

THE “ART OF TRELLIS DECODING” IS FIXED-PARAMETER TRACTABLE

JISU JEONG, EUN JUNG KIM, AND SANG-IL OUM

ABSTRACT. Given n subspaces of a finite-dimensional vector space over a fixed finite field \mathbb{F} , we wish to find a linear layout V_1, V_2, \dots, V_n of the subspaces such that $\dim((V_1 + V_2 + \dots + V_i) \cap (V_{i+1} + \dots + V_n)) \leq k$ for all i ; such a linear layout is said to have width at most k . When restricted to 1-dimensional subspaces, this problem is equivalent to computing the trellis-width (or minimum trellis state-complexity) of a linear code in coding theory and computing the path-width of an \mathbb{F} -represented matroid in matroid theory.

We present a fixed-parameter tractable algorithm to construct a linear layout of width at most k , if it exists, for input subspaces of a finite-dimensional vector space over \mathbb{F} . As corollaries, we obtain a fixed-parameter tractable algorithm to produce a path-decomposition of width at most k for an input \mathbb{F} -represented matroid of path-width at most k , and a fixed-parameter tractable algorithm to find a linear rank-decomposition of width at most k for an input graph of linear rank-width at most k . In both corollaries, no such algorithms were known previously. Our approach is based on dynamic programming combined with the idea developed by Bodlaender and Kloks (1996) for their work on path-width and tree-width of graphs.

It was previously known that a fixed-parameter tractable algorithm exists for the decision version of the problem for matroid path-width; a theorem by Geelen, Gerards, and Whittle (2002) implies that for each fixed finite field \mathbb{F} , there are finitely many forbidden \mathbb{F} -representable minors for the class of matroids of path-width at most k . An algorithm by Hliněný (2006) can detect a minor in an input \mathbb{F} -represented matroid of bounded branch-width. However, this indirect approach would not produce an actual path-decomposition. Our algorithm is the first one to construct such a path-decomposition and does not depend on the finiteness of forbidden minors.

1. INTRODUCTION

In telecommunication, a message sent through a communication channel is exposed to noise which corrupts the original message. Decoding the original message from a received signal stream is a major issue in coding theory. A linear code is an error-correcting code, for which the alphabet is an element of a finite field \mathbb{F} and its codewords form a subspace of a fixed vector space over \mathbb{F} . Many algorithms for decoding a linear (block or convolutional) code rely on a graphic interpretation of a linear code called the trellis. A *trellis* for a linear code C of length n is a layered directed graph such that the following hold:

Date: March 1, 2017.

Key words and phrases. parameterized complexity, branch-width, path-width, matroid, trellis-width, trellis state-complexity, linear code, linear rank-width, linear clique-width.

The first and last authors are supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science, ICT & Future Planning (2011-0011653).

This paper was presented in part at the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016) [25].




FIGURE 1. Equivalent linear codes can produce trellises of different sizes [53].

- The vertex set is partitioned into V_0, V_1, \dots, V_n .
- Every edge starts at a vertex in V_{i-1} and ends at a vertex in V_i for some $i \in \{1, 2, \dots, n\}$.
- Every vertex is on a path from a vertex in V_0 to a vertex in V_n .
- Every edge is labeled by an element of \mathbb{F} .
- Each path from V_0 to V_n corresponds to a codeword in C .

The weight of an edge is defined by the square of the error, that is the difference between the label on the edge and the received signal at position i . Many decoding algorithms, such as the monumental algorithm by Viterbi [54], for a linear (block or convolutional) code compute a minimum-weight path from V_0 to V_n on the trellis. Such a path corresponds to the most likely original message.

Obtaining a small trellis is crucial in reducing the complexity of a trellis-based decoding algorithm, because its space complexity depends on $(\log \text{ of }) \max_{0 \leq i \leq n} |V_i|$. By permuting the coordinates of a linear code, one may produce an equivalent code with a smaller trellis. Massey [32] asks whether one can find a good permutation of the coordinates to make a smallest trellis, calling it the “art of trellis decoding” of a linear code. The example in Figure 1 from the excellent survey of Vardy [53] on trellises depicts two equivalent linear codes. The first one corresponds to the linear code C generated by the basis $(1, 0, 0, 0, 0, 1)$, $(0, 1, 0, 1, 0, 0)$, and $(0, 0, 1, 0, 1, 0)$ and the second code C' is generated by $(1, 1, 0, 0, 0, 0)$, $(0, 0, 1, 1, 0, 0)$, and $(0, 0, 0, 0, 1, 1)$, obtained by a permutation $\pi = (2, 3, 6)$ on coordinates. The *trellis-width*, also called the *trellis-state complexity*, of a linear code measures the minimum bit size of memory registers required for the decoding after reordering the coordinates.

The *trellis-width of a linear code* can be stated equivalently in terms of vectors over a finite field; if a linear code is generated by a matrix $(v_1 \ v_2 \ \dots \ v_n)$, then its trellis-width can be determined by finding a desired linear layout of the vectors v_1, v_2, \dots, v_n as follows.

Input: n vectors v_1, \dots, v_n in a vector space over a fixed finite field \mathbb{F} , an integer k .

Parameter: k .

Problem: Determine whether there exists a permutation σ of $\{1, 2, \dots, n\}$ such that

$$\dim \langle v_{\sigma(1)}, v_{\sigma(2)}, \dots, v_{\sigma(i)} \rangle \cap \langle v_{\sigma(i+1)}, v_{\sigma(i+2)}, \dots, v_{\sigma(n)} \rangle \leq k$$

for all $1 \leq i \leq n - 1$ and find a linear layout $v_{\sigma(1)}, v_{\sigma(2)}, \dots, v_{\sigma(n)}$ if one exists.

The minimum k having such a linear layout is called the *trellis-width* of a linear code generated by a matrix $(v_1 \ v_2 \ \cdots \ v_n)$. This problem and its equivalent variants have been studied under various names in the coding theory literature, such as MAXIMUM PARTITION RANK PERMUTATION [23], MAXIMUM WIDTH [24], and TRELLIS STATE-COMPLEXITY [53]. In matroid theory, such a linear layout is called a *path-decomposition* (of width at most k) of a matroid M represented by vectors v_1, v_2, \dots, v_n and the minimum such k is called the *path-width* of the matroid M .

Kashyap [28] proved that it is NP-complete to decide whether such a linear layout exists when k is a part of the input. To prove its hardness, he first observed that the trellis-width is actually equal to the path-width of a matroid represented by the same set of vectors. Path-width has been widely studied in structural graph theory, notably by Robertson and Seymour [41], and is also investigated in the context of matroid theory, for instance, in [18]. This link allowed him to deduce the NP-completeness of trellis-width from the NP-completeness of the path-width of graphs.

When k is a fixed constant, then the problem is solvable in polynomial time, thanks to a simple and general algorithm by Nagamochi [33]. His algorithm can decide in time $O(n^{ck})$ whether such a linear layout exists, even if the input matroid is given by an independence oracle.

Our contribution. We present a *fixed-parameter tractable* algorithm to decide whether such a linear layout exists when k is a parameter.¹ In other words, we have an algorithm that runs in time $f(k)n^c$ for some function f and a constant c when n is the number of elements. Our algorithm does not only decide whether such a linear layout exists, but also output such a linear layout if one exists.

More precisely we will prove the following theorem.

Theorem 1.1. *Let \mathbb{F} be a fixed finite field. There exists, for some function f , an $O(f(k)n^3)$ -time algorithm that, for input vectors v_1, v_2, \dots, v_n in \mathbb{F}^r for some $r \leq n$ and a parameter k , either finds a linear layout $v_{\sigma(1)}, v_{\sigma(2)}, \dots, v_{\sigma(n)}$ for a permutation σ on $\{1, 2, \dots, n\}$ such that*

$$\dim\langle v_{\sigma(1)}, v_{\sigma(2)}, \dots, v_{\sigma(i)} \rangle \cap \langle v_{\sigma(i+1)}, v_{\sigma(i+2)}, \dots, v_{\sigma(n)} \rangle \leq k$$

for all $i = 1, 2, \dots, n - 1$ or confirms that no such linear layout exists.

Equivalently we have the following.

Theorem 7.2. *Let \mathbb{F} be a fixed finite field. There is an algorithm that, for an input n -element matroid given by its matrix representation over \mathbb{F} having at most n rows and a parameter k , decides in time $O(f(k)n^3)$ for some function f whether its path-width is at most k and if so, outputs a path-decomposition of width at most k .*

The existence of a fixed-parameter tractable algorithm for the decision version of the problem is not new, but no prior algorithm was able to construct a linear layout of width at most k even if one exists. For the existence of a decision algorithm we use branch-width. Branch-width of a graph or a matroid is a width parameter, introduced by Robertson and

¹In 2012, van Bevern, Downey, Fellows, Gaspers, and Rosamond [51] posted a manuscript on arXiv claiming such a fixed-parameter tractable algorithm. However, their reduction to the problem of *hypergraph cutwidth*, which is proven to be fixed-parameter tractable, had a subtle bug. So their journal paper does not mention trellis-width [52].

Seymour [42], that measures how easy it is to decompose a graph or a matroid into a tree-like structure by cutting through a separation of small order. In particular, the path-width of a matroid is always greater than or equal to its branch-width. Geelen, Gerards, and Whittle [16, 17] proved that for a fixed finite field \mathbb{F} , \mathbb{F} -representable matroids of bounded branch-width are well-quasi-ordered by the matroid minor relation. (Recently they announced a stronger theorem that does not require bounded branch-width [14].) As a corollary, we deduce that there are finitely many \mathbb{F} -representable minor obstructions for the class of matroids of path-width at most k , because all those minor obstructions have branch-width at most $k + 1$. Hliněný [20] proved that for a fixed matroid N , it is possible to decide in time $f(k)n^3$ whether an input n -element matroid given by a matrix representation over \mathbb{F} having branch-width at most k has a minor isomorphic to N . As path-width is always greater than or equal to branch-width, by checking the existence of obstructions as a minor, we can decide whether the path-width of an input matroid given by a matrix representation over \mathbb{F} is at most k . But this approach has the following weaknesses.

- There are a lot of minor obstructions. Koutsonas, Thilikos, and Yamazaki [29] proved that there are at least $(k!)^2$ minor obstructions for matroid path-width at most k in any field \mathbb{F} . Furthermore there is no algorithm known to generate all \mathbb{F} -representable minor obstructions.²
- Even if we know the complete list of minor obstructions, it is still non-trivial to construct a linear layout of width at most k when it exists.³ This is because non-existence of forbidden minors in the input matroid does not provide any hint of how to construct a linear layout.

The current situation for matroid path-width is somewhat similar to the status for path-width of graphs 25 years ago. For a fixed constant k , the problem of deciding whether a graph has path-width at most k has been studied by various researchers. In 1983, Ellis, Sudborough, and Turner [11] presented an $O(n^{2k^2+4k+8})$ -time algorithm. Robertson and Seymour [43] proved the existence of a fixed-parameter tractable $O(f(k)n^2)$ -time algorithm based on the finiteness of minor obstructions, thus only solving the decision problem. Fellows and Langston [12] developed a *self-reduction* technique to convert such a decision algorithm into a construction algorithm and therefore a path-decomposition of a graph witnessing small path-width can be found in time $f(k)n^c$. However $f(k)$ depends on the number of minor obstructions, which is only known to be finite. For matroid path-width, we do not know any self-reduction algorithm that converts a decision algorithm into a construction algorithm.⁴ Bodlaender and Kloks [5] proved the first constructive algorithm for path-width and tree-width of graphs based on dynamic programming, developing important techniques to be used in other papers [3, 6, 31, 47, 48, 49].

Extension to subspace arrangements. In fact, we present an algorithm for a more general problem as described below.

²Kanté and Kwon have claimed that they proved an upper bound of the size of each minor obstruction as a function of k and $|\mathbb{F}|$ in a manuscript posted on arXiv on December 2014, but later we were told by the authors that the proof is incorrect and not yet fixed [private communication, 2015].

³Hliněný [21] claimed that this is possible by the self-reduction technique.

⁴For matroid branch-width, Hliněný and Oum [22] devised a self-reduction algorithm to convert a decision algorithm to a construction algorithm, but this was based on a lemma on *titanic* sets, originated by Robertson and Seymour [42], which works well with branch-width, not path-width.

Input: n subspaces V_1, V_2, \dots, V_n of a vector space \mathbb{F}^r over a fixed finite field \mathbb{F} , an integer k .

Parameter: k .

Problem: Determine whether there exists a permutation σ of $\{1, 2, \dots, n\}$ such that

$$\dim\langle V_{\sigma(1)}, V_{\sigma(2)}, \dots, V_{\sigma(i)} \rangle \cap \langle V_{\sigma(i+1)}, V_{\sigma(i+2)}, \dots, V_{\sigma(n)} \rangle \leq k$$

for all $1 \leq i \leq n - 1$ and find a linear layout $V_{\sigma(1)}, V_{\sigma(2)}, \dots, V_{\sigma(n)}$ if one exists.

A set of subspaces is called a *subspace arrangement* and let us call minimum such k the *path-width* of a subspace arrangement. Clearly if we restrict input subspaces to have dimension at most 1, then this problem is equivalent to the problem on vectors. The following theorem is our main result.

Theorem 6.3. *Let \mathbb{F} be a fixed finite field. Given, as an input, n subspaces of \mathbb{F}^r for some r and a parameter k , in time $O(f(k)n^3 + rm^2)$ for some function f , we can either find a linear layout V_1, V_2, \dots, V_n of the subspaces such that*

$$\dim(V_1 + V_2 + \dots + V_i) \cap (V_{i+1} + V_{i+2} + \dots + V_n) \leq k$$

for all $i = 1, 2, \dots, n - 1$, or confirm that no such linear layout exists, where each V_i is given by its spanning set of d_i vectors and $m = \sum_{i=1}^n d_i$.

Our contribution to rank-width of graphs. The linear rank-width of a graph is a linearized variant of the rank-width introduced by Oum and Seymour [37]. The linear rank-width of an n -vertex graph is the smallest k such that the vertices can be arranged in a linear order v_1, v_2, \dots, v_n so that the rank of the $i \times (n - i)$ 0-1 matrix (over the binary field) representing the adjacency between $\{v_1, v_2, \dots, v_i\}$ and $\{v_{i+1}, v_{i+2}, \dots, v_n\}$ is at most k for all $i = 1, 2, \dots, n - 1$. Deciding whether the linear rank-width of a graph is at most k is NP-complete when k is a part of the input; Proposition 3.1 in [34] shows that the problem of determining the linear rank-width of a bipartite graph is equivalent to the problem of determining the path-width of a binary matroid, which is shown to be NP-complete by Kashyap [28]. However, if k is a fixed parameter, then the following two theorems imply that there exists an $O(f(k)n^3)$ -time algorithm to decide whether the linear rank-width is at most k :

- (1) Graphs of bounded rank-width are well-quasi-ordered by the vertex-minor relation, shown by Oum [36]⁵. Suppose that F_k is a minimal set of graphs such that a graph G has linear rank-width at most k if and only if no graph in F_k is isomorphic to a vertex-minor of G . Then it can be easily seen that graphs in F_k have linear rank-width exactly $k + 1$, implying that their rank-width are at most $k + 1$. By the well-quasi-ordering theorem of graphs of bounded rank-width, we deduce that F_k is finite.
- (2) If F is a fixed graph, then one can decide whether the input n -vertex graph has a vertex-minor isomorphic to F or confirm that the rank-width is larger than k in time $O(g(k)n^3)$ implied by Courcelle and Oum [10].

Jeong, Kwon, and Oum [26] showed that $|F_k| \geq 2^{\Omega(3^k)}$ but there is no known upper bound on the size of graphs in F_k . Furthermore, there is no known algorithm to generate F_k and even if we know the list F_k , it still does not produce a linear rank-decomposition.

⁵This theorem is a generalization of the well-quasi-ordering theorem of Geelen, Gerards, and Whittle [16] for \mathbb{F} -representable matroids of bounded branch-width where \mathbb{F} is a fixed finite field.

The following theorem presents the first fixed-parameter tractable algorithm that provides a linear rank-decomposition of width at most k for graphs of linear rank-width at most k . It will be proved by reducing this problem to the problem on path-width for a subspace arrangement consisting of 2-dimensional subspaces arising from a graph, discussed in Section 8.

Theorem 8.3. *For an input n -vertex graph and a parameter k , we can decide in time $O(f(k)n^3)$ for some function f whether its linear rank-width is at most k and if so, find a linear rank-decomposition of width at most k .*

We would like to mention that the above theorem has an application to linear clique-width. The linear clique-width of a graph is a variant of clique-width introduced by Courcelle and Olariu [9], which was the main motivation to introduce rank-width of graphs. In Subsection 8.4, we will present definitions of the linear clique-width and the linear k -expression, which is a certificate showing a graph has linear clique-width at most k . Fellows, Rosamond, Rotics, and Szeider [13] showed that it is NP-complete to decide whether a graph has linear clique-width at most k when k is a part of the input. However, when k is a fixed parameter, then by Theorem 8.3, we obtain the following approximation algorithm for linear clique-width, which was not known previously.

Corollary 8.5. *For an input n -vertex graph G and a parameter k , we can find a linear $(2^k + 1)$ -expression of G confirming that G has linear clique-width at most $2^k + 1$ or certify that G has linear clique-width larger than k in time $O(f(k)n^3)$ for some function f .*

Computing path-width exactly when branch-width is bounded. We also present a polynomial-time algorithm that can

- compute the path-width of an \mathbb{F} -represented matroid of bounded branch-width,
- compute the linear rank-width of a graph of bounded rank-width.

This is analogous to the polynomial-time algorithm to compute path-width of a graph of bounded tree-width by Bodlaender and Kloks [5] and generalizes polynomial-time algorithms to compute the path-width of the cycle matroid of outerplanar graphs by Koutsonas, Thilikos, and Yamazaki [29] and compute the linear rank-width of graphs of rank-width 1 by Adler, Kanté, and Kwon [1].

Proof overview. We assume that we are given a branch-decomposition of small width and then use dynamic programming to solve the problem based on the tree-like structure given by the branch-decomposition. Courcelle's Theorem [8] reveals what typically makes a graph problem amenable for dynamic programming on tree-like structure. In his seminal work, Courcelle proved that every graph property that can be expressed as a monodic second-order formula ϕ admits a finite number of *states* in the following sense: when two graphs G and H are *glued* along w common vertices, we can decide whether the glued graph satisfies the formula or not by looking at their respective *behavior* on the w vertices with respect to the formula. As the number of possible behaviors over all graphs, or *states*, on w vertices with respect to ϕ is bounded by a fixed function $f(w, |\phi|)$, remembering a small set of states available for G on w vertices is sufficient to decide whether G glued with H shall satisfy ϕ or not. Especially when an input graph can be recursively decomposed along vertex separators of size at most w , a fixed graph property ϕ can be tested on the input in linear time. Courcelle's Theorem also applies to graph optimization problems. While Courcelle's Theorem provides a powerful meta-algorithm, the downside is that the

function f that plays as a hidden constant in $O(n)$ is gigantic in general, or the algorithm might not be constructible. The novelty of non-trivial dynamic programming algorithm lies in devising an encoding scheme of a behavior of a graph (with a partial solution) on a small separator.

Our proof strategy is similar to the one employed by Bodlaender and Kloks [5]. The algorithm of Bodlaender and Kloks is based on dynamic programming on tree-decomposition of small width. We also use the idea of *typical sequences* introduced by Bodlaender and Kloks [5] to encode and compress the amount of information needed to track partial solutions.⁶ One of our key contribution is to come up with an encoding of partial solutions with respect to a fixed subspace. In a nutshell, we propose an encoding scheme of a linear layout of a subspace arrangement \mathcal{V} with respect to a subspace B as follows: for every linear layout σ of a subspace arrangement \mathcal{W} with $\text{span}(\mathcal{V}) \cap \text{span}(\mathcal{W}) \subseteq B$,

if two linear layouts σ_1 and σ_2 of \mathcal{V} have the same encoding Γ with respect to B , then the width of a layout π_1 of $\mathcal{V} \dot{\cup} \mathcal{W}$ extending σ_1 and σ is equal to the width of a layout π_2 of $\mathcal{V} \dot{\cup} \mathcal{W}$ extending σ_2 and σ , as long as π_1 and π_2 mix σ_i and σ the same way for both $i = 1, 2$ in relation to Γ .

This property⁷ allows us to safely forget the specific linear layout of \mathcal{V} that gives rise to an encoding Γ , and remember Γ on B only. Also, due to this property, it is enough to consider the possibilities of mixing two encodings instead of considering all possible ways of mixing two linear layouts: the set of achievable widths remains the same in both cases. What we mean by *mixing* shall become clear in Subsection 3.4 with the notion of *lattice path*. Unlike usual encodings of partial solutions on graphs, our encoding needs to handle subspaces and keep track of how subspaces interact. We develop an encoding scheme, operations and relations on encodings to this aim. We also present a framework to perform dynamic programming on branch-decompositions of a subspace arrangement.⁸

The paper is organized as follows. Section 2 will review basic concepts and a framework to perform dynamic programming on branch-decompositions. Section 3 will introduce B -trajectories, which will be a key concept for our algorithm. Section 4 will define full sets consisting of B -trajectories and show several propositions on combining two full sets coming from child nodes of a node to obtain the full set. We will present our algorithm in Section 5 based on the operations discussed in Section 4. Since our algorithm is based on dynamic programming, up to Section 5 we assume that some branch-decomposition is given. Section 6 will discuss two methods of providing this branch-decomposition. Section 7 presents a corollary to path-width of matroids and Section 8 will discuss corollaries to linear rank-width and linear clique-width of graphs. Section 9 will present an exact polynomial-time algorithm to compute path-width of a matroid when the input has small

⁶A similar attempt to adapt the method of Bodlaender and Kloks to a wide class of width parameters was made by Berthomé et al. [2] (see also the Ph.D. thesis of Soares [46, Theorem 11]). One may ask whether their algorithm implies our theorems. However, their theorem requires a “nice decomposition” of small width “compatible” with the partition function corresponding to branch-width of matroids or subspace arrangements to be given as an input in order to run dynamic programming. Since their definition of “compatible nice decompositions” is very strong, the existence of a compatible nice decomposition is not guaranteed, even if we have bounded path-width for matroids or subspace arrangements. Thus it is unlikely that their theorem contains ours.

⁷We do not formally prove this property, but it can be easily deduced from the definition of B -trajectory and Lemma 3.18.

⁸Hliněný [20] presented such a framework for represented matroids.

branch-width and its corollary to linear rank-width of graphs. Section 10 concludes this paper by presenting some open problems on matroids.

2. PRELIMINARIES

We assume readers are familiar with linear algebra. For a set X of vectors, we write $\langle X \rangle$ to denote the *span* of X , which is the set of all vectors that are linear combinations of vectors in X . We say that X is a *spanning set* of a vector space S if $\langle X \rangle = S$. For simplicity, we also write $\langle v_1, v_2, \dots, v_n \rangle = \langle \{v_1, v_2, \dots, v_n\} \rangle$ for a sequence of vectors v_1, v_2, \dots, v_n .

A *subspace arrangement* \mathcal{V} over a field \mathbb{F} is a finite set of subspaces V_1, V_2, \dots, V_n of a finite-dimensional vector space over \mathbb{F} . Formally a subspace arrangement is an indexed family $\mathcal{V} = \{V_i\}_{i \in E}$ of subspaces for a finite set E . For two subspaces A, B , we write $A + B = \{a + b : a \in A, b \in B\}$. For a subspace arrangement \mathcal{V} , let us write $\langle \mathcal{V} \rangle$ to denote $\sum_{V \in \mathcal{V}} V$.

2.1. Path-width and branch-width of a subspace arrangement. Let \mathcal{V} be a subspace arrangement over a field \mathbb{F} . A *linear layout* of \mathcal{V} is a permutation $\sigma = V_1, V_2, \dots, V_n$ of \mathcal{V} . The width of a linear layout $\sigma = V_1, V_2, \dots, V_n$ is

$$\max_{1 \leq i < n} \dim(V_1 + V_2 + \dots + V_i) \cap (V_{i+1} + V_{i+2} + \dots + V_n).$$

The *path-width* of \mathcal{V} is the minimum width of all possible linear layouts of \mathcal{V} . (If $|\mathcal{V}| \leq 1$, then the width of its trivial linear layout of \mathcal{V} is 0 and the path-width is 0.)

We now define the branch-width of \mathcal{V} . We say a tree is *subcubic* if all its internal vertices have degree 3. A *branch-decomposition* of \mathcal{V} is a pair (T, \mathcal{L}) of a subcubic tree T and a bijection \mathcal{L} from the set of leaves of T to \mathcal{V} . For a set X of leaves of T , let $\mathcal{L}(X) = \{\mathcal{L}(x) : x \in X\}$. Each edge e induces a partition (A_e, B_e) of the leaves given by $T - e$ and the *width* of e is defined to be

$$\dim \langle \mathcal{L}(A_e) \rangle \cap \langle \mathcal{L}(B_e) \rangle.$$

The *width* of a branch-decomposition (T, \mathcal{L}) is the maximum width of all edges of T . The *branch-width* of \mathcal{V} is the minimum width of all possible branch-decompositions of \mathcal{V} . (If $|\mathcal{V}| \leq 1$, then there is no branch-decomposition and we define the branch-width of \mathcal{V} to be 0.)

Proposition 2.1. *Let \mathcal{V} be a subspace arrangement. For a subset X of \mathcal{V} , let $f(X) = \dim \langle X \rangle \cap \langle \mathcal{V} - X \rangle$. Then f satisfies the following.*

- (i) $f(X) = f(\mathcal{V} - X)$ for all $X \subseteq \mathcal{V}$. (symmetric)
- (ii) $f(X) + f(Y) \geq f(X \cap Y) + f(X \cup Y)$ for all $X, Y \subseteq \mathcal{V}$. (submodular)

Proof. (i) is trivial. Let us prove (ii). It is enough to prove the following.

$$\begin{aligned} & \dim \langle X \rangle + \dim \langle \mathcal{V} - X \rangle + \dim \langle Y \rangle + \dim \langle \mathcal{V} - Y \rangle \\ & \geq \dim \langle X \cap Y \rangle + \dim \langle \mathcal{V} - (X \cap Y) \rangle + \dim \langle X \cup Y \rangle + \dim \langle \mathcal{V} - (X \cup Y) \rangle. \end{aligned}$$

This follows from the facts that $\dim(\langle A \rangle + \langle B \rangle) = \dim \langle A \cup B \rangle$, $\dim(\langle A \rangle \cap \langle B \rangle) \geq \dim \langle A \cap B \rangle$, and $\dim \langle A \rangle + \dim \langle B \rangle = \dim(\langle A \rangle + \langle B \rangle) + \dim \langle A \rangle \cap \langle B \rangle$ for all $A, B \subseteq \mathcal{V}$. \square

Proposition 2.2. *If \mathcal{V} is a subspace arrangement of path-width at most k , then*

$$\dim X \cap \langle \mathcal{V} - \{X\} \rangle \leq 2k$$

for all $X \in \mathcal{V}$ and the branch-width of \mathcal{V} is at most $2k$.

Proof. Let V_1, V_2, \dots, V_n be a linear layout of \mathcal{V} of width at most k . Clearly $\dim V_1 \cap (\mathcal{V} - \{V_1\}) \leq k$ and $\dim V_n \cap (\mathcal{V} - \{V_n\}) \leq k$. For $i \in \{2, 3, \dots, n-1\}$,

$$\begin{aligned} \dim(V_1 + \dots + V_i) \cap (V_{i+1} + \dots + V_n) &\leq k, \\ \dim(V_i + V_{i+1} + \dots + V_n) \cap (V_1 + \dots + V_{i-1}) &\leq k. \end{aligned}$$

By Proposition 2.1, we deduce that $\dim V_i \cap \langle \mathcal{V} - \{V_i\} \rangle \leq 2k$ because $\{V_1, V_2, \dots, V_i\} \cap \{V_i, V_{i+1}, \dots, V_n\} = \{V_i\}$. This proves that $\dim X \cap \langle \mathcal{V} - \{X\} \rangle \leq 2k$ for all $X \in \mathcal{V}$.

Let T be a tree obtained from the path graph P on n vertices by attaching one leaf to each internal node of the path. Note that T has n leaves. Let \mathcal{L} be a function that maps all leaves of T to V_1, V_2, \dots, V_n bijectively in the order of the path. Then (T, \mathcal{L}) is a branch-decomposition of \mathcal{V} . For every edge e in P , the width of e is clearly at most k by the assumption and so (T, \mathcal{L}) has width at most $2k$. \square

2.2. Sequence of integers. Bodlaender and Kloks [5] introduced typical sequences for their works on path-width and tree-width. For an integer $n > 0$, let $\alpha = (a_1, a_2, \dots, a_n)$ be a finite sequence of integers of length n .

For a sequence α , we define its *typical sequence* $\tau(\alpha)$ to be the sequence obtained by iterating the following *compression* operations until no further operations can be applied:

- If two consecutive entries are equal, then remove one.
- If there exist i and j such that $j - i > 1$ and $a_i \leq a_k \leq a_j$ for all $k \in \{i+1, i+2, \dots, j-1\}$ or $a_i \geq a_k \geq a_j$ for all $k \in \{i+1, i+2, \dots, j-1\}$, then remove all entries from the $(i+1)$ -th entry to the $(j-1)$ -th entry.

A sequence α is called *typical* if $\tau(\alpha) = \alpha$. For example, if $\alpha = (1, 3, 2, 5, 2, 2, 4, 4, 3)$, then $\tau(\alpha) = (1, 5, 2, 4, 3)$.

Bodlaender and Kloks [5] not only showed that $\tau(\alpha)$ is uniquely defined but also proved that the length of typical sequences of bounded integers and the number of distinct typical sequences of bounded integers are bounded as follows. These lemmas are fundamental to our applications.

Lemma 2.3 (Bodlaender and Kloks [5, Lemma 3.3]). *The length of a typical sequence of integers consisting of $\{0, 1, \dots, k\}$ is at most $2k + 1$.*

The following lemma was originally proved for the number of typical sequences only. In fact, the proof is constructive and it is not difficult to see that all sequences can be generated with polynomial delay. Here, we will base our algorithm on a weaker result.

Lemma 2.4 (Bodlaender and Kloks [5, Lemma 3.5]). *There are at most $\frac{8}{3}2^{2k}$ distinct typical sequences of integers consisting of $\{0, 1, \dots, k\}$. Furthermore, the set of all typical sequences can be generated in $\text{poly}(k) \cdot 2^{2k}$ steps.*

2.3. Manipulating subspaces. We present the data structure to store subspaces and review algorithms performing elementary operations on subspaces.

Let \mathcal{B} be an ordered basis of a subspace $S \subseteq \mathbb{F}^r$ of dimension d . Let $M_S \in \mathbb{F}^{r \times d}$ be an $r \times d$ matrix whose column vectors are vectors in \mathcal{B} .

In our application, d will be usually a lot smaller than r and so we will use the coordinate system with respect to \mathcal{B} in order to represent a vector in S . In other words, for each vector $w \in S \subseteq \mathbb{F}^r$, we write $[w]_{\mathcal{B}}$ to denote the coordinate vector with respect to \mathcal{B} as a column vector. It is easy to obtain the coordinates $w \in \mathbb{F}^r$ by multiplying M_S with $[w]_{\mathcal{B}}$.

By using this coordinate system, a d' -dimensional subspace S' of S can be represented by a $d \times d'$ matrix $M_{S,S'} \in \mathbb{F}^{d \times d'}$ in such a way that the range of $M_S M_{S,S'}$ is exactly S' . If so, then we say that $M_{S,S'}$ represents S' .

It is well known that one can solve a system of linear equations in time $O(n^3)$ by Gauss-Jordan elimination when the number of variables and the number of equations are at most n . We summarize the time complexity of computing a matrix representing a subspace as follows.

Lemma 2.5. *Let S be a d -dimensional subspace of \mathbb{F}^r and S' be a subspace of S . Let \mathcal{B} be an ordered basis of S . Given a matrix $M_{S,S'}$ representing S' and a vector x in S , when x is given as a coordinate vector $[x]_{\mathcal{B}}$ in \mathbb{F}^d with respect to \mathcal{B} , we can decide $x \in S'$ in time $O(d^3)$.*

Proof. We can decide $x \in S'$ by checking whether $[x]_{\mathcal{B}}$ is a linear combination of columns in $M_{S,S'}$. \square

Lemma 2.6. *Let S be a d -dimensional subspace of \mathbb{F}^r . Given, as an input, matrices M_{S,S_i} representing subspaces S_i of S for $i = 1, 2$, we can*

- (1) *decide whether $S_1 \subseteq S_2$ in time $O(d^3)$,*
- (2) *compute a matrix $M_{S,S'}$ representing $S' = S_1 + S_2$ in time $O(d^3)$, and*
- (3) *compute a matrix $M_{S,S''}$ representing $S'' = S_1 \cap S_2$ in time $O(d^3)$.*

Proof. (1) This can be answered by solving a matrix equation $M_{S,S_2} X = M_{S,S_1}$; this is feasible if and only if $S_1 \subseteq S_2$. (2) This can be done by finding a column basis of $(M_{S,S_1} \ M_{S,S_2})$. (3) We find a matrix $\begin{pmatrix} M_x \\ M_y \end{pmatrix}$ to represent the space of all vectors $\begin{pmatrix} x \\ y \end{pmatrix}$ satisfying $(M_{S,S_1} \ -M_{S,S_2}) \begin{pmatrix} x \\ y \end{pmatrix} = 0$ where $x \in \mathbb{F}^{\dim S_1}$ and $y \in \mathbb{F}^{\dim S_2}$. Then $M_{S,S''}$ is obtained from $M_{S,S_1} M_x$ by deleting linearly dependent column vectors. \square

2.4. Dynamic programming on branch-decompositions. For many discrete objects, tree-like structures allow us to view the given input structure as a composition of simple objects in a tree-like structure, making it very efficient to run algorithms based on dynamic programming. Tree-decompositions of graphs have been widely used as a decomposition tree for the dynamic programming algorithm.

For matroids, Hliněný [20] introduced “matroid parse trees” for the purpose of dynamic programming on matroids of bounded branch-width, represented over a fixed finite field. We are going to provide a similar yet different framework for the same purpose that is easier for us to handle.

Let (T, \mathcal{L}) be a branch-decomposition of a subspace arrangement \mathcal{V} having width at most θ . We may assume that T is a rooted binary tree by picking an arbitrary edge e and subdividing e to create a degree-2 vertex called the *root*. By orienting each edge of T towards the root, we may further assume that T is a rooted binary directed tree. For a node v of T , let \mathcal{V}_v be the set of all subspaces in \mathcal{V} associated with v and its descendants by \mathcal{L} . We define the *boundary space* B_v as $\langle \mathcal{V}_v \rangle \cap \langle \mathcal{V} - \mathcal{V}_v \rangle$. From now on, when we call (T, \mathcal{L}) a branch-decomposition, T is always a rooted binary directed tree.

The following easy lemma is essential for our dynamic programming. It is well known that $\dim(X) + \dim(Y) = \dim(X + Y) + \dim(X \cap Y)$ for finite-dimensional subspaces X, Y .

Lemma 2.7. *Let v be an internal node of T and w_1, w_2 be two children of v . Let B, B_1, B_2 be the boundary space of v, w_1, w_2 , respectively. Then*

$$\begin{aligned} \langle \mathcal{V}_{w_1} \rangle \cap B_1 &= B_1 = \langle \mathcal{V}_{w_1} \rangle \cap (B_1 + B_2), \\ \langle \mathcal{V}_{w_2} \rangle \cap B_2 &= B_2 = \langle \mathcal{V}_{w_2} \rangle \cap (B_1 + B_2), \\ B &\subseteq B_1 + B_2, \\ (\langle \mathcal{V}_{w_1} \rangle + B_1 + B_2) \cap (\langle \mathcal{V}_{w_2} \rangle + B_1 + B_2) &= B_1 + B_2. \end{aligned}$$

Proof. Let $M_w = \langle \mathcal{V}_w \rangle$ for each node w of T . Since $M_{w_1} = \langle \mathcal{V}_{w_1} \rangle$ and $B_1 = \langle \mathcal{V}_{w_1} \rangle \cap \langle \mathcal{V} - \mathcal{V}_{w_1} \rangle$, trivially $M_{w_1} \cap B_1 = B_1 \subseteq M_{w_1} \cap (B_1 + B_2)$. Note that $\mathcal{V}_{w_1} \subseteq \mathcal{V} - \mathcal{V}_{w_2}$ and $\mathcal{V}_{w_2} \subseteq \mathcal{V} - \mathcal{V}_{w_1}$. Then

$$\begin{aligned} \dim M_{w_1} \cap (B_1 + B_2) &= \dim M_{w_1} + \dim(B_1 + B_2) - \dim(M_{w_1} + B_2) \\ &= \dim B_1 - \dim B_1 \cap B_2 + \dim M_{w_1} \cap B_2 \\ &= \dim B_1 \end{aligned}$$

because $M_{w_1} + B_1 + B_2 = M_{w_1} + B_2$ and $M_{w_1} \cap B_2 = M_{w_1} \cap M_{w_2} = B_1 \cap B_2$. Thus, $B_1 = M_{w_1} \cap (B_1 + B_2)$.

Now we claim that $B \subseteq B_1 + B_2$. Let $b \in B = \langle \mathcal{V}_v \rangle \cap \langle \mathcal{V} - \mathcal{V}_v \rangle$. Clearly $b \in \langle \mathcal{V}_v \rangle = \langle \mathcal{V}_{w_1} \rangle + \langle \mathcal{V}_{w_2} \rangle$. Let $x \in \langle \mathcal{V}_{w_1} \rangle, y \in \langle \mathcal{V}_{w_2} \rangle$ be vectors such that $x + y = b$. Since $b \in \langle \mathcal{V} - \mathcal{V}_v \rangle \subseteq \langle \mathcal{V} - \mathcal{V}_{w_1} \rangle$ and $y \in \langle \mathcal{V}_{w_2} \rangle \subseteq \langle \mathcal{V} - \mathcal{V}_{w_1} \rangle$, we deduce that $x \in \langle \mathcal{V} - \mathcal{V}_{w_1} \rangle$. Similarly, $y \in \langle \mathcal{V} - \mathcal{V}_{w_2} \rangle$. Thus, $b = x + y$ is in $B_1 + B_2$.

Note that $(M_{w_1} + B_1 + B_2) \cap (M_{w_2} + B_1 + B_2) = (M_{w_1} + B_2) \cap (M_{w_2} + B_1)$ and $B_1 \cap B_2 = M_{w_1} \cap M_{w_2}$. Thus

$$\begin{aligned} &\dim(M_{w_1} + B_2) \cap (M_{w_2} + B_1) \\ &= \dim(M_{w_1} + B_2) + \dim(M_{w_2} + B_1) - \dim(M_{w_1} + B_1 + M_{w_2} + B_2) \\ &= \dim(M_{w_1}) + \dim(B_2) - \dim(M_{w_1} \cap B_2) \\ &\quad + \dim(M_{w_2}) + \dim(B_1) - \dim(M_{w_2} \cap B_1) - \dim(M_{w_1} + M_{w_2}) \\ &= \dim(M_{w_1}) + \dim(B_2) - \dim M_{w_1} \cap M_{w_2} \cap \langle \mathcal{V} - \mathcal{V}_{w_2} \rangle \\ &\quad + \dim(M_{w_2}) + \dim(B_1) - \dim M_{w_2} \cap M_{w_1} \cap \langle \mathcal{V} - \mathcal{V}_{w_1} \rangle \\ &\quad - \dim(M_{w_1} + M_{w_2}) \\ &= \dim(M_{w_1}) + \dim(M_{w_2}) + \dim(B_1) + \dim(B_2) \\ &\quad - 2 \dim(M_{w_1} \cap M_{w_2}) - \dim(M_{w_1} + M_{w_2}) \\ &= \dim(B_1) + \dim(B_2) - \dim(M_{w_1} \cap M_{w_2}) \\ &= \dim(B_1 + B_2). \end{aligned}$$

As $B_1 + B_2 \subseteq (M_{w_1} + B_1 + B_2) \cap (M_{w_2} + B_1 + B_2)$, we deduce that $(M_{w_1} + B_1 + B_2) \cap (M_{w_2} + B_1 + B_2) = B_1 + B_2$. \square

3. TRAJECTORIES

Given a branch-decomposition (T, \mathcal{L}) of a subspace arrangement \mathcal{V} over \mathbb{F} whose width is at most θ , let us assume that the boundary space B_v has been computed for every node v of T . Consider \mathcal{V}_v at a node v of T . To design an algorithm based on dynamic programming, we need to keep a bounded (in terms of θ , k , and $|\mathbb{F}|$) amount of information that encodes all “good” linear layouts of \mathcal{V}_v . Suppose that $\sigma = V_1, V_2, \dots, V_m$ is a linear layout of \mathcal{V}_v . To be a “good” partial solution, we need $\dim(V_1 + \dots + V_{i-1}) \cap (V_i + \dots + V_m) \leq k$ for all $2 \leq i \leq m$. What information do we need to keep in order to determine whether this can be extended to a linear layout of \mathcal{V} of width at most k ? Lemmas 3.16 and 3.18 will ensure that we only need to keep

$$\begin{aligned} L_i &:= (V_1 + \dots + V_{i-1}) \cap B_v, \\ R_i &:= (V_i + \dots + V_m) \cap B_v \end{aligned}$$

and the extra dimension of the intersection not shown in B_v , that is

$$\lambda_i := \dim(V_1 + \dots + V_{i-1}) \cap (V_i + \dots + V_m) - \dim(V_1 + \dots + V_{i-1}) \cap (V_i + \dots + V_m) \cap B_v$$

for each $2 \leq i \leq m$. This motivates the definition of a *statistic* that is a triple (L, R, λ) of two subspaces L, R of bounded dimension and one nonnegative integer λ . For instance, each $i \in \{1, 2, \dots, m+1\}$ induces a statistic

$$a_i := (L_i, R_i, \lambda_i)$$

where we define $L_1 = \{0\} = R_{m+1}$, $R_1 = (V_1 + \dots + V_m) \cap B_v = L_{m+1}$, and $\lambda_1 = \lambda_{m+1} = 0$. Then the linear layout $\sigma = V_1, V_2, \dots, V_m$ induces a sequence of statistics

$$\Gamma = a_1, a_2, \dots, a_{m+1},$$

which we call the *canonical B-trajectory* of σ .

While the number of *B-trajectories* can be arbitrarily large, it can be shown that the number of *compact* ones is bounded in terms of θ , k , and $|\mathbb{F}|$. Hence, at each node v of T , we carry over *only* compact B_v -trajectories for dynamic programming.

3.1. B-trajectories. For a vector space B , a *statistic* is a triple $a = (L, R, \lambda)$ of two subspaces L, R of B and a nonnegative integer λ . We write $L(a) = L$, $R(a) = R$, and $\lambda(a) = \lambda$. For two statistics a and b , we write $a \leq b$ if

$$L(a) = L(b), \quad R(a) = R(b), \quad \lambda(a) \leq \lambda(b).$$

A *B-trajectory* of length $n \geq 1$ is a sequence $\Gamma = a_1, a_2, \dots, a_n$ of statistics such that $R(a_1) = L(a_n)$, $L(a_i) \subseteq L(a_{i+1}) \subseteq B$ and $B \supseteq R(a_i) \supseteq R(a_{i+1})$ for all $i = 1, 2, \dots, n-1$. The *width* of Γ is $\max_{1 \leq i \leq n} \lambda(a_i)$. We write $|\Gamma|$ to denote the length n of Γ . We write $\Gamma(i)$ to denote the i -th statistic $(L(a_i), R(a_i), \lambda(a_i))$. A *B-trajectory* Γ^* is an *extension* of a *B-trajectory* Γ if Γ^* is obtained from Γ by repeating some of its entries. The set of all extensions of Γ is denoted as $E(\Gamma)$.

The *compactification* $\tau(\Gamma)$ of $\Gamma = a_1, a_2, \dots, a_n$ is a *B-trajectory* obtained from Γ by applying the following operations until no further operation can be applied.

- Remove an entry a_i if $a_{i-1} = a_i$.
- If there exist i and j such that $j - i > 1$, $L(a_i) = L(a_j)$, $R(a_i) = R(a_j)$ and $\lambda(a_i) \leq \lambda(a_k) \leq \lambda(a_j)$ for all $k \in \{i+1, i+2, \dots, j-1\}$ or $\lambda(a_i) \geq \lambda(a_k) \geq \lambda(a_j)$

for all $k \in \{i+1, i+2, \dots, j-1\}$, then remove $a_{i+1}, a_{i+2}, \dots, a_{j-1}$ from the sequence.

It is clear that the width of $\tau(\Gamma)$ and Γ are equal. We say that Γ is *compact* if $\tau(\Gamma) = \Gamma$.

As the notion of B -trajectories encapsulates sequences of integers, the notion of compact B -trajectories captures typical sequences of integers. For example, when $B = \{0\}$, a B -trajectory Γ is of the form $(\{0\}, \{0\}, \lambda_1), (\{0\}, \{0\}, \lambda_2), \dots, (\{0\}, \{0\}, \lambda_n)$. Then Γ is compact if and only if the sequence $\lambda_1, \lambda_2, \dots, \lambda_n$ is typical. Due to [5], we can show that if B is a vector space over \mathbb{F} of dimension θ , then the length of compact B -trajectories of width at most k and the number of distinct compact B -trajectories of width at most k are bounded from above as follows.

Lemma 3.1. *Let B be a vector space over \mathbb{F} of dimension θ . If $\Gamma = a_1, a_2, \dots, a_m$ is a compact B -trajectory, then the number of distinct pairs $(L(a_i), R(a_i))$ is at most $2\theta + 1$.*

Proof. Let $(L_1, R_1), (L_2, R_2), \dots, (L_p, R_p)$ be distinct pairs of subspaces of B such that $L_1 \subseteq L_2 \subseteq \dots \subseteq L_p$ and $R_1 \supseteq R_2 \supseteq \dots \supseteq R_p$. Then $\dim(L_1) \leq \dots \leq \dim(L_p) \leq \theta$ and $\theta \geq \dim(R_1) \geq \dots \geq \dim(R_p)$. As they are all distinct, $\dim(L_{i+1}) > \dim(L_i)$ or $\dim(R_{i+1}) < \dim(R_i)$ for all i . It follows that

$$-\theta \leq \dim(L_1) - \dim(R_1) < \dim(L_2) - \dim(R_2) < \dots < \dim(L_p) - \dim(R_p) \leq \theta.$$

This implies that $p \leq 2\theta + 1$. □

Lemma 3.2. *Let B be a vector space over \mathbb{F} of dimension θ . Every compact B -trajectory of width k has length at most $(2\theta + 1)(2k + 1)$.*

Proof. Let $\Gamma = a_1, a_2, \dots, a_m$ be a compact B -trajectory. By Lemma 3.1, there are at most $2\theta + 1$ distinct pairs $(L(a_i), R(a_i))$ in Γ . Each maximal consecutive subsequence $\lambda(a_i), \lambda(a_{i+1}), \dots, \lambda(a_j)$ with $L(a_i) = L(a_j)$ and $R(a_i) = R(a_j)$ forms a typical sequence of integers, whose length is bounded by $2k + 1$ by Lemma 2.3. This proves the claim. □

For a vector space B over \mathbb{F} , let $U_k(B)$ be the set of all compact B -trajectories of width at most k . The following lemma will find an upper bound of $|U_k(B)|$. Its proof uses lattice paths. A *lattice path* from $(1, 1)$ to (x, y) is a path $P = v_1, v_2, \dots, v_n$ on the plane such that $v_1 = (1, 1)$, $v_n = (x, y)$, $v_i \in \mathbb{Z} \times \mathbb{Z}$ and $v_{i+1} - v_i \in \{(1, 0), (0, 1)\}$ for all $i = 1, 2, \dots, n-1$.

Lemma 3.3. *Let B be a vector space over \mathbb{F} of dimension θ . Then $U_k(B)$ contains at most $2^{9\theta+2} |\mathbb{F}|^{\theta(\theta-1)} 2^{2(2\theta+1)k}$ elements. Furthermore, the set $U_k(B)$ can be generated in $\text{poly}(\theta, |\mathbb{F}|, k) \cdot |U_k(B)|$ steps.*

Proof. Let $q = |\mathbb{F}|$. First let us assume that $\theta > 0$. The number of length- $(\theta + 1)$ chains $V_0 \subset V_1 \subset \dots \subset V_\theta$ of subspaces of B is precisely

$$\binom{\theta}{\theta-1}_q \binom{\theta-1}{\theta-2}_q \binom{\theta-2}{\theta-3}_q \cdots \binom{1}{0}_q = \prod_{d=1}^{\theta-1} \binom{d+1}{d}_q,$$

where $\binom{d+1}{d}_q = \frac{q^{d+1}-1}{q-1}$, known as the Gaussian binomial coefficients. Since $\binom{d+1}{d}_q = q^d + q^{d-1} + \dots + 1 \leq 2q^d$, we deduce that the number of chains of length $\theta + 1$ is at most $2^{\theta-1} q^{\theta(\theta-1)/2}$.

Each B -trajectory $\Gamma = a_1, a_2, \dots, a_m$ induces a set $\{(L(a_i), R(a_i)) : i = 1, 2, \dots, m\}$ and there exist length- $(\theta + 1)$ chains $L_0 \subset L_1 \subset L_2 \subset \dots \subset L_\theta$, $R_0 \supset R_1 \supset R_2 \supset \dots \supset$

R_θ of subspaces of B such that $L(a_i) \in \{L_0, L_1, \dots, L_\theta\}$ and $R(a_i) \in \{R_0, R_1, \dots, R_\theta\}$. Furthermore there exists a lattice path $(x_0, y_0), (x_1, y_1), \dots, (x_{2\theta}, y_{2\theta})$ from $(0, 0)$ to (θ, θ) such that $(L(a_i), R(a_i)) \in \{(L_{x_i}, R_{y_i}) : i = 0, 1, \dots, 2\theta\}$.

By the previous observation, there are at most $2^{\theta-1}q^{\theta(\theta-1)/2}$ ways of choosing each of $\{L_i\}$ and $\{R_i\}$ and there are $\binom{2\theta}{\theta}$ ways of selecting a lattice path $P = v_0, v_1, v_2, \dots, v_{2\theta}$ from $v_0 = (0, 0)$ to $v_{2\theta} = (\theta, \theta)$ and let $(x_i, y_i) = v_i$. Since $\{(L(a_i), R(a_i)) : i = 1, 2, \dots, m\}$ is a subset of $\{(L_{x_i}, R_{y_i}) : i = 0, 1, \dots, 2\theta\}$, there are at most

$$(2^{\theta-1}q^{\theta(\theta-1)/2})^2 \binom{2\theta}{\theta} 2^{2\theta+1} \leq 2^{6\theta-1}q^{\theta(\theta-1)}$$

distinct sets $\{(L(a_i), R(a_i)) : i = 1, 2, \dots, m\}$ induced by B -trajectories Γ . It is not difficult to see that all such sets $\{(L(a_i), R(a_i)) : i = 1, 2, \dots, m\}$ can be generated in $\text{poly}(k) \cdot 2^{6\theta-1}q^{\theta(\theta-1)}$ for some fixed polynomial function in k .

Each maximal consecutive subsequence a_i, \dots, a_j of a B -trajectory $\Gamma = a_1, \dots, a_n$ with $L(a_i) = L(a_j)$ and $R(a_i) = R(a_j)$ induces a ‘typical sequence’ $\lambda(a_i), \dots, \lambda(a_j)$. There are at most $\frac{8}{3}2^{2k}$ distinct such typical sequences and they can be generated in $\text{poly}(k) \cdot 2^{2k}$ steps by Lemma 2.4.

Therefore, the number of compact B -trajectories of width at most k is at most

$$2^{6\theta-1}q^{\theta(\theta-1)} \left(\frac{8}{3}2^{2k}\right)^{2\theta+1} \leq 2^{9\theta+1}q^{\theta(\theta-1)}2^{2(2\theta+1)k}.$$

If $\theta = 0$, then the number of compact B -trajectories is precisely the number of typical sequences, that is at most $\frac{8}{3}2^{2k} \leq 2^{2 \cdot 2^{2k}}$. The running time for producing $U_k(B)$ follows immediately. \square

Let Γ_1 and Γ_2 be two B -trajectories. We write $\Gamma_1 \leq \Gamma_2$ if Γ_1 and Γ_2 have the same length, say n , and $\Gamma_1(i) \leq \Gamma_2(i)$ for all $1 \leq i \leq n$. We say that $\Gamma_1 \preceq \Gamma_2$ if there are extensions $\Gamma_1^* \in E(\Gamma_1)$ and $\Gamma_2^* \in E(\Gamma_2)$ such that $\Gamma_1^* \leq \Gamma_2^*$.

Lemma 3.4. *Let Γ, Δ be two B -trajectories. If $\Gamma \preceq \Delta$, then the width of Γ is less than or equal to the width of Δ .*

Proof. This is trivial from definitions. \square

A lattice path with diagonal steps from $(1, 1)$ to (x, y) on $\mathbb{Z} \times \mathbb{Z}$ is a sequence v_1, v_2, \dots, v_m of distinct points in $\mathbb{Z} \times \mathbb{Z}$ such that $v_1 = (1, 1)$, $v_m = (x, y)$, and $v_{i+1} - v_i \in \{(1, 0), (0, 1), (1, 1)\}$ for all $i = 1, 2, \dots, m-1$. Obviously the number of lattice paths with diagonal steps from $(1, 1)$ to (x, y) depends only on x and y . The following lemma indicates that comparing two B -trajectories can be done without going through infinitely many extensions.

Lemma 3.5. *Let Γ, Δ be two B -trajectories. Then $\Gamma \preceq \Delta$ if and only if there exists a lattice path $P = v_1, v_2, \dots, v_m$ with diagonal steps from $(1, 1)$ to $(|\Gamma|, |\Delta|)$ such that $\Gamma(x_i) \leq \Delta(y_i)$ when $v_i = (x_i, y_i)$ for each $i = 1, 2, \dots, m$.*

Proof. Suppose that there exist $\Gamma^* \in E(\Gamma)$ and $\Delta^* \in E(\Delta)$ such that $\Gamma^* \leq \Delta^*$. We may assume that $m = |\Gamma^*| (= |\Delta^*|)$ is chosen to be minimum. Let $1 = x_1 \leq x_2 \leq \dots \leq x_m = |\Gamma|$ be integers such that $\Gamma^*(i) = \Gamma(x_i)$ for all $i \in \{1, 2, \dots, m\}$ and $x_{i+1} - x_i \leq 1$ for all $i \in \{1, 2, \dots, m-1\}$. Similarly choose integers $1 = y_1 \leq y_2 \leq \dots \leq y_m = |\Delta|$ so that $\Delta^*(i) = \Delta(y_i)$ for all $i \in \{1, 2, \dots, m\}$ and $y_{i+1} - y_i \leq 1$ for all $i \in \{1, 2, \dots, m-1\}$. Let $v_i = (x_i, y_i)$ for $1 \leq i \leq m$.

If $x_{i+1} = x_i$ and $y_{i+1} = y_i$, then we can remove $\Gamma^*(i+1)$ and $\Delta^*(i+1)$ from Γ^* and Δ^* , respectively, contradicting our assumption that m is chosen to be minimum. Thus we conclude that $v_{i+1} - v_i \in \{(1, 0), (0, 1), (1, 1)\}$ for all i . This proves the forward implication.

To prove the converse, let us define $\Gamma^*(i) = \Gamma(x_i)$ and $\Delta^*(i) = \Delta(y_i)$ for all $i = 1, 2, \dots, m$. Then $\Gamma^* \in E(\Gamma)$, $\Delta^* \in E(\Delta)$ and $\Gamma^* \leq \Delta^*$ and therefore $\Gamma \preceq \Delta$. \square

We say that $\Gamma_1 \cong \Gamma_2$ if $\Gamma_1 \preceq \Gamma_2$ and $\Gamma_2 \preceq \Gamma_1$.

Lemma 3.6. *The binary relations \leq and \preceq on B -trajectories are transitive.*

Proof. By definition, it is easy to see that \leq is transitive.

Suppose $\Gamma_1 \preceq \Gamma_2$ and $\Gamma_2 \preceq \Gamma_3$. Then there are extensions $\Gamma_1^* \in E(\Gamma_1)$ and $\Gamma_2^* \in E(\Gamma_2)$ such that $\Gamma_1^* \leq \Gamma_2^*$. It is easy to see that $\Gamma_2^* \preceq \Gamma_3$ and therefore there exist extensions $\Gamma_2^{**} \in E(\Gamma_2^*)$ and $\Gamma_3^* \in E(\Gamma_3)$ such that $\Gamma_2^{**} \leq \Gamma_3^*$. We can easily find an extension Γ_1^{**} of Γ_1 such that $\Gamma_1^{**} \leq \Gamma_2^{**} \leq \Gamma_3^*$ and therefore $\Gamma_1 \preceq \Gamma_3$. \square

Lemma 3.7. *For a B -trajectory Γ , there exist $\Gamma_1, \Gamma_2 \in E(\tau(\Gamma))$ such that*

$$\Gamma_1 \leq \Gamma \leq \Gamma_2.$$

Proof. We proceed by induction on $|\Gamma|$. Let $\Gamma = a_1, a_2, \dots, a_m$. If $\tau(\Gamma) = \Gamma$, then this is trivial. If $a_{i-1} = a_i$ for some i , then we apply the induction hypothesis to $\Gamma' = a_1, a_2, \dots, a_{i-1}, a_{i+1}, a_{i+2}, \dots, a_m$ and deduce this lemma. Thus, we may assume that there exist i and j such that $i+1 < j$, $L(a_i) = L(a_j)$, $R(a_i) = R(a_j)$ and $\lambda(a_i) \leq \lambda(a_k) \leq \lambda(a_j)$ for all $k \in \{i+1, i+2, \dots, j-1\}$ or $\lambda(a_i) \geq \lambda(a_k) \geq \lambda(a_j)$ for all $k \in \{i+1, i+2, \dots, j-1\}$. Note that $L(a_i) = L(a_k) = L(a_j)$ and $R(a_i) = R(a_k) = R(a_j)$ for all $k \in \{i+1, i+2, \dots, j-1\}$.

Let $\Gamma' = a_1, a_2, \dots, a_i, a_j, a_{j+1}, a_{j+2}, \dots, a_m$. By definition, $\tau(\Gamma') = \tau(\Gamma)$ and therefore by the induction hypothesis, there exist $\Gamma'_1, \Gamma'_2 \in E(\tau(\Gamma))$ such that $\Gamma'_1 \leq \Gamma' \leq \Gamma'_2$. Let

$$\begin{aligned} \Gamma'_1 &= b_1, b_2, \dots, b_i, b_j, b_{j+1}, b_{j+2}, \dots, b_m \text{ and} \\ \Gamma'_2 &= c_1, c_2, \dots, c_i, c_j, c_{j+1}, c_{j+2}, \dots, c_m. \end{aligned}$$

(We intentionally skip some indices for the convenience of the proof.)

If $\lambda(a_i) \leq \lambda(a_j)$, then define $b_{i+1}, b_{i+2}, \dots, b_{j-1}$ to be equal to b_i and define $c_{i+1}, c_{i+2}, \dots, c_{j-1}$ to be equal to c_j . Let $\Gamma_1 = b_1, b_2, \dots, b_m$ and $\Gamma_2 = c_1, c_2, \dots, c_m$. Clearly $\Gamma_1, \Gamma_2 \in E(\tau(\Gamma))$. For all $i < k < j$, $b_k = b_i \leq a_i \leq a_k \leq a_j \leq c_j = c_k$. This proves that $\Gamma_1 \leq \Gamma \leq \Gamma_2$.

Similarly, we can find appropriate Γ_1, Γ_2 if $\lambda(a_i) > \lambda(a_j)$. \square

The following corollary is an immediate consequence of Lemma 3.7.

Corollary 3.8. *For all B -trajectories Γ , $\tau(\Gamma) \cong \Gamma$.*

The next corollary follows from Corollary 3.8 and the transitivity of \preceq .

Corollary 3.9. *For two B -trajectories Γ_1 and Γ_2 , $\Gamma_1 \preceq \Gamma_2$ if and only if $\tau(\Gamma_1) \preceq \tau(\Gamma_2)$.*

3.2. Realizable B -trajectories. For a linear layout $\sigma = V_1, V_2, \dots, V_m$ of a subspace arrangement \mathcal{V} , the *canonical B -trajectory of σ* is the B -trajectory $\Gamma_\sigma = a_1, a_2, \dots, a_{m+1}$

such that

$$\begin{aligned} L(a_i) &= (V_1 + V_2 + \cdots + V_{i-1}) \cap B, \\ R(a_i) &= (V_i + V_{i+1} + \cdots + V_m) \cap B, \\ \lambda(a_i) &= \dim(V_1 + V_2 + \cdots + V_{i-1}) \cap (V_i + V_{i+1} + \cdots + V_m) \\ &\quad - \dim(V_1 + V_2 + \cdots + V_{i-1}) \cap (V_i + V_{i+1} + \cdots + V_m) \cap B \end{aligned}$$

for all $i = 1, 2, \dots, m+1$. Note that $L(a_1) = \{0\}$, $R(a_{m+1}) = \{0\}$, $\lambda(a_1) = \lambda(a_{m+1}) = 0$ and $\lambda(a_i) \geq 0$.

We say that a B -trajectory Γ is *realizable in \mathcal{V}* if some extension of Γ is the canonical B -trajectory of some linear layout of \mathcal{V} .

3.3. Projection. For subspaces B and B' , we will define how to obtain a B' -trajectory from a B -trajectory. For a statistic $a = (L, R, \lambda)$, we write

$$a|_{B'} = (L \cap B', R \cap B', \lambda + \dim(L \cap R) - \dim(L \cap R \cap B')).$$

For a B -trajectory $\Gamma = a_1, a_2, \dots, a_n$, its *projection* $\Gamma|_{B'}$ is the B' -trajectory $a_1|_{B'}, a_2|_{B'}, \dots, a_n|_{B'}$. Observe that the width of Γ is less than or equal to the width of $\Gamma|_{B'}$.

Lemma 3.10. *Let B be a subspace of \mathbb{F}^r . If a, b are statistics such that $a \leq b$, then $a|_B \leq b|_B$.*

Proof. Let $a' = a|_B$ and $b' = b|_B$. It is easy to see that

$$\begin{aligned} L(a') &= L(a) \cap B = L(b) \cap B = L(b'), \\ R(a') &= R(a) \cap B = R(b) \cap B = R(b'), \\ \lambda(a') &= \lambda(a) + \dim L(a) \cap R(a) - \dim L(a) \cap R(a) \cap B \\ &\leq \lambda(b) + \dim L(b) \cap R(b) - \dim L(b) \cap R(b) \cap B \\ &= \lambda(b'). \end{aligned}$$

Therefore, $a|_B \leq b|_B$. □

Lemma 3.11. *Let B, B' be subspaces of \mathbb{F}^r and let Γ, Δ be B -trajectories. If $\Gamma \preceq \Delta$, then $\Gamma|_{B'} \preceq \Delta|_{B'}$.*

Proof. Without loss of generality, we may assume that $\Gamma \leq \Delta$. Lemma 3.10 easily implies that $\Gamma|_{B'} \leq \Delta|_{B'}$. □

Lemma 3.12. *Let \mathcal{V} be a subspace arrangement of subspaces of \mathbb{F}^r and let σ be a linear layout of \mathcal{V} . Let B be a subspace of \mathbb{F}^r and let B' be a subspace of B . If Γ is the canonical B -trajectory of σ , then $\Gamma|_{B'}$ is the canonical B' -trajectory of σ .*

Proof. Let $\sigma = V_1, V_2, \dots, V_m$. Let Γ' be the canonical B' -trajectory of σ . For all $1 \leq i \leq m+1$, we observe that

$$\begin{aligned} L(\Gamma'(i)) &= (V_1 + V_2 + \cdots + V_{i-1}) \cap B' = L(\Gamma(i)) \cap B', \\ R(\Gamma'(i)) &= (V_i + V_{i+1} + \cdots + V_m) \cap B' = R(\Gamma(i)) \cap B', \end{aligned}$$

and $\lambda(\Gamma'(i)) = \dim(V_1 + V_2 + \cdots + V_{i-1}) \cap (V_i + V_{i+1} + \cdots + V_m) - \dim L(\Gamma'(i)) \cap R(\Gamma'(i)) = \lambda(\Gamma(i)) + \dim L(\Gamma(i)) \cap R(\Gamma(i)) - \dim L(\Gamma(i)) \cap R(\Gamma(i)) \cap B'$. Therefore, $\Gamma'(i) = \Gamma(i)|_{B'}$ and so $\Gamma' = \Gamma|_{B'}$. □

Lemma 3.13. *Let \mathcal{V} be a subspace arrangement of subspaces of \mathbb{F}^r . Let B be a subspace of \mathbb{F}^r and let B' be a subspace of B . If Γ is a realizable B -trajectory in \mathcal{V} , then $\Gamma|_{B'}$ is realizable in \mathcal{V} .*

Proof. Let $\Gamma^* \in E(\Gamma)$ be the canonical B -trajectory of a linear layout σ of \mathcal{V} . By Lemma 3.12, $\Gamma^*|_{B'}$ is the canonical B' -trajectory of σ . Observe now that $\Gamma^*|_{B'}$ is an extension of $\Gamma|_{B'}$ and therefore $\Gamma|_{B'}$ is realizable in \mathcal{V} . \square

3.4. Sum of two B -trajectories. For two statistics a and b , let $a * b$ denote the statistic $(L(a) + L(b), R(a) + R(b), \lambda(a) + \lambda(b))$. For a statistic a and a nonnegative integer n , let $a + n$ be the statistic $(L(a), R(a), \lambda(a) + n)$.

For two B -trajectories Γ_1, Γ_2 and a lattice path $P = v_1, v_2, \dots, v_{|\Gamma_1|+|\Gamma_2|-1}$ from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$, we define the *sum* $\Gamma_1 +_P \Gamma_2$ as the B -trajectory Γ of length $|\Gamma_1| + |\Gamma_2| - 1$ such that if $v_i = (x_i, y_i)$, then

$$\Gamma(i) = \Gamma_1(x_i) * \Gamma_2(y_i) + (\dim R(\Gamma_1(1)) \cap R(\Gamma_2(1)) - \dim(L(\Gamma_1(x_i)) + R(\Gamma_1(x_i))) \cap (L(\Gamma_2(y_i)) + R(\Gamma_2(y_i))))$$

for all $1 \leq i \leq |\Gamma_1| + |\Gamma_2| - 1$. The *sum set* $\Gamma_1 \oplus \Gamma_2$ of two B -trajectories Γ_1 and Γ_2 is the set of all sums $\Gamma_1 +_P \Gamma_2$ for all possible lattice paths P from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$.

Lemma 3.14. *Let Γ_1, Γ_2 be two B -trajectories and $\Gamma \in \Gamma_1 \oplus \Gamma_2$. Then the width of Γ_1 and Γ_2 are at most the width of Γ .*

Proof. Let $P = (x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$ be a lattice path from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$ such that $\Gamma = \Gamma_1 +_P \Gamma_2$. It is enough to show that $\lambda(\Gamma_1(x_i)) \leq \lambda(\Gamma(i))$ for all $1 \leq i \leq n$.

We will show $\dim R(\Gamma_1(1)) \cap R(\Gamma_2(1)) - \dim(L(\Gamma_1(x_i)) + R(\Gamma_1(x_i))) \cap (L(\Gamma_2(y_i)) + R(\Gamma_2(y_i))) \geq 0$ for all $1 \leq i \leq n$. By the definition of B -trajectories, $R(\Gamma_1(1)) = L(\Gamma_1(|\Gamma_1|)) \supseteq L(\Gamma_1(x_i))$ and $R(\Gamma_1(1)) \supseteq R(\Gamma_1(x_i))$. Thus, $R(\Gamma_1(1)) \supseteq (L(\Gamma_1(x_i)) + R(\Gamma_1(x_i)))$ and similarly $R(\Gamma_2(1)) \supseteq (L(\Gamma_2(y_i)) + R(\Gamma_2(y_i)))$. Therefore, we conclude $\dim R(\Gamma_1(1)) \cap R(\Gamma_2(1)) - \dim(L(\Gamma_1(x_i)) + R(\Gamma_1(x_i))) \cap (L(\Gamma_2(y_i)) + R(\Gamma_2(y_i))) \geq 0$. Since $\lambda(\Gamma_2(y_i))$ is a nonnegative integer, by the definition of the sum, $\lambda(\Gamma_1(x_i)) \leq \lambda(\Gamma(i))$ for all $1 \leq i \leq n$. \square

Lemma 3.15. *Let Γ_1, Γ_2 be B -trajectories and $\Gamma_1^* \in E(\Gamma_1)$ and $\Gamma_2^* \in E(\Gamma_2)$. Then for every lattice path P from $(1, 1)$ to $(|\Gamma_1^*|, |\Gamma_2^*|)$, there is a lattice path Q from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$ such that $\Gamma_1^* +_P \Gamma_2^*$ is an extension of $\Gamma_1 +_Q \Gamma_2$.*

Proof. By induction, it is enough to prove the statement when $|\Gamma_1^*| = |\Gamma_1| + 1$ and $\Gamma_2 = \Gamma_2^*$. Suppose that $\Gamma_1(i) = \Gamma_1^*(i)$ for all $1 \leq i \leq p$ and $\Gamma_1(i-1) = \Gamma_1^*(i)$ for all $p+1 \leq i \leq |\Gamma_1^*|$.

Let $P = v_1, v_2, \dots, v_m$ with $m = |\Gamma_1^*| + |\Gamma_2| - 1$. Let (x_i, y_i) be the coordinate of v_i . Let s be the maximum such that $x_s = p$. Let t be the minimum such that $x_t = p+2$. Let us consider the subpath $P' = v_s, v_{s+1}, \dots, v_t$ of P . Then the points on P' are

$$(p, y_s), (p+1, y_{s+1}), (p+1, y_{s+2}), \dots, (p+1, y_{t-1}), (p+2, y_t)$$

where $y_s = y_{s+1} < y_{s+2} < \dots < y_{t-2} < y_{t-1} = y_t$.

Let us construct $Q = w_1, w_2, \dots, w_{m-1}$ as follows:

$$w_i = \begin{cases} v_i & \text{if } i \leq s, \\ v_{i+1} - (1, 0) & \text{if } i > s. \end{cases}$$

Let (x'_i, y'_i) be the coordinate of w_i . Then for $i > s$, $x'_i = x_{i+1} - 1 \geq p$ and $y'_i = y_{i+1}$. Let $\Gamma = \Gamma_1 +_Q \Gamma_2$ and $\Gamma' = \Gamma_1^* +_P \Gamma_2$. For $i \leq s$, $\Gamma'(i) = \Gamma(i)$. For $i > s$,

$$\begin{aligned}
\Gamma'(i) &= \Gamma_1^*(x_i) * \Gamma_2(y_i) + (\dim R(\Gamma_1^*(1)) \cap R(\Gamma_2(1))) \\
&\quad - \dim(L(\Gamma_1^*(x_i)) + R(\Gamma_1^*(x_i))) \cap (L(\Gamma_2(y_i)) + R(\Gamma_2(y_i))) \\
&= \Gamma_1(x_i - 1) * \Gamma_2(y_i) + (\dim R(\Gamma_1(1)) \cap R(\Gamma_2(1))) \\
&\quad - \dim(L(\Gamma_1(x_i - 1)) + R(\Gamma_1(x_i - 1))) \cap (L(\Gamma_2(y_i)) + R(\Gamma_2(y_i))) \\
&= \Gamma_1(x'_{i-1}) * \Gamma_2(y'_{i-1}) + (\dim R(\Gamma_1(1)) \cap R(\Gamma_2(1))) \\
&\quad - \dim(L(\Gamma_1(x'_{i-1})) + R(\Gamma_1(x'_{i-1}))) \cap (L(\Gamma_2(y'_{i-1})) + R(\Gamma_2(y'_{i-1}))) \\
&= \Gamma(i - 1).
\end{aligned}$$

Thus we conclude that Γ' is an extension of Γ . \square

3.5. Join. We first introduce three lemmas on subspaces, which are essential for join operations.

Lemma 3.16. *Let \mathcal{V}_1 and \mathcal{V}_2 be subspace arrangements of subspaces of \mathbb{F}^r and let B be a subspace of \mathbb{F}^r . If $(\langle \mathcal{V}_1 \rangle + B) \cap (\langle \mathcal{V}_2 \rangle + B) = B$, then*

$$(X_1 \cap B) + (X_2 \cap B) = (X_1 + X_2) \cap B$$

for each subspace X_1 of $\langle \mathcal{V}_1 \rangle$ and each subspace X_2 of $\langle \mathcal{V}_2 \rangle$.

Proof. It is enough to prove that if $(X_1 + B) \cap (X_2 + B) = B$, then $(X_1 \cap B) + (X_2 \cap B) = (X_1 + X_2) \cap B$. Clearly, $(X_1 \cap B) + (X_2 \cap B) \subseteq (X_1 + X_2) \cap B$.

Note that $X_1 \cap X_2 \subseteq B$ because $(X_1 \cap X_2) \subseteq (X_1 + B) \cap (X_2 + B) = B$. Now,

$$\begin{aligned}
&\dim(X_1 \cap B + X_2 \cap B) \\
&= \dim X_1 \cap B + \dim X_2 \cap B - \dim X_1 \cap X_2 \\
&= \dim(X_1) + \dim(B) + \dim(X_2) + \dim(B) \\
&\quad - \dim(X_1 + B) - \dim(X_2 + B) - \dim X_1 \cap X_2 \\
&= \dim(X_1) + \dim(B) + \dim(X_2) + \dim(B) \\
&\quad - \dim(X_1 + X_2 + B) - \dim(X_1 + B) \cap (X_2 + B) - \dim X_1 \cap X_2 \\
&= \dim(X_1) + \dim(B) + \dim(X_2) - \dim(X_1 + X_2 + B) - \dim X_1 \cap X_2 \\
&= \dim(X_1 + X_2) + \dim(B) - \dim(X_1 + X_2 + B) \\
&= \dim(X_1 + X_2) \cap B.
\end{aligned}$$

\square

Lemma 3.17. *For finite-dimensional vector spaces $X_1, X_2, Y_1,$ and $Y_2,$*

$$\begin{aligned}
&\dim(X_1 + X_2) \cap (Y_1 + Y_2) \\
&= \dim X_1 \cap Y_1 + \dim X_2 \cap Y_2 + \dim(X_1 + Y_1) \cap (X_2 + Y_2) - \dim X_1 \cap X_2 - \dim Y_1 \cap Y_2.
\end{aligned}$$

Proof. Recall that for two finite-dimensional vector spaces X and Y , $\dim(X) + \dim(Y) = \dim(X + Y) + \dim(X \cap Y)$. Using this fact, we deduce as follows.

$$\begin{aligned}
& \dim(X_1 + X_2) \cap (Y_1 + Y_2) \\
&= \dim(X_1 + X_2) + \dim(Y_1 + Y_2) - \dim(X_1 + X_2 + Y_1 + Y_2) \\
&= \dim X_1 + \dim X_2 - \dim X_1 \cap X_2 + \dim Y_1 + \dim Y_2 - \dim Y_1 \cap Y_2 \\
&\quad - \dim(X_1 + X_2 + Y_1 + Y_2) \\
&= \dim(X_1 + Y_1) + \dim X_1 \cap Y_1 + \dim(X_2 + Y_2) + \dim X_2 \cap Y_2 \\
&\quad - \dim X_1 \cap X_2 - \dim Y_1 \cap Y_2 - \dim(X_1 + X_2 + Y_1 + Y_2) \\
&= \dim(X_1 + Y_1) \cap (X_2 + Y_2) + \dim X_1 \cap Y_1 + \dim X_2 \cap Y_2 \\
&\quad - \dim X_1 \cap X_2 - \dim Y_1 \cap Y_2. \quad \square
\end{aligned}$$

Lemma 3.18. *Let \mathcal{V}_1 and \mathcal{V}_2 be subspace arrangements of subspaces of \mathbb{F}^r and B be a subspace of \mathbb{F}^r . If $(\langle \mathcal{V}_1 \rangle + B) \cap (\langle \mathcal{V}_2 \rangle + B) = B$, then*

$$\begin{aligned}
& \dim(X_1 + X_2) \cap (Y_1 + Y_2) - \dim(X_1 + X_2) \cap (Y_1 + Y_2) \cap B \\
&= \dim X_1 \cap Y_1 - \dim X_1 \cap Y_1 \cap B + \dim X_2 \cap Y_2 - \dim X_2 \cap Y_2 \cap B \\
&\quad + \dim(X_1 + Y_1) \cap (X_2 + Y_2) \cap B - \dim(X_1 \cap B + Y_1 \cap B) \cap (X_2 \cap B + Y_2 \cap B)
\end{aligned}$$

for subspaces X_1, Y_1 of $\langle \mathcal{V}_1 \rangle$ and X_2, Y_2 of $\langle \mathcal{V}_2 \rangle$.

Proof. From Lemma 3.16, observe that

$$\begin{aligned}
& \dim(X_1 + X_2) \cap (Y_1 + Y_2) - \dim(X_1 + X_2) \cap (Y_1 + Y_2) \cap B \\
&= \dim(X_1 + X_2) \cap (Y_1 + Y_2) - \dim(X_1 \cap B + X_2 \cap B) \cap (Y_1 \cap B + Y_2 \cap B).
\end{aligned}$$

We apply Lemma 3.17 to X_1, X_2, Y_1, Y_2 and $X_1 \cap B, X_2 \cap B, Y_1 \cap B, Y_2 \cap B$, respectively, which yields

$$\begin{aligned}
\dim(X_1 + X_2) \cap (Y_1 + Y_2) &= \dim X_1 \cap Y_1 + \dim X_2 \cap Y_2 \\
&\quad - \dim X_1 \cap X_2 - \dim Y_1 \cap Y_2 + \dim(X_1 + Y_1) \cap (X_2 + Y_2),
\end{aligned}$$

and

$$\begin{aligned}
& \dim(X_1 \cap B + X_2 \cap B) \cap (Y_1 \cap B + Y_2 \cap B) \\
&= \dim X_1 \cap Y_1 \cap B + \dim X_2 \cap Y_2 \cap B - \dim X_1 \cap X_2 \cap B - \dim Y_1 \cap Y_2 \cap B \\
&\quad + \dim(X_1 \cap B + Y_1 \cap B) \cap (X_2 \cap B + Y_2 \cap B).
\end{aligned}$$

Since $(X_1 + Y_1) \cap (X_2 + Y_2) \subseteq (X_1 + Y_1 + B) \cap (X_2 + Y_2 + B) \subseteq (\langle \mathcal{V}_1 \rangle + B) \cap (\langle \mathcal{V}_2 \rangle + B) = B$, we have $(X_1 + Y_1) \cap (X_2 + Y_2) = (X_1 + Y_1) \cap (X_2 + Y_2) \cap B$. Similarly, $X_1 \cap X_2 = X_1 \cap X_2 \cap B$, $Y_1 \cap Y_2 = Y_1 \cap Y_2 \cap B$. Thus the equality of the statement follows. \square

Proposition 3.19. *Let $\mathcal{V}_1, \mathcal{V}_2$ be subspace arrangements of subspaces of \mathbb{F}^r and B be a subspace of \mathbb{F}^r such that $(\langle \mathcal{V}_1 \rangle + B) \cap (\langle \mathcal{V}_2 \rangle + B) = B$. If Γ is realizable in $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$, then there exist B -trajectories Γ_1, Γ_2 and a lattice path P from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$ such that*

- (1) Γ_i is realizable in \mathcal{V}_i for $i = 1, 2$,
- (2) $\Gamma_1 +_P \Gamma_2 \preceq \Gamma$.

Proof. Let $\Gamma^* \in E(\Gamma)$ be the canonical B -trajectory of a linear layout $\sigma = V_1, V_2, \dots, V_m$ of $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$. For $i = 1, 2$, let $\sigma_i = W_1^i, W_2^i, \dots, W_{m_i}^i$ be a subsequence of σ by taking all $V_j \in \mathcal{V}_i$ and let Γ_i be the canonical B -trajectory of σ_i . Clearly, $m_1 + m_2 = m$, and σ_i is a linear layout of \mathcal{V}_i and so Γ_i is realizable in \mathcal{V}_i . Also, observe that $W_j^i \subseteq \langle \mathcal{V}_i \rangle$ for all i and j .

Let $P = v_1, v_2, \dots, v_m$ be a lattice path from $(1, 1)$ to $(m_1 + 1, m_2 + 1)$ such that for $i = 1, 2, \dots, m$,

$$v_{i+1} = \begin{cases} v_i + (1, 0) & \text{if } V_i \in \mathcal{V}_1, \\ v_i + (0, 1) & \text{if } V_i \in \mathcal{V}_2. \end{cases}$$

We claim that $\Gamma_1 +_P \Gamma_2 = \Gamma^*$, which implies that $\Gamma_1 +_P \Gamma_2 \preceq \Gamma$. As Γ^* has length $m + 1$ and Γ_i has length $m_i + 1$ for $i = 1, 2$, both sides have the same length. Let a_j be the j -th statistic of $\Gamma_1 +_P \Gamma_2$ for B . It remains to prove that $\Gamma(j) = a_j$ for all $1 \leq j \leq m + 1$.

Let $v_j = (x_j, y_j)$. Let $L_1 = W_1^1 + W_2^1 + \dots + W_{x_j-1}^1$, $L_2 = W_1^2 + W_2^2 + \dots + W_{y_j-1}^2$, $R_1 = W_{x_j}^1 + W_{x_j+1}^1 + \dots + W_{m_1}^1$, and $R_2 = W_{y_j}^2 + W_{y_j+1}^2 + \dots + W_{m_2}^2$. By the construction and Lemma 3.16, $x_j - 1 + y_j - 1 = j - 1$ and $L(a_j) = L(\Gamma_1(x_j)) + L(\Gamma_2(y_j)) = L_1 \cap B + L_2 \cap B = (L_1 + L_2) \cap B$. Since $L_1 + L_2 = V_1 + V_2 + \dots + V_{j-1}$, we deduce that $L(a_j) = L(\Gamma(j))$. Similarly, $R(a_j) = R(\Gamma(j))$.

It remains to prove that $\lambda(\Gamma(j)) = \lambda(a_j)$. First observe that $\lambda(\Gamma_1(x_j)) = \dim L_1 \cap R_1 - \dim L_1 \cap R_1 \cap B$ and $\lambda(\Gamma_2(y_j)) = \dim L_2 \cap R_2 - \dim L_2 \cap R_2 \cap B$. Also, $R(\Gamma_1(1)) = (L_1 + R_1) \cap B$ and $R(\Gamma_2(1)) = (L_2 + R_2) \cap B$. By definition, L_1 and R_1 are subspaces of $\langle \mathcal{V}_1 \rangle$, and L_2 and R_2 are subspaces of $\langle \mathcal{V}_2 \rangle$. Therefore, by Lemma 3.18,

$$\begin{aligned} & \lambda(\Gamma(j)) \\ &= \dim(L_1 + L_2) \cap (R_1 + R_2) - \dim(L_1 + L_2) \cap (R_1 + R_2) \cap B \\ &= \dim L_1 \cap R_1 - \dim L_1 \cap R_1 \cap B + \dim L_2 \cap R_2 - \dim L_2 \cap R_2 \cap B \\ &\quad + \dim(L_1 + R_1) \cap (L_2 + R_2) \cap B - \dim(L_1 \cap B + R_1 \cap B) \cap (L_2 \cap B + R_2 \cap B) \\ &= \lambda(\Gamma_1(x_j)) + \lambda(\Gamma_2(y_j)) + \dim R(\Gamma_1(1)) \cap R(\Gamma_2(1)) \\ &\quad - \dim(L(\Gamma_1(x_j)) + R(\Gamma_1(y_j))) \cap (L(\Gamma_2(x_j)) + R(\Gamma_2(y_j))) \\ &= \lambda(a_j). \end{aligned} \quad \square$$

Proposition 3.20. *Let $\Gamma_1, \Gamma'_1, \Gamma_2, \Gamma'_2$ be B -trajectories. If $\Gamma'_1 \preceq \Gamma_1$ and $\Gamma'_2 \preceq \Gamma_2$, then for all $\Gamma \in \Gamma_1 \oplus \Gamma_2$, there exists $\Gamma' \in \Gamma'_1 \oplus \Gamma'_2$ such that $\Gamma' \preceq \Gamma$.*

To prove this proposition, we use the following lemmas.

Lemma 3.21. *Let $\Gamma_1, \Gamma'_1, \Gamma_2$ be B -trajectories such that $\Gamma'_1 \leq \Gamma_1$. Let P be a lattice path from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$. Then $\Gamma'_1 +_P \Gamma_2 \leq \Gamma_1 +_P \Gamma_2$.*

Proof. Observe that $\Gamma'_1 +_P \Gamma_2$ is well defined because $|\Gamma'_1| = |\Gamma_1|$. Let $\Gamma' = \Gamma'_1 +_P \Gamma_2$ and $\Gamma = \Gamma_1 +_P \Gamma_2$. It is enough to prove that for $1 \leq j \leq |\Gamma| = |\Gamma_1| + |\Gamma_2| - 1$, we have $\Gamma'(j) \leq \Gamma(j)$. Let (x_j, y_j) be the j -th point of P . Let $a = \Gamma_1(x_j)$, $a' = \Gamma'_1(x_j)$, and $b = \Gamma_2(y_j)$. First, observe that $L(\Gamma'(j)) = L(a') + L(b) = L(a) + L(b) = L(\Gamma(j))$ because $L(a') = L(a)$ and similarly $R(\Gamma'(j)) = R(\Gamma(j))$. Finally, $\lambda(\Gamma'(j)) = \lambda(a') + \lambda(b) + \dim R(\Gamma'_1(1)) \cap R(\Gamma_2(1)) - \dim(L(a') + R(a')) \cap (L(b) + R(b)) \leq \lambda(a) + \lambda(b) + \dim R(\Gamma_1(1)) \cap R(\Gamma_2(1)) - \dim(L(a) + R(a)) \cap (L(b) + R(b)) = \lambda(\Gamma(j))$. This proves the lemma. \square

Lemma 3.22. *Let Γ_1 and Γ_2 be two B -trajectories and let P be a lattice path from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$. For two extensions $\Gamma'_1 \in E(\Gamma_1)$, $\Gamma'_2 \in E(\Gamma_2)$, there exists a lattice path P' from $(1, 1)$ to $(|\Gamma'_1|, |\Gamma'_2|)$ such that $\Gamma'_1 +_{P'} \Gamma'_2$ is an extension of $\Gamma_1 +_P \Gamma_2$.*

Proof. It is enough to prove the statement when $|\Gamma'_1| = |\Gamma_1| + 1$ and $\Gamma'_2 = \Gamma_2$, from which we can inductively deduce our main statement. The lattice path P is denoted as v_1, v_2, \dots, v_m with $v_i = (x_i, y_i)$. Let a be a statistic of Γ_1 such that Γ'_1 is obtained by repeating a once. We arbitrarily choose $1 \leq p \leq m$ such that $\Gamma_1(x_p) = a$. Notice that $\Gamma'_1(i) = \Gamma_1(i - 1)$ for $i > x_p$. Take $P' = v'_1, v'_2, \dots, v'_{m+1}$ with $v'_i = (x'_i, y'_i)$ as follows:

$$v'_i = \begin{cases} (x_i, y_i) & \text{if } i \leq p, \\ (x_{i-1} + 1, y_{i-1}) & \text{if } i \geq p + 1. \end{cases}$$

We claim that $\Gamma' = \Gamma'_1 +_{P'} \Gamma_2$ is an extension of $\Gamma_1 +_P \Gamma_2$. For $1 \leq i \leq p$,

$$\begin{aligned} \Gamma'(i) &= \Gamma'_1(x'_i) * \Gamma_2(y'_i) + (\dim R(\Gamma'_1(1)) \cap R(\Gamma_2(1))) \\ &\quad - \dim(L(\Gamma'_1(x'_i)) + R(\Gamma'_1(x'_i))) \cap (L(\Gamma_2(y'_i)) + R(\Gamma_2(y'_i))) \\ &= \Gamma_1(x_i) * \Gamma_2(y_i) + (\dim R(\Gamma_1(1)) \cap R(\Gamma_2(1))) \\ &\quad - \dim(L(\Gamma_1(x_i)) + R(\Gamma_1(x_i))) \cap (L(\Gamma_2(y_i)) + R(\Gamma_2(y_i))). \end{aligned}$$

For $p + 1 \leq i \leq m + 1$,

$$\begin{aligned} \Gamma'(i) &= \Gamma'_1(x_{i-1} + 1) * \Gamma_2(y_{i-1}) + (\dim R(\Gamma'_1(1)) \cap R(\Gamma_2(1))) \\ &\quad - \dim(L(\Gamma'_1(x_{i-1} + 1)) + R(\Gamma'_1(x_{i-1} + 1))) \cap (L(\Gamma_2(y_{i-1})) + R(\Gamma_2(y_{i-1}))) \\ &= \Gamma_1(x_{i-1}) * \Gamma_2(y_{i-1}) + (\dim R(\Gamma_1(1)) \cap R(\Gamma_2(1))) \\ &\quad - \dim(L(\Gamma_1(x_{i-1})) + R(\Gamma_1(x_{i-1}))) \cap (L(\Gamma_2(y_{i-1})) + R(\Gamma_2(y_{i-1}))). \end{aligned}$$

Therefore, Γ' can be obtained from $\Gamma_1 +_P \Gamma_2$ by repeating the statistic $\Gamma_1(x_p) + \Gamma_2(y_p)$. \square

Proof of Proposition 3.20. Since $\Gamma_1 \oplus \Gamma_2 = \Gamma_2 \oplus \Gamma_1$ and \preceq is transitive by Lemma 3.6, it is enough to prove it for the case that $\Gamma'_2 = \Gamma_2$. Let $\Gamma''_1 \in E(\Gamma'_1)$ and $\Gamma^*_1 \in E(\Gamma_1)$ such that $\Gamma''_1 \leq \Gamma^*_1$. Let P be a lattice path from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$ such that $\Gamma = \Gamma_1 +_P \Gamma_2$. From Lemma 3.22, there is a lattice path P' from $(1, 1)$ to $(|\Gamma^*_1|, |\Gamma_2|)$ such that $\Gamma^*_1 +_{P'} \Gamma_2$ is an extension of Γ . Note that $\Gamma^*_1 +_{P'} \Gamma_2 \preceq \Gamma$. By Lemma 3.21, $\Gamma''_1 +_{P'} \Gamma_2 \leq \Gamma^*_1 +_{P'} \Gamma_2$. By Lemma 3.15, there is $\Gamma' \in \Gamma''_1 \oplus \Gamma_2$ such that $\Gamma''_1 +_{P'} \Gamma_2$ is an extension of Γ' . \square

Proposition 3.23. *Let $\mathcal{V}_1, \mathcal{V}_2$ be subspace arrangements of subspaces of \mathbb{F}^r and let B be a subspace of \mathbb{F}^r such that $(\langle \mathcal{V}_1 \rangle + B) \cap (\langle \mathcal{V}_2 \rangle + B) = B$. Let Γ_1 and Γ_2 be two B -trajectories realizable in \mathcal{V}_1 and \mathcal{V}_2 , respectively. Then every $\Gamma \in \Gamma_1 \oplus \Gamma_2$ is realizable in $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$.*

Let us delay proving Proposition 3.23 by presenting a helpful lemma for the proof first.

Lemma 3.24. *Let $\mathcal{V}_1, \mathcal{V}_2$ be subspace arrangements of subspaces of \mathbb{F}^r and let B be a subspace of \mathbb{F}^r such that $(\langle \mathcal{V}_1 \rangle + B) \cap (\langle \mathcal{V}_2 \rangle + B) = B$. For $i = 1, 2$, let Γ_i be the canonical B -trajectory of some linear layout of \mathcal{V}_i . For every lattice path P from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$, there exists a linear layout of $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$ whose canonical B -trajectory is $\Gamma_1 +_P \Gamma_2$.*

Proof. Let $P = v_1, \dots, v_{|\Gamma_1|+|\Gamma_2|-1}$ be a lattice path from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$ with $v_j = (x_j, y_j)$. Let $\sigma_1 = S_1, S_2, \dots, S_{|\Gamma_1|-1}$ be the linear layout of \mathcal{V}_1 whose canonical B -trajectory is Γ_1 and $\sigma_2 = T_1, T_2, \dots, T_{|\Gamma_2|-1}$ be the linear layout of \mathcal{V}_2 whose canonical B -trajectory is Γ_2 . We build a linear layout σ from σ_1 and σ_2 as follows: for

$$i = 1, 2, \dots, |\Gamma_1| + |\Gamma_2| - 2,$$

$$\sigma[i] = \begin{cases} S_{x_i} & \text{if } v_{i+1} - v_i = (1, 0), \\ T_{y_i} & \text{if } v_{i+1} - v_i = (0, 1), \end{cases}$$

where $\sigma[i]$ is the i -th subspace in σ . It is clear that σ is a linear layout of $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$. Let Γ be the canonical B -trajectory of σ .

We claim that $\Gamma = \Gamma_1 +_P \Gamma_2$. Let $i \in \{1, 2, \dots, |\Gamma|\}$. We shall show that $\Gamma(i) = (\Gamma_1 +_P \Gamma_2)(i)$. Let $L_1 = S_1 + \dots + S_{x_{i-1}}$, $R_1 = S_{x_i} + \dots + S_{|\Gamma_1|-1}$, $L_2 = T_1 + \dots + T_{y_{i-1}}$ and $R_2 = T_{y_i} + \dots + T_{|\Gamma_2|-1}$. Note that $S_1, S_2, \dots, S_{|\Gamma_1|-1} \subseteq \langle \mathcal{V}_1 \rangle$ and $T_1, T_2, \dots, T_{|\Gamma_2|-1} \subseteq \langle \mathcal{V}_2 \rangle$, and thus $L_1, R_1 \subseteq \langle \mathcal{V}_1 \rangle$ and $L_2, R_2 \subseteq \langle \mathcal{V}_2 \rangle$. This allows us to use Lemma 3.16. We deduce from Lemma 3.16 that

$$\begin{aligned} L(\Gamma(i)) &= (L_1 + L_2) \cap B = L_1 \cap B + L_2 \cap B \\ &= L(\Gamma_1(x_i)) + L(\Gamma_2(y_i)) = L(\Gamma_1(x_i) + \Gamma_2(y_i)) \end{aligned}$$

where the last equality follows from the definition of the sum $\Gamma_1 +_P \Gamma_2$. Similarly, $R(\Gamma(i)) = R(\Gamma_1(x_i) + \Gamma_2(y_i)) = R((\Gamma_1 +_P \Gamma_2)(i))$.

From $(\langle \mathcal{V}_1 \rangle + B) \cap (\langle \mathcal{V}_2 \rangle + B) = B$, it holds that for every subspace X_1 of $\langle \mathcal{V}_1 \rangle$ and subspace X_2 of $\langle \mathcal{V}_2 \rangle$, $X_1 \cap X_2 \subseteq (X_1 + B) \cap (X_2 + B) \subseteq B$. Since $R(\Gamma_1(1)) = (L_1 + R_1) \cap B$ and $R(\Gamma_2(1)) = (L_2 + R_2) \cap B$, by Lemma 3.18, we have

$$\begin{aligned} \lambda(\Gamma(i)) &= \dim(L_1 + L_2) \cap (R_1 + R_2) - \dim(L_1 + L_2) \cap (R_1 + R_2) \cap B \\ &= \dim L_1 \cap R_1 - \dim L_1 \cap R_1 \cap B + \dim L_2 \cap R_2 - \dim L_2 \cap R_2 \cap B \\ &\quad + \dim(L_1 + R_1) \cap (L_2 + R_2) \cap B \\ &\quad - \dim(L_1 \cap B + R_1 \cap B) \cap (L_2 \cap B + R_2 \cap B) \\ &= \lambda(\Gamma_1(x_i)) + \lambda(\Gamma_2(y_i)) + \dim R(\Gamma_1(1)) \cap R(\Gamma_2(1)) \\ &\quad - \dim(L(\Gamma_1(x_i)) + R(\Gamma_1(x_i))) \cap (L(\Gamma_2(y_i)) + R(\Gamma_2(y_i))) \\ &= \lambda((\Gamma_1 +_P \Gamma_2)(i)). \end{aligned}$$

Hence, we conclude that $\Gamma = \Gamma_1 +_P \Gamma_2$. \square

Now we are ready to finish the proof of Proposition 3.23.

Proof of Proposition 3.23. For $i = 1, 2$, let $\Gamma_i^* \in E(\Gamma_i)$ be the canonical B -trajectory of some linear layout of \mathcal{V}_i . Let P be a lattice path from $(1, 1)$ to $(|\Gamma_1|, |\Gamma_2|)$ such that $\Gamma = \Gamma_1 +_P \Gamma_2$. By Lemma 3.22, there exists a lattice path P' from $(1, 1)$ to $(|\Gamma_1^*|, |\Gamma_2^*|)$ such that $\Gamma_1^* +_{P'} \Gamma_2^*$ is an extension of Γ . Lemma 3.24 implies that $\Gamma_1^* +_{P'} \Gamma_2^*$ is the canonical B -trajectory of a linear layout of $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$. Therefore, Γ is realizable in $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$. \square

4. THE FULL SET FOR DYNAMIC PROGRAMMING

Let \mathcal{V} be a subspace arrangement of subspaces of \mathbb{F}^r and B be a subspace of \mathbb{F}^r . The *full set* of \mathcal{V} of width k with respect to B , denoted by $\text{FS}_k(\mathcal{V}, B)$, is the set of all compact B -trajectories Γ of width at most k such that there exists a B -trajectory Δ realizable in \mathcal{V} with $\Delta \preceq \Gamma$.

For a set \mathcal{R} of B -trajectories and $B \subseteq B'$, we define $\text{up}_k(\mathcal{R}, B')$ as the set of all compact B' -trajectories Γ of width at most k such that there exists $\Delta \in \mathcal{R}$ with $\Delta \preceq \Gamma$ (considering Δ as a B' -trajectory). For subspaces $B' \subseteq B$ and a set \mathcal{R} of B -trajectories, let $\mathcal{R}|_{B'}$ be

the set of all B' -trajectories $\Gamma|_{B'}$ where $\Gamma \in \mathcal{R}$. For two sets $\mathcal{R}_1, \mathcal{R}_2$ of B -trajectories, we define $\mathcal{R}_1 \oplus \mathcal{R}_2 = \bigcup_{\Gamma_1 \in \mathcal{R}_1, \Gamma_2 \in \mathcal{R}_2} \Gamma_1 \oplus \Gamma_2$.

We give an overview of this section. Each subsection describes how we are going to compute a full set. Subsection 4.1 is about the case that $|\mathcal{V}| = 1$. Subsection 4.2 discusses how to obtain $\text{FS}_k(\mathcal{V}, B')$ from a given $\text{FS}_k(\mathcal{V}, B)$ for a subspace B of B' such that $\langle \mathcal{V} \rangle \cap B' \subseteq B$, which we call *expanding*. Subsection 4.3 is the opposite; it presents the construction of $\text{FS}_k(\mathcal{V}, B')$ from a given $\text{FS}_k(\mathcal{V}, B)$ when B' is a subspace of B . This process is called *shrinking*. Subsection 4.4 explains how we obtain $\text{FS}_k(\mathcal{V}, B)$ when we are given $\text{FS}_k(\mathcal{V}_1, B)$ and $\text{FS}_k(\mathcal{V}_2, B)$ for some subspace arrangements \mathcal{V}_1 and \mathcal{V}_2 such that \mathcal{V} is the disjoint union of \mathcal{V}_1 and \mathcal{V}_2 . This process is called *joining*.

4.1. A full set for a single subspace.

Proposition 4.1. *Let B, V be subspaces in \mathbb{F}^r such that $B \subseteq V$. Then $\text{FS}_k(\{V\}, B) = \text{up}_k(\{\Delta\}, B)$ where Δ is a B -trajectory $(\{0\}, B, 0), (B, \{0\}, 0)$.*

Proof. Since $\{V\}$ has the unique linear layout and $V \cap B = B$, Δ is the unique canonical B -trajectory of the linear layout. Thus, for a compact B -trajectory Γ of width at most k , $\Gamma \in \text{FS}_k(\{V\}, B)$ if and only if $\Delta \preceq \Gamma$. \square

4.2. A full set for an expanding operation.

Proposition 4.2. *Let \mathcal{V} be a subspace arrangement of subspaces of \mathbb{F}^r . Let B, B' be subspaces of \mathbb{F}^r such that $B \subseteq B'$. If $\langle \mathcal{V} \rangle \cap B' \subseteq B$, then $\text{FS}_k(\mathcal{V}, B') = \text{up}_k(\text{FS}_k(\mathcal{V}, B), B')$.*

Proof. Because $\langle \mathcal{V} \rangle \cap B' \subseteq B$, for every subspace W of $\langle \mathcal{V} \rangle$, we have $W \cap B' \subseteq B$. This implies that for every linear layout σ of \mathcal{V} , the canonical B -trajectory of σ is also the canonical B' -trajectory of σ and vice versa. Thus, Γ is a realizable B -trajectory if and only if it is a realizable B' -trajectory. The conclusion follows easily, as \preceq is transitive by Lemma 3.6. \square

4.3. A full set for a shrinking operation.

Proposition 4.3. *Let \mathcal{V} be a subspace arrangement of subspaces of \mathbb{F}^r and B, B' be subspaces of \mathbb{F}^r . If $B' \subseteq B$, then $\text{FS}_k(\mathcal{V}, B') = \text{up}_k(\text{FS}_k(\mathcal{V}, B)|_{B'}, B')$.*

Proof. We first prove $\text{FS}_k(\mathcal{V}, B') \subseteq \text{up}_k(\text{FS}_k(\mathcal{V}, B)|_{B'}, B')$. Let $\Gamma \in \text{FS}_k(\mathcal{V}, B')$ and Δ be a B' -trajectory realizable in \mathcal{V} such that $\Delta \preceq \Gamma$. Let σ be a linear layout of \mathcal{V} whose canonical B' -trajectory is an extension of Δ . Let Δ' be the canonical B -trajectory of σ . Since Γ is a compact B' -trajectory of width at most k , it is enough to show that (i) $\tau(\Delta')|_{B'} \preceq \Gamma$ and (ii) $\tau(\Delta')$ is in $\text{FS}_k(\mathcal{V}, B)$. For (i), $\Delta'|_{B'}$ is an extension of Δ because it is the canonical B' -trajectory of π . Thus, $\Delta'|_{B'} \preceq \Delta \preceq \Gamma$. By Corollary 3.8 and Lemma 3.11, $\tau(\Delta')|_{B'} \preceq \Delta'|_{B'}$ and we conclude that $\tau(\Delta')|_{B'} \preceq \Gamma$ by Lemma 3.6. For (ii), by Lemma 3.4, we know that $(\text{the width of } \tau(\Delta')) = (\text{the width of } \Delta') \leq (\text{the width of } \Delta'|_{B'}) \leq (\text{the width of } \Gamma) \leq k$. By Corollary 3.8, we have $\Delta' \preceq \tau(\Delta')$. Since Δ' is realizable in \mathcal{V} , $\tau(\Delta')$ is in $\text{FS}_k(\mathcal{V}, B)$.

Next, let us show $\text{FS}_k(\mathcal{V}, B') \supseteq \text{up}_k(\text{FS}_k(\mathcal{V}, B)|_{B'}, B')$. Choose an arbitrary Γ from $\text{up}_k(\text{FS}_k(\mathcal{V}, B)|_{B'}, B')$ and let $\Delta \in \text{FS}_k(\mathcal{V}, B)$ such that $\Delta|_{B'} \preceq \Gamma$. By the definition of the full set $\text{FS}_k(\mathcal{V}, B)$, there exists a B -trajectory Δ' that is realizable in \mathcal{V} and $\Delta' \preceq \Delta$. By Lemma 3.13, $\Delta'|_{B'}$ is realizable in \mathcal{V} . By Lemma 3.11, $\Delta'|_{B'} \preceq \Delta|_{B'} \preceq \Gamma$. Recall that Γ is compact and of width at most k , and thus $\Gamma \in \text{FS}_k(\mathcal{V}, B')$. \square

4.4. A full set for a join operation.

Proposition 4.4. *Let $\mathcal{V}_1, \mathcal{V}_2$ be subspace arrangements of subspaces of \mathbb{F}^r and let B be a subspace of \mathbb{F}^r . If $(\langle \mathcal{V}_1 \rangle + B) \cap (\langle \mathcal{V}_2 \rangle + B) = B$, then*

$$\text{FS}_k(\mathcal{V}_1 \dot{\cup} \mathcal{V}_2, B) = \text{up}_k(\text{FS}_k(\mathcal{V}_1, B) \oplus \text{FS}_k(\mathcal{V}_2, B), B).$$

Proof. We prove the “ \supseteq ”-direction first. Let $\Gamma \in \text{up}_k(\text{FS}_k(\mathcal{V}_1, B) \oplus \text{FS}_k(\mathcal{V}_2, B), B)$. By definition, there exists $\Delta \in \Delta_1 \oplus \Delta_2$ for some $\Delta_1 \in \text{FS}_k(\mathcal{V}_1, B)$, $\Delta_2 \in \text{FS}_k(\mathcal{V}_2, B)$ satisfying $\Delta \preceq \Gamma$. Again by definition, for $i = 1, 2$, there exists a B -trajectory Δ'_i realizable in \mathcal{V}_i such that $\Delta'_i \preceq \Delta_i$. By applying Proposition 3.20 to Δ_1, Δ_2 , and Δ'_1, Δ'_2 , we deduce that there exists $\Delta' \in \Delta'_1 \oplus \Delta'_2$ such that $\Delta' \preceq \Delta$. Proposition 3.23 implies that all elements of $\Delta'_1 \oplus \Delta'_2$ are realizable in $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$, and so is Δ' . From the fact that Δ' is realizable in $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$ and $\Delta' \preceq \Delta \preceq \Gamma$, we conclude that Γ belongs to the set $\text{FS}_k(\mathcal{V}_1 \dot{\cup} \mathcal{V}_2, B)$.

Now we prove the “ \subseteq ”-direction. Consider an arbitrary Γ from $\text{FS}_k(\mathcal{V}_1 \dot{\cup} \mathcal{V}_2, B)$. By the definition of the full set, there exists a B -trajectory Δ realizable in $\mathcal{V}_1 \dot{\cup} \mathcal{V}_2$ such that $\Delta \preceq \Gamma$. By Proposition 3.19, there exist two B -trajectories Δ_1, Δ_2 and a lattice path P from $(1, 1)$ to $(|\Delta_1|, |\Delta_2|)$ such that

- (1) Δ_i is realizable in \mathcal{V}_i for $i = 1, 2$,
- (2) $\Delta_1 +_P \Delta_2 \preceq \Delta$.

We apply Proposition 3.20 to B -trajectories Δ_1, Δ_2 and $\tau(\Delta_1), \tau(\Delta_2)$. Since $\tau(\Delta_1) \preceq \Delta_1$ and $\tau(\Delta_2) \preceq \Delta_2$, there exists $\Delta' \in \tau(\Delta_1) \oplus \tau(\Delta_2)$ such that $\Delta' \preceq \Delta_1 +_P \Delta_2$, and thus $\Delta' \preceq \Delta \preceq \Gamma$ by Lemma 3.6. Since $\Delta_1 +_P \Delta_2 \preceq \Delta \preceq \Gamma$ and Γ has width at most k , the width of Δ_1 and Δ_2 are at most k by Lemmas 3.4 and 3.14. So it follows that both $\tau(\Delta_1)$ and $\tau(\Delta_2)$ have width at most k by Corollary 3.8 and Lemma 3.4. Since Δ_i is realizable in \mathcal{V}_i and $\Delta_i \preceq \tau(\Delta_i)$, we have $\tau(\Delta_i) \in \text{FS}_k(\mathcal{V}_i, B)$ for $i = 1, 2$. To summarize, we have $\Delta' \in \tau(\Delta_1) \oplus \tau(\Delta_2)$ for some $\tau(\Delta_i) \in \text{FS}_k(\mathcal{V}_i, B)$, $i = 1, 2$, satisfying $\Delta' \preceq \Gamma$ and therefore, $\Gamma \in \text{up}_k(\text{FS}_k(\mathcal{V}_1, B) \oplus \text{FS}_k(\mathcal{V}_2, B), B)$. \square

5. AN ALGORITHM

We are ready to describe our main algorithm to solve the following problem. We first describe how subspaces are given as an input. For an $r \times m$ matrix M and $X \subseteq \{1, 2, \dots, r\}$, $Y \subseteq \{1, 2, \dots, m\}$, let us write $M[X, Y]$ to denote the submatrix of M induced by the rows in X and the columns in Y and $M[Y] := M[\{1, 2, \dots, r\}, Y]$.

Input: An $r \times m$ matrix M over a fixed finite field \mathbb{F} with a partition $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$ of $\{1, 2, \dots, m\}$, an integer k .

Parameter: k .

Problem: For each $1 \leq i \leq n$, let V_i be the column space of $M[I_i]$ and let $\mathcal{V} = \{V_i : 1 \leq i \leq n\}$ be a subspace arrangement.

Decide whether there exists a permutation σ of $\{1, 2, \dots, n\}$ such that

$$\dim(V_{\sigma(1)} + V_{\sigma(2)} + \dots + V_{\sigma(i)}) \cap (V_{\sigma(i+1)} + V_{\sigma(i+2)} + \dots + V_{\sigma(n)}) \leq k$$

for all $i = 1, 2, \dots, n-1$ and if it exists, then output such a permutation σ .

If such a permutation σ exists, then we say that (M, \mathcal{I}, k) is a YES instance. Otherwise it is a NO instance. In this case, a sequence $V_{\sigma(1)}, V_{\sigma(2)}, \dots, V_{\sigma(n)}$ is a linear layout of \mathcal{V} having width at most k .

In this section, we assume that a branch-decomposition (T, \mathcal{L}) of \mathcal{V} of width at most θ is given as a part of the input. By Proposition 2.2, there exists a branch-decomposition

of width $\theta \leq 2k$ if the input instance has path-width at most k . Finding such a branch-decomposition and the corresponding running time shall be discussed in detail in Section 6. Let us assume that T is a rooted binary tree by picking an arbitrary edge e and subdividing e to create a degree-2 vertex called the root.

5.1. Preprocessing. Our aim is to obtain a uniform bound on the running time of our algorithm in terms of n for each fixed k . However, it is possible that m or r is very large or V_i has huge dimension. Thus, our first step is to preprocess the input so that both r and m are bounded by functions of k and n .

Let us say that an $r \times m$ matrix M is *of the standard form* if the submatrix induced by the first r columns is the $r \times r$ identity matrix. Clearly in every matrix of the standard form, row vectors are linearly independent.

Lemma 5.1 (Row Reduction Lemma). *Let \mathbb{F} be a finite field. Given an $r \times m$ matrix M over \mathbb{F} with a partition $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$ of $\{1, 2, \dots, m\}$, we can find, in time $O(rm^2)$, an $r' \times m$ matrix M' over \mathbb{F} with a partition $\mathcal{I}' = \{I'_1, I'_2, \dots, I'_n\}$ of $\{1, 2, \dots, m\}$ such that*

- (i) M' is of the standard form,
- (ii) $r' \leq r$ and $|I'_i| = |I_i|$ for all $i \in \{1, 2, \dots, n\}$,
- (iii) For all k , (M, \mathcal{I}, k) is a YES instance with a permutation σ of $\{1, 2, \dots, n\}$ if and only if (M', \mathcal{I}', k) is a YES instance with σ .

Proof. Let r' be the rank of M . This is easily achieved by applying elementary row operations and removing dependent rows to make the $r' \times r'$ identity submatrix and then permuting columns of M and adjusting \mathcal{I} accordingly so that the first r' columns form the $r' \times r'$ identity matrix. Note that $|I_i| = |I'_i|$ for all i . \square

Lemma 5.2 (Column Reduction Lemma). *Let \mathbb{F} be a finite field and let θ be an integer. Let M be an $r \times m$ matrix M over \mathbb{F} of the standard form with a partition $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$ of $\{1, 2, \dots, m\}$. Let V_i be the column space of $M[I_i]$ and let $\mathcal{V} = \{V_1, V_2, \dots, V_n\}$.*

In time $O(\theta r m n)$, we can either find an $r \times m'$ matrix over \mathbb{F} and a partition $\mathcal{I}' = \{I'_1, I'_2, \dots, I'_n\}$ of $\{1, 2, \dots, m'\}$ such that

- (i) *the column vectors of $M'[I'_i]$ are linearly independent for every $i = 1, 2, \dots, n$,*
- (ii) $|I'_i| \leq \min(|I_i|, \theta)$ for all $i = 1, 2, \dots, n$,
- (iii) *For all k , (M, \mathcal{I}, k) is a YES instance with a permutation σ of $\{1, 2, \dots, n\}$ if and only if (M', \mathcal{I}', k) is a YES instance with σ .*

or find $V_i \in \mathcal{V}$ such that $\dim(V_i \cap \langle \mathcal{V} - \{V_i\} \rangle) > \theta$.

Proof. Let $Z = \{1, 2, \dots, r\}$. Let $\bar{V}_i = \sum_{j \neq i, j \in \{1, 2, \dots, n\}} V_j$ and $\bar{I}_i = \{1, 2, \dots, m\} - I_i$.

For this algorithm, we will either find i such that $\dim V_i \cap \bar{V}_i > \theta$ or reduce the instance M and \mathcal{I} to a smaller instance M' and \mathcal{I}' whose subspaces V'_i (the column space of $M'[I'_i]$) have dimension at most θ by replacing V_i with $V_i \cap \bar{V}_i$.

Let $V'_i = V_i \cap \bar{V}_i$ and let $\mathcal{V}' = \{V'_1, V'_2, \dots, V'_n\}$. We first claim that for a subset $X \subseteq \{1, 2, \dots, n\}$ and $\bar{X} = \{1, 2, \dots, n\} - X$,

$$\left(\sum_{i \in X} V_i \right) \cap \left(\sum_{j \in \bar{X}} V_j \right) = \left(\sum_{i \in X} V'_i \right) \cap \left(\sum_{j \in \bar{X}} V'_j \right).$$

As $V'_i \subseteq V_i$ for all i , it is enough to show that $(\sum_{i \in X} V_i) \cap (\sum_{j \in \bar{X}} V_j) \subseteq (\sum_{i \in X} V'_i) \cap (\sum_{j \in \bar{X}} V'_j)$. Suppose $a \in (\sum_{i \in X} V_i) \cap (\sum_{j \in \bar{X}} V_j)$. Let $a_i \in V_i$ so that $a = \sum_{i \in X} a_i = \sum_{j \in \bar{X}} a_j$. If $i \in X$, then $a_i = \sum_{j \in \bar{X}} a_j - \sum_{j \in X - \{i\}} a_j \in \bar{V}_i$ and therefore $a_i \in V'_i$. Similarly $a_i \in V'_i$ for $i \in \bar{X}$. Thus $a \in (\sum_{i \in X} V'_i) \cap (\sum_{j \in \bar{X}} V'_j)$. This claim proves that for each permutation σ of $\{1, 2, \dots, n\}$, $V_{\sigma(1)}, V_{\sigma(2)}, \dots, V_{\sigma(n)}$ is a linear layout of width at most k if and only if $V'_{\sigma(1)}, V'_{\sigma(2)}, \dots, V'_{\sigma(n)}$ is a linear layout of width at most k .

Now we describe how to output an $r \times m'$ matrix M' with a partition $\mathcal{I}' = \{I'_1, I'_2, \dots, I'_n\}$ of $\{1, 2, \dots, m'\}$ such that V'_i is the column space of $M'[I'_i]$ for each $i \in \{1, 2, \dots, n\}$ or find i such that $\dim V'_i > \theta$.

For each $i \in \{1, 2, \dots, n\}$, we find the column basis $X_i \subseteq \bar{I}_i - Z$ of $M[Z \cap I_i, \bar{I}_i - Z]$ and the column basis $Y_i \subseteq I_i - Z$ of $M[Z - I_i, I_i - Z]$ and let

$$M'_i = \begin{pmatrix} M[Z \cap I_i, X_i] & 0 \\ 0 & M[Z - I_i, Y_i] \end{pmatrix}.$$

Let

$$M' = \begin{pmatrix} I'_1 & I'_2 & \cdots & I'_n \\ M'_1 & M'_2 & \cdots & M'_n \end{pmatrix}.$$

We claim that the column space of M'_i is exactly $V'_i = V_i \cap \bar{V}_i$. This can be seen from the following representation of M after permuting rows and columns:

$$M = \begin{array}{c} Z \cap I_i \\ Z \cap I_i \\ Z - I_i \end{array} \begin{pmatrix} Z \cap I_i & Z - I_i & I_i - Z & \bar{I}_i - Z \\ \begin{matrix} 1 & & & \\ & \ddots & & \\ & & 1 & \\ & & & \ddots & 1 \end{matrix} & \begin{matrix} 0 & & & \\ & \mathbf{A} & & \\ & & \mathbf{B} & \\ & & & \mathbf{C} & \mathbf{D} \end{matrix} & \end{pmatrix}.$$

The column spaces of $M[I_i]$ and $M[\bar{I}_i]$ do not change if we replace \mathbf{A} , \mathbf{D} with 0 and therefore for the computation of V'_i , we may assume that $\mathbf{A} = 0$ and $\mathbf{D} = 0$. Then $V'_i = V_i \cap \bar{V}_i$ is equal to the column space of the matrix $\begin{pmatrix} 0 & \mathbf{B} \\ \mathbf{C} & 0 \end{pmatrix}$. As we take the column basis X_i of \mathbf{B} and Y_i of \mathbf{C} , the column space of M'_i is equal to V'_i . Since $|X_i| + |Y_i| = \dim V'_i$, we can find i with $\dim V_i \cap \langle \mathcal{V} - \{V_i\} \rangle > \theta$ if $|X_i| + |Y_i| > \theta$.

Thus our algorithm can output M' with $\mathcal{I}' = \{I'_1, I'_2, \dots, I'_n\}$ or find i such that $\dim V'_i > \theta$. Let us now estimate its running time. It takes $O(\theta|Z \cap I_i| \cdot |\bar{I}_i - Z|)$ time to compute X_i and $O(\theta|Z - I_i| \cdot |I_i - Z|)$ time to compute Y_i or verify that $|X_i| + |Y_i| > \theta$ by applying the elementary row operations. (Note that we only need to apply at most θ pivots because if the rank is big, then $\dim V'_i$ is big.) Since

$$\theta \sum_{i=1}^n (|Z \cap I_i| \cdot |\bar{I}_i - Z| + |Z - I_i| \cdot |I_i - Z|) \leq \theta \sum_{i=1}^n r(m-r) \leq \theta r m n,$$

the total running time of the algorithm is at most $O(\theta r m n)$. \square

We apply the row reduction lemma first, the column reduction lemma second, and the row reduction lemma last to reduce the input instance M and \mathcal{I} to an equivalent instance M' and \mathcal{I}' , or find i such that $\dim V_i \cap \langle \mathcal{V} - V_i \rangle > \theta$, in time $O(rm^2) + O(\theta r m n) + O(rm^2) =$

$O(rm(m + \theta n))$. If there exists i such that $\dim V_i \cap \langle \mathcal{V} - V_i \rangle > \theta$, then we verify that the branch-width of \mathcal{V} is larger than θ . Thus from now on, we may assume the following for the input instance M and \mathcal{I} .

- $|I_i| \leq \theta$ and the columns of $M[I_i]$ are linearly independent for each i .
- The number r of rows of M is at most θn .

As usual, we let $V_i \subseteq \mathbb{F}^r$ be the column space of $M[I_i]$ and $\mathcal{V} = \{V_1, V_2, \dots, V_n\}$ be the subspace arrangement given by M and \mathcal{I} . Then by the above assumptions, we have the following.

- $\dim V_i \leq \theta$ for all i .
- $m \leq \theta n$.
- $r \leq \theta n$.

If $\dim V_i = 0$ for all i , then the path-width of \mathcal{V} is 0 and therefore we may assume that $\dim V_i > 0$ for some i and the path-width of \mathcal{V} is at least 1. Therefore, we may further assume that $k > 0$, $I_i \neq \emptyset$ and $\dim V_i > 0$ for all i by discarding such a subspace from \mathcal{V} .

We are going to solve the problem for the subspace arrangement \mathcal{V} given by M and \mathcal{I} , which are preprocessed. Once we find a linear layout σ of \mathcal{V} of width at most k , it is straightforward to obtain a linear layout of the original input of width at most k .

5.2. Preparing bases of boundary spaces. Given a branch-decomposition (T, \mathcal{L}) of width at most θ , in order to compute the full set $\text{FS}_k(\mathcal{V}_v, B_v)$ for every node v of T , we first need to figure out B_v . By Proposition 5.3, we can compute bases of the boundary space at each node and the sum of two boundary spaces of children at each internal node of T in time $O(\theta r m n)$.

Proposition 5.3. *Let \mathbb{F} be a fixed finite field and θ be a nonnegative integer. Let M be an $r \times m$ matrix of the standard form over \mathbb{F} given with a partition $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$ of $\{1, 2, \dots, m\}$. Let V_i be the column space of $M[I_i]$ and $\mathcal{V} = \{V_i : i = 1, 2, \dots, n\}$ be a subspace arrangement. Given a branch-decomposition (T, \mathcal{L}) of \mathcal{V} having width θ , we can compute the following for all nodes v of T , in time $O(\theta r m n)$.*

- A basis \mathcal{B}_v of the boundary space B_v .
- If v has two children w_1 and w_2 , then
 - a basis \mathcal{B}'_v of $B_{w_1} + B_{w_2}$ such that $\mathcal{B}_v \subseteq \mathcal{B}'_v$,
 - for each $i \in \{1, 2\}$, a transition matrix T_{w_i} from \mathcal{B}_{w_i} to \mathcal{B}'_v such that $T_{w_i} \cdot [x]_{\mathcal{B}_{w_i}} = [x]_{\mathcal{B}'_v}$ for all $x \in B_{w_i}$.

Proof. Let $Z = \{1, 2, \dots, r\}$. For each node v of T , we need to compute the basis \mathcal{B}_v of B_v . Let I_v be the union of all I_i such that $\mathcal{L}^{-1}(V_i)$ is v or the descendants of v in T . Let $\overline{I}_v = \{1, 2, \dots, m\} - I_v$. Then B_v is the intersection of the column spaces of $M[I_v]$ and $M[\overline{I}_v]$. Now observe the following representation of M after permuting rows and columns, similar to the proof of Lemma 5.2:

$$M = \begin{matrix} & \begin{matrix} Z \cap I_v & Z - I_v & I_v - Z & \overline{I}_v - Z \end{matrix} \\ \begin{matrix} Z \cap I_v \\ Z - I_v \end{matrix} & \begin{pmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & & 1 & & \\ & 0 & & & \ddots & \\ & & & & & 1 \end{pmatrix} \end{matrix}.$$

The column space of $M[I_v]$ is independent of \mathbf{A} and so we may assume $\mathbf{A} = 0$. Similarly we may assume $\mathbf{D} = 0$ as it does not affect the column space of $M[\overline{I}_v]$. Then B_v is indeed equal to the column space of $\begin{pmatrix} \mathbf{0} & \mathbf{B} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$. Furthermore $\dim B_v \leq \theta$ and therefore a basis \mathcal{B}_v of B_v can be found in time $O(\theta|Z \cap I_v| \cdot |\overline{I}_v - Z| + \theta|Z - I_v| \cdot |I_v - Z|)$. As $|Z \cap I_v|, |Z - I_v| \leq r$, and $|\overline{I}_v - Z| + |I_v - Z| \leq m - r \leq m$, we deduce that in time $O(\theta r m n)$, we can compute \mathcal{B}_v for all nodes v of T .

If v is an internal node of T with two children w_1 and w_2 , we need to compute a basis \mathcal{B}'_v of $B_{w_1} + B_{w_2}$ such that $\mathcal{B}_v \subseteq \mathcal{B}'_v$. Let M be a matrix of r rows and $|\mathcal{B}_v| + |\mathcal{B}_{w_1}| + |\mathcal{B}_{w_2}| (\leq 3\theta)$ columns of the form

$$M = \begin{pmatrix} P & Q & R \\ \mathcal{B}_v & \mathcal{B}_{w_1} & \mathcal{B}_{w_2} \end{pmatrix},$$

where each of \mathcal{B}_v , \mathcal{B}_{w_1} , and \mathcal{B}_{w_2} is considered as a matrix whose column vectors are the vectors in each of those sets. By applying the elementary row operations to M , we can find, in time $O(\theta r^2) \leq O(\theta r m)$, a column basis X of M . Note that $r \leq m$ because M is of a standard form. As $B_v \subseteq B_{w_1} + B_{w_2}$, we can choose X such that $P \subseteq X$. Then the column vectors of $M[X]$ form a basis \mathcal{B}'_v of $B_{w_1} + B_{w_2}$. Since we do this for each internal node v , it takes the time $O(\theta r m n)$ to compute \mathcal{B}'_v for all internal nodes v of T .

It remains to compute the transition matrices T_{w_1} and T_{w_2} for each internal node v with two children w_1 and w_2 . Let $i \in \{1, 2\}$. Our goal is to find a matrix T_{w_i} such that

$$T_{w_i} \cdot [x]_{\mathcal{B}_{w_i}} = [x]_{\mathcal{B}'_v}$$

for all $x \in \mathcal{B}_{w_i}$. Let $\mathcal{B}_{w_i} = \{b_1, b_2, \dots, b_\ell\} \subseteq \mathbb{F}^r$. Then if $x = b_j$, then $[x]_{\mathcal{B}_{w_i}} = e_j \in \mathbb{F}^\ell$. Thus the j -th column vector of T_{w_i} is equal to the coordinate of b_j with respect to \mathcal{B}'_v . In other words, we have

$$(\mathcal{B}'_v) \cdot T_{w_i} = (\mathcal{B}_{w_i})$$

where T_{w_i} is a $|\mathcal{B}'_v| \times |\mathcal{B}_{w_i}|$ matrix. This matrix equation can be solved in time $O(\theta r^2) \leq O(\theta r m)$. Thus we can compute T_{w_1} and T_{w_2} for all v in time $O(\theta r m n)$. \square

By Proposition 5.3, the following information can be computed in time $O(\theta r m n)$. Therefore, we assume that the following are given from Subsection 5.3 to Subsection 5.6.

- A branch-decomposition (T, \mathcal{L}) of \mathcal{V} of width at most θ .
- A basis \mathcal{B}_v of the boundary space B_v at every node v of T .
- A basis \mathcal{B}'_v of $B'_v = B_{w_1} + B_{w_2}$ extending \mathcal{B}_v at every internal node v of T having two children w_1 and w_2 .
- For each $i \in \{1, 2\}$, a transition matrix T_{w_i} from \mathcal{B}_{w_i} to \mathcal{B}'_v such that

$$T_{w_i} \cdot [x]_{\mathcal{B}_{w_i}} = [x]_{\mathcal{B}'_v} \text{ for all } x \in B_{w_i}$$

at every internal node v of T having two children w_1 and w_2 .

5.3. Data structure for the full sets. Before describing the algorithm, we present our data structure to store a B -trajectory in the full set $\text{FS}_k(\mathcal{V}, B)$ when we have a precomputed basis \mathcal{B} of B . For a B -trajectory $\Gamma = a_1, a_2, \dots, a_m$ with $a_i = (L_i, R_i, \lambda_i)$, we need to store subspaces L_i and R_i of B . Let $d = \dim B$.

If we want to represent a d' -dimensional subspace S of B such as L_i and R_i , a naive method is to pick a basis and make a matrix. As our vector space is a subspace of \mathbb{F}^r , one might use an $r \times d'$ matrix to represent S . However, our r depends on the input and may grow very large even if $\dim B$ is bounded.

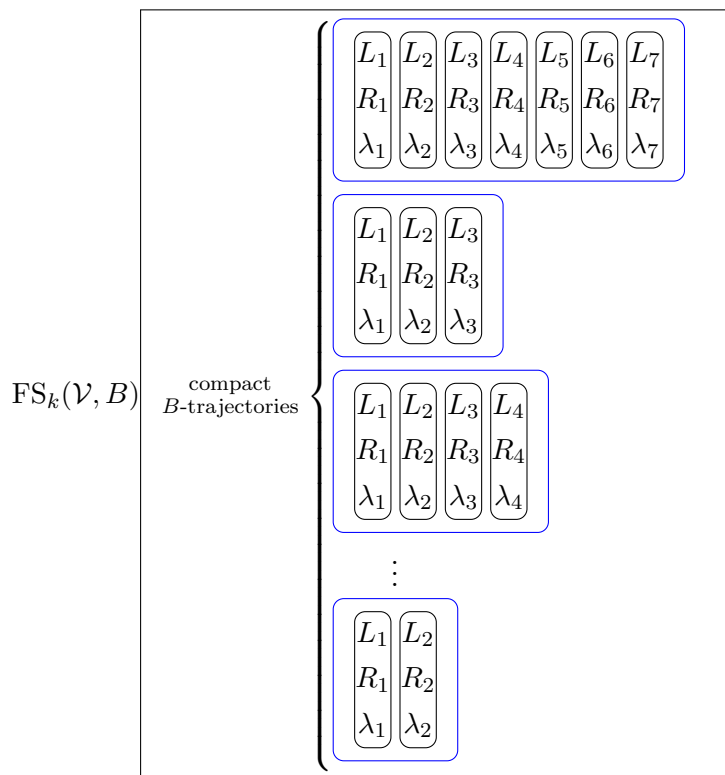


FIGURE 2. An illustration of a full set: λ_i are nonnegative integers and L_i, R_i are represented by matrices with $(\dim B)$ rows and linearly independent columns, each representing a vector in \mathbb{F}^r by using a basis \mathcal{B} of B .

In Subsection 2.3, we say that a $d \times d'$ matrix $M_{B,S}$ represents a d' -dimensional subspace S of B if the range of $M_B M_{B,S}$ is exactly S . We store subspaces L_i and R_i using a $d \times (\dim L_i)$ matrix M_{B,L_i} and a $d \times (\dim R_i)$ matrix M_{B,R_i} . This will ensure that each compact B -trajectory of width at most k can be stored in bounded amount of space when $\dim B$ is at most 2θ . Also, by Lemmas 2.5 and 2.6, many common operations on subspaces L_i and R_i can be done in time $O(d^3)$ when M_{B,L_i} and M_{B,R_i} are given. See Figure 2 for an illustration of a full set.

5.4. Computing the full sets. The procedure $\text{FULL-SET}(\mathcal{V}, k, (T, \mathcal{L}))$ aims to construct the full set $FS_k(\mathcal{V}, \{0\})$ at the root node of T . We compute a set \mathcal{F}_v at each node v , which will be shown to be $FS_k(\mathcal{V}_v, B_v)$ later. To this end, $\text{FULL-SET}(\mathcal{V}, k, (T, \mathcal{L}))$ recursively chooses a node v of T , farthest from the root (ties broken arbitrarily), such that \mathcal{F}_v is not computed yet.

The core of $\text{FULL-SET}(\mathcal{V}, k, (T, \mathcal{L}))$ is the *join* operation at an internal node v which combines the two full sets obtained at its children w_1 and w_2 . For this, the two full sets, which are represented with respect to the individual boundary spaces B_{w_1} and B_{w_2} of w_1 and w_2 , respectively, are adjusted via the *expand* operation. After that, two full sets which are represented with respect to the same subspace $B_{w_1} + B_{w_2}$ become amenable for join. After the join operation, the full set is again adjusted via the *shrink* operation

and now represented with respect to the boundary space B_v at node v . In this way, all subspaces under consideration have dimension at most 2θ .

Algorithm 1 Constructing the full sets

```

1: procedure FULL-SET( $\mathcal{V}, k, (T, \mathcal{L})$ )
2:   repeat
3:     choose an unmarked node  $v$  farthest from the root
4:     if  $v$  is a leaf then                                     ▷ initialization
5:       set  $\mathcal{F}_v$  as in Proposition 4.1
6:     else if  $v$  is internal with two children  $w_1, w_2$  then
7:       recall  $B'_v = B_{w_1} + B_{w_2}$ 
8:        $\mathcal{F}_v^{(i)} \leftarrow \text{up}_k(\mathcal{F}_{w_i}, B'_v)$  for  $i = 1, 2$        ▷ expand
9:        $\mathcal{F}'_v \leftarrow \text{up}_k(\mathcal{F}_v^{(1)} \oplus \mathcal{F}_v^{(2)}, B'_v)$            ▷ join
10:       $\mathcal{F}_v \leftarrow \text{up}_k(\mathcal{F}'_v|_{B_v}, B_v)$                        ▷ shrink
11:     end if
12:     mark  $v$ 
13:   until all nodes in  $T$  are marked
14: end procedure

```

Recall that $U_k(B)$ is the set of all compact B -trajectories of width at most k and that $|U_k(B)|$ is at most $2^{9d+2}|\mathbb{F}|^{d(d-1)}2^{2(2d+1)k}$ by Lemma 3.3, where $d = \dim(B)$. Clearly, $\text{FS}_k(\mathcal{V}, B) \subseteq U_k(B)$ and therefore we consider the elements of $U_k(B)$ as candidates for $\text{FS}_k(\mathcal{V}, B)$. To evaluate the number of steps carried out by $\text{FULL-SET}(\mathcal{V}, k, (T, \mathcal{L}))$, we need the following lemmas.

Lemma 5.4. *There are at most 3^{m+n-2} distinct lattice paths with diagonal steps from $(1, 1)$ to (m, n) .*

Proof. For every lattice path with diagonal steps v_1, v_2, \dots, v_ℓ from $(1, 1)$ to (m, n) , we have $\ell \leq m + n - 1$. At each i , we have three possible choices. \square

Lemma 5.5. *Let B and B' be subspaces of dimension at most 2θ given with bases \mathcal{B} and \mathcal{B}' , respectively, such that $B \subseteq B'$. Let T be the transition matrix from \mathcal{B} to \mathcal{B}' so that $T \cdot [x]_{\mathcal{B}} = [x]_{\mathcal{B}'}$ for all $x \in B$. Then, for a B -trajectory Δ represented with respect to \mathcal{B} and a B' -trajectory Γ represented with respect to \mathcal{B}' , we can decide $\Delta \preceq \Gamma$ in time $\theta^3(|\Delta| + |\Gamma|)3^{|\Delta|+|\Gamma|} \cdot O(1)$.*

Proof. First we transform the representation of Δ into a representation with respect to \mathcal{B}' . This can be done by multiplying T to each of the matrices representing bases of subspaces appearing in Δ . Since both the number of rows and the number of columns are at most 2θ and each subspace in Δ has dimension at most 2θ , we can transform the representation of Δ in time $|\Delta|\theta^3 \cdot O(1)$.

Now we are prepared with Δ and Γ , both represented with respect to \mathcal{B}' . In order to test $\Delta \preceq \Gamma$, we need to consider all possible lattice paths from $(1, 1)$ to $(|\Delta|, |\Gamma|)$. By Lemma 5.4, there are at most $3^{|\Delta|+|\Gamma|-2}$ such paths. For each path, we need to compare an extension of Δ with an extension of Γ . Comparing two subspaces of B' takes time $\theta^3 \cdot O(1)$ and there are $2(|\Delta| + |\Gamma| - 1)$ comparisons of subspaces to make for each path. Thus, we can decide $\Delta \preceq \Gamma$ in time $|\Delta|\theta^3 \cdot O(1) + \theta^3(|\Delta| + |\Gamma|)3^{|\Delta|+|\Gamma|} \cdot O(1) = \theta^3(|\Delta| + |\Gamma|)3^{|\Delta|+|\Gamma|} \cdot O(1)$. \square

Lemma 5.6. *Let k be an integer. Let B, B' be subspaces such that $B \subseteq B'$ and let \mathcal{R} be a set of B -trajectories. To compute $\text{up}_k(\mathcal{R}, B')$, it is enough to perform at most $|U_k(B')| \cdot |\mathcal{R}|$ comparisons, each deciding whether $\Delta \preceq \Gamma$ for some $\Delta \in \mathcal{R}$ and $\Gamma \in U_k(B')$.*

Proof. Recall that $\Gamma \in \text{up}_k(\mathcal{R}, B)$ if and only if $\Gamma \in U_k(B)$ and there exists $\Delta \in \mathcal{R}$ such that $\Delta \preceq \Gamma$. Therefore, for each pair $\Gamma \in U_k(B)$ and $\Delta \in \mathcal{R}$ we verify whether $\Delta \preceq \Gamma$. \square

Below, we explain how to perform each operation in $\text{FULL-SET}(\mathcal{V}, k, (T, \mathcal{L}))$ and argue (in the proof of Proposition 5.7) that each operation takes at most $\text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{151\theta k} |\mathbb{F}|^{12\theta^2}$ steps.

Initialization: At each leaf v of T , which is handled at line 5, Proposition 4.1 states that the set \mathcal{F}_v is equal to $\text{up}_k(\{\Delta\}, B_v)$ where $\Delta = (\{0\}, B_v, 0), (B_v, \{0\}, 0)$.

By Lemma 5.6, computing \mathcal{F}_v requires at most $|U_k(B_v)|$ comparisons for $\Delta \preceq \Gamma$ between Δ and $\Gamma \in U_k(B_v)$.

Expand: At line 8, for $i = 1, 2$, we compute the set $\text{up}_k(\mathcal{F}_{w_i}, B'_v)$. By Lemma 5.6, this requires at most $|\mathcal{F}_{w_i}| \cdot |U_k(B'_v)| \leq |U_k(B_{w_i})| \cdot |U_k(B'_v)|$ comparisons between $\Delta \in \mathcal{F}_{w_i} \subseteq U_k(B_{w_i})$ and $\Gamma \in U_k(B'_v)$.

Join: At line 9, computing $\mathcal{F}'_v = \text{up}_k(\mathcal{F}_v^{(1)} \oplus \mathcal{F}_v^{(2)}, B'_v)$, where $B'_v = B_{w_1} + B_{w_2}$, consists of two steps: (a) to compute $\mathcal{F}_v^{(1)} \oplus \mathcal{F}_v^{(2)}$, and (b) to take $\text{up}_k(\mathcal{F}_v^{(1)} \oplus \mathcal{F}_v^{(2)}, B'_v)$. Note that, for (b), the set $\text{up}_k(\mathcal{F}_v^{(1)} \oplus \mathcal{F}_v^{(2)}, B'_v)$ can be computed in at most $|U_k(B'_v)|^3$ comparisons by Lemma 5.6 since $|\mathcal{F}_v^{(1)} \oplus \mathcal{F}_v^{(2)}| \leq |U_k(B'_v)|^2$.

Shrink: At line 10, we compute the set $\text{up}_k(\mathcal{F}'_v|_{B_v}, B_v)$ in two steps: (a) to compute $\mathcal{F}'_v|_{B_v}$, and (b) to compute $\text{up}_k(\mathcal{F}'_v|_{B_v}, B_v)$. Here, by Lemma 5.6, the set $\text{up}_k(\mathcal{F}'_v|_{B_v}, B_v)$ can be computed by comparing at most $|U_k(B_v)| \cdot |U_k(B_v)|$ pairs Δ and Γ , where $\Delta \in \mathcal{F}'_v|_{B_v} \subseteq U_k(B_v)$ and $\Gamma \in U_k(B_v)$.

Proposition 5.7. *Let $\theta \geq 1$. The procedure $\text{FULL-SET}(\mathcal{V}, k, (T, \mathcal{L}))$ runs in $O(\text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{151\theta k} |\mathbb{F}|^{12\theta^2} \cdot n)$ steps, when θ is the width of (T, \mathcal{L}) .*

Proof. In each operation, when we compare two trajectories Δ and Γ to decide $\Delta \preceq \Gamma$, we know, by Lemma 3.2, $|\Delta| \leq 2(4\theta + 1)(2k + 1) \leq 28\theta k + 2$ and $|\Gamma| \leq (4\theta + 1)(2k + 1) \leq 14\theta k + 1$; the worst case happens at the join operation. Thus, by Lemma 5.5, one comparison can be done in time $\theta^3(|\Delta| + |\Gamma|)3^{|\Delta| + |\Gamma|} \cdot O(1) \leq \theta^3(42\theta k + 3) \cdot 3^{42\theta k + 3} \cdot O(1) \leq \text{poly}(\theta, k) \cdot 3^{42\theta k} < \text{poly}(\theta, k) 2^{67\theta k}$. Note that $3^{42} < 2^{67}$.

Recall that $\dim(B_v) \leq \theta$, $\dim(B'_v) \leq 2\theta$, and $\dim(B_{w_i}) \leq \theta$ for $i = 1, 2$. By Lemma 3.3,

$$\begin{aligned} |U_k(B_v)| &\leq 2^{9\theta+2} |\mathbb{F}|^{\theta(\theta-1)} 2^{2(2\theta+1)k} \leq 2^{15\theta k+2} |\mathbb{F}|^{\theta^2}, \\ |U_k(B'_v)| &\leq 2^{18\theta+2} |\mathbb{F}|^{2\theta(2\theta-1)} 2^{2(4\theta+1)k} \leq 2^{28\theta k+2} |\mathbb{F}|^{4\theta^2}, \text{ and} \\ |U_k(B_{w_i})| &\leq 2^{9\theta+2} |\mathbb{F}|^{\theta(\theta-1)} 2^{2(2\theta+1)k} \leq 2^{15\theta k+2} |\mathbb{F}|^{\theta^2} \text{ for } i = 1, 2. \end{aligned}$$

When B is a vector space with $\dim(B) \leq 2\theta$, Lemma 3.3 implies that $U_k(B)$ can be generated in $\text{poly}(\theta, |\mathbb{F}|, k) \cdot |U_k(B)| \leq \text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{28\theta k+2} |\mathbb{F}|^{4\theta^2}$ steps.

Let us analyze the time complexity for each operation.

Initialization: Since computing \mathcal{F}_v requires at most $|U_k(B_v)|$ comparisons, it can be computed in time $\text{poly}(\theta, k) 2^{67\theta k} \cdot 2^{15\theta k+2} |\mathbb{F}|^{\theta^2} \leq \text{poly}(\theta, k) \cdot 2^{82\theta k} |\mathbb{F}|^{\theta^2}$.

Expand: The set $\text{up}_k(\mathcal{F}_{w_i}, B'_v)$ can be computed in time $\text{poly}(\theta, k)2^{67\theta k} \cdot |U_k(B_{w_i})| \cdot |U_k(B'_v)| \leq \text{poly}(\theta, k)2^{67\theta k} \cdot 2^{15\theta k+2}|\mathbb{F}|^{\theta^2} \cdot 2^{28\theta k+2}|\mathbb{F}|^{4\theta^2} \leq \text{poly}(\theta, k) \cdot 2^{110\theta k}|\mathbb{F}|^{5\theta^2}$.

Join: Recall that $\mathcal{F}_v^{(1)} \oplus \mathcal{F}_v^{(2)} = \bigcup_{\Delta_1 \in \mathcal{F}_v^{(1)}, \Delta_2 \in \mathcal{F}_v^{(2)}} \Delta_1 \oplus \Delta_2$. Since $\mathcal{F}_v^{(i)} \subseteq U_k(B'_v)$ for $i = 1, 2$ after an expanding operation, there are at most $|U_k(B'_v)|^2$ pairs Δ_1, Δ_2 to consider. For each pair Δ_1 and Δ_2 , computing $\Delta_1 \oplus \Delta_2$ requires considering all lattice paths from $(1, 1)$ to $(|\Delta_1|, |\Delta_2|)$. The number of such paths is at most $2^{|\Delta_1|+|\Delta_2|}$ and thus, (a) takes at most $|U_k(B'_v)|^2 \cdot 2^{|\Delta_1|+|\Delta_2|}$ steps. By Lemma 3.2, $|U_k(B'_v)|^2 \cdot 2^{|\Delta_1|+|\Delta_2|} \leq (2^{28\theta k+2}|\mathbb{F}|^{4\theta^2})^2 \cdot 2^{2(4\theta+1)(2k+1)} \leq 2^{84\theta k+6}|\mathbb{F}|^{8\theta^2}$.

For (b), the set $\text{up}_k(\mathcal{F}_v^{(1)} \oplus \mathcal{F}_v^{(2)}, B'_v)$ can be constructed in time $\text{poly}(\theta, k)2^{67\theta k} \cdot |U_k(B'_v)|^3 \leq \text{poly}(\theta, k)2^{67\theta k} \cdot (2^{28\theta k+2}|\mathbb{F}|^{4\theta^2})^3 \leq \text{poly}(\theta, k) \cdot 2^{151\theta k}|\mathbb{F}|^{12\theta^2}$. Therefore, the join operation can be performed in $\text{poly}(\theta, k) \cdot 2^{151\theta k}|\mathbb{F}|^{12\theta^2}$ steps.

Shrink: For (a), we consider every $\Delta \in \mathcal{F}'_v$ and take $\Delta|_{B_v}$. For each $\Delta \in \mathcal{F}'_v$, computing $\Delta|_{B_v}$ can be done in time $\theta^3|\Delta| \cdot O(1)$ (see Subsection 2.3). Since $\mathcal{F}'_v \subseteq U_k(B'_v)$, there are at most $|U_k(B'_v)|$ elements to consider. By Lemma 3.2, $|\Delta| \leq (4\theta + 1)(2k + 1)$. It follows that (a) can be done in $\theta^3(4\theta + 1)(2k + 1) \cdot |U_k(B'_v)| \cdot O(1) \leq \text{poly}(\theta, k) \cdot 2^{28\theta k}|\mathbb{F}|^{4\theta^2}$ steps.


To compute $\text{up}_k(\mathcal{F}'_v|_{B_v}, B_v)$, we need to compare two trajectories at most $|U_k(B_v)| \cdot |U_k(B'_v)|$ times. Thus, (b) takes at most $\text{poly}(\theta, k)2^{67\theta k} \cdot |U_k(B_v)| \cdot |U_k(B'_v)| \leq \text{poly}(\theta, k)2^{67\theta k} \cdot (2^{15\theta k+2}|\mathbb{F}|^{\theta^2})^2 \leq \text{poly}(\theta, k) \cdot 2^{97\theta k}|\mathbb{F}|^{2\theta^2}$ steps. Hence the shrink operation can be done in $\text{poly}(\theta, k) \cdot 2^{97\theta k}|\mathbb{F}|^{4\theta^2}$ steps.

Therefore, the running time of one iteration of the repeat-loop is $\text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{28\theta k+2}|\mathbb{F}|^{4\theta^2} + \text{poly}(\theta, k) \cdot 2^{82\theta k}|\mathbb{F}|^{\theta^2} + \text{poly}(\theta, k) \cdot 2^{110\theta k}|\mathbb{F}|^{5\theta^2} + \text{poly}(\theta, k) \cdot 2^{151\theta k}|\mathbb{F}|^{12\theta^2} + \text{poly}(\theta, k) \cdot 2^{97\theta k}|\mathbb{F}|^{4\theta^2} \leq \text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{151\theta k}|\mathbb{F}|^{12\theta^2}$. Notice that the tree T contains n leaf nodes and $n - 1$ internal nodes. This completes the proof. \square

Proposition 5.8. *The path-width of \mathcal{V} is at most k if and only if $\mathcal{F}_{\text{root}} \neq \emptyset$ at the root node of T .*

Proof. First, we claim that for each node v of T , we have $\mathcal{F}_v = \text{FS}_k(\mathcal{V}_v, B_v)$, and furthermore $\mathcal{F}'_v = \text{FS}_k(\mathcal{V}_v, B'_v)$ and $\mathcal{F}_v^{(i)} = \text{FS}_k(\mathcal{V}_{w_i}, B'_v)$ for $i = 1, 2$ when v has two children w_1 and w_2 . We prove this by induction on the number of steps executed by $\text{FULL-SET}(\mathcal{V}, k, (T, \mathcal{L}))$. The claim holds for \mathcal{F}_v produced at line 5 by Proposition 4.1. At line 8, Lemma 2.7 implies $\langle \mathcal{V}_v \rangle \cap (B_{w_1} + B_{w_2}) = B_v$. Since the condition of Proposition 4.2 holds, it follows that $\mathcal{F}_v^{(i)} = \text{FS}_k(\mathcal{V}_{w_i}, B'_v)$ for $i = 1, 2$. For \mathcal{F}'_v produced at line 9, notice that $(\langle \mathcal{V}_{w_1} \rangle + B_{w_1} + B_{w_2}) \cap (\langle \mathcal{V}_{w_2} \rangle + B_{w_1} + B_{w_2}) = B_{w_1} + B_{w_2}$ by Lemma 2.7. Hence, the condition of Proposition 4.4 is satisfied. By the induction hypothesis and Proposition 4.4, we have $\mathcal{F}'_v = \text{FS}_k(\mathcal{V}_v, B'_v)$. At line 10, observe that $B_v \subseteq B_{w_1} + B_{w_2}$ by Lemma 2.7. The condition of Proposition 4.3 holds, which implies $\mathcal{F}_v = \text{FS}_k(\mathcal{V}_v, B_v)$. This completes the proof of our claim.

From the above claim, it immediately follows that the forward implication holds. For the opposite direction, let $\Gamma \in \mathcal{F}_{\text{root}} = \text{FS}_k(\mathcal{V}_{\text{root}}, B_{\text{root}})$ be a compact B_{root} -trajectory of width at most k . Note that $\mathcal{V}_{\text{root}} = \mathcal{V}$ and $B_{\text{root}} = \{0\}$. By definition, there exists a $\{0\}$ -trajectory Δ realizable in \mathcal{V} such that $\Delta \preceq \Gamma$. This means that there exists a linear layout $\sigma = V_1, V_2, \dots, V_m$ of \mathcal{V} whose canonical $\{0\}$ -trajectory a_1, a_2, \dots, a_{m+1} is an extension of Δ . Since Γ is of width at most k and $\Delta \preceq \Gamma$, Lemma 3.4 implies that Δ is indeed of width

FIGURE 3. Subtrees of T' for an internal node v with two children w_1, w_2 .

at most k . Recall that $L(a_i)$ and $R(a_i)$ are subspaces of $\{0\}$, and we have

$$\begin{aligned} \lambda(a_i) &= \dim(V_1 + V_2 + \cdots + V_{i-1}) \cap (V_i + V_{i+1} + \cdots + V_m) \\ &\quad - \dim L(a_i) \cap R(a_i) \\ &= \dim(V_1 + V_2 + \cdots + V_{i-1}) \cap (V_i + V_{i+1} + \cdots + V_m) \end{aligned}$$

for every $i = 1, \dots, m+1$. Therefore, σ is a linear layout of \mathcal{V} of width at most k . \square

5.5. Backtracking to construct a linear layout. In this subsection, we illustrate how to construct an optimal linear layout of \mathcal{V} if $\text{FS}_k(\mathcal{V}, \{0\})$ is nonempty. First let T' be a tree obtained from T as follows:

- The vertex set of T' is given as the disjoint union of $\{v^1, v^2, v^J, v^U, v^S, v\}$ for each internal node v of T and $\{v^L, v\}$ for each leaf node v of T .
- For each internal node v of T , $v^1v^J, v^2v^J, v^Jv^U, v^Uv^S, v^Sv$ are edges of T' induced on $\{v^1, v^2, v^J, v^U, v^S, v\}$.
- For each leaf node v of T , v^Lv is an edge of T' .
- For each internal node v of T and its two children w_1 and w_2 , T' has edges w_1v^1 and w_2v^2 .

This tree T' is called the *composition tree* of T . See Figure 3. There are four kinds of nodes in T' .

- Vertices v^U, v^1, v^2 of T' for all internal nodes v of T and vertices v of T' for all nodes v of T are *up nodes*.
- Vertices v^J of T' for all internal nodes v of T are *join nodes*.
- Vertices v^L of T' for all leaf nodes v of T are *leaf nodes*.
- Vertices v^S of T' for all internal nodes v of T are called *shrink nodes*.

The composition tree T' describes how our dynamic programming works as follows. The nodes of T' can be grouped so that each group corresponds to the computational steps for a full set in Algorithm 1. For a leaf node v of T , v^L and v correspond to Line 5. For an internal node v of T , v^1 and v^2 correspond to Line 8, v^J and v^U correspond to Line 9, and v^S and v correspond to Line 10.

We choose $\Gamma \in \text{FS}_k(\mathcal{V}, \{0\})$ with the minimum width. The width of Γ is going to be equal to the path-width of \mathcal{V} . By backtracking how Γ is placed in $\text{FS}_k(\mathcal{V}, \{0\})$ in our algorithm, we can label each node x of T' by a B_x -trajectory Γ_x for some B_x and other necessary information satisfying the following.

- $B_x = B_v$ if $x = v$, $x = v^S$, or $x = v^L$, and $B_x = B'_v$ if $x = v^U$, $x = v^J$, $x = v^1$, or $x = v^2$ for some node v of T .
- If x is a shrink node with a child y , then $\Gamma_x = \Gamma_y|_{B_x}$.

- If x is a up node with a child y , then $\Gamma_y \preceq \Gamma_x$ and there is a lattice path P_x with diagonal steps as given in Lemma 3.5. Then there is a sequence $1 = x_1 \leq x_2 \leq \dots \leq x_{|\Gamma_x|} = |\Gamma_y|$ of integers such that P_x goes through points (x_j, j) for all j . This sequence $x_1, x_2, \dots, x_{|\Gamma_x|}$ is also stored at x .
- If x is a join node with two children y and z , then $\Gamma_x = \Gamma_y +_{P_x} \Gamma_z$ for some lattice path P_x . This lattice path P_x is stored at the node x .
- If x is a leaf node v^L for a leaf node v of T , then

$$\Gamma_x = \begin{cases} (\{0\}, B, 0), (B, \{0\}, 0) & \text{if } B_v \neq \{0\}, \\ (\{0\}, \{0\}, 0) & \text{if } B_v = \{0\}. \end{cases}$$

- If x is the root node of T' , then $\Gamma_x = \Gamma \in \text{FS}_k(\mathcal{V}, \{0\})$.

We modify $\text{FULL-SET}(\mathcal{V}, k, (T, \mathcal{L}))$ so that whenever an operation is carried out, for every Γ which is placed in $\mathcal{F}_v, \mathcal{F}'_v, \mathcal{F}_v^{(1)}$, or $\mathcal{F}_v^{(2)}$, a *certificate* for Γ is stored so that later when we find $\Gamma \in \text{FS}_k(\mathcal{V}, \{0\})$, we can construct the *labeled composition tree* T' by backtracking.

Algorithm 2 Print a linear layout of width at most k

```

1: procedure PRINTORDER( $x, i$ )
2:   if  $x$  is a leaf and  $i = 1$  then
3:     print the space  $V$  in  $\mathcal{V}$  such that  $(\mathcal{L}^{-1}(V))^L = x$ 
4:   else if  $x$  is a join node with two children  $y$  and  $z$  then
5:     assume that  $P_x = v_1, v_2, \dots, v_t$ . Let  $v_i = (x_i, y_i)$ 
6:     if  $v_{i+1} - v_i = (1, 0)$  then
7:       call PRINTORDER( $y, x_i$ )
8:     else
9:       call PRINTORDER( $z, y_i$ )
10:    end if
11:   else if  $x$  is a up node with a child  $y$  then
12:     recall the sequence  $1 = x_1 \leq x_2 \leq \dots \leq x_{|\Gamma_x|} = |\Gamma_y|$  so that the lattice path  $P_x$ 
with diagonal steps goes through points  $(x_j, j)$ 
13:     call PRINTORDER( $y, j$ ) for all  $j = x_i, x_i + 1, \dots, x_{i+1} - 1$ 
14:   else
15:     for a child  $y$  of  $x$ , call PRINTORDER( $y, i$ )
16:   end if
17: end procedure
18: procedure ORDER
19:   print all  $V \in \mathcal{V}$  if  $B_x = \{0\}$  and  $\mathcal{L}(x) = V$ 
20:   call PRINTORDER( $x, i$ ) for the root  $x$  of  $T'$  and  $1 \leq i < |\Gamma_x|$ 
21: end procedure

```

Proposition 5.9. *The algorithm ORDER in Algorithm 2 correctly finds a linear layout of \mathcal{V} of width at most k in time $O(\theta kn)$ when a labeled composition tree T' obtained from a branch-decomposition (T, \mathcal{L}) of width at most θ and a $\{0\}$ -trajectory $\Gamma \in \text{FS}_k(\mathcal{V}, \{0\})$ are given.*

Proof. We sketch the proof for the correctness. Consider a B -trajectory $\Gamma = a_1, a_2, \dots, a_{m+1}$ realizable in a subspace arrangement \mathcal{V} . Then there exists the canonical B -trajectory

$\Gamma' = a'_1, a'_2, \dots, a'_{n+1} \in E(\Gamma)$ of a linear layout $V_{\sigma(1)}, V_{\sigma(2)}, \dots, V_{\sigma(n)}$ of \mathcal{V} . By definition, we can say that V_i is between a'_i and a'_{i+1} . Since Γ' is an extension of Γ , we can naturally say that V_i is between a_j and a_{j+1} if $a'_i = a_j$ and $a'_{i+1} \in \{a_j, a_{j+1}\}$. Let $X_j = \{V_i : V_i \text{ is between } a_j \text{ and } a_{j+1} \text{ for all } 1 \leq i \leq n\}$ for all $1 \leq j \leq m$. Then X_1, X_2, \dots, X_m is a partition of \mathcal{V} . Thus, we can regard each realizable B -trajectory to be equipped with such an ordered partition of \mathcal{V} .

For each Γ in the full set, there is a realizable Δ with $\Delta \preceq \Gamma$ and so we can also give such an ordered partition to Γ . The only place where we mix two ordered partitions is at the join node of T' and we have the complete description how to combine linear layouts by Lemma 3.24. (For a leaf x with $B_x = \{0\}$ and $\mathcal{L}(x) = V$, it does not matter where this subspace $V \in \mathcal{V}$ is placed in the linear layout and so we simply put it in the beginning of the linear layout.)

Now let us discuss the running time. Each node x of T' is visited $|\Gamma_x| - 1$ times because $1 \leq i < |\Gamma_x|$ is always ensured in PRINTORDER. If a node x of T' is neither a join node nor a shrink node, then x is labeled by a compact B_x -trajectory and $\dim B_x \leq 2\theta$ and therefore $|\Gamma_x| \leq (4\theta + 1)(2k + 1)$ by Lemma 3.2.

If x is a join node with two children y and z , then Γ_x might not be compact, but $\Gamma_x = \Gamma_y +_{P_x} \Gamma_z$ for a compact B_y -trajectory Γ_y and a compact B_z -trajectory Γ_z . Since $\dim B_y \leq \theta$ and $\dim B_z \leq \theta$, both Γ_y and Γ_z have length at most $(2\theta + 1)(2k + 1)$. Then Γ_x has length at most $2(2\theta + 1)(2k + 1) - 1$.

If x is a shrink node with a child y , then $\Gamma_x = \Gamma_y|_{B_x}$ and Γ_y is a B_y -trajectory with $\dim B_y \leq 2\theta$ and so $|\Gamma_x| = |\Gamma_y| \leq (4\theta + 1)(2k + 1)$.

In all cases, we deduce that $|\Gamma_x| = \theta k \cdot O(1)$. Thus the algorithm visits each node of T' at most $\theta k \cdot O(1)$ times and T' has $O(n)$ nodes and therefore the running time of this algorithm is bounded by $O(\theta kn)$. \square

5.6. Summary. The following theorem summarizes the procedure and the time complexity for an algorithm that solves the problem introduced at the beginning of Section 5.

Theorem 5.10. *Let k, θ , and r be nonnegative integers and let \mathbb{F} be a fixed finite field. Let $\mathcal{V} = \{V_1, V_2, \dots, V_n\}$ be a subspace arrangement of subspaces of \mathbb{F}^r and (T, \mathcal{L}) be a branch-decomposition of \mathcal{V} of width at most θ . One can find in time $O(rm(m + \theta n) + \text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{151\theta k} |\mathbb{F}|^{12\theta^2} n)$ a linear layout of \mathcal{V} of width at most k , or confirm that no such linear layout exists, where each V_i is given by its spanning set of d_i vectors and $m = \sum_{i=1}^n d_i$.*

We remark that the algorithm can output a linear layout of the minimum width if there is a linear layout of width at most k , by choosing a $\{0\}$ -trajectory of the minimum width in the full set.

Proof. First, apply the preprocessing steps discussed in Subsection 5.1 using Lemmas 5.1 and 5.2 in time $O(rm(m + \theta n))$ in order to simplify the input. If we find $V \in \mathcal{V}$ with $\dim(V \cap (\mathcal{V} - \{V\})) > k$ in Lemma 5.2, then we confirm that no such linear layout exists. For convenience, let \mathcal{V} denote the subspace arrangement after the preprocessing and let $|\mathcal{V}| = n$. If $n \leq 1$, then the path-width of \mathcal{V} is 0 and so we can output an arbitrary linear layout when $k \geq 0$. Thus we may assume that $n \geq 2$. Also, we may assume that $k > 0$ and $\theta > 0$ as described in Subsection 5.1.

Then we run the algorithm due to Proposition 5.3 in time $O(\theta rmn)$ so that we are ready to run Algorithm 1.

By Proposition 5.7, Algorithm 1 can be executed in $O(\text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{151\theta k} |\mathbb{F}|^{12\theta^2} \cdot n)$ steps. If Algorithm 1 outputs an empty full set for the root node, then no linear layout of width at most k exists. If Algorithm 1 outputs a nonempty full set for the root node, then we construct a linear layout of width at most k using Algorithm 2, which takes time $O(\theta kn)$ by Proposition 5.9. When Algorithm 2 gives a linear layout of width at most k , we can easily convert it into a linear layout of the input subspace arrangement (before the preprocessing) in linear time.

Therefore, the total time complexity is $O(rm(m+\theta n)+\theta rmn+\text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{151\theta k} |\mathbb{F}|^{12\theta^2} \cdot n + \theta kn) = O(rm(m+\theta n) + \text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{151\theta k} |\mathbb{F}|^{12\theta^2} \cdot n)$. \square

6. HOW TO PROVIDE AN APPROXIMATE BRANCH-DECOMPOSITION

Up to the previous section, we assumed that we are given a branch-decomposition of width at most θ upon which our algorithm is based. But how do we obtain such a branch-decomposition in the beginning? We present two methods to accomplish this goal. We assume that an input subspace arrangement contains at least two subspaces otherwise it has path-width 0.

6.1. Iterative compression. Our first method is to use our algorithm iteratively. This strategy is sometimes called iterative compression; roughly speaking, for a subspace arrangement $\mathcal{V} = \{V_i\}_{i=1}^n$ and $2 \leq \ell < n$, we will iteratively apply Algorithms 1 and 2 to a subset $\{V_i\}_{i=1}^{\ell}$ of \mathcal{V} and then modify the output to obtain a branch-decomposition of small width of $\{V_i\}_{i=1}^{\ell+1}$, which is needed for the next step. This requires to run Algorithms 1 and 2 $O(n)$ times and has an advantage of making this paper self-contained.

For this approach, we need the following simple lemma whose trivial proof we omit. A tree is a *caterpillar tree* if it has a path P such that every vertex is either a vertex in P or is adjacent to some vertex in P .

Lemma 6.1. *Let $\mathcal{V} = \{V_i\}_{i=1}^n$ be a subspace arrangement over \mathbb{F} with $n \geq 3$ and $\dim V_i \leq 2k$ for all $i = 1, 2, \dots, n$. If V_1, V_2, \dots, V_{n-1} is a linear layout of width at most k , then V_1, V_2, \dots, V_n is a linear layout of width at most $3k$.*

We argue that for an input subspace arrangement \mathcal{V}_{in} of subspaces of \mathbb{F}^r , in time $O(r(|\mathcal{V}_{\text{in}}|^3 + m^2) + f(k)|\mathcal{V}_{\text{in}}|^2)$ we can either find a linear layout of width at most k or confirm that no such linear layout exists, where each subspace in \mathcal{V}_{in} is given by its spanning set of d_i vectors and $m = \sum_{i=1}^{|\mathcal{V}_{\text{in}}|} d_i$. Given a subspace arrangement \mathcal{V}_{in} of subspaces of \mathbb{F}^r , we run the preprocessing described in Subsection 5.1 with $\theta = 2k$. Note that the preprocessing takes time $O(rm(m+2k|\mathcal{V}_{\text{in}}|)) = O(rm^2)$. During the preprocessing, the algorithm due to Lemma 5.2 may find a subspace $V \in \mathcal{V}_{\text{in}}$ such that $\dim(V \cap \langle \mathcal{V}_{\text{in}} - \{V\} \rangle) > 2k$. If so, we can confirm that \mathcal{V}_{in} has path-width larger than k by Proposition 2.2. Otherwise, let \mathcal{V} be the resulting subspace arrangement with $|\mathcal{V}| \leq |\mathcal{V}_{\text{in}}|$. If \mathcal{V} contains at most one element, then every linear layout of \mathcal{V} has width 0 and thus every linear layout of \mathcal{V}_{in} has width 0. Thus, we may assume that $\mathcal{V} = \{V_1, V_2, \dots, V_n\}$ contains at least two subspaces, that is, $2 \leq n \leq |\mathcal{V}_{\text{in}}|$.

Now we iterate over $\ell = 2, \dots, n$ and we may assume that we are given a linear layout $\sigma_{\ell-1}$ of width at most k for a subspace arrangement $\{V_1, \dots, V_{\ell-1}\}$. If this is not the case, we declare that \mathcal{V} allows no linear layout of width at most k and terminate. For the base case when $\ell = 2$, we have the unique linear layout σ_1 of $\{V_1\}$ whose width is

0. By Lemma 6.1, we can construct a linear layout of width at most $3k$ for a subspace arrangement $\{V_1, \dots, V_\ell\}$. This can be converted into a branch-decomposition (T, \mathcal{L}) of $\{V_1, \dots, V_\ell\}$ whose width is at most $3k$ by taking a subcubic caterpillar tree T whose leaves are mapped to V_1, V_2, \dots, V_ℓ following the order of the linear layout. And then, we run the algorithm due to Proposition 5.3 to compute bases of the boundary spaces. When we prepare a branch-decomposition (T, \mathcal{L}) with bases of the boundary spaces, we run Algorithm 1 with $\{V_1, \dots, V_\ell\}$, k , and (T, \mathcal{L}) . If the computed \mathcal{F}_{root} is empty, then by Proposition 5.8, the path-width of \mathcal{V} is larger than k and therefore the path-width of the input \mathcal{V}_{in} is also larger than k . Otherwise we run Algorithm 2 with an element of \mathcal{F}_{root} having the minimum width to obtain an output linear layout σ_ℓ whose width is at most k . Until $\ell = n$, we repeat this process by increasing ℓ . Once ℓ is equal to n , if it is not terminated yet, then we obtain a linear layout σ_n of \mathcal{V} of width at most k , that can be converted to a linear layout of \mathcal{V}_{in} of width at most k by Lemmas 5.1 and 5.2. See Figure 4 for a flowchart of the algorithm using an iterative compression.

Here, we discuss the total running time. Given (T, \mathcal{L}) of $\{V_1, \dots, V_\ell\}$ whose width is at most $3k$, we first prepare bases of the boundary spaces by running the algorithm of Subsection 5.2. This takes $O(\theta rm |\mathcal{V}_{in}|) = O(r \cdot k^2 |\mathcal{V}_{in}|^2)$. Then, we run Algorithm 1 and Algorithm 2 to produce a linear layout of width at most k if one exists. The second step can be done in time $\text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{151\theta k} |\mathbb{F}|^{12\theta^2} \cdot n$ by Propositions 5.7 and 5.9. As we iterate this procedure for $\ell = 2, \dots, n$, the running time follows.

6.2. Using exact algorithms for partitioned matroids. One can relate a subspace arrangement \mathcal{V} with a partitioned matroid as introduced in Hliněný and Oum [22]. A *partitioned matroid* is a pair (M, \mathcal{P}) of a matroid M and a partition \mathcal{P} of $E(M)$; the precise definition of matroids will be introduced in Subsection 7.2. Hliněný and Oum [22] defined the branch-width of a partitioned matroid in order to devise an algorithm to construct a branch-decomposition of an \mathbb{F} -represented matroid of width at most k if it exists. Indeed, if (M, \mathcal{P}) is the partitioned matroid and M is \mathbb{F} -represented, then we can construct a subspace arrangement \mathcal{V} by taking the span of the vectors in P for each part $P \in \mathcal{P}$ as an element of \mathcal{V} . One can also convert \mathcal{V} into a partitioned matroid easily by taking a basis for each $V \in \mathcal{V}$. Thus the problem of computing branch-width of (M, \mathcal{P}) is equivalent to that of computing branch-width of \mathcal{V} . Thus, we can identify \mathbb{F} -represented partitioned matroids with subspace arrangements over \mathbb{F} .

Then an algorithm by Hliněný and Oum for partitioned matroids naturally translates for a subspace arrangement.

Theorem 6.2 (Hliněný and Oum [22]). *Let \mathbb{F} be a fixed finite field. There is, for some function f , an $O(f(\theta)m^3)$ -time algorithm that, for an input subspace arrangement $\mathcal{V} = \{V_i\}_{i=1}^n$ of subspaces of \mathbb{F}^r with $r \leq m = \sum_{i=1}^n \dim V_i$ and an integer θ , either*

- *outputs a branch-decomposition of \mathcal{V} of width at most θ*
- *or confirms the branch-width of \mathcal{V} is larger than θ .*

Proposition 2.2 states that if the path-width of \mathcal{V} is at most k , then the branch-width of \mathcal{V} is at most $2k$. Thus by applying Theorem 6.2 for $\theta = 2k$, we obtain a branch-decomposition of width at most $2k$ unless \mathcal{V} has path-width larger than k . By running the algorithm in Section 5 once with this branch-decomposition, we deduce an $O(n^3)$ -time algorithm to find a linear layout of width at most k if it exists for an input subspace arrangement $\mathcal{V} = \{V_i\}_{i=1}^n$ over a fixed finite field \mathbb{F} . (Note that after the preprocessing in




FIGURE 4. A flowchart for an iterative compression in Subsection 6.1.

Subsection 5.1, $m \leq \theta n = 2kn$ and therefore $O(m^3) = O(n^3)$.) This completes the proof of the following theorem.

Theorem 6.3. *Let \mathbb{F} be a fixed finite field. Given, as an input, n subspaces of \mathbb{F}^r for some r and a parameter k , in time $O(f(k)n^3 + rm^2)$ for some function f , we can either find a linear layout V_1, V_2, \dots, V_n of the subspaces such that*

$$\dim(V_1 + V_2 + \dots + V_i) \cap (V_{i+1} + V_{i+2} + \dots + V_n) \leq k$$

for all $i = 1, 2, \dots, n-1$, or confirm that no such linear layout exists, where each V_i is given by its spanning set of d_i vectors and $m = \sum_{i=1}^n d_i$.

This approach gives a better asymptotic running time in terms of n , but the hidden constant as a function of k is bigger here. The algorithm in Theorem 6.2 is based on the fact that forbidden minors for matroid branch-width at most k has at most $(6^{k+1} - 1)/5$ elements, proved by Geelen, Gerards, Robertson and Whittle [15].⁹ Though this bound allows the algorithm to search all forbidden \mathbb{F} -representable minors in “constant” time, it would take extremely long, as a function of k , to search all of them because there are too many \mathbb{F} -representable matroids to consider. Thus the approach in Subsection 6.1 would be a lot easier to be implemented and used, though it gives an algorithm that runs in time $O(n^4 + rm^2)$, worse than $O(n^3 + rm^2)$ of Theorem 6.3. We warn that the algorithm in Subsection 6.1 is still far from being practical because the upper bound on the size of a full set is gigantic and so the hidden constant is huge, but a lot smaller than that of Theorem 6.3.

7. APPLICATION TO PATH-WIDTH OF \mathbb{F} -REPRESENTED MATROIDS

So far our paper focused on path-width of a subspace arrangement and presented the algorithm in a self-contained manner. In this section we will discuss how to avoid using Theorem 6.2 of Hliněný and Oum by adapting a simpler algorithm of Hliněný while keeping the running time to be still $O(n^3)$ for the path-width of n -element matroids.

7.1. Path-width and branch-width of a connectivity function. We will first present the definition of path-width and branch-width of a connectivity function so that we can use in the next section. An integer-valued function f on the set 2^E of all subsets of a finite set E is a *connectivity function* on E if

- (1) $f(X) = f(E - X)$ for all $X \subseteq E$ (*symmetric*),
- (2) $f(X) + f(Y) \geq f(X \cap Y) + f(X \cup Y)$ for all $X, Y \subseteq E$ (*submodular*),
- (3) $f(\emptyset) = 0$.

Let f be a connectivity function on E . A *linear layout* of E is a permutation $\sigma = e_1, e_2, \dots, e_n$ of E . The *width* of a linear layout $\sigma = e_1, e_2, \dots, e_n$ of E with respect to f is $\max_{1 \leq i \leq n-1} f(\{e_1, e_2, \dots, e_i\})$. The *path-width* of f is the minimum width of all possible linear layouts of E . (If $|E| \leq 1$, then the path-width is 0.)

A *branch-decomposition* of E is a pair (T, \mathcal{L}) of a subcubic tree T and a bijection \mathcal{L} from all leaves of T to E . Every edge e of T induces a partition (A_e, B_e) of E given by components of $T - e$ and \mathcal{L} . The *width* of e is defined as $f(A_e)$. The *width* of a branch-decomposition (T, \mathcal{L}) is the maximum width of all edges of T . The *branch-width* of f is

⁹In [15], their definition of λ_M is one bigger than ours and therefore their bound $(6^k - 1)/5$ is translated into our bound $(6^{k+1} - 1)/5$.

the minimum width of all possible branch-decompositions of E . (If $|E| \leq 1$, then there is no branch-decomposition and we say f has branch-width 0.)

7.2. Matroids and vectors. A *matroid* M is a pair (E, \mathcal{I}) of a finite set E and a set \mathcal{I} of subsets of E satisfying the following.

- The empty set is in \mathcal{I} .
- If $X \in \mathcal{I}$, then every subset of X is contained in \mathcal{I} .
- If $X, Y \in \mathcal{I}$ and $|X| > |Y|$, then there exists an element $e \in X \setminus Y$ such that $Y \cup \{e\} \in \mathcal{I}$.

We write $E(M) := E$. For a subset X of E , the *rank* $r_M(X)$ of X is the maximum size of an independent subset of X . Let $\lambda_M(X) = r_M(X) + r_M(E - X) - r_M(E)$ for all subsets X of E . It is well known that λ_M is a connectivity function on E [39, Lemma 8.2.9]. The *path-width* of a matroid M is the path-width of λ_M and a *path-decomposition* of width k is a linear layout of $E(M)$ whose width with respect to λ_M is k . Similarly *branch-width* and *branch-decompositions* of a matroid M are branch-width and branch-decompositions, respectively, of λ_M .

For a matrix A over a field \mathbb{F} , let $E(A)$ be the set of all column vectors of A and $\mathcal{I}(A)$ be the set of all linearly independent subsets of $E(A)$. Then $M(A) = (E(A), \mathcal{I}(A))$ is a matroid, called the *vector matroid*. If a matroid M admits a matrix A over \mathbb{F} such that $M = M(A)$, then we say M is \mathbb{F} -*representable* or M is *representable* over \mathbb{F} and such a matrix A is called an \mathbb{F} -*representation* of M or a *matrix representation* of M over \mathbb{F} . A matroid is *binary* if it is representable over the field of two elements. When $M = M(A)$, then it is easy to observe

$$\lambda_M(X) = \dim(\langle X \rangle \cap \langle E(A) - X \rangle).$$

We remark that in general we cannot hope to find an \mathbb{F} -representation of a matroid efficiently, even if we assume such a representation exists [45, 50], unless $\mathbb{F} = GF(2)$. So we will need that a representation of a matroid, or equivalently a multiset of vectors, is given as an input. Let us say that a matroid is \mathbb{F} -*represented* if it is given with an \mathbb{F} -representation.

7.3. A better algorithm for matroid path-width. Recall that we give an algorithm that, for a parameter k and an input subspace arrangement \mathcal{V} with its branch-decomposition of small width, decides in time $O(n^3)$ whether its path-width is at most k and if so, outputs a path-decomposition of width at most k . In Subsection 6.2, we used Theorem 6.2 that gives an algorithm by Hliněný and Oum to obtain a branch-decomposition of bounded width. However, that algorithm by Hliněný and Oum [22] uses the huge but finite list of forbidden minors for the class of matroids of branch-width at most k and so in practice it would be too slow to implement it. We give an alternative in Subsection 6.1 yielding more direct and implementable algorithm but slower, only giving the time $O(n^4)$.

For the path-width of matroids, we can instead use the following $O(n^3)$ -time algorithm by Hliněný, which does not depend on the existence of finitely many forbidden minors, than the algorithm due to Theorem 6.2.

Theorem 7.1 (Hliněný [19]). *Let k be a fixed constant, \mathbb{F} be a finite field. Let \mathcal{V} be a subspace arrangement of n 1-dimensional subspaces of \mathbb{F}^r with $r \leq n$. There exists, for some function f , an $O(f(k)n^3)$ -time algorithm that, for a given \mathcal{V} , either*

- outputs a branch-decomposition of \mathcal{V} of width at most $3k$
- or confirms the branch-width of \mathcal{V} is larger than k .

We associate a matroid M represented by vectors v_1, v_2, \dots, v_n in a vector space over a finite field \mathbb{F} with a subspace arrangement $\mathcal{V} = \{\langle v_1 \rangle, \langle v_2 \rangle, \dots, \langle v_n \rangle\}$. Then a linear layout of \mathcal{V} of width k is precisely a path-decomposition of M having width k . This relation allows us to deduce the following. In addition, we do not need to apply the preprocessing step discussed in Subsection 5.1, as long as we remove all *loops* and *coloops*¹⁰ from M .

Theorem 7.2. *Let \mathbb{F} be a fixed finite field. There is an algorithm that, for an input n -element matroid given by its matrix representation over \mathbb{F} having at most n rows and a parameter k , decides in time $O(f(k)n^3)$ for some function f whether its path-width is at most k and if so, outputs a path-decomposition of width at most k .*

8. APPLICATION TO LINEAR RANK-WIDTH AND LINEAR CLIQUE-WIDTH OF GRAPHS

As it is discussed in Section 1, our full theorem on subspace arrangements provides a nice corollary to linear rank-width of graphs when applied to subspaces of dimension at most two. As linear rank-width is closely related to linear clique-width, this will also give a fixed-parameter tractable approximation algorithm for linear clique-width as well. We will review the definitions and discuss how to obtain the desired results from Theorem 6.3.

8.1. Linear rank-width and rank-width of graphs. Let $G = (V, E)$ be a graph and A_G be its adjacency matrix, which is the $V(G) \times V(G)$ matrix over the binary field whose (i, j) -entry is 1 if and only if i and j are adjacent in G . For an $X \times Y$ matrix $M = (m_{ij})_{i \in X, j \in Y}$ and subsets $A \subseteq X$ and $B \subseteq Y$, let $M[A, B]$ be the $A \times B$ submatrix $(m_{ij})_{i \in A, j \in B}$ of M . The *cut-rank function* of G , denoted by ρ_G , is defined to be

$$\rho_G(S) = \text{rank}(A_G[S, V(G) \setminus S])$$

for all subsets S of $V(G)$. Oum and Seymour [38] showed that the cut-rank function is a connectivity function on $V(G)$ (which is defined in Subsection 7.1). In Subsection 7.1, we discussed how to define path-width and branch-width of a connectivity function. We define *linear rank-width* and *rank-width* of a graph G to be path-width and branch-width, respectively, of the cut-rank function ρ_G . For a graph G , a *linear rank-decomposition* of width k is a linear layout of $V(G)$ whose width with respect to ρ_G is k . For a graph G , a *rank-decomposition* of width k is a branch-decomposition of ρ_G having width k .

8.2. From linear rank-width of graphs to path-width of subspace arrangements.

We now explain how to relate the linear rank-width of a graph with the path-width of a subspace arrangement.

Let G be a graph with $V(G) = \{1, 2, \dots, n\}$. Let e_1, e_2, \dots, e_n be the standard basis of $GF(2)^n$. Each vertex i is associated with a vector $v_i \in GF(2)^n$ such that

$$v_i = \sum_{j \text{ is adjacent to } i \text{ in } G} e_j.$$

Let $V_i = \langle e_i, v_i \rangle$ for each i . Let $\mathcal{V}_G = \{V_1, V_2, \dots, V_n\}$ be the subspace arrangement associated with G . The following lemma is equivalent to [22, Lemma 7.1] and easily implies that the path-width of \mathcal{V}_G is precisely twice the linear rank-width of G .

¹⁰In \mathbb{F} -representable matroids, a loop corresponds to a zero vector and a coloop corresponds to a vector that is not spanned by the other vectors.

Lemma 8.1. *For $X \subseteq V(G)$, $\dim \left(\sum_{i \in X} V_i \right) \cap \left(\sum_{j \in V(G) - X} V_j \right) = 2\rho_G(X)$.*

Proof. Observe that $\dim \left(\sum_{i \in X} V_i \right) \cap \left(\sum_{j \in V(G) - X} V_j \right) = \dim \sum_{i \in X} V_i + \dim \sum_{i \in V(G) - X} V_i - n$. Since $\dim \sum_{i \in X} V_i = |X| + \text{rank}(A_G[V(G) - X, X])$ and $\dim \sum_{i \in V(G) - X} V_i = n - |X| + \text{rank}(A_G[X, V(G) - X])$, the conclusion follows. \square

Therefore, a linear rank-decomposition of width k of a graph G precisely corresponds to a linear layout of \mathcal{V}_G of width $2k$ and a rank-decomposition of width k of a graph G is exactly a branch-decomposition of width $2k$ of \mathcal{V}_G .

8.3. An algorithm for linear rank-width. In order to run the algorithm in Theorem 6.3, we need a branch-decomposition of small width to be given. Instead of using methods in Section 6, we may use the following algorithm by Oum [35] for rank-width.

Theorem 8.2 (Oum [35]). *For fixed k , there is an $O(f(k)n^3)$ -algorithm that, for an input n -vertex graph G , either obtains a rank-decomposition of G of width at most $3k - 1$ or confirms that the rank-width of G is larger than k .*

We will describe an algorithm that finds a linear rank-decomposition of width at most k of an input graph if it exists. For an input graph G with $V(G) = \{1, 2, \dots, n\}$, we first run the algorithm due to Theorem 8.2 in time $O(f(k)n^3)$. If the algorithm confirms that the rank-width of G is larger than k , then we deduce that the linear rank-width of G is also larger than k . Otherwise, we obtain a rank-decomposition (T, \mathcal{L}) of G of width at most $3k - 1$. Let $\mathcal{V}_G = \{V_1, V_2, \dots, V_n\}$ be the subspace arrangement associated with G , which is defined in Subsection 8.2. Let \mathcal{L}' be a bijection from the set of all leaves of T to \mathcal{V}_G such that $\mathcal{L}'(v) = V_{\mathcal{L}(v)}$ for every leaf v of T . Then (T, \mathcal{L}') is a branch-decomposition of a subspace arrangement \mathcal{V}_G . By Lemma 8.1, the width of (T, \mathcal{L}') is at most $6k - 2$. We then apply the algorithm proven in Theorem 5.10 with, as an input, a subspace arrangement \mathcal{V}_G , a parameter $2k$, a branch-decomposition (T, \mathcal{L}') of width at most $6k - 2$. If \mathcal{F}_{root} is empty, then by Proposition 5.8 and Lemma 8.1, we can conclude that the linear rank-width of G is larger than k . Otherwise, we run Algorithm 2 by choosing an element with the minimum width in \mathcal{F}_{root} . Note that once we obtain a linear layout of \mathcal{V}_G whose path-width at most $2k$, then it is easy to convert a linear rank-decomposition of G of width at most k . We remark that there is no need to apply the preprocessing discussed in Subsection 5.1 since the dimension of V_i is at most 2 for each i . See Figure 5 for a flowchart of this algorithm.

Theorem 8.3. *For an input n -vertex graph and a parameter k , we can decide in time $O(f(k)n^3)$ for some function f whether its linear rank-width is at most k and if so, find a linear rank-decomposition of width at most k .*

8.4. Application to linear clique-width. For an integer k , a k -expression is an algebraic expression that consists of the following operations on graphs whose vertices are labeled by integers in $\{1, 2, \dots, k\}$:

- \cdot_i (a graph with a single vertex labeled by $i \in \{1, 2, \dots, k\}$).
- $G_1 \oplus G_2$ (the disjoint union of two vertex-labeled graphs G_1 and G_2).
- $\eta_{i,j}(G)$ with $i \neq j$ (adding an edge from each pair of a vertex of label i and a vertex of label j).
- $\rho_{i \rightarrow j}(G)$ (relabeling all vertices of label i to label j).

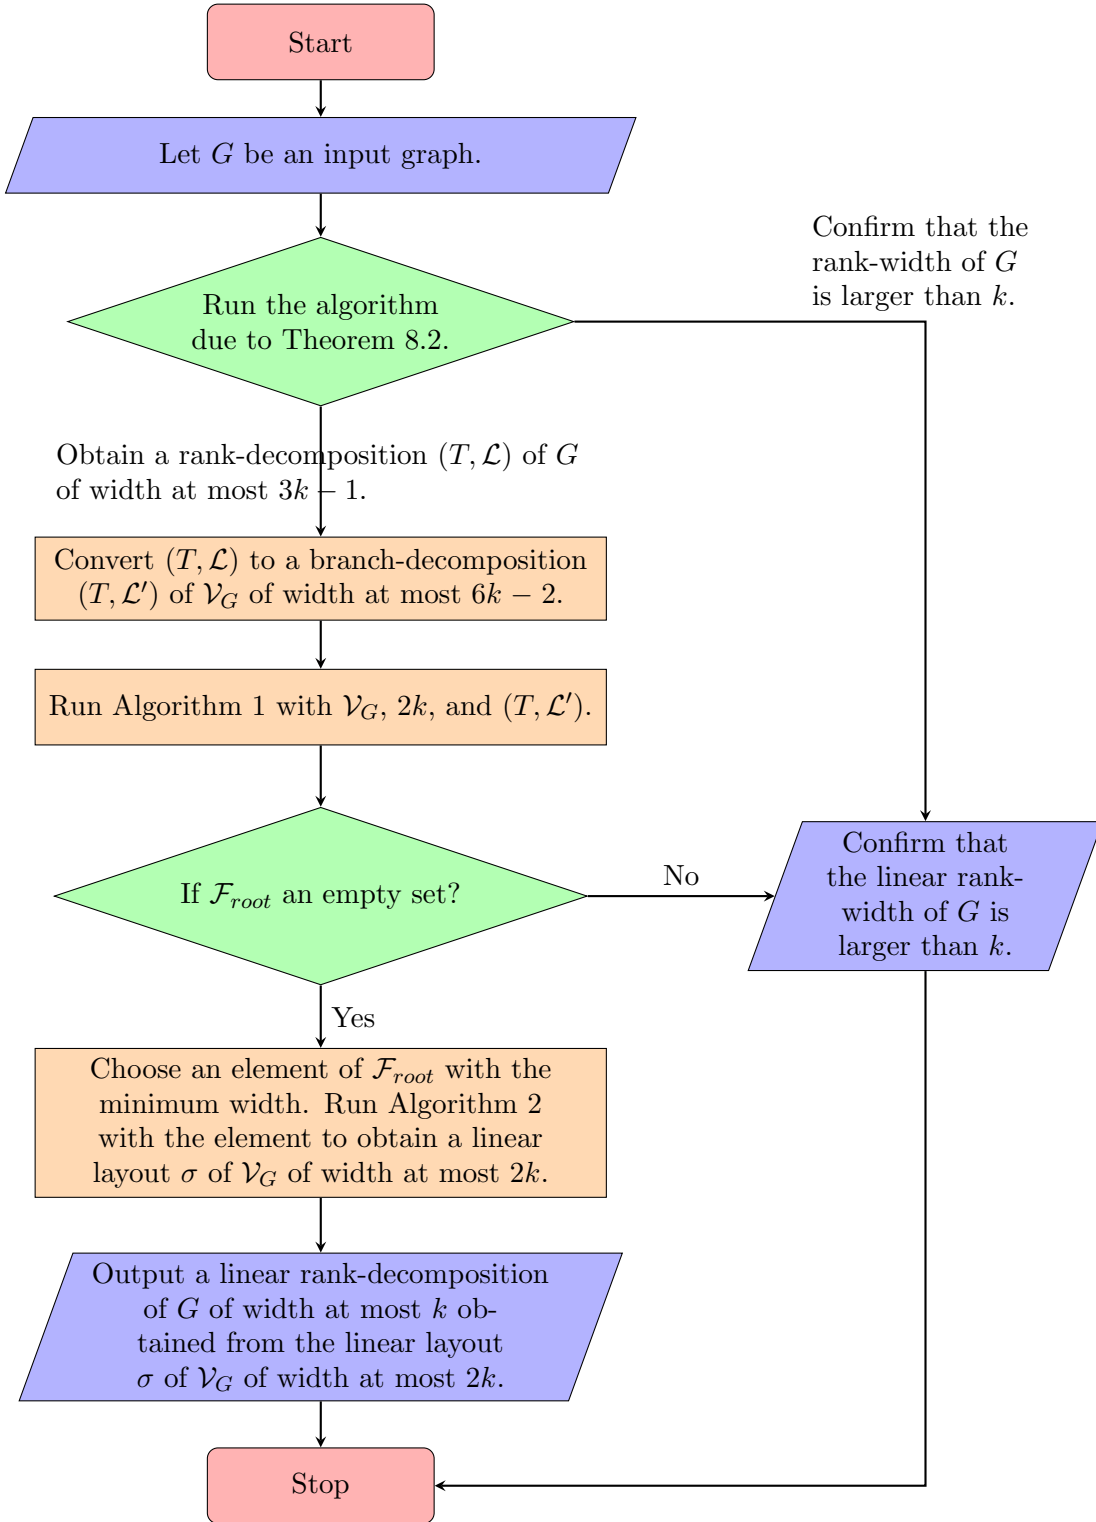


FIGURE 5. A flowchart of the algorithm due to Theorem 8.3.

A k -expression is *linear* if one operand of each disjoint union \oplus has at most one vertex. The *clique-width* of a graph G is the minimum k such that there is a k -expression representing G after ignoring labels of the vertices [9]. Similarly the *linear clique-width* of a graph G is defined as the minimum k such that there exists a linear k -expression representing G after ignoring labels of the vertices.

By modifying the proof of Oum and Seymour [37, Proposition 6.3], it is straightforward to prove the following lemma. We omit its easy proof.

Lemma 8.4. *If a graph G has linear rank-width k , then the linear clique-width of G is at least k and at most $2^k + 1$. Furthermore, for each fixed k , a linear rank-decomposition of G of width k can be converted to a linear $(2^k + 1)$ -expression representing G in linear time.*

Now we provide an algorithm that makes an approximate linear expression of an input graph, which was unknown. For an input n -vertex graph G , we run the algorithm due to Theorem 8.2. If the linear rank-width of G is larger than k , then by Lemma 8.4, the linear clique-width of G is also larger than k . If the linear rank-width of G is at most k and thus the algorithm outputs a linear rank-decomposition of width at most k , then by Lemma 8.4, we obtain a linear $(2^k + 1)$ -expression of G .

Corollary 8.5. *For an input n -vertex graph G and a parameter k , we can find a linear $(2^k + 1)$ -expression of G confirming that G has linear clique-width at most $2^k + 1$ or certify that G has linear clique-width larger than k in time $O(f(k)n^3)$ for some function f .*

9. COMPUTING PATH-WIDTH EXACTLY WHEN THE INPUT HAS BOUNDED BRANCH-WIDTH

In this section, we present, for each fixed integer θ , a polynomial-time algorithm to compute path-width exactly with an optimal linear layout when the input subspace arrangement has branch-width at most θ . Our idea is similar to the idea of Bodlaender and Kloks [5] who showed a polynomial-time algorithm to compute path-width of graphs of tree-width at most θ .

The running time of this algorithm is polynomial because our algorithm in Theorem 6.3 has the running time $O(c^k n^3)$ for some c depending only on the underlying field \mathbb{F} and θ . Therefore if we can make sure that $k \leq \theta \lfloor \log_2 n \rfloor$, then we can run the algorithm with $k = \theta \lfloor \log_2 n \rfloor$ and deduce the exact path-width from the full set.

The following proposition shows why the path-width is at most $\theta \lfloor \log_2 n \rfloor$. Its proof follows proofs of a theorem of Kwon [30] (also in Adler, Kanté, and Kwon [1]) on an inequality between rank-width and linear rank-width of graphs and a theorem of Bodlaender, Gilbert, Hafsteinsson, and Kloks [4] on an inequality between tree-width and path-width of graphs.

Proposition 9.1. *Let $\mathcal{V} = \{V_1, V_2, \dots, V_n\}$ be a subspace arrangement. If the branch-width of \mathcal{V} is k , then the path-width of \mathcal{V} is at most $k \lfloor \log_2 n \rfloor$.*

Proof. Let (T, \mathcal{L}) be a branch-decomposition of width k . By picking an edge e of T and subdividing e to add a root r of degree 2, we may make T a rooted binary tree. For each internal node v of T , let A_v be the subset of \mathcal{V} that corresponds to all leaves below the left child of v by \mathcal{L} and let B_v be the subset of \mathcal{V} that corresponds to all leaves below the right child of v by \mathcal{L} . We may assume that $|A_v| \geq |B_v|$ for all internal nodes v of T .

By permuting labels, we may assume that for all $1 \leq i < j \leq n$, if v is the common ancestor of V_i and V_j with the maximum distance from r , then V_i is a descendant of the left child of v and V_j is a descendant of the right child of v . We claim that this linear layout V_1, V_2, \dots, V_n has width at most $k \lfloor \log_2 n \rfloor$.

To see this, we need to prove that

$$\dim(V_1 + \dots + V_{i-1}) \cap (V_i + V_{i+1} + \dots + V_n) \leq k \lfloor \log_2 n \rfloor$$

for each $i \in \{2, 3, \dots, n\}$. Let P be the unique path $w_1 = r, w_2, \dots, w_t$ in T from r to a leaf w_t with $\mathcal{L}(w_t) = V_i$.

Now observe that $|A_{w_j}| \geq |B_{w_j}|$ and therefore $|B_{w_j}| \leq \frac{1}{2}(|A_{w_j}| + |B_{w_j}|)$. Thus if w_{j+1} is the right child of w_j , then

$$|A_{w_{j+1}}| + |B_{w_{j+1}}| = |B_{w_j}| \leq \frac{|A_{w_j}| + |B_{w_j}|}{2}.$$

If w_{j+1} is the left child of w_j , then trivially

$$|A_{w_{j+1}}| + |B_{w_{j+1}}| = |A_{w_j}| < |A_{w_j}| + |B_{w_j}|.$$

Note that $|A_{w_{t-1}}| + |B_{w_{t-1}}| \geq 1$ because $V_i \in A_{w_{t-1}} \cup B_{w_{t-1}}$.

Let $P = \{j \in \{1, 2, \dots, t-1\} : w_{j+1} \text{ is the right child of } w_j\}$. Then

$$1 \leq |A_{w_{t-1}}| + |B_{w_{t-1}}| \leq \frac{1}{2^{|P|}}(|A_1| + |B_1|) = \frac{1}{2^{|P|}}n.$$

This implies that $|P| \leq \lfloor \log_2 n \rfloor$. Let $X = \bigcup_{j \in P} A_{w_j}$. Then by the definition of P , $X = \{V_1, V_2, \dots, V_{i-1}\}$. Since $|P| \leq \lfloor \log_2 n \rfloor$ and $\dim \langle A_{w_j} \rangle \cap \langle \mathcal{V} - A_{w_j} \rangle \leq k$, by Proposition 2.1, we deduce that $\dim \langle X \rangle \cap \langle \mathcal{V} - X \rangle \leq |P|k \leq k \lfloor \log_2 n \rfloor$. This completes the proof. \square

By using this inequality, we can prove the following main theorem.

Theorem 9.2. *Let θ be a constant and \mathbb{F} be a fixed finite field. We can compute the path-width of a given subspace arrangement \mathcal{V} over \mathbb{F} if the branch-width of \mathcal{V} is at most θ in polynomial time. Furthermore, we can find an optimal linear layout of \mathcal{V} .*

Proof. Let $\mathcal{V} = \{V_1, V_2, \dots, V_n\}$. Since the branch-width of \mathcal{V} is at most θ , using an algorithm by Hliněný and Oum in Theorem 6.2, we obtain a branch-decomposition of \mathcal{V} of width at most θ in polynomial time.

We will run an algorithm in Theorem 5.10 for $k = \theta \lfloor \log_2 n \rfloor$ whose time complexity is $O(m(m + \theta n) + \text{poly}(\theta, |\mathbb{F}|, k) \cdot 2^{151\theta k} |\mathbb{F}|^{12\theta^2} \cdot n)$. This is a polynomial in the size of the input.

By Proposition 9.1, the path-width of \mathcal{V} is at most $\theta \lfloor \log_2 n \rfloor$ and therefore the full set $\text{FS}_k(\mathcal{V}, \{0\})$ is non-empty. Find a $\{0\}$ -trajectory $\Gamma \in \text{FS}_k(\mathcal{V}, \{0\})$ with minimum width w . Then w is the path-width of \mathcal{V} and we can find a corresponding linear layout by using the algorithm in Algorithm 2. \square

9.1. Path-width of matroids having bounded branch-width. Computing the path-width of a given \mathbb{F} -represented matroid is NP-hard, shown by Kashyap [28]. However, if we restrict our input matroid to a small class, then it may be possible to compute path-width in polynomial time. For instance, Koutsonas, Thilikos, and Yamazaki [29] presented a polynomial-time algorithm to compute path-width exactly for the cycle matroids of outerplanar graphs. We remark that all cycle matroids of outerplanar graphs have very small branch-width.

Theorem 9.2 implies that there exists a polynomial-time algorithm to compute path-width with an optimal linear layout if the input matroid has bounded branch-width.

Corollary 9.3. *Let θ be a constant and \mathbb{F} be a fixed finite field. We can compute the path-width of an \mathbb{F} -represented matroid M having branch-width at most θ in polynomial time. Furthermore, we can find an optimal linear layout of M .*

9.2. Linear rank-width of graphs having bounded rank-width. Computing the linear rank-width of a given graph is also NP-hard by the result of Kashyap [28] and Oum [34]. Adler, Kanté, and Kwon [1] proved that the linear rank-width of graphs of rank-width at most 1 can be computed in polynomial time. They also asked the following question:

Can we compute in polynomial time the linear rank-width of graphs of rank-width at most k , for fixed $k \geq 2$?

We answer this question in the affirmative by deducing Corollary 9.4 from Theorem 9.2 by using the method in Section 8.

Corollary 9.4. *Let θ be a constant. We can compute the linear rank-width of a graph G having rank-width at most θ in polynomial time. Furthermore, we can find an optimal linear rank-decomposition of G .*

Proof. Let G be a graph of rank-width at most θ with $V(G) = \{1, 2, \dots, n\}$. Let $\mathcal{V}_G = \{V_1, V_2, \dots, V_n\}$ be the subspace arrangement associated with G , which is defined in Subsection 8.2. By Theorem 9.2, we can obtain a linear layout $V_{\sigma(1)}, V_{\sigma(2)}, \dots, V_{\sigma(n)}$ of \mathcal{V}_G having optimal path-width. Then by Lemma 8.1, a linear rank-decomposition $\sigma(1), \sigma(2), \dots, \sigma(n)$ of G has optimal linear rank-width. \square

10. DISCUSSIONS

Can we decide whether the path-width of an input matroid is at most k for a fixed k when the input matroid is general? Our algorithm only works if the input matroid is \mathbb{F} -represented. First we have to clarify our question. Binary strings of length $\text{poly}(n)$ cannot represent all n -element matroids, because there are too many n -element matroids [40]. To study algorithmic problems on general matroids, typically we assume that a matroid is given by an independence oracle, which answers whether a given set is independent in the matroid in the constant time.

The algorithm by Nagamochi [33] can still decide in time $O(n^{ck})$ for some c whether the path-width of an input matroid is at most k , even if the input matroid is given by its independence oracle. However, the exponent of n depends on k and so his algorithm is not fixed-parameter tractable when k is a parameter.

Question 1. *Does there exist, for each fixed k , an algorithm that runs in time $O(n^c)$ for an input n -element matroid given with its independence oracle whether its path-width is at most k for some c ?*

As a weaker question, we ask the following. A matroid is *binary* if it is representable over the field of two elements.

Question 2. *Does there exist, for each fixed k , an algorithm that runs in time $O(n^c)$ for an input n -element matroid given with its independence oracle and confirms one of the following three outcomes for some c ?*

- The path-width is at most k .
- The path-width is larger than k .
- The input matroid is not binary.

One may attempt to construct a binary representation of the input matroid and then apply Theorem 7.2. However, no polynomial-time algorithm can verify whether an input matroid given by its independence oracle is binary [44] and so that strategy is not going to work.

By using the result of this paper, one can make a slightly weaker algorithm as follows.

Proposition 10.1. *For each fixed k , there exists an algorithm that runs in time $O(n^3)$ for an input n -element matroid given by an independence oracle and confirms one of the following.*

- Its path-width is at most k .
- Either its path-width is larger than k or the matroid is not binary.

Proof. Let M be the input matroid. By assuming the input matroid is binary, we can construct a potential binary representation, representing a binary matroid M' . This representation can be constructed in time $O(n^2)$ by picking a base and finding fundamental circuits for each element outside of the chosen base. The problem is that we do not yet know whether $M = M'$. But if M is binary, then $M = M'$.

We apply our algorithm in Theorem 7.2 with M' to decide whether the path-width of M' is at most k . If the path-width of M' is at most k , then we also obtain a linear layout of M' having width at most k . It is now easy to verify whether this linear layout also has width at most k in M and if not, then $M \neq M'$ and so M is not binary.

If our algorithm confirms that M' has path-width larger than k , then either M has path-width larger than k or M is non-binary. \square

Lastly, it is natural to ask whether the running time of our algorithm in Theorem 7.2 can be improved. We assume that the input matroid is given as an $n \times n$ matrix over \mathbb{F} .

Question 3. *Does there exist an algorithm that, given an n -element \mathbb{F} -represented matroid and a nonnegative integer $k \geq 2$, decides whether its path-width is at most k in time $O(n^{3-\epsilon})$ for some $\epsilon > 0$?*

It can be readily verified that the preprocessing step of Section 5.1 can be implemented with $O(n^\omega)$ field operations using fast matrix multiplication [7], where $\omega < 2.38$. Recall that the algorithm of Theorem 7.2 uses the approximation algorithm by Hliněný (see Theorem 7.1, also [19]) for computing branch-decomposition. The latter algorithm runs in time $O(n^3)$ and it appears that there is an inherent bottleneck in the approach of [19] for obtaining a truly subcubic algorithm.

For $k = 0$ and 1, the path-width of a \mathbb{F} -represented matroid M can be computed in $O(n^\omega)$ -time. When $k = 0$, the path-width of M equals 0 if and only if every element of M is either a loop or a coloop. This is true if and only if the number of non-zero columns in the matrix representing M equals the rank of M , which can be checked in $O(n^\omega)$ -time.

We consider the case $k = 1$. For a fixed basis of M , we construct a potential binary representation M' as in Proposition 10.1 in time $O(n^\omega)$. Notice that M' is isomorphic to M , written as $M' \cong M$, if M is binary. It is known that the path-width of M' equals the linear rank-width of the fundamental graph $G(M')$ of M' , whose vertices correspond to the columns of M' , see Proposition 3.1 in [34]. One can test whether $G(M')$ has linear

rank-width at most one in time $O(n + m)$, where n and m are the number of vertices and edges of $G(M')$ respectively [27]. We argue the following.

- If $G(M')$ has linear rank-width at least two, then M has path-width at least two regardless of whether $M \cong M'$ or not. This is because M is a binary matroid if and only if M does not contain the uniform matroid $U_{2,4}$ as a minor. As $U_{2,4}$ has path-width 2, $M \not\cong M'$ implies that M contains $U_{2,4}$ as a minor, thus the path-width of M is at least two. In case $M \cong M'$, the path-width of M equals the path-width of M' , which is in turn equals the linear rank-width of $G(M')$.
- If $G(M')$ has linear rank-width at most one, then consider the corresponding layout σ of M . If the width of σ on M is at most one, we are done. If not, observe that $M \not\cong M'$; indeed $M \cong M'$ implies that an identical layout on M and M' results in the same width on both M and M' . Especially, M is non-binary, and thus has path-width at least two.

REFERENCES

- [1] I. Adler, M. M. Kanté, and O.-j. Kwon. Linear rank-width of distance-hereditary graphs I. A polynomial-time algorithm. *Algorithmica*, pages 1–36, 2016.
- [2] P. Berthomé, T. Bouvier, F. Mazoit, N. Nisse, and R. P. Soares. An unified FPT algorithm for width of partition functions. [Technical Report] RR-8372, Unpublished, <https://hal.inria.fr/hal-00865575v2/document>, 2013.
- [3] H. L. Bodlaender, M. R. Fellows, and D. M. Thilikos. Derivation of algorithms for cutwidth and related graph layout parameters. *J. Comput. System Sci.*, 75(4):231–244, 2009.
- [4] H. L. Bodlaender, J. R. Gilbert, H. Hafsteinsson, and T. Kloks. Approximating treewidth, pathwidth, frontsize, and shortest elimination tree. *J. Algorithms*, 18(2):238–255, 1995.
- [5] H. L. Bodlaender and T. Kloks. Efficient and constructive algorithms for the pathwidth and treewidth of graphs. *J. Algorithms*, 21(2):358–402, 1996.
- [6] H. L. Bodlaender and D. M. Thilikos. Computing small search numbers in linear time. In R. Downey, M. Fellows, and F. Dehne, editors, *Parameterized and Exact Computation. Proc. First International Workshop, IWPEC 2004, Bergen, Norway, September 14-17, 2004.*, volume 3162 of *Lecture Notes in Comput. Sci.*, pages 37–48. Springer, 2004.
- [7] H. Y. Cheung, T. C. Kwok, and L. C. Lau. Fast matrix rank algorithms and applications. *J. ACM*, 60(5):31:1–31:25, 2013.
- [8] B. Courcelle. The monadic second-order logic of graphs. i. recognizable sets of finite graphs. *Information and Computation*, 85(1):12 – 75, 1990.
- [9] B. Courcelle and S. Olariu. Upper bounds to the clique width of graphs. *Discrete Appl. Math.*, 101(1-3):77–114, 2000.
- [10] B. Courcelle and S. Oum. Vertex-minors, monadic second-order logic, and a conjecture by Seese. *J. Combin. Theory Ser. B*, 97(1):91–126, 2007.
- [11] J. A. Ellis, I. H. Sudborough, and J. S. Turner. Graph separation and search number. In *Twenty-first annual Allerton conference on communication, control, and computing (Monticello, Ill., 1983)*, pages 224–233. Univ. Illinois, Urbana, IL, 1983.
- [12] M. R. Fellows and M. A. Langston. On search, decision, and the efficiency of polynomial-time algorithms. *J. Comput. System Sci.*, 49(3):769–779, 1994.
- [13] M. R. Fellows, F. A. Rosamond, U. Rotics, and S. Szeider. Clique-width is NP-complete. *SIAM J. Discrete Math.*, 23(2):909–939, 2009.
- [14] J. Geelen, B. Gerards, and G. Whittle. Solving Rota’s conjecture. *Notices Amer. Math. Soc.*, 61(7):736–743, 2014.
- [15] J. F. Geelen, A. M. H. Gerards, N. Robertson, and G. Whittle. On the excluded minors for the matroids of branch-width k . *J. Combin. Theory Ser. B*, 88(2):261–265, 2003.
- [16] J. F. Geelen, A. M. H. Gerards, and G. Whittle. Branch-width and well-quasi-ordering in matroids and graphs. *J. Combin. Theory Ser. B*, 84(2):270–290, 2002.

- [17] J. F. Geelen, A. M. H. Gerards, and G. Whittle. A correction to our paper “Branch-width and well-quasi-ordering in matroids and graphs”. Manuscript, 2006.
- [18] R. Hall, J. Oxley, and C. Semple. The structure of 3-connected matroids of path width three. *European J. Combin.*, 28(3):964–989, 2007.
- [19] P. Hliněný. A parametrized algorithm for matroid branch-width. *SIAM J. Comput.*, 35(2):259–277, loose erratum (electronic), 2005.
- [20] P. Hliněný. Branch-width, parse trees, and monadic second-order logic for matroids. *J. Combin. Theory Ser. B*, 96(3):325–351, 2006.
- [21] P. Hliněný. Simpler self-reduction algorithm for matroid path-width. Submitted. arXiv:1605.09520, 2016.
- [22] P. Hliněný and S. Oum. Finding branch-decompositions and rank-decompositions. *SIAM J. Comput.*, 38(3):1012–1032, 2008.
- [23] G. B. Horn and F. R. Kschischang. On the intractability of permuting a block code to minimize trellis complexity. *IEEE Trans. Inform. Theory*, 42(6, part 1):2042–2048, 1996. Codes and complexity.
- [24] K. Jain, I. Măndoiu, and V. V. Vazirani. The “art of trellis decoding” is computationally hard—for large fields. *IEEE Trans. Inform. Theory*, 44(3):1211–1214, 1998.
- [25] J. Jeong, E. J. Kim, and S. Oum. Constructive algorithm for path-width of matroids. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016)*, pages 1695–1704, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics.
- [26] J. Jeong, O.-j. Kwon, and S. Oum. Excluded vertex-minors for graphs of linear rank-width at most k . *European J. Combin.*, 41:242–257, 2014.
- [27] M. M. Kanté, E. J. Kim, O. Kwon, and C. Paul. FPT algorithm and polynomial kernel for linear rank-width one vertex deletion. Submitted. arXiv:1504.05905, 2015.
- [28] N. Kashyap. Matroid pathwidth and code trellis complexity. *SIAM J. Discrete Math.*, 22(1):256–272, 2008.
- [29] A. Koutsonas, D. M. Thilikos, and K. Yamazaki. Outerplanar obstructions for matroid pathwidth. *Discrete Math.*, 315–316:95–101, 2014.
- [30] O.-j. Kwon. *On the structural and algorithmic properties of linear rank-width*. PhD thesis, KAIST, 2015.
- [31] A. Maheshwari and N. Zeh. I/O-efficient algorithms for graphs of bounded treewidth. *Algorithmica*, 54(3):413–469, 2009.
- [32] J. L. Massey. Foundation and methods of channel encoding. In *Proc. Int. Conf. Information Theory and Systems*, volume 65, pages 148–157, 1978.
- [33] H. Nagamochi. Linear layouts in submodular systems. In K.-M. Chao, T.-S. Hsu, and D.-T. Lee, editors, *ISAAC '12*, volume 7676 of *Lecture Notes in Comput. Sci.*, pages 475–484. Springer Berlin Heidelberg, 2012.
- [34] S. Oum. Rank-width and vertex-minors. *J. Combin. Theory Ser. B*, 95(1):79–100, 2005.
- [35] S. Oum. Approximating rank-width and clique-width quickly. *ACM Trans. Algorithms*, 5(1):Art. 10, 20, 2008.
- [36] S. Oum. Rank-width and well-quasi-ordering. *SIAM J. Discrete Math.*, 22(2):666–682, 2008.
- [37] S. Oum and P. Seymour. Approximating clique-width and branch-width. *J. Combin. Theory Ser. B*, 96(4):514–528, 2006.
- [38] S. Oum and P. Seymour. Testing branch-width. *J. Combin. Theory Ser. B*, 97(3):385–393, 2007.
- [39] J. Oxley. *Matroid theory*, volume 21 of *Oxford Graduate Texts in Mathematics*. Oxford University Press, Oxford, second edition, 2011.
- [40] M. J. Piff and D. J. A. Welsh. The number of combinatorial geometries. *Bull. London Math. Soc.*, 3:55–56, 1971.
- [41] N. Robertson and P. Seymour. Graph minors. I. Excluding a forest. *J. Combin. Theory Ser. B*, 35(1):39–61, 1983.
- [42] N. Robertson and P. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B*, 52(2):153–190, 1991.
- [43] N. Robertson and P. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory Ser. B*, 63(1):65–110, 1995.
- [44] P. D. Seymour. Recognizing graphic matroids. *Combinatorica*, 1(1):75–78, 1981.

- [45] P. D. Seymour and P. N. Walton. Detecting matroid minors. *J. London Math. Soc. (2)*, 23(2):193–203, 1981.
- [46] R. Soares. *Pursuit-Evasion, decompositions and convexity on graphs*. PhD thesis, Université de Nice - Sophia-Antipolis, 2013.
- [47] D. M. Thilikos. Algorithms and obstructions for linear-width and related search parameters. *Discrete Appl. Math.*, 105(1-3):239–271, 2000.
- [48] D. M. Thilikos, M. Serna, and H. L. Bodlaender. Cutwidth. I. A linear time fixed parameter algorithm. *J. Algorithms*, 56(1):1–24, 2005.
- [49] D. M. Thilikos, M. Serna, and H. L. Bodlaender. Cutwidth. II. Algorithms for partial w -trees of bounded degree. *J. Algorithms*, 56(1):25–49, 2005.
- [50] K. Truemper. On the efficiency of representability tests for matroids. *European J. Combin.*, 3(3):275–291, 1982.
- [51] R. van Bevern, R. G. Downey, M. R. Fellows, S. Gaspers, and F. A. Rosamond. How applying Myhill-Nerode methods to hypergraphs helps mastering the art of trellis decoding. arXiv:1211.1299v1, 2012.
- [52] R. van Bevern, R. G. Downey, M. R. Fellows, S. Gaspers, and F. A. Rosamond. Myhill-Nerode methods for hypergraphs. *Algorithmica*, 73(4):696–729, 2015.
- [53] A. Vardy. Trellis structure of codes. In *Handbook of coding theory, Vol. I, II*, pages 1989–2117. North-Holland, Amsterdam, 1998.
- [54] A. J. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans. Information Theory*, 13(2):260–269, 1967.

(Kim) CNRS, LAMSADE, PLACE DU MARECHAL DE LATRE DE TASSIGNY, 75775 PARIS CEDEX 16, FRANCE

(Jeong, Oum) DEPARTMENT OF MATHEMATICAL SCIENCES, KAIST, 291 DAEHAK-RO YUSEONG-GU DAEJEON, 34141 SOUTH KOREA

E-mail address: `jjisu@kaist.ac.kr`

E-mail address: `eunjungkim78@gmail.com`

E-mail address: `sangil@kaist.edu`