DIMACS Technical Report 2002-27

Extending the Balas-Yu Bounds on the Number of Maximal Independent Sets in Graphs to Hypergraphs and Lattices [1]

by

Endre Boros,[2] Khaled Elbassioni[3] Vladimir Gurvich[2]
Leonid Khachiyan[3]

[1]This research was supported by the National Science Foundation (Grant IIS-0118635), and by the Office of Naval Research (Grant N00014-92-J-1375). The second and third authors are also grateful for the partial support by DIMACS, the National Science Foundation's Center for Discrete Mathematics and Theoretical Computer Science.
[2]RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway NJ 08854-8003; ({boros, gurvich}@rutcor.rutgers.edu).
[3]Department of Computer Science, Rutgers University, 110 Frelinghuysen Road, Piscataway, New Jersey 08854-8004; ({elbassio@paul, leonid@cs}.rutgers.edu.

ABSTRACT

A result of Balas and Yu (1989) states that the number of maximal independent sets of a graph G is at most $\delta^p + 1$, where δ is the number of pairs of vertices in G at distance 2, and p is the cardinality of a maximum induced matching in G. In this paper, we give an analogue of this result for hypergraphs and, more generally, for subsets of vectors $\mathcal{B}$ in the product of n lattices $\mathcal{L} = \mathcal{L}_1 \times \cdots \times \mathcal{L}_n$, where the notion of an induced matching in G is replaced by a certain binary tree each internal node of which is mapped into $\mathcal{B}$. We show that our bounds may be nearly sharp for arbitrarily large hypergraphs and lattices. As an application, we prove that the number of maximal infeasible vectors $x \in \mathcal{L} = \mathcal{L}_1 \times \cdots \times \mathcal{L}_n$ for a system of polymatroid inequalities $f_1(x) \geq t_1, \ldots, f_r(x) \geq t_r$ does not exceed $\max\{Q, \beta^{\log t/c(2Q,\beta)}\}$, where β is the number of minimal feasible vectors for the system, $Q = |\mathcal{L}_1| + \ldots + |\mathcal{L}_n|$, $t = \max\{t_1, \ldots, t_r\}$, and $c(\rho, \beta)$ is the unique positive root of the equation $2^c(\rho^{c/\log\beta} - 1) = 1$. This bound is nearly sharp for the Boolean case $\mathcal{L} = \{0,1\}^n$, and it allows for the efficient generation of all minimal feasible sets to a given system of polymatroid inequalities with quasi-polynomially bounded right-hand sides $t_1, \ldots, t_m$.

1 Introduction

A result of Balas and Yu (1989) shows that a graph with many maximal independent sets must have a large induced matching. More precisely, the following bounds hold.

Theorem 1 ([3]) *Let $G = (V, E)$ be a connected graph, $\delta = \delta(G)$ be the number of pairs of vertices in G at distance 2, and $p = p(G)$ be the cardinality of a maximum induced matching in G. Then*

$$2^p \leq |\mathcal{I}(G)| \leq \delta^p + 1, \tag{1}$$

where $\mathcal{I}(G)$ is the family of all maximal independent sets in G.

Slightly weaker results were obtained independently by Alekseev (1991) and Prisner (1992). Let us note that the lower bound in (1) is obvious, and that $\delta \leq (n-1)(n-2)/2$, where $n = |V|$.

In order to generalize inequalities (1) to hypergraphs and lattices, it will be convenient to restate the notion of an induced matching in a graph in the following way. Let $G = (V, E)$ be a graph and let $\mathcal{M} = pK_2$ be an induced matching in G with edges $e_i = \{u_i, v_i\}$, $i = 1, ..., p$. Consider a rooted uniform binary tree $\mathbf{T}$ of depth p, and the mapping ϕ which assigns the first edge e_1 to the root $\mathbf{r} = \mathbf{r}(\mathbf{T})$ of $\mathbf{T}$, the second edge e_2 to the two children of $\mathbf{r}$, the third edge e_3 to the four grandchildren of $\mathbf{r}$, up to the last edge e_p assigned to the 2^{p-1} internal nodes preceding the leaves of $\mathbf{T}$. Let ℓ be one of the 2^p leaves of $\mathbf{T}$, then there is a unique path $\mathcal{P} = \{w_0, w_1, \ldots, w_{p-1}, w_p\}$ from the root $w_0 = \mathbf{r}(\mathbf{T})$ to $w_p = \ell$, where each node w_{i+1} is a child of w_i, for $i = 0, \ldots, p-1$. Let us denote by $LC(w)$ and $RC(w)$ the left and right children of an internal node w in $\mathbf{T}$, and define for each path $\mathcal{P} = \{w_0, w_1, \ldots, w_{p-1}, w_p\}$ the following vertex set:

$$S(\mathcal{P}) = \left(\bigcup_{w : LC(w) \in \mathcal{P}} u(w) \right) \cup \left(\bigcup_{w : RC(w) \in \mathcal{P}} v(w) \right).$$

where $u(w)$ and $v(w)$ are the endpoints of the edge $e = \phi(w)$ assigned to w. Then, since $\mathcal{M}$ is an *induced* matching, $S(\mathcal{P})$ is an independent vertex set in G and furthermore, if we arbitrarily extend, for each leaf ℓ, the corresponding set $S(\mathcal{P})$ to a *maximal* independent set in G, we will obtain 2^p distinct maximal independent sets of G. This gives the lower bound on $|\mathcal{I}(G)|$ in Theorem 1; the upper bound of the theorem states that the size of $\mathcal{I}(G)$ does not exceed $\delta^{\log |\mathbf{L}(\mathbf{T})|} + 1$, where $\mathbf{L}(\mathbf{T})$ is the set of leaves of $\mathbf{T}$.

We can generalize the above construction from graphs to hypergraphs as follows. Let V be a finite set of cardinality $|V| = n$, and let $\mathcal{H} \subseteq 2^V$ be a hypergraph on V. Let $\mathbf{T}$ be a binary (but not necessarily uniform) tree $\mathbf{T}$ each internal node w of which has two children $LC(w)$ and $RC(w)$. Denote by $\mathbf{N}(\mathbf{T})$ the set of internal nodes of $\mathbf{T}$ and consider a mapping $\phi : \mathbf{N}(\mathbf{T}) \mapsto \mathcal{H} \times 2^V \times 2^V$, which assigns a hyperedge $H(w) \in \mathcal{H}$ and two vertex sets $L(w), R(w) \in 2^V$ to each internal node $w \in \mathbf{N}(\mathbf{T})$. We will call the mapping ϕ *proper* if

(i) $H(w) \subseteq L(w) \cup R(w)$ for each internal node $w \in \mathbf{N}(\mathbf{T})$ and (ii) for each path $\mathcal{P}$ from the root to a leaf of $\mathbf{T}$, the vertex set

$$S(\mathcal{P}) = \left(\bigcup_{w:LC(w)\in\mathcal{P}} L(w) \right) \cup \left(\bigcup_{w:RC(w)\in\mathcal{P}} R(w) \right).$$

is independent, i.e. contains no hyperedge of $\mathcal{H}$. In particular, our mapping above of the uniform binary tree of depth p into the edges of the induced matching pK_2 is proper if we define $L(w) = \{u(w)\}$ and $R(w) = \{v(w)\}$ to be the endpoints of the edge assigned to w.

Let $\phi : \mathbf{N}(\mathbf{T}) \to \mathcal{H} \times 2^V \times 2^v$ be a proper mapping. If ϕ assigns a hyperedge $H(w) \in \mathcal{H}$ to some node of $\mathbf{T}$, and $H(w)$ contains another hyperedge H' of $\mathcal{H}$, then we can replace $H(w)$ by H' without violating properties (i) and (ii). We can thus replace the given hypergraph $\mathcal{H}$ by the hypergraph $\min(\mathcal{H})$ consisting of all minimal hyperedges of $\mathcal{H}$, and hence assume without loss of generality that $\mathcal{H}$ is Sperner, i.e., no hyperedge of $\mathcal{H}$ contains another hyperedge. In addition, after the transformation $L(w) \leftarrow H(w) \cap L(w),\ R(w) \leftarrow H(w) \setminus (H(w) \cap L(w))$ we again obtain a proper mapping. Hence we can also assume without loss of generality that $L(w)$ and $R(w)$ not only cover, but partition the hyperedge $H(w)$. Note that by (ii), both sets $L(w)$ and $R(w)$ must be non-empty, i.e., $\{L(w), R(w)\}$ must be a proper partition of $H(w)$. As before, it is easily seen that extending each set $S(\mathcal{P})$ to a maximal independent set, we obtain $|\mathbf{L}(\mathbf{T})|$ distinct maximal independent sets.

Given a hypergraph $\mathcal{H} \subseteq 2^V$, let $\gamma = \gamma(\mathcal{H})$ be the number of leaves in the largest binary tree for which a proper mapping for $\mathcal{H}$ exists. Denote by $\mathcal{I}(\mathcal{H})$ the family of maximal independent sets for $\mathcal{H}$, and let $\beta = \beta(\mathcal{H}) = |\mathcal{H}|$ and $\alpha = \alpha(\mathcal{H}) = |\mathcal{I}(\mathcal{H})|$. Theorem 2 below provides an analogue of the bounds in (1) to hypergraphs $\mathcal{H}$, in terms of the value of $\gamma(\mathcal{H})$:

Theorem 2 ([4]) *For any hypergraph $\mathcal{H} \subseteq 2^V$ of size $\beta \geq 2$, it holds that*

$$\gamma \leq \alpha \leq \beta^{\log\gamma/c(n,\beta)}, \tag{2}$$

where $n = |V|$, and $c(n,\beta)$ is the unique positive root of the equation

$$2^c(n^{c/\log\beta} - 1) = 1. \tag{3}$$

In addition, $\alpha \leq n$ holds if $\beta = 1$.

As mentioned earlier, the lower bound $\gamma \leq \alpha$ is obvious. This bound is sharp, e.g., for $G = pK_2$. As for the upper bound of (2), let us first remark that by (3), $1 = n^{-c(n,\beta)/\log\beta} + (n\beta)^{-c(n,\beta)/\log\beta} \geq 2(n\beta)^{-c(n,\beta)/\log\beta}$, and hence $\beta^{1/c(n,\beta)} \leq n\beta$. Consequently, for $\beta \geq 2$ we can replace the upper bound of (2) by the simpler but weaker inequality

$$\alpha \leq (n\beta)^{\log\gamma}. \tag{4}$$

In fact, (4) holds even in case of $\beta = 1$, because if the hypergraph $\mathcal{H}$ consists only of a single hyperedge $X \subseteq V$, then $\alpha \leq |X| \leq n$ follows immediately by the relation $\alpha = |\mathcal{I}(\mathcal{H})|$. On the other hand, for large β the upper bound of Theorem 2 becomes increasingly stronger than (4). For instance, $c(n,n) = \log(1+\sqrt{5}) - 1 > .694$, $c(n,n^2) > 1.102$, and $c(n,n^\sigma) \sim \log\sigma$ for large σ.

As shown in [4], there are arbitrarily large graphs for which $\alpha \geq \beta^{.551 \log\gamma/c(n,\beta)}$. Before proceeding further, we show here that for hypergraphs of unbounded edge size, the upper bound of Theorem 2 is, in fact, nearly sharp.

Example 1.1 Let us start with the following simple observation. Suppose we are given a hypergraph $\mathcal{H} \subseteq 2^V$ and a proper mapping $\phi : \mathbf{N}(\mathbf{T}) \mapsto \mathcal{H} \times 2^V \times 2^V$. Let $\mathbf{r}$ be the root of $\mathbf{T}$, and let $L(\mathbf{r}) \cup R(\mathbf{r}) = H(\mathbf{r})$ be the partition of the hyperedge $H(\mathbf{r})$ assigned to $\mathbf{r}$ by ϕ. Consider the hypergraph $\mathcal{H}_L \subseteq 2^{V\setminus L(\mathbf{r})}$ (respectively $\mathcal{H}_R \subseteq 2^{V\setminus R(\mathbf{r})}$) obtained by deleting all vertices in $L(\mathbf{r})$ (respectively, $R(\mathbf{r})$) from V, replacing each of the hyperedges $H \in \mathcal{H}$ by the set difference $H \setminus L(\mathbf{r})$ (respectively, $H \setminus R(\mathbf{r})$), and then leaving only the minimal of the resulting hyperedges. Then the proper mapping ϕ naturally splits into two separate proper mappings ϕ_L and ϕ_R defined for $\mathcal{H}_L$ and $\mathcal{H}_R$ on the subtrees rooted at the left and right children of $\mathbf{r}$, respectively. In particular, if $\mathcal{H} = \binom{[m]}{d}$ is the complete d-uniform hypergraph on $m \geq d$ vertices, then $\mathcal{H}_L$ and $\mathcal{H}_R$ are also complete $|R(\mathbf{r})|$-uniform and $|L(\mathbf{r})|$-uniform hypergraphs. This shows that $\gamma(\mathcal{H}) = d$. Similarly, if $\mathcal{H}$ consists of k disjoint d_i-uniform complete hypergraphs $\mathcal{H}_i$, $i = 1, \dots, k$, then the same argument gives $\gamma(\mathcal{H}) = d_1 d_2 \cdots d_k$.

Now consider the hypergraph $\mathcal{H} = k \times \binom{[m]}{d}$ consisting of k pairwise disjoint copies of the complete d-uniform hypergraph on m vertices. For this hypergraph, we have $n = km$, $\gamma(\mathcal{H}) = d^k$, $\beta = k\binom{m}{d}$, and $\alpha = \binom{m}{d-1}^k$. Thus, letting $d = k$ and $m = 2^k$, we obtain as $k \to \infty$:

$$\alpha = 2^{(1-o(1))k^3}, \quad \beta = 2^{(1-o(1))k^2}, \quad c(n,\beta) = (1-o(1))\log k,$$

and consequently $\alpha = \beta^{(1-o(1))\log\gamma/c(n,\beta)}$. □

Each hypergraph on n vertices can be viewed as a subset of the Boolean cube $\{0,1\}^n$. In this paper, we further generalize Theorem 2 from hypergraphs to subsets $\mathcal{B}$ of vectors in the product $\mathcal{L} = \mathcal{L}_1 \times \cdots \times \mathcal{L}_n$ of n arbitrary finite lattices $\mathcal{L}_1, \dots, \mathcal{L}_n$. Let us denote by $\preceq$ the precedence relation in $\mathcal{L}$ and also in $\mathcal{L}_1, \dots, \mathcal{L}_n$, i.e. if $p = (p_1, \dots, p_n) \in \mathcal{L}$ and $q = (q_1, \dots, q_n) \in \mathcal{L}$, then $p \preceq q$ in $\mathcal{L}$ if and only if $p_1 \preceq q_1$ in $\mathcal{L}_1$, $\dots$, $p_n \preceq q_n$ in $\mathcal{L}_n$. Let us also use, as customary, $\vee$ and $\wedge$ to denote the join and meet operators over $\mathcal{L}$. Generalizing standard terminology of the theory of hypergraphs, an element $a \in \mathcal{L}$ is said to be *independent* of $\mathcal{B} \subseteq \mathcal{L}$ if $a \not\succeq b$ for all $b \in \mathcal{B}$, and is said to be *maximal* independent of $\mathcal{B}$ if it is maximal with this property. Denote by $\mathcal{I}(\mathcal{B})$ the set of maximal independent elements of $\mathcal{B}$. Clearly, if $\mathcal{L}_i = \{0,1\}$ for all $i \in V \stackrel{\text{def}}{=} \{1, \dots, n\}$, the families $\mathcal{B}$ and $\mathcal{I}(\mathcal{B})$ correspond respectively to a hypergraph and its maximal independent sets. Furthermore, for any lattice $\mathcal{L}$, we can naturally extend the notion of proper mappings to subsets $\mathcal{B}$ of $\mathcal{L}$. As before, let $\mathbf{T}$ be a rooted binary tree in which every internal node $w \in \mathbf{N}(\mathbf{T})$ has two children. Given $\mathcal{B} \subseteq \mathcal{L}$, let us consider now mappings $\phi : \mathbf{N}(\mathbf{T}) \mapsto \mathcal{B} \times \mathcal{L} \times \mathcal{L}$ that assign an element $b = b(w) \in \mathcal{B}$, and two elements $L(w), R(w) \in \mathcal{L}$ such that $L(w) \vee R(w) \succeq b(w)$, to each node $w \in \mathbf{N}(\mathbf{T})$. Given a path $\mathcal{P}$ from the root of $\mathbf{T}$ to a leaf, define the join

$$s(\mathcal{P}) = s_\phi(\mathcal{P}) = \left(\bigvee_{w:LC(w)\in\mathcal{P}} L(w) \right) \vee \left(\bigvee_{w:RC(w)\in\mathcal{P}} R(w) \right), \tag{5}$$

and call the mapping ϕ proper if for every such path $\mathcal{P}$, the element $s(\mathcal{P})$ is independent of $\mathcal{B}$. Finally, let $\gamma = \gamma(\mathcal{B})$ be the number of leaves in the largest binary tree for which a proper mapping for $\mathcal{B}$ exists. As our main result, we shall prove the following.

Theorem 3 *Let $\mathcal{L} = \mathcal{L}_1 \times \cdots \times \mathcal{L}_n$ be the product of n lattices, and let $\mathcal{B} \subseteq \mathcal{L}$ be any subset of $\mathcal{L}$ of size $|\mathcal{B}| \geq 1$. Then*

$$\gamma(\mathcal{B}) \leq |\mathcal{I}(\mathcal{B})| \leq \max\left\{Q, |\mathcal{B}|^{\log \gamma(\mathcal{B})/c(2Q,|\mathcal{B}|)}\right\}, \tag{6}$$

where $Q = \sum_{i=1}^n |\mathcal{L}_i|$.

We next discuss a corollary of Theorem 3 related to polymatroid inequalities on lattices. Let $\mathcal{L}$ be a lattice and let $f : \mathcal{L} \mapsto \mathbb{Z}_+$ be an integer-valued function over the elements of $\mathcal{L}$. The function f is called *polymatroid* if

- f is *monotone*: $x \succeq y$ implies $f(x) \geq f(y)$,
- f is *submodular*: $f(x \vee y) + f(x \wedge y) \leq f(x) + f(y)$ holds for all $x, y \in \mathcal{L}$
- $f(\mathbf{0}) = 0$, where $\mathbf{0} \in \mathcal{L}$ is the minimum element of $\mathcal{L}$.

Given r polymatroid functions $f_1, \ldots, f_r : \mathcal{L} \mapsto \mathbb{Z}_+$ and r positive integer thresholds $t_1, \ldots, t_r \in \mathbb{Z}_+$, consider the system of polymatroid inequalities:

$$f_i(x) \geq t_i, \qquad i = 1, \ldots, r, \tag{7}$$

over the elements $x \in \mathcal{L}$. Let us denote by $\mathcal{F} \subseteq \mathcal{L}$ the set of minimal feasible vectors for (7), then $\mathcal{I}(\mathcal{F})$ is the set of all maximal infeasible vectors for (7). Theorem 3 implies the following bound.

Theorem 4 *Consider a system of r polymatroid inequalities (7) over the product $\mathcal{L} = \mathcal{L}_1 \times \ldots \times \mathcal{L}_n$ of n lattices. Let $\mathcal{X} \subseteq \mathcal{F}$ be an arbitrary non-empty set of minimal feasible solutions for the system. Then*

$$|\mathcal{I}(\mathcal{X}) \cap \{x \in \mathcal{L} \mid x \textit{ is infeasible for (7)}\}| \leq r \cdot \max\left\{Q, |\mathcal{X}|^{\log t/c(2Q,|\mathcal{X}|)}\right\}, \tag{8}$$

where $t = \max\{t_1, \ldots, t_r\}$ and $Q = |\mathcal{L}_1| + \ldots + |\mathcal{L}_n|$. In particular, for $\mathcal{X} = \mathcal{F}$ we obtain

$$|\mathcal{I}(\mathcal{F})| \leq r \cdot \max\{Q, |\mathcal{F}|^{\log t/c(2Q,|\mathcal{F}|)}\}. \tag{9}$$

The following example shows that the bounds of Theorem 4 may be nearly sharp for systems of polymatroid inequalities in binary variables, with arbitrarily large r, n and t.

Example 1.2 Let M be a linear space of dimension D over some field $\mathbf{F}$, and let $S_1, \ldots, S_n \subseteq M$ be n subsets of M. Given a set $X \subseteq V = \{1, \ldots, n\}$, define $f(X)$ to be the dimension of the linear subspace spanned by the sets $S_i,\ i \in X$. Then the function $f(X)$ is polymatroid and the set $\mathcal{F}$ of all minimal solutions to the polymatroid inequality

$$f(X) = \dim\left(Span\langle\bigcup_{i\in X} S_i\rangle\right) \geq D. \tag{10}$$

is the set of all "generalized bases", i.e., minimal subfamilies of $S_1, \ldots, S_n$ spanning the entire space M.

Given natural numbers k, m, and d, let us now consider k linear spaces $M_1, \ldots, M_k$ over $\mathbf{F}$, of dimension d each, and let $M = M_1 \otimes \cdots \otimes M_k$ be the space of all k-linear forms $\psi(y_1, \ldots, y_k)$, where $y_i \in M_i,\ i = 1, \ldots, k$. Thus, the dimension of M is $D = d^k$. Suppose further that we can find m linear forms in general position in each space M_i, i.e., some linear forms $l_{ij}(y_i),\ j = 1, \ldots, m$, such that any d of them form a basis. For instance, this is true for all d and m if the characteristic of $\mathbf{F}$ is zero. Now let us define $n = mk$ sets

$$S_{ij} = M_1 \otimes \cdots \otimes M_{i-1} \otimes l_{ij} \otimes M_{i+1} \otimes \cdots \otimes M_k, \quad i = 1, \ldots, k;\ j = 1, \ldots, m,$$

i.e., S_{ij} is the set of all k-linear forms in M which can be written as $l_{ij}(y_i)\psi'$, where ψ' is a $(k-1)$-linear form of the remaining variables $y_1, \ldots, y_{i-1}, y_{i+1}, \ldots, y_k$. Note that each S_{ik} is a subspace of dimension d^{k-1}. It is easy to see each generalized basis composed of S_{ij} consists of some d sets $S_{ij_1}, \ldots, S_{ij_d}$ selected for a fixed i, and vice versa. This is due to the fact that the forms l_{ij} are in general position. Thus, the hypergraph $\mathcal{F}$ of all minimal solutions to (10) is exactly $k \times \binom{[m]}{d}$, the disjoint union of k complete d-uniform hypergraphs, as in Example 1.1. Letting $d = k,\ m = 2^k$ and $t = D = d^k$, we again obtain as $k \to \infty$:

$$\begin{aligned}&|\mathcal{I}(\mathcal{F})| = 2^{(1-o(1))k^3}, \quad |\mathcal{F}| = 2^{(1-o(1))k^2}, \quad Q = 2n = 2^{(1+o(1))k},\\ &c(2Q, |\mathcal{F}|) = (1-o(1))\log k, \quad \log t = k \log k\end{aligned}$$

and consequently

$$|\mathcal{I}(\mathcal{F})| = |\mathcal{F}|^{(1-o(1))\log t/c(2Q,\beta)} \tag{11}$$

which should be compared with the bound in (9) for $r = 1$. Note that for the above example, the polymatroid function $f : 2^V \mapsto \mathbb{Z}_+$ in (10) can be written as follows:

$$f(X) = \begin{cases} d^k, & \text{if } |X \cap V_i| \geq d \text{ for some } i = 1, \ldots, k \\ d^k - \prod_{i=1}^k (d - |X \cap V_i|), & \text{otherwise,} \end{cases} \tag{12}$$

where $V = V_1 \cup \cdots \cup V_k$, and $V_1, \ldots, V_k$ are disjoint vertex sets of size m each. In fact, one can show that the bound of Theorem 4 is nearly sharp even with arbitrarily large number of inequalities r. To see this, let us consider $2k$ disjoint sets $V_1^0, V_1^1, V_2^0, V_2^1, \ldots, V_k^0, V_k^1$, of m vertices each, whose union is V. For each binary vector $y \in \{0,1\}^k$, let $f^y : 2^V \mapsto \mathbb{Z}_+$ be the function defined by (12) on the disjoint vertex sets $V_1^{y_1} \cup \cdots \cup V_k^{y_k}$. Consider the system of $r = 2^k$ inequalities

$$f^y(X) \geq d^k, \quad y \in \{0,1\}^k, \tag{13}$$

over the sets $X \subseteq V$. It is easy to see that any minimal feasible set for (13) is of the form $X^0 \cup X^1$, where $|X^0| = |X^1| = d-1$ and $X^0 \subseteq V_i^0, X^1 \subseteq V_i^1$ for some $i \in \{1,\ldots,k\}$. This gives $|\mathcal{F}| = k\binom{m}{d}$. On the other hand, any maximal infeasible set for (13) has the form $\bigcup_{i=1}^k (X_i^{y_i} \cup V_i^{1-y_i})$ where $y \in \{0,1\}^k$, $X_i^{y_i} \subseteq V_i^{y_i}$, and $|X_i^{y_i}| = d-1$. Consequently, we have $|\mathcal{I}(\mathcal{F})| = 2^k\binom{m}{d-1}^k$, and selecting $m = 2^k$ and $d = k$, we again get (11) as $k \to \infty$. $\square$

We prove Theorems 3, 4 and 5 in the next two sections. Then, in the last section of the paper, we discuss an algorithmic application of Theorem 4 dealing with the enumeration of all minimal feasible vectors $x \in \mathcal{L} = \mathcal{L}_1 \times \ldots \times \mathcal{L}_n$ for a system of polymatroid inequalities. Specifically, we assume that each factor-lattice $\mathcal{L}_i$ is explicitly represented by its precedence digraph, each of the r polymatroid function in (7) is defined by a polynomial-time evaluation oracle, and consider the following problem of generating all the elements of $\mathcal{F}$ incrementally:

GEN$(\mathcal{F}, \mathcal{X})$**:** *Given a system of polymatroid inequalities (7) and a collection $\mathcal{X} \subseteq \mathcal{F}$ of minimal feasible vectors for (7), either find a new minimal feasible vector $x \in \mathcal{F} \setminus \mathcal{X}$ for (7), or show that the given partial list is complete: $\mathcal{X} = \mathcal{F}$.*

Clearly, we can incrementally generate all elements of $\mathcal{F}$ by initializing $\mathcal{X} = \emptyset$ and then iteratively solving problem $GEN(\mathcal{F}, \mathcal{X})$ a number of $|\mathcal{F}| + 1$ times. In general, problem GEN$(\mathcal{F}, \mathcal{X})$ may be NP-hard already for a single polymatroid inequality $f_1(x) \geq t_1$ in n Boolean variables $x \in \{0,1\}^n$, provided that t_1 is exponentially large in n (see [8]). However, if $t = \max\{t_1, \ldots, t_r\} \leq 2^{\mathrm{polylog}(nr)}$, i.e. the right-hand sides in (7) are bounded by a quasi-polynomial function in the dimension of the system, then Theorem 4 can be used to reduce GEN$(\mathcal{F}, \mathcal{X})$ to the following *dualization problem*:

DUAL$(\mathcal{L}, \mathcal{B}, \mathcal{A})$**:** *Given a set $\mathcal{B} \subseteq \mathcal{L}$ in a lattice product $\mathcal{L} = \mathcal{L}_1 \times \cdots \times \mathcal{L}_n$ and a collection of maximal independent elements $\mathcal{A} \subseteq \mathcal{I}(\mathcal{B})$, either find a new maximal independent element $x \in \mathcal{I}(\mathcal{B}) \setminus \mathcal{A}$, or prove that the given collection is complete: $\mathcal{B} = \mathcal{I}(\mathcal{A})$.*

When $\mathcal{L} = \{0,1\}^n$, i.e. all factor lattices $\mathcal{L}_i$ are just chains $\{0,1\}$, then the above dualization problem turns into the dualization problem for hypergraphs and can be solved in quasi-polynomial time $poly(n) + (|\mathcal{B}| + |\mathcal{X}|)^{o(\log(|\mathcal{B}|+|\mathcal{X}|))}$ (see [10]). As shown in [9], the dualization problem can also be solved in quasi-polynomial time for the product $\mathcal{L} = \mathcal{L}_1 \times \cdots \times \mathcal{L}_n$ of factor-lattices of bounded (or poly-logarithmically bounded) widths $W(\mathcal{L}_i)$. Hence we arrive at the following result.

Theorem 5 *Consider a system of polymatroid inequalities (7), over the elements $x \in \mathcal{L} = \mathcal{L}_1 \times \cdots \times \mathcal{L}_n$ of the product of lattices of poly-logarithmically bounded width:* $\max\{W(\mathcal{L}_1), \ldots, W(\mathcal{L}_n)\} \leq \mathrm{polylog}(nr)$*, and in which the right hand sides are bounded by a quasi-polynomial in the dimension of the system:*

$$\max\{t_1, \ldots, t_m\} \leq 2^{\mathrm{polylog}(nr)}.$$

Then problem GEN($\mathcal{F}, \mathcal{X}$) is solvable in quasi-polynomial time.

Some applications of Theorem 5 are briefly discussed in the last section of the paper.

2 Proof of Theorem 3

The lower bound is obvious: let $\mathbf{T}$ be a binary tree with a proper mapping for $\mathcal{B}$, then by definition, the join element $s(\mathcal{P})$ given by (5) for a path $\mathcal{P}$ from the root of $\mathbf{T}$ to a leaf must be an independent element. Furthermore, the joins corresponding to two different paths $\mathcal{P}$ and $\mathcal{P}'$ are distinct. To see this, let $w \in \mathbf{N}(\mathbf{T})$ be the least common ancestor of the leaves ℓ, ℓ' corresponding respectively to these paths in $\mathbf{T}$, and assume that ℓ belongs to the right sub-tree descending from w. Then, again by the definition of a proper mapping, we must have $s(\mathcal{P}) \not\succeq L(w)$, since otherwise $s(\mathcal{P}) \succeq b(w)$ would follow by $L(w) \vee R(w) \succeq b(w)$. On the other hand, we have $s(\mathcal{P}') \succeq L(w)$, implying that, indeed, $s(\mathcal{P})$ and $s(\mathcal{P}')$ are different. Now, extending the elements $s(\mathcal{P})$, for all root-leaf paths $\mathcal{P}$, to maximal independent elements, we obtain a set of cardinally $|\mathbf{L}(\mathbf{T})|$ of maximal independent elements of $\mathcal{B}$.

To prove the upper bound of (6), we shall construct, for any $\mathcal{B} \subseteq \mathcal{L}$ of size $|\mathcal{B}| \geq 2$, a proper mapping with $|\mathbf{L}(\mathbf{T})|^{\log|\mathcal{B}|/c(2Q,|\mathcal{B}|)} \geq |\mathcal{I}(\mathcal{B})|$. Since for $|\mathcal{B}| = 1$, we have trivially $|\mathcal{I}(\mathcal{B})| \leq Q$, the upper bound in Theorem 3 will follow from Lemma 1 below. Throughout this section, we will use the notation $\beta = |\mathcal{B}|$ and $Q = \sum_{i=1}^n |\mathcal{L}_i|$.

Lemma 1 *Let $\mathcal{B} \subseteq \mathcal{L}$ be a subset of size $\beta \geq 2$, and let $\mathcal{A} \subseteq \mathcal{I}(\mathcal{B})$ be a subset of maximal independent elements of $\mathcal{B}$. Then there exists a binary tree $\mathbf{T}$ and a proper mapping $\phi : \mathbf{N}(\mathbf{T}) \mapsto \mathcal{B} \times \mathcal{L} \times \mathcal{L}$, such that for each path $\mathcal{P}$ from the root of $\mathbf{T}$ to a leaf, the join element $s_\phi(\mathcal{P})$ defined by (5) satisfy $s_\phi(\mathcal{P}) \preceq a$ for some $a \in \mathcal{A}$, and such that*

$$|\mathbf{L}(\mathbf{T})| \geq |\mathcal{A}|^{c(2Q,\beta)/\log\beta}.$$

To prove Lemma 1, we shall need a few more definitions. Call a subset $\mathcal{A} \subseteq \mathcal{L}$ *an antichain* if for every two distinct elements $a, a' \in \mathcal{A}$, neither $a \preceq a'$ nor $a' \preceq a$. For $\mathcal{A} \subseteq \mathcal{L}$, $i \in V = \{1, \ldots, n\}$, and $x \in \mathcal{L}_i$, let $\mathcal{A}_i(x) \stackrel{\text{def}}{=} \{a \in \mathcal{A} \mid a_i = x\}$ and for $b \in \mathcal{B}$ and $i \in V$, let $\mathcal{A}_b^{V \setminus i} \stackrel{\text{def}}{=} \{a \in \mathcal{A} \mid a_j \succeq b_j \text{ for all } j \in V \setminus i\}$.

Lemma 2 *Let $\epsilon > 0$ be a given positive number, and let $\mathcal{A} \subseteq \mathcal{L}$ be an antichain in $\mathcal{L}$ of size $|\mathcal{A}| \geq 1 + 1/\epsilon$. Then there exists an $i \in V$ and two distinct elements $x, y \in \mathcal{L}_i$ such that*

$$|\mathcal{A}_i(x)| \geq \frac{|\mathcal{A}|}{(1+\epsilon)Q} \qquad \text{and} \qquad |\mathcal{A}_i(y)| \geq \frac{|\mathcal{A}|}{(1+\epsilon)Q}. \tag{14}$$

Proof. For $i \in V$, define

$$X_i = \{x \in \mathcal{L}_i \ : \ |\mathcal{A}_i(x)| \geq \frac{|\mathcal{A}|}{(1+\epsilon)Q}\}, \qquad X = \bigcup_{i=1}^{n} X_i.$$

Define further

$$\mathcal{A}(X) = \{a \in \mathcal{A} \mid a_i \in X_i \text{ for all } i \in [n]\}.$$

Then it follows that

$$|\mathcal{A}(X)| > |\mathcal{A}|\frac{\epsilon}{1+\epsilon}. \tag{15}$$

Indeed, we have

$$\begin{aligned} |\mathcal{A} \setminus \mathcal{A}(X)| &= |\{a \in \mathcal{A} \ : \ a_i \notin X_i \text{ for some } i \in [n]\}| = |\bigcup_{i=1}^{n}\{a \in \mathcal{A} \ : \ a_i \notin X_i\}| \\ &\leq \sum_{i=1}^{n}\sum_{x \notin X_i} |\mathcal{A}_i(x)| < \sum_{i=1}^{n} |\mathcal{L}_i \setminus X_i| \frac{|\mathcal{A}|}{(1+\epsilon)Q} \leq \frac{|\mathcal{A}|}{(1+\epsilon)}, \end{aligned}$$

by the definition of X_i, establishing (15).

Let us note next that $|X_i| \geq 2$ for some $i = 1, \ldots, n$. If this was not the case, then since $\mathcal{A}(X)$ is an antichain, it follows that $|\mathcal{A}(X)| = 1$ implying by (15) that $|\mathcal{A}| < 1 + 1/\epsilon$, in contradiction to our assumptions. Thus there exist an $i \in V$, and distinct $x, y \in X_i \subseteq \mathcal{L}_i$ which satisfy (14). $\square$

Lemma 3 *For every subset $\mathcal{B} \subseteq \mathcal{L}$ and every $\mathcal{A} \subseteq \mathcal{I}(\mathcal{B})$, $|\mathcal{A}| \geq 1 + 1/\epsilon$ where $\epsilon > 0$, there exist $b \in \mathcal{B}$, $i \in V$, and two distinct elements $x, y \in \mathcal{L}_i$ such that*

(i) x, y satisfy (14), and $x \vee y \succeq b_i$.

(ii) $|\mathcal{A}_b^{V \setminus i} \cap \mathcal{A}_i(x)| \geq |\mathcal{A}|/((1+\epsilon)Q|\mathcal{B}|)$.

Proof. By Lemma 2, there exist $i \in V$ and $x, y \in \mathcal{L}_i$ such that (14) is satisfied, and either $x \prec y$ or x, y are incomparable in $\mathcal{L}_i$. Letting $z = x \vee y$ and noting that $z \succ x$, we conclude that, for every $a \in \mathcal{A}_i(x)$, there exists an element $b \in \mathcal{B}$ such that $b_i \preceq z$, and $b_j \preceq a_j$ for all $j \neq i$, by the maximality of the independent element a, i.e., $\mathcal{A}_i(x) = \bigcup_{b \in \mathcal{B}}\{a \in \mathcal{A}_i(x) \mid a_i \prec z, a_j \succeq b_j \text{ for all } j \neq i\}$. From this, it follows that there must exist an element $b \in \mathcal{B}$ such that

$$|\{a \in \mathcal{A}_i(x) \mid a_i \prec z, a_j \succeq b_j \text{ for all } j \neq i\} \geq |\mathcal{A}_i(x)|/|\mathcal{B}| \geq \frac{|\mathcal{A}|}{(1+\epsilon)Q|\mathcal{B}|}.$$

This immediately gives $|\mathcal{A}_b^{V \setminus i} \cap \mathcal{A}_i(x)| \geq \frac{|\mathcal{A}|}{(1+\epsilon)Q|\mathcal{B}|}$, and the lemma follows. $\square$

Proof of Lemma 1. Let $\alpha = |\mathcal{A}|$. If $\alpha = 0$, there is nothing to prove. If $\alpha = 1$ then the statement is trivial: a tree with a single leaf, and no internal nodes satisfies the lemma.

Otherwise, we shall use Lemma 3 with $\epsilon = 1$ and $\mathcal{A} \subseteq \mathcal{I}(\mathcal{B})$ to construct a proper mapping, by induction on $\alpha \geq 2$.

Let $b \in \mathcal{B}$, $i \in V$, and $x, y \in \mathcal{L}_i$ be elements that satisfy the conditions of Lemma 3. If x, y are incomparable in $\mathcal{L}_i$, let us assume without loss of generality that $x \prec y$. To build a binary tree $\mathbf{T}$ on $\mathcal{B}$, let us associate to its root $\mathbf{r}$ the element $b(\mathbf{r}) = b$, and the two elements $L(\mathbf{r}), R(\mathbf{r}) \in \mathcal{L}$ defined by

$$\begin{aligned} L(\mathbf{r}) &= (\mathbf{0}_1, \ldots, \mathbf{0}_{i-1}, y, \mathbf{0}_{i+1}, \ldots, \mathbf{0}_n), \\ R(\mathbf{r}) &= (b_1, \ldots, b_{i-1}, x, b_{i+1}, \ldots, b_n), \end{aligned}$$

where $\mathbf{0} = (\mathbf{0}_1, \ldots, \mathbf{0}_n)$ is the minimum element of $\mathcal{L}$. Clearly $L(\mathbf{r}) \vee R(\mathbf{r}) \succeq b$ since $x \vee y \succeq b_i$. Define further $\mathcal{A}' = \{a \in \mathcal{A} \mid a \succeq L(\mathbf{r})\}$, and $\mathcal{A}'' = \{a \in \mathcal{A} \mid a \succeq R(\mathbf{r})\}$ and note that

(a) $\mathcal{A}' \cap \mathcal{A}'' = \emptyset$: if there is an $a \in \mathcal{A}' \cap \mathcal{A}''$ then $a \succeq b(\mathbf{r})$ would follow by $a \succeq L(\mathbf{r})$, $a \succeq R(\mathbf{r})$, and $L(\mathbf{r}) \vee R(\mathbf{r}) \succeq b(\mathbf{r})$,

(b) $|\mathcal{A}'| > 0$ and $|\mathcal{A}''| > 0$, by (14),

(c) $\mathcal{A}' \subseteq \mathcal{I}(\mathcal{B})$ and $\mathcal{A}'' \subseteq \mathcal{I}(\mathcal{B})$.

It follows by (a) and (b) that $\mathcal{A}' \subset \mathcal{A}$ and $\mathcal{A}'' \subset \mathcal{A}$. It follows furthermore by Lemma 3 that $|\mathcal{A}'| \geq \frac{\alpha}{2Q}$ and $|\mathcal{A}''| \geq \frac{\alpha}{2\beta Q}$. Therefore, we can conclude from (c) by induction that there exist binary trees $\mathbf{T}'$ and $\mathbf{T}''$ of sufficiently large number of leaves:

$$|\mathbf{L}(\mathbf{T}')| \geq |\mathcal{A}'|^{\frac{c(2Q,\beta)}{\log \beta}} \qquad \text{and} \qquad |\mathbf{L}(\mathbf{T}'')| \geq |\mathcal{A}''|^{\frac{c(2Q,\beta)}{\log \beta}}, \tag{16}$$

and proper mappings $\phi' : \mathbf{N}(\mathbf{T}') \mapsto \mathcal{B} \times \mathcal{L} \times \mathcal{L}$ and $\phi'' : \mathbf{N}(\mathbf{T}'') \mapsto \mathcal{B} \times \mathcal{L} \times \mathcal{L}$, such that if we join $\mathbf{T}'$ and $\mathbf{T}''$ as the two children of the root w of $\mathbf{T}$, we get

$$\begin{aligned} |\mathbf{L}(\mathbf{T})| &= |\mathbf{L}(\mathbf{T}')| + |\mathbf{L}(\mathbf{T}'')| \\ &\geq |\mathcal{A}'|^{c(2Q,\beta)/\log \beta} + |\mathcal{A}''|^{c(2Q,\beta)/\log \beta} \\ &\geq \left(\frac{\alpha}{2Q}\right)^{c(2Q,\beta)/\log \beta} + \left(\frac{\alpha}{2Q\beta}\right)^{c(2Q,\beta)/\log \beta} \\ &= \alpha^{c/\log \beta}\left[(2Q)^{-c/\log \beta} + (2Q\beta)^{-c/\log \beta}\right] \\ &= \alpha^{c/\log \beta}, \end{aligned}$$

where the last equality holds by (3). The mapping $\phi : \mathbf{N}(\mathbf{T}) \mapsto \mathcal{B} \times \mathcal{L} \times \mathcal{L}$ will be defined in the obvious way: $\phi(\mathbf{r}) = (b(\mathbf{r}), L(\mathbf{r}), R(\mathbf{r}))$, $\phi(u) = \phi'(u)$ for all $u \in \mathbf{N}(\mathbf{T}')$, and $\phi(u) = \phi''(u)$ for all $u \in \mathbf{N}(\mathbf{T}'')$. Clearly, condition (b) above implies also by induction that, for any path $\mathcal{P}$ from the root of $\mathbf{T}$ to a leaf, we have $s_\phi(\mathcal{P}) \preceq a$ for some $a \in \mathcal{A}$.

Let us finally verify that the mapping ϕ, built in the above way, is indeed proper. Let $\mathcal{P}$ be any path from the root $\mathbf{r}$ of $\mathbf{T}$ to a leaf ℓ in the left sub-tree $\mathbf{T}'$. Let $\mathcal{P}'$ be the sub-path of $\mathcal{P}$ from the root of $\mathbf{T}'$ to ℓ. By induction, there is an $a \in \mathcal{A}'$ such that $a \succeq s_{\phi'}(\mathcal{P}')$. Then $a \succeq L(\mathbf{r}) \vee s_{\phi'}(\mathcal{P}') = s_\phi(\mathcal{P})$, implying that $s_\phi(\mathcal{P})$ is independent of $\mathcal{B}$. A similar argument will also show that the join element $s_\phi(\mathcal{P})$ for any path $\mathcal{P}$ from the root $\mathbf{r}$ of $\mathbf{T}$ to a leaf in the right sub-tree $\mathbf{T}''$ is independent, and the lemma follows. $\square$

3 Proof of Theorems 4 and 5

Proof of Theorem 4. Fix $i \in \{1,\dots,r\}$, and let $\mathcal{B} \subseteq \{x \in \mathcal{L} \mid f_i(x) \geq t_i\}$ be a non-empty subset of feasible vectors for the i-th inequality of (7). To prove the theorem, it is enough to show that

$$|\mathcal{I}(\mathcal{B}) \cap \{x \in \mathcal{L} \mid f_i(x) < t_i\}| \leq \max\left\{Q, |\mathcal{B}|^{\log t_i/c(2Q,|\mathcal{B}|)}\right\}, \tag{17}$$

from which (8) follows by substituting $\mathcal{B} = \mathcal{X}$ (which is a set of feasible vectors for every inequality in the system), summing the inequalities (17) over $i = 1,\dots,r$, and noting that

$$\mathcal{I}(\mathcal{X}) \cap \{x \in \mathcal{L} \mid x \text{ is infeasible for (7)}\} = \bigcup_{i=1}^{r}(\mathcal{I}(\mathcal{X}) \cap \{x \in \mathcal{L} \mid f_i(x) < t_i\}).$$

To prove (17), let us invoke Lemma 1 with $\mathcal{B} \subseteq \mathcal{L}$, and $\mathcal{A} \stackrel{\text{def}}{=} \mathcal{I}(\mathcal{B}) \cap \{x \in \mathcal{L} \mid f_i(x) < t_i\} \subseteq \mathcal{I}(\mathcal{B})$, and obtain a proper mapping $\phi : \mathbf{N}(\mathbf{T}) \mapsto \mathcal{B} \times \mathcal{L} \times \mathcal{L}$ of the elements of $\mathcal{B}$ to the nodes of the binary tree $\mathbf{T}$. For a node $w \in \mathbf{N}(\mathbf{T}) \cup \mathbf{L}(\mathbf{T})$, let $\mathcal{P}_w$ be the path from the root of $\mathbf{T}$ to w, let $s(\mathcal{P}_w)$ be the join element defined by (5) on $\mathcal{P}_w$, and let $\mathbf{T}(w)$ denote the sub-tree of $\mathbf{T}$ rooted at w. Then we shall show by induction that

$$f_i(s(\mathcal{P}_w)) \leq t_i - |\mathbf{L}(\mathbf{T}(w))| \tag{18}$$

holds for every node w of the binary tree $\mathbf{T}$. Since f is non-negative, it follows that

$$|\mathbf{L}(\mathbf{T}(w))| \leq t_i$$

which, if applied to the root of $\mathbf{T}$, gives $|\mathbf{L}(\mathbf{T})| \leq t_i$. Now we conclude by Lemma 1 that

$$|\mathcal{I}(\mathcal{B}) \cap \{x \in \mathcal{L} \mid f_i(x) < t_i\}| \leq \max\left\{Q, |\mathcal{B}|^{\log|\mathbf{L}(\mathbf{T})|/c(2Q,|\mathcal{B}|)}\right\},$$

proving (17).

To see (18), let us apply induction by the size of $\mathbf{L}(\mathbf{T}(w))$. Clearly, if $w = \ell$ is a leaf of $\mathbf{T}$, then $|\mathbf{L}(\mathbf{T}(\ell))| = 1$, and (18) follows from the monotonicity of f_i and the fact that $s(\mathcal{P}_\ell) \preceq a$ for some $a \in \mathcal{A}$. Let us assume now that w is an internal node of $\mathbf{T}$ with u and v as its immediate successors. Then $|\mathbf{L}(\mathbf{T}(w))| = |\mathbf{L}(\mathbf{T}(u))| + |\mathbf{L}(\mathbf{T}(v))|$, and $s(\mathcal{P}_w) \preceq s(\mathcal{P}_u) \wedge s(\mathcal{P}_v)$. By our inductive hypothesis, and since f is monotone and submodular, we have the inequalities

$$\begin{aligned} f_i(s(\mathcal{P}_u) \vee s(\mathcal{P}_v)) + f_i(s(\mathcal{P}_w)) &\leq f_i(s(\mathcal{P}_u) \vee s(\mathcal{P}_v)) + f_i(s(\mathcal{P}_u) \wedge s(\mathcal{P}_v)) \\ &\leq f_i(s(\mathcal{P}_u)) + f_i(s(\mathcal{P}_v)) \\ &\leq t_i - |\mathbf{L}(\mathbf{T}(u))| + t_i - |\mathbf{L}(\mathbf{T}(v))| \\ &= 2t_i - |\mathbf{L}(\mathbf{T}(w))|. \end{aligned}$$

Since ϕ is a proper mapping, we have $s(\mathcal{P}_u) \vee s(\mathcal{P}_v) \succeq b(w) \in \mathcal{B}$, and thus $f_i(s(\mathcal{P}_u) \vee s(\mathcal{P}_v)) \geq f_i(b(w)) \geq t_i$ by the monotonicity of f_i, and by our assumption that $\mathcal{B} \subseteq \{x \in \mathcal{L} \mid f_i(x) \geq t_i\}$. Thus, from the above inequality we get $t_i + f_i(s(\mathcal{P}_w)) \leq f_i(s(\mathcal{P}_u) \vee s(\mathcal{P}_v)) + f_i(s(\mathcal{P}_w)) \leq 2t_i - |\mathbf{L}(\mathbf{T}(w))|$, from which (18) follows. □

Using Theorem 4 we immediately obtain the following result.

Corollary 1 *For any system of polymatroid inequalities (7), described by a polynomial-time feasibility oracle, and having quasi-polynomially right-hand sides: $t_1, \ldots, t_r \leq 2^{\mathrm{polylog}(nr)}$, problem GEN($\mathcal{F}, \mathcal{X}$) is reducible, in quasi-polynomial time to dualization in products of lattices.*

Proof. It is easy to see that the first minimal feasible vector of $\mathcal{F}$ can be found (or $\mathcal{F} = \emptyset$ can be recognized) by evaluating (7) a number of $Q + 1$-times. Thus we may assume that we are given a partial list $\mathcal{X} \subseteq \mathcal{F}$ of size $|\mathcal{X}| \geq 1$. Starting with $\mathcal{B} = \mathcal{X}$ and $\mathcal{A} = \emptyset$, we solve the dualization problem DUAL($\mathcal{L}, \mathcal{B}, \mathcal{A}$) repeatedly, until either a vector $x \in \mathcal{I}(\mathcal{X})$ that is feasible for (7) is found, or the whole set $\mathcal{I}(\mathcal{X})$ has been generated. In the former case, we obtain a new element in $\mathcal{F} \setminus \mathcal{X}$. In the latter case, we conclude that all the elements of $\mathcal{I}(\mathcal{X})$ are infeasible for the system (7), and therefore, the given list is complete, i.e. $\mathcal{X} = \mathcal{F}$. In both cases, we are assured by Theorem 4 that the number of generated elements before termination does not exceed $|\mathcal{X}|^{\mathrm{polylog}(nr)}$. □

Combining Corollary 2 with the fact that the dualization problem on products of lattices of poly-logarithmically bounded width can be solved in quasi-polynomial time, we readily obtain Theorem 5. □

4 Applications

Let us conclude by two examples of polymatroid systems defined on products of lattices.

Example 4.1 *(Maximal boxes containing specified numbers of points)* Given a set of n-dimensional points $\mathcal{S} \subseteq \mathbb{R}^n$ and a coloring $C : \mathcal{S} \mapsto \{1, 2, \ldots, r\}$ of the point set, suppose that we want to generate all maximal n-dimensional boxes that contain at most t_1 points of $\mathcal{S}$ of the first color, at most t_2 points of the second color, ..., and at most t_r points of the r-th color, where $0 \leq t_1, t_2, \ldots, t_r \leq |\mathcal{S}|$ are given integer thresholds. We shall assume without loss of generality that the generated boxes *minimally* bound the points inside them, i.e. there must exist a point of $\mathcal{S}$ on each of the n sides of each generated box. This problem can be described as of generating minimal feasible solutions of a system of polymatroid inequalities over a product of $2n$ chains (or more precisely, n join semi-lattices). Indeed, consider the set of projection points $\mathcal{S}_i \stackrel{\mathrm{def}}{=} \{p_i \in \mathbb{R} \mid p \in \mathcal{S}\}$, for $i = 1, \ldots, n$. Clearly, the lower and upper end-points a_i, b_i of each candidate box in the i-th dimension belong to the set $\mathcal{S}_i$. Thus, letting $\mathcal{C}_{2i-1} \stackrel{\mathrm{def}}{=} \mathcal{S}_i$ and $\mathcal{C}_{2i} \stackrel{\mathrm{def}}{=} \mathcal{S}_i^*$ be the two chains whose elements are $\mathcal{S}_i$, ordered in increasing and decreasing orders respectively, we conclude that the projection $x_i = [a_i, b_i]$ of each box in the i-th dimension belongs to the join semi-lattice $\mathcal{L}_i \stackrel{\mathrm{def}}{=} \{(a_i, b_i) \in \mathcal{C}_{2i-1} \times \mathcal{C}_{2i} \mid a_i \leq b_i\}$. (Correspondingly $\mathcal{L}_i \cup \{\mathbf{0}_i\}$ is the *lattice of intervals* whose elements are the different intervals defined by the projection points $\mathcal{S}_i$. The meet of any two intervals is their *intersection*, and the join is their *span*, i.e., the minimum interval containing both of them. The minimum element $\mathbf{0}_i$ of $\mathcal{L}_i$ is the empty interval.) Clearly, the semi-lattice product $\mathcal{L} = \mathcal{L}_1 \times \cdots \times \mathcal{L}_n$ defines the set of all possible boxes that may satisfy the required conditions. Let $f_i : \mathcal{C} \mapsto \mathbb{Z}_+$, for $i = 1, \ldots, r$, be the functions defined by

$$f_i(x) = \begin{cases} |\{p \in \mathcal{S} \mid C(p) = i \text{ and } p \text{ is } \textit{not} \text{ contained inside the box } x\}|, & \text{if } x \in \mathcal{L} \\ |\{p \in \mathcal{S} \mid C(p) = i\}|, & \text{otherwise,} \end{cases}$$

over the elements of $\mathcal{C} \stackrel{\text{def}}{=} (\mathcal{C}_1 \times \mathcal{C}_2) \times (\mathcal{C}_3 \times \mathcal{C}_4) \times \cdots \times (\mathcal{C}_{2n-1} \times \mathcal{C}_{2n})$. Then the function $f_i(\cdot)$ is polymatroid over the elements of the *dual lattice* $\mathcal{C}^*$ (that is, the lattice $\mathcal{C}^*$ with the same set of elements as $\mathcal{C}$, but such that $x \prec y$ in $\mathcal{C}^*$ whenever $x \succ y$ in $\mathcal{C}$), we have an evaluation oracle for f_i, and the minimal elements $x \in \mathcal{C}^*$, feasible for the system

$$f_i(x) \geq t_i, \quad i = 1, \ldots, r,$$

correspond to the maximal boxes that contain t_i points from color i, for $i = 1, \ldots, r$ (plus exactly $\sum_{i=1}^{n}(|\{p \in \mathcal{S} \mid C(p) = i\}| - 1)$ maximal *artificial* boxes, i.e., boxes with end points $a_i > b_i$). Since the dualization problem on products of chains (where the maximum width is 1) can be solved in quasi-polynomial time, it follows by Theorem 5 that those maximal boxes can be generated in incremental quasi-polynomial time. This problem has some applications in quantitative data mining [14]. □

Example 4.2 *(Minimal infrequent elements in databases with semi-lattice attributes)* In many database applications, the data attributes assume values ranging over products of lattices or semi-lattices of small width, e.g. quantitative attributes [14], taxonomies (hierarchical databases) [13], and lattices of small sizes in logical analysis of data [7].

The notion of frequent sets in data mining [1] has a natural generalization over products of semi-lattices. Formally, consider a database $\mathcal{D} \subseteq \mathcal{L} = \mathcal{L}_1 \times \cdots \times \mathcal{L}_n$ of transactions, each of which is an n-dimensional vector of attributes over $\mathcal{L}$. For an element $x \in \mathcal{L}$, denote by

$$\mathcal{S}(x) = \mathcal{S}_{\mathcal{D}}(x) \stackrel{\text{def}}{=} \{p \in \mathcal{D} \mid p \succeq x\},$$

the set of transactions in $\mathcal{D}$ that *support* x. Note that, by this definition, the function $|\mathcal{S}(.)| : \mathcal{L} \mapsto \mathbb{Z}_+$ is an anti-monotone *supermodular* function, and hence, the function $f : \mathcal{L} \mapsto \mathbb{Z}_+$, defined by $f(x) = |\mathcal{D}| - |\mathcal{S}(x)|$ is polymatroid.

Given $\mathcal{D} \subseteq \mathcal{L}$ and an integer threshold t, an element $x \in \mathcal{L}$ is said to be t-frequent if it is supported by at least t transactions in the database, i.e. if $|\mathcal{S}_{\mathcal{D}}(x)| \geq t$. Conversely, $x \in \mathcal{L}$ is said to be t-infrequent if $|\mathcal{S}_{\mathcal{D}}(x)| < t$. Since the function $|\mathcal{S}_{\mathcal{D}}(.)|$ is anti-monotone, we may restrict our attention only to *maximal frequent* and *minimal infrequent* elements. Denote by $\mathcal{F}_{\mathcal{D},t}$ the set of all minimal t-infrequent elements of $\mathcal{L}$ with respect to the database $\mathcal{D}$. Then $\mathcal{I}(\mathcal{F}_{\mathcal{D},t})$ is the set of all maximal t-frequent elements. It is clear that $\mathcal{F}_{\mathcal{D},t}$ is the set of minimal feasible solutions of the polymatroid inequality $f(x) \geq |\mathcal{D}| - t + 1$. Consequently, for products of lattices with bounded width, this set can be generated in incremental quasi-polynomial time by Theorem 5. The special case of the above result for databases $\mathcal{D}$ of binary attributes can be found in [5, 6]. $\square$

We remark finally that many other examples of polymatroid functions and systems in the Boolean case can be found in [11] and [4].

References

[1] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen and A. I. Verkamo, Fast discovery of association rules, in U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth and R. Uthurusamy eds., *Advances in Knowledge Discovery and Data Mining*, pp. 307-328, AAAI Press, Menlo Park, California, 1996.

[2] V.E. Alekseev. On the number of dead-end independent sets in graphs from hereditary classes, *Combinatoric-algebraic methods in discrete optimization* **6** (1991) 5-8 (Nizhni Novgorod, in Russian).

[3] E. Balas and C. S. Yu, On graphs with polynomially solvable maximal-weight clique problem, *Networks* **19** (1989) 247–253.

[4] E. Boros, K. Elbassioni, V. Gurvich and L. Khachiyan, An inequality for polymatroid functions and its applications, DIMACS Technical Report 2001-14, Rutgers University, **http://dimacs.rutgers.edu/TechnicalReports/2001.html**, to appear in *Discrete Applied Mathematics.*

[5] E. Boros, V. Gurvich, L. Khachiyan and K. Makino, Dual bounded generating problems: partial and multiple transversals of a hypergraph. *SIAM Journal on Computing* 30 (6) (2001) pp. 2036-2050.

[6] E. Boros, V. Gurvich, L. Khachiyan and K. Makino, On the complexity of generating maximal frequent and minimal infrequent sets, in *19th Int. Symp. on Theoretical Aspects of Computer Science, (STACS)*, March 2002, LNCS 2285, pp. 133–141.

[7] E. Boros, P.L. Hammer and J.N. Hooker, Predicting cause-effect relationships from incomplete discrete observations, *SIAM Journal on Discrete Mathematics*, **7** (1994), 481–491.

[8] K. Elbassioni, Incremental Algorithms for Enumerating Extremal Solutions of Monotone Systems of Submodular inequalities and Their Applications, Ph.D. Thesis, Department of Computer Science, Rutgers University, New Jersey, 2002.

[9] K. Elbassioni, An algorithm for dualization in products of lattices, DIMACS Technical Report 2002-26, Rutgers University, **http://dimacs.rutgers.edu/TechnicalReports/2002.html**, to appear in *Proc. 10th Annual European Symposium on Algorithms (ESA 2002)*, 17-21 September, 2002.

[10] M. L. Fredman and L. Khachiyan, On the complexity of dualization of monotone disjunctive normal forms, *Journal of Algorithms*, 21 (1996) pp. 618-628.

[11] L. Lovász, Submodular functions and convexity, in *Mathematical Programming: The State of the Art*, Bonn 1982, pp. 235-257, (Springer Verlag, 1983).

[12] E. Prisner, Graphs with few cliques, *Graph theory, combinatorics and algorithms*, **1,2** (Kalamazoo, MI, 1992) 945-956, Wiley-Intersi. Publ., Wiley, New York, 1995.

[13] R. Srikant and R. Agrawal, Mining generalized association rules, in *Proceedings of the 21st International Conference on Very Large Data Bases*, pp. 407–419, 1995.

[14] R. Srikant and R. Agrawal, Mining quantitative association rules in large relational tables, in *Proceedings of the ACM-SIGMOD 1996 Conference on Management of Data*, pp. 1–12, 1996.