Counting Vanishing Matrix-Vector Products*

Cornelius Brand
Algorithms and Complexity Group, TU Wien, Austria

Viktoriia Korchemna
Algorithms and Complexity Group, TU Wien, Austria

Michael Skotnica
Department of Applied Mathematics, Charles University, Prague, Czech Republic

Kirill Simonov
Hasso Plattner Institute, University of Potsdam, Germany

Abstract

Consider the following parameterized counting variation of the classic subset sum problem, which arises notably in the context of higher homotopy groups of topological spaces: Let $\mathbf{v} \in \mathbb{Q}^d$ be a rational vector, $(T_1, T_2 \dots T_m)$ a list of $d \times d$ rational matrices, $S \in \mathbb{Q}^{h \times d}$ a rational matrix not necessarily square and k a parameter. The goal is to compute the number of ways one can choose k matrices $T_{i_1}, T_{i_2}, \dots, T_{i_k}$ from the list such that $ST_{i_k} \cdots T_{i_1}\mathbf{v} = \mathbf{0} \in \mathbb{Q}^h$.

In this paper, we show that this problem is #W[2]-hard for parameter k. As a consequence, computing the k-th homotopy group of a d-dimensional 1-connected topological space for $d > 3$ is #W[2]-hard for parameter k. We also discuss a decision version of the problem and its several modifications for which we show W[1]/W[2]-hardness. This is in contrast to the parameterized k-sum problem, which is only W[1]-hard (Abboud-Lewi-Williams, ESA'14). In addition, we show that the decision version of the problem without parameter is an undecidable problem, and we give a fixed-parameter tractable algorithm for matrices of bounded size over finite fields, parameterized the matrix dimensions and the order of the field.

1 Introduction

Topology is one of the most important and active areas of mathematics, emerging from vast generalizations of geometry (see, e.g., [12] for a gentle introduction along this path). In full generality, it studies fundamental properties of *topological spaces*, which generalize a broad array of geometric objects (including manifolds, Hilbert spaces, algebraic varieties and even embeddings of graphs). The concept of a topological space allows to speak in a very general manner about the "shape" of a space, and a prime goal of topology consists in classifying spaces according to their shapes. For instance, it is intuitively obvious that a mug with a handle and a football should belong to distinct classes of shapes, for instance because one has a hole in it and the other, preferably, does not. Whether or not, then, a mug with sharp edges and a doughnut should belong to the same class is a different question, and good reasons exist for choosing either way of answering it.

*C.B. and V.K. were supported by Austrian Science Fund (FWF, project Y1329), K.S. was supported by DFG Research Group ADYN via grant DFG 411362735, M.S. acknowledges support by the project "Grant Schemes at CU" (reg. no. CZ.02.2.69/0.0/0.0/19_073/0016935) and GAČR grant 22-19073S.

Thus, clearly, any such classification depends on the precise way in which the classes are defined and the structures provided on top of purely topological information (such as differential information, i.e., about "sharp edges"); one particularly important way of doing so is to make a single class out of all those shapes that can be deformed into each other according to specific rules retaining. The usual notion of equivalence under deformation of shapes corresponding to general topological spaces is furnished by *homotopy*, which, very roughly speaking, identifies any two shapes that can be obtained from one another through arbitrary deformations without "tearing" or "cutting" (and hence identifying the mug with the doughnut, while differentiating both from the football).

Associated to this notion are the so-called *homotopy groups* of a topological space, denoted π_k, for $k \geq 1$. The most intuitive of them is the group π_1, which is often called the *fundamental group* of the space. It captures certain data about the different ways that loops (that is, closed curves in the space) can pass through the space. The higher homotopy groups ($k > 1$) correspond to ways of routing higher-dimensional "loops" in the space, and Whitehead's Theorem provides a crucial equivalence between the structure of homotopy groups and the homotopy class of a broad category of topological spaces called CW-complexes [19, 20]. The present paper deals with an intermediate problem related to the computation of homotopy groups, which allows to show lower bounds for the complexity of computing the higher homotopy groups of a topological space.

Before speaking about computational tasks associated with topological spaces, one needs to define how a topological space is even represented. While the generality of the concept may make it seem hard to come up with such a representation in general, the usual path taken in computational topology is as follows: Many topological spaces can be described by finite structures, e.g., by abstract simplicial complexes, which are simply collections of point sets closed under taking subsets, and it hence suffices to provide the maximal subsets of a simplicial complex to specify it in full. Such structure can then be used as an input for a computer and therefore, it is natural to ask how hard it is to compute these homotopy groups of a given topological space, represented by an abstract simplicial complex.

Novikov in 1955 [16] and independently Boone in 1959 [6] showed undecidability of the word problem for groups. Their result also implies undecidability of computing the fundamental group. In fact, even determining whether the fundamental group of a given topological space is trivial is undecidable.

On the other hand, for 1-connected spaces (for those, whose π_1 is trivial) it is known that their π_k for $k > 1$ are finitely generated abelians group which are always isomorphic to groups of the form $\mathbb{Z}^n \oplus \mathbb{Z}_{p_1} \oplus \mathbb{Z}_{p_2} \oplus \cdots \oplus \mathbb{Z}_{p_m}$, where $p_1, \ldots, p_m$ are powers of prime numbers.[1] An algorithm for computing π_k of a 1-connected space, where $k > 1$, was first introduced by Brown in 1957 [7].

In 1989, Anick [4] proved that computing the rank of π_k, that is, the number of direct summands isomorphic to $\mathbb{Z}$ (represented by n in the expression above) is #P-hard for 4-dimensional 1-connected spaces.[2] Another computational problem called VEST, which we define below, was used in Anick's proof as an intermediate step. Briefly said, #P-hardness of VEST implies #P-hardness of computing the rank of π_k, which is the motivation for studying the problem in the present article.

Vector Evaluated After a Sequence of Transformations (VEST). The input of this problem defined by Anick [4] is a vector $\mathbf{v} \in \mathbb{Q}^d$, a list $(T_1, \ldots, T_m)$ of rational $d \times d$ matrices and a rational matrix $S \in \mathbb{Q}^{h \times d}$ where $d, m, h \in \mathbb{N}$.

[1]Note that $\mathbb{Z}^n$ is a direct sum of n copies of $\mathbb{Z}$ while $\mathbb{Z}_{p_i}$ is a finite cyclic group of order p_i.

[2]When k is a part of the input and represented in unary.

For an instance of VEST let an M-*sequence* be a sequence of integers $M_1, M_2, M_3, \ldots$, where

$$M_k := |\{(i_1, \ldots, i_k) \in \{1, \ldots, m\}^k ; ST_{i_k} \cdots T_{i_1}\mathbf{v} = \mathbf{0}\}|.$$

Given an instance of VEST and $k \in \mathbb{N}$, the goal is to compute M_k.

From an instance of VEST, it is possible to construct a corresponding algebraic structure called $123H$-*algebra* in polynomial time whose *Tor-sequence* is equal to the M-sequence of the original instance of a VEST. This is stated in [4, Theorem 3.4] and it follows from [2, Theorem 1.3] and [3, Theorem 7.6].

Given a presentation of a $123H$-algebra, one can construct a corresponding 4-dimensional simplicial complex in polynomial time whose sequence of ranks (rk π_2, rk $\pi_3, \ldots$) is related to the Tor-sequence of the $123H$-algebra. In particular, it is possible to compute that Tor-sequence from the sequence of ranks using an FPT algorithm. (To be defined in the next paragraph). This follows from [18] and [8]. To sum up, hardness of computing M_k of VEST implies hardness of computing π_k.

Parameterized Complexity and the W-hierarchy Parameterized complexity classifies decision or counting computational problems with respect to a given parameter(s). For instance, one can ask if there exists an independent set of size k in a given graph or how many independent sets of size k (for counting version) are in a given graph, respectively, where k is the parameter. From this viewpoint, we can divide problems into several groups which form the W-*hierarchy*.

$$\mathsf{FPT} \subseteq \mathsf{W}[1] \subseteq \mathsf{W}[2] \subseteq \cdots \subseteq \mathsf{XP}$$

The class FPT consists of decision problems solvable in time $f(k)n^c$, where $f(k)$ is a computable function of the parameter k, n is the size of input and c is a constant, while the class XP consists of decision problems solvable in time $cn^{f(k)}$. The class $\mathsf{W}[1]$ consists of all problems which admit a parameterized reduction to the satisfiability problem of a boolean circuit of constant depth with AND, OR and NOT gates such that there is at most 1 gate of higher input size than 2 on each path from the input gate to the final output gate (this number of larger gates is called *weft*), where the parameter is the number of input gates set to TRUE. Here, a parameterized reduction from a parameterized problem A to a parameterized problem B is an algorithm that, given an instance (x, k) of A, in time $f(k)n^c$ produces an equivalent instance (x, k') of B such that $k' \leq g(k)$, for some computable functions $f(\cdot)$, $g(\cdot)$, and a constant c. See Figure 1 for an example of a reduction showing $\mathsf{W}[1]$-completeness of finding independent set of size k.

The class $\mathsf{W}[i]$ then consists of problems that admit a parameterized reduction to the satisfiablity problem of a boolean circuit of a constant depth and weft at most i, parameterized by the number of input gates set to TRUE.

It is only known that $\mathsf{FPT} \subsetneq \mathsf{XP}$, while the other inclusions in the W-hierarchy are not known to be strict. However, it is strongly believed that $\mathsf{FPT} \subsetneq \mathsf{W}[1]$. *Therefore, one cannot expect existence of an algorithm solving a* $\mathsf{W}[1]$*-hard problem in time* $f(k)n^c$ *where* $f(k)$ *is a computable function of* k *and* c *is a constant.* For the detailed presentation of W-hierarchy and parameterized complexity in general we refer the reader to [13].

Analogously, one can define classes FPT and XP for counting problems. That is, a class of counting problems solvable in time $f(k)n^c$ or $cn^{f(k)}$, respectively. Problems for which there is a parameterized counting reduction to a problem of counting solutions for a boolean circuit of constant depth and weft at most i then form class $\#\mathsf{W}[i]$. Note that there are decision problems from FPT whose counting versions are $\#\mathsf{W}[1]$-hard, e.g., counting paths or cycles of length k

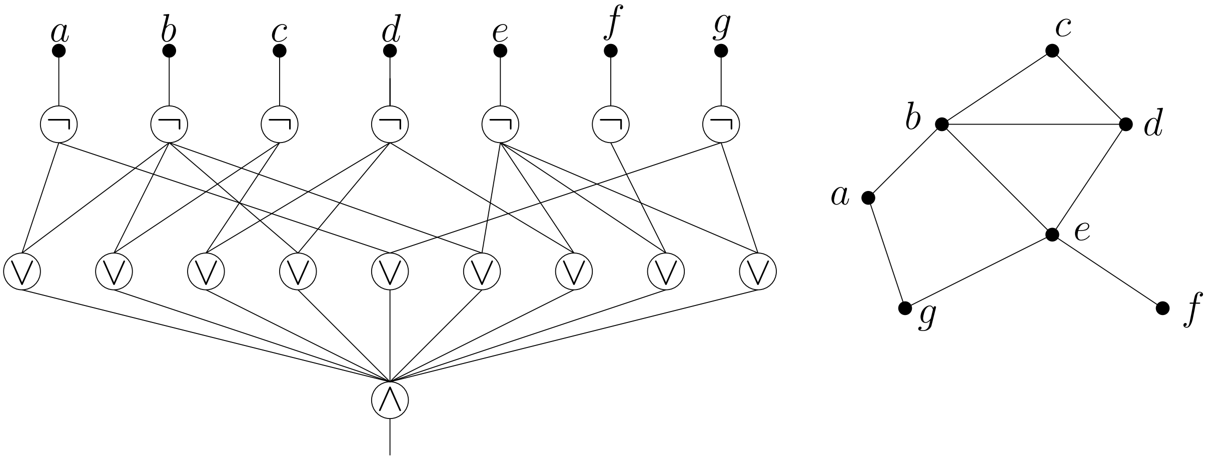

Figure 1: A boolean circuit solving the problem of existence of an independent set of size k in the graph on the left. There is an independent set of size k in the graph if and only if the boolean circuit outputs TRUE for an input consisting of exactly k true values.

parameterized by k [14]. Similarly to the decision case, *if a counting problem is shown to be* #W[i]*-hard for some* i *one should not expect existence of an algorithm solving this problem in time* $f(k)n^c$. For more details on parameterized counting we refer the reader to [14].

In our case, the number k of the homotopy group π_k plays the role of the parameter. In 2014 Čadek et al. [9] proved that computing π_k (and thus, also computing the rank of π_k) is in XP parameterized by k.

A lower bound for the complexity from the parameterized viewpoint was obtained by Matoušek in 2013 [15]. He proved that computing M_k of a VEST instance is #W[1]-hard. This also implies #W[1]-hardness for the original problem of computing the rank of higher homotopy groups π_k (for 4-dimensional 1-connected spaces) for parameter k. Matoušek's proof also works as a proof for #P-hardness and it is shorter and considerably easier than the original proof of Anick in [4].

In this paper, we strengthen the result of Matoušek and show that computing M_k of a VEST instance is #W[2]-hard. Our proof is even simpler than the previous proof of #W[1]-hardness.

Theorem 1. *Given a* VEST *instance, computing* M_k *is* #W[2]*-hard when parameterized by* k.

Theorem 1 together with the result of Anick [4] implies the following.

Corollary 2. *Computing the rank of the* k*-th homotopy group of a* d*-dimensional 1-connected space for* $d > 3$ *is* #W[2]*-hard for parameter* k.

Remark 3. Note that computing M_k of a VEST instance is an interesting natural self-contained problem even without the topological motivation. We point out that our reduction showing #W[2]-hardness of this problem uses only 0, 1 values in the matrices and the initial vector $\mathbf{v}$. Moreover, each matrix will have at most one 1 in each row. Therefore, such construction also shows #W[2]-hardness of computing M_k of a VEST instance in the $\mathbb{Z}_2$ setting. That is, for the case when $T_1, T_2, \ldots T_m \in \mathbb{Z}_2^{d\times d}, S \in \mathbb{Z}_2^{h\times d}$ and $\mathbf{v} \in \mathbb{Z}_2^d$.

The Decision Version of VEST We also provide a comprehensive overview of the parameterized complexity of VEST as a decision problem, where given an instance of VEST one needs to determine whether $M_k > 0$. In addition to the standard variant of the problem, we consider several modifications of VEST: when the matrices have constant size, when the matrix S is the identity matrix, when we omit the initial vector and the target is identity/zero matrix etc.

Unfortunately, even considering the simplifications above, we show that nearly all versions in our consideration are W[1]- or W[2]-hard. The following table is an overview of our results.

Size of matrices	a) $\mathbf{v}$ and S	b) only $\mathbf{v}$		c) only S	d) no $\mathbf{v}$, no S
1. 1×1	P	P	$\mathbf{0}$	P	P
			I	W[1]-hard	W[1]-hard
2. 2×2	W[1]-hard	W[1]-hard	$\mathbf{0}$	W[1]-hard	W[1]-hard
			I		
3. input size	W[2]-hard	W[2]-hard	$\mathbf{0}$	W[2]-hard	W[2]-hard
			I	W[1]-hard	W[1]-hard

The first column stands for the standard VEST while the second stands for the VEST without the special matrix S or alternatively, for the case when S is the identity matrix. Therefore, the hardness results for the first column follow from the second.

The third and the fourth columns are without the initial vector $\mathbf{v}$. In this case, it is natural to assume the following two targets for the result of the sought matrix product: the zero matrix (the rows labeled by $\mathbf{0}$) and the identity matrix (the rows labeled by I). Again, the hardness results for the third column follow from the fourth.

Regarding the 1×1 case, the only nontrivial case is when the target is $I = 1$. The W[1]-hardness results for the 1×1 case also implies W[1]-hardness for the 2×2 case and the input size case when the target is the identity matrix.

Therefore, in Section 3 we prove hardness for

- "1 d) I" (Theorem 6),
- "2 b)" (Theorem 8),
- "2 d) $\mathbf{0}$" (Theorem 7).

The #W[2]-hardness for "3 c)" follows from the proof of Theorem 1 (see Remark 5) and we show that "3 b)" and "3 d) $\mathbf{0}$" are equivalent to "3 a)" under parameterized reduction (Theorems 9, 10).

Fixed-Parameter Tractability over Finite Fields Reductions from the previous section show that VEST remains hard even on highly restricted instances, such as binary matrices with all the ones located along the main diagonal, or matrices of a constant size. However, it turns out that combination of this two restrictions – on the field size and the matrix sizes – makes even the counting version of VEST tractable.

We proceed by lifting tractability to the matrices of unbounded size but with all non-zero entries occurring in at most the p first rows.

Theorem 4. *Given an instance of* VEST *and* $k \in \mathbb{N}$, *computing* M_k *is* FPT *when parameterized by* $|\mathbb{F}|$ *and* p, *if all but the first* p *rows of the input matrices are zeros.*

The problem remains FPT with respect to $|\mathbb{F}|$ and p even if the task is to find the minimal k for which the vanishing sequence of length k exists, or to report that there is no such k.

Undecidability of VEST Without Parameter In contrast, we show in the last section (Section 5) that for $\mathbb{F} = \mathbb{Q}$ the problem of determining whether there exists k such that $M_k > 0$ for an instance of VEST is an undecidable problem (even for the case where $T_1, \ldots, T_m$ are of size 4×4).

$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Figure 2: The submatrix of T_u consisting of rows and columns $u_1, \ldots, u_4$. The rest of the non-diagonal entries of T_u are zeros. The diagonal entries $T_u^{w_1,w_1}$ for $w \in N[u]$ are zeros, the rest of the diagonal entries are ones.

2 The Proof of #W[2]-hardness of VEST

In this section, we prove that computing M_k of a VEST is #W[2]-hard (Theorem 1). Our reduction is from the problem of counting dominating sets of size k which is known to be #W[2]-complete (see [14]) and which we recall in the paragraph below.

For a graph $G(V, E)$ and its vertex $v \in V$ let $N[v]$ denote the closed neighborhood of a vertex v. That is, $N[v] := \{u \in V; \{u, v\} \in E\} \cup \{v\}$. A *dominating set* of a graph $G(V, E)$ is a set $U \subseteq V$ such that for each v there is $u \in U$ such that $v \in N[u]$.

Number of dominating sets of size k	
Input:	A graph $G(V, E)$ and a parameter k.
Question:	How many dominating sets of size k are in G?

Proof of Theorem 1. As we said, we show an FPT counting reduction from the problem of counting dominating sets of size k to VEST.

Let $G = (V, E)$ be the input graph and let $n = |V|$. The corresponding instance of VEST will consist of n matrices $\{T_u : u \in V\}$ of size $4n \times 4n$, one for each vertex, and matrix S of the same size. Whence, the initial vector $\mathbf{v}$ must be of size $4n$. For each vertex $u \in V$, we introduce four new coordinates $u_1, \ldots, u_4$ and set $\mathbf{v}_{u_1} = 1, \mathbf{v}_{u_2} = \mathbf{v}_{u_3} = 0$ and $\mathbf{v}_{u_4} = 1$.

We define the matrices $\{T_u : u \in V\}$ and S by describing their behavior. Let $\mathbf{x}$ be a vector which is going to be multiplied with a matrix T_u (that is, some intermediate vector obtained from $\mathbf{v}$ after potential multiplications). The matrix T_u sets $\mathbf{x}_{w_1}$ to zero for each $w \in N[u]$, which corresponds to domination of vertices in $N[u]$ by the vertex u, and also sets $\mathbf{x}_{u_2}$ to $\mathbf{x}_{u_3}$ and $\mathbf{x}_{u_3}$ to $\mathbf{x}_{u_4}$. The rest of the entries of $\mathbf{x}$ including $\mathbf{x}_{u_4}$ are kept, see Figure 2.

The matrix S then nullifies coordinates u_3, u_4 and keeps the coordinates u_1 and u_2 for each $u \in V$. In other words, S is diagonal such that $S^{u_1,u_1} = S^{u_2,u_2} = 1$ and $S^{u_3,u_3} = S^{u_4,u_4} = 0$.

The parameter remains equal to k.

For correctness, let $u^1, \ldots, u^k$ be any vertices from V, and let $\mathbf{r}$ be the vector obtained from $\mathbf{v}$ after multiplying by the matrices $T_{u^1}, \ldots, T_{u^k}$ (observe that the order of multiplication does not matter since all T_u, $u \in V$, pairwise commute). By construction, for every vertex $u \in V$, the entry $\mathbf{r}_{u_1} = 0$ if and only u is dominated by some u^i, $i \in [k]$, and $\mathbf{r}_{u_2} = 0$ if and only if T_u appears among $T_{u^1}, \ldots, T_{u^k}$ at most once. Indeed, if T_u is selected once then $\mathbf{r}_{u_2} = \mathbf{v}_{u_3} = 0$ while if it is selected more than once then $\mathbf{r}_{u_2} = \mathbf{v}_{u_4} = 1$. If T_u is not among $T_{u^1}, \ldots, T_{u^k}$ then $\mathbf{r}_{u_2} = \mathbf{v}_{u_2} = 0$.

Therefore, $\mathbf{r} = T_{u^1} \ldots T_{u^k}\mathbf{v}$ is a zero vector if and only if $u^1, \ldots, u^k$ are pairwise distinct and form the dominating set in G. This provides a one-to-one correspondence between subsets of matrices yielding the solution of VEST and dominating sets of size k in G. It remains to note that every such subset of matrices gives rise to $k!$ sequences that have to be counted in M_k. Hence, $M_k = k!D_k$ where D_k is the number of dominating sets of size k in G. The reduction is clearly FPT since the construction does not use parameter k and is polynomial in size of the input. ☐

Remark 5. Note that the decision version of the problem of DOMINATING SETS OF SIZE k is $\mathsf{W}[2]$-hard. For showing $\mathsf{W}[2]$-hardness of the decision version of VEST we need not deal with the repetition of matrices. In particular, we do not need the special coordinates u_2, u_3, u_4 and therefore, the corresponding instance of VEST can consist only of diagonal 0, 1 matrices of size $n \times n$.

3 Modifications of VEST

In this section, we prove hardness for the variants of the decision version of VEST we have discussed in the introduction. First of all, we recall a well-known $\mathsf{W}[1]$-hard k-SUM problem. See also [1].

k-SUM

Input:	A set A of integers and a parameter k.
Question:	Is it possible to choose k distinct integers from A such that their sum is equal to zero?

We note that in the versions of k-SUM studied in the literature the goal is to pick *distinct* elements of the input set in order to achieve 0 or eventually another number. However, the motivation for VEST, to the contrary, does not suggest that the matrices chosen for the product have to be distinct. Thus, in order to model VEST by k-SUM, it is more natural to also allow repetition of numbers. For our particular proofs, we will use the following version with target number 1.

AT-MOST-k-SUM WITH REPETITIONS AND TARGET 1

Input:	A set A of integers and parameter k.
Question:	Is it possible to choose *at most* k integers from A (possibly with repetition) such that their sum is equal to 1?

We are not aware of any previous studies on parameterized complexity of AT-MOST-k-SUM WITH REPETITIONS AND TARGET 1, nor does it seem that there exists a simple parameterized reduction from the original variant of the problem to the one with repetitions. Therefore, in Appendix (see A, Theorem 15) we prove $\mathsf{W}[1]$-hardness of this problem directly. Our reduction is from the problem of k-EXACT COVER, which is known to be $\mathsf{W}[1]$-hard (see [11]).

k-EXACT COVER

Input:	A universe U, a collection $\mathcal{C}$ of subsets of U and a parameter k.
Question:	Can U be partitioned into k sets from $\mathcal{C}$?

When we assume multiplication instead of addition the following problem arises.

k-PRODUCT WITH REPETITIONS

Input:	A set A of rational numbers and a parameter k.
Question:	Is it possible to choose k numbers from A (possibly with repetitions) such that their product is equal to 1?

$\mathsf{W}[1]$-hardness for this problem might be a folklore result but we present a complete proof using a reduction from k-EXACT COVER.

Theorem 6. *k-PRODUCT WITH REPETITIONS is $\mathsf{W}[1]$-hard parameterized by k.*

Proof. We show a parameterized reduction from k-Exact Cover. For each element $u \in U$ we associate one prime p_u, then for each $C \in \mathcal{C}$ we set $i_C := p \prod_{c \in C} p_c$ where p is a prime which is not used for any element from U and $s := \frac{1}{p^k \prod_{u \in U} p_u}$.

The integers i_C for each $C \in \mathcal{C}$ and s then form the input for $(k+1)$-Product with Repetitions

If $C_1, C_2, \ldots, C_k \in \mathcal{C}$ is a solution of k-Exact Cover then $s \prod_{i=1}^{k} i_{C_i} = 1$.

Conversely, let $q_1, q_2, \ldots, q_{k+1}$ be a solution of the constructed $(k+1)$-Product with Repetitions. First of all, note that s must be chosen precisely once. Indeed, all numbers except for s are greater than 1 and thus, s must be chosen at least once. If it were chosen more than once it would not be possible to cancel a power of p^k in the denominator since the numerator would contain at most p^{k-1}. Therefore, the product of $q_1, q_2, \ldots, q_{k+1}$ is of the form $s i_{C_{j_k}} i_{C_{j_{k-1}}} \cdots i_{C_{j_1}} = 1$ which means that each prime representing an element of U in the denominator is canceled. In other words, each element of U is covered. Note also that since s is chosen precisely once there cannot be any repetition within $i_{C_{j_k}} i_{C_{j_{k-1}}} \ldots i_{C_{j_1}}$.

The reduction is parameterized since we only need the parameter k for k multiplications of $\frac{1}{p}$ and first $n+1$ primes, where $n = |U|$, can be generated in time $\mathcal{O}(n^3)$ using, e.g., the Sieve of Eratosthenes for $(n+1)^2$. This follows from the fact, that the first n primes lie among $1, \ldots, n^2$. For more details we refer the reader to Lemma 16 in Appendix (A). □

Let us now call the variant of VEST without S and $\mathbf{v}$ Matrix k-Product with Repetitions. As we have mentioned in the introduction we consider two cases regarding the target matrix. Namely, the Identity matrix and the Zero matrix:

Matrix k-Product with Repetitions resulting to Zero Matrix

Input:	A list of $d \times d$ rational matrices and a parameter k.
Question:	Is it possible to choose k matrices from the list (possibly with repetitions) such that their product is the $d \times d$ zero matrix?

Matrix k-Product with Repetitions resulting to Identity Matrix

Input:	A list of $d \times d$ rational matrices and a parameter k.
Question:	Is it possible to choose k matrices from the list (possibly with repetitions) such that their product is the $d \times d$ identity matrix?

Note that Matrix k-Product with Repetitions resulting to Identity Matrix for 1×1 matrices is exactly k-Product with Repetitions. Therefore $\mathsf{W}[1]$-hardness for Matrix k-Product with Repetitions resulting to Identity Matrix for all matrix sizes follows from Theorem 6.

Regarding Matrix k-Product with Repetitions resulting to Zero Matrix, we can easily see that it is solvable in linear time for 1×1 matrices. Indeed, it is sufficient to check whether $T_i = 0$ for some i. However, already for 2×2 matrices the problem becomes hard.

Theorem 7. Matrix k-Product with Repetitions resulting to Zero Matrix *is* $\mathsf{W}[1]$-*hard for parameter* k *even for* 2×2 *integer matrices.*

Proof. We reduce from At-Most-k-Sum with Repetitions and Target 1. For every integer x let us define

$$U_x := \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}.$$

It is easy to see that $U_xU_y = U_{x+y}$. Let $\mathcal{I}$ be an instance of At-Most-k-Sum with Repetitions and Target 1 with the set of integers A and parameter k. We create an equivalent instance $\mathcal{I}'$ of Matrix $(k+2)$-product with Repetitions with the set of matrices $\{U_a : a \in A\} \cup \{X\}$, where

$$X = \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix}.$$

For correctness, assume that $\mathcal{I}$ is a YES-instance and $a_1, \dots, a_\ell \in A$ are such that $\ell \leq k$ and $\sum_{i=1}^{\ell} a_i = 1$. Consider the following product of $\ell + 2$ matrices:

$$X \cdot \prod_{i=1}^{l} U_{a_i} \cdot X = X \cdot U_{\sum_{i=1}^{\ell} a_i} \cdot X = XU_1X = \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} = \mathbf{0}.$$

For the other direction, assume that $\mathcal{I}'$ is a YES-instance. Let ℓ, $1 \leq \ell \leq k+2$, be the minimal integer such that there are matrices $T_1, \dots, T_\ell$ from $\{U_a : a \in A\} \cup \{X\}$ with $T_\ell T_{\ell-1} \cdots T_1 = \mathbf{0} \in \mathbb{Q}^{2\times 2}$. Since the matrix X is idempotent (i.e. $X^2 = X$), it does not appear two times in a row, otherwise we could reduce the length of the product. Notice that X should appear at least once, since the determinants of all U_a are non-zero. Assume that there is precisely one occurrence of X, then the product has form:

$$U_rXU_s = \begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} -r & -rs + r \\ -1 & 1 - s \end{pmatrix} \neq \mathbf{0}.$$

Hence, X appears at least twice. Let us fix any two consequent occurrences and consider the partial product between them:

$$XU_rX = \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ r-1 & 1-r \end{pmatrix} = (1-r) \cdot X.$$

If $r \neq 1$, we would get a shorter product resulting in zero, which contradicts to minimality of ℓ. Hence $r = 1$, so the product of U_a that appear between two occurrences of X is equal to U_1. Since there are at most k of such U_a and the sum of corresponding indices a is equal to 1, we obtain a solution to $\mathcal{I}$. $\square$

We can use similar approach to establish hardness of the VEST problem without S (or alternatively when S is the identity matrix). Recall that here the task is to obtain not necessarily a zero matrix but any matrix which contains a given vector $\mathbf{v}$ in a kernel.

Theorem 8. VEST *is* W[1]*-hard for parameter k even for 2×2 integer matrices and when S is the identity matrix.*

The proof of this theorem is very similar to the proof of Theorem 7 and it can be found in Appendix (see A).

At the end of this section, we show that VEST is equivalent to VEST without S (in other words, when $S = I_d$) and to Matrix k-Product with Repetitions resulting to Zero Matrix.

Theorem 9. *There is a parameterized reduction from* VEST *to the special case of* VEST *where S is the identity matrix, and the other way around.*

$$
\left(\begin{array}{c|ccc} & 0 & \dots & 0 \\ S & \vdots & \ddots & \vdots \\ & 0 & \dots & 0 \end{array}\right) \in \mathbb{Q}^{h\times h}, \quad \left(\begin{array}{ccc|ccc} & & & 0 & \dots & 0 \\ & T_i & & \vdots & \ddots & \vdots \\ & & & 0 & \dots & 0 \\ \hline 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & 0 & \dots & 0 \end{array}\right) \in \mathbb{Q}^{h\times h}, \quad \left(\begin{array}{c} \mathbf{v} \\ \hline 0 \\ \vdots \\ 0 \end{array}\right) \in \mathbb{Q}^{h}.
$$

Figure 3: A figure showing how to make all matrices square in the proof of Theorem 9 when $h > d$.

$$
\mathbf{v}' = \left(\begin{array}{c} \mathbf{v} \\ \hline k \\ 1 \end{array}\right), S' = \left(\begin{array}{ccc|cc} & & & 0 & 0 \\ & S & & \vdots & \vdots \\ & & & 0 & 0 \\ \hline 0 & \dots & 0 & 10 & 0 \\ 0 & \dots & 0 & 0 & 0 \end{array}\right), T_i' = \left(\begin{array}{ccc|cc} & & & 0 & 0 \\ & T_i & & \vdots & \vdots \\ & & & 0 & 0 \\ \hline 0 & \dots & 0 & 1 & -1 \\ 0 & \dots & 0 & 0 & 1 \end{array}\right).
$$

Figure 4: A construction forcing the matrix S' to be selected last in the proof of Theorem 9.

Proof. One direction is trivial since the case when $S = I$ is just a special case of VEST.

Regarding the other, let $\left(S \in \mathbb{Q}^{h\times d}, T_1, T_2, \dots, T_m \in \mathbb{Q}^{d\times d}, \mathbf{v} \in \mathbb{Q}^d, k\right)$ be an instance of VEST. First, we observe that without loss of generality we can suppose that S is a square matrix (in other words, $h = d$). Indeed, if $h < d$ then we just add $d-h$ zero lines to S. If $h > d$ we add $h-d$ zero columns to S, $h-d$ zero entries to $\mathbf{v}$ and $h-d$ zero lines as well as $h-d$ zero columns to each T_i. See Figure 3.

Now, we add 2 dimensions: To the vector $\mathbf{v}$ we add k on the $(d+1)$-st position and 1 on the $(d+2)$-nd position. To each matrix matrix T_i we add a 2×2 submatrix which subtracts the $(d+2)$-nd component of a vector from the $(d+1)$-st. To the matrix S we add a submatrix which nullifies the $(d+2)$-nd component and multiplies the $(d+1)$-th component by 10. Let $S', T_1', T_2' \dots, T_m'$ denote the resulting $(d+2)\times(d+2)$ matrices and $\mathbf{v}'$ denote the resulting $(d+2)$-dimensional vector. See Figure 4. The new parameter is set to $k+1$.

If there is a solution of the original problem, that is, there are k matrices $T_{i_1}, \dots, T_{i_k}$ such that $ST_{i_k}T_{i_{k-1}}\cdots T_{i_1}\mathbf{v} = \mathbf{0}$, then $S'T'_{i_k}T'_{i_{k-1}}\cdots T'_{i_1}\mathbf{v}' = \mathbf{0}$, since 1 is k times subtracted from the $(d+1)$-st component of $\mathbf{v}'$ and the $(d+2)$-nd component is then nullified by S'.

Conversely, if there are $k+1$ matrices $Y_1, Y_2, \dots, Y_{k+1}$, where each Y_i is either S' or T_j' for some j, such that $\mathbf{r} = Y_{k+1}Y_k\cdots Y_1\mathbf{v}' = \mathbf{0}$ then Y_{k+1} must be equal to S' and the rest of the matrices are of type T_j', otherwise $\mathbf{r}_{d+1} \neq 0$ or $\mathbf{r}_{d+2} \neq 0$. Indeed, at first k matrices of type T_j' must be selected to nullify the $(d+1)$-st component: if $Y_i = S'$ for some $i \le k$, this would increase the non-zero $(d+1)$-st component, so there would be no way to nullify it by remaining matrices $Y_{i+1}, \dots, Y_{k+1}$. At the same time, S' should be necessarily selected once to nullify the $(d+2)$-nd component, so $Y_{k+1} = S'$. Therefore, by restricting the matrices $Y_1, \dots, Y_k$ to the first d coordinates we obtain a solution to VEST with matrix S. □

Theorem 10. VEST *and* Matrix k-Product with Repetitions resulting to Zero Matrix *are equivalent under parameterized reduction.*

Note that one implication is relatively straightforward. Regarding the other, the idea is to again "simulate" the special matrix S and the vector $\mathbf{v}$ by an ordinary matrix and force them to be selected as the leftmost and the rightmost, respectively. For the complete proof please see A.

4 Fixed-Parameter Tractability of VEST over Finite Fields

While most of the hardness results for VEST and its variations in the previous section use constant-sized matrices, the entries of this matrices can be arbitrarily large. Here, we study the variation of the problem when all the matrices have entries from some finite field. Notice that restricting the field size by itself does not make the problem tractable: recall the reduction from dominating set from Section 1 which also works over $\mathbb{Z}_2$. However, along with a bound on the matrix sizes this makes the problem tractable.

Lemma 11. *Computing M_k for a given instance of* VEST *over finite field $\mathbb{F}$ is* FPT *when parameterized by the size of $\mathbb{F}$ and the size of matrices.*

Proof. Let $\mathcal{M}_{\mathbb{F}}^d$ be the set of all $d \times d$ matrices with entries from $\mathbb{F}$, then $|\mathcal{M}_{\mathbb{F}}^d| = |\mathbb{F}|^{d^2}$. For every $X \in \mathcal{M}_{\mathbb{F}}^d$ and every integer $i \in [k]$ we will compute a value $a_X^i \in \mathbb{N}_0$ equal to the number of sequences of i matrices from the input such that their product is equal to X. In particular, this allows to obtain $M_k = \sum_{X \in \mathcal{M}_{\mathbb{F}}^d : SX\mathbf{v}=\mathbf{0}} a_X^k$.

For $i = 1$ the computation can be done simply by traversing the input matrices. Assume that a_X^i have been computed for all the matrices X and all $i \in [j]$. We initiate by setting $a_X^{j+1} = 0$ for every $X \in \mathcal{M}_{\mathbb{F}}^d$. Then, for every pair (X, q), where $X \in \mathcal{M}_{\mathbb{F}}^d$ and $q \in [m]$, we increment $a_{XT_q}^{j+1}$ by a_X^j. In the end we will then have a correctly computed value $a_Y^{j+1} = \sum_{q=1}^m \sum_{X \,:\, XT_q = Y} a_X^j$. $\square$

Our next step is to consider the matrices of unbounded size, but with at most p first rows containing non-zero entries. In particular, if $\mathbb{F} = \mathbb{Z}_2$, we can associate to every such matrix T a graph with the vertex set $[d]$ such that there exists an edge between the vertices i and j, $i \leq j$, if and only if $T^{i,j} = 1$. Conversely, a graph with the vertex set $[d]$ can be represented by such a matrix if and only if the vertices in $[p]$ form it's vertex cover.

Observe that every matrix T with at most p first non-zero rows has the following form:

$$T = \left(\begin{array}{c|ccc} A & & B & \\ \hline 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{array}\right), \text{ where } A \text{ is } p \times p \text{ matrix and } B \text{ is } p \times (d-p) \text{ matrix.}$$

Further, we will denote matrices of this form by $A|B$. Consider the product of two such matrices $T_1 = A_1|B_1$ and $T_2 = A_2|B_2$:

$$\left(\begin{array}{c|ccc} A_1 & & B_1 & \\ \hline 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{array}\right) \left(\begin{array}{c|ccc} A_2 & & B_2 & \\ \hline 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{array}\right) = \left(\begin{array}{c|ccc} A_1A_2 & & A_1B_2 & \\ \hline 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{array}\right) = (A_1A_2)|(A_1B_2).$$

Corollary 12. $\prod_{i=1}^k (A_i|B_i) = (\prod_{i=1}^k A_i)|(\prod_{i=1}^{k-1} A_i \cdot B_k) = (XA_k)|(XB_k)$*, where* $X = \prod_{i=1}^{k-1} A_i$*. In particular, the product does not depend on B_i for $i < k$.*

Theorem 4. *Given an instance of* VEST *and* $k \in \mathbb{N}$, *computing* M_k *is* FPT *when parameterized by* $|\mathbb{F}|$ *and* p, *if all but the first* p *rows of the input matrices are zeros.*

Proof. We slightly modify the definition of $a_X^i \in \mathbb{N}_0$ from the proof of Theorem 11. Now, for every $i \in [k]$ and every matrix $X \in \mathcal{M}_{\mathbb{F}}^p$, let a_X^i be the number of sequences of i matrices $T_j = A_j|B_j$ from the input such that corresponding product of A_j is equal to X.

The values of a_X^i for every $i \in [k-1]$ can be computed same as in the proof of Theorem 11. Given this information, we can count the sequences of length k that nullify $\mathbf{v}$. Indeed, by Corollary 12, the number of such sequences with the last matrix $T_j = A_j|B_j$ is precisely $b_j = \sum_{X \in \mathcal{M}_{\mathbb{F}}^p \colon S\cdot(XA_j)|(XB_j)\cdot\mathbf{v}=\mathbf{0}} a_X^{k-1}$. M_k is then equal to $\sum_{j=1}^m b_j$. □

We remark that the algorithm for computing a_X^i from the last proof can be exploited to determine minimal k such that $M_k > 0$, or to report that there is no such k. For this, let us run the algorithm with $k = 1$, then with $k = 2$ and so on. If after some iteration $k = j + 1$ we obtain that $M_i = 0$ for all $i \in [j]$ and there is no $X \in \mathcal{M}_{\mathbb{F}}^p$ such that $a_X^1 = \cdots = a_X^j = 0$ and $a_X^{j+1} \neq 0$, we may conclude that $M_k = 0$ for all $k \in N$, since every product of length more than j can be obtained as a product of length at most j, and none of the latter nulify $\mathbf{v}$. Otherwise, there exists at least one $X \in \mathcal{M}_{\mathbb{F}}^p$ such that $a_X^1 = \cdots = a_X^j = 0$ and $a_X^{j+1} \neq 0$. Note that every $X \in \mathcal{M}_{\mathbb{F}}^p$ can play this role only for one value of k. Therefore, it always suffices to make $|\mathcal{M}_{\mathbb{F}}^p|$ iterations of the algorithm.

5 Undecidability of VEST

In this section, we show that determining whether there exists $k \in \mathbb{N}$ such that $M_k > 0$ for an instance of VEST is an undecidable problem. The reduction is from Post's Correspondence Problem which is known to be undecidable. See [17].

(Binary) Post's Correspondence Problem	
Input:	m pairs $(v_1, w_2), (v_2, w_2), \ldots, (v_m, w_m)$ of words over alphabet $\{0, 1\}$.
Question:	Is possible to choose k pairs $(v_{i_1}, w_{i_k}), (v_{i_2}, w_{i_2}) \ldots, (v_{i_k}, w_{i_k})$, for some $k \in \mathbb{N}$, such that $v_{i_1} v_{i_2} \cdots v_{i_k} = w_{i_1} w_{i_2} \cdots w_{i_k}$?

For a word $v \in \{0, 1\}^*$ let $|v|$ be its length and let $(v)_2$ be the integer value of v interpreting it as a binary number. Let us define the following matrix for a binary word v.

$$T_v = \begin{pmatrix} 2^{|v|} - (v)_2 & (v)_2 \\ 2^{|v|} - (v)_2 - 1 & (v)_2 + 1 \end{pmatrix}, \text{ then the following holds:}$$

Lemma 13. *Let* v, w *be binary words. Then,* $T_v T_w = T_{wv}$ *where* wv *is the concatenation of* w *and* v.

Note that the construction of T_v is a based on [10][Satz 28, p. 157] which we are aware of thanks to Günter Rote. For the complete proof of Lemma 13 please see A.

Reduction Given an instance of Post's Correspondence Problem we describe what an instance of VEST may look like. For each pair (v, w) we define

$$T_{(v,w)} = \left(\begin{array}{c|c} T_v & \begin{matrix} 0 & 0 \\ 0 & 0 \end{matrix} \\ \hline \begin{matrix} 0 & 0 \\ 0 & 0 \end{matrix} & T_w \end{array}\right),$$

we set the initial vector $\mathbf{v} := (0,1,0,1)^T$ and $S := (1,0,-1,0)$. The undecidability of VEST then follows from the following lemma.

Lemma 14. *Let $(v_{i_1}, w_{i_1}), (v_{i_2}, w_{i_2}) \dots, (v_{i_k}, w_{i_k})$ be k pairs of binary words. Then*

$$ST_{(v_{i_k}, w_{i_k})} T_{(v_{i_{k-1}}, w_{i_{k-1}})} \cdots T_{(v_{i_1}, w_{i_1})} \mathbf{v} = \mathbf{0}$$

if and only if $v_{i_1} v_{i_2} \cdots v_{i_k} = w_{i_1} w_{i_2} \cdots w_{i_k}$.

Proof. By Lemma 13 $T_{(v_{i_k}, w_{i_k})} T_{(v_{i_{k-1}}, w_{i_{k-1}})} \cdots T_{(v_{i_1}, w_{i_1})} = T_{(v_{i_1} v_{i_2} \cdots v_{i_k}, w_{i_1} w_{i_2} \cdots w_{i_k})}$. The vector $\mathbf{v}$ selects the second column of the submatrix $T_{v_{i_1} v_{i_2} \cdots v_{i_k}}$ and the second column of the submatrix $T_{w_{i_1} w_{i_2} \cdots w_{i_k}}$. In other words, the result is equal to

$$\left((v_{i_1} v_{i_2} \cdots v_{i_k})_2, (v_{i_1} v_{i_2} \cdots v_{i_k})_2 + 1, (w_{i_1} w_{i_2} \cdots w_{i_k})_2, (w_{i_1} w_{i_2} \cdots w_{i_k})_2 + 1\right)^T.$$

The final result after multiplying S with the vector above is the following 1-dimensional vector $(v_{i_1} v_{i_2} \cdots v_{i_k})_2 - (w_{i_1} w_{i_2} \cdots w_{i_k})_2$. □

References

[1] Amir Abboud, Kevin Lewi, and Ryan Williams. Losing weight by gaining edges. In Andreas S. Schulz and Dorothea Wagner, editors, *Algorithms - ESA 2014 - 22th Annual European Symposium, Wroclaw, Poland, September 8-10, 2014. Proceedings*, volume 8737 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2014.

[2] David J. Anick. Diophantine equations, Hilbert series, and undecidable spaces. *Annals of Mathematics*, 122:87–112, 1985.

[3] David J. Anick. Generic algebras and CW complexes. *Algebraic topology and algebraic K-theory*, pages 247–321, 1987.

[4] David J. Anick. The computation of rational homotopy groups is #$\wp$-hard. Computers in geometry and topology, Proc. Conf., Chicago/Ill. 1986, Lect. Notes Pure Appl. Math. 114. pages 1–56, 1989.

[5] Rosser Barkley. Explicit bounds for some functions of prime numbers. *American Journal of Mathematics*, 63(1):211–232, 1941.

[6] William W. Boone. The word problem. *Annals of mathematics*, 70:207–265, 1959.

[7] Edgar H. Brown. Finite computability of Postnikov complexes. *Annals of Mathematics*, 65:1, 1957.

[8] Martin Čadek, Marek Krčál, Jiří Matoušek, Lukáš Vokřínek, and Uli Wagner. Extendability of continuous maps is undecidable. *Discrete & Computational Geometry*, 51(1):24–66, 2014.

[9] Martin Čadek, Marek Krčál, Jiří Matoušek, Lukáš Vokřínek, and Uli Wagner. Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. *SIAM Journal on Computing*, 43(5):1728–1780, 2014.

[10] Volker Claus. *Stochastische Automaten.* Vieweg+Teubner Verlag, 1971. (in German).

[11] Rod G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theoretical Computer Science*, 141(1):109–131, 1995.

[12] Graham Flegg. *From geometry to topology.* Courier Corporation, 2001.

[13] Jörg Flum and Martin Groge. *Parameterized Complexity Theory.* Springer, 2004.

[14] Jörg Flum and Martin Grohe. The parameterized complexity of counting problems. *SIAM Journal on Computing*, 33(4):892–922, 2004.

[15] Jiří Matoušek. Computing higher homotopy groups is W[1]-hard. *arXiv preprint arXiv:1304.7705*, 2013.

[16] Pyotr S. Novikov. On the algorithmic unsolvability of the word problem in group theory. *Trudy Mat. Inst. Steklov*, 44:1–143, 1955. (in Russian).

[17] Emil L. Post. A variant of a recursively unsolvable problem. *Journal of Symbolic Logic*, 12(2):55–56, 1947.

[18] Jan-Erik Roos. Relations between the poincaré-betti series of loop spaces and of local rings. In *Séminaire d'Algèbre Paul Dubreil*, pages 285–322. Springer, 1979.

[19] J. H. C. Whitehead. Combinatorial homotopy. i. *Bulletin of the American Mathematical Society*, 55(3):213–245, 1949.

[20] J. H. C. Whitehead. Combinatorial homotopy. II. *Bulletin of the American Mathematical Society*, 55(5):453 – 496, 1949.

A Appendix

Our aim in the appendix is to give a complete proof of

- W[1]-hardnes of At-Most-k-Sum with Repetitions and Target 1 (Theorem 15),
- Theorem 8,
- Theorem 10,
- two auxiliary lemmas. Namely, Lemma 16 and Lemma 13.

Theorem 15. At-Most-k-Sum with Repetitions and Target 1 *is* W[1]*-hard when parameterized by k.*

Proof. Consider an instance $(U, \mathcal{C}, k)$ of Unique Hitting Set. Intuitively, we would like to model the sets in $\mathcal{C}$ as their characteristic vectors over $|U|$ dimensions, where each dimension corresponds to an element from U, and the vector representing a set $C \in \mathcal{C}$ is set to one exactly in those dimensions which correspond to the elements contained in the set which is represented by the vector. To model this in an instance of At-Most-k-Sum with Repetitions and Target 1, we will represent said characteristic vectors as numbers in base $(k + 2)$.

Formally, let $m = |U|$, $U = \{u_1, \ldots, u_m\}$, and $x = k+2$. For each $C \in \mathcal{C}$, we add an element $a_C = -\left(x^{m+1} + \sum_{j;u_j \in C} x^j\right)$ to the set A of numbers. Then we also add to A the number $y := kx^{m+1} + \sum_{j=0}^{m} x^j$ and we set the new parameter to $k+1$. Note that the numbers in A are bounded by x^{m+2}, thus can be represented by $O(m \log k)$ bits, and $|A| = |\mathcal{C}|$, meaning that the reduction can be done in polynomial time. It remains to verify that the produced instance of At-Most-k-Sum with Repetitions and Target 1 is equivalent to the original instance of k-Exact Cover.

First, let $C_1, \ldots, C_k \in U$ be a solution to k-Exact Cover. We claim that $\{y, a_{C_1}, \ldots, a_{C_k}\} \subset A$ is a solution to the instance $(A, k+1)$ of At-Most-k-Sum with Repetitions and Target 1. Indeed, by construction and since each element of U is covered exactly once, we have $a_{C_1} + \cdots + a_{C_k} = -\left(kx^{m+1} + \sum_{j=1}^{m} x^j\right) = -y + x^0 = -y + 1$. Therefore, $y + a_{C_1} + \cdots + a_{C_k} = 1$.

In the other direction, consider a solution $a_1, \ldots, a_t \in A$ to At-Most-k-Sum with Repetitions and Target 1 where $t \leq k+1$. First of all, we observe that y must be chosen precisely once. The sum $\sum_{j=1}^{t} a_t = 1 = x^0$ and y is the only number with a coefficient ($= 1$) of x^0. Therefore, y can be chosen $(\ell x + 1)$ times where $\ell \in \mathbb{N}_0$. However $t < x = k+2$. Whence, $\ell = 0$. In other words, y is chosen precisely once and without loss of generality, we suppose that $a_1 = y$.

Next, we show that $t = k+1$. The number y which is chosen precisely once has k as the coefficient of x^{m+1} which has to be nullified. The only option how to do that is to choose k numbers other than y. (Such numbers are negative and have 1 as a coefficient of x^{m+1}.)

Finally, from the equality $\sum_{j=2}^{k+1} a_j = -y + 1 = -kx^{m+1} - \sum_{j=1}^{m} x^j$ we conclude that no $-x^i$ for $i \leq m$ is contained in more than one a_j as a summand since $k < k+2 = x$. By the same argument we observe that each $-x^i$ is contained in some a_j as a summand. Indeed, addition of at most k terms $-x^i$ cannot affect coefficient of x^{i+1}. Therefore, each $-x^i$ for $i \leq m$ is contained in precisely one a_j and thus, $\{C; a_C \in \{a_2, \ldots, a_{k+1}\}\}$ is a desired k-exact cover. ☐

Proof of Theorem 8. As in the proof of Theorem 7, we proceed by reduction from At-Most-k-Sum with Repetitions and Target 1. Let $\mathcal{I}$ be an arbitrary instance of the problem with the set of integers A and parameter k. We create an equivalent instance $\mathcal{I}'$ of VEST with parameter $k+1$, vector $v = (0,1)^T$ and the set of matrices $\{U_a : a \in A\} \cup \{X\}$, where U_a and X are defined same as in the proof of Theorem 7. We set S equal to the identity matrix.

For correctness, assume that $\mathcal{I}$ is a YES-instance and $a_1, \ldots, a_\ell \in A$ are such that $\ell \leq k$ and $\sum_{i=1}^{\ell} a_i = 1$. We apply the following $\ell + 1$ matrices to nullify v:

$$X \cdot \prod_{i=1}^{\ell} U_{a_i} \cdot v = XU_1 v = \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

For another direction, assume that $\mathcal{I}'$ is a YES-instance. Let ℓ, $1 \leq \ell \leq k+1$, be the minimal integer such that $T_\ell \cdots T_1 v = (0,0)^T$ for some $T_1, \ldots, T_\ell$ from $\{U_a : a \in A\} \cup \{X\}$. Since the determinants of all U_a are non-zero, $T_i = X$ for some $i \in [\ell]$. Observe that $Xv = v$, so by minimality of ℓ we have that $T_1 \neq X$. Let i be the minimal index such that $T_i = X$, $2 \leq i \leq \ell$. Then $T_{i-1} \cdots T_1 = U_s$ for some integer s. Let us apply first i matrices to v:

$$\begin{aligned} T_i \cdots T_1 v = XU_s \cdot v &= \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} s \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 1-s \end{pmatrix}. \end{aligned}$$

If $s \neq 1$, we get a multiple of v, which is in contradiction to minimality of l. So $T_{i-1}, \ldots, T_1 = U_1$, which is a product of at most k matrices of the form U_a with $a \in A$. The sum of corresponding indices a is then equal to 1, resulting in a solution to $\mathcal{I}$. $\square$

Proof of Theorem 10.

1. "Parameterized reduction from Matrix k-Product with Repetitions resulting to Zero Matrix to VEST"
   For each matrix $T_i \in \mathbb{Q}^{d\times d}$ we introduce a block matrix $T'_i \in \mathbb{Q}^{d^2\times d^2}$ whose each block is T_i. We set $\mathbf{v} = (e_1, e_2, \ldots, e_d)^T \in \mathbb{Q}^{d^2}$ where each e_i is the d-dimensional unit vector with 1 on its i-th coordinate and S to the d^2-dimensional identity matrix. Therefore, $T_{i_k}T_{i_{k-1}} \cdots T_{i_1} = R$ if and only if $ST'_{i_k} \cdots T'_{i_1}\mathbf{v} = (R_{*,1}, R_{*,2}, \ldots, R_{*,d})^T$ where $R_{*,j}$ is the j-th column of the matrix R.

2. "Parameterized reduction from VEST to Matrix k-Product with Repetitions resulting to Zero Matrix"
   We first reduce VEST to the version of VEST without S as we did in the proof of Theorem 9. Thus, we assume that our input consists of the initial vector $\mathbf{v}$, square matrices $S', T'_1, \ldots, T'_m$, where S' represents the original special matrix S, and the parameter is $k+1$. Let us recall that S' has to be selected precisely once as the leftmost matrix otherwise the resulting vector cannot be zero by the construction from the proof of Theorem 9.

   Now, we create an instance of Matrix $(k+3)$-Product with Repetitions. Let $T_{\mathbf{v}}$ be a matrix containing the vector $\mathbf{v}$ in the first column and zero otherwise. The idea is to use the matrix $T_{\mathbf{v}}$ instead of the vector $\mathbf{v}$ and force such matrix to be selected as the rightmost after S' and k matrices of type T'_i by adding some blocks. We use the construction from the proof of Theorem 7. Namely, we use matrices X and U_{-2} and U_{2k+1} as submatrices. By the same argument as in the proof of Theorem 7 the only way how to make the zero matrix by multiplying $k+3$ matrices from $\{X, U_{-2}, U_{2k+1}\}$ is to choose X twice, as the leftmost and the rightmost matrix, k-times U_{-2} and once U_{2k+1} as intermediate matrices. Therefore, we can add X to $T_{\mathbf{v}}$ and to the identity matrix as block submatrices, U_{2k+1} to S' (since S' must be selected precisely once) and U_{-2} to T'_i. It remains to force the order of $T_{\mathbf{v}}$ and the identity matrix enriched by X. For this, we add submatrices A, B such that $AB = 0$ while $BA \neq 0, AA \neq 0, BB \neq 0$. We add A to the identity matrix enriched by X, B to $T_{\mathbf{v}}$ enriched by X and identity matrices to the rest. See Figure 5. The following settings for A and B, respectively, work.

$$A = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, B = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

$\square$

Lemma 16. *Let p_n denote the n-th prime. Then $p_n \leq n^2$ for $n \geq 2$.*

Proof. Let $\pi(x)$ denote the number of primes less than or equal to x. The lemma follows, e.g., from the following claims:

- $p_n < n(\ln n + \ln\ln n)$ for $6 \leq n \leq e^{95}$ (see [5, Theorem 28]),
- $\frac{x}{\ln x + 2} \leq \pi(x)$ for $x \geq 55$ (see [5, Theorem 29.A]),

$$
T_i'' = \left(\begin{array}{ccc|ccccc}
 & & & 0 & 0 & 0 & 0 & 0 \\
 & T_i' & & \vdots & \vdots & \vdots & \vdots & \vdots \\
 & & & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & \dots & 0 & & & & 0 & 0 \\
0 & \dots & 0 & & I_3 & & 0 & 0 \\
0 & \dots & 0 & & & & 0 & 0 \\ \hline
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{|c}{} \\
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{|c}{U_{-2}}
\end{array}\right), S'' = \left(\begin{array}{ccc|ccccc}
 & & & 0 & 0 & 0 & 0 & 0 \\
 & S' & & \vdots & \vdots & \vdots & \vdots & \vdots \\
 & & & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & \dots & 0 & & & & 0 & 0 \\
0 & \dots & 0 & & I_3 & & 0 & 0 \\
0 & \dots & 0 & & & & 0 & 0 \\ \hline
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{|c}{} \\
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{|c}{U_{2k+1}}
\end{array}\right),
$$

$$
T_{\mathbf{v}}' = \left(\begin{array}{ccc|ccccc}
 & & & 0 & 0 & 0 & 0 & 0 \\
 & T_{\mathbf{v}} & & \vdots & \vdots & \vdots & \vdots & \vdots \\
 & & & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & \dots & 0 & & & & 0 & 0 \\
0 & \dots & 0 & & B & & 0 & 0 \\
0 & \dots & 0 & & & & 0 & 0 \\ \hline
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{|c}{} \\
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{|c}{X}
\end{array}\right), H = \left(\begin{array}{ccc|ccccc}
 & & & 0 & 0 & 0 & 0 & 0 \\
 & I_d & & \vdots & \vdots & \vdots & \vdots & \vdots \\
 & & & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & \dots & 0 & & & & 0 & 0 \\
0 & \dots & 0 & & A & & 0 & 0 \\
0 & \dots & 0 & & & & 0 & 0 \\ \hline
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{|c}{} \\
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{|c}{X}
\end{array}\right).
$$

Figure 5: The instance of Matrix k-Product with Repetitions resulting to Zero Matrix obtained after the reduction from VEST in the proof of Theorem 10.

- $p_2 = 3$, $p_3 = 5$, $p_4 = 7$, $p_5 = 11$.

$\square$

Proof of Lemma 13. First of all, note that $2^{|v|}(w)_2 + (v)_2 = (wv)_2$. Using this observation, we compute all the entries of the matrix

$$
T_v T_w = \begin{pmatrix} 2^{|v|} - (v)_2 & (v)_2 \\ 2^{|v|} - (v)_2 - 1 & (v)_2 + 1 \end{pmatrix} \begin{pmatrix} 2^{|w|} - (w)_2 & (w)_2 \\ 2^{|w|} - (w)_2 - 1 & (w)_2 + 1 \end{pmatrix}.
$$

Thus,

$$\begin{aligned}
(T_vT_w)^{1,1} =& \left(2^{|v|} - (v)_2\right)\left(2^{|w|} - (w)_2\right) + (v)_2\left(2^{|w|} - (v)_2 - 1\right) \\
=&2^{|wv|} - 2^{|v|}(w)_2 - 2^{|w|}(v)_2 + (v)_2(w)_2 + 2^{|w|}(v)_2 - (v)_2(w)_2 - (v)_2 \\
=&2^{|wv|} - 2^{|v|}(w)_2 - (v)_2 \\
=&2^{|wv|} - (wv)_2, \\
(T_vT_w)^{1,2} =& \left(2^{|v|} - (v)_2\right)(w)_2 + (v)_2\left((w)_2 + 1\right) \\
=&2^{|v|}(w)_2 - (v)_2(w)_2 + (v)_2(w)_2 + (v)_2 \\
=&2^{|v|}(w)_2 + (v)_2 \\
=&(wv)_2, \\
(T_vT_w)^{2,1} =& \left(2^{|v|} - (v)_2 - 1\right)\left(2^{|w|} - (w)_2\right) + \left((v)_2 + 1\right)\left(2^{|w|} - (w)_2 - 1\right) \\
=&2^{|wv|} - 2^{|v|}(w)_2 - 2^{|w|}(v)_2 + (v)_2(w)_2 - 2^{|w|} + (w)_2 \\
&+ 2^{|w|}(v)_2 - (v)_2(w)_2 - (v)_2 + 2^{|w|} - (w)_2 - 1 \\
=&2^{|wv|} - 2^{|v|}(w)_2 - (v)_2 - 1 \\
=&2^{|wv|} - (wv)_2 - 1, \\
(T_vT_w)^{2,2} =& \left(2^{|v|} - (v)_2 - 1\right)(w)_2 + \left((v)_2 + 1\right)\left((w)_2 + 1\right) \\
=&2^{|v|}(w)_2 - (v)_2(w)_2 - (w)_2 + (v)_2(w)_2 + (v)_2 + (w)_2 + 1 \\
=&2^{|v|}(w)_2 + (v)_2 + 1 \\
=&(wv)_2 + 1.
\end{aligned}$$

□