Lower Bounds for Three Algorithms for Transversal Hypergraph Generation ⋆

Matthias Hagen [1]

Friedrich-Schiller-Universität Jena, Institut für Informatik, D–07737 Jena

Abstract

The computation of all minimal transversals of a given hypergraph in output-polynomial time is a long standing open question known as Transversal Hypergraph Generation. One of the first attempts on this problem—the sequential method [Ber89]—is not output-polynomial as was shown by Takata [Tak07]. Recently, three new algorithms improving the sequential method were published and experimentally shown to perform very well in practice [BMR03,DL05,KS05]. Nevertheless, a theoretical worst-case analysis has been pending. We close this gap by proving lower bounds for all three algorithms. Thereby, we show that none of them is output-polynomial.

Key words: Analysis of algorithms, Computational complexity, Dualization, Transversal hypergraph

1 Introduction

Transversal Hypergraph Generation is the problem to compute, for a given hypergraph $\mathcal{H}$ with vertex set V, the transversal hypergraph $Tr(\mathcal{H})$ that consists of all minimal subsets of V having a non-empty intersection with each hyperedge of $\mathcal{H}$. This problem has many applications in such different fields like artificial intelligence and logic [EG95,EG02], computational biology [Dam06,KSG07,HKS08], computational geometry [KBE$^+$07,KBEG08b],

⋆ A preliminary version of this work was presented at WG 2007 [Hag07].
Email address: hagen@cs.uni-jena.de (Matthias Hagen).
[1] Partially supported by a Landesgraduiertenstipendium Thüringen and the Deutsche Forschungsgemeinschaft (DFG) through project OPAL (optimal algorithms for hard problems in computational biology), NI-369/2.

cryptography [KP04], database theory [MR92], data mining [GKM+03], distributed computing [GB85], e-commerce [ZB06], machine learning [GKMT97], mathematical programming [KBEG08a,Kha00], mobile communication systems [SS98], semantic web [KLM07], topology [DHSW03], and XML [Tri08], to name but a few. For a more detailed list of related problems and respective references see [Hag08, Chapter 3].

Due to the importance of TRANSVERSAL HYPERGRAPH GENERATION there have been various approaches to solve it. But since the size of $Tr(\mathcal{H})$ may be exponential in the size of $\mathcal{H}$, we cannot find an algorithm that runs in time polynomial in the size of the input $\mathcal{H}$. Therefore, another notion of efficient solvability has to be used. An algorithm is said to be output-polynomial if its running time is bounded polynomially in the size of the input and output [JPY88]. Finding an output-polynomial algorithm for TRANSVERSAL HYPERGRAPH GENERATION is a long standing open problem [Pap97]. Moreover, note that the decision version of the problem—given two hypergraphs, decide if one is the transversal hypergraph of the other—is one of the very few problems that currently cannot be classified as polynomial or NP- resp. coNP-hard. The best known algorithms run in time $n^{o(\log n)}$ [FK96,Elb08] or use $O(\log^2 n)$ many nondeterministic bits [EGM03,KS03].

One of the earliest approaches is the sequential method [Ber89]. It computes the transversal hypergraph by iteratively combining transversals of specific subhypergraphs of the input in a brute-force manner. The worst-case analysis of the sequential method took many years until Takata showed that it is not output-polynomial [Tak07]. So far, this is the only proven nontrivial lower bound for any algorithm for TRANSVERSAL HYPERGRAPH GENERATION.

In recent years, several improvements of the sequential method have been published. We focus on the DL-algorithm of Dong and Li [DL05], the BMR-algorithm of Bailey, Manoukian, and Ramamohanarao [BMR03], and the KS-algorithm of Kavvadias and Stavropoulos [KS05]. All three algorithms have been empirically tested on practical instances. Especially the BMR-algorithm performs very well on instances from the data mining field. But while the practical performance of the algorithms has been examined, a theoretical worst-case analysis of their running times has been pending. We close this gap by giving nontrivial lower bounds for all three algorithms. Furthermore, the bounds show that none of the three algorithms is output-polynomial.

The paper is organized as follows. Section 2 contains some basic definitions, a brief recapitulation of the sequential method and its analysis by Takata. In Section 3 we show the DL- and the BMR-algorithm not to be output-polynomial. Section 4 contains the analysis of the KS-algorithm. Some concluding remarks follow in Section 5.

2 Basic Definitions and the Sequential Method

A *hypergraph* $\mathcal{H} = (V, E)$ consists of a set V of vertices and a finite family E of subsets of V—the edges. If there is no danger of ambiguity, we also use the edge set to refer to $\mathcal{H}$. The *size* of $\mathcal{H}$ is the total number of occurrences of vertices in the edges. A *transversal* of $\mathcal{H}$ is a set $t \subseteq V$ that has a non-empty intersection with each edge of $\mathcal{H}$. A transversal t is *minimal* if no proper subset of t is a transversal. The set of all minimal transversals of $\mathcal{H}$ forms the *transversal hypergraph* $Tr(\mathcal{H})$. A hypergraph $\mathcal{H}$ is *simple* if it does not contain two hyperedges e, f with $e \subseteq f$. By $\min(\mathcal{H})$ we denote the simple hypergraph consisting of the minimal hyperedges of $\mathcal{H}$ with respect to set inclusion. Since $\min(\mathcal{H})$ can be easily computed in polynomial time and $Tr(\mathcal{H}) = Tr(\min(\mathcal{H}))$ holds for every hypergraph $\mathcal{H}$, we concentrate on Transversal Hypergraph Generation for simple hypergraphs. But even for simple hypergraphs the size of the transversal hypergraph may be exponential. Hence, there cannot be an algorithm computing the transversal hypergraph in polynomial time. A suitable notion of efficient solvability for such kind of problems is that of output-polynomial time [JPY88]. An algorithm is said to be *output-polynomial* if its running time is bounded polynomially in the sum of the sizes of the input and output.

Given simple hypergraphs $\mathcal{H} = \{e_1, e_2, \ldots, e_m\}$ and $\mathcal{H}' = \{e'_1, e'_2, \ldots, e'_{m'}\}$ there are two different "unions", namely

$$\begin{aligned} \mathcal{H} \cup \mathcal{H}' &= \{e_1, e_2, \ldots, e_m, e'_1, e'_2, \ldots, e'_{m'}\} \text{ and} \\ \mathcal{H} \vee \mathcal{H}' &= \{e_i \cup e'_j : i = 1, 2, \ldots, m,\ j = 1, 2, \ldots, m'\}. \end{aligned}$$

Proposition 2.1 ([Ber89]) *Let $\mathcal{H}$ and $\mathcal{H}'$ be two simple hypergraphs. Then $Tr(\mathcal{H} \cup \mathcal{H}') = \min(Tr(\mathcal{H}) \vee Tr(\mathcal{H}'))$.*

The sequential method [Ber89] uses Proposition 2.1 to generate the transversal hypergraph as follows. For a hypergraph $\mathcal{H} = \{e_1, e_2, \ldots, e_m\}$ let $\mathcal{H}_i = \{e_1, e_2, \ldots, e_i\}$, $i = 1, 2, \ldots, m$. We then have

$$Tr(\mathcal{H}_i) = \min(Tr(\mathcal{H}_{i-1}) \vee Tr(\{e_i\})) = \min(Tr(\mathcal{H}_{i-1}) \vee \{\{v\} : v \in e_i\})$$

and $Tr(\mathcal{H}) = Tr(\mathcal{H}_m)$. This implies a straightforward iterative computation process—the sequential method. A pseudocode listing is given in Algorithm 1. Despite the simplicity of the sequential method it took a couple of decades until Takata [Tak07] presented a nontrivial lower bound using the following inductively defined family of hypergraphs.

Algorithm 1 The Sequential Method

$Tr(\mathcal{H}_1) \leftarrow \{\{v\} : v \in e_1\}$
for $i \leftarrow 2, \ldots, m$ **do**
  $Tr(\mathcal{H}_i) \leftarrow \min(Tr(\mathcal{H}_{i-1}) \vee \{\{v\} : v \in e_i\})$
end for
output $Tr(\mathcal{H}_m)$

$$\mathcal{G}_0 = \{\{v_1\}\} \quad \text{and}$$
$$\mathcal{G}_i = (\mathcal{A} \cup \mathcal{B}) \vee (\mathcal{C} \cup \mathcal{D}),$$
where $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}$ are vertex-disjoint copies of $\mathcal{G}_{i-1}$.

Takata showed the sequential method not to be output-polynomial based on the following observations.

Lemma 2.2 ([Tak07]) *We have* $|V_{\mathcal{G}_i}| = 4^i$, $|\mathcal{G}_i| = 2^{2(2^i-1)}$, $|Tr(\mathcal{G}_i)| = 2^{2^i-1}$. *For* $i \geq 2$ *and any* $e \in \mathcal{G}_i$*, it holds that* $|Tr(\mathcal{G}_i \backslash \{e\}) \backslash Tr(\mathcal{G}_i)| \geq 2^{(i-2)2^i+2}$.

From Lemma 2.2 it follows that, independent of the edge ordering, the penultimate (intermediate) result computed by the sequential method on input $\mathcal{G}_i$ is superpolynomial in the size of the input and output (cf. the original paper [Tak07] for more details).

Very recently, Boros et. al. [BEM08] proved a subexponential $n^{\sqrt{n}}$ upper bound on the running time of the Berge-multiplication.

3 The Algorithms of Dong and Li, and Bailey, Manoukian and Ramamohanarao

The border-differential algorithm of Dong and Li [DL05] comes from the data mining field and is intended for mining emerging patterns. The analogy to the generation of hypergraph transversals was already pointed out by Bailey, Manoukian, and Ramamohanarao [BMR03]. A pseudocode listing of the DL-algorithm is given in Algorithm 2.

The algorithm was experimentally evaluated on many practical data mining cases [DL05] whereas a theoretical analysis of the running time was left open. For this purpose the conversion of the algorithm to the hypergraph setting is very fruitful. The only observable difference between the sequential method and the DL-algorithm is that the DL-algorithm takes special care on how to perform the minimization of $Tr(\mathcal{H}_{i-1}) \vee \{\{v\} : v \in e_i\}$. But as Takata's analysis showed, the minimization is not the bottleneck of the sequential method. Thus, we can extend Takata's analysis of the sequential method in a straightforward way to the DL-algorithm and get the same lower bound.

Algorithm 2 The DL-Algorithm

$Tr(\mathcal{H}_1) \leftarrow \{\{v\} : v \in e_1\}$
for $i \leftarrow 2, \ldots, m$ **do**
  $T_{guaranteed} \leftarrow \{t \in Tr(\mathcal{H}_{i-1}) : t \cap e_i \neq \emptyset\}$
  $e_i^{covered} \leftarrow \{v \in e_i : \{v\} \in T_{guaranteed}\}$
  $Tr(\mathcal{H}_{i-1})' \leftarrow Tr(\mathcal{H}_{i-1}) \setminus T_{guaranteed}$
  $e_i' \leftarrow e_i \setminus e_i^{covered}$
  for all $t' \in Tr(\mathcal{H}_{i-1})'$ in increasing cardinality order **do**
    for all $v \in e_i'$ **do**
      if $t' \cup \{v\}$ is not superset of any $t \in T_{guaranteed}$ **then**
        $T_{guaranteed} \leftarrow T_{guaranteed} \cup \{t' \cup \{v\}\}$
      end if
    end for
  end for
  $Tr(\mathcal{H}_i) \leftarrow T_{guaranteed}$
end for
output $Tr(\mathcal{H}_m)$

Theorem 3.1 *The DL-algorithm is not output-polynomial. Its running time is at least $n^{\Omega(\log \log n)}$, where n denotes the size of the input and output.*

Nevertheless, for hypergraphs with only a few edges of small size the DL-algorithm has been shown experimentally to perform well [DL05]. This property is exploited by the BMR-algorithm [BMR03] (cf. Algorithm 3 for the listing) as it uses the DL-algorithm as a subroutine that computes all minimal transversals for small hypergraphs (line 16 of the listing). The BMR-algorithm on input $\mathcal{H}$ is invoked by the top-level call with the set E of edges of $\mathcal{H}$ and an empty set V_{part}. The global variable T is initially empty.

Before calling the DL-algorithm, the BMR-algorithm ensures that the hypergraph has only few edges of small size. If this is not yet the case, the BMR-algorithm reduces the number of edges and their size by recursively deriving smaller hypergraphs from $\mathcal{H}$ (line 14). This is achieved by partitioning the edge set and masking out vertices that are more frequent than the actual partitioning vertex v_i (lines 5 to 10). If the hypergraph is small, the DL-algorithm computes all minimal transversals (line 16). These transversals are expanded by the current partitioning vertices V_{part} (line 17) since the result is a transversal of $\mathcal{H}$. The global variable T contains all the minimal transversals of the hypergraph $\mathcal{H}$ when the algorithm stops.

A bottleneck for the running time of the BMR-algorithm is that possibly many of the recursively computed transversals—the set T' in the listing—actually are not minimal for the input hypergraph $\mathcal{H}$. We concentrate on this issue and construct a family $\mathcal{G}_i'$ of hypergraphs for which the BMR-algorithm computes too many such non-minimal transversals to run in output-polynomial time.

Algorithm 3 The BMR-Algorithm

Input: a simple hypergraph, given by the set E of its hyperedges, and a set V_{part} of partitioning vertices

$V \leftarrow$ set of all vertices in E
order vertices by increasing number of occurrences in $E \Rightarrow [v_1, \dots, v_k]$
for $i \leftarrow 1, \dots, k$ **do**
    $E_{part} \leftarrow \emptyset$
    $V \leftarrow V \setminus \{v_i\}$
    for all $e \in E$ **do**
        if $v_i \notin e$ **then**
            $E_{part} \leftarrow \min(E_{part} \cup \{e \setminus V\})$
        end if
    end for
    $V_{part} \leftarrow V_{part} \cup \{v_i\}$
    $a \leftarrow$ average edge cardinality of E_{part} multiplied by $|E_{part}|$
    if $|E_{part}| \geq 2$ and $a \geq 50$ **then**
        recursively call the BMR-algorithm on input E_{part}, V_{part}
    else
        compute $Tr(E_{part})$ via the DL-algorithm
        $T' \leftarrow Tr(E_{part}) \vee \{V_{part}\}$
        $T \leftarrow \min(T \cup T')$
    end if
    $V_{part} \leftarrow V_{part} \setminus \{v_i\}$
end for
return T

Let $\mathcal{G}'(i) = \{e_i, f_i\}$, where $e_i = \{v_{i^2-i+1}, \dots, v_{i^2}\}$ and $f_i = \{v_{i^2+1}, \dots, v_{i^2+i}\}$. We inductively define

$$\mathcal{G}'_1 = \{\{v_1\}, \{v_2\}\}, \text{ and}$$
$$\mathcal{G}'_i = (\mathcal{G}'_{i-1} \cup \{\{w_i\}\}) \vee \mathcal{G}'(i), \text{ for } i \geq 2.$$

Note that $\mathcal{G}'_{i-1}$, $\{\{w_i\}\}$ and $\mathcal{G}'(i)$ are pairwise vertex-disjoint simple hypergraphs for $i \geq 2$. To calculate the size of $\mathcal{G}'_i$ and of $Tr(\mathcal{G}'_i)$ we have to solve the recurrences $|\mathcal{G}'_i| = 2 \cdot |\mathcal{G}'_{i-1}| + 2$ and $|Tr(\mathcal{G}'_i)| = |Tr(\mathcal{G}'_{i-1})| + i^2$. With the initial conditions $|\mathcal{G}'_1| = 2$ and $|Tr(\mathcal{G}'_1)| = 1$ we obtain

$$|\mathcal{G}'_i| = 2^{i+1} - 2 \quad \text{and} \quad |Tr(\mathcal{G}'_i)| = \frac{2i^3 + 3i^2 + i}{6}$$

by iteration. As for the number $|V_{\mathcal{G}'_i}|$ of vertices of $\mathcal{G}'_i$, we have $|V_{\mathcal{G}'_i}| = i^2 + 2i - 1$.

The BMR-algorithm iteratively partitions the input hypergraph to obtain smaller hypergraphs where the transversal generation is feasible. The partitioning depends on the vertex frequencies. Hence, we first have to analyze the frequencies of the vertices in $\mathcal{G}'_i$.

Lemma 3.2 *For $i \geq 2$ let $\#_v(i,j)$ and $\#_w(i,j)$ respectively denote the number of occurrences of vertices v_j and w_j in $\mathcal{G}'_i$. Then*

$$\begin{aligned}
&\#_w(i,j) = 0, \quad \textit{for } j > i,\\
&\#_v(i,j) = 0, \quad \textit{for } j > i^2 + i,\\
&\#_w(i,2) = \#_v(i,1) = \#_v(i,2),\\
&\#_w(i,j) > \#_w(i,j+1), \quad \textit{for } 2 \leq j < i,\\
&\#_v(i,j) = \#_v(i,k), \quad \textit{for } l^2 - l + 1 \leq j \leq k \leq l^2 + l, \ \textit{with } 1 \leq l \leq i,\\
&\#_v(i,j) < \#_v(i,k), \quad \textit{for } 1 \leq j < l^2 - l + 1 \leq k \leq l^2 + l, \ \textit{with } 2 \leq l \leq i.
\end{aligned}$$

PROOF.

(1) We have the obvious equations

$$\#_w(i,j) = 0, \text{ for } j > i, \qquad \text{and} \qquad \#_v(i,j) = 0, \text{ for } j > i^2 + i,$$

as neither w_j, for $j > i$, nor v_j, for $j > i^2 + 1$, are vertices of $\mathcal{G}'_i$.

(2) Another easy case is

$$\#_w(i,2) = \#_v(i,1) = \#_v(i,2), \text{ for } i \geq 2,$$

as it is not difficult to show that all three values are equal to 2^{i-1}.

(3) The next inequality

$$\#_w(i,j) > \#_w(i,j+1), \text{ for } 2 \leq j < i,$$

also is straightforward as we have $\#_w(i,j) = 2^{i-j+1}$ for $2 \leq j \leq i$.

(4) We next consider

$$\#_v(i,j) = \#_v(i,k), \text{ for } l^2 - l + 1 \leq j \leq k \leq l^2 + l, \text{ with } 1 \leq l \leq i.$$

The proof is by induction on i. Let $i = 2$. In this case, from the definition of $\mathcal{G}'_2$, we have $2 = \#_v(2,1) = \#_v(2,2)$, and $3 = \#_v(2,3) = \#_v(2,4) = \#_v(2,5) = \#_v(2,6)$. So let the equation hold for $i = m - 1$. We will show it for $i = m$. From the definition of $\mathcal{G}'_m$ we have $\#_v(m,j) = 2 \cdot \#_v(m-1,j)$ for every $j < m^2 - m + 1$. Hence, for $2 \leq l < m$ the equation follows from our assumption.

From the definition of $\mathcal{G}'_m$ we also have $\#_v(m,j) = |\mathcal{G}'_{m-1}| + 1$ for $m^2 - m + 1 \leq j \leq m^2 + m$. Hence, the equation follows for $l = m$.

(5) The last inequality to prove is

$$\#_v(i,j) < \#_v(i,k) \text{ for } 1 \le j < l^2 - l + 1 \le k \le l^2 + l, \text{ with } 2 \le l \le i.$$

Again, the induction is on i. Let $i = 2$. From the definition of $\mathcal{G}'_2$ we have $\#_v(2,1) = \#_v(2,2) = 2 < 3 = \#_v(2,3) = \#_v(2,4) = \#_v(2,5) = \#_v(2,6)$. Let us assume that the inequality holds for $i = m - 1$. We have to prove it for $i = m$.

First, we consider the case $l < m$. From the definition of $\mathcal{G}'_m$ we have $\#_v(m,j) = 2 \cdot \#_v(m-1,j)$ for all $j < l^2 - l + 1$ and $\#_v(m,k) = 2 \cdot \#_v(m-1,k)$ for all $l^2 - l + 1 \le k \le l^2 + l$. Together with the assumption this yields the inequality for the case $l < m$.

Secondly, we have to examine the case $l = m$. Let us consider the vertex v_{m^2-m}, the vertex from $\mathcal{G}'_{m-1}$ in $\mathcal{G}'_m$ with the largest index. From the case $l < m$ we known that v_{m^2-m} is one of the most frequent vertices of $\mathcal{G}'_{m-1}$ in $\mathcal{G}'_m$. To complete the proof it suffices to show $\#_v(m, m^2 - m) < \#_v(m, m^2 - m + 1)$ as we know from Equation (4) and the already established "$l < m$"-case. From the definition of $\mathcal{G}'_m$ and $\mathcal{G}'_{m-1}$ we have

$$\begin{aligned} \#_v(m, m^2 - m) &= 2 \cdot (|\mathcal{G}'_{m-2}| + 1), \text{ and} \\ \#_v(m, m^2 - m + 1) &= |\mathcal{G}'_{m-1}| + 1. \end{aligned}$$

With $|\mathcal{G}'_i| = 2^{i+1} - 2$ this gives

$$\begin{aligned} \#_v(m, m^2 - m) &= 2^m - 2, \text{ and} \\ \#_v(m, m^2 - m + 1) &= 2^m - 1. \end{aligned}$$

Hence, we have $\#_v(m, m^2 - m) < \#_v(m, m^2 - m + 1)$.

Thus, the proof of Lemma 3.2 is completed. □

From Lemma 3.2 it follows that the vertices from $\mathcal{G}'(i)$ are the last vertices in the vertex ordering computed by the BMR-algorithm on input $\mathcal{G}'_i$. This is crucial for the next step of our analysis in which we examine the recursive calls produced by the BMR-algorithm on input $\mathcal{G}'_i$.

Lemma 3.3 *For $i \ge 4$, the BMR-algorithm on input $\mathcal{G}'_i$ recursively calls the BMR-algorithm at least $2i$ times with a modified $\mathcal{G}'_{i-1} \cup \{\{w_i\}\}$ as input. Here, modified means that all edges of $\mathcal{G}'_{i-1} \cup \{\{w_i\}\}$ may additionally include at most half of the vertices of $\mathcal{G}'(i)$.*

PROOF. We only examine the last $2i$ vertices processed by the BMR-algorithm. From Lemma 3.2 we know that these are exactly the vertices from

$\mathcal{G}'(i)$—contained in the edges e_i and f_i. Let $v'_1, v'_2, \ldots, v'_{2i}$ be any ordering of these vertices. We consider the BMR-algorithm on that ordering.

Let the j-th vertex v'_j, $1 \le j \le 2i$, from the above ordering be the current partitioning vertex (line 3 of the BMR-algorithm). After partitioning (lines 5 to 10), the remaining hypergraph has the form

$$(\mathcal{G}'_{i-1} \cup \{\{w_i\}\}) \vee \{\{v'_1, \ldots, v'_{j-1}\} \cap x_i\},$$

where $x_i = f_i$ if $v'_j \in e_i$, and $x_i = e_i$ if $v'_j \in f_i$. Hence, the remaining hypergraph always is a $\mathcal{G}'_{i-1} \cup \{\{w_i\}\}$ with at most half of the vertices from $\mathcal{G}'(i)$ in every edge.

Altogether, for each of the last $2i$ vertices the minimal transversals of a modified $\mathcal{G}'_{i-1} \cup \{\{w_i\}\}$ have to be computed. Note that a modified $\mathcal{G}'_3 \cup \{\{w_4\}\}$ has 15 edges of average size at least 5.4 and thus $a \ge 81$ (line 12). Hence, for $i \ge 4$ the last $2i$ vertices invoke recursive calls of the BMR-algorithm with a modified $\mathcal{G}'_{i-1} \cup \{\{w_i\}\}$ as input. □

With Lemma 3.3 at hand we can analyze the number of non-minimal transversals computed by the BMR-algorithm.

Lemma 3.4 *Let $i \ge 4$. For the number $\eta(i)$ of non-minimal transversals computed during a run of the BMR-algorithm on input $\mathcal{G}'_i$ we have $\eta(i) \ge 2^{i-1} \cdot i!$.*

PROOF. From Lemma 3.3 it follows that there are $2i$ recursive calls with a modified $\mathcal{G}'_{i-1} \cup \{\{w_i\}\}$ as input. Such a recursive call produces at least all of the minimal and some non-minimal transversals of $\mathcal{G}'_{i-1} \cup \{\{w_i\}\}$ augmented by the current partitioning vertex as transversals for $\mathcal{G}'_i$. But since at least the partitioning vertex is dispensable in these transversals, none of them is minimal for $\mathcal{G}'_i$ and thus will not be part of the final output. There are at least $\eta(i-1) + |Tr(\mathcal{G}'_{i-1})|$ such non-minimal transversals per recursive call. Hence, we have to solve the recurrence

$$\begin{aligned}\eta(i) &\ge 2i \cdot (\eta(i-1) + |Tr(\mathcal{G}'_{i-1})|) \\ &\ge 2i \cdot \eta(i-1).\end{aligned}$$

As for the initial condition we have the following.

Claim 3.5 $\eta(3) = 34$.

PROOF. We have

$$\mathcal{G}'_1 = \{\{v_1\}, \{v_2\}\}, \qquad \mathcal{G}'_2 = \{\{v_1, v_3, v_4\}, \{v_2, v_3, v_4\}, \{w_2, v_3, v_4\}, \{v_1, v_5, v_6\}, \{v_2, v_5, v_6\}, \{w_2, v_5, v_6\}\},$$

$$\begin{aligned}\mathcal{G}'_3 = \{&\{v_1, v_3, v_4, v_7, v_8, v_9\}, \{v_2, v_3, v_4, v_7, v_8, v_9\}, \{w_2, v_3, v_4, v_7, v_8, v_9\},\\ &\{v_1, v_5, v_6, v_7, v_8, v_9\}, \{v_2, v_5, v_6, v_7, v_8, v_9\}, \{w_2, v_5, v_6, v_7, v_8, v_9\},\\ &\{w_3, v_7, v_8, v_9\},\\ &\{v_1, v_3, v_4, v_{10}, v_{11}, v_{12}\}, \{v_2, v_3, v_4, v_{10}, v_{11}, v_{12}\},\\ &\{w_2, v_3, v_4, v_{10}, v_{11}, v_{12}\}, \{v_1, v_5, v_6, v_{10}, v_{11}, v_{12}\},\\ &\{v_2, v_5, v_6, v_{10}, v_{11}, v_{12}\}, \{w_2, v_5, v_6, v_{10}, v_{11}, v_{12}\},\\ &\{w_3, v_{10}, v_{11}, v_{12}\}\}.\end{aligned}$$

We examine the BMR-algorithm with $\mathcal{G}'_3$ as input. Without loss of generality we assume that the order in which the BMR-algorithm processes the vertices is $w_3, w_2, v_1, v_2, v_3, \ldots, v_{12}$. When using w_3, w_2, or v_1 as partitioning vertex, nothing happens since the resulting hypergraph is empty.

The next partitioning vertex is v_2 and there remains the hypergraph with the three edges $\{w_3\}$, $\{w_2\}$, and $\{v_1\}$. The DL-algorithm is invoked and outputs one minimal transversal, which is augmented by v_2. The resulting transversal is minimal for $\mathcal{G}'_3$.

When using v_3 or v_4 as partitioning vertex, there remains the hypergraph with the four edges $\{w_3\}$, $\{w_2\}$, $\{v_1\}$, and $\{v_2\}$. The DL-algorithm computes the minimal transversal of this hypergraph, which is augmented by v_3 and respectively v_4. Obviously, the resulting transversals are not minimal since they contain the minimal transversal $\{w_3, w_2, v_1, v_2\}$. Hence, the BMR-algorithm has computed two non-minimal transversals of $\mathcal{G}'_3$.

When using v_5 or v_6 as partitioning vertex, there remains the hypergraph with the four edges $\{w_3\}$, $\{w_2, v_3, v_4\}$, $\{v_1, v_3, v_4\}$, and $\{v_2, v_3, v_4\}$. The DL-algorithm computes the three minimal transversals $\{w_3, v_3\}$, $\{w_3, v_4\}$, and $\{w_3, w_2, v_1, v_2\}$ and augments them by v_5 and respectively v_6. This yields four minimal transversals of $\mathcal{G}'_3$ and another two not minimal transversals of $\mathcal{G}'_3$.

When using v_7, v_8, or v_9 as partitioning vertex, there remains a $\mathcal{G}'_2 \cup \{\{w_3\}\}$. Each time, the DL-algorithm is invoked to compute all five minimal transversals of $\mathcal{G}'_2 \cup \{\{w_3\}\}$. Each such computed minimal transversal of $\mathcal{G}'_2 \cup \{\{w_3\}\}$ is augmented by the current partitioning vertex. The resulting transversal is not minimal for $\mathcal{G}'_3$ since already the minimal transversals of $\mathcal{G}'_2 \cup \{\{w_3\}\}$ are minimal for $\mathcal{G}'_3$. Hence, the algorithm produces 15 non-minimal transversals for the vertices v_7, v_8, and v_9.

As for the vertices v_{10}, v_{11}, and v_{12} there remains a $(\mathcal{G}'_2 \cup \{\{w_3\}\}) \vee \{v_7, v_8, v_9\}$ after partitioning. For each such modified $\mathcal{G}'_2 \cup \{\{w_3\}\}$ the DL-algorithm as a subroutine is invoked to compute the minimal transversals since $a = 40 < 50$

(line 12 of the BMR-algorithm). For each such call the DL-algorithm produces all five minimal transversals of $\mathcal{G}_2' \cup \{\{w_3\}\}$ plus the three minimal transversals $\{v_7\}$, $\{v_8\}$, $\{v_9\}$. Each such computed transversal is augmented by the current partitioning vertex. This yields nine minimal transversals of $\mathcal{G}_3'$ and another 15 non-minimal transversals.

Altogether, the BMR-algorithm with input $\mathcal{G}_3'$ computes 34 non-minimal transversals. This yields $\eta(3) = 34$. □

Hence, $\eta(3) \geq 2^2 \cdot 3!$ and we get $\eta(i) \geq 2^{i-1} \cdot i!$ by iteration. □

Putting all pieces together we are able to give a superpolynomial lower bound on the running time of the BMR-algorithm.

Theorem 3.6 *The BMR-algorithm is not output-polynomial. Its running time is at least $n^{\Omega(\log\log n)}$, where n denotes the size of the input and output.*

PROOF. We consider the BMR-algorithm on input $\mathcal{G}_i'$. By $m_i = |V_{\mathcal{G}_i'}| \cdot (|\mathcal{G}_i'| + |Tr(\mathcal{G}_i')|)$ we denote an upper bound on the size of the input and output. For $i \geq 22$ we have

$$m_i = (i^2 + 2i - 1) \cdot \left(2^{i+1} - 2 + \frac{2i^3 + 3i^2 + i}{6}\right) \leq 2^{3i}.$$

The running time of the BMR-algorithm on input $\mathcal{G}_i'$ is at least $\eta(i)$, the number of non-minimal transversals computed. Since $i \geq \frac{\log m_i}{3}$ and $i! \geq \left(\frac{i}{e}\right)^i$, we get $\eta(i) \geq 2^{i-1} \cdot i! = m_i^{\Omega(\log\log m_i)}$. □

4 The Algorithm of Kavvadias and Stavropoulos

A first drawback of the sequential method or the BMR-algorithm that Kavvadias and Stavropoulos [KS05] observe is the memory requirement. Since newly computed transversals have to be checked for minimality against the previously computed minimal transversals, all the previously generated minimal transversals have to be stored. The KS-algorithm tries to overcome this potentially exponential memory requirement by two techniques. The first is to combine vertices that belong exactly to the same hyperedges.

Definition 4.1 (generalized vertex, [KS05]) *Let $\mathcal{H}$ be a hypergraph with vertex set V. The set $X \subseteq V$ is a generalized vertex of $\mathcal{H}$ if all vertices in X belong to exactly the same hyperedges of $\mathcal{H}$.*

A transversal possibly containing generalized vertices will be referred to as *generalized* transversal. While adding edge e_i, and hence generating the minimal generalized transversals of $\mathcal{H}_i$ out of the minimal generalized transversals of $\mathcal{H}_{i-1}$, the generalized vertices have to be updated according to e_i. Kavvadias and Stavropoulos characterize the following three types of generalized vertices X of a minimal generalized transversal t of $\mathcal{H}_{i-1}$.

- type α: $X \cap e_i = \emptyset$. Hence, X is a generalized vertex of $\mathcal{H}_i$.
- type β: $X \subset e_i$. Hence, X is a generalized vertex of $\mathcal{H}_i$.
- type γ: $X \cap e_i \neq \emptyset$ and $X \not\subset e_i$. Here, X is divided into $X_1 = X \setminus (X \cap e_i)$ and $X_2 = X \cap e_i$. Both X_1 and X_2 are generalized vertices of $\mathcal{H}_i$.

Let $\kappa_\alpha(t,i)$, $\kappa_\beta(t,i)$, and $\kappa_\gamma(t,i)$ denote the number of generalized vertices of type α, β, and γ in t according to e_i. When edge e_i is added, the minimal generalized transversal t of $\mathcal{H}_{i-1}$ has to be split into $2^{\kappa_\gamma(t,i)}$ generalized transversals of $\mathcal{H}_{i-1}$—the so-called *offsprings* of t—since all combinations of newly generalized vertices have to be generated. If $\kappa_\beta(t,i) \neq 0$, all these newly generated offsprings are also minimal transversals of $\mathcal{H}_i$. But if $\kappa_\beta(t,i) = 0$, there is a special offspring t_0 of t that contains all the X_1-parts of the γ-type generalized nodes of t. Hence, $t_0 \cap e_i = \emptyset$ and t_0 has to be augmented by a vertex from e_i to be a transversal of $\mathcal{H}_i$. All the other offsprings of t already are minimal transversals of $\mathcal{H}_i$ since they contain at least one X_2-part of a generalized vertex from t.

The second technique to overcome the potentially exponential memory requirement is based on the observation that the sequential method is a form of breadth-first search through a "tree" of minimal transversals. At the ith-level of the "tree" the nodes are the minimal transversals of the partial hypergraph $\mathcal{H}_i$. The descendants of a minimal transversal t at level i are the minimal transversals of $\mathcal{H}_{i+1}$ that include t. Note that, since a node at level $i+1$ may have several ancestors at level i, the structure is not really a tree but very tree-like. The bottom level consists of the minimal transversals of $\mathcal{H}$. When cycling through this "tree" breadth-first, one has to wait very long for the first minimal transversal to be output and some nodes are visited several times because they have more than one ancestor. To overcome the long time that may pass till the first minimal transversal is output, the KS-algorithm uses a depth-first strategy. And to really cycle through a tree and not a tree-like structure with some cycles, Kavvadias and Stavropoulos introduce the notion of so-called appropriate vertices.

Definition 4.2 (appropriate vertex, [KS05]) *Let $\mathcal{H} = \{e_1, \ldots, e_m\}$ be a hypergraph with vertex set V and let t be a minimal transversal of the partial hypergraph $\mathcal{H}_i$ of $\mathcal{H}$. A generalized vertex $v \subseteq V \setminus t$ at level i is an appropriate vertex for t if no other vertex in $t \cup \{v\}$ except v can be removed and the remaining set still be a transversal of $\mathcal{H}_i$. The set appr(t,e) contains all* *appropriate vertices for t in edge e.*

Note that the special offspring t_0 of a minimal generalized transversal t of $\mathcal{H}_{i-1}$ has to be augmented by a vertex from $appr(t, e_i)$ only. All the other vertices from e_i can be skipped. Expanding only with appropriate vertices ensures that no non-minimal transversals are generated and avoids regenerations. Another advantage is that the previously described transversal "tree" structure becomes a real tree (cf. the original paper [KS05] for more details).

All the described techniques—generalized vertices, depth-first strategy, appropriate vertices—together with the main idea of the sequential method—processing the edges one after the other—are used in the KS-algorithm (cf. Algorithm 4 for the listing) and yield a space requirement (the size of the output does not count) that is only polynomial in the input size [KS05].

After computing a first transversal (only one since it consists of a generalized vertex), the recursive AddNextHyperedge procedure is called (note that t' is a global variable). Due to the usage of generalized vertices the expansion of t is divided into two parts according to the presence (line 18 of the listing) or absence (line 25) of a generalized vertex of type β in t. If the minimal transversal t of $\mathcal{H}_{i-1}$ contains a generalized vertex of type β, all its offsprings intersect e_i and hence are minimal transversals of $\mathcal{H}_i$ (lines 19 to 24). If t does not contain a type β vertex, all its offsprings except t_0 intersect e_i and hence are minimal for $\mathcal{H}_i$ (lines 26 to 28). The offspring t_0 has to be augmented by every appropriate vertex (line 30).

The effect as shown by Kavvadias and Stavropoulos is that a newly generated transversal is minimal and that regenerations are avoided [KS05]. Since the KS-algorithm uses a depth-first strategy, it does not have to store all the minimal transversals of the subhypergraph $\mathcal{H}_{i-1}$ to compute the minimal transversals of $\mathcal{H}_i$. This yields a space requirement of the KS-algorithm that is polynomial in the input size $|\mathcal{H}|$ [KS05].

As for the running time, the KS-algorithm is experimentally shown [KS05] to be competitive to the sequential method, the BMR-algorithm, and an implementation of Algorithm A of Fredman and Khachiyan [BEGK03,FK96]. We will show that the KS-algorithm is not output-polynomial.

First, we note that there are situations in which the KS-algorithm cannot find an appropriate vertex. Consider for example the hypergraph

$$\mathcal{H} = \{\{v_1, v_5\}, \{v_2, v_5\}, \{v_3, v_6\}, \{v_4, v_6\}, \{v_5, v_6\}\}.$$

Having processed all but the last edge, there are no generalized vertices left. We concentrate on the path down the transversal tree that corresponds to choosing v_1, v_2, v_3, and v_4. The intermediate transversal is $t = \{v_1, v_2, v_3, v_4\}$.

Algorithm 4 The KS-Algorithm

```
express e_1 as a set of one generalized vertex
compute the transversal t = Tr(e_1)
AddNextHyperedge(t, e_2)

procedure AddNextHyperedge(t, e_i)
    update the set of generalized vertices
    express t and e_i as sets of generalized vertices of level i
    l ← 1
    while GenerateNextTransversal(t, l) do
        if e_i is the last hyperedge then
            output t' without using generalized vertices
        else
            AddNextHyperedge(t', e_{i+1})
            l ← l + 1
        end if
    end while
end procedure

function GenerateNextTransversal(t, l)
    if κ_β(t, i) ≠ 0 then
        if l ≤ 2^{κ_γ(t,i)} then
            t' ← the l-th offspring of t
            return true
        else
            return false
        end if
    else if κ_β(t, i) = 0 then
        if l ≤ 2^{κ_γ(t,i)} − 1 then
            t' ← the l-th offspring of t except t_0
            return true
        else if 2^{κ_γ(t,i)} ≤ l ≤ 2^{κ_γ(t,i)} − 1 + |appr(t, e_i)| then
            t' = t_0 augmented by the (l − 2^{κ_γ(t,i)} + 1)-th vertex of appr(t, e_i)
            return true
        end if
    else
        return false
    end if
end function
```

The only edge left is $\{v_5, v_6\}$. But the KS-algorithm cannot find an appropriate vertex in this edge for t. Hence, there are dead ends in the tree, namely leaves that do not contain a minimal transversal of the input $\mathcal{H}$. The next step is to find hypergraphs with too many such dead ends.

Lemma 4.3 *For $i \geq 3$, the number of dead ends the KS-algorithm has to visit for any of Takata's hypergraphs $\mathcal{G}_i$ as input is at least $2^{(i-2)2^i+1}$, independent of the edge ordering.*

PROOF. Consider the hypergraph family $\mathcal{G}_i$ of Takata defined in Section 2. First note that when the KS-algorithm adds the last edge of $\mathcal{G}_i$, there are no *proper* generalized vertices left (generalized vertices that are not singleton sets). We want to argue that the same already holds for the penultimate step, hence, that $\mathcal{G}_i \setminus \{e\}$ has no proper generalized vertex, for any edge $e \in \mathcal{G}_i$. Assume otherwise that after processing all of $\mathcal{G}_i \setminus \{e\}$'s edges there remains a proper generalized vertex $X \subseteq V$. As $\mathcal{G}_i$ has no proper generalized vertices, we have $X \subseteq e$. Let e be composed of the $\mathcal{A}$ and $\mathcal{C}$ component in $\mathcal{G}_i$'s definition (the argumentation is analogous for the other cases) and consider two different vertices $v, u \in X$. If both v and u are vertices in the $\mathcal{A}$ component we have a contradiction as already $\mathcal{A}$ contains an edge f that contains v but not u. This edge appears in $|\mathcal{C}| + |\mathcal{D}|$ edges of $\mathcal{G}_i \setminus \{e\}$. Hence, not both v and u can be vertices in X as they would have been split according to the f copies in prior steps of the KS-algorithm's run (an analogous argumentation shows that not both are in $\mathcal{C}$).

The remaining possibility (minus symmetry) is that v is from $\mathcal{A}$ and u is from $\mathcal{C}$. But note that $\mathcal{G}_i \setminus \{e\}$ contains an edge f with $v \in f$ but f is composed of the $\mathcal{A}$ and $\mathcal{D}$ part of $\mathcal{G}_i$. Again, this shows that not both v and u can be vertices in X as they would have been split in a prior step.

Altogether, we now know that before processing the last edge, there cannot be proper generalized vertices in $\mathcal{G}_i \setminus \{e\}$. From Lemma 2.2 it follows that, whatever ordering of the edges is chosen, there are at least $2^{(i-2)2^i+2}$ nodes in the penultimate level of the transversal tree described above Definition 4.2. The bottom level of the tree obviously contains $|Tr(\mathcal{G}_i)|$ many nodes—one for each minimal transversal. Since $|Tr(\mathcal{G}_i)| = 2^{2^i-1}$ (cf. Lemma 2.2), there is a decrease in the number of nodes from the penultimate level to the bottom level for $i \geq 3$. This decrease can only be caused by dead ends in the penultimate level. Hence, for $i \geq 3$ there are at least $2^{(i-2)2^i+2} - 2^{2^i-1} \geq 2^{(i-2)2^i+1}$ many dead ends in the penultimate level. □

Using Lemma 4.3 we can show that the KS-algorithm is not output-polynomial.

Theorem 4.4 *The KS-algorithm is not output-polynomial. Its running time is at least $n^{\Omega(\log \log n)}$, where n denotes the size of the input and output.*

PROOF. We consider the KS-algorithm on input $\mathcal{G}_i$. By $m_i = |V_{\mathcal{G}_i}| \cdot (|\mathcal{G}_i| + |Tr(\mathcal{G}_i)|)$ we denote an upper bound on the size of $\mathcal{G}_i$ and $Tr(\mathcal{G}_i)$. From Lemma 2.2 we have $m_i = 4^i \cdot (2^{2(2^i-1)} + 2^{2^i-1})$, which results in $m_i \leq 2^{2^{i+2}}$.

Let $\hat{\eta}(i)$ denote the number of dead end situations visited by the KS-algorithm on input $\mathcal{G}_i$. The time, the KS-algorithm needs to compute $Tr(\mathcal{G}_i)$, is at least the number of dead end situations visited. Since the KS-algorithm visits the transversal tree depth-first, it visits all the dead end situations in the penultimate level of the tree. With Lemma 4.3 we have $\hat{\eta}(i) \geq 2^{(i-2)2^i+1}$ for $i \geq 3$. Thus, to analyze the running time we will show that $\hat{\eta}(i)$ is superpolynomial in m_i. It suffices to show that $2^{(i-2)2^i} > (2^{2^{i+2}})^c$, for any constant c. This is equivalent to $i - 2 > 4c$, for any constant c. Since this obviously holds for large enough i, we have proven that $\hat{\eta}(i)$ is superpolynomial in m_i, namely $\hat{\eta}(i) = m_i^{\Omega(\log \log m_i)}$. □

5 Concluding Remarks

We have proven superpolynomial lower bounds for the DL-, the BMR-, and the KS-algorithm in terms of the size of the input and output. Thus, like the underlying sequential method, these three algorithms are not output-polynomial.

We are not aware of any other nontrivial lower bounds for algorithms generating the transversal hypergraph although we suppose that none of the known algorithms is output-polynomial. Extending the existing lower bounds to other algorithms seems to be not that straightforward.

Consider for instance the multiplication method suggested by Takata [Tak07]. Very recently Elbassioni proved a quasi-polynomial upper bound on the running time [Elb06]. But giving a superpolynomial lower bound for the multiplication method requires the construction of new hypergraphs. Takata's hypergraphs $\mathcal{G}_i$ and our hypergraphs $\mathcal{G}'_i$ are solved too fast by the multiplication method.

There are also no nontrivial lower bounds known for Algorithms A and B of Fredman and Khachiyan [FK96]. Though Gurvich and Khachiyan [GK97] note that it should be possible to give a superpolynomial lower bound for Algorithm A using hypergraphs very similar to the $\mathcal{G}_i$, the proof is still open. Giving a lower bound for Algorithm B—considered to be the fastest known transversal hypergraph algorithm—seems to be even more involved.

Acknowledgments

I thank Martin Mundhenk and the anonymous referees for their valuable comments and suggestions that helped to improve the paper.

References

[BEGK03] Endre Boros, Khaled M. Elbassioni, Vladimir Gurvich, and Leonid Khachiyan. Extending the Balas-Yu bounds on the number of maximal independent sets in graphs to hypergraphs and lattices. *Mathematical Programming*, 98(1-3):355–368, 2003.

[BEM08] Endre Boros, Khaled M. Elbassioni, and Kazuhisa Makino. On Berge multiplication for monotone Boolean dualization. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part I: Tack A: Algorithms, Automata, Complexity, and Games*, volume 5125 of *Lecture Notes in Computer Science*, pages 48–59. Springer, 2008.

[Ber89] Claude Berge. *Hypergraphs*, volume 45 of *North-Holland Mathematical Library*. North-Holland, 1989.

[BMR03] James Bailey, Thomas Manoukian, and Kotagiri Ramamohanarao. A fast algorithm for computing hypergraph transversals and its application in mining emerging patterns. In *Proceedings of the 3rd IEEE International Conference on Data Mining (ICDM 2003), 19-22 December 2003, Melbourne, Florida, USA*, pages 485–488. IEEE Computer Society, 2003.

[Dam06] Peter Damaschke. Parameterized enumeration, transversals, and imperfect phylogeny reconstruction. *Theoretical Computer Science*, 351(3):337–350, 2006.

[DHSW03] Jean-Guillaume Dumas, Frank Heckenbach, B. David Saunders, and Volkmar Welker. Computing simplicial homology based on efficient Smith Normal Form algorithms. In Michael Joswig and Nobuki Takayama, editors, *Algebra, Geometry, and Software Systems [outcome of a Dagstuhl seminar]*, pages 177–206. Springer, 2003.

[DL05] Guozhu Dong and Jinyan Li. Mining border descriptions of emerging patterns from dataset pairs. *Knowledge and Information Systems*, 8(2):178–202, 2005.

[EG95] Thomas Eiter and Georg Gottlob. Identifying the minimal transversals of a hypergraph and related problems. *SIAM Journal on Computing*, 24(6):1278–1304, 1995.

[EG02] Thomas Eiter and Georg Gottlob. Hypergraph transversal computation and related problems in logic and AI. In Sergio Flesca, Sergio Greco, Nicola Leone, and Giovambattista Ianni, editors, *Logics in Artificial Intelligence, European Conference, JELIA 2002, Cosenza, Italy, September, 23-26, Proceedings*, volume 2424 of *Lecture Notes in Computer Science*, pages 549–564. Springer, 2002.

[EGM03] Thomas Eiter, Georg Gottlob, and Kazuhisa Makino. New results on monotone dualization and generating hypergraph transversals. *SIAM Journal on Computing*, 32(2):514–537, 2003.

[Elb06] Khaled M. Elbassioni. On the complexity of the multiplication method for monotone CNF/DNF dualization. In Yossi Azar and Thomas Erlebach, editors, *Algorithms - ESA 2006, 14th Annual European Symposium, Zurich, Switzerland, September 11-13, 2006, Proceedings*, volume 4168 of *Lecture Notes in Computer Science*, pages 340–351. Springer, 2006.

[Elb08] Khaled M. Elbassioni. On the complexity of monotone dualization and generating minimal hypergraph transversals. *Discrete Applied Mathematics*, 32(2):171–187, 2008.

[FK96] Michael L. Fredman and Leonid Khachiyan. On the complexity of dualization of monotone disjunctive normal forms. *Journal of Algorithms*, 21(3):618–628, 1996.

[GB85] Hector Garcia-Molina and Daniel Barbará. How to assign votes in a distributed system. *Journal of the ACM*, 32(4):841–860, 1985.

[GK97] Vladimir Gurvich and Leonid Khachiyan. On the frequency of the most frequently occurring variable in dual monotone DNFs. *Discrete Mathematics*, 169(1-3):245–248, 1997.

[GKM+03] Dimitrios Gunopulos, Roni Khardon, Heikki Mannila, Sanjeev Saluja, Hannu Toivonen, and Ram Sewak Sharm. Discovering all most specific sentences. *ACM Transactions on Database Systems*, 28(2):140–174, 2003.

[GKMT97] Dimitrios Gunopulos, Roni Khardon, Heikki Mannila, and Hannu Toivonen. Data mining, hypergraph transversals, and machine learning. In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, May 12-14, 1997, Tucson, Arizona*, pages 209–216. ACM Press, 1997.

[Hag07] Matthias Hagen. Lower bounds for three algorithms for the transversal hypergraph generation. In Andreas Brandstädt, Dieter Kratsch, and Haiko Müller, editors, *Graph-Theoretic Concepts in Computer Science, 33rd International Workshop, WG 2007, Dornburg, Germany, June 21-23, 2007. Revised Papers*, volume 4769 of *Lecture Notes in Computer Science*, pages 316–327. Springer, 2007.

[Hag08] Matthias Hagen. *Algorithmic and Computational Complexity Issues of* Monet. PhD thesis, Friedrich-Schiller-Universität Jena, 2008. To appear.

[HKS08] Utz-Uwe Haus, Steffen Klamt, and Tamon Stephen. Computing knock-out strategies in metabolic networks. *Journal of Computational Biology*, 15(3):259–268, 2008.

[JPY88] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. On generating all maximal independent sets. *Information Processing Letters*, 27(3):119–123, 1988.

[KBE+07] Leonid Khachiyan, Endre Boros, Khaled M. Elbassioni, Vladimir Gurvich, and Kazuhisa Makino. Dual-bounded generating problems: Efficient and inefficient points for discrete probability distributions and sparse boxes for multidimensional data. *Theoretical Computer Science*, 379(3):361–376, 2007.

[KBEG08a] Leonid Khachiyan, Endre Boros, Khaled M. Elbassioni, and Vladimir Gurvich. Generating all minimal integral solutions to AND-OR systems of monotone inequalities: Conjunctions are simpler than disjunctions. *Discrete Applied Mathematics*, 156(11):2020–2034, 2008.

[KBEG08b] Leonid Khachiyan, Endre Boros, Khaled M. Elbassioni, and Vladimir Gurvich. On enumerating minimal dicuts and strongly connected subgraphs. *Algorithmica*, 50(1):159–172, 2008.

[Kha00] Leonid Khachiyan. Transversal hypergraphs and families of polyhedral cones. In *International Conference on Advances in Convex Analysis and Global Optimization, Honoring the memory of C. Carathéodory (1873-1950), June 5-9, 2000, Pythagorion, Samos, Greece, Proceedings*, pages 6–8, 2000.

[KLM07] Naouel Karam, Serge Linckels, and Christoph Meinel. Semantic composition of lecture subparts for a personalized e-learning. In Enrico Franconi, Michael Kifer, and Wolfgang May, editors, *The Semantic Web: Research and Applications, 4th European Semantic Web Conference, ESWC 2007, Innsbruck, Austria, June 3-7, 2007, Proceedings*, volume 4519 of *Lecture Notes in Computer Science*, pages 716–728. Springer, 2007.

[KP04] Dogan Kesdogan and Lexi Pimenidis. The hitting set attack on anonymity protocols. In Jessica J. Fridrich, editor, *Information Hiding, 6th International Workshop, IH 2004, Toronto, Canada, May 23-25, 2004, Revised Selected Papers*, volume 3200 of *Lecture Notes in Computer Science*, pages 326–339. Springer, 2004.

[KS03] Dimitris J. Kavvadias and Elias C. Stavropoulos. Monotone Boolean dualization is in $\mathsf{coNP}[\log^2 n]$. *Information Processing Letters*, 85(1):1–6, 2003.

[KS05] Dimitris J. Kavvadias and Elias C. Stavropoulos. An efficient algorithm for the transversal hypergraph generation. *Journal of Graph Algorithms and Applications*, 9(2):239–264, 2005.

[KSG07] Steffen Klamt, Julio Saez-Rodriguez, and Ernst Dieter Gilles. Structural and functional analysis of cellular networks with *CellNetAnalyzer*. *BMC Systems Biology*, 1(2), 2007.

[MR92] Heikki Mannila and Kari-Jouko Räihä. On the complexity of inferring functional dependencies. *Discrete Applied Mathematics*, 40(2):237–243, 1992.

[Pap97] Christos H. Papadimitriou. NP-completeness: A retrospective. In Pierpaolo Degano, Roberto Gorrieri, and Alberto Marchetti-Spaccamela, editors, *Automata, Languages and Programming, 24th International Colloquium, ICALP'97, Bologna, Italy, 7-11 July 1997, Proceedings*, volume 1256 of *Lecture Notes in Computer Science*, pages 2–6. Springer, 1997.

[SS98] Saswati Sarkar and Kumar N. Sivarajan. Hypergraph models for cellular mobile communication systems. *IEEE Transactions on Vehicular Technology*, 47(2):460–471, 1998.

[Tak07] Ken Takata. A worst-case analysis of the sequential method to list the minimal hitting sets of a hypergraph. *SIAM Journal on Discrete Mathematics*, 21(4):936–946, 2007.

[Tri08] Thu Trinh. Using transversals for discovering XML functional dependencies. In Sven Hartmann and Gabriele Kern-Isberner, editors, *Foundations of Information and Knowledge Systems, 5th International Symposium, FoIKS 2008, Pisa, Italy, February 11-15, 2008, Proceedings*, volume 4932 of *Lecture Notes in Computer Science*, pages 199–218. Springer, 2008.

[ZB06] Zhou Zhu and James Bailey. Fast discovery of interesting collections of web services. In *2006 IEEE / WIC / ACM International Conference on Web Intelligence (WI 2006), 18-22 December 2006, Hong Kong, China*, pages 152–160. IEEE Computer Society, 2006.