find a decomposition with acyclic splits or return FAIL if not possible in polynomial time.




Figure 19: The Σ -tree of Σ .

Example 14 (Running example). First, we compute a decomposition of Σ in terms of acyclic splits. We obtain the Σ -tree illustrated in Figure 19.

Then, we apply Theorem 5 bottom-up to construct the the set \mathcal{M} of meet-irreducible elements of \mathcal{C} . This part is shown in Figure 20. For readability, we highlighted at each step which closed sets are part of \mathcal{C}_2 and also the two types of meet-irreducible elements of Theorem 5.

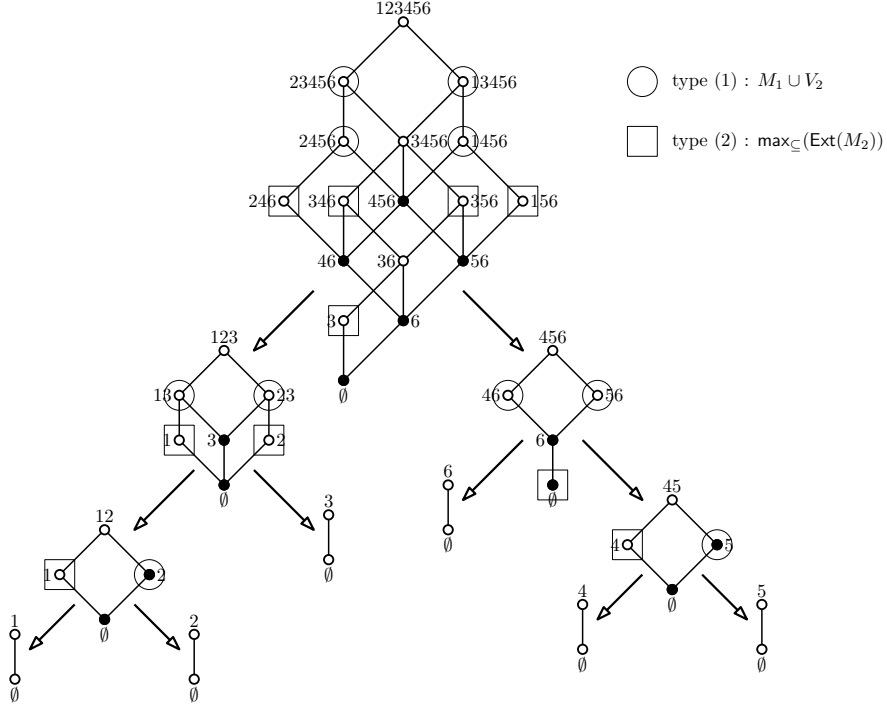


Figure 20: Recursive computation of \mathcal{M} using a decomposition by acyclic splits.

To conclude, we derive a class of implicational bases where our strategy can be applied to obtain the meet-irreducible elements in output quasi-polynomial time.

Theorem 6. *Let Σ be an implicational base over V . Assume there exists a full partition V_1, \dots, V_k of V such that for every implication $A \rightarrow b \in \Sigma$, $A \subseteq V_i$ and $b \in V_j$ for some $1 \leq i < j \leq k$. Then CCM can be solved in output-quasipolynomial time.*

Proof. Observe that Σ is acyclic in this case. Then, Σ can be hierarchically decomposed by $k - 1$ acyclic splits such that the implicational base on the left of the i -th split is $\Sigma[V_i] = \emptyset$ and the right-one $\Sigma[\bigcup_{j>i} V_j]$. Then, MAXEXT reduces to hypergraph dualization, and we can compute

\mathcal{M} from Σ in output-quasipolynomial time using the algorithm of Fredman and Khachiyan [21]. \square

The class of closure systems associated to these implicational bases generalizes both distributive closure systems and ranked convex geometries [13] since an implicational base is ranked when it further satisfies the condition that $A \subseteq V_i$ implies $b \in V_{i+1}$.

5 Discussions and open problems

We conclude the paper with some discussions and open questions for future work. Splits and more notably acyclic splits are decomposition methods based on the syntax of implications. However, two equivalent implicational bases may not share the same (acyclic) splits. In fact, it is even possible to find two equivalent implicational bases where one has an acyclic split, and not the other. This is demonstrated by the following example.

Example 15. Let $V = \{1, 2, 3, 4\}$ and $\Sigma = \{1 \rightarrow 4, 124 \rightarrow 3, 3 \rightarrow 4\}$. The unique possible split is $(124, 3)$ which is not acyclic. Observe that Σ is the Duquenne-Guigues base of the closure system. However, the implicational base $\Sigma' = \{1 \rightarrow 4, 12 \rightarrow 3, 3 \rightarrow 4\}$, which is clearly equivalent to Σ has an acyclic split being $(12, 34)$.

Note that the Duquenne-Guigues base is not of interest for finding acyclic splits as it can hide possible acyclic splits, as suggested by the previous example. In fact, the example suggests considering only minimum implicational bases whose left-sides are as small as possible. However, several such bases may exist and finding the right-one might be an expensive task, whence the following question.

Question 1. *Is it possible to decide whether a closure system has an acyclic split in polynomial time from an implicational base?*

A similar question holds for the case of meet-irreducible elements:

Question 2. *Is it possible to recognize an acyclic split in polynomial time from a set of meet-irreducible elements?*

In Corollary 3, we give a first step towards a characterization of acyclic splits from meet-irreducible elements. The statement in Corollary 3 does consider the representation of closed sets by meet-irreducible elements. Nonetheless, this characterization needs to be checked on every closed set of \mathcal{C}_2 . In order to recognize an acyclic split from a set of meet-irreducible elements only, an idea would be to replace the statement by this one:

for every $M_2, M'_2 \in \mathcal{M}_2$ such that $M_2 \subseteq M'_2$, $\text{Ext}(M_2) : V_1 \subseteq \text{Ext}(M'_2) : V_1$.

Unfortunately, this latter condition is not sufficient, as demonstrated by the next example.

Example 16. Let $V_1 = \{4, 5\}$, $V_2 = \{1, 2, 3\}$ and consider the closure systems \mathcal{C}_1 and \mathcal{C}_2 given in Figure 21.



Figure 21: The closure systems \mathcal{C}_1 and \mathcal{C}_2 .




Figure 22: The closure system \mathcal{C} , failing Corollary 3.

An implicational base for \mathcal{C}_1 is $\Sigma_1 = \emptyset$ and $\Sigma_2 = \{12 \rightarrow 3, 13 \rightarrow 2, 23 \rightarrow 1\}$ is an implicational base for \mathcal{C}_2 . We have $\mathcal{M}_1 = \{4, 5\}$ and $\mathcal{M}_2 = \{1, 2, 3\}$. Now let $V = V_1 \cup V_2$ and consider the closure system \mathcal{C} of Figure 22 and the pair (V_1, V_2) .

We have $\mathcal{M} = \{1234, 1235\} \cup \{14, 24, 25, 34, 35\}$. As \mathcal{M}_2 is an antichain, the condition given above is satisfied. However, Corollary 3 fails because $\max_{\subseteq}(\text{Ext}(\emptyset): V_1) \not\subseteq \text{Ext}(1): V_1$. Hence, (V_1, V_2) is not an acyclic split for \mathcal{C} .

When (V_1, V_2) is an acyclic split of \mathcal{C} and V_2 is a singleton element, the construction of \mathcal{C} can be interpreted as the duplication of an ideal of \mathcal{C}_1 . This puts the light on a possible link between (acyclic) splits and lower-bounded lattices [22, 3]. In particular, we know from [3] that the non-left-unit part of the D -base of a lower bounded lattice is acyclic. As left-unit implications play no role in the existence of splits, there should exist a H-decomposition of the D -base by “almost acyclic” splits.

Example 17. Let $V = \{1, 2, 3\}$ and $\Sigma = \{12 \rightarrow 3, 3 \rightarrow 1\}$. The associated closure system is (lower) bounded. Its D -base is precisely Σ . It has no acyclic split when we consider $3 \rightarrow 1$, but it has a split $(12, 3)$ which becomes acyclic once $3 \rightarrow 1$ is removed.

Thus, we are naturally lead to the next question.

Question 3. *Can implicational bases of lower-bounded closure systems be characterized by the existence of a particular Σ -tree?*

Answering this question would allow extending Theorem 5 to take into account unitary implications creating cycles.

Acknowledgments The second author is funded by the CNRS, France, ProFan project. This research is also supported by the French government IDEXISITE initiative 16-IDEX-0001 (CAP 20-25).

References

- [1] ADARICHEVA, K. V., AND NATION, J. B. On implicational bases of closure systems with unique critical sets. *Discrete Applied Mathematics* 162 (2014), 51–69.
- [2] ADARICHEVA, K. V., AND NATION, J. B. Discovery of the D-basis in binary tables based on hypergraph dualization. *Theoretical Computer Science* 658 (2017), 307–315.
- [3] ADARICHEVA, K. V., NATION, J. B., AND RAND, R. Ordered direct implicational basis of a finite closure system. *Discrete Applied Mathematics* 161, 6 (2013), 707–723.
- [4] AUSIELLO, G., D’ATRI, A., AND SACCA, D. Minimal representation of directed hypergraphs. *SIAM Journal on Computing* 15, 2 (1986), 418–431.

- [5] BABIN, M. A., AND KUZNETSOV, S. O. Computing premises of a minimal cover of functional dependencies is intractable. *Discrete Applied Mathematics* 161, 6 (2013), 742–749.
- [6] BABIN, M. A., AND KUZNETSOV, S. O. Dualization in lattices given by ordered sets of irreducibles. *Theoretical Computer Science* 658 (2017), 316–326.
- [7] BEAUDOU, L., MARY, A., AND NOURINE, L. Algorithms for k-meet-semidistributive lattices. *Theoretical Computer Science* 658 (2017), 391–398.
- [8] BERTET, K., DEMKO, C., VIAUD, J.-F., AND GUÉRIN, C. Lattices, closures systems and implication bases: A survey of structural aspects and algorithms. *Theoretical Computer Science* 743 (2018), 93–109.
- [9] BERTET, K., AND MONJARDET, B. The multiple facets of the canonical direct unit implicational basis. *Theoretical Computer Science* 411, 22-24 (2010), 2155–2166.
- [10] DASGUPTA, S. A cost function for similarity-based hierarchical clustering. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (2016), pp. 118–127.
- [11] DAVEY, B. A., AND PRIESTLEY, H. A. *Introduction to Lattices and Order*. Cambridge university press, 2002.
- [12] DEFRAIN, O., NOURINE, L., AND UNO, T. On the dualization in distributive lattices and related problems. *Discrete Applied Mathematics* 300 (2021), 85–96.
- [13] DEFRAIN, O., NOURINE, L., AND VILMIN, S. Translating between the representations of a ranked convex geometry. *Discrete Mathematics* 344, 7 (2021), 112399.
- [14] DEMETROVICS, J., LIBKIN, L., AND MUCHNIK, I. B. Functional dependencies in relational databases: A lattice point of view. *Discrete Applied Mathematics* 40, 2 (1992), 155–185.
- [15] DISTEL, F., AND SERTKAYA, B. On the complexity of enumerating pseudo-intents. *Discrete Applied Mathematics* 159, 6 (2011), 450–466.
- [16] DOIGNON, J.-P., AND FALMAGNE, J.-C. *Knowledge Spaces*. Springer Science & Business Media, 2012.
- [17] DUNG, P. M. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence* 77, 2 (1995), 321–357.
- [18] EDELMAN, P. H., AND JAMISON, R. E. The theory of convex geometries. *Geometriae dedicata* 19, 3 (1985), 247–270.
- [19] EITER, T., AND GOTTLOB, G. Identifying the minimal transversals of a hypergraph and related problems. *SIAM Journal on Computing* 24, 6 (1995), 1278–1304.
- [20] ELAROSSI, M., NOURINE, L., AND RADJEF, M. Lattice point of view for argumentation framework.
- [21] FREDMAN, M. L., AND KHACHIVAN, L. On the complexity of dualization of monotone disjunctive normal forms. *Journal of Algorithms* 21, 3 (1996), 618–628.
- [22] FREESE, R., JEŽEK, J., AND NATION, J. B. *Free Lattices*, vol. 42. American Mathematical Soc., 1995.
- [23] GANTER, B., AND WILLE, R. *Formal Concept Analysis: Mathematical Foundations*. Springer Science & Business Media, 2012.

- [24] GRÄTZER, G. A. *Lattice Theory: Foundation*. Springer Science & Business Media, 2011.
- [25] GUIGUES, J.-L., AND DUQUENNE, V. Familles minimales d’implications informatives résultant d’un tableau de données binaires. *Mathématiques et Sciences humaines* 95 (1986), 5–18.
- [26] HABIB, M., AND NOURINE, L. Representation of lattices via set-colored posets. *Discrete Applied Mathematics* 249 (2018), 64–73.
- [27] HAMMER, P. L., AND KOGAN, A. Quasi-acyclic propositional Horn knowledge bases: Optimal compression. *IEEE Transactions on knowledge and data engineering* 7, 5 (1995), 751–762.
- [28] JOHNSON, D. S., YANNAKAKIS, M., AND PAPADIMITRIOU, C. H. On generating all maximal independent sets. *Information Processing Letters* 27, 3 (1988), 119–123.
- [29] KAUTZ, H. A., KEARNS, M. J., AND SELMAN, B. Reasoning with characteristic models. In *AAAI* (1993), vol. 93, Citeseer, pp. 34–39.
- [30] KAVVADIAS, D. J., SIDERI, M., AND STAVROPOULOS, E. C. Generating all maximal models of a Boolean expression. *Information Processing Letters* 74, 3-4 (2000), 157–162.
- [31] KHARDON, R. Translating between Horn representations and their characteristic models. *Journal of Artificial Intelligence Research* 3 (1995), 349–372.
- [32] KUZNETSOV, S. O. On the intractability of computing the duquenne-guigues base. *Journal of Universal Computer Science* 10, 8 (2004), 927–933.
- [33] LAWLER, E. L., LENSTRA, J. K., AND RINNOOY KAN, A. Generating all maximal independent sets: NP-hardness and polynomial-time algorithms. *SIAM Journal on Computing* 9, 3 (1980), 558–565.
- [34] LIBKIN, L. Direct product decompositions of lattices, closures and relation schemes. *Discrete Mathematics* 112, 1-3 (1993), 119–138.
- [35] LUCCHESI, C. L., AND OSBORN, S. L. Candidate keys for relations. *Journal of Computer and System Sciences* 17, 2 (1978), 270–279.
- [36] MANNILA, H., AND RÄIHÄ, K.-J. *The Design of Relational Databases*. Addison-Wesley Longman Publishing Co., Inc., 1992.
- [37] MANNILA, H., AND RÄIHÄ, K.-J. Algorithms for inferring functional dependencies from relations. *Data & Knowledge Engineering* 12, 1 (1994), 83–99.
- [38] MARKOWSKY, G. The factorization and representation of lattices. *Transactions of the American Mathematical Society* 203 (1975), 185–200.
- [39] NATION, J. B. Unbounded semidistributive lattices. *Algebra and Logic* 39, 1 (2000), 50–53.
- [40] OBIEDKOV, S. A., AND DUQUENNE, V. Attribute-incremental construction of the canonical implication basis. *Annals of Mathematics and Artificial Intelligence* 49, 1 (2007), 77–99.
- [41] SHOCK, R. C. Computing the minimum cover of functional dependencies. *Information Processing Letters* 22, 3 (1986), 157–159.
- [42] TARJAN, R. E., AND VAN LEEUWEN, J. Worst-case analysis of set union algorithms. *Journal of the ACM (JACM)* 31, 2 (1984), 245–281.

- [43] WILD, M. A theory of finite closure spaces based on implications. *Advances in Mathematics* 108, 1 (1994), 118–139.
- [44] WILD, M. Computations with finite closure systems and implications. In *International Computing and Combinatorics Conference* (1995), Springer, pp. 111–120.
- [45] WILD, M. Optimal implicational bases for finite modular lattices. *Quaestiones Mathematicae* 23, 2 (2000), 153–161.
- [46] WILD, M. The joy of implications, aka pure Horn formulas: Mainly a survey. *Theoretical Computer Science* 658 (2017), 264–292.
- [47] ZANUTTINI, B. Sur des propriétés structurelles des formules de horn. In *9es Journées d’Intelligence Artificielle Fondamentale (IAF 2015)* (2015).