Generating Maximal Independent Sets for Hypergraphs with Bounded Edge-Intersections ⋆

E. Boros[1], K. Elbassioni[1], V. Gurvich[1], and L. Khachiyan[2]

[1] RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway NJ 08854-8003; {boros,elbassio,gurvich}@rutcor.rutgers.edu

[2] Department of Computer Science, Rutgers University, 110 Frelinghuysen Road, Piscataway NJ 08854-8003; leonid@cs.rutgers.edu

Abstract. Given a finite set V, and integers $k \geq 1$ and $r \geq 0$, denote by $\mathbb{A}(k, r)$ the class of hypergraphs $\mathcal{A} \subseteq 2^V$ with (k, r)-bounded intersections, i.e. in which the intersection of any k distinct hyperedges has size at most r. We consider the problem $MIS(\mathcal{A}, \mathcal{I})$: given a hypergraph $\mathcal{A}$ and a subfamily $\mathcal{I} \subseteq \mathcal{I}(\mathcal{A})$, of its maximal independent sets (MIS) $\mathcal{I}(\mathcal{A})$, either extend this subfamily by constructing a new MIS $I \in \mathcal{I}(\mathcal{A}) \setminus \mathcal{I}$ or prove that there are no more MIS, that is $\mathcal{I} = \mathcal{I}(\mathcal{A})$. We show that for hypergraphs $\mathcal{A} \in \mathbb{A}(k, r)$ with $k + r \leq const$, problem MIS$(\mathcal{A}, \mathcal{I})$ is NC-reducible to problem MIS$(\mathcal{A}', \emptyset)$ of generating a single MIS for a partial subhypergraph $\mathcal{A}'$ of $\mathcal{A}$. In particular, for this class of hypergraphs, we get an incremental polynomial algorithm for generating all MIS. Furthermore, combining this result with the currently known algorithms for finding a single maximal independent set of a hypergraph, we obtain efficient parallel algorithms for incrementally generating all MIS for hypergraphs in the classes $\mathbb{A}(1, c)$, $\mathbb{A}(c, 0)$, and $\mathbb{A}(2, 1)$, where c is a constant. We also show that, for $\mathcal{A} \in \mathbb{A}(k, r)$, where $k + r \leq const$, the problem of generating all MIS of $\mathcal{A}$ can be solved in incremental polynomial-time with space polynomial only in the size of $\mathcal{A}$.

1 Introduction

Let $\mathcal{A} \subseteq 2^V$ be a hypergraph (set family) on a finite vertex set V. A vertex set $I \subseteq V$ is called *independent* if I contains no hyperedge of $\mathcal{A}$. Let $\mathcal{I}(\mathcal{A}) \subseteq 2^V$ denote the family of all maximal independent sets (MIS) of $\mathcal{A}$. We assume that $\mathcal{A}$ is given by the list of its hyperedges and consider problem $MIS(\mathcal{A})$ of incrementally generating all sets in $\mathcal{I}(\mathcal{A})$. Clearly, this problem can be solved by performing $|\mathcal{I}(\mathcal{A})| + 1$ calls to the following problem:

$MIS(\mathcal{A}, \mathcal{I})$: *Given a hypergraph $\mathcal{A}$ and a collection $\mathcal{I} \subseteq \mathcal{I}(\mathcal{A})$ of its maximal independent sets, either find a new maximal independent set $I \in \mathcal{I}(\mathcal{A}) \setminus \mathcal{I}$, or prove that the given collection is complete, $\mathcal{I} = \mathcal{I}(\mathcal{A})$.*

⋆ This research was supported in part by the National Science Foundation, grant IIS-0118635. The research of the first and third authors was also supported in part 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.

Note that if $I \in \mathcal{I}(\mathcal{A})$ is an independent set, the complement $B = V \setminus I$ is a *transversal* to $\mathcal{A}$, that is $B \cap A \neq \emptyset$ for all $A \in \mathcal{A}$, and vice versa. Hence $\{B \mid B = V \setminus I,\ I \in \mathcal{I}(\mathcal{A})\} = \mathcal{A}^d$, where $\mathcal{A}^d \stackrel{\text{def}}{=} \{B \mid B \text{ is a minimal transversal to } \mathcal{A}\}$ is the *transversal* or *dual* hypergraph of $\mathcal{A}$. For this reason, MIS($\mathcal{A}, \mathcal{I}$) can be equivalently stated as the *hypergraph dualization problem*:

$DUAL(\mathcal{A}, \mathcal{B})$: *Given a hypergraph $\mathcal{A}$ and a collection $\mathcal{B} \subseteq \mathcal{A}^d$ of minimal transversals to $\mathcal{A}$, either find a new minimal transversal $B \in \mathcal{A} \setminus \mathcal{B}$ or show that $\mathcal{B} = \mathcal{A}$.*

This problem has applications in combinatorics, graph theory, artificial intelligence, reliability theory, database theory, integer programming, and learning theory (see, e.g. [5, 9]). It is an open question whether problem DUAL($\mathcal{A}, \mathcal{B}$), or equivalently MIS($\mathcal{A}, \mathcal{I}$), can be solved in polynomial time for arbitrary hypergraphs. The fastest currently known algorithm [11] for DUAL($\mathcal{A}, \mathcal{B}$) is quasi-polynomial and runs in time $O(nm) + m^{o(\log m)}$, where $n = |V|$ and $m = |\mathcal{A}| + |\mathcal{B}|$.

It was shown in [6, 9] that in the case of hypergraphs of bounded dimension, $\dim(\mathcal{A}) \stackrel{\text{def}}{=} \max_{A \in \mathcal{A}} |A| \leq const$, problem MIS($\mathcal{A}, \mathcal{I}$) can be solved in polynomial time. Moreover, [4] shows that the problem can be efficiently solved in parallel, MIS($\mathcal{A}, \mathcal{I}$) $\in NC$ for $\dim(\mathcal{A}) \leq 3$ and MIS($\mathcal{A}, \mathcal{I}$) $\in RNC$ for $\dim(\mathcal{A}) = 4, 5...$ Let us also mention that for graphs, $\dim(\mathcal{A}) \leq 2$, all MIS can be generated with polynomial delay, see [13] and [19].

In [8], a total polynomial time generation algorithm was obtained for the hypergraphs of bounded degree, $\deg(\mathcal{A}) \stackrel{\text{def}}{=} \max_{v \in V} |\{A : v \in A \in \mathcal{A}\}| \leq const$. This result was recently strengthened in [10], where a polynomial delay algorithm was obtained for a wider class of hypergraphs.

In this paper we consider the class $\mathbb{A}(k, r)$ of hypergraphs with (k, r)-*bounded intersections*: $\mathcal{A} \in \mathbb{A}(k, r)$ if the intersection of each (at least) k distinct hyperedges of $\mathcal{A}$ is of cardinality at most r. We will always assume that $k \geq 1$ and $r \geq 0$ are fixed integers whose sum is bounded, $k + r \leq c = const$. Note that

$$\dim(\mathcal{A}) \leq r \quad \text{iff} \quad \mathcal{A} \in \mathbb{A}(1, r) \qquad \text{and} \qquad \deg(\mathcal{A}) < k \quad \text{iff} \quad \mathcal{A} \in \mathbb{A}(k, 0),$$

and hence, the class $\mathbb{A}(k, r)$ contains both the bounded-dimension and bounded-degree hypergraphs as subclasses. It will be shown that problem MIS($\mathcal{A}, \mathcal{I}$) can be solved in polynomial time for hypergraphs with (k, r)-bounded intersections. It is not difficult to see that for any hypergraph $\mathcal{A} \in \mathbb{A}(k, r)$ the following property holds for every vertex-set $X \subseteq V$: X is contained in a hyperedge of $\mathcal{A}$ whenever each subset of X of cardinality at most $c = k + r$ is contained in a hyperedge of $\mathcal{A}$. [Indeed, suppose that X is a minimal subset of V not contained in any hyperedge of $\mathcal{A}$, and that every subset of X of cardinality at most $k + r$ is contained in a hyperedge of $\mathcal{A}$. Note that $|X| \geq k + r + 1$. Let $e_1, \ldots, e_k$ be distinct elements of X. Then the exist distinct hyperedges $A_1, \ldots, A_k \in \mathcal{A}$ such that $X \setminus \{e_i\} \subseteq A_i$, for $i = 1, \ldots, k$. Now we get a contradiction to the fact that $\mathcal{A} \in \mathbb{A}(k, r)$ since $|A_1 \cap \ldots \cap A_k| \geq r + 1$.] Hypergraphs $\mathcal{A} \subseteq 2^V$ with this property were introduced by Berge [3] under the name of *c-conformal hypergraphs*, and clearly define a wider class of hypergraphs than $\mathbb{A}(k, r)$ with $k + r = c$. In fact, we will prove our result for this wider class of *c-conformal* hypergraphs.

Theorem 1. *For the c-conformal hypergraphs, $c \le const$, and in particular for $\mathcal{A} \in \mathbb{A}(k,r)$, $k+r \le c = const$, problem MIS($\mathcal{A},\mathcal{I}$) is polynomial and hence $\mathcal{I}(\mathcal{A})$, the set all MIS of $\mathcal{A}$, can be generated in incremental polynomial time.*

Theorem 1 is a corollary of the following stronger theorem which will be proved in Section 2.

Theorem 2. *For any c-conformal hypergraph $\mathcal{A}$, where c is a constant, problem MIS($\mathcal{A},\mathcal{I}$) is NC-reducible to MIS($\mathcal{A}',\emptyset$), where $\mathcal{A}'$ is a partial sub-hypergraph of $\mathcal{A}$.*

In Section 2, we also derive some further consequences of Theorem 2, related to the parallel complexity of problem MIS($\mathcal{A},\mathcal{I}$) for certain classes of hypergraphs.

Let us note that our algorithm of generating $I(\mathcal{A})$ based on Theorem 1 is incremental, since it requires solving problem MIS($\mathcal{A},\mathcal{I}$) iteratively $|I(\mathcal{A})|+1$ times. Thus, this algorithm may require space exponential in the size of the input hypergraph $N = N(\mathcal{A}) = \sum_{A \in \mathcal{A}} |A|$. A generation algorithm for $\mathcal{I}(\mathcal{A})$ is said to work in *polynomial space* if the total space required by the algorithm to output all the elements of $\mathcal{I}(\mathcal{A})$ is polynomial in N. In Section 3, we prove the following.

Theorem 3. *For the hypergraphs of bounded intersections, $\mathcal{A} \in \mathbb{A}(k,r)$, where $k+r \le const$, all MIS of $\mathcal{A}$ can be enumerated in incremental polynomial time and with polynomial space.*

Finally, we conclude in Section 4, with a third algorithm for generating all maximal independent sets of a hypergraph $A \in \mathbb{A}(k,r)$, $k+r \le const$.

2 NC-reduction for *c*-Conformal Hypergraphs

The results of [4] show that, for hypergraphs of bounded dimension $\mathcal{A}(1,c)$, there is an NC-reduction from MIS($\mathcal{A},\mathcal{I}$) to MIS($\mathcal{A}',\emptyset$), where $\mathcal{A}'$ is a partial sub-hypergraph of $\mathcal{A}$. In other words, the problem of extending in parallel a given list of MIS of $\mathcal{A}$ can be reduced to the problem of generating in parallel a single MIS for a partial sub-hypergraph of $\mathcal{A}$. In this section we extend this reduction to the class of c-conformal hypergraphs, when c is a constant.

2.1 *c*-Conformal Hypergraphs

Given a hypergraph $\mathcal{A} \subseteq 2^V$, we say that $\mathcal{A}$ is *Sperner* if no hyperedge of $\mathcal{A}$ contains another hyperedge. By definition, for every hypergraph $\mathcal{A}$, its MIS hypergraph $\mathcal{I}(\mathcal{A})$ is Sperner. Let us inverse the operator $\mathcal{I}$. Given a Sperner hypergraph $\mathcal{B} \subseteq 2^V$, introduce the hypergraph $\mathcal{A} = \mathcal{I}^{-1}(\mathcal{B}) \subseteq 2^V$ whose hyperedges are all minimal subsets $A \subseteq V$ which are not contained in any hyperedge of $\mathcal{B}$, that is $A \subseteq B$ for no $A \in \mathcal{A}, B \in \mathcal{B}$ and $A' \subseteq B$ for some $B \in \mathcal{B}$ for each proper subset $A' \subset A \in \mathcal{A}$. The hypergraph $\mathcal{A} = \mathcal{I}^{-1}(\mathcal{B})$ is Sperner by definition, too. It is also easy to see that $\mathcal{B}$ is the MIS hypergraph of $\mathcal{A}$. In other words, for Sperner hypergraphs $\mathcal{B} = \mathcal{I}(\mathcal{A})$ if and only if $\mathcal{A} = \mathcal{I}^{-1}(\mathcal{B})$. In [3], Berge introduced the class of c-conformal hypergraphs and characterized them in several equivalent ways as follows.

Proposition 1 ([3]). *For each hypergraph $\mathcal{A} \subseteq 2^V$ the following statements are equivalent: (i) $\mathcal{A}$ is c-conformal; (ii) The transposed hypergraph $\mathcal{A}^T$ (whose incidence matrix is the transposed incidence matrix of $\mathcal{A}$) satisfies the $(c-1)$-dimensional Helly property: a subset of hyperedges from $\mathcal{A}^T$ has a common vertex whenever every at most c hyperedges of this subset have one; (iii) For each partial hypergraph $\mathcal{A}' \subseteq \mathcal{A}$ having $c+1$ edges, the set $\{x \in V \mid d_{\mathcal{A}'}(x) \geq c\}$ of vertices of degree at least c in $\mathcal{A}'$, is contained in an edge of $\mathcal{A}$.*

It is not difficult to see that we can add to the above list the following equivalent characterization:
(iv) $\dim(\mathcal{I}^{-1}(\mathcal{A})) \leq c$.

Note also that (iii) gives a polynomial-time membership test for c-conformal hypergraphs, for a fixed constant c. Thus even though, given a hypergraph $\mathcal{A}$, the precise computation of $\dim(\mathcal{I}^{-1}(\mathcal{A}))$ is an NP-complete problem (it can be reduced from *stability number for graphs*), verifying condition (iv) is polynomial for every fixed c by Proposition 1.

Given a hypergraph $\mathcal{A} \subseteq 2^V$, let us introduce the *complementary* hypergraph $\mathcal{A}^c = \{V \setminus A \mid A \in \mathcal{A}\}$ whose hyperedges are complementary to the hyperedges of $\mathcal{A}$. It is easy to see that $\mathcal{A}^{dd} = \mathcal{A}$, $\mathcal{A}^{cc} = \mathcal{A}$ for each Sperner hypergraph $\mathcal{A}$. In other words, both operations, duality and complementation, are involutions. It is also clear that $\mathcal{A}^{dc} = \mathcal{I}(\mathcal{A})$ and $\mathcal{A}^{cd} = \mathcal{I}^{-1}(\mathcal{A})$.

A vertex set S is called a *sub-transversal* of $\mathcal{A}$ if $S \subseteq B$ for some minimal transversal $B \in \mathcal{A}^d$. Our proof of Theorem 2 makes use of a characterization of sub-transversals suggested in [6].

2.2 Characterization of Sub-transversals to a Hypergraph

Given a hypergraph $\mathcal{A} \subseteq 2^V$, a subset $S \subseteq V$, and a vertex $v \in S$, let $\mathcal{A}_v(S) = \{A \in \mathcal{A} \mid A \cap S = \{v\}\}$ denote the family of all hyperedges of $\mathcal{A}$ whose intersection with S is exactly v. Let further $\mathcal{A}_0(S) = \{A \in \mathcal{A} \mid A \cap S = \emptyset\}$ denote the partial hypergraph consisting of the hyperedges of $\mathcal{A}$ disjoint from S. A selection of $|S|$ hyperedges $\{A_v \in \mathcal{A}_v(S) \mid v \in S\}$ is called *covering* if there exists a hyperedge $A \in \mathcal{A}_0(S)$ such that $A \subseteq \bigcup_{v \in S} A_v$.

Proposition 2 (cf. [6]). *Let $S \subseteq V$ be a non-empty vertex set in a hypergraph $\mathcal{A} \in 2^V$.*

i) *If S is a sub-transversal for $\mathcal{A}$ then there exists a non-covering selection $\{A_v \in \mathcal{A}_v(S) \mid v \in S\}$ for S.*

ii) *Given a non-covering selection $\{A_v \in \mathcal{A}_v(S) \mid v \in S\}$ for S, we can extend S to a minimal transversal of $\mathcal{A}$ by solving problem MIS$(\mathcal{A}', \emptyset)$ for the induced partial hypergraph $\mathcal{A}' = \{A \cap U \mid A \in \mathcal{A}_0(S)\} \subseteq 2^U$, where $U = V \setminus \bigcup_{v \in S} A_v$.*

Unfortunately, finding a non-covering selection for S (or equivalently, testing if S is a sub-transversal) is NP-hard if the cardinality of S is not bounded (see [4]). However, if the size of S is bounded by a constant then there are only polynomially many selections $\{A_v \in \mathcal{A}_v(S) \mid v \in S\}$ for S. All of these selections, including the non-covering ones, can be easily enumerated in polynomial time (moreover, it can be done in parallel).

Corollary 1. *For any fixed c there is an NC algorithm which, given a hypergraph* $\mathcal{A} \subseteq 2^V$ *and a set* S *of at most* c *vertices, determines whether* S *is a sub-transversal to* $\mathcal{A}$ *and if so finds a non-covering selection* $\{A_v \in \mathcal{A}_v(S) \mid v \in S\}$.

Note that this Corollary holds for hypergraphs of arbitrary dimension.

2.3 Proof of Theorem 2

We prove the theorem for the equivalent problem DUAL$(\mathcal{A}, \mathcal{B})$. We may assume without loss of generality that $\mathcal{A}$ is Sperner. Our reduction consists of the following steps:
Step 1. By definition, each set $B \in \mathcal{B}$ is a minimal transversal to $\mathcal{A}$. This implies that each set $A \in \mathcal{A}$ is transversal to $\mathcal{B}$. Check whether each $A \in \mathcal{A}$ is a *minimal* transversal to $\mathcal{B}$. If not, a new element in $\mathcal{A}^d \setminus \mathcal{B}$ can be found be calling problem MIS$(\mathcal{A}', \emptyset)$, for some induced partial hypergraph $\mathcal{A}'$ of $\mathcal{A}$. We may assume therefore that each set in $\mathcal{A}$ is a minimal transversal to $\mathcal{B}$, i.e. $\mathcal{A} \subseteq \mathcal{B}^d$. Recall that $\mathcal{A}^{dd} = \mathcal{A}$ for each Sperner hypergraph $\mathcal{A}$. Therefore, if $B \neq \mathcal{A}^d$ then $\mathcal{A} \neq \mathcal{B}^d$, and thus $\mathcal{B}^d \setminus \mathcal{A} \neq \emptyset$. Hence we arrive at the following duality criterion: $\mathcal{A}^d \setminus \mathcal{B} \neq \emptyset$ iff there is a sub-transversal S to $\mathcal{B}$ such that

$$S \subseteq A \;\text{ for no }\; A \in \mathcal{A}. \tag{1}$$

Hence we can apply the sub-transversal test only to S such that

$$|S| \leq \dim(\mathcal{I}^{-1}(\mathcal{A})). \tag{2}$$

Step 2 (Duality test.) For each set S satisfying (1), (2) and the condition that

$$A \not\subseteq S \text{ for all } A \in \mathcal{A}, \tag{3}$$

check whether or not

$$S \text{ is a sub-transversal to } \mathcal{B}. \tag{4}$$

We need the assumption that $\dim(\mathcal{I}^{-1}(\mathcal{A}))$ is bounded to guarantee that this step is polynomial (and moreover, is in NC). Recall that by Proposition 2, S satisfies (4) iff there is a selection

$$\{B_v \in \mathcal{B}_v(S) \mid v \in S\} \tag{5}$$

which covers no set $B \in \mathcal{B}_0(S)$. Here as before, $\mathcal{B}_0(S) = \{B \in \mathcal{B} \mid B \cap S = \emptyset\}$ and $\mathcal{B}_v(S) = \{B \in \mathcal{B} \mid B \cap S = \{v\}\}$ for $v \in S$.

If conditions (1)-(4) cannot be met, we conclude that $\mathcal{B} = \mathcal{A}^d$ and halt.
Step 3. Suppose we have found a non-covering selection (5) for some set S satisfying (1)-(4). Then it is easy to see that the set $Z = S \bigcup \left[V \setminus \bigcup_{v \in S} B_v\right]$ is independent in $\mathcal{A}$. Furthermore, Z is transversal to $\mathcal{B}$, because selection (5) is non-covering. Let $\mathcal{A}' = \{A \cap U \mid A \in \mathcal{A}\}$, where $U = V \setminus Z$, and let T be a minimal transversal to $\mathcal{A}'$. (As before, we can let $T = U \setminus \mathit{output}(MIS(\mathcal{A}', \emptyset)$.) Since Z is an independent set of $\mathcal{A}$, we have $T \cap A \neq \emptyset$ for all $A \in \mathcal{A}$, that is T is transversal to $\mathcal{A}$. Clearly, T is minimal, that is $T \in \mathcal{A}^d$. It remains to argue that T is a *new* minimal transversal to $\mathcal{A}$, that is $T \notin \mathcal{B}$. This follows from the fact that Z is transversal to $\mathcal{B}$ and disjoint from T. □

Note that Theorem 2 does not imply that $\text{MIS}(\mathcal{A}, \mathcal{I}) \in NC$ because the parallel complexity of the resulting problem $\text{MIS}(\mathcal{A}', \emptyset)$ is not known. The question whether it is in NC in general (for arbitrary hypergraphs) was raised in [14]. The affirmative answers were obtained in [1, 7, 15, 17] for the following special cases: For hypergraphs of bounded dimension, $\mathcal{A} \in \mathbb{A}(1, c)$, it is known that $\text{MIS}(\mathcal{A}', \emptyset) \in NC$ for $c \leq 3$, and $\text{MIS}(\mathcal{A}', \emptyset) \in RNC$ for $c = 4, 5, \ldots$, see [2, 15]. Furthermore, it was shown in [17, 18] that $\text{MIS}(\mathcal{A}', \emptyset) \in NC$ for the so-called *linear* hyperedges, in which each two hyperedges intersect in at most one vertex, that is for $\mathcal{A}' \in \mathbb{A}(2, 1)$. Finally, it follows from [12] that $\text{MIS}(\mathcal{A}', \emptyset) \in NC$ for hypergraphs of bounded degree, that is for $\mathcal{A}' \in \mathbb{A}(c, 0)$. Combining the above results with Theorem 2, we obtain the following corollary.

Corollary 2. *Problem MIS($\mathcal{A}, \mathcal{I}$) is in RNC for $\mathcal{A} \in \mathbb{A}(1, c)$, where c is a constant (hypergraphs of bounded dimension). Furthermore, MIS($\mathcal{A}, \mathcal{I}$) is in NC for $\mathcal{A} \in \mathbb{A}(1, c)$, $c \leq 3$ (hypergraphs of $\dim \leq 3$), for $\mathcal{A} \in \mathbb{A}(c, 0)$, where c is a constant (hypergraphs of bounded degree), and for $\mathcal{A} \in \mathbb{A}(2, 1)$ (linear hypergraphs).*

Yet, for a hypergraph $\mathcal{A}$ satisfying $\dim(\mathcal{I}^{-1}(\mathcal{A})) \leq \mathit{const}$, or even more specifically for $\mathcal{A} \in \mathbb{A}(k, r)$, $k + r \leq \mathit{const}$, we only have an NC-reduction of $\text{MIS}(\mathcal{A}, \mathcal{I})$ to $\text{MIS}(\mathcal{A}', \emptyset)$, where the parallel complexity of the latter problem is not known.

3 Polynomial Space Algorithm for Generating $\mathcal{A}^d$

For $i = 1, \ldots, n$ denote by $[i : n]$ the set $\{i, i+1, \ldots, n\}$, where $[n+1 : n]$ is assumed to be the empty set. Given a hypergraph $\mathcal{A} \subseteq 2^{[n]}$, we shall say that $X \subseteq [n]$ is an i-minimal transversal for $\mathcal{A}$ if $X \supseteq [i : n]$, X is a transversal of $\mathcal{A}$, and $X \setminus \{j\}$ is not a transversal for all $j \in X \cap [1 : i-1]$. Thus, $n+1$-minimal transversals are just the minimal transversals of $\mathcal{A}$. For $i = 1, \ldots, n$, let $\mathcal{A}^{d_i}$ be the family of i-minimal transversals for $\mathcal{A}$.

Given $i \in [n]$ and $X \in \mathcal{A}^{d_i}$, let $\mathcal{A}_i(X)$ be the hypergraph

$$\mathcal{A}_i(X) = \{A \setminus \{i\} \ : \ A \in \mathcal{A}, \ A \cap X = \{i\}\}.$$

Proposition 3 (see [10, 16]).
(i) $|\mathcal{A}^{d_i}| \leq |\mathcal{A}^d|$, *for* $i = 1, \ldots, n+1$.
(ii) $|\mathcal{A}_i(X)^d| \leq |\mathcal{A}^{d_{i+1}}|$, *for* $i \in [n]$ *and* $X \in \mathcal{A}^{d_i}$.

Now consider the following generalization of an algorithm in [19] for generating maximal independent sets in graphs (see also [13] and [16]). Given $i \in [n]$, and $X \in \mathcal{A}^{d_i}$, we assume in the algorithm that the minimal transversals $\mathcal{A}_i(X)^d$ are computed by calling a process $P(i, X)$ that invokes the same algorithm recursively on the partial hypergraph $\mathcal{A}_i(X)$. We further assume that, once $P(i, X)$ finds an element $Y \in \mathcal{A}_i(X)^d$, it returns control to the calling process GEN$(\mathcal{A}, i, X)$. When called for the next time, $P(i, X)$ returns the next element of $\mathcal{A}_i(X)^d$ that has not been generated yet, if such an element exists.

Algorithm GEN$(\mathcal{A}, i, X)$**:**
Input: A hypergraph $\mathcal{A}$, an index $i \in [n]$, and an i-minimal transversal $X \in \mathcal{A}^{d_i}$.
Output: All minimal transversals of $\mathcal{A}$.

1. if $i = n + 1$ then
2.   output X;
3. else
4.   if $X \setminus \{i\}$ is a transversal of $\mathcal{A}$ then
5.   GEN$(\mathcal{A}, i + 1, X \setminus \{i\})$;
6.   else
7.   GEN$(\mathcal{A}, i + 1, X)$;
8.   for each minimal transversal $Y \in \mathcal{A}_i(X)^d$ (found recursively) do
9.   if $X \cup Y \setminus \{i\} \in \mathcal{A}^{d_{i+1}}$ then
10.   Compute the *lexico. largest* set $Z \subseteq X \cup Y$ such that $Z \in \mathcal{A}^{d_i}$;
11.   if $Z = X$ then
12.   GEN$(\mathcal{A}, i + 1, X \cup Y \setminus \{i\})$;

Lemma 1. *When called with $i = 1$ and $X = [n]$, Algorithm GEN$(\mathcal{A}, i, X)$ outputs all minimal transversals of $\mathcal{A}$ with no repetitions.*

Proof. Consider the recursion tree $\mathbf{T}$ traversed by the algorithm. Label each node of tree by the pair (i, X) which represents the input to the algorithm at this node. Clearly i represents the level of node (i, X) in the tree (where the root of $\mathbf{T}$ is at level 1). By induction on $i = 1, \ldots, n + 1$, we can verify the following statement:

$$\mathcal{A}^{d_i} = \{X \subseteq [n] \ : \ (i, X) \in \mathbf{T}\}. \tag{6}$$

Indeed, this trivially holds at $i = 1$. Assume now that (6) holds for a specific $i \in [n - 1]$. It is easy to see that any node $(i + 1, X) \in \mathbf{T}$ generated at level $i + 1$ of the tree must have $X \in \mathcal{A}^{d_{i+1}}$. Thus it remains to verify that $\mathcal{A}^{d_{i+1}} \subseteq \{X \ : \ (i + 1, X) \in \mathbf{T}\}$. To see this, let X' be an arbitrary element of $\mathcal{A}^{d_{i+1}}$. Note first that if $X' \ni i$ then $X' \setminus \{i\}$ is not a transversal of $\mathcal{A}$ and $X' \in \mathcal{A}^{d_i}$, and therefore by induction we have a node $(i, X') \in \mathbf{T}$. Consequently, we get a node $(i+1, X') \in \mathbf{T}$ as a child of $(i, X') \in \mathbf{T}$, by Step 7 of the algorithm. Let us therefore assume that $X' \not\ni i$. Note that X' must contain a subset $X \setminus \{i\}$, for some $X \in \mathcal{A}^{d_i}$. This is because $X' \cup \{i\}$ is a transversal and therefore it contains an i-minimal transversal X of $\mathcal{A}$. Among all the sets X satisfying this property, let Z be the lexicographically largest. Now, if $Z \setminus \{i\}$ is a transversal of $\mathcal{A}$, then $Z \setminus \{i\} = X'$ and Step 5 will create a node $(i + 1, X') \in \mathbf{T}$ as the only child of

$(i, Z) \in \mathbf{T}$. On the other hand, if $Z \setminus \{i\}$ is not a transversal, then X' can be written as $X' = Z \cup Y \setminus \{i\}$, for some $Y \in \mathcal{A}_i(Z)^d$. But then node $(i+1, X')$ will be generated as a child of $(i, Z) \in \mathbf{T}$ by Step 12 of the Algorithm. This completes the proof of (6). Finally, it follows from Step 10 that each node in the tree is generated as the child of exactly one other node. Consequently each leaf is visited, and hence each set $X \in \mathcal{A}^d$ is output, only once and the lemma follows. □

The next lemma states that, for hypergraphs $\mathcal{A}$ of (k, r)-bounded intersections, Algorithm GEN is a polynomial-space, *output*-polynomial time algorithm for generating all minimal transversals of $\mathcal{A}$.

Lemma 2. *The time taken by Algorithm GEN until it outputs the last minimal transversal of a hypergraph* $\mathcal{A} \in \mathbb{A}(k, r)$ *is* $O(n^{k+r-1}|\mathcal{A}^d|^{r+1})$*, and the total space required is* $O(N^{r+1})$*.*

Proof. For a hypergraph $\mathcal{A} \in \mathbb{A}(k, r)$, let $T(\mathcal{A})$ and $M(\mathcal{A})$ be respectively the time and space required by Algorithm GEN to output the last minimal transversal of $\mathcal{A}$. Note that the algorithm basically performs depth-first search on the tree $\mathbf{T}$ (whose leaves are the elements of $\mathcal{A}^d$), and only generates nodes of $\mathbf{T}$ as needed during the search. Since each node of the tree $\mathbf{T}$, which is not a leaf, has at least one child, the time between two successive outputs generated by the algorithm does not exceed the time required to generate the children of nodes along a complete path of the tree $\mathbf{T}$ from the root to a leaf. But, as can be seen from the algorithm, for a given node $v = (i, X)$ in $\mathbf{T}$, where $i \in [n]$ and $X \in \mathcal{A}^{d_i}$, the time required to generate all the children of v, is bounded by the time to output all the elements of $\mathcal{A}_i(X)^d$. Since the depth of the tree is $n + 1$, we get the recurrence

$$T(\mathcal{A}) \le n|\mathcal{A}^d| \max\{T(\mathcal{A}_i(X)) \ : \ i \in [n], \quad X \in \mathcal{A}^{d_i}\}. \tag{7}$$

Note that $\mathcal{A}_i(X) \in \mathbb{A}(k, r-1)$. Furthermore, by Proposition 3, we have $|\mathcal{A}_i(X)^d| \le |\mathcal{A}^d|$, and thus (7) gives $T(\mathcal{A}) \le (n|\mathcal{A}^d|)^r T(\mathcal{A}')$, for some sub-hypergraph $\mathcal{A}' \in \mathbb{A}(k, 0)$ of $\mathcal{A}$ which satisfies $|(\mathcal{A}')^d| \le |\mathcal{A}^d|$. Now, we observe that for any $i \in [n]$ and $X \in (\mathcal{A}')^{d_i}$, we have $|\mathcal{A}_i'(X)| \le k - 1$, and hence it follows that $T(\mathcal{A}') = O(n^{k-1}|(\mathcal{A}')^d|)$. The bound on the running time follows.

Now let us consider the total memory required by the algorithm. Since, for each recursion tree (corresponding to a (sub-)hypergraph that is to be dualized), the algorithm maintains only the path from the root to a leaf of the tree, we get the recurrence $M(\mathcal{A}) \le N \max\{M(\mathcal{A}_i(X)) \ : \ i \in [n], \quad X \in \mathcal{A}^{d_i}\}$. This recurrence again gives $M(\mathcal{A}) \le N^r M(\mathcal{A}')$, for some sub-hypergraph $\mathcal{A}' \in \mathbb{A}(k, 0)$ of $\mathcal{A}$. But $M(\mathcal{A}') = O(N)$ and the bound on the space follows. □

Now Theorem 3 follows by combining Lemma 2 with the following reduction.

Proposition 4. *Let* $\mathcal{A} \subseteq 2^{[n]}$ *be a hypergraph. Suppose that there is an algorithm* P *that generates all minimal transversals of* $\mathcal{A}$ *in time* $p(n, |\mathcal{A}^d|)$ *and space* $q(N(\mathcal{A}))$*, for some polynomials* $p(\cdot, \cdot)$ *and* $q(\cdot)$*. Then for any integer* k*, we can* *generate at least k minimal transversals of* $\mathcal{A}$ *in time* $2n(p(n,k)+1)$ *and space* $q(N(\mathcal{A}))$.

Note that it is implicit in the proof of Lemma 2 that, for both graphs $\mathcal{A} \in \mathbb{A}(1,2)$ and hypergraphs of bounded degree $\mathcal{A} \in \mathbb{A}(c,0)$, Algorithm GEN is in fact a polynomial delay and polynomial space algorithm for generating $\mathcal{A}^d$. In particular, Theorem 3 implies the following previously known results [10, 13, 19].

Corollary 3. *For graphs,* $\mathcal{A} \in \mathbb{A}(1,2)$*, and also for the hypergraphs of bounded degree,* $\mathcal{A} \in \mathbb{A}(c,0)$*, all minimal transversals of* $\mathcal{A}$ *can be enumerated with polynomial delay and polynomial space.*

4 Generating $\mathcal{A}^d$ Using the Supergraph Approach

Let $\mathcal{A} \subseteq 2^V$ be a hypergraph. In this section, we sketch another algorithm to list all minimal transversals of $\mathcal{A}$. The algorithm works by building a *strongly connected directed supergrah* $\mathcal{G} = (\mathcal{A}^d, \mathcal{E})$ on the set of minimal transversals, in which a pair of vertices (X, X') forms an edge in $\mathcal{E}$ if and only if X' can be obtained from X by deleting an element from $X \setminus X'$, adding a minimal subset of elements from $X' \setminus X$ to obtain a transversal, and finally reducing the resulting set to a minimal feasible solution in a specified way (say in reverse-lexicographic order). In other words, $(X, X') \in \mathcal{E}$ if and only if $X' \subseteq X \cup Z \setminus \{e\}$, for some $e \in X \setminus X'$ and $Z \subseteq X' \setminus X$, such that Z is minimal with the property that $X \cup Z \setminus \{e\}$ is a transversal.

The strong connectivity of $\mathcal{G}$ can be proved as follows. Given two vertices $X_0, X_l \in \mathcal{A}^d$ of $\mathcal{G}$, there exists a set $\{X_1, \ldots, X_{l-1}\}$ of elements of $\mathcal{F}$, where for all $i = 1, \ldots, l$, X_i is obtained from X_{i-1} by deleting an element $e_i \in X_{i-1} \setminus X_l$ (thus making $X_{i-1} \setminus \{e_i\}$ non-transversal), adding a minimal subset of elements $Z_i \subseteq X_l \setminus X_{i-1}$ to obtain a transversal $X_{i-1} \setminus \{e_i\} \cup Z_i$, and finally, reducing the resulting set to a minimal transversal $X_i \subseteq X_{i-1} \cup Z_i \setminus \{e_i\}$. Note that, for $i = 1, \ldots, l$, $|X_i \setminus X_l| < |X_{i-1} \setminus X_l|$ and therefore $l \leq |X_0 \setminus X_l|$. In other words, $\mathcal{G}$ has *diameter* at most n.

The minimal transversals of $\mathcal{A}$ can thus be generated by performing breadth-first search on the vertices of $\mathcal{G}$, starting from an arbitrary vertex. Such a procedure can be executed in incremental polynomial time if the neighbourhood of every vertex in $\mathcal{G}$ can also be generated in (incremental) polynomial time. Given a hypergraph $\mathcal{A} \in \mathcal{A}(k,r)$, and a minimal transversal $X \in \mathcal{A}^d$, all neighbours of X in $\mathcal{G}$ can be generated in time $O(n^{k+r}|\mathcal{A}^d|^{r+1})$. Indeed, for any $e \in X$, all minimal subsets of vertices Z, such that $X \setminus \{e\} \cup Z$ is a transversal of $\mathcal{A}$, can be obtained by finding all minimal transversals for the hypergraph $\mathcal{A}_e(X) = \{A \setminus \{e\} \;:\; A \in \mathcal{A},\; A \cap X = \{e\}\}$. But as noted before, $\mathcal{A}_e(X) \in \mathcal{A}(k, r-1)$ and $|\mathcal{A}_e(X)^d| \leq |\mathcal{A}^d|$. We conclude therefore, as in the proof of Lemma 2, that the time required to produce all the neighbours of X by applying the algorithm recursively on each of the hypergraphs $\mathcal{A}_e$, for $e \in X$, is $O(n^{k+r}|\mathcal{A}^d|^{r+1})$.

References

1. N. Alon, L. Babai, A. Itai, A fast randomized parallel algorithm for the maximal independent set problem, *J. Algorithms* 7 (1986), pp. 567–583
2. P. Beame, M. Luby, Parallel search for maximal independence given minimal dependence, *Proc. 1st SODA Conference* (1990), pp. 212–218.
3. C. Berge, *Hypergraphs*, North Holland Mathematical Library, Vol. 445, Elsevier-North Holland, Amestrdam, 1989.
4. E. Boros, K. Elbassioni, V. Gurvich, and L. Khachiyan, An efficient incremental algorithm for generating all maximal independent sets in hypergraphs of bounded dimension, *Parallel Processing Letters*, 10 (2000), pp. 253–266.
5. E. Boros, K. Elbassioni, V. Gurvich, L. Khachiyan and K.Makino, Dual-bounded generating problems: All minimal integer solutions for a monotone system of linear inequalities, *SIAM Journal on Computing*, 31 (5) (2002) pp. 1624–1643.
6. E. Boros, V. Gurvich, and P.L. Hammer, Dual subimplicants of positive Boolean functions, *Optimization Methods and Software*, 10 (1998), pp. 147–156.
7. E. Dahlhaus, M. Karpinski, P. Kelsen, An efficient parallel algorithm for computing a maximal independent set in a hypergraph of dimension 3, *Information Processing Letters* 42(6) (1992), pp. 309–313.
8. C. Domingo, N. Mishra and L. Pitt, Efficient read-restricted monotone CNF/DNF dualization by learning with membership queries, *Machine learning* 37 (1999), pp. 89–110.
9. T. Eiter and G. Gottlob, Identifying the minimal transversals of a hypergraph and related problems, *SIAM Journal on Computing*, 24 (1995), pp. 1278–1304.
10. T. Eiter, G. Gottlob and K. Makino, New results on monotone dualization and generating hypergraph transversals, in *Proc. 34th Annual ACM STOC Conference*, (2002), pp. 14–22.
11. M. L. Fredman and L. Khachiyan, On the complexity of dualization of monotone disjunctive normal forms. *J. Algorithms*, 21 (1996), pp. 618–628.
12. O. Garrido, P. Kelsen, A. Lingas, A simple NC-algorithm for a maximal independent set in a hypergraph of polylog arboricity, *Information Processing Letters* 58(2) (1996), pp. 55–58.
13. D. S. Johnson, M. Yannakakis and C. H. Papadimitriou, On generating all maximal independent sets, *Information Processing Letters*, 27 (1988), pp. 119–123.
14. R. Karp and A. Wigderson, A fast parallel algorithm for the maximal independent set problem, *JACM* 32 (1985), pp. 762–773.
15. P. Kelsen, On the parallel complexity of computing a maximal independent set in a hypergraph, *Proc. the 24-th Anual ACM STOC Conference*, (1992), pp. 339–350.
16. 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), pp. 558–565.
17. T. Luczak , E. Szymanska, A parallel randomized algorithm for finding a maximal independent set in a linear hypergraph, *J. Algorithms*, v.25 (1997) n.2, pp. 311–320.
18. E. Szymanska, Derandomization of a Parallel MIS Algorithm in a Linear Hypergraph, in *Proc. 4th International Workshop on Randomization and Approximation Techniques in Computer Science*, (2000), pp. 39–52.
19. S. Tsukiyama, M. Ide, H. Ariyoshi and I. Shirakawa, A new algorithm for generating all maximal independent sets, *SIAM Journal on Computing*, 6 (1977), pp. 505–517.