

Polynomial Delay and Space Discovery of Connected and Acyclic Sub-Hypergraphs in a Hypergraph

Kunihiro Wasa¹, Takeaki Uno², Kouichi Hirata³, and Hiroki Arimura¹

¹ Graduate School of IST, Hokkaido University, Sapporo, Japan
 {wasa, arim}@ist.hokudai.ac.jp

² National Institute of Informatics, Tokyo, Japan
 uno@nii.jp

³ Dept. of Artificial Intelligence, Kyushu Institute of Technology, Iizuka, Japan
 hirata@ai.kyutech.ac.jp

Abstract. In this paper, we study the problem of finding all tree-like substructure contained in a hypergraph, with potential applications to substructure mining from relational data. We employ the class of connected and Berge acyclic sub-hypergraphs as definition of tree-like substructures, which is the most restricted notion of acyclicities for hypergraphs. Then, we present an efficient depth-first algorithm that finds all connected and Berge acyclic sub-hypergraphs S in a hypergraph \mathcal{H} with m hyperedges and n vertices in $O(nm^2)$ time per solution (delay) using $O(N)$ space, where $N = ||\mathcal{H}||$ is the total input size. To achieve efficient enumeration, we use the notion of the maximum border set. This result gives the first polynomial delay and time algorithm for enumeration of connected and Berge-acyclic sub-hypergraphs. We also present an incremental enumeration algorithm that finds all solutions S in $O(\Delta MB(S)\tau(m)) = O(rd \cdot \tau(m))$ delay using $O(N)$ space and preprocessing, whose delay depends only on the difference of solutions, where S is the enumerated sub-hypergraph, $\Delta MB(S)$ is the number of newly added hyperedges to the maximum border of S , r and d are the rank and degree of \mathcal{H} , respectively, and $\tau(m) = ((\log \log m)^2 / \log \log \log m)$.

1 Introduction

In data mining, it is a well-studied problem to discover all interesting substructures of a given discrete structure under various notions of substructures. Particularly, examples of such substructure discovery are frequent itemset mining [25, 28, 29], sequence mining [1, 2, 19], trees and graph mining [3, 13, 14, 27], and kernel-like similarity computation [16] to name a few.

In this paper, we study the problem of enumerating all connected and acyclic sub-hypergraphs contained in an input hypergraph for the notions of acyclicity, called Berge-acyclicity [6], which is at the bottom of hierarchy of acyclicities given by Fagin [9]. A *hypergraph* is a pair $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ of a collection \mathcal{V} of vertices and a collection \mathcal{E} of hyperedges (See Fig. 1 for example), where a hyperedge is any finite set $e \subseteq \mathcal{V}$ of vertices. Essentially, a hypergraph is a representation of

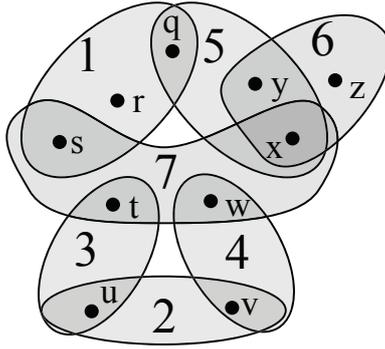


Fig. 1. An example of a hypergraph $\mathcal{H}_1 = (\mathcal{V}_1, \mathcal{E}_1)$ with the vertex set $\mathcal{V}_1 = \{p, q, r, \dots, x, y, z\}$ and the hyperedge set $\mathcal{E}_1 = \{1, 2, 3, 4, 5, 6, 7\}$.

a *set collection* \mathcal{E} consisting of groups of objects taken from a common universe \mathcal{V} . For example, the followings are examples of such set collections: transaction databases, author groups in bibliographic data, co-citation networks in social networks, and interaction graphs for genes and proteins in bioinformatics [17]. In such networks, discovery of substructures such as connected components, connected subtrees, cliques, quasi-cliques, and dense subgraphs have been extensively studied in the context of network mining [17, 21, 24].

Recently, discovery of *subtrees* in a graph, which are connected and acyclic edge subsets, attract much attention [10, 26]. By generalizing the notion of subtrees in a graph to a hypergraph, we consider discovery of the class of *connected acyclic sub-hypergraphs* appearing in a hypergraph. Particularly, among several definitions of acyclicities in a hypergraph, we employ the most restricted one, called *Berge-acyclicity* [6], where a hypergraph is Berge-acyclic if and only if it contains no cycle of hyperedges. For example, in the example of Fig. 1, the hyperedge subset $S = \{1, 3, 4, 7\}$ is connected Berge-acyclic sub-hypergraph. It is known that Berge-acyclicity locates the bottom of the degrees among α -, β -, and γ -acyclicities [9]. Now, our goal is to devise an efficient algorithm for the connected, Berge-acyclic sub-hypergraph mining problem in *polynomial delay* (time per solution) and *polynomial space*, which means that the algorithm achieves high-throughput and small memory footprint computation in large scale applications.

Main results: Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be an input hypergraph consisting of n vertices and m hyperedges. Then, we present efficient depth-first algorithms BERGEMINE that finds all connected Berge-acyclic sub-hypergraphs S contained in \mathcal{H} without duplicates in $O(m^2n)$ delay and $O(N)$ words of space (Theorem 1). To achieve polynomial delay and space complexity, our algorithm searches for all solutions in the depth-first manner on a tree-shaped search space without using any extra memory for table-lookup. This search space is designed based on a characterization of Berge-acyclic sub-hypergraphs given by us, with which we proposed an efficient and complete pruning strategy.

Next, we present the modified algorithm, called FASTBERGEMINE, that finds all S in $O(\Delta MB(S)\tau(m)) = O(rd \cdot \tau(m))$ delay using $O(N)$ space and preprocessing, where S is the enumerated sub-hypergraph, $\Delta MB(S)$ is the number of newly added hyperedges to the maximum border of S , r and d are the rank and degree of \mathcal{H} , respectively, and $\tau(m) = ((\log \log m)^2 / \log \log \log m)$. The algorithm uses incremental computation of the maximum border set. Since it has the delay that depends only on the size of each discovered subset S and its neighbors $N(S)$, it will be more efficient for the large inputs in the real world.

Related work: For the class of α -acyclic sub-hypergraphs [9], Hirata *et al.* [12] presented an efficient algorithm that finds one of the maximal connected and acyclic sub-hypergraphs in an input hypergraph in linear time in the total input size. Extending this work, Daigo and Hirata [8] presented a polynomial delay and space algorithm that finds all connected and acyclic sub-hypergraphs in an input hypergraph. For the class of Berge-acyclic sub-hypergraphs, Lovász [18] showed a polynomial time algorithm that finds one of the maximal connected and Berge-acyclic sub-hypergraphs in an input hypergraph.

As closely related work, Ferreira *et al.* [11] presented an efficient algorithm for finding all distinct subtrees of size k in an input graph in $O(k)$ time (time per solution) and space, and Wasa *et al.* [26] the improved version in constant delay when an input is a tree. However, their approaches cannot be directly applicable to our problem.

In the case that the maximum size of hyperedges, the *rank*, is restricted to two, the problem coincides to the well-known spanning tree problem for undirected graphs. For the problem, Tarjan and Read [23] first presented an $O(ns)$ time and $O(n)$ space algorithm in 1960's, where n is the number of edges in G . Recently, Shioura, Tamura, and Uno [20] presented $O(n+s)$ time and $O(n)$ space algorithm. Unfortunately, it is not easy to extend the algorithms for spanning tree enumeration to subtree enumeration.

Organization of this paper: In Sec. 2, we give basic definitions and notations on hypergraphs and our data mining problem. In Sec. 3, we present the basic depth-first algorithm BERGEMINE for the problem. In Sec. 4, we present the modified version of the algorithm, FASTBERGEMINE, using incremental computation. Finally, in Sec. 5, we conclude.

2 Preliminaries

In this section, we give the definitions and notations for hypergraphs. For the definitions not found here, please consult appropriate textbooks (e.g., [6, 7]). For a set A , we denote by $|A|$ the cardinality of A . For a collection $\mathcal{X} \subseteq 2^A$ of subsets of A , $\|\mathcal{X}\| = \sum_{S \in \mathcal{X}} |S|$ denotes the *total size* of \mathcal{X} .

2.1 Hypergraphs

Intuitively, a hyper graph is a structure defined by a set collection $\mathcal{E} \subseteq 2^{\mathcal{V}}$ over some finite domain \mathcal{V} of objects. Formally, a *hypergraph* over a set of vertices \mathcal{V} is any pair $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ consists of the following components:

- A set of *vertices* $\mathcal{V} = \mathcal{V}(\mathcal{H}) = \{1, \dots, n\}$, $n \geq 0$, and
- A set of *hyperedges* $\mathcal{E} = \mathcal{E}(\mathcal{H}) = \{e_1, \dots, e_m\}$, $m \geq 0$,

where for every $1 \leq i \leq m$, e_i is a subset of \mathcal{V} , called a hyperedge, and its index i is called the *edge ID* of e_i . Since $\mathcal{E} \subseteq 2^{\mathcal{V}}$, the number m of hyperedges can be exponential in n . The *total size* of \mathcal{H} , denoted by $\|\mathcal{H}\| = N$, is the sum of the sizes of its hyperedges, that is, $\|\mathcal{H}\| = \|\mathcal{E}\| = \sum_{e \in \mathcal{E}} |e|$. By definition, $\|\mathcal{E}\| \leq O(mn)$.

For vertex x and hyperedges e, f , we say that e is *incident to* x (or e *contains* x) if $x \in e$ holds, and that e is a *neighbor* of f if $f \cap e \neq \emptyset$ holds. Then, $N(x) = \{f \in \mathcal{E} \mid x \in f\}$ is the set of all hyperedges incident to x , and $NE(e) = \{f \in \mathcal{E} \mid e \cap f \neq \emptyset\}$. Then, the *rank* of \mathcal{H} , denoted by $r = \text{rank}(\mathcal{H}) = \max_{f \in \mathcal{E}} |f|$, is the maximum size of its hyperedges. The *degree* of \mathcal{H} , denoted by $d = \text{deg}(\mathcal{H}) = \max_{x \in \mathcal{V}} |N(x)|$, is the maximum number of incident edges. The *hyperedge degree* of \mathcal{H} , denoted by $D = \text{hyperedge-deg}(\mathcal{H}) = \max_{e \in \mathcal{E}} |NE(e)|$, is the maximum number of the neighbors. Clearly, we see that $r = \text{rank}(\mathcal{H}) \leq n$, $d = \text{deg}(\mathcal{H}) \leq m$, and $D = \text{hyperedge-deg}(\mathcal{H}) \leq \min\{m, rd\}$.

In this paper, a *sub-hypergraph* of \mathcal{H} is just a subset $S = \{e_{i_1}, \dots, e_{i_k}\} \subseteq \mathcal{E}$ ($k \leq m$) of hyperedges of \mathcal{H} .⁴ We use the terms a hyperedge subset and a sub-hypergraphs interchangeability in what follows. Actually, a subset S induces a hypergraph $\mathcal{H}[S] = (S, \mathcal{E}[S])$, where $\mathcal{E}[S] = \{e \in \mathcal{E} \mid e \subseteq S\}$. The *neighbor set* of S is $N(S) = \{f \in \mathcal{E} \setminus S \mid e \in S, e \cap f \neq \emptyset\}$. Clearly, $\|S + N(S)\| \leq \|\mathcal{E}\| = O(mn)$.

In what follows, we refer to vertices as x, y, \dots , hyperedges as e, f, \dots , and hyperedge subsets as S, T, \dots , possibly subscripted. For convenience, we often represent a set of elements $\{a, b, c\}$ by a juxtaposition abc if it is clear from context.

Example 1. In Fig. 1, we show an example of a hypergraph $\mathcal{H}_1 = (\mathcal{V}_1, \mathcal{E}_1)$ consisting of eleven vertices $\mathcal{V}_1 = \{p, q, r, \dots, x, y, z\}$ and seven hyperedges $\mathcal{E}_1 = \{e_1, e_2, \dots, e_7\}$ such that $e_1 = qrs, e_2 = uv, e_3 = tu, e_4 = vw, e_5 = qxy, e_6 = xyz$, and $e_7 = stwx$.

2.2 Connected and Berge-acyclic sub-hypergraphs

A *path* between hyperedges e and $f \in \mathcal{E}$ in $S \subseteq \mathcal{E}$ is a sequence $\pi = (e_1 = e, e_2, \dots, e_k = f)$ ($k \geq 1$) of hyperedges that satisfies the condition $e_i \cap e_{i+1} \neq \emptyset$ for every $1 \leq i \leq k - 1$.

Definition 1. A hyperedge subset S is *connected* if any pair of hyperedges e and f has some path between them in S .

⁴ The definition of a sub-hypergraph in this paper is also referred to as a *partial hypergraph* in literatures.

By definition, the empty set and singleton set of hyperedges are connected. We can easily test the connectivity of a given subset S in $O(|S|)$ time.

Example 2. In the hypergraph \mathcal{H}_1 in Example 1, the subsets $S_1 = 1567 = \{1, 5, 6, 7\}$ and $S_2 = 1347$ are connected. On the other hand, the subset $S_3 = 135$ is not connected since there is no path between the edges 1 and 3, and also the edges 5 and 3.

Definition 2 ([6]). *In a hypergraph \mathcal{H} , a Berge-cycle (or simply a cycle) of length k is a sequence $\pi = (e_1, x_1, \dots, e_k, x_k)$ ($k \geq 2$) that satisfies the following conditions (i)–(iii):*

- (i) e_1, \dots, e_k are mutually distinct hyperedges.
- (ii) x_1, \dots, x_k are mutually distinct vertices.
- (iii) For each $1 \leq i \leq k - 1$, $x_i \in e_i \cap e_{i+1}$ holds, and $x_k \in e_k \cap e_1$ holds.

In the above definition, we say that the set $\{e_1, \dots, e_k\}$ of hyperedges forms a Berge-cycle. Intuitively, a Berge-cycle is a path of length more than or equal to two that starts from some hyperedge and returns to the start.

Example 3. In the hypergraph \mathcal{H}_1 in Example 1, the hyperedge subset $S_4 = 157$ forms a Berge-cycle $\pi_4 = (1, q, 5, x, 7, s)$ of length three. From Lemma 1, we also see that the pair of hyperedges $S_5 = 56$ forms a Berge-cycle $\pi_5 = \langle 5, x, 6, y \rangle$ of length two since hyperedges 5 and 6 share common vertices x and y .

From the construction of minimum length cycle S_5 in the above example, we have the following lemma, which is well-known providing a fundamental property of Berge-acyclicity.

Lemma 1 (Berge [6]). *If two hyperedges e and f contain mutually distinct vertices x and y in common, i.e., $x, y \in e \cap f$, then they form a Berge-cycle.*

Proof. Take a path $\pi = (e, x, f, y)$ as a Berge-cycle. ■

Definition 3 (Berge-acyclic subgraph [6]). *A sub-hypergraph S is Berge-acyclic if it contains no Berge-cycles.*

By definition, the empty set and any singleton sets of hyperedges are Berge-acyclic. From the next lemma, Berge-acyclicity is closed under subsets.

Lemma 2. *If a non-empty subset S is Berge-acyclic, then any subset S' ($S' \subseteq S$) is also Berge-acyclic.*

From Lemma 1 above, we see that Berge-acyclicity has strong restriction compared to other notions of hypergraph acyclicities. Actually, there is a hierarchy of acyclicities for hypergraphs, called the *degrees of acyclicities* of Fagin [9], that consists of α -acyclicity, β -acyclicity, γ -acyclicity, and Berge-acyclicities. In this hierarchy, α -acyclicity is most general, while Berge-acyclicity is most restricted.

In what follows, we denote by $\mathcal{AC} = \mathcal{AC}(\mathcal{H})$ the class of *all connected, and Berge-acyclic sub-hypergraphs* in an input hypergraph \mathcal{H} . Now, we state our data mining problem.

$$A_1 = \begin{matrix} & q & r & s & t & u & v & w & x & y & z \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \end{matrix} & \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \end{pmatrix} \end{matrix}$$

Fig. 2. The incident matrix A_1 of the hypergraph \mathcal{H}_1 in Fig. 1, where each row indicates a hyperedge and each column a vertex.

Definition 4 (Connected and Berge-acyclic sub-hypergraph mining problem in a hypergraph). *Given an input hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$, the task is to find all connected, and Berge-acyclic sub-hypergraphs $S \subseteq \mathcal{E}$ in \mathcal{H} belonging to the class $\mathcal{AC}(\mathcal{H})$ without duplicates.*

Example 4. Consider the hypergraph \mathcal{H}_1 in Example 1 again. Then, the subset $S_1 = 1567$ is not a connected and Berge-acyclic subset in $\mathcal{AC}(\mathcal{H}_1)$ because it is connected but cyclic. On the other hand, the subset $S_2 = 1347$ is a connected and Berge-acyclic subset in $\mathcal{AC}(\mathcal{H}_1)$.

The model of computation: An enumeration algorithm \mathcal{A} receives an instance of size n and outputs all of m solutions without duplicates (See, e.g. [4]). For polynomials $p(\cdot, \cdot), q(\cdot), r(\cdot)$, \mathcal{A} is of *output* $O(p(n, m))$ -time if the total time of \mathcal{A} is bounded by polynomial in n and m . \mathcal{A} is of $O(q(n))$ -delay using preprocessing $r(n)$ if the *delay*, which is the maximum computation time between two consecutive outputs, is bounded by $q(n)$ after preprocessing in $r(n)$ time. Clearly, if \mathcal{A} is polynomial delay, \mathcal{A} is also output polynomial time. \mathcal{A} is of *polynomial space* if the maximum size of its working space is bounded by some polynomial $p(n)$.

2.3 Other definitions and properties

Leaves and connection counts: Let $S \subseteq \mathcal{E}$ be a subset of hyperedges, or a sub-hypergraph of \mathcal{H} . A hyperedge e *connects* S if the intersection $e \cap V(S)$ is not empty. Any vertex x in the intersection is called a *connection point*. Then, the *connection count* of e relative to S is defined by $\text{cnt}(e, S) = |e \cap V(S)| \geq 0$. In the next section, we give a characterization of connected and Berge-acyclic sub-hypergraphs using the connection count.

A *leaf* of a subset S is a hyperedge $e \in S$ such that $\text{cnt}(e, S - e) = 1$, that is, that has a single connection point in S except itself. Clearly, the empty subset \emptyset has no leaf at all, and any singleton $S = \{e\}$ has the hyperedge e as its only leaf. We denote by $L(S)$ the *set of all leaves* of S . Actually, $L(S) = \{e \in S \mid \text{cnt}(e, S \setminus \{e\}) = 1\}$.

Representation of hypergraphs: Using the incident relation $x \in e$, a hypergraph \mathcal{H} with n vertices and m hyperedges can be represented as an $n \times m$ binary matrix $A = (a_{i,j}) \in \{0,1\}^{m \times n}$, called the *incident matrix* of \mathcal{H} . For every $1 \leq i \leq n$ and $1 \leq j \leq m$, $a_{i,j} = 1$ if and only if $x_i \in e_j$ holds. In an incident matrix $A = (a_{i,j})$, each row $1 \leq i \leq n$ represents the incident set $N(x_i) \subseteq \mathcal{E}$, while each column $1 \leq j \leq m$ represents the corresponding hyperedge $e_j \subseteq \mathcal{V}$. In Fig. 2, we show an example of an incident matrix.

Data structure: In our algorithms in Sec. 3 and Sec. 4, we use a dynamic data structure \mathcal{D} similar to the *DLX* (also known as “*Dancing Links*”) data structure by Knuth [15] for dynamically maintaining a hyperedge subset S . Our data structure \mathcal{D} stores the incident matrix of a set collection $D \subseteq \mathcal{E}$ in linear words of space in $\|D\|$ supporting the following operations: (i) retrieval of a hyperedge $e = e_i$ by an edge ID i , (ii) retrieval of the neighbor $N(x, D)$ by a vertex x , and (iii) insert/delete of elements to/from an edge or a neighbor set. Using dynamic predecessor dictionaries [7] (such as the hash table or the `map` collection of C++/STL or Java), we can execute the above operations in sublinear time $t = \tau(k)$, where we have $\tau(k) = \log k$ if we use ordinary binary tree, and $\tau(k) = O((\log \log k)^2 / \log \log \log k)$ for $k = \max\{n, m\}$ if we use the dynamic data structure of Beame and Fich [5]. The details are omitted here.

3 The Basic algorithm

In this section, we show the basic version of our DFS mining algorithm BERGEMINE that finds all connected, and Berge-acyclic sub-hypergraphs in \mathcal{H} in polynomial delay and space. In what follows, we write $S - e$ for $S' \setminus \{e\}$.

To devise efficient depth-first search algorithm, we need a systematic way to reduce the search for larger subsets to smaller subsets. The next lemma is essential to our algorithm.

Lemma 3. *Let $S' \subseteq \mathcal{E}$ is a subset such that $|S'| \geq 2$. If S' is connected and Berge-acyclic, then $\text{cnt}(e, S' - e) = 1$ holds for some $e \in S'$. Furthermore, $S = S' - e$ is connected and Berge-acyclic, too, and has size $|S| < |S'|$.*

Proof. We can show the lemma by induction on $|S'|$. If $|S'| = 2$, the claim is obvious since S' consists of two edges. Otherwise, assume that $|S'| > 2$. Since S' is connected, $\text{cnt}(e, S - e) \geq 1$ always holds for any $e \in S'$. Furthermore, if $\text{cnt}(e, S - e) \leq 1$ holds for some $e \in S'$, then we are done. Therefore, we assume that $\text{cnt}(e, S - e) \geq 2$ holds for any $e \in S'$. Consider this case. Then, we split S' by removing e , and consider the connected components S_1, \dots, S_k of $S' - e$, where $k \geq 1$. There are two cases. (i) If e connects to some S_i at at least two points, then $S_i \cup \{e\}$, and thus S' , immediately has a cycle, and we are done (ii) Otherwise, using induction hypothesis, we can show that there exists an edge f in some component, say S_1 , such that $\{e, f\} \cup R$ forms a cycle for some $R \in \{\{e\}, S_1 - f, S_2, \dots, S_k\}$ (details are omitted), and we are done. Hence, by contradiction, the lemma follows. ■

Algorithm 1 A basic algorithm BERGEMINE for mining all connected, Berge-acyclic sub-hypergraphs based on the reverse search

```

1: procedure BERGEMINE( $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ : input hypergraph )
2:   BASICREC( $\emptyset, \mathcal{H}$ );

3: procedure BASICREC( $S$ : sub-hypergraph,  $\mathcal{H}$ : input hypergraph)
4:   Output  $S$ ;
5:    $Border(S) \leftarrow \{ f \in (\mathcal{E}(\mathcal{H}) \setminus S) \mid cnt(f, S) = 1 \}$ ;
6:   for each  $f \in Border(S)$  do                                      $\triangleright$  Generation of children
7:      $S' \leftarrow S \cup \{f\}$ ;
8:     if  $f = \max L(S')$  then
9:       BASICREC( $S', \mathcal{H}$ );

```

From Lemma 3 above, Starting from any connected and Berge-acyclic subhypergraph S with more than one edges, we can obtain a series of sub-hypergraphs $\mathcal{R} = S_0 = S \supset S_1 \supset \dots \supset S_\ell = \{e\}$ of length $\ell = |S| - 1 \geq 0$. Our DFS algorithm reverses this process by starting from any singleton set $\{e\}$, $e \in \mathcal{E}$, and by iteratively expanding the current subset $S \subseteq \mathcal{E}$ by adding new hyperedge $e \in \mathcal{E} \setminus S$ in a systematic manner using backtracking.

However, there is one problem with this approach. The above DFS search process may generate the same subset by exponentially many different paths. One easy way to avoid this duplication is to use table-lookup. When we discovered a new subset S , we lookup a hash table H to decide if $S \in H$. If so, we skip S , and otherwise, we output S and register it to H . This modification yields a polynomial delay, but exponential space mining algorithm.

We solve this problem by pruning of redundant path by careful design of the tree-shaped search space described as follows. Recall the previous key lemma, Lemma 3. In the lemma, the possible source of redundancy is more than one choice of a leaf $e \in S$ of S to delete. An idea to solve this is to restrict the deletion in reduction (and the addition in generation) to the *maximum* leaf of S . This ensures the reduction sequence $\mathcal{R} = S_0, \dots, S_\ell$ for S to be unique to each S . We call such a unique sequence the *maximum elimination sequence* for a sub-hypergraph S , and denote by $\mathcal{MES}(S)$.

Lemma 4. $\mathcal{MES}(S)$ is the unique signature of each connected and Berge-acyclic sub-hypergraph $S \subseteq \mathcal{E}$.

From this lemma, we can generate S in a unique way by generating $\mathcal{MES}(S)$ instead. Now, we describe our algorithm.

Definition 5. Let S' be any connected and Berge-acyclic subhypergraph S' such that $|S'| \geq 2$. Then, the *parent* of S' is the set $\mathcal{P}(S') = S' - f$, where f is the leaf of S' such that $cnt(f, S' - f) = 1$ having the maximum edge ID among all leaves, that is, $f = \max(L(S'))$. This condition is called the *maximum leaf condition*. In this case, we call S' a *child* of $S = \mathcal{P}(S')$.

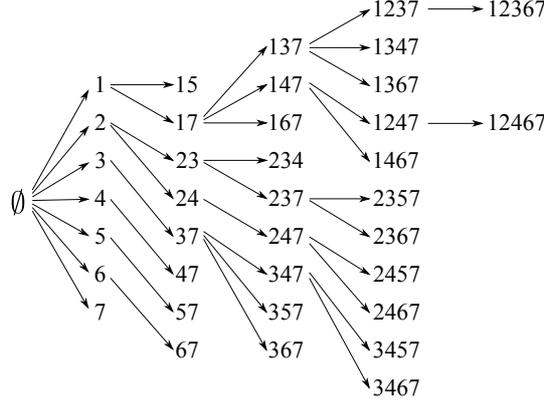


Fig. 3. The family tree for \mathcal{H}_1 in Fig. 1.

Then, we define our tree-shaped search space. Let \mathcal{H} be an input hypergraph. The *family tree* for the class \mathcal{CA} of connected, and Berge-acyclic sub-hypergraphs of \mathcal{H} is a multi-rooted DAG $\mathcal{T} = (\mathcal{CA}, \mathcal{P}, \mathcal{I})$, where

- \mathcal{CA} is the vertex set of \mathcal{T} that consists of all connected, and Berge-acyclic sub-hypergraphs in an input hypergraph \mathcal{H} .
- \mathcal{P} defines the set of reverse edges of \mathcal{T} that assign the parent $\mathcal{P}(S')$ to a child S' .
- \mathcal{I} is the set of all single subsets as the root nodes of \mathcal{T} .

The next lemma says that the family tree is actually a tree-shaped search root.

Lemma 5. *For any input hypergraph \mathcal{H} , the family tree for \mathcal{CA} on \mathcal{H} is a spanning forest that contains all connected and Berge-acyclic subsets in \mathcal{CA} as its nodes.*

Proof. From Lemma 3, it immediately follows that $\mathcal{P}(S')$ is always connected, and Berge-acyclic, and has size strictly smaller than S' . Since each path of \mathcal{T} is a \mathcal{MES} for some element of \mathcal{CA} , \mathcal{T} is connected at some root in \mathcal{I} . On the other hand, since each reverse edge strictly reduces the size of S , \mathcal{T} contains no cycle. Hence, the lemma is proved. ■

Example 5. In Fig. 3, we show the family tree for \mathcal{H}_1 in Fig. 1. For example, the parent $\mathcal{P}(S_2)$ of S_2 is $\mathcal{P}(P_2) = 137$. Then, there exists the reverse edge from $\mathcal{P}(S_2)$ to its child S_2 . $\mathcal{P}(S_2)$ has other children $S_6 = 1237$ and $S_7 = 1367$. Edges $(\mathcal{P}(S_2), S_6)$ and $(\mathcal{P}(S_2), S_7)$ are also the members of the set of reverse edge in the family tree.

In Algorithm 1, we show our basic DFS algorithm BERGEMINE and its recursive subprocedure BASICREC that finds all connected, and Berge-acyclic sub-hypergraphs in \mathcal{H} in depth-first manner. This algorithm is a simple backtracking algorithm, working as follows. Starting from each singleton subset $\{e\}$ in \mathcal{I} , the

Algorithm 2 The algorithm for computing the border set of a sub-hypergraph

```

1: procedure COMPUTEBORDER( $S$ : sub-hypergraph,  $\mathcal{V}, \mathcal{E}$ )
2:   Output: Border( $S$ ) =  $\{ f \in (\mathcal{E} \setminus S) \mid cnt(f, S) = 1 \}$ .
3:   Mark all vertices of  $\mathcal{V}(S)$ ;
4:    $Border \leftarrow \emptyset$ ;
5:   for each  $e \in \mathcal{E}(S)$  do
6:     Count the number  $cnt(e, S)$  of all marked vertices in  $e$ ;
7:     if  $cnt(e, S) = 1$  then
8:        $Border \leftarrow Border \cup \{e\}$ ;
9:   return  $Border$ ;

```

algorithm searches the family tree \mathcal{T} for connected, and Berge-acyclic subsets by expanding the parent subset S by adding a new leaf f to obtain a child $S' = S \cup \{f\}$. In the expansion, it apply pruning for redundant subsets using the definition of a correct child based on the maximum leaf condition of the child. If expansion is no longer possible, it backtrack to the parent.

To compute the border set, we use the procedure COMPUTEBORDER in Algorithm 2.

Lemma 6. *The algorithm COMPUTEBORDER in Fig. 2 computes the border set of an hyperedge subset S $O(\|S\|) = O(nm)$ time using $O(n)$ additional space.*

We give the time and space complexity of the basic algorithm below.

Theorem 1 (main result). *The algorithm BERGEMINE of Fig. 1 finds all connected Berge-acyclic sub-hypergraphs contained in an input hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ without duplicates in $O(Nm) = O(nm^2)$ delay and $O(N)$ words of space, where $n = |\mathcal{V}|$, $m = |\mathcal{E}|$, and $N = \|\mathcal{E}\|$ are the numbers of vertices and hyperedges, and the total size of the hyperedges in \mathcal{H} .*

From the theorem, we have the following corollary.

Corollary 1. *The class of all connected Berge-acyclic sub-hypergraphs contained in an input hypergraph \mathcal{H} can be enumerated in polynomial delay and polynomial space in the size of input \mathcal{H} .*

4 The Modified Algorithm

In this section, we show a modified version of our depth-first mining algorithm that finds all connected and Berge-acyclic sub-hypergraphs in an input hypergraph \mathcal{H} in $O(|f|\tau(m) + |B|)$ time using $O(N)$ space and preprocessing, where $N = \|\mathcal{E}(\mathcal{H})\|$. This algorithm is *adaptive* since its time complexity only depends on the size of the discovered sub-hypergraph S , rather than the whole input. This adaptively is quite important in mining a large hypergraph. In what follows, we use the dynamic data structure of Beame and Fich [5] with operation time $\tau(m) = O((\log \log m)^2 / \log \log \log m)$.

The basic idea of our modified algorithm is an incremental maintenance of the subset $MaxBorder(S)$ of hyperedge candidates to insert, called the maximal border hyperedges.

Definition 6. The maximal border of a sub-hypergraph S is the set of hyperedges defined by:

$$MaxBorder(S) = \{ e \in \mathcal{E} \setminus S \mid cnt(e, S) = 1, e = \max L(S \cup \{e\}) \}, \quad (1)$$

that is, $MaxBorder(S)$ consists of all hyperedges e of \mathcal{H} satisfying the next conditions: (i) e is a border of S (i.e., $cnt(e, S) = 1$), and (ii) e is the maximum leaf of $S' = S \cup \{e\}$ among all leaves when it is added to S .

In Algorithm 4, we show our modified depth-first algorithm FASTBERGEMINE as well as its recursive subprocedure FASTREC for mining all connected, Berge-acyclic sub-hypergraphs incrementally. By using $MaxBorder(S)$, in our depth-first mining algorithm FASTBERGEMINE, we can generate any children $S' = S \cup \{f\}$ from a parent S by just selecting any hyperedge $e \in MaxBorder(S)$ without testing the pruning condition for duplication because the condition is already included by the definition of the maximal border set. In other words, we are eager to make selection of border candidates and test for duplication at the same time in advance.

Therefore, it remains how to efficiently compute the maximal border set. Surprisingly, we can show that this is done in almost optimal time complexity in amortized analysis using a procedure similar to the α -acyclicity test by (Tarjan and Yannakakis [22]). The key to the algorithm is the following recurrence relation for the maximum and the second maximum leaves when we update a parent S by adding a new maximum border $f \in MaxBorder(S)$ to generate a children $S' = S \cup \{f\}$.

Lemma 7. *Let us denote by $max(S)$ and $2max(S)$ the maximum and the second maximum leaves of a parent set $S \subseteq \mathcal{E}$. Then, the maximum and the second maximum leaves $max(S')$ and $2max(S')$ of a child $S' = S \cup \{f\}$ satisfy the following recurrence:*

- If f connects $\ell_{\max} = maxLeaf(S)$:
 - If $f > 2max(S)$, then $max(S') \leftarrow f$ and $2max(S') \leftarrow max(S)$ hold.
 - Otherwise, $max(S') \leftarrow max(S)$ and $2max(S') \leftarrow 2max(S)$ hold.
- Otherwise:
 - If $f > max(S)$, then $max(S') \leftarrow f$ and $2max(S') \leftarrow max(S)$ hold.
 - Otherwise, $max(S') \leftarrow max(S)$ and $2max(S') \leftarrow 2max(S)$.

Proof. In each case, the proof immediately follows from the case analysis using the definitions of max , $2max$, and the maximal border set. ■

From Lemma 7 above, we can update $max(S)$ and $2max(S)$ incrementally in constant time. Now, we show the algorithm UPDATEMAXBORDER in Algorithm 3 that incrementally updates the new border set $MaxBorder(S \cup \{f\})$ from the older one given the border edge f to add, S , and $MB = MaxBorder(S)$.

Algorithm 3 The algorithm for computing the border set of a sub-hypergraph.

```

1: procedure UPDATEMAXBORDER( $f$ : hyperedge,  $S, B, R$ : hyperedge subsets,  $\mathcal{H}$ : hypergraph )
   Pre-conditions:  $S' = S \cup \{f\}$ ,  $f \in R$ ,  $B = MB(S)$ , and  $R = \mathcal{E}(\mathcal{H}) \setminus S$ .
   Output:  $MB(S') = \{f \in (\mathcal{E} \setminus S') \mid cnt(f, S') = 1, f = \max L(S')\}$ .
   Global variable: A dynamic data structure  $\mathcal{D}$  for storing a hyperedge subset in linear space supporting membership, insert, and delete in sublinear time  $t = \tau(m)$ .
   //Step 1: Update the maximum leaves.
2:    $S' = S \cup \{f\}$ ;
3:   Update the maximum leaves  $max(S')$  and  $2max(S')$ 
     from  $f$ ,  $max(S')$ , and  $2max(S')$  according to Lemma 7.           ▷ in  $O(1)$  time
   //Step 2: Update the maximum border set.
4:    $MB(S') \leftarrow \emptyset$ ;
5:   //Step 2.1: Existing borders other than  $f$ 
6:   for each  $e$  in  $MB(S) \setminus \{f\}$  do                                     ▷  $O(|MB(S)|)$  times
7:     Add  $e$  to  $MB(S')$  if  $e = maxL(S' \cup \{e\})$ .
8:   //Step 2.2: New borders connecting to  $f$ 
9:   for each vertex  $x \in f$  do                                           ▷  $O(|f|)$  times to  $S'$ 
10:    for each hyperedge id  $e \in N(x, \mathcal{D})$  do                             ▷ Charge  $O(1)$  time to  $e$  in  $\mathcal{D}$ 
11:       $cnt[e] \leftarrow cnt[e] + 1$ ;
12:      if  $cnt[e] = 1$  then                                               ▷  $cnt$  increased from 0 to 1
13:        Add  $e$  to the candidate set  $\mathcal{D}$ ;                               ▷ Charge  $O(\tau(m))$  time to  $e$ 
14:        Add  $e$  to  $MB(S')$  if  $e = maxL(S' \cup \{e\})$ .
15:      else if  $cnt[e] = 2$  then                                         ▷  $cnt$  increased from 1 to 2
16:        Remove  $e$  from candidate set  $\mathcal{D}$ ;                               ▷ Charge  $O(\tau(m))$  time to  $e$ 
17:      //Note: each hyperedge is processed at most twice overall
18:    end for
19:  return  $MB$ ;

```

For efficient update, the algorithm uses a dynamic data structure \mathcal{D} for storing a set \mathcal{D} of candidate hyperedges, which is similar to the *DLX* (also known as “*Dancing Links*”) data structure by Knuth [15] as described in Sec. 2. Then, we have the next lemma.

Lemma 8. *Let $S \subseteq \mathcal{E}$ be a sub-hypergraph and $f \in R = (\mathcal{E} \setminus S)$ be a maximum border hyperedge of S . Given f , B , and R , the algorithm UPDATEMAXBORDER in Algorithm 3 computes the set $MaxBorder(S \cup \{f\})$ of all maximum border hyperedges of $S' = S \cup \{f\}$ in $O(\Delta MB(S)\tau(m) + |B|)$ time using $O(N)$ space and $O(N)$ preprocessing (at once in the initialization), where $\Delta MB(S)$ is the number of newly added hyperedges to the maximum border of S , $N = \|\mathcal{E}(\mathcal{H})\|$, $\tau(m) = ((\log \log m)^2 / \log \log \log m)$, r and d are the rank and degree of \mathcal{H} , respectively, and $B = MaxBorder(S)$ is the set of all maximum borders of S .*

Proof. Consider Algorithm 3. During the computation of the recursive mining procedure, we maintain the pointers $max(S)$, $2max(S)$, and the dynamic data structure \mathcal{D} . From Lemma 7, Step 1 correctly updates the maximum and 2nd maximum leaves in S in constant time. When a new border f is added to S , the

Algorithm 4 The modified algorithm FASTBERGEMINE for mining all connected, Berge-acyclic sub-hypergraphs based on the reverse search

```

1: procedure FASTBERGEMINE( $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ : input hypergraph )
2:   for each hyperedge  $e \in \mathcal{E}(\mathcal{H})$  do
3:      $MB_e \leftarrow \{ f \in \mathcal{E}(\mathcal{H}) \mid (|f \cap e| = 1) \}$ ;
4:      $R_e \leftarrow \mathcal{E}(\mathcal{H}) \setminus \{e\}$ ;
5:     FASTREC( $\{e\}, MB_e, R_e, \mathcal{H}$ );

6: procedure FASTREC( $S, MB, R \subseteq \mathcal{E}(\mathcal{H}), \mathcal{H}$ : hypergraph)
7:   Invariant:  $MB = \text{MaxBorder}(S)$  and  $R = \mathcal{E}(\mathcal{H}) \setminus S$  hold.
8:   Output  $S$ ;
9:   for each border hyperedge  $f \in MB$  do ▷ Generation of children
10:     $S' \leftarrow S \cup \{f\}$ ;
11:     $R' \leftarrow R \setminus \{f\}$ ;
12:    Incrementally compute  $MB' = \text{MaxBorder}(S', \mathcal{H})$  from  $f, MB,$  and  $R$ ;
13:    FASTREC( $S', MB', R', \mathcal{H}$ );
14:    Restore the changes on  $MB'$ ;

```

only borders to be changed are (i) f is removed, and (ii) all neighboring hyperedges of f , and (ii) all existing border edges that has a non-empty intersection to f other than its connection point to S . Step 2 handles these cases correctly. For time analysis of Step 2, we observe that during computation from the root hypothesis \emptyset to the current set S , any hyperedge e in \mathcal{H} will be processed at most twice after initialization, that is, it is incremented to $\text{cnt}(e) = 1$ at the first time, and it is incremented to $\text{cnt}(e) = 2$ the second time. Then, it is removed from \mathcal{D} forever (otherwise a back tracking occurs). We can show that the amortized cost for Step 2 to obtain each the maximum border of S' is at most $O(\Delta MB(S)\tau(m) + |B|)$, where $\tau(m) = ((\log \log m)^2 / \log \log \log m)$. ■

From Lemma 8, we show the main theorem of this paper.

Theorem 2 (The adaptive delay by the modified mining algorithm).
The algorithm FASTBERGEMINE of Fig. 4 finds all connected Berge-acyclic sub-hypergraphs contained in an input hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ without duplicates in $t = O(\Delta MB(S)\tau(m)) = O(rd \cdot \tau(m))$ amortized delay (amortized time per solution) using $O(N)$ space and $O(N)$ preprocessing, where f is the added edge, B is the maximum border of S , $N = \|\mathcal{E}\|$ is the total size of the hyperedges in \mathcal{H} , r and d are the rank and degree of \mathcal{H} , respectively, and $\tau(m) = ((\log \log m)^2 / \log \log \log m)$.

Proof. The correctness of the algorithm FASTBERGEMINE is obvious from that of the basic algorithm and the definition of *MaxBorder*. From Lemma 8, at each iteration for solution, the computation time of the maximum border is $t = O(\Delta MB(S)\tau(m) + |B|) = O(rd \cdot \tau(m) + |B|)$. By using appropriate charging scheme to each child of S' , we can remove the cost $O(|B|)$ since the addition of any maximum border hyperedge to S' always yields a proper solution. Moreover, we have $\Delta MB(S) = O(|f| \sum_{x \in f} |N(x)|) = O(rd)$. Therefore, the time complexity becomes $t = O(\Delta MB(S)\tau(m)) = O(rd \cdot \tau(m))$ as claimed. ■

From the theorem, we have the following corollary.

Corollary 2. *The class of all connected Berge-acyclic sub-hypergraphs contained in an input hypergraph \mathcal{H} can be enumerated in amortized delay that depends only on the number of newly added maximum border hyperedges of a discovered subset S using polynomial space and preprocessing in the input size $||\mathcal{E}(\mathcal{H})||$.*

It is an open question whether there exists some enumeration algorithm whose amortized delay depends only on the number of difference of a discovered subset.

5 Conclusion

In this paper, we considered the problem of finding all all connected Berge-acyclic sub-hypergraphs contained in an input hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ without duplicate with applications to generalization of itemset mining from transaction databases and also discovering connected substructures from datasets in the form of sets of sets. As main results, we presented an efficient DFS algorithm for the problem that achieves polynomial delay and space complexity. We also presented an improved algorithm that has adaptive delay depending only on the size of discovered sub-hypergraph.

In this paper, we focused on only the theoretical aspect of the problem. One of the most important future researches is implementation and empirical evaluation of the proposed algorithms on artificial and real datasets. It is also an important problem to find suitable application of this problem in knowledge discovery problems in the real world including knowledge discovery from mobility data or social networks. We want to study these problems in future.

Acknowledgements.

The authors would like thank anonymous reviewers for their comments which improved the correctness and the presentation of this paper very much, and thank Shin-ichi Nakano, Kunihiro Sadakane, and Tetsuji Kuboyama for helpful discussions and comments. This research was partly supported by MEXT Grant-in-Aid for Scientific Research (A), 24240021, FY2012—2015, and Grant-in-Aid for JSPS Fellows (25·1149).

References

1. H. Arimura and T. Uno. An efficient polynomial space and polynomial delay algorithm for enumeration of maximal motifs in a sequence. *Journal of Combinatorial Optimization*, 13:243–262, 2006.
2. H. Arimura and T. Uno. Mining maximal flexible patterns in a sequence. In *New Frontiers in Artificial Intelligence*, LNCS 4914, pages 307–317, 2007.

3. H. Arimura and T. Uno. Polynomial-delay and polynomial-space algorithms for mining closed sequences, graphs, and pictures in accessible set systems. In *Proceedings of the SIAM Int'l Conf. on Data Mining 2009 (SDM'09)*, pages 1087–1098, 2009.
4. D. Avis and K. Fukuda. Reverse search for enumeration. *Discrete Applied Math.*, 65:21–46, 1993.
5. P. Beame and F. E. Fich. Optimal bounds for the predecessor problem and related problems. *Journal of Computer and System Sciences*, 65(1):38 – 72, 2002.
6. C. Berge and E. Minieka. *Graphs and hypergraphs*, volume 7. North-Holland, 1973.
7. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, 2 edition, 2001.
8. T. Daigo and K. Hirata. On Generating All Maximal Acyclic Subhypergraphs with Polynomial Delay. *SOFSEM 2009: Theory and Practice of Computer ...*, pages 181–192, 2009.
9. R. Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *Journal of the ACM*, 30(3):514–550, July 1983.
10. R. Ferreira, R. Grossi, A. Marino, and N. Pisanti. Optimal Listing of Cycles and st-Paths in Undirected Graphs. 2012.
11. R. Ferreira, R. Grossi, and R. Rizzi. Output-sensitive listing of bounded-size trees in undirected graphs. *Algorithms ESA 2011*, pages 275–286, 2011.
12. K. Hirata, M. Kuwabara, and M. Harao. On finding acyclic subhypergraphs. *Fundamentals of Computation Theory*, pages 491–503, 2005.
13. A. Inokuchi, T. Washio, and H. Motoda. An apriori-based algorithm for mining frequent substructures from graph data. In *Proc. PKDD'00, LNCS 1910*, pages 13–23. Springer, 2000.
14. S. KAWASOE, H. SAKAMOTO, H. ARIMURA, and S. ARIKAWA. Efficient substructure discovery from large semi-structured data. *IEICE TRANSACTIONS on Information and Systems*, 87(12):2754–2763, 2004.
15. D. E. Knuth. Dancing links. *eprint arXiv:cs/0011047*, Nov. 2000.
16. T. Kuboyama, K. Hirata, and K. F. Aoki-Kinoshita. An efficient unordered tree kernel and its application to glycan classification. In *Proc. PAKDD'08, LNCS 5012*, pages 184–195. Springer, 2008.
17. X.-L. Li, S.-H. Tan, C.-S. Foo, S.-K. Ng, et al. Interaction graph mining for protein complexes using local clique merging. *Genome Informatics*, 16(2):260, 2005.
18. L. Lovász. Matroid matching and some applications. *Journal of Combinatorial Theory, Series B*, 236:208–236, 1980.
19. J. Pei, J. Han, B. Mortazavi-Asl, J. Wang, H. Pinto, Q. Chen, U. Dayal, and M.-C. Hsu. Mining sequential patterns by pattern-growth: The prefixspan approach. *IEEE TKDE*, 16(11):1424–1440, 2004.
20. A. Shioura, A. Tamura, and T. Uno. An optimal algorithm for scanning all spanning trees of undirected graphs. *SIAM J. Comput.*, 26(3):678–692, 1997.
21. A. Silva, W. Meira Jr, and M. J. Zaki. Structural correlation pattern mining for large graphs. In *Proceedings of the Eighth Workshop on Mining and Learning with Graphs*, pages 119–126. ACM, 2010.
22. R. Tarjan and M. Yannakakis. Simple Linear-Time Algorithms to Test Chordality of Graphs, Test Acyclicity of Hypergraphs, and Selectively Reduce Acyclic Hypergraphs. *SIAM Journal on computing*, 13(3), 1984.
23. R. E. Tarjan and R. C. Read. Bounds on backtrack algorithms for listing cycles, paths, and spanning trees. *Networks*, 5(3):237–252, 1975.
24. T. Uno and H. Arimura. Ambiguous frequent itemset mining and polynomial delay enumeration. In *Proc. PAKDD2008, LNAI 5012*, pages 357–368. Springer, 2008.

25. T. Uno, T. Asai, Y. Uchida, and H. Arimura. An efficient algorithm for enumerating closed patterns in transaction databases. In *Proc. the 7th Discovery Science (DS'04)*, volume 3245 of *LNCS*, pages 16–31. Springer, 2004.
26. K. Wasa, Y. Kaneta, T. Uno, and H. Arimura. Constant time enumeration of bounded-size subtrees in trees and its application. In *Computing and Combinatorics*, pages 347–359. Springer, 2012.
27. X. Yan and J. Han. gspan: Graph-based substructure pattern mining. In *Data Mining, 2002. ICDM 2003. Proceedings. 2002 IEEE International Conference on*, pages 721–724. IEEE, 2002.
28. M. J. Zaki. Scalable algorithms for association mining. *IEEE Transactions on Knowledge and Data Engineering*, 12(3):372–390, May/Jun 2000.
29. M. J. Zaki and C.-J. Hsiao. Efficient algorithms for mining closed itemsets and their lattice structure. *IEEE Transactions on Knowledge and Data Engineering*, 17(4):462–478, 2005.