DIMACS Technical Report 2001-12
April 2001

Dual-Bounded Generating Problems: All Minimal Integer Solutions for a Monotone System of Linear Inequalities [1,2]

by

E. Boros[3] K. Elbassioni[4] V. Gurvich[5] L. Khachiyan[6]
K. Makino[7]

[1] An extended abstract of this paper containing some of the results will appear in *Proceedings of the 28th International Colloquium on Automata, Languages and Programming* (*ICALP2001*), Crete, Greece, July 8-12, 2001.

[2] The research of the first and third authors was supported in part by the Office of Naval Research (Grant N00014-92-J-1375), and the National Science Foundation (Grant DMS 98-06389). The research of the third and forth authors was supported in part by the National Science Foundation (Grant CCR-9618796). Visits of the third author to Rutgers University were also supported by DIMACS, the National Science Foundation's Center for Discrete Mathematics and Theoretical Computer Science.

[3] RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway NJ 08854-8003; (boros@rutcor.rutgers.edu).

[4] Department of Computer Science, Rutgers University, 110 Frelinghuysen Road, Piscataway, New Jersey 08854-8004; (elbassio@paul.rutgers.edu).

[5] RUTCOR, Rutgers University; 640 Bartholomew Road, Piscataway NJ 08854-8003; (gurvich@rutcor.rutgers.edu)

[6] Department of Computer Science, Rutgers University, 110 Frelinghuysen Road, Piscataway, New Jersey, 08854-8004; (leonid@cs.rutgers.edu)

[7] Division of Systems Science, Graduate School of Engineering Science, Osaka University, Toyonaka, Osaka, 560-8531, Japan; (makino@sys.es.osaka-u.ac.jp)

ABSTRACT

We consider the problem of enumerating all minimal integer solutions of a monotone system of linear inequalities. We first show that for any monotone system of r linear inequalities in n variables, the number of maximal infeasible integer vectors is at most rn times the number of minimal integer solutions to the system. This bound is accurate up to a $polylog(r)$ factor and leads to a polynomial-time reduction of the enumeration problem to a natural generalization of the well-known dualization problem for hypergraphs, in which dual pairs of hypergraphs are replaced by dual collections of integer vectors in a box. We provide a quasi-polynomial algorithm for the latter dualization problem. These results imply, in particular, that the problem of incrementally generating all minimal integer solutions to a monotone system of linear inequalities can be done in quasi-polynomial time.

1 Introduction

Consider a system of r linear inequalities in n integer variables

$$Ax \geq b, \quad x \in \mathcal{C} = \{x \in \mathbb{Z}^n \mid 0 \leq x \leq c\}, \tag{1}$$

where A is a rational $r \times n$-matrix, b is a rational r-vector, and c is a non-negative integral n-vector some or all of whose components may be infinite. We assume that (1) is a monotone system of inequalities: if $x \in \mathcal{C}$ satisfies (1) then any vector $y \in \mathcal{C}$ such that $y \geq x$ is also feasible. For instance, (1) is monotone if the matrix A is non-negative. Let us denote by $\mathcal{F} = \mathcal{F}_{A,b,c}$ the set of all minimal feasible integral vectors for (1), i.e., $y \in \mathcal{F}$ if there is no solution x of (1) such that $x \leq y,\ x \neq y$. Then, we have

$$\{x \in \mathcal{C} \mid Ax \geq b\} = \bigcup_{y \in \mathcal{F}} \{\, x \in \mathcal{C} \mid x \geq y\}.$$

In this paper, we are concerned with the problem of incrementally generating $\mathcal{F}$:

$GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$**:** *Given a monotone system (1) and a set $\mathcal{X} \subseteq \mathcal{F}_{A,b,c}$ of minimal feasible vectors for (1), either find a new minimal integral vector satisfying (1), or show that $\mathcal{X} = \mathcal{F}_{A,b,c}$.*

The entire set $\mathcal{F} = \mathcal{F}_{A,b,c}$ can be constructed by initializing $\mathcal{X} = \emptyset$ and iteratively solving the above problem $|\mathcal{F}| + 1$ times.

If A is a binary matrix, and b, c are vectors of all ones, then $\mathcal{F}$ is the set of (characteristic vectors of) all minimal transversals to the hypergraph defined by the rows of A. In this case, problem $GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$ turns into the well-known *hypergraph dualization problem*: incrementally enumerate all the minimal transversals (equivalently, all the maximal independent sets) for a given hypergraph The case where A is binary, c is the vector of all ones and b is arbitrary, is equivalent with the generation of so-called *multiple* transversals [4]. If A is integral and $c = +\infty$, the generation of $\mathcal{F}$ can also be regarded as the computation of the Hilbert basis for the ideal $\{x \in \mathbb{Z}^n \mid Ax \geq b,\ x \geq 0\}$. One more application of problem $GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$ is related to stochastic programming, more precisely to the generation of minimal *p-efficient* points for a given probability distribution of a discrete random variable $\xi \in \mathbb{Z}^n$. An integer vector $y \in \mathbb{Z}^n$ is called p-efficient, if $Prob(\xi \leq y) \geq p$. It is known that for every probability distribution and every $p > 0$ there are finitely many minimal p-efficient points and furthermore, for r-concave probability distributions these points are exactly the minimal integral points of a corresponding convex monotone system (see, e.g., [16]).

Let $J^* = \{j \mid c_j = \infty\}$ and $J_* = \{1, \dots, n\} \setminus J^*$ be, respectively, the sets of unbounded and bounded integer variables in (1). Consider an arbitrary vector $x = (x_1, \dots, x_n) \in \mathcal{F}_{A,b,c}$ such that $x_j > 0$ for some $j \in J^*$. Then it is easy to see that

$$x_j \leq \max_{i: a_{ij} > 0} \left\lceil \frac{b_i - \sum_{k \in J_*} \min\{0, a_{ik}\} c_k}{a_{ij}} \right\rceil < +\infty. \tag{2}$$

[Indeed, let x' be the vector obtained by decreasing the j^{th} component of x by 1, then $x' \in \mathcal{C}$ is infeasible for (1) and hence $b_i - a_{ij} \leq a_i x - a_{ij} = a_i x' < b_i$ for some $i \in \{1, \ldots, r\}$, implying $a_{ij} > 0$. Thus $a_i x < b_i + a_{ij}$. Since $a_{ij}x_j + \sum_{k \in J_*} \min\{0, a_{ik}\}c_k \leq a_{ij}x_j + \sum_{k \in J_*} \min\{0, a_{ik}\}x_k \leq a_{ij}x_j + \sum_{k \in J_*} a_{ik}x_k \leq a_i x$, where the last inequality follows from the non-negativity of the restriction of A on J^*, we have $a_{ij}x_j + \sum_{k \in J_*} \min\{0, a_{ik}\}c_k < b_i + a_{ij}$. This implies (2).]

Since the bounds of (2) are easy to compute, and since appending these bounds to (1) does not change the set $\mathcal{F}_{A,b,c}$, we shall assume in the sequel that all components of the non-negative vector c are finite, even though this may not be the case for the original system. This assumption does not entail any loss of generality and allows us to consider $\mathcal{F}_{A,b,c}$ as a system of integral vectors in a finite box. We shall also assume that the input monotone system (1) is feasible, i.e., $\mathcal{F}_{A,b,c} \neq \emptyset$. For a finite and non-negative c this is equivalent to $Ac \geq b$. In addition, we say that system (1) is non-trivial if $\mathcal{F}_{A,b,c} \neq \mathcal{C}$, i.e., $0 \notin \mathcal{F}_{A,b,c}$.

Let $\mathcal{A}$ be a collection of integral vectors in $\mathcal{C}$ and let $\mathcal{A}^+ = \{x \in \mathcal{C} \mid x \geq a \text{ for some } a \in \mathcal{A}\}$ and $\mathcal{A}^- = \{x \in \mathcal{C} \mid x \leq a \text{ for some } a \in \mathcal{A}\}$ denote the ideal and filter generated by $\mathcal{A}$. Any element in $\mathcal{C} \setminus \mathcal{A}^+$ is called *independent of* $\mathcal{A}$. Let $\mathcal{I}(\mathcal{A})$ be the set of all maximal independent elements for $\mathcal{A}$, then for any finite box $\mathcal{C}$ we have the decomposition:

$$\mathcal{A}^+ \cap \mathcal{I}(\mathcal{A})^- = \emptyset, \qquad \mathcal{A}^+ \cup \mathcal{I}(\mathcal{A})^- = \mathcal{C}. \tag{3}$$

In particular, if $\mathcal{A}$ is the set $\mathcal{F} = \mathcal{F}_{A,b,c}$ of all minimal feasible integral vectors for (1), then the ideal $\mathcal{F}^+$ is the solution set of (1), while the filter $\mathcal{C} \setminus \mathcal{F}^+$ is generated by the set $\mathcal{I}(\mathcal{F})$ of all maximal infeasible integral vectors for (1):

$$\{x \in \mathcal{C} \mid Ax \not\geq b\} = \bigcup_{y \in \mathcal{I}(\mathcal{F})} \{y\}^-.$$

It is known that the problem of incrementally generating all maximal infeasible vectors for (1) is NP-hard even if c is the vector of all ones and the matrix A is binary:

Proposition 1 (c.f. [11]) *Given a binary matrix A and a set $\mathcal{X} \subseteq \mathcal{I}(\mathcal{F}_{A,b,1})$ of maximal infeasible Boolean vectors for $Ax \geq b$, $x \in \{0,1\}^n$, it is NP-complete to decide if the set $\mathcal{X}$ can be extended, that is if $\mathcal{I}(\mathcal{F}_{A,b,1}) \setminus \mathcal{X} \neq \emptyset$.*

It was conjectured in [10] that problem $GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$ cannot be solved in polynomial time either, unless $P = NP$. In this paper we show that $GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$ can be solved in quasi-polynomial time.

Theorem 1 *Problem $GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$ can be solved in time poly$(|input|) + t^{o(\log t)}$, where $t = \max\{n, r, |\mathcal{X}|\}$.*

To prove this result, we first bound the number of maximal infeasible vectors for (1) in terms of the dimension of the system and the number of minimal feasible vectors.

Theorem 2 *Suppose that the monotone system (1) is feasible, i.e., $Ac \geq b$. Then for any non-empty set $\mathcal{X} \subseteq \mathcal{F}_{A,b,c}$ we have*

$$|\mathcal{I}(\mathcal{X}) \cap \mathcal{I}(\mathcal{F}_{A,b,c})| \leq r \sum_{x \in \mathcal{X}} p(x), \tag{4}$$

where $p(x)$ is the number of positive components of x. In particular, $|\mathcal{I}(\mathcal{X}) \cap \mathcal{I}(\mathcal{F}_{A,b,c})| \leq rn|\mathcal{X}|$, which for $\mathcal{X} = \mathcal{F}_{A,b,c}$ implies the inequality $|\mathcal{I}(\mathcal{F}_{A,b,c})| \leq rn|\mathcal{F}_{A,b,c}|$.

It should be mentioned that the bounds of Theorem 2 are sharp for $r = 1$, e.g., for the inequality $x_1 + \ldots + x_n \geq n$ in binary variables. For large r, these bounds are accurate up to a factor poly-logarithmic in r. To see this, let $n = 2k$ and consider the monotone system of $r = 2^k$ inequalities of the form

$$x_{i_1} + x_{i_2} + \cdots + x_{i_k} \geq 1, \;\; i_1 \in \{1, 2\}, \;\; i_2 \in \{3, 4\}, \ldots , i_k \in \{2k - 1, 2k\},$$

where $x = (x_1, \ldots , x_n) \in \mathcal{C} = \{x \in \mathbb{Z}^n \mid 0 \leq x \leq c\}$. For any positive integral vector c, this system has 2^k maximal infeasible integral vectors and only k minimal feasible integral vectors, i.e.,

$$|\mathcal{I}(\mathcal{F}_{A,b,c})| = \frac{rn}{2(\log r)^2}|\mathcal{F}_{A,b,c}|.$$

Needless to say that in general, $|\mathcal{F}_{A,b,c}|$ cannot be bounded by a polynomial in r, n, and $|\mathcal{I}(\mathcal{F}_{A,b,c})|$. For instance, for $n = 2k$ the system of k inequalities $x_1 + x_2 \geq 1, \quad x_3 + x_4 \geq 1, \ldots, x_{2k-1} + x_{2k} \geq 1$ has 2^k minimal feasible binary vectors and only k maximal infeasible binary vectors.

We prove Theorem 2 in Section 2, and then use this theorem in the next section to reduce problem $GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$ to a natural generalization of the hypergraph dualization problem. Our generalized dualization problem replaces hypergraphs by collections of integer vectors in a box.

Theorem 3 *$GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$ is polynomial-time reducible to the following problem:*

$DUAL(\mathcal{C}, \mathcal{A}, \mathcal{B})$**:** *Given an integral box $\mathcal{C}$, a family of vectors $\mathcal{A} \subseteq \mathcal{C}$, and a collection of maximal independent elements $\mathcal{B} \subseteq \mathcal{I}(\mathcal{A})$, either find a new maximal independent element $x \in \mathcal{I}(\mathcal{A}) \setminus \mathcal{B}$, or prove that $\mathcal{B} = \mathcal{I}(\mathcal{A})$.*

Note that for $\mathcal{C} = \{0, 1\}^n$, problem $DUAL(\mathcal{C}, \mathcal{A}, \mathcal{B})$ turns into the hypergraph dualization problem. Other applications of the dualization problem on boxes can be found in [1, 7, 13]. In Section 4 we extend the hypergraph dualization algorithms of [8] to problem $DUAL(\mathcal{C}, \mathcal{A}, \mathcal{B})$ and show that the latter problem can be solved in quasi-polynomial time:

Theorem 4 *Given two sets $\mathcal{A}$, and $\mathcal{B} \subseteq \mathcal{I}(\mathcal{A})$ in an integral box $\mathcal{C} = \{x \in \mathbb{Z}^n \mid 0 \leq x \leq c\}$, problem $DUAL(\mathcal{C}, \mathcal{A}, \mathcal{B})$ can be solved in $poly(n, m) + m^{o(\log m)}$ time, where $m = |\mathcal{A}| + |\mathcal{B}|$.*

Clearly, Theorem 1 follows from Theorems 3 and 4. The special cases of Theorems 2 and 3 for Boolean systems $c = (1, \ldots, 1)$ can be found in [4]. If c is bounded and the number of non-zero coefficients per inequality in (1) is fixed, the results of [3] also imply that problem $GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$ can be efficiently solved in parallel.

Let us add finally that even though by Proposition 1 generating all maximal infeasible vectors for (1) is NP-hard, this problem can be solved efficiently if the number of inequalities in (1) is fixed. Specifically, for $r = const$ the size of $\mathcal{F}_{A,b,c}$ can be bounded by a polynomial in n and $|\mathcal{I}(\mathcal{F}_{A,b,c})|$ and consequently, all elements of $\mathcal{I}(\mathcal{F}_{A,b,c})$ can be generated in quasi-polynomial time. In fact, for $r = const$ the problem of generating $\mathcal{I}(\mathcal{F}_{A,b,c})$ as well as the problem of generating $\mathcal{F}_{A,b,c}$ can be solved separately in incremental polynomial time.

Theorem 5 *Suppose that the monotone system (1) is nontrivial, i.e.,* $0 \notin \mathcal{F}_{A,b,c}$*. Then for any non-empty subset* $\mathcal{Y} \subseteq \mathcal{I}(\mathcal{F}_{A,b,c})$ *we have*

$$\left|\mathcal{I}^{-1}(\mathcal{Y}) \cap \mathcal{F}_{A,b,c}\right| \leq \Big(\sum_{y \in \mathcal{Y}} q(y)\Big)^r, \tag{5}$$

where $\mathcal{I}^{-1}(\mathcal{Y})$ *is the set of all minimal integral vectors of the ideal* $\mathcal{C} \setminus \mathcal{Y}^-$ *and* $q(y)$ *is the number of components* y_l *such that* $y_l < c_l$*. In particular,* $|\mathcal{I}^{-1}(\mathcal{Y}) \cap \mathcal{F}_{A,b,c}| \leq \Big(n|\mathcal{Y}|\Big)^r$*, which for* $\mathcal{Y} = \mathcal{I}(\mathcal{F}_{A,b,c})$ *implies* $|\mathcal{F}_{A,b,c}| \leq \Big(n|\mathcal{I}(\mathcal{F}_{A,b,c})|\Big)^r$*.*

Theorem 6 *If the number of rows in (1) is bounded, then problems* $GEN(\mathcal{F}_{A,b,c}, \mathcal{X})$ *and* $GEN(\mathcal{I}(\mathcal{F}_{A,b,c}), \mathcal{Y})$ *can be solved in polynomial time.*

The remainder of the paper consists of the proofs of Theorems 2, 3, 4, 5, and 6 in Sections 2, 3, 4, 5 and 6, respectively.

2 Bounding the number of maximal infeasible vectors

In this section we prove Theorem 2. We first need some notations and definitions.

Let $\mathcal{C} = \{x \in \mathbb{Z}^n \mid 0 \leq x \leq c\}$ be a box and let $f : \mathcal{C} \to \{0, 1\}$ be a discrete binary function. The function f is called *monotone* if $f(x) \geq f(y)$ whenever $x \geq y$ and $x, y \in \mathcal{C}$. We denote by $T(f)$ and $F(f)$ the sets of all true and all false vectors of f, i.e.,

$$T(f) = \{x \in \mathcal{C} | f(x) = 1\} = (\min[f])^+, \qquad F(f) = \{x \in \mathcal{C} | f(x) = 0\} = (\max[f])^-,$$

where $\min[f]$ and $max[f]$ are the sets of all minimal true and all maximal false vectors of f, respectively.

Let $\sigma \in \mathbb{S}_n$ be a permutation of the coordinates and let x, y be two n-vectors. We say that y is a *left-shift* of x and write $y \succeq_\sigma x$ if the inequalities

$$\sum_{j=1}^{k} y_{\sigma_j} \geq \sum_{j=1}^{k} x_{\sigma_j}$$

hold for all $k = 1, \ldots, n$. A discrete binary function $f : \mathcal{C} \to \{0,1\}$ is called *2-monotonic with respect to* σ if $f(y) \geq f(x)$ whenever $y \succeq_\sigma x$ and $x, y \in \mathcal{C}$. Clearly, $y \geq x$ implies $y \succeq_\sigma x$ for any $\sigma \in \mathbb{S}_n$, so that any 2-monotonic function is monotone.

The function f will be called *regular* if it is 2-monotonic with respect to the identity permutation $\sigma = (1, 2, ..., n)$. Any 2-monotonic function can be transformed into a regular one by appropriately re-indexing its variables. To simplify notations, we shall state Lemma 1 below for regular functions, i.e., we fix $\sigma = (1, 2, ..., n)$ in this lemma.

For a given subset $\mathcal{A} \subseteq \mathcal{C}$ let us denote by $\mathcal{A}^*$ all the vectors which are left-shifts of some vectors of $\mathcal{A}$, i.e., $\mathcal{A}^* = \{y \in \mathcal{C} \mid y \succeq x \text{ for some } x \in \mathcal{A}\}$. Clearly, $T(f) = (\min[f])^*$ for a regular function f (in fact, the subfamily of *right-most* vectors of $\min[f]$ would be enough to use here.)

Given monotone discrete functions f and g, we call g a *regular majorant* of f, if $g(x) \geq f(x)$ for all $x \in \mathcal{C}$, and g is regular. Clearly, $T(g) \supseteq (\min[f])^*$ must hold in this case, and the discrete function h defined by $T(h) = (\min[f])^*$ is the unique minimal regular majorant of f.

For a vector $x \in \mathcal{C}$, and for an index $1 \leq k \leq n$, let the vectors $x^{(k]}$ and $x^{[k)}$ be defined by

$$x_j^{(k]} = \begin{cases} x_j & \text{for } j \leq k, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$x_j^{[k)} = \begin{cases} x_j & \text{for } j \geq k, \\ 0 & \text{otherwise.} \end{cases}$$

Let us denote by e the n-vector of all 1's, let e_j denote the j^{th} unit vector, $j = 1, ..., n$, and let $p(x)$ denote the number of positive components of the vector $x \in \mathcal{C}$.

Lemma 1 *Given a monotone discrete binary function $f : \mathcal{C} \to \{0,1\}$ such that $f \not\equiv 0$, and a regular majorant $g \geq f$, we have the inequality*

$$|F(g) \cap \max[f]| \leq \sum_{x \in \min[f]} p(x). \tag{6}$$

Proof. Let us denote by h the unique minimal regular majorant of f. Then we have $F(g) \cap \max[f] \subseteq F(h) \cap \max[f]$, and hence it is enough to show the statement for $g = h$, i.e. when $T(g) = (\min[f])^*$.

For a vector $y \in \mathcal{C} \setminus \{c\}$ let us denote by $l = l_y$ the index of the last component which is less than c_l, i.e., $l = \max\{j \mid y_j < c_j\} \in \{1, \ldots, n\}$. We claim that for every $y \in F(h) \cap \max[f]$ there exists an $x \in \min[f]$ such that

$$y = x^{(l-1]} + (x_l - 1)e_l + c^{[l+1)}, \tag{7}$$

where $l = l_y$. To see this claim, first observe that $y \neq c$ because $y \in F(f)$ and $f \not\equiv 0$. Second, for any j with $y_j < c_j$ we have $y + e_j \in T(f)$, by the definition of a maximal false point. Hence there exists a minimal true-vector $x \in \min[f]$ such that $x \leq y + e_l$ for $l = l_y$. We must have $x^{(l-1]} = y^{(l-1]}$, since if $x_i < y_i$ for some $i < l$, then $y \geq x + e_i - e_l \succeq x$ would hold, i.e. $y \succeq x$ would follow, implying $y \in (\min[f])^*$ and yielding a contradiction with $y \in F(h) = \mathcal{C} \setminus (\min[f])^*$. Finally, the definition of $l = l_y$ implies that $y^{[l+1)} = c^{[l+1)}$. Hence, our claim and the equality (7) follow.

The above claim implies that

$$F(h) \cap \max[f] \subseteq \{x^{(l-1]} + (x_l - 1)e_l + c^{[l+1)} \mid x \in \min[f],\ x_l > 0\},$$

and hence (6) and thus the lemma follow. $\square$

Lemma 2 *Let $f : \mathcal{C} \to \{0, 1\}$ be a monotone discrete binary function such that $f \not\equiv 0$ and*

$$x \in T(f) \quad \Rightarrow \quad \alpha x \stackrel{\text{def}}{=} \alpha_1 x_1 + \ldots \alpha_n x_n \geq \beta, \tag{8}$$

where $\alpha = (\alpha_1, \ldots, \alpha_n)$ is a given real vector and β is a real threshold. Then

$$|\{x \in \mathcal{C} \mid \alpha x < \beta\} \cap \max[f]| \leq \sum_{x \in \min[f]} p(x).$$

Proof. Suppose that some of the weights $\alpha_1, \ldots, \alpha_n$ are negative, say $\alpha_1 < 0, \ldots, \alpha_k < 0$ and $\alpha^{[k+1)} \geq 0$. Since $\alpha x \geq \beta$ for any $x \in T(f)$ and since f is monotone, we have $x \in T(f) \Rightarrow \alpha^{[k+1)} x \geq \beta - \alpha^{(k]} c^{(k]}$. For any $x \in \mathcal{C}$ we also have $\{x \mid \alpha x < \beta\} \subseteq \{x \mid \alpha^{[k+1)} x < \beta - \alpha^{(k]} c^{(k]}\}$. Hence it suffices to prove the lemma for the non-negative weight vector $\alpha^{[k+1)}$ and the threshold $\beta - \alpha^{(k]} c^{(k]}$. In other words, we can assume without loss of generality that the original weight vector α is non-negative.

Let $\sigma \in \mathbb{S}^n$ be a permutation such that $\alpha_{\sigma_1} \geq \alpha_{\sigma_2} \geq \cdots \geq \alpha_{\sigma_n} \geq 0$. Then the threshold function

$$g(x) = \begin{cases} 1 & \text{if } \alpha x \geq \beta, \\ 0 & \text{otherwise.} \end{cases}$$

is 2-monotonic with respect to σ. By (8), we have $g \geq f$ for all $x \in \mathcal{C}$, i.e., g is a majorant of f. In addition, $F(g) = \{x \in \mathcal{C} \mid \alpha x < \beta\}$, and hence Lemma 2 follows from Lemma 1. $\square$

We are now ready to show inequality (4) and finish the proof of Theorem 2. Given a non-empty set $\mathcal{X} \subseteq \mathcal{F}_{A,b,c}$, consider the monotone discrete function $f : \mathcal{C} \to \{0, 1\}$ defined by the condition $\min[f] = \mathcal{X}$. Since (1) is monotone, any true vector of f also satisfies (1):

$$x \in T(f) \quad \Rightarrow a_{k1} x_1 + \ldots + a_{kn} x_n \geq b_k$$

for all $k = 1, \dots, r$. In addition, $f \not\equiv 0$ because $\mathcal{X} \neq \emptyset$. Thus, by Lemma 2 we have the inequalities

$$|\{x \mid a_{k1}x_1 + \ldots + a_{kn}x_n < b_k\} \cap \max[f]| \leq \sum_{x \in \mathcal{X}} p(x) \tag{9}$$

for each $k = 1, ..., r$. Now, from $\max[f] = \mathcal{I}(\mathcal{X})$ we deduce that

$$\mathcal{I}(\mathcal{F}_{A,b,c}) \cap \mathcal{I}(\mathcal{X}) \subseteq \bigcup_{k=1}^{r} \{x \mid a_{k1}x_1 + \ldots + a_{kn}x_n < b_k\} \cap \max[f],$$

and thus (4) and the theorem follows by (9). □

3 Generating minimal integer solutions via integral dualization

The proof of Theorem 3 has two ingredients. First, we show that given a monotone system (1), the sets $\mathcal{I}(\mathcal{F}_{A,b,c})$ and $\mathcal{F}_{A,b,c}$ can be *jointly* enumerated by iteratively solving the dualization problem $DUAL(\mathcal{C}, \mathcal{A}, \mathcal{B})$ introduced in Theorem 3. Second, we invoke Theorem 2 and argue that since the number of maximal infeasible vectors is relatively small, the generation of $\mathcal{F}_{A,b,c}$ polynomially reduces to the joint generation of $\mathcal{I}(\mathcal{F}_{A,b,c})$ and $\mathcal{F}_{A,b,c}$.

3.1 Joint generation of dual subsets in an integral box

Let $\mathcal{F} = \mathcal{F}_{A,b,c}$ be the set of minimal integral vectors for (1), and consider the following problem of jointly generating all points of $\mathcal{F}$ and $\mathcal{I}(\mathcal{F})$:

$GEN(\mathcal{F}, \mathcal{I}(\mathcal{F}), \mathcal{A}, \mathcal{B})$**:** *Given two explicitly listed collections $\mathcal{A} \subseteq \mathcal{F}$ and $\mathcal{B} \subseteq \mathcal{I}(\mathcal{F})$, either find a new point in $(\mathcal{F} \setminus \mathcal{A}) \cup (\mathcal{I}(\mathcal{F}) \setminus \mathcal{B})$, or prove that these collections are complete: $(\mathcal{A}, \mathcal{B}) = (\mathcal{F}, \mathcal{I}(\mathcal{F}))$.*

Proposition 2 *Problem $GEN(\mathcal{F}, \mathcal{I}(\mathcal{F}), \mathcal{A}, \mathcal{B})$ can be solved in time $poly(n, |\mathcal{A}|, |\mathcal{B}|, \log \|c\|_\infty)$ + T_{dual}, where T_{dual} denotes the time required to solve problem $DUAL(\mathcal{C}, \mathcal{A}, \mathcal{B})$.*

Proof. The reduction is via the following Algorithm $\mathcal{J}$:

Step 1. Check whether $\mathcal{B} \subseteq \mathcal{I}(\mathcal{A})$. If there is an $x \in \mathcal{B} \setminus \mathcal{I}(\mathcal{A})$, then $x \notin \mathcal{F}^+$ because $x \in \mathcal{B} \subseteq \mathcal{I}(\mathcal{F})$. This and the inclusion $\mathcal{A} \subseteq \mathcal{F}$ imply that $x \notin \mathcal{A}^+$. Since $x \notin \mathcal{I}(\mathcal{A})$, we can find a coordinate $j \in \{1, \dots, n\}$ for which $y = x + e_j \notin \mathcal{A}^+$. By the maximality of x in $\mathcal{C} \setminus \mathcal{F}^+$, y belongs to $\mathcal{F}^+$ and therefore, there must exist a $z \in \mathcal{F}$ such that $z \leq y$. Since $z \notin \mathcal{A}^+$, we have $z \in \mathcal{F} \setminus \mathcal{A}$, i.e., z is a *new* minimal integral vector in $\mathcal{F}$ which can be found in $poly(n, |\mathcal{A}|, |\mathcal{B}|, \log \|c\|_\infty)$ time by performing coordinate binary searches on the box $\{z \in \mathbb{Z}^n \mid 0 \leq z \leq y\}$.

Step 2 is similar to the previous step: we check whether $\mathcal{A} \subseteq \mathcal{I}^{-1}(\mathcal{B})$, where $\mathcal{I}^{-1}(\mathcal{B})$ is the set of integral vectors minimal in $\mathcal{C} \setminus \mathcal{B}^-$. If $\mathcal{A}$ contains an element that is not minimal in $\mathcal{C} \setminus \mathcal{B}^-$, we can find a new point in $\mathcal{I}(\mathcal{F}) \setminus \mathcal{B}$ and halt.

Step 3. Suppose that $\mathcal{B} \subseteq \mathcal{I}(\mathcal{A})$ and $\mathcal{A} \subseteq \mathcal{I}^{-1}(\mathcal{B})$. Then $(\mathcal{A}, \mathcal{B}) = (\mathcal{F}, \mathcal{I}(\mathcal{F})) \Leftrightarrow \mathcal{B} = \mathcal{I}(\mathcal{A})$. (To see this, assume that $\mathcal{B} = \mathcal{I}(\mathcal{A})$, and suppose on the contrary that there is an $x \in \mathcal{F} \setminus \mathcal{A}$. Since $x \notin \mathcal{A} = \mathcal{I}^{-1}(\mathcal{B})$ and $x \notin \mathcal{B}^- \subseteq \mathcal{I}(\mathcal{F})^-$, there must exist a $y \in \mathcal{I}^{-1}(\mathcal{B}) = \mathcal{A} \subseteq \mathcal{F}$ such that $y \leq x$. Hence we get two distinct elements $x, y \in \mathcal{F}$ such that $y \leq x$, which contradicts the definition of $\mathcal{F}$. The existence of an $x \in \mathcal{I}(\mathcal{F}) \setminus \mathcal{B}$ leads to a similar contradiction.) To check the condition $\mathcal{B} = \mathcal{I}(\mathcal{A})$, we solve problem $DUAL(\mathcal{C}, \mathcal{A}, \mathcal{B})$. If $\mathcal{B} \neq \mathcal{I}(\mathcal{A})$, we obtain a new point $x \in \mathcal{I}(\mathcal{A}) \setminus \mathcal{B}$. By (3), either $x \in \mathcal{F}^+$, or $x \in \mathcal{I}(\mathcal{F})^-$ and we can decide which of these two cases holds by checking the feasibility of x for (1). In the first case, we obtain a new point $y \in \{x\}^- \cap (\mathcal{F} \setminus \mathcal{A})$ by performing binary searches on the coordinates of the box $\{y \in \mathbb{Z}^n \mid 0 \leq y \leq x\}$. In the second case, a new point in $\{x\}^+ \cap (\mathcal{I}(\mathcal{F}) \setminus \mathcal{B})$ can be obtained by searching the box $\{y \in \mathbb{Z}^n \mid x \leq y \leq c\}$. $\square$

Let $\mathcal{F} \subseteq \mathcal{C}$ be an arbitrary antichain, i.e., a system of integral vectors such that $x \not\leq y$ for any two distinct elements $x, y \in \mathcal{F}$. It is easy to see that Algorithm $\mathcal{J}$ and Proposition 2 can be used for any class of antichains $\mathcal{F}$ defined by a polynomial-time membership oracle for $\mathcal{F}^+$.

3.2 Uniformly dual-bounded antichains

Extending the definition of dual-bounded hypergraphs in [4], we say that (a class of antichains) $\mathcal{F} \subseteq \mathcal{C}$ is *uniformly dual-bounded* if there exists a polynomial p such that

$$|\mathcal{I}(\mathcal{F}) \cap \mathcal{I}(\mathcal{X})| \leq p(|\mathcal{X}|),$$

for any nonempty subset $\mathcal{X} \subseteq \mathcal{F}$.

Proposition 3 *Suppose that $\mathcal{F}$ is uniformly dual-bounded and there exists a polynomial-time membership oracle for $\mathcal{F}^+$. Then problem $GEN(\mathcal{F})$ is polynomial-time reducible to problem $DUAL(\mathcal{C}, \mathcal{A}, \mathcal{B})$.*

Proof. Given a set $\mathcal{X}$ in $\mathcal{F}$, we repeatedly run Algorithm $\mathcal{J}$ until it either produces a new element in $\mathcal{F} \setminus \mathcal{X}$ or proves that $\mathcal{X} = \mathcal{F}$ by generating the entire family $\mathcal{I}(\mathcal{F})$. By Step 1, as long as Algorithm $\mathcal{J}$ outputs elements of $\mathcal{I}(\mathcal{F})$, these elements also belong to $\mathcal{I}(\mathcal{X})$, and hence the total number of such elements does not exceed $p(|\mathcal{X}|)$. $\square$

By Theorem 2, the set of minimal integral solutions to any monotone system of linear inequalities is uniformly-dual bounded, and hence Theorem 3 is a corollary of Proposition 3.

4 Dualization in products of chains

Let $\mathcal{C} \stackrel{\text{def}}{=} \mathcal{C}_1 \times \ldots \times \mathcal{C}_n$ be an integer box defined by the product of n chains $\mathcal{C}_i = [l_i : u_i]$ where $l_i, u_i \in \mathbb{Z}$ are, respectively, the lower and upper bounds of chain $\mathcal{C}_i$. Given an antichain $\mathcal{A} \subseteq \mathcal{C}$, and an antichain $\mathcal{B} \subseteq \mathcal{I}(\mathcal{A})$, we say that $\mathcal{B}$ is *dual to* $\mathcal{A}$ if $\mathcal{B} = \mathcal{I}(\mathcal{A})$, i.e., $\mathcal{B}$ contains all the maximal elements of $\mathcal{C} \setminus \mathcal{A}^+$. If $\mathcal{C}$ is the unit cube, we obtain the familiar notion of dual hypergraphs, where $\mathcal{I}(\mathcal{A})$ becomes the complementary set of the transversal hypergraph of $\mathcal{A}$. In the following two subsections, we will show how to extend the hypergraph dualization algorithms of [8] to arbitrary systems $\mathcal{A}$ of integral vectors in a box $\mathcal{C}$.

As in [8], we shall analyze the running time of the algorithms in terms of the "volume" $v = v(\mathcal{A}, \mathcal{B}) \stackrel{\text{def}}{=} |\mathcal{A}||\mathcal{B}|$ of the input problem. In general, a given problem will be decomposed into a number of subproblems which we solve recursively. Since we have assumed that $\mathcal{B} \subseteq \mathcal{I}(\mathcal{A})$, (3) implies that the following condition holds for the original problem and all subsequent subproblems:

$$a \not\leq b, \quad \text{for all } a \in \mathcal{A}, b \in \mathcal{B}. \tag{10}$$

Let $C(v) = C(v(\mathcal{A}, \mathcal{B}))$ denote the number of subproblems that have to be solved in order to solve the original problem, and let $m \stackrel{\text{def}}{=} |\mathcal{A}| + |\mathcal{B}|$, and $[n] \stackrel{\text{def}}{=} \{1, \ldots, n\}$. We start with the following proposition that provides the base case for recursion.

Proposition 4 *Suppose* $\min\{|\mathcal{A}|, |\mathcal{B}|\} \leq$ *const, then problem DUAL*$(\mathcal{C}, \mathcal{A}, \mathcal{B})$ *is solvable in polynomial time.*

Proof. Let us assume without loss of generality that $\mathcal{B} = \{b^1, \ldots, b^k\}$, for some constant k. For every $t \in [n]^k$ and $i \in [n]$, let $I_i^t = \{j \in [k] \mid t_j = i\}$. Then $\mathcal{C} = \mathcal{A}^+ \cup \mathcal{B}^-$ *if and only if* for every $t \in [n]^k$ for which

$$b_i^j \neq u_i, \text{ for all } i \in [n], j \in I_i^t, \tag{11}$$

there exists an $a \in \mathcal{A}$ such that

$$a_i \leq \max\{b_i^j + 1 \mid j \in I_i^t\}, \text{ if } I_i^t \neq \emptyset, \text{ and } a_i = l_i \text{ otherwise.} \tag{12}$$

To see this, assume first that $\mathcal{C} = \mathcal{A}^+ \cup \mathcal{B}^-$ and consider any $t \in [n]^k$ such that (11) holds. Let $x \in \mathcal{C}$ be defined by taking $x_i = \max\{b_i^j + 1 \mid j \in I_i^t\}$ if $I_i^t \neq \emptyset$, and $x_i = l_i$ otherwise. Then $x \in \mathcal{C} \setminus \mathcal{B}^-$ and hence $x \in \mathcal{A}^+$, implying that there is an $a \in \mathcal{A}$ satisfying (12). On the other hand, let us assume that for every $t \in [n]^k$ satisfying (11), there is an $a \in \mathcal{A}$ for which (12) holds. Consider an $x \in \mathcal{C} \setminus \mathcal{B}^-$, then there must exist, for every $j \in [k]$, a $t_j \in [n]$, such that $x_{t_j} \geq b_{t_j}^j + 1$. Clearly $t = (t_1, \ldots, t_k) \in [n]^k$ satisfies (11), and therefore, there is an $a \in \mathcal{A}$ such that $a_i \leq \max\{b_i^j + 1 \mid j \in I_i^t\} \leq x_i$ if $I_i^t \neq \emptyset$, and $a_i = l_i$ otherwise. This gives $x \in \mathcal{A}^+$. $\square$

Remark 1 Having found an $x \in \mathcal{C} \setminus (\mathcal{A}^+ \cup \mathcal{B}^-)$, it is always possible to extend it to a maximal point with the same property in $O(nm \log m)$ time as follows. Let $\mathcal{Q}_i = \{a_i - 1 \mid a \in$

$\mathcal{A}\} \cup \{x_i, u_i\}$, $i = 1, \ldots, n$, and assume that this list is kept in sorted order for each i. For $i = 1, \ldots, n$, we iterate $x_i \leftarrow \max\{z \in \mathcal{Q}_i \mid (x_1, \ldots, x_{i-1}, z, x_{i+1}, \ldots, x_n\} \notin \mathcal{A}^+$. Then the resulting point x is maximal in $\mathcal{C} \setminus (\mathcal{A}^+ \cup \mathcal{B}^-)$.

4.1 Algorithm A

Assume that $\mathcal{A}, \mathcal{B}$ satisfy (10), and let $\mathcal{C} = \mathcal{C}_1 \times \ldots \times \mathcal{C}_n$, where $\mathcal{C}_i = [l_i : u_i]$. We say that a coordinate $i \in [n]$ is *essential* for a point $a \in \mathcal{A}$ ($b \in \mathcal{B}$), if $a_i > l_i$ (respectively, $b_i < u_i$). Let us denote by $E(x)$ the set of essential coordinates of a point $x \in \mathcal{A} \cup \mathcal{B}$. The following lemma generalizes a well-known fact for dual Boolean functions (cf. [8]).

Lemma 3 *If $\mathcal{A}, \mathcal{B}$ are dual in $\mathcal{C}$, then there exists a point $x \in \mathcal{A} \cup \mathcal{B}$ with few essential coordinates: $|E(x)| \leq \log m$, where $m = |\mathcal{A}| + |\mathcal{B}|$.*

Proof. Let $z \in \mathcal{C}$ be the vector obtained by picking each coordinate z_i randomly from $\{l_i, u_i\}$, $i = 1, \ldots, n$, and consider the random variable $N(z) \stackrel{\text{def}}{=} |\{a \in \mathcal{A} \mid z \geq a\}| + |\{b \in \mathcal{B} \mid z \leq b\}|$. Then the expected value of $N(z)$ is given by

$$\begin{aligned} \mathbb{E}[N(z)] &= \textstyle\sum_{a \in \mathcal{A}} \Pr\{z \geq a\} + \sum_{b \in \mathcal{B}} \Pr\{z \leq b\} \\ &= \textstyle\sum_{a \in \mathcal{A}} \prod_{i \in E(a)} \Pr\{z_i = u_i\} + \sum_{b \in \mathcal{B}} \prod_{i \in E(b)} \Pr\{z_i = l_i\} \\ &= \textstyle\sum_{a \in \mathcal{A}} \left(\frac{1}{2}\right)^{|E(a)|} + \sum_{b \in \mathcal{B}} \left(\frac{1}{2}\right)^{|E(b)|}. \end{aligned} \tag{13}$$

Clearly if $\mathbb{E}[N(z)] < 1$, we can find an $x \in \mathcal{C} \setminus (\mathcal{A}^+ \cup \mathcal{B}^-)$ in polynomial time as follows: for $i = 1, \ldots, n$, use (13) to compute the expectations of $N(x_1, \ldots, x_{i-1}, l_i, z_{i+1}, \ldots, z_n)$ and $N(x_1, \ldots, x_{i-1}, u_i, z_{i+1}, \ldots, z_n)$ over $(z_{i+1}, \ldots, z_n) \in \{l_i, u_i\}^{n-i}$, and select the value of $x_i \in \{l_i, u_i\}$ so as to minimize the corresponding expectation.

Let us therefore assume that $\mathbb{E}[N(z)] \geq 1$, and let $r = \min\{|E(z)| \; : \; z \in \mathcal{A} \cup \mathcal{B}\}$, then (13) implies that

$$1 \leq \mathbb{E}[N(z)] \leq (|\mathcal{A}| + |\mathcal{B}|) \left(\frac{1}{2}\right)^r = m \left(\frac{1}{2}\right)^r.$$

The lemma follows. $\square$

Next we show that an essential coordinate with high frequency exists for either $\mathcal{A}$ or $\mathcal{B}$.

Lemma 4 *Let $\mathcal{A}, \mathcal{B}$ be a pair of dual subsets of $\mathcal{C}$ with $|\mathcal{A}||\mathcal{B}| \geq 1$. Then there exists a coordinate $i \in [n]$, which is essential in $\mathcal{A}$ or in $\mathcal{B}$ with frequency at least $1/\log m$, where $m = |\mathcal{A}| + |\mathcal{B}|$.*

Proof. By Lemma 3, $\mathcal{A} \cup \mathcal{B}$ contains an element x with a logarithmically small number of essential coordinates. Suppose without loss of generality that $x \in \mathcal{A}$. From the duality condition (10), we know that for every $b \in \mathcal{B}$, there is an $i \in E(b) \cap E(x)$ such that $b_i < x_i$. Letting $\mathcal{B}_i^x \stackrel{\text{def}}{=} \{b \in \mathcal{B} \mid b_i < x_i\}$ for $i \in E(x)$, we conclude that

$$|\mathcal{B}| = |\bigcup_{i \in E(x)} \mathcal{B}_i^x| \leq \sum_{i \in E(x)} |\mathcal{B}_i^x|,$$

and therefore there is an $i \in [n]$ which is essential for at least $|\mathcal{B}|/|E(x)| \geq |\mathcal{B}|/\log m$ elements of $\mathcal{B}$. $\square$

We are now ready to state the first dualization algorithm.

Step 1. If $|\mathcal{A}||\mathcal{B}| \leq 1$, then the dualization problem can be solved in $O(n^2 m)$ using Proposition 4.

Step 2. For each $k \in [n]$: if $a_k < l_k$ for some $a \in \mathcal{A}$ ($b_k > u_k$ for some $b \in \mathcal{B}$), set $a_k \leftarrow l_k$ (respectively, set $b_k \leftarrow u_k$). Note that the duality condition (10) continues to hold after such replacements. Thus in $O(nm)$ time, we can assume that $\mathcal{A}, \mathcal{B} \subseteq \mathcal{C}$.

Step 3. Check if there is an $x \in \mathcal{A} \cup \mathcal{B}$ with at most $\log m$ essential coordinates. If no such x can be found, a new point in $\mathcal{C} \setminus (\mathcal{A}^+ \cup \mathcal{B}^-)$ can be obtained as described in the proof of Lemma 3. Otherwise, assume without loss of generality, that $x = a^o \in \mathcal{A}$ and find an $i \in E(a^o)$ for which $|\{b \in \mathcal{B} \mid b_i < a^o_i\}| \geq 1/\log m$. Let us suppose, again without any loss of generality, that $i = 1$ and set $\mathcal{C}'_1 \leftarrow [a^o_1 : u_1]$, $\mathcal{C}''_1 \leftarrow [l_1 : a^o_1 - 1]$. Let further

$$\mathcal{A}'' = \{a \in \mathcal{A} \mid a_1 < a^o_1\}, \qquad \mathcal{B}' = \{b \in \mathcal{B} \mid b_1 \geq a^o_1\}.$$

Observe that $|\mathcal{A}''| \leq |\mathcal{A}| - 1$ since $a^o \notin \mathcal{A}''$, and that $|\mathcal{B}'| \leq (1 - 1/\log m)|\mathcal{B}|$ by our choice of a^o and i.

Step 4. Denoting by $\mathcal{C}' = \mathcal{C}'_1 \times \mathcal{C}_2 \times \ldots \times \mathcal{C}_n$, and $\mathcal{C}'' = \mathcal{C}''_1 \times \mathcal{C}_2 \times \ldots \times \mathcal{C}_n$ the two half-boxes of $\mathcal{C}$ induced by the above partitioning, it is then easy to see that $\mathcal{A}$ and $\mathcal{B}$ are dual in $\mathcal{C}$ *if and only if*

$$\mathcal{A}, \mathcal{B}' \text{ are dual in } \mathcal{C}' \tag{14}$$

and

$$\mathcal{A}'', \mathcal{B} \text{ are dual in } \mathcal{C}''. \tag{15}$$

Thus by applying the algorithm recursively to these two subproblems, we reduce the computation on a problem of size $v = |\mathcal{A}||\mathcal{B}|$ to computing the solution for two subproblems (14)-(15) of volumes

$$\begin{aligned} v(\mathcal{A}, \mathcal{B}') &= |\mathcal{A}||\mathcal{B}'| \leq |\mathcal{A}|(1-\epsilon)|\mathcal{B}| = (1-\epsilon)v, \text{ and} \\ v(\mathcal{A}'', \mathcal{B}) &= |\mathcal{A}''||\mathcal{B}| = (|\mathcal{A}| - 1)|\mathcal{B}| \leq v - 1, \end{aligned}$$

where $\epsilon = 1/\log m$. This leads to the recurrence

$$C(v) \leq 1 + C((1-\epsilon)v) + C(v-1),$$

which was shown in [8] to evaluate to $C(v) \leq (3 + 2v\epsilon)^{\log v/\epsilon}$, implying that the running time of the algorithm is $O(m^{4\log^2 m})$.

4.2 Algorithm B

Given an integral box $\mathcal{C} = \mathcal{C}_1 \times \ldots \times \mathcal{C}_n$, where $\mathcal{C}_i = [l_i : u_i]$, and subsets of integral vectors $\mathcal{A}, \mathcal{B}$ that satisfy the necessary duality condition (10), we proceed as follows:

Step 1. If $\min\{|\mathcal{A}|, |\mathcal{B}|\} \leq 2$, the duality of $\mathcal{A}$ and $\mathcal{B}$ can be tested in $O(n^3 m)$ using Proposition 4.
Step 2. For each $k \in [n]$:
1. if $a_k > u_k$ for some $a \in \mathcal{A}$ ($b_k < l_k$ for some $b \in \mathcal{B}$), then a (respectively, b) can be clearly discarded from further consideration;
2. if $a_k < l_k$ for some $a \in \mathcal{A}$ ($b_k > u_k$ for some $b \in \mathcal{B}$), we set $a_k \leftarrow l_k$ (respectively, $b_k \leftarrow u_k$).

Thus we may assume for next steps that $\mathcal{A}, \mathcal{B} \subseteq \mathcal{C}$.
Step 3. Let $a^o \in \mathcal{A}$, $b^o \in \mathcal{B}$. By (10), there exists an $i \in [n]$, such that $a_i^o > b_i^o$. Let us assume, without any loss of generality, that $i = 1$ and set $\mathcal{C}_1' \leftarrow [a_1^o : u_1]$, $\mathcal{C}_1'' \leftarrow [l_1 : a_1^o - 1]$. (Alternatively, we may set $\mathcal{C}_1'' \leftarrow [l_1 : b_1^o]$ and $\mathcal{C}_1' \leftarrow [b_1^o + 1 : u_1]$.) Define

$$\begin{array}{rclrclrcl} \mathcal{A}'' & = & \{a \in \mathcal{A} \mid a_1 < a_1^o\}, & \mathcal{A}' & = & \mathcal{A} \setminus \mathcal{A}'', & \epsilon_1^{\mathcal{A}} & = & \frac{|\mathcal{A}'|}{|\mathcal{A}|}, \\ \mathcal{B}' & = & \{b \in \mathcal{B} \mid b_1 \geq a_1^o\}, & \mathcal{B}'' & = & \mathcal{B} \setminus \mathcal{B}', & \epsilon_1^{\mathcal{B}} & = & \frac{|\mathcal{B}''|}{|\mathcal{B}|}. \end{array}$$

Observe that $\epsilon_1^{\mathcal{A}} > 0$ and $\epsilon_1^{\mathcal{B}} > 0$ since $a^o \in \mathcal{A}'$ and $b^o \in \mathcal{B}''$.
Step 4. Define

$$\epsilon(v) = 1/\chi(v), \qquad \textit{where} \;\; \chi^\chi = v = v(\mathcal{A}, \mathcal{B}).$$

If $\min\{\epsilon_1^{\mathcal{A}}, \epsilon_1^{\mathcal{B}}\} > \epsilon(v)$, we use the decomposition rule given above, which amounts to solving recursively two subproblems (14), (15) of respective volumes:

$$\begin{array}{rclclclcl} v(\mathcal{A}, \mathcal{B}') & = & |\mathcal{A}||\mathcal{B}'| & = & |\mathcal{A}|(1-\epsilon_1^{\mathcal{B}})|\mathcal{B}| & = & (1-\epsilon_1^{\mathcal{B}})v(\mathcal{A}, \mathcal{B}) & \leq & (1-\epsilon(v))v, \text{ and} \\ v(\mathcal{A}'', \mathcal{B}) & = & |\mathcal{A}''||\mathcal{B}| & = & (1-\epsilon_1^{\mathcal{A}})|\mathcal{A}||\mathcal{B}| & = & (1-\epsilon_1^{\mathcal{A}})v(\mathcal{A}, \mathcal{B}) & \leq & (1-\epsilon(v))v. \end{array}$$

This gives rise to the recurrence

$$C(v) \leq 1 + C((1-\epsilon_1^{\mathcal{B}})v) + C((1-\epsilon_1^{\mathcal{A}})v) \leq 1 + 2C((1-\epsilon(v))v).$$

Step 5. Let us now suppose that $\epsilon_1^{\mathcal{B}} \leq \epsilon(v)$. In this case, we begin by solving subproblem (14). If $\mathcal{A}, \mathcal{B}'$ are not dual in $\mathcal{C}'$, we get a point x maximal in $\mathcal{C}' \setminus [\mathcal{A}^+ \cup (\mathcal{B}')^-]$, and we are done. Otherwise we claim that

$$\mathcal{A}'', \mathcal{B} \text{ are dual in } \mathcal{C}'' \iff \forall a \in \widetilde{\mathcal{A}} : \mathcal{A}'', \mathcal{B}'' \text{ are dual in } \mathcal{C}''(a), \tag{16}$$

where $\widetilde{\mathcal{A}} = \{a \in \mathcal{A} \mid a_1 \leq a_1^o\}$, and $\mathcal{C}''(a) = \mathcal{C}_1'' \times [a_2 : u_2] \times \ldots \times [a_n : u_n]$.

Proof of (16): The forward direction does not use (14). Suppose that there is an $x \in \mathcal{C}''(a) \setminus [(\mathcal{A}'')^+ \cup (\mathcal{B}'')^-]$ for some $a \in \widetilde{\mathcal{A}}$, then $x_i \geq a_i$, for $i = 2, \ldots, n$. If $x \in (\mathcal{B}')^-$, i.e., $x \leq b$ for some $b \in \mathcal{B}'$, then by the definition of $\mathcal{B}'$, $b_1 \geq a_1^o$. On the other hand, $a \in \widetilde{\mathcal{A}}$ implies that $a_1 \leq a_1^o$. But then,

$$(a_1, a_2, \ldots, a_n) \leq (a_1^o, x_2, \ldots, x_n) \leq (b_1, b_2, \ldots, b_n),$$

which contradicts the assumed duality condition (10). This shows that $x \in \mathcal{C}'' \setminus [(\mathcal{A}'')^+ \cup (\mathcal{B}' \cup \mathcal{B}'')^-]$.

For the other direction, let $x \in \mathcal{C}'' \setminus [(\mathcal{A}'')^+ \cup \mathcal{B}^-]$. Since $x \notin (\mathcal{B}')^-$ and $x = (x_1, x_2, \ldots, x_n) < y \stackrel{\text{def}}{=} (a_1^o, x_2, \ldots, x_n)$, the vector y is also not covered by $\mathcal{B}'$. Thus $y \in \mathcal{C}' \setminus (\mathcal{B}')^-$. We conclude therefore, assuming (14), that $y \in \mathcal{A}^+$, i.e., there is an $a \in \mathcal{A}$ such that $a \leq y$. But this implies that $a \in \widetilde{\mathcal{A}}$ and hence that $x \in \mathcal{C}''(a) \setminus [(\mathcal{A}'')^+ \cup (\mathcal{B}'')^-]$ for some $a \in \widetilde{\mathcal{A}}$. $\square$

It follows by (16) that, once we discover that (14) holds, we can reduce the solution of subproblem (15) to solving $|\widetilde{\mathcal{A}}|$ subproblems, each of which has a volume of $v(|\mathcal{A}''|, |\mathcal{B}''|) \leq \epsilon_1^{\mathcal{B}} v(\mathcal{A}, \mathcal{B})$. Thus we obtain the recurrence

$$C(v) \leq 1 + C((1 - \epsilon_1^{\mathcal{B}})v) + |\mathcal{A}| C(\epsilon_1^{\mathcal{B}} v). \tag{17}$$

Step 6. Finally, if $\epsilon_1^{\mathcal{A}} \leq \epsilon(v) < \epsilon_1^{\mathcal{B}}$, we solve subproblem (15), and if we discover that $\mathcal{A}'', \mathcal{B}$ are dual in $\mathcal{C}''$, we obtain the following decomposition rule, symmetric to (16):

$$\mathcal{A}, \mathcal{B}' \text{ are dual in } \mathcal{C}' \iff \forall b \in \widetilde{\mathcal{B}} : \mathcal{A}', \mathcal{B}' \text{ are dual in } \mathcal{C}'(b),$$

where $\widetilde{\mathcal{B}} = \{b \in \mathcal{B} \mid b_1 \geq a_1^o - 1\}$, and $\mathcal{C}'(b) = \mathcal{C}_1' \times [l_2 : b_2] \times \ldots \times [l_n : b_n]$. This reduces our original problem into one subproblem of volume $\leq (1 - \epsilon_1^{\mathcal{A}})v$, plus $|\widetilde{\mathcal{B}}|$ subproblems, each of volume at most $\epsilon_1^{\mathcal{A}} v$, thus giving the recurrence

$$C(v) \leq 1 + C((1 - \epsilon_1^{\mathcal{A}})v) + |\mathcal{B}| C(\epsilon_1^{\mathcal{A}} v),$$

which is the symmetric version of (17).

Using induction on $v \geq 9$, it can be shown that the above recurrences imply that $R(v) \leq v^{\chi(v)}$ (see [8]). As $\chi(m^2) < 2\chi(m)$ and $v(\mathcal{A}, \mathcal{B}) < m^2$, we get $\chi(v) < \chi(m^2) < 2\chi(m) \sim 2 \log m / \log \log m$. Let us also note that every step above can be implemented in $O(n^3 m)$ time, independent of the chains sizes $|\mathcal{C}_i|$. This establishes the bound stated in Theorem 4.

5 Bounding the number of minimal feasible solutions

To prove inequality (5) of Theorem 5, let us consider an arbitrary non-empty antichain $\mathcal{Y} \subseteq \mathcal{I}(\mathcal{F}_{A,b,c})$. For any $y \in \mathcal{I}(\mathcal{F}_{A,b,c})$ we can find an index $i = \rho(y) \in [r] \stackrel{\text{def}}{=} \{1, \ldots, r\}$ such that y violates the i^{th} inequality of the system, i.e., $a_{(i)}y < b_i$, where $a_{(i)}$ and b_i denote the i^{th} row and component of A and b, respectively.

Consider a vector $x \in \mathcal{I}^{-1}(\mathcal{Y}) \cap \mathcal{F}_{A,b,c}$ and let x_l be a positive component of x. Then there exists a vector $y_{(l)} \in \mathcal{Y}$ such that $y_{(l)} \geq x - e_l$. Let $i = \rho(y_{(l)})$ and assume without loss of generality that

$$(a_{(i)})_1 \;\geq\; (a_{(i)})_2 \;\geq\; \ldots \;\geq\; (a_{(i)})_n. \tag{18}$$

We claim that $x^{(l]} = z^{(l]}_{(l)}$, where

$$z_{(l)} \;\;=\;\; y^{(l]} + e_l. \tag{19}$$

It follows from $y_{(l)} \geq x - e_l$ that $z^{(l]}_{(l)} \geq x^{(l]}$. If $(z_{(l)})_l > x_l$, then $(y_{(l)})_l \geq x_l$, which implies $y_{(l)} \geq x$, a contradiction. Thus $(z_{(l)})_l = x_l$ holds. Moreover, if $(z_{(l)})_j > x_j$ for some $j < l$, then we have $(y_{(l)})_j \geq x_j + 1$. By (18), $a_{(i)}(y_{(l)} - e_j + e_l) < b_i$, i.e., $y_{(l)} - e_j + e_l$ is infeasible for (1). However, $y_{(l)} - e_j + e_l \geq x$ by $y_{(l)} \geq x - e_l$ and hence $y_{(l)} - e_j + e_l$ must be feasible. This shows that $x^{(l]} = z^{(l]}_{(l)}$ and consequently leads to the representation

$$x \;\;=\;\; \bigvee_{l \in [n]:\, x_l > 0} z_{(l)}, \tag{20}$$

where for vectors $v, u \in \mathcal{C}$ we let $v \vee u$ denote the component-wise maximum of v and u.

Not all of the vectors $z_{(l)}$ are necessary for this representation. Suppose that $\rho(y_{(l)}) = \rho(y_{(l')}) = i$ for some positive components x_l and $x_{l'}$ of x, and $l' < l$. Then (20) remains valid if we drop $z_{(l')}$, the vector with the smaller index l'. In other words, by sorting the i^{th} row of A and then selecting among the vectors $y_{(l)} \in \rho^{-1}(i)$ the one with the highest l, we obtain at most r vectors $z_{(i)} = z_{(l_i)}$ such that

$$x \;\;=\;\; \bigvee_{i \in [r]} z_{(i)}. \tag{21}$$

The latter representation readily implies (5). $\square$

6 Polynomial generation of $\mathcal{F}_{A,b,c}$ and $\mathcal{I}(\mathcal{F}_{A,b,c})$

Theorems 5 implies that for $r \leq \mathit{const}$ the antichain $\mathcal{I}(\mathcal{F}_{A,b,c})$ is uniformly dual-bounded and consequently, $\mathcal{I}(\mathcal{F}_{A,b,c})$ can be generated in incremental quasi-polynomial time via Algorithm $\mathcal{J}$ presented in Section 3. In this section we show that for bounded r, the antichains $\mathcal{I}(\mathcal{F}_{A,b,c})$ and $\mathcal{F}_{A,b,c}$ can, in fact, be generated in incremental polynomial time. Since $|\mathcal{F}_{A,b,c}|$ and $|\mathcal{I}(\mathcal{F}_{A,b,c})|$ are (uniformly) polynomially related by Theorems 2 and 5, the required result will follow from Proposition 2, provided that Step 3 of Algorithm $\mathcal{J}$, the dualization step, can be done in polynomial time. Thus, it is enough to show that problem $DUAL(\mathcal{C}, \mathcal{A}, \mathcal{B})$

can be solved in polynomial time, if $\mathcal{A} \subseteq \mathcal{F}_{A,b,c}$ is a subset of the minimal solutions of (1), with bounded r.

For $i \in [r] = \{1, \dots, r\}$, let $\sigma_{(i)} \in \mathbb{S}_n$ be a permutation of the coordinates such that

$$(a_{(i)})_{\sigma_{(i)1}} \;\geq\; (a_{(i)})_{\sigma_{(i)2}} \;\geq\; \dots \;\geq\; (a_{(i)})_{\sigma_{(i)n}}. \tag{22}$$

Given $\mathcal{A} \subseteq \mathcal{F}_{A,b,c}$ and $\mathcal{B} \subseteq \mathcal{I}(\mathcal{F}_{A,b,c})$, we may assume that $0 \notin \mathcal{A}$ (otherwise, $\mathcal{F}_{A,b,c} = \{0\}$ and $\mathcal{I}(\mathcal{F}_{A,b,c}) = \emptyset$) and $\mathcal{B} \neq \emptyset$ (otherwise, Proposition 4 can be used to generate a point $x \in \mathcal{C} \setminus (\mathcal{A}^+ \cup \mathcal{B}^-)$.) Now we proceed in two basic steps:

Step 1. For every $y \in \mathcal{B}$ and for each pair of indices $\sigma_{(i)j}$ and $\sigma_{(i)h}$ with $y_{\sigma_{(i)j}} > 0$, $y_{\sigma_{(i)h}} < c_{\sigma_{(i)h}}$ and $j < h$, check if there exists a $y' \in \mathcal{B}$ such that

$$y' \;\geq\; y - e_{\sigma_{(i)j}} + e_{\sigma_{(i)h}}. \tag{23}$$

Note that $y - e_{\sigma_{(i)j}} + e_{\sigma_{(i)h}}$ is infeasible, and hence is an independent element of $\mathcal{A}$. If no such y' can be found, we generate a new maximal independent vector $y' \in \mathcal{I}(\mathcal{A}) \setminus \mathcal{B}$, satisfying (23), and halt.

Step 2. For every collection $\{y_{(i)} \in \rho_{\mathcal{B}}^{-1}(i) \mid i \in [r]\}$, where $\rho_{\mathcal{B}}^{-1}(i) \stackrel{\text{def}}{=} \{y \in \mathcal{B} \mid \rho(y) = i\}$, and for every $(l_i \mid i \in [r], (y_{(i)})_{l_i} < c_{l_i}) \in [n]^r$, construct the vector $x = \bigvee_{i \in [r]} z_{(i)}$, where $z_{(i)} (= z_{(l_i)})$ is given by (19) (according to the permutation $\sigma_{(i)}$ and using $y_{(i)} = y_{(l_i)}$.) If $x \notin \mathcal{A}^+ \cup \mathcal{B}^-$, then a new maximal independent vector can be generated.

Clearly, the above two steps run in $O((n|\mathcal{B}|)^r)$ time, which is polynomially bounded for constant r. It is also clear that if the algorithm outputs a point $x \in \mathcal{C}$, then $x \notin \mathcal{A}^+ \cup \mathcal{B}^-$, so it remains to verify that the algorithm indeed outputs such a point if $\mathcal{A}^+ \cup \mathcal{B}^- \neq \mathcal{C}$. To see this, let x be a minimal vector in $\mathcal{C} \setminus (\mathcal{A}^+ \cup \mathcal{B}^-)$. From our assumptions, it follows that $x \neq 0$, and thus, there exists and index l with $x_l > 0$. By the minimality of x, there exists a vector $y\,(= y_{(l)}) \in \mathcal{B}$ such that $y \geq x - e_l$. Assume without loss of generality, again, that (18) holds, and consider an index $j < l$. If $(y_{(l)})_j > x_j$, we get $y_{(l)} - e_j + e_l \geq x$, and therefore a new maximal independent point $x' \geq x$ must have been output in Step 1 of the algorithm (c.f. (23)). On the other hand, if for this $y\,(= y_{(l)})$, we have $x^{(l]} = z_{(l)}^{(l]}$, then we can conclude that x satisfies (21), and consequently, it must have been created in Step 2. □

We mention in closing that all the results of the paper remain valid for systems of 2-monotonic inequalities in integer variables. The case of a single Boolean 2-monotonic inequality is discussed in [2, 5, 6, 9, 12, 14, 15].

References

[1] J. C. Bioch, Dualization, decision lists and identification of monotone discrete functions, Annals of Mathematics and Artificial Intelligence 24 (1998) 69-91.

[2] P. Bertolazzi and A. Sassano, An $O(nm)$ algorithm for regular set-covering problems, *Theoretical Computer Science* 54 (1987), 237-247.

[3] E. Boros, K. Elbassioni, V. Gurvich and L. Khachiyan, An incremental RNC algorithm for generating all maximal independent sets in hypergraphs of bounded dimension, DIMACS Technical Report 2000-21, Rutgers University. To appear in *Parallel Processing Letters*.

[4] E. Boros, V. Gurvich, L. Khachiyan and K. Makino, Generating partial and multiple transversals of a hypergraph, In: *Automata, Languages and Programming, 27th International Colloquium, ICALP 2000* (Montanari, J.D.P. Rolim and E. Welzl, eds.), pp. 588-599 (Springer Verlag, Berlin, Heidelberg, New York, July 2000). An extended version of this paper is to appear in *SIAM J. Computing*.

[5] E. Boros, P. L. Hammer, T. Ibaraki and K. Kawakami, Polynomial time recognition of 2-monotonic positive Boolean functions given by an oracle, *SIAM Journal on Computing* 26 (1997) 93-109.

[6] Y. Crama, Dualization of regular Boolean functions, *Discrete Applied Mathematics* 16 (1987) 79-85.

[7] Y. Crama, P. L. Hammer and T. Ibaraki, Cause-effect relationships and partially defined boolean functions, *Annals of Operations Research* 16 (1988) 299-326.

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

[9] P. L. Hammer, U.N. Peled, and M.A. Pollatschek, An algorithm to dualize a regular switching function, *IEEE Trans. on Computers* **28** (1979), 238-243.

[10] E. Lawler, J. K. Lenstra and A. H. G. Rinnooy Kan, Generating all maximal independent sets: NP-hardness and polynomial-time algorithms, *SIAM Journal on Computing*, 9 (1980) 558-565.

[11] K. Makino and T. Ibaraki, Interor and exterior functions of Boolean functions, *Discrete Applied Mathematics*, 69 (1996) 209-231.

[12] K. Makino and T. Ibaraki, A fast and simple algorithm for identifying 2-monotonic positive Boolean functions, *Journal of Algorithms*, 26 (1998) 291-305.

[13] O. L. Mangasarian, Mathematical programming in machine learning, in G. Di. Pillo and F. Giannessi eds, Nonlinear Optimization and Applications (Plenum Publishing, New York, 1996) 283-295.

[14] U. N. Peled and B. Simeone, Polynomial-time algorithm for regular set-covering and threshold synthesis, *Discrete Applied Mathematics* 12 (1985) 57-69.

[15] U. N. Peled and B. Simeone, An $O(nm)$-time algorithm for computing the dual of a regular Boolean function, *Discrete Applied Mathematics* 49 (1994) 309-323.

[16] A. Prékopa, *Stochastic Programming*, (Kluwer, Dordrecht, 1995).