

Variants of the Determinant Polynomial and the VP-completeness

Prasad Chaugule, Nutan Limaye, and Shourya Pandey

Indian Institute of Technology, Bombay, India
{prasad,nutan **,shouryap}@cse.iitb.ac.in

Abstract. The determinant is a canonical VBP-complete polynomial in the algebraic complexity setting. In this work, we introduce two variants of the determinant polynomial which we call $\mathbf{StackDet}_n(X)$ and $\mathbf{CountDet}_n(X)$ and show that they are VP and VNP complete respectively under p -projections. The definitions of the polynomials are inspired by a combinatorial characterisation of the determinant developed by Mahajan and Vinay (SODA 1997). We extend the combinatorial object in their work, namely *clow sequences*, by introducing additional edge labels on the edges of the underlying graph. The idea of using edge labels is inspired by the work of Mengel (MFCS 2013).

Keywords: Algebraic Circuits · VP-Completeness · Determinant Family.

1 Introduction

In an influential paper of Valiant [12], a complexity theoretic view of algebraic computation was presented. This work led to a classification of polynomials based on the ease of computing them. Consequently, complexity classes such as VF, VBP, VP and VNP were defined and investigated in many follow-up papers. These algebraic classes were designed with the intention of mimicking Boolean complexity classes. It was believed that they would give rise to equally interesting, but potentially easier to resolve questions. For example, the question of separating the classes VP and VNP turned out to be very interesting, like its Boolean counterpart, namely the famous question of separating NP from P.

While there are many parallels between these two worlds, over the years, many crucial differences between them have also surfaced. Specifically, in the Boolean world, many naturally occurring problems have been found to be *complete* for the classes NP and P¹. Although many naturally occurring polynomials

** Funded by SERB Project no. MTR/2017/000909

¹ A problem P is said to be complete for a Boolean complexity class \mathcal{C} if $P \in \mathcal{C}$ and any problem P' in \mathcal{C} reduces to P in polynomial time.

are known to be complete² for VNP, until very recently no natural polynomial was known to be complete for VP.

The process of finding many complete problems for a complexity class is crucial in many ways. For one, each complete problem presents a potentially different way of understanding the class. It also makes the complexity class rich and robust. In this work, we contribute to the class of VP-complete polynomials.

Until as recently as 2014, hardly any natural VP-complete polynomials were known. In Durand et. al. [2] and Mahajan et al. [5], many interesting and fairly natural families of polynomials were shown to be VP-complete. In [1], a few more polynomials complete for VP were presented. All these polynomials were based on counting graph homomorphisms³.

In this work we define two fairly simple to state variants of the determinant polynomial and show that they are VP and VNP complete. As the determinant is known to be complete for the class VBP⁴ (a class known to be contained in VP), this gives a satisfactory way of using the same base polynomial, namely the determinant polynomial, whose generalisations capture the class VP and VNP.

The determinant polynomial is a central object of study in algebraic complexity theory. Classically, the determinant has been studied for many centuries by mathematicians, physicists, numerical analysts and computer scientists.

The determinant is known to be *easy to compute*. In this respect, it enjoys a rather rare place in computation; it is an extremely useful quantity which is also efficiently computable. The classical efficient algorithms for the determinant are typically variants of the Gaussian elimination method. In last three to four decades, other approaches for computing the determinant have also been proposed. One such example is an innovative approach proposed by Mahajan and Vinay [6], which gave the first combinatorial characterization of the determinant that yielded an efficient algorithm.

In this work, we take our inspiration from this combinatorial characterization of the determinant polynomial and define two variants of the determinant which we call **StackDet_n** and **CountDet_n**. We show that they are complete for the classes VP and VNP, respectively.

The main proof idea comes from a paper of Mengel [8], which introduces characterisations of VP and VNP using *Algebraic Branching Programs (ABPs)*

² A polynomial $P_n(X)$ is said to be complete for an algebraic complexity class \mathcal{A} if $P_n(X)$ can be computed in \mathcal{A} and any polynomial $P'_m(Y)$ can be obtained from $P_n(X)$ by setting the variables in X to variables in Y or field constants. For formal definitions see Section 2

³ See also [3] for interesting variants of homomorphism polynomials.

⁴ An algebraic branching program (ABP) is a directed layered acyclic graph with a source s and a sink t . The edges are labelled with formal variables or field constants. The weight of an s to t path π is the product of the weights on the edges of π . The polynomial computed by the ABP is the sum of weights of all the s to t paths. A family f_n with $s(n)$ number of variables and degree $d(n)$ where both $s(n)$ and $d(n)$ are polynomially bounded in n is said to be in VBP iff there exist algebraic branching program of size polynomially bounded in n which computes f_n . For more details see [11].

with memory. In that work, informally speaking, it is shown that when ABPs are appended with *stack-like* memory, then they capture the class VP, and when they are appended with *counter-like* memory, they characterise the class VNP. We use these ideas and combine them with the combinatorial characterisation of the determinant to define our polynomial families.

The proof that shows that the determinant polynomial is complete for VBP can be adapted in a very straightforward way along with the *ABP with memory characterisations* of VP and VNP from the work of [8], to obtain polynomial families that are hard for these classes. However, like many other classes of polynomials (see for instance polynomial families from [10] and [7]), they are circuit-description dependent. From the work started by Durand et al. the quest has been to find circuit-description independent polynomial families complete for VP. We are able to achieve that here. The polynomial families we obtain here are circuit-description independent as desired and are variants of the determinant polynomial, which make them substantially different from the previous works [2], [5], [1].

1.1 Our Results

Before going into the details of our results, we recall the combinatorial characterisation of the determinant. Let Y be an $m \times m$ matrix, with (i, j) th entry equal to $y_{i,j}$. It is known that the determinant of Y is sum of signed cycle covers of the directed graph represented by Y . This is one of the many combinatorial definitions of the determinant, but as is, it is not known to give rise to an efficient computational procedure. Mahajan and Vinay generalized cycle covers using a notion of *clow sequences* and proved that the sum of signed clow sequences also equals the determinant. They then proved that the signed sum of clow sequences is efficiently computable.

StackDet_m and CountDet_m. We also use sum of signed clow sequences to define our polynomial. In our case, the graph has some additional edge labels. For **StackDet_m** (for **CountDet_m**), the labels come from a *stack alphabet* (*counter alphabet*, resp.). Based on these labels, we get two types of clow sequences; those which are *stack-realizable* (*counter-realizable*) and those which are not. The polynomial sums only the prior clow sequences. We show that **StackDet_m** is VP-complete and that **CountDet_m** is VNP-complete. We state our main theorem.

Theorem 1. *StackDet_n(X) is VP-complete and CountDet_n(X) is VNP-complete over any field under p-projections.*

2 Preliminaries

In this paper, a graph is always a directed graph unless specified otherwise.

Let $G = (V, E)$ be a graph where $V = [n]$. A walk $(u_1, u_2, \dots, u_{k+1})$ in G is

called a closed walk, or a clow, if $u_1 = u_{k+1}$, u_1 is the least numbered vertex in the walk and for any $2 \leq i \leq k$, $u_i \neq u_1$. The vertex u_1 is called *the head of the clow*. We use $\deg(\mathcal{C})$ to denote the number of edges in \mathcal{C} (counted with multiplicity), i.e. in this case k .

Definition 1 (A clow sequence [6]). A clow sequence $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \dots, \mathcal{C}_\ell \rangle$ in a graph $G = (V, E)$ is an ordered tuple of clows such that $\text{Head}(\mathcal{C}_1) < \text{Head}(\mathcal{C}_2) < \text{Head}(\mathcal{C}_3) < \dots < \text{Head}(\mathcal{C}_\ell)$ and $\deg(\widehat{\mathcal{C}}) = \sum_{i=1}^{\ell} \deg(\mathcal{C}_i) = n$, where $V = [n]$. The sign of a clow sequence, $\text{sign}(\widehat{\mathcal{C}})$, is $(-1)^{n+\ell}$.

Definition 2 (Stack graphs and Counter graphs). A stack graph is a directed graph $G = (V, E, \Sigma, \phi)$, where V is a set of vertices, E is a set of edges. The set Σ is a symbol set. The function ϕ labels every edge of the graph with either $\text{Push}(s)$, $\text{Pop}(s)$ for some $s \in \Sigma$ or with No-op ⁵. A counter graph $G = (V, E, \Sigma, \phi)$ is very similar to the stack graph except in this case, the function ϕ labels every edge of the graph with either $\text{Read}(s)$, $\text{Write}(s)$ for some $s \in \Sigma$ or with No-op .

We call $\text{Push}(s)$, $\text{Pop}(s)$ as the stack operations whereas $\text{Read}(s)$, $\text{Write}(s)$ as the counter operations. No-op is both a stack as well as a counter operation. Let $s_1 = [a_1, a_2, \dots, a_m]$ and $s_2 = [b_1, b_2, \dots, b_n]$ be two sequences of stack operations (or counter operations) then concatenation of s_1 followed by s_2 (denoted as $s_1 \square s_2$) is the ordered sequence $[a_1, a_2, \dots, a_m, b_1, b_2, \dots, b_n]$. It is easy to extend this definition of concatenation of two sequences to any number of sequences. Let $\mathcal{W} = (u_1, \dots, u_{k+1})$ be a walk of length k in a stack graph (or counter graph) G . We define $\text{Seq}[\mathcal{W}]$ to be the sequence of stack operations (counter operations, respectively) along the edges in this walk, i.e. $[\phi(u_1, u_2), \phi(u_2, u_3), \dots, \phi(u_k, u_{k+1})]$.

Definition 3 (Stack-realizable sequence). A stack-realizable sequence of operations is a sequence of stack operations which can be inductively formed using the following rules :

- The empty sequence is a stack-realizable sequence.
- If P is a stack-realizable sequence then $\text{Push}(s) \square P \square \text{Pop}(s)$ is stack-realizable $\forall s \in \Sigma$.
- If P is a stack-realizable sequence then $\text{No-op} \square P$ and $P \square \text{No-op}$ are also stack-realizable.
- If P and Q are stack-realizable sequences then $P \square Q$ is a stack-realizable sequence.

For example, $[\text{Push}(a), \text{Push}(b), \text{Pop}(b), \text{Push}(c), \text{No-op}, \text{Pop}(c), \text{Pop}(a), \text{No-op}]$ is a stack-realizable sequence, whereas $[\text{Push}(a), \text{Pop}(b)]$ is not.

Definition 4 (Counter-realizable sequence). A sequence of counter operations P is said to be counter-realizable if for every $s \in \Sigma$, $\text{Write}(s)$ and $\text{Read}(s)$ occur equal number of times in P and for every prefix P' of P , the number of times $\text{Write}(s)$ occurs in P' is at least as much as the number of times $\text{Read}(s)$ appears in P' .

⁵ No-op stands for No-operation

A directed walk \mathcal{W} in a stack graph (or counter graph) G is called *stack-realizable walk* (or counter-realizable walk, respectively) if and only if $\text{Seq}[\mathcal{W}]$ is stack-realizable (or counter-realizable, respectively).

Definition 5 (A realizable clog sequence). A clog sequence $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \dots, \mathcal{C}_\ell \rangle$ of a stack graph (or counter graph) G is called *stack-realizable* (or counter-realizable, respectively) if and only if $\text{Seq}[\mathcal{C}_1] \square \text{Seq}[\mathcal{C}_2] \square \dots \square \text{Seq}[\mathcal{C}_\ell]$ is a stack-realizable sequence (or counter-realizable, respectively).

Let X be a set of variables. Consider a stack graph (or a counter graph) $G = (V, E, \Sigma, \phi)$ with an edge-labeling function $\mathcal{L} : E \rightarrow X \cup \mathbb{F}$. For some $1 \leq j \leq \ell$ and clog $\mathcal{C}_j = (u_1, u_2, \dots, u_{k+1})$, $\text{mon}(\mathcal{C}_j)$ denotes monomial formed by multiplying the labels of the edges in \mathcal{C}_j , i.e. $\text{mon}(\mathcal{C}_j) = \prod_{i=1}^k \mathcal{L}((u_i, u_{i+1}))$ and $\text{mon}(\widehat{\mathcal{C}}) = \prod_{i=1}^\ell \text{mon}(\mathcal{C}_i)$.

Before going into the details of the definition of the Stack Determinant, we recall the definition of a determinant of a graph G as stated in [6].

Definition 6 (Determinant). Consider a directed graph $G = (V, E)$. Let $\mathcal{L} : E \rightarrow X \cup \mathbb{F}$, where X is the set of variables. The determinant polynomial $\text{Det}(X)$ is defined as follows : $\text{Det}(X) = \sum_{\substack{\text{All clog} \\ \text{sequences } \widehat{\mathcal{C}} \text{ of degree } |V|}} \text{sign}(\widehat{\mathcal{C}}) \text{mon}(\widehat{\mathcal{C}})$.

Definition 7 (General Stack Determinant). Consider a stack graph $G = (V, E, \Sigma, \phi)$. Let $\mathcal{L} : E \rightarrow X \cup \mathbb{F}$, where X is the set of variables. Let $\phi : E \rightarrow \bigcup_{a \in \Sigma} \{\text{Push}(a), \text{Pop}(a)\} \cup \{\text{No-op}\}$ be any edge map. The general stack determinant polynomial $\text{GenStackDet}(X)$ is defined as follows :

$$\text{GenStackDet}(X) = \sum_{\substack{\text{All stack realizable clog} \\ \text{sequences } \widehat{\mathcal{C}} \text{ of degree } |V|}} \text{sign}(\widehat{\mathcal{C}}) \text{mon}(\widehat{\mathcal{C}}).$$

In Definition 7, we now fix our graph family G_n and instantiate the function ϕ to Φ and define the stack determinant family $\text{StackDet}_n(X)$.

Definition 8 (Stack Determinant). Let $\Sigma = \{s_1, \dots, s_n\}$. Consider a stack graph $G_n = (V, E, \Sigma, \Phi)$ with $V = [4n]$ and $E = \{e_{i,j} = (i, j) | 1 \leq i, j \leq 4n\}$. Let $\mathcal{L}((i, j)) = x_{i,j}$. We define the function $\Phi : E \rightarrow \bigcup_{a \in \Sigma} \{\text{Push}(a), \text{Pop}(a)\} \cup \{\text{No-op}\}$ such that for all $i \in [n]$, $\Phi(e_{k=4(i-1)+1, k+1}) = \text{Push}(s_i)$ and $\Phi(e_{k+2, k+3}) = \text{Pop}(s_i)$. All the other remaining edges are mapped to No-op. The stack determinant polynomial $\text{StackDet}_n(X)$ is defined as follows :

$$\text{StackDet}_n(X) = \sum_{\substack{\text{All stack realizable clog} \\ \text{sequences } \widehat{\mathcal{C}} \text{ of degree } |V|}} \text{sign}(\widehat{\mathcal{C}}) \text{mon}(\widehat{\mathcal{C}})$$

Definition 9 (General Counter Determinant). Consider a counter graph $G = (V, E, \Sigma, \phi)$. Let $\mathcal{L} : E \rightarrow X \cup \mathbb{F}$, where X is the set of variables. Let $\phi : E \rightarrow \bigcup_{a \in \Sigma} \{\text{Write}(a), \text{Read}(a)\} \cup \{\text{No-op}\}$ be any edge map. The general counter determinant polynomial $\text{GenCountDet}(X)$ is defined as follows :

$$\text{GenCountDet}(X) = \sum_{\substack{\text{All counter realizable clog} \\ \text{sequences } \widehat{\mathcal{C}} \text{ of degree } |V|}} \text{sign}(\widehat{\mathcal{C}}) \text{mon}(\widehat{\mathcal{C}}).$$

In Definition 9, we now instantiate the function ϕ to Φ and redefine the counter determinant (with respect to Φ).

Definition 10 (Counter Determinant). Let $\Sigma = \{s_1, \dots, s_n\}$. Consider a counter graph $G_n = (V, E, \Sigma, \Phi)$ with $V = [4n]$ and $E = \{e_{i,j} = (i,j) | 1 \leq i, j \leq 4n\}$. Let the function $\Phi : E \rightarrow \bigcup_{a \in \Sigma} \{\text{Write}(a), \text{Read}(a)\} \cup \{\text{No-op}\}$ be such that for all $i \in [n]$, $\Phi(e_{k=4(i-1)+1, k+1}) = \text{Write}(s_i)$ and $\Phi(e_{k+2, k+3}) = \text{Read}(s_i)$. All the other remaining edges are mapped to **No-op**. Let $\mathcal{L}((i,j)) = x_{i,j}$. The counter determinant polynomial $\text{CountDet}_n(X)$ is defined as follows :

$$\text{CountDet}_n(X) = \sum_{\substack{\text{All counter realizable clow} \\ \text{sequences } \hat{\mathcal{C}} \text{ of degree } |V|}} \text{sign}(\hat{\mathcal{C}}) \text{mon}(\hat{\mathcal{C}}).$$

Definition 11. A polynomial family $\{f_n\}$ is said to be a projection of a family $\{g_n\}$, denoted as $\{f_n\} \leq \{g_n\}$, if for every f_n (where $n \in \mathbb{N}$), there exist some $m \in \mathbb{N}$ where f_n can be computed by g_m by setting the variables of g_m to either the variables of f_n or the field constants. If m is polynomially bounded in n , it is said to be a p -projection, denoted by $\{f_n\} \leq_p \{g_n\}$.

Definition 12. A p -bounded family $\{f_n\}$ is complete for class \mathcal{C} , if $f_n \in \mathcal{C}$ and for every $\{g_n\} \in \mathcal{C}$, $\{g_n\} \leq_p \{f_n\}$.

3 Upper bounds for variants of determinant family

In this section, we discuss the upperbounds of $\text{StackDet}_n(X)$ and $\text{CountDet}_n(X)$. Before going into the details of the upper bound proof, we recall the definition of the stack branching program (SBP) and the definition of the random access branching program (RABP) and the characterization of the classes VP and VNP using SBP and RABP respectively [8].

Definition 13 (SBP [8]). A stack branching program $G = (V, E)$ (over Σ) is an algebraic branching program with a function $\phi : E \rightarrow \bigcup_{a \in \Sigma} \{\text{Push}(a), \text{Pop}(a)\} \cup \{\text{No-op}\}$. The polynomial computed by G is $f_G = \sum_{\mathcal{P}} \text{mon}(\mathcal{P})$, where the sum is over all the stack-realizable s - t paths in G . The size of a stack branching program G is the number of vertices in it, i.e., $|V|$

Definition 14 (RABP [8]). A random access branching program $G = (V, E)$ (over Σ) is an algebraic branching program with an additional function $\phi : E \rightarrow \bigcup_{a \in \Sigma} \{\text{Write}(a), \text{Read}(a)\} \cup \{\text{No-op}\}$. The polynomial computed by G is $f_G = \sum_{\mathcal{P}} \text{mon}(\mathcal{P})$, where the sum is over all the counter-realizable s - t paths in G . The size of a random access branching program G is the number of vertices in it.

Lemma 1 ([8]). A family $\{f_n\}$ is in VP if and only if there exist a stack branching program family \mathcal{S}_n of size $\text{poly}(n)$ to compute $\{f_n\}$. A family $\{f_n\}$ is in VNP if and only if there exist a random access branching program family \mathcal{R}_n of size $\text{poly}(n)$ to compute $\{f_n\}$.

The upper bound proofs is motivated by the ABP upper bound for the Determinant polynomial proved by [6]. The determinant is known to be equal to the sum of signed cflow sequences. This combinatorial definition of the determinant was used in [6] to obtain an ABP upper bound. Those familiar with the proof of [6] may notice that the definitions of $\text{StackDet}_n(X)$ and $\text{CountDet}_n(X)$ are inspired by this definition of the determinant. We observe that, just like the combinatorial definition of the determinant is used to obtain an ABP upper bound in [6], our definition of $\text{StackDet}_n(X)$ and $\text{CountDet}_n(X)$ allow us to compute them using an SBP and RABP, respectively.

Construction of an SBP computing $\text{StackDet}_n(X)$: Let $G_n = (V, E, \Sigma, \Phi)$ and \mathcal{L} be as in the definition of $\text{StackDet}_n(X)$. Consider the complete directed graph $G'_n = (V, E)$, i.e G_n without the stack symbols and labels. Let A_n denote the adjacency matrix of this graph under the labelling \mathcal{L} , i.e. $A_n[i, j] = x_{i,j}$. From the result of [6], we get an ABP, say \mathcal{B}_n (of size $\text{poly}(n)$), that computes the determinant of A_n . From \mathcal{B}_n we obtain an SBP \mathcal{S}_n , by simply defining the function ϕ . We inherit ϕ from the Φ defined in the stack graph G_n as follows. Let B_n be the graph underlying the ABP \mathcal{B}_n . In B_n some edges are labelled with X variables, while some other edges are labelled with field constants. The function ϕ for all edges which are labelled with field constants is set to **No-op**. Consider any edge (p, q) in B_n that is labelled with an X variable. Suppose the edge is labelled $x_{i,j}$, then we let $\phi((p, q)) = \Phi((i, j))$.

The following statement can now be proved in a straightforward way, which finishes the proof of the upper bound.

Claim. Let $\hat{\mathcal{C}}$ be any cflow sequence in G_n . The SBP \mathcal{S}_n has a stack-realizable path from s to t with weight $\text{sign}(\hat{\mathcal{C}}) \cdot \text{mon}(\hat{\mathcal{C}})$ if and only if $\hat{\mathcal{C}}$ is a stack-realizable cflow sequence of degree $|V|$.

Proof. Let us start by recalling the construction of an ABP for the determinant polynomial from [6]. First recall that the determinant polynomial Det_n is defined as follows in [6].

$$\text{Det}_n(G'_n) = \sum_{\hat{\mathcal{C}} \text{ a cflow sequence of degree } |V|} \text{sign}(\hat{\mathcal{C}}) \text{mon}(\hat{\mathcal{C}})$$

It was shown that there exist an algebraic branching program \mathcal{B}_n (with s as the source vertex and t as the sink vertex and two special nodes t^+ and t^-) of size $\mathcal{O}(n^3)$ which computes $\text{Det}_n(G)$. The ABP \mathcal{B}_n has the following properties.

- For every cflow sequence $\hat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k \rangle$ of degree $|V|$ and positive signature, there exists a unique $s-t$ path \mathcal{P} in \mathcal{B}_n such that path \mathcal{P} is obtained by unwinding the cflows in the cflow sequence $\hat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k \rangle$ into paths, $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_k$, respectively and then stitching these paths together in order \mathcal{P}_1 followed by \mathcal{P}_2 and so on upto \mathcal{P}_k and then followed by a single edge \hat{e} labelled by $+1$ from t^+ to t . For negative signature, it is similar; except the last edge is labelled -1 and is from t^- to t .

- The variable labels on the edges (except the last edge) in $s - t$ path \mathcal{P} in \mathcal{B}_n are consistent with the variable labels on the edges in the closed walks in the clog sequence \mathcal{C} of G_n .
- There are no $s - t$ paths in \mathcal{B}_n other than the kind of paths stated above.

As the SBP \mathcal{S}_n has the same underlying graph as \mathcal{B}_n , i.e. B_n , ignoring Φ , we get a bijection between clog sequences of the stack graph G_n and s to t paths in B_n .

Stack graph S_n is obtained by specifying ϕ along with B_n . Note that the set of s to t paths in S_n and B_n continue to be the same. In S_n some paths become stack-realizable under the function ϕ . Consider a stack-realizable path \mathcal{P} in S_n . It has a corresponding clog sequence $\hat{\mathcal{C}}$ associated with it in G_n . As the labels of \mathcal{P} are consistent with those on $\hat{\mathcal{C}}$, we get that $\hat{\mathcal{C}}$ is a stack-realizable clog sequence.

Conversely, if we start with a stack-realizable clog sequence of G_n , we will find an s to t stack-realizable path in S_n . This finishes the proof. \square

Remark 1. To show that $\text{CountDet}_n(X)$ is in VNP, we can construct RABP \mathcal{R}_n which computes $\text{CountDet}_n(X)$ using arguments very similar as in the case of the construction of SBP \mathcal{S}_n computing $\text{StackDet}_n(X)$.

4 $\text{StackDet}_n(X)$ is hard for VP

In this section we prove that $\text{StackDet}_n(X)$ is VP-hard. We start by proposing two simple approaches for proving the hardness and discuss why they do not seem to work directly.

- The first way is to mimic the construction used to show that the determinant polynomial is VBP hard. Start with a stack branching program P computing f . P has designated nodes s and t . Add an extra vertex, say α , and add edges from t to α and from α to s . Also add self-loops on all the vertices of P except α . Then do the following.
 - (a) Firstly observe that the stack-realizable clog sequences of this graph can be partitioned into two sets, say \mathcal{G} and \mathcal{B} . Prove that the clog sequences in \mathcal{B} pairwise cancel each other and their weights add up to zero.
 - (b) Moreover, show that the signed clog sequences in \mathcal{G} are in one-to-one correspondence with the monomials of f .
 - (c) Finally prove that the sum of signed clog sequences in \mathcal{G} is equal to StackDet_n .

While (a) and (b) above can be proved, (c) does not seem to be true. This is because we do not have any control over the map ϕ used in P . Note that in the definition of StackDet_n Φ is a fixed map, whereas, in P , ϕ depends on the polynomial f . For instance, it is possible that stack symbols repeat themselves several times in ϕ , while in Φ they do not as per the definition. To obtain a graph *along with the Φ* as defined in StackDet_n does not seem feasible in this straightforward proof idea.

- A possible fix to the above problem is to update the definition of $\mathbf{StackDet}_n$ so that it allows for a ϕ that arises from the underlying stack branching program P that computes f . Unfortunately, that leads to polynomial families that are circuit-description dependent.

We will work on the first approach above. Our proof steps consist of the additional effort required to make this approach work. We state the main three steps in the proof outline.

Step 1 : Let \mathcal{U}_m be a universal circuit ([10, 11, 2]) of size $\text{poly}(m)$ computing an m -variate, degree $\text{poly}(m)$ polynomial $f_m(Y) \in \mathbf{VP}$. We obtain a *universal block circuit* $\tilde{\mathcal{U}}_m$, which has some more structure than \mathcal{U}_m and computes $f_m(Y)$.

Step 2 : We take the directed graph underlying the circuit $\tilde{\mathcal{U}}_m$ and transform it into another graph G_N with N vertices, where $N = \text{poly}(m)$ and $N = 4n$ for some parameter n . The graph G_N has the following properties.

- All the cycle covers of G_N have the same sign (say +ve sign w.l.o.g.).

- All the cycle covers can be classified into two categories: *good cycle covers*, say \mathcal{G} , and *bad cycle covers*, say \mathcal{B} ; and the sum of weights of the good cycle covers equals $f_m(Y)$. (We will define these notions formally below.)

Step 3 : From G_N we obtain a stack graph H_N with the following properties. The set of stack-realizable cflow sequences in H_N which are cycle covers, equals \mathcal{G} . Moreover, the sum of signed weights of stack-realizable cycle covers in H_N equals the sum of signed weights of cycle covers in \mathcal{G} , i.e. equal to $f_m(Y)$. and the sum of signed weights of stack-realizable cflow sequences that are not cycle covers equals 0. Overall, the sum of signed weights of stack-realizable cflow sequences of H_N equals $f_m(Y)$.

We can now interpret H_N as a complete graph, where $\mathcal{L}((i, j)) = 0$ if (i, j) is not an edge in H_N . We will show that the polynomial $\mathbf{StackDet}_n$ defined with respect to H_N equals $f_m(Y)$.

The Step 1 and 2 above are obtained using the ideas of Block Trees from [1]. Step 3 above uses the cancellation trick from [6], but now in the context of stack-realizable cflow sequences (instead of cflow sequences) and with respect to an SBP (instead of an ABP).

4.1 VP-hardness of $\mathbf{StackDet}_n(X)$ Step 1

Recall that from the constructions in [10, 11, 2], we can assume the following properties about the universal circuit. The circuit \mathcal{U}_m has m variables, size $s(m)$ and each even layer is a + gate, while each odd layer is a \times gate. The output gate is a \times gate. The \times gates are multiplicatively disjoint⁶ and have fan-in bounded by 2. The input gates have fanin 0, fanout 1. The total depth⁷ of the circuit

⁶ A multiplication gate α with children gates α_ℓ and α_r in an arithmetic circuit C is called multiplicatively disjoint if the subcircuits rooted at α_ℓ and α_r are disjoint.

⁷ The depth of the circuit is the length of the longest input gate to output gate path.

is $2c\lceil \log m \rceil + 1$, where c is some fixed constant. Say it computes a polynomial $f_m(Y)$ of degree $\text{poly}(m)$ ⁸.

We now create a circuit $\tilde{\mathcal{U}}_m$, which will have the same depth, each even layer will again consist of $+$ gates and each odd layer of \times gates. It will continue to be multiplicatively disjoint and its size will be $\text{poly}(s(m))$. It is created as follows:

Block structure. In the j th layer of $\tilde{\mathcal{U}}_m$ we create $t(j) = 2^{\lfloor \frac{j}{2} \rfloor}$ many blocks. The blocks on the j layer are denoted by $B_1^{(j)}, B_2^{(j)}, \dots, B_{t(j)}^{(j)}$.

Gates. If j is odd - Let g_1, \dots, g_r be the \times gates appearing in \mathcal{U}_m in layer j . In $\tilde{\mathcal{U}}_m$, each block B has one copy of g_1, \dots, g_r .

If j is even - Let g_1, \dots, g_r be the $+$ gates appearing in \mathcal{U}_m in layer j . Each block B in j th layer in $\tilde{\mathcal{U}}_m$ has $s(m)$ sub-blocks. Each sub-block has one copy of g_1, \dots, g_r . (That is, there are $s(m)$ copies of each gate in each block and there are $t(j)$ many blocks. So each gate is copied $t(j) \cdot s(m)$ times. Note that this is polynomially bounded in $\text{poly}(m)$.)

Wires: Let g be a $+$ gate in layer j with children g_1, g_2, \dots, g_r in \mathcal{U}_m . Then the copy of g in $B_i^{(j)}$ has copies of g_1, \dots, g_r from block $B_i^{(j+1)}$ as its children for each $i \in t(j)$. Let g be a \times gate in layer j with children $g_{\text{left}}, g_{\text{right}}$ in \mathcal{U}_m . Also among the different gates that use g_{left} , let g be the k th such gate. Then (the unique) copy of g in $B_i^{(j)}$ has k th copy of g_{left} from block $B_{2i-1}^{(j+1)}$ as its child. Similarly, among the gates that use g_{right} , let g be the k' th such gate. Then the copy of g in $B_i^{(j)}$ has k' th copy of g_{right} from block $B_{2i}^{(j+1)}$ as its child. Finally, we only keep the minimal circuit, i.e. we remove gates that eventually do not feed into the output gate.

This completes the description of $\tilde{\mathcal{U}}_m$. The construction is exactly the same as the construction of D'_n in [1]. We call this the universal block circuit.

Claim. The polynomial computed by $\tilde{\mathcal{U}}_m$ is $f_m(Y)$ and the size of the circuit is polynomial in $s(m)$, say $p(m)$, which in turn is polynomial in m .

We skip the proof details (See [9] for proof details).




Fig. 1. \mathcal{U}_m computing $(x_1x_2 + x_1x_3)(x_2x_3)$ and the corresponding $\tilde{\mathcal{U}}_m$

⁸ This description is slightly different as compared to the one in [2], but it is easy to see that we can get this form for a universal circuit using ideas from [10].

4.2 VP-hardness of $\text{StackDet}_n(X)$ Step 2

We now consider the graph underlying the universal block circuit created in Step 1. We direct all the edges in this graph from top (i.e. from the output gate) to bottom (to the input gates). The top-most layer has a single vertex, which is the output gate. Each layer j has $t(j)$ -many blocks. We denote this directed graph by $V_{p(m)}$, where $p(m)$ is the number of vertices in this graph. We take two views of this underlying graph; a *coarse* view and a *fine* view. The fine view is simply the whole graph $V_{p(m)}$, while the coarse view is the graph formed by the block structure.

Block Tree. For the coarse view, we think of each block of $V_{p(m)}$ as a vertex. We call these *block vertices*. Two blocks vertices B, B' are said to be connected if and only if $\exists u \in B$ and $v \in B'$ such that there is an edge between u and v in $V_{p(m)}$. We refer to (B, B') as a *block edge*. By observing the connections in $V_{p(m)}$, it is easy to see that the coarse view results into a tree. We call this tree $T_{\Delta(m)}$, where $\Delta(m)$ denotes the number of leaf nodes in the tree. Let B be a vertex in T_{Δ} . If B is on an even layer, then it has only one child. We call these the *unary blocks*. If it is on an odd layer then it has two children. We call these blocks *binary blocks*. A path formed by block edges is called a *block path*.

When m is clear from the context, we use V_p and T_{Δ} to talk about these two graphs.

Construction of G_N .

- For any binary block B and any vertex $u \in B$, we do the following. Let B_ℓ and B_r be the two children of B in T_{Δ} . Let $u_\ell \in B_\ell$ and $u_r \in B_r$ such that (u, u_ℓ) and (u, u_r) are edges in V_p . We sub-divide the edge (u, u_r) into (u, z_u) and (z_u, u_r) . We delete the edge (u, z_u) from the graph, but retain the edge (z_u, u_r) . For any node u in a binary block, we use $\text{Couple}(u)$ to denote the pair of edges $\{(u, u_\ell), (z_u, u_r)\}$. ($\text{Couple}(u)$ is not defined for a u in a unary block.)
- Note that this creates a new graph which is disconnected. If we look at the coarse view of this new graph then it is a collection of Δ block paths, let us call them $\mathcal{P}_1, \dots, \mathcal{P}_{\Delta}$. Each block path contains exactly one leaf node of T_{Δ} . We will assume that the block paths are numbered such that the i th leaf node of T_{Δ} belongs to \mathcal{P}_i .
- We add two more vertices for each block path. We add a source vertex s_i and a sink vertex t_i for each $i \in [\Delta]$. We also add edges from s_i to all the vertices in the first block in the block path \mathcal{P}_i . The vertices in the last block in any block path are vertices corresponding to input gates in \tilde{U}_m and hence are labelled with input variables Y . Let u be a vertex in the leaf block of the path \mathcal{P}_i labelled $y \in Y$. We add a directed edge (u, t_i) and label it with y . (We do this for each vertex in every leaf block of all block paths.) The graphs thus obtained are called $\mathcal{R}_1, \dots, \mathcal{R}_{\Delta}$.
 - We now identify t_i with s_{i+1} for $1 \leq i \leq \Delta - 1$. We use \mathcal{R} to denote the graph thus formed and θ_i to denote the vertex formed by identifying t_i with s_{i+1} for $1 \leq i \leq \Delta - 1$. Additionally, we want to ensure that the number of vertices in the resultant graph is a multiple of 4 (This will help

in defining a stack graph in the next step). To ensure this, we add three⁹ additional vertices $\alpha_1, \alpha_2, \alpha_3$ and the following directed edges to obtain a graph D_N : $(t_\Delta, \alpha_3), (\alpha_3, \alpha_2), (\alpha_2, \alpha_1), (\alpha_1, s_1)$. We add self-loops on all the vertices except on α_1, α_2 and α_3 . The edges which are not labelled with variables from Y are labelled 1.

The graph thus obtained is denoted by G_N , where N is the number of vertices in it. It is easy to note that $N = \text{poly}(p(m))$ which is $\text{poly}(m)$. We have also ensured that $N = 4n$ for some parameter n .

Definition 15. We say that a cycle cover $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \dots, \mathcal{C}_k \rangle$ of G_N is a good cycle cover if for any vertex u appearing in $\widehat{\mathcal{C}}$ for which $\text{Couple}(u)$ is defined, either both the edges in $\text{Couple}(u)$ are present in $\widehat{\mathcal{C}}$ or neither is. All the other cycle covers are called bad cycle covers. Let \mathcal{G} denote the set of all good cycle covers of G_N and \mathcal{B} denote the set of all the bad cycle covers.

Claim. All the cycle covers of G_N have the same sign. Moreover, the sum of weights of good cycle covers equals $f_m(Y)$.

Proof. Recall the graphs $\mathcal{R}_1, \dots, \mathcal{R}_\Delta$ that we created from $\mathcal{P}_1, \dots, \mathcal{P}_\Delta$.

Consider any path π from s_i to t_i in \mathcal{R}_i . The first edge of π must be from s_i to a vertex belonging to the first block, and the last edge of π must be from a vertex belonging to the last block to the vertex t_i . All intermediate edges must connect adjacent blocks. So, the number of edges in π is one more than the number of blocks in \mathcal{R}_i . Therefore all paths from s_i to t_i in \mathcal{R}_i have the same number of edges, say p_i .

Consider any path Π from s_1 to t_Δ . For any $2 \leq i \leq \Delta$, the vertex s_i must belong to Π (because deleting s_i disconnects the graph into two components, where s_1 and t_Δ belong to different components). This means Π can be viewed as a composition of the paths $\pi_1, \pi_2, \dots, \pi_\Delta$, where π_i is a path from s_i to t_i for all $1 \leq i \leq \Delta$. This path π_i is also a path in \mathcal{R}_i , so it has length p_i . Therefore the path Π has length $p_1 + p_2 + \dots + p_\Delta$, which we call q , say. In all, any path from s_1 to t_Δ has the same length q .

Let $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k \rangle$ be a cycle cover of G_N , and consider a cycle of the cycle cover $\widehat{\mathcal{C}}$ that α_1 belongs to, say \mathcal{C}_1 . The only incoming edge to α_1 is via t_Δ , and the only edge outgoing from α_1 is to α_2 . This means the edges (t_Δ, α_1) and (α_1, α_2) belong to \mathcal{C}_1 . The only outgoing edge from α_2 is to α_3 , and the only outgoing edge from α_3 is to s_1 . Therefore, the edges (α_2, α_3) and (α_3, s_1) also belong to \mathcal{C}_1 . So, \mathcal{C}_1 contains a path from t_Δ to s_1 via α_1, α_2 and α_3 . The remaining part of \mathcal{C}_1 is a path from s_1 to t_Δ . This path does not use the vertices α_1, α_2 , and α_3 , so it is also a path in \mathcal{R} . As shown before, any such path from s_1 to t_Δ has length $q = p_1 + p_2 + \dots + p_\Delta$. Therefore \mathcal{C}_1 is a cycle of length $q + 4$.

Consider a cycle $\mathcal{C}_j \neq \mathcal{C}_1$ in the cycle cover $\widehat{\mathcal{C}}$. This cycle cannot use the vertices α_1, α_2 and α_3 . Furthermore, if \mathcal{C}_j is not a loop, then it is a cycle in \mathcal{R} ,

⁹ As Δ is a power of 2, it is easy to note that adding three new vertices will always make the total number of vertices of graph G_N a multiple of 4.

which contradicts the fact that \mathcal{R} is a DAG. Therefore \mathcal{C}_j is a loop. In all, the cycle cover $\widehat{\mathcal{C}}$ has exactly one cycle \mathcal{C}_1 of length $q+4$ passing through α_1, α_2 , and α_3 , and $N - q - 4$ loops covering the vertices not present in the cycle \mathcal{C}_1 . Either way, the sign of any cycle cover $\widehat{\mathcal{C}}$ is fixed. It is also easy to see from the above discussion that there is a one-to-one correspondence between a path Π from s_1 to t_Δ in \mathcal{R} and cycle covers of G_N .

We will now show that the good cycle covers of G_N have a one-to-one correspondence with the parse trees of $\tilde{\mathcal{U}}_m$.

Let \mathcal{T} be any parse tree of $\tilde{\mathcal{U}}_m$. For any vertex u corresponding to a \times gate of $\tilde{\mathcal{U}}_m$, such that u_l is the left child and u_r is the right child of u in \mathcal{T} , split the edge (u, u_r) into (u, z_u) and (z_u, u_r) and delete edge (u, z_u) . This splits \mathcal{T} into Δ paths $Q_1, Q_2, \dots, Q_\Delta$, where Q_i belongs to \mathcal{R}_i for each $i \in [\Delta]$. These Δ paths (when concatenated appropriately) trace out a path Π in D_N . This path can be completed into a cycle \mathcal{C}_1 in G_N . This cycle \mathcal{C}_1 along with self-loops on all the other vertices outside of \mathcal{C}_1 , forms a cycle cover $\widehat{\mathcal{C}}$ of G_N . Note that, the way this cycle cover was created, for each u in a binary block of V_p , either both edges of $\text{Couple}(u)$ are present in $\widehat{\mathcal{C}}$ or neither edge of $\text{Couple}(u)$ is present in $\widehat{\mathcal{C}}$. Therefore $\widehat{\mathcal{C}}$ is a good cycle cover. It is easy to see that the cycle cover has weight equal to the monomial computed by \mathcal{T} in $\tilde{\mathcal{U}}_m$.

For the converse, we show that a good cycle cover of G_N can be traced back to a unique parse tree of $\tilde{\mathcal{U}}_m$. Let $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k \rangle$ be a good cycle cover of G_N . Let \mathcal{C}_1 be the big cycle and the rest of the cycles in the cover be self-loops. Let E_1 denote the edges that \mathcal{C}_1 shares with graphs $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_\Delta$. As this is a good cycle cover, for each vertex u in \mathcal{C}_1 for which $\text{Couple}(u)$ is defined, edges (u, u_ℓ) and (z_u, u_r) are both present in \mathcal{C}_1 . We will identify z_u with u for all such vertices. This will give rise to a unique parse tree of $\tilde{\mathcal{U}}_m$. \square

Remark 2. To be able to sum over only the good cycle covers, we need a mechanism to ensure that either both or none of the edges in $\text{Couple}(u)$ are selected in any cycle cover $\widehat{\mathcal{C}}$. In Valiant's work [12] for instance, this is ensured by using an *iff graph gadget*. If we can come up with such a gadget (in the determinant setting) then we will be able to show that Det_n is VP-complete, thereby showing $\text{VP} = \text{VBP}$. We ensure *coupling* using the stack symbols.

4.3 VP-hardness of $\text{StackDet}_n(X)$ Step 3

We would like to modify the graph G_N so that we filter out good cycle covers, while killing all the bad cycle covers. We achieve this using stack symbols. Specifically, we create a stack graph H_N from G_N to achieve this.

Construction of H_N . For a vertex u for which $\text{Couple}(u)$ is defined, we set $\phi((u, u_\ell)) = \text{Push}(s_u)$ and $\phi((z_u, u_r)) = \text{Pop}(s_u)$. For all the other edges, ϕ is set to **No-op**.

Claim. Consider the stack graph H_N constructed as above.

- The sum of signed weights of stack-realizable cycle covers in H_N equals the




Fig. 2. Graphs $\mathcal{P}_1, \dots, \mathcal{P}_\Delta$ and $\mathcal{R}_1, \dots, \mathcal{R}_\Delta$.

signed sum of weights of cycle covers in \mathcal{G} , i.e. equal to $f_m(Y)$.

- The sum of signed weights of stack-realizable cflow sequences that are not cycle covers equals 0.

Proof. Part 1. From the proof of Claim in Section 4.2, we have that there is a bijection between parse trees of \tilde{U}_m and good cycle covers of G_N . To prove the first part of the claim, we will show that there is a bijective map from a good cycle covers of G_N to stack-realizable cycle covers of H_N .

We start with some notations. Let $\hat{\mathcal{C}}$ be a good cycle cover in G_N . Let $\mathcal{T}_{\hat{\mathcal{C}}}$ be the unique parse tree corresponding to $\hat{\mathcal{C}}$. Let $\hat{\mathcal{C}} = \langle \mathcal{C}_1, \dots, \mathcal{C}_k \rangle$ and \mathcal{C}_1 be the long cycle, while all other \mathcal{C}_i s be self-loops. (Any good cycle cover has this structure as we established in the proof of Claim in Section 4.2.) Let $U_{\hat{\mathcal{C}}} = \{u_1, \dots, u_\tau\}$ be the subset of vertices in \mathcal{C}_1 for which **Couple** is defined. Note that the output gate, let us call it u^* , of $\mathcal{T}_{\hat{\mathcal{C}}}$ belongs to $U_{\hat{\mathcal{C}}}$.

We say that a vertex $u \in U_{\hat{\mathcal{C}}}$ has **rank** k , denoted as $\mathbf{rank}(u)$, if it appears at distance $2k - 1$ from the leaves in $\mathcal{T}_{\hat{\mathcal{C}}}$. (Note that, vertices in $U_{\hat{\mathcal{C}}}$ appear at only odd distance from the leaves in $\mathcal{T}_{\hat{\mathcal{C}}}$.)

For $u \in U_{\hat{\mathcal{C}}}$ such that $\mathbf{rank}(u) = 1$, u_ℓ and u_r are leaves, i.e. nodes corresponding to input gates. For a vertex u in $U_{\hat{\mathcal{C}}}$ such that $\mathbf{rank}(u) > 1$, let u_ℓ and u_r be its two children in $\mathcal{T}_{\hat{\mathcal{C}}}$. Let u' be u_ℓ 's unique child in $\mathcal{T}_{\hat{\mathcal{C}}}$ and let u'' be the unique child of u_r in $\mathcal{T}_{\hat{\mathcal{C}}}$. Note that $u', u'' \in U_{\hat{\mathcal{C}}}$ and $\mathbf{rank}(u') = \mathbf{rank}(u'') = \mathbf{rank}(u) - 1$.

Let $\Pi_{\hat{\mathcal{C}}}$ be the unique path traced out by \mathcal{C}_1 in \mathcal{R} . (Recall, \mathcal{R} is the graph obtained by concatenating \mathcal{R}_i for $i \in [\Delta]$ as described in the construction.)

For a vertex $u \in U_{\hat{\mathcal{C}}}$, such that $\mathbf{rank}(u) = 1$, we use $\Pi_{[u]}$ to denote the subpath of $\Pi_{\hat{\mathcal{C}}}$ from u to u_r . Given the structure of the subtree rooted at u in $\mathcal{T}_{\hat{\mathcal{C}}}$, and assuming that u_r appears in \mathcal{R}_{i+1} for some $i \in [\Delta - 1]$, we get that $\Pi_{[u]} = (u, u_\ell) \cdot (u_\ell, \theta_i) \cdot (\theta_i, z_u) \cdot (z_u, u_r)$. (Recall that θ_i was the vertex obtained by identifying t_i of \mathcal{R}_i with s_{i+1} of \mathcal{R}_{i+1} for $i \in [\Delta - 1]$.)

On the other hand, for $u \in U_{\hat{\mathcal{C}}}$ and $\mathbf{rank}(u) > 1$ such that u_r appears in \mathcal{R}_{i+1} for some $i \in [\Delta - 1]$, we use $\Pi_{[u]}$ to denote the subpath of Π corresponding to the entire subtree rooted at u in $\mathcal{T}_{\hat{\mathcal{C}}}$. Specifically, for the given structure of the subtree rooted at u in $\mathcal{T}_{\hat{\mathcal{C}}}$, $\Pi_{[u]} = (u, u_\ell) \cdot (u_\ell, u') \cdot \Pi_{[u']} \cdot (\theta_i, z_u) \cdot (z_u, u_r) \cdot$

$(u_r, u'') \cdot \Pi_{[u'']}$. We will now prove the following statement.

$$\text{For any } u \in U_{\widehat{C}}, \text{Seq}[\Pi_{[u]}] \text{ is stack-realizable in } H_N. \quad (1)$$

If we are able to show this, then in particular for $u^* \in U_{\widehat{C}}$ we will get that $\Pi_{[u^*]}$ is stack-realizable. This will then imply that $(s_1, u^*) \cdot \Pi_{[u^*]} \cdot (\theta_\Delta, t_\Delta)$ is also stack-realizable, because both (s_1, u^*) and $(\theta_\Delta, t_\Delta)$ are **No-op** edges.

We prove (1) by induction on $\text{rank}(u)$. Suppose $\text{rank}(u) = 1$ and say $u_r \in \mathcal{R}_{i+1}$, then as noted above, $\Pi_{[u]} = (u, u_\ell) \cdot (u_\ell, \theta_i) \cdot (\theta_i, z_u) \cdot (z_u, u_r)$. From our function ϕ defined for H_N , we see that $\text{Seq}[\Pi_{[u]}] = \text{Push}(s_u) \square \text{No-op} \square \text{No-op} \square \text{Pop}(s_u)$. Therefore it is stack-realizable.

Suppose $\text{rank}(u) = k > 1$ and say that $u_r \in \mathcal{R}_{i+1}$. In this case, as noted above, we have $\Pi_{[u]} = (u, u_\ell) \cdot (u_\ell, u') \cdot \Pi_{[u']} \cdot (\theta_i, z_u) \cdot (z_u, u_r) \cdot (u_r, u'') \cdot \Pi_{[u'']}$. From this, we see that $\text{Seq}[\Pi_{[u]}] = \text{Push}(s_u) \square \text{No-op} \square \text{Seq}[\Pi_{[u']}] \square \text{No-op} \square \text{Pop}(s_u) \square \text{No-op} \square \text{Seq}[\Pi_{[u'']}]$. As $\text{rank}(u')$, $\text{rank}(u'') < k$, by induction hypothesis we have that $\text{Seq}[\Pi_{[u']}]$ and $\text{Seq}[\Pi_{[u'']}]$ are stack-realizable. Therefore, we get that $\text{Seq}[\Pi_{[u]}]$ is also stack-realizable.

It is not hard to argue that bad cycle covers of G_N get mapped to cycle covers of H_N , which are not stack-realizable by a similar argument.

Part 2 Recall that in the proof of Claim in Section 4.2 we showed that any cycle cover of G_N consists of one big cycle and a collection of self-loops. Similarly, it is easy to see that in H_N any clow sequence has a certain structure: except for one clow, which will be of length $\geq p + 4$, all other clows in the clow sequence are self-loops.

We first note that this unique long clow will contain the vertex α_1 . Suppose it does not contain α_1 , then no other clow in the clow sequence can cover α_1 (as all other clows are self-loops and α_1 does not have a self-loop on it). But suppose α_1 is not covered by any clow in the clow sequence, then the degree of such a clow sequence is strictly less than $|V|$.

Under the ordering in which vertex α_1 gets the lowest number, say 1, the long clow will be the first clow in the sequence, say \mathcal{C}_1 and α_1 will be its head.

We will adopt ideas from [6] in order to argue that the sum of weights of stack-realizable clow sequences which are not cycle covers is 0 in H_N . Like in [6], we define an involution on the signed clow sequences. (Recall that an involution is a bijective map ψ such that ψ^2 is identity.) The map ψ will have the property that any stack-realizable clow sequence \widehat{C} which is not a cycle cover, is paired off with another stack-realizable clow sequence \widehat{C}' which is again not a cycle cover and the monomials corresponding to \widehat{C} and \widehat{C}' are the same, but $\text{sign}(\widehat{C}') = -\text{sign}(\widehat{C})$. For clow sequence \widehat{C} which is a cycle cover, the map ψ maps it to itself, i.e. it is identity for cycle covers.

Let \widehat{C} be a stack-realizable clow sequence, which is not a cycle cover. We start walking along the edges of \mathcal{C}_1 starting from the head. One of the following two cases will happen first.

- **Case 1.** Either we will encounter a vertex v in \mathcal{C}_1 such that there exists a $\mathcal{C}_i \in \widehat{C}$ for $i > 1$, such that \mathcal{C}_i is a self-loop at vertex v .

- **Case 2.** Or we will encounter a vertex u that has $\beta \geq 1$ self-loops in \mathcal{C}_1 .

First note that, if $\widehat{\mathcal{C}}$ is not a cycle cover then one of the two cases must occur.

Suppose Case 1 occurs. In this case, consider $\widehat{\mathcal{C}}'$ obtained from $\widehat{\mathcal{C}}$ by merging cycle \mathcal{C}_i with \mathcal{C}_1 , by attaching it at v in \mathcal{C}_1 . We will define ψ of $\widehat{\mathcal{C}}$ to be this $\widehat{\mathcal{C}}'$. It is easy to see that if $\widehat{\mathcal{C}}$ is a stack-realizable cflow sequence, then so is $\widehat{\mathcal{C}}'$. Both have the same set of edges. And $\widehat{\mathcal{C}}'$ has one less component than $\widehat{\mathcal{C}}$, i.e. their signs are opposite.

On the other hand, suppose Case 2 occurs. In this case, consider $\widehat{\mathcal{C}}'$ obtained from $\widehat{\mathcal{C}}$ by detaching one of the β -many self-loops from u and adding that as a separate cycle in $\widehat{\mathcal{C}}'$. To observe that $\widehat{\mathcal{C}}'$ thus obtained is a stack-realizable cflow sequence, we first note that there is no other cflow in $\widehat{\mathcal{C}}'$ with the same head as this newly added self-loop. This is easy to see, because if say there was already a cflow in $\widehat{\mathcal{C}}'$ with u as its head, then we would have been in Case 1 above.

We also observe that if $\widehat{\mathcal{C}}$ is stack-realizable, then detaching a self-loop, which is a **No-op** edge, will ensure that $\widehat{\mathcal{C}}'$ is also stack-realizable. Here again, $\widehat{\mathcal{C}}$ and $\widehat{\mathcal{C}}'$ have the same set of edges and $\widehat{\mathcal{C}}$ has one less component than $\widehat{\mathcal{C}}'$, i.e. they have opposite signs.

Note that in both the cases above, if $\psi(\widehat{\mathcal{C}}) = \widehat{\mathcal{C}}'$ then $\psi(\widehat{\mathcal{C}}') = \widehat{\mathcal{C}}$. Hence, we have the desired involution. \square

If we now show an ordering of the vertices of graph H_N such that the function ϕ in graph H_N can be exactly mapped to the function Φ as defined in Definition 8 then the VP-hardness immediately follows. It is easy to show that such an ordering always exist.

Ordering of the vertices of graph H_N We will now show that there is an ordering of the vertices of H_N , which gives a graph $G_n = (V, E, \Sigma, \Phi)$ as defined in Definition 8 and a labelling function \mathcal{L} as defined in Definition 8, such that $f_m(Y)$ can be obtained as a projection of $\mathbf{StackDet}_n(X)$ defined with respect to G_n , which finishes the proof.

We now come up with such an ordering. We start by ordering vertices $\theta_1, \dots, \theta_{\Delta-1}$ and α_1, α_2 and α_3 . Note that these vertices must appear in any cycle, which is not a self-loop. If we start traversing any such cycle from α_1 , then we will visit these vertices in the following order $\langle \alpha_1, s_1, \theta_1, \dots, \theta_{\Delta-1}, t_{\Delta}, \alpha_3, \alpha_2, \alpha_1 \rangle$. We number these vertices in the reverse order, i.e. α_1 gets numbered 1, α_2 gets 2, α_3 is numbered 3, t_{Δ} is numbered 4 and so on till s_1 is numbered $\Delta + 4$. This numbering ensures that all the edges that appear between these vertices get **No-op** label on them.

Now, let $u_1, u_2, \dots, u_{\tau}$ be the vertices for which **Couple** is defined. Let $\mathbf{Couple}(u_i) = \{(u_i, u_{i\ell}), (z_{u_i}, u_{ir})\}$. For every $i \in [\tau]$, let the four vertices $u_i, u_{i\ell}, z_{u_i}, u_{ir}$ be numbered as $4(i-1) + 1 + (\Delta + 4)$, $4(i-1) + 2 + (\Delta + 4)$, $4(i-1) + 3 + (\Delta + 4)$ and $4(i-1) + 4 + (\Delta + 4)$ respectively¹⁰. It is easy to check that such an ordering always gives distinct numbers to all the vertices of the graph and this ordering is consistent with Φ from Definition 8.

¹⁰ It is not too hard to see that $\Delta + 4$ is a multiple of 4. (As Δ is a power of 2.)

The labelling function \mathcal{L} retains the labels of all the edges of H_N as they are. For any two vertices u, v in H_N , such that there is no edge in H_N between u and v , we add such an edge in G_n , but set $\mathcal{L}((u, v)) = 0$. This labelling function now ensures that when we consider the StackDet_n polynomial with respect to G_n we obtain $f_m(Y)$.

5 VNP-Hardness of the $\text{CountDet}_n(X)$

In this section we first show that $\text{CountDet}_n(X)$ is hard for VNP. We show the VNP-hardness in two cases ¹¹. We will first show that the permanent polynomial ¹² can be computed as a projection of $\text{CountDet}_n(X)$. This will prove that $\text{CountDet}_n(X)$ is VNP-hard over fields of characteristic $\neq 2$. To show it's hardness over fields of characteristic 2, we will show that it can compute another polynomial, namely EC_m^* , as a projection, where $n = \text{poly}(m)$. This polynomial was shown to be VNP-complete over fields of characteristic 2 in [4].

5.1 Details regarding VNP-hardness of $\text{CountDet}_n(X)$ when $\text{char} \neq 2$

We will first show that the Permanent polynomial can be computed as a projection of $\text{CountDet}_n(X)$. This will prove that $\text{CountDet}_n(X)$ is VNP-hard over fields of $\text{char} \neq 2$.

Let $Y = \{y_{1,1}, y_{1,2}, \dots, y_{m,m}\}$. We will show that $\text{Perm}_m(Y)$ can be obtained as a projection of $\text{CountDet}_n(X)$, where $n = \text{poly}(m)$. To prove this, we create a counter graph H_N , such that $N = \text{poly}(m)$ and the following properties hold.

- All the counter-realizable cycle covers in H_N have the same sign.
- Moreover, the sum of the weights of the counter-realizable clog sequences which are cycle covers, equals Perm_m and the sum of the signed weights of the counter-realizable clog sequences which are not cycle covers = 0.

Then by simple re-ordering of the vertices of H_N and adding edges to make it a complete graph G_n , as in the definition of $\text{CountDet}_n(X)$, we get that $\text{Perm}_m(Y)$ can be obtained as a projection of $\text{CountDet}_n(X)$.

In order to describe the construction of H_N , we first create $2m$ smaller counter graphs, W_1, \dots, W_m and R_1, \dots, R_m . For each $i \in [m]$, $W_i = (V_i^w, E_i^w, \Sigma_i^w, \phi_i^w)$ is as follows.

$$- V_i^w = \{s_i^w, t_i^w\} \cup \{u_{i,1}, \dots, u_{i,m}\} \cup \{v_{i,1}, \dots, v_{i,m}\}. E_i^w = \bigcup_{j \in [m]} \{(s_i^w, u_{i,j})\} \cup \bigcup_{j \in [m]} \{(v_{i,j}, t_i^w)\} \cup \bigcup_{j \in [m]} \{(u_{i,j}, v_{i,j})\}. \Sigma_i^w = \{\alpha_{i,1}, \alpha_{i,2}, \dots, \alpha_{i,m}\}.$$

¹¹ Note that the VNP-hardness can also possibly be shown in a single step via constructing a RABP \mathcal{R}_n which computes a polynomial family \mathcal{P}_n (known to be hard for VNP for all fields) and converting it to a counter graph G_n such that CountDet_n defined over G_n computes \mathcal{P}_n . However, constructing an RABP \mathcal{R}_n (and converting it then to a counter graph G_n) such that the function ϕ in G_n getting exactly mapped to the function Φ defined in Definition 10 is not immediate. Therefore, we show the hardness in two steps.

¹² Recall that $\text{Perm}_m(Y) = \sum_{\sigma: \text{permutation of } [m]} \prod_{i \in [m]} y_{i, \sigma(i)}$.

- For each $j \in [m]$, $\phi_i^w((u_{i,j}, v_{i,j})) = \mathbf{Write}(\alpha_{i,j})$. ϕ_i^w is **No-op** for all other edges in E_i^w .

Similarly, for each $i \in [m]$, $R_i = (V_i^r, E_i^r, \Sigma_i^r, \phi_i^r)$ can be described as follows.

- $V_i^r = \{s_i^r, t_i^r\} \cup \{a_{i,1}, \dots, a_{i,m}\} \cup \{b_{i,1}, \dots, b_{i,m}\}$. $E_i^r = \bigcup_{j \in [m]} \{(s_i^r, a_{i,j})\} \cup \bigcup_{j \in [m]} \{(b_{i,j}, t_i^r)\} \cup \bigcup_{j \in [m]} \{(a_{i,j}, b_{i,j})\}$. $\Sigma_i^r = \{\alpha_{1,i}, \alpha_{2,i}, \dots, \alpha_{m,i}\}$. For each $j \in [m]$, $\phi_i^r((a_{i,j}, b_{i,j})) = \mathbf{Read}(\alpha_{j,i})$. ϕ_i^r is **No-op** for all other edges in E_i^r .

Let H'_N be the graph formed by identifying t_i^w with s_{i+1}^w for $1 \leq i \leq m-1$ and by identifying t_m^w with s_1^r and also identifying t_i^r with s_{i+1}^r for $1 \leq i \leq m-1$. We also add labels on the edges of H'_N . We define $\mathcal{L}((u_{i,j}, v_{i,j})) = y_{i,j}$ for $i, j \in [m]$. For all other edges, \mathcal{L} is set to 1. We first make the following observation about H'_N .

Claim. For each monomial in \mathcal{M} in $\mathbf{Perm}_m(Y)$, there is a unique counter-realizable path π from s_1^w to t_m^r in H'_N such that $\prod_{e \in \pi} \mathcal{L}(e) = \mathcal{M}$. For any counter-realizable path π from s_1^w to t_m^r in H'_N , $\prod_{e \in \pi} \mathcal{L}(e)$ corresponds to a unique monomial of $\mathbf{Perm}_m(Y)$.

Proof. From our construction of H'_N , it is easy to see that the vertices $t_1^w, t_2^w, \dots, t_m^w$ and the vertices $t_1^r, t_2^r, \dots, t_{m-1}^r$ are all cut-vertices in H'_N , and deleting any one of them disconnects the vertices s_1^w and t_m^r . This means any path π from s_1^w to t_m^r passes through the vertices t_i^w for $1 \leq i \leq m$ and t_i^r for $1 \leq i \leq m-1$. Therefore, π can be viewed as a composition of the $2m$ paths $\pi_1^w, \pi_2^w, \dots, \pi_m^w, \pi_1^r, \pi_2^r, \dots, \pi_m^r$ in that order, where π_i^w is the subpath of π between s_i^w and t_i^w for $1 \leq i \leq m$, and π_i^r is the subpath of π between s_i^r and t_i^r for $1 \leq i \leq m$. In fact, for any such $2m$ paths, their composition (in that order) is a path from s_1^w to t_m^r in H'_N .

We now proceed with the proof of the claim. Any monomial \mathcal{M} in $\mathbf{Perm}_m(Y)$ is of the form $\prod_{i=1}^m y_{i,\sigma(i)}$, where σ is a permutation of $[m]$. The path π is constructed as follows: take π_i^w to be the path $s_i^w, u_{i,\sigma(i)}^w, v_{i,\sigma(i)}^w, t_i^w$, and π_i^r to be the path $s_i^r, u_{i,\sigma^{-1}(i)}^r, v_{i,\sigma^{-1}(i)}^r, t_i^r$, both for $1 \leq i \leq m$. Consider the sequence of counter operations along π , other than the **No-op** operations. The only edges that have such operations are the edges $(u_{i,\sigma(i)}^w, v_{i,\sigma(i)}^w)$ for $1 \leq i \leq m$ and $(u_{i,\sigma^{-1}(i)}^r, v_{i,\sigma^{-1}(i)}^r)$ for $1 \leq i \leq m$. This implies that the counter operations encountered in π are $\mathbf{Write}(\alpha_{1,\sigma(1)}), \mathbf{Write}(\alpha_{2,\sigma(2)}), \dots, \mathbf{Write}(\alpha_{m,\sigma(m)})$, $\mathbf{Read}(\alpha_{\sigma^{-1}(1),1}), \mathbf{Read}(\alpha_{\sigma^{-1}(2),2}), \dots, \mathbf{Read}(\alpha_{\sigma^{-1}(m),m})$ in that order. Now, σ is a permutation of $[m]$, so the pairs $(\sigma^{-1}(j), j)$ for $1 \leq j \leq m$ are a permutation of the pairs $(j, \sigma(j))$ for $1 \leq j \leq m$. Therefore, the m symbols read are exactly the m symbols written, possibly in a different order. Since the write operations all come before the read operations, this sequence of counter operations is indeed a counter-realizable sequence. Moreover, the only edges that have labels other than 1 are the edges $(u_{i,\sigma(i)}^w, v_{i,\sigma(i)}^w)$ for $1 \leq i \leq m$, and these edges have labels $y_{i,\sigma(i)}$. Therefore, π is a counter-realizable path from s_1^w to t_m^r in H'_N computing the monomial $\prod_{e \in \pi} \mathcal{L}(e) = \prod_{i=1}^m y_{i,\sigma(i)} = \mathcal{M}$.

Conversely, let π be a counter-realizable path from s_1^w to t_m^r in H'_N . For each $1 \leq i \leq m$, the path π_i^w is a path from s_i^w to t_i^w . Any such path clearly is of the form $s_i^w, u_{i,f_i}^w, v_{i,f_i}^w, t_i^w$ for some $1 \leq f_i \leq m$. Similarly, for each $1 \leq i \leq m$, the path π_i^r is a path from s_i^r to t_i^r of the form $s_i^r, u_{g_i,i}^r, v_{g_i,i}^r, t_i^r$ for some $1 \leq g_i \leq m$. We represent the j_i s and k_i s using two functions $f, g : [m] \rightarrow [m]$ defined as $f(i) = j_i$ for all $i \in [m]$ and $g(i) = k_i$ for all $i \in [m]$. From the previous paragraph, the sequence of counter operations along π other than **No-op** is $\text{Write}(\alpha_{1,f(1)}), \text{Write}(\alpha_{2,f(2)}), \dots, \text{Write}(\alpha_{m,f(m)}), \text{Read}(\alpha_{g(1),1}), \text{Read}(\alpha_{g(2),2}), \dots, \text{Read}(\alpha_{g(m),m})$ in that order. This sequence is counter-realizable, because π is a counter-realizable path.

For each $1 \leq j \leq m$, the operation $\text{Read}(\alpha_{g(j),j})$ appears in the sequence of counter operations. This means $\text{Write}(\alpha_{g(j),j})$ is an operation earlier in the sequence. The only such write operation appearing in the sequence is $\text{Write}(\alpha_{g(j),f(g(j))})$, so $f(g(j)) = j$. Similarly, for each $1 \leq j \leq m$, the operation $\text{Write}(\alpha_{j,f(j)})$ appears in the sequence of counter operations, which means $\text{Read}(\alpha_{j,f(j)})$ appears later in the sequence. The only such read operation appearing in the sequence is $\text{Read}(\alpha_{g(f(j)),f(j)})$, so $g(f(j)) = j$. Therefore $f(g(j)) = g(f(j)) = j$ for all $j \in [m]$, so f and g are both permutations of $[m]$ and are inverses of each other. We rewrite f as σ and g as σ^{-1} . The monomial computed by π is, therefore, $\prod_{e \in \pi} \mathcal{L}(e) = \prod_{i=1}^m y_{i,\sigma(i)}$, which is a monomial of $\text{Perm}_m(Y)$. \square

We now construct the graph H_N from graph H'_N . If m is odd, then we add a vertex α and edges (α, s_1^w) and (t_m^r, α) and if m is even, then we add three vertices $\alpha_1, \alpha_2, \alpha_3$ and edges (α_1, s_1^w) (α_2, α_1) , (α_3, α_2) and (t_m^r, α_3) . This ensures that, $N = 4n$ for some parameter n , where N is the number of vertices in H_N . We set the weights of all the extra added edges as 1 and label it with **No-op**. We now add self-loops on all the vertices with weight 1 except the α vertices. All the self-loop edges have the label of **No-op** on it. Consider the counter graph H_N constructed as above, we will argue that the sum of weights of counter-realizable cycle covers in H_N equals the $\text{Perm}_m(Y)$ and the sum of signed weights of counter-realizable clog sequences that are not cycle covers equals 0. Without loss of generality, we assume that m is odd. Similarly, we can extend our arguments for even m .

We already know from our claim that there exists a bijection between the set of monomials in $\text{Perm}_m(Y)$ and the set of counter-realizable paths between s_1^w to t_m^r in graph H'_N . It is therefore sufficient to show a bijection between the set of counter-realizable paths between s_1^w to t_m^r in graph H'_N and the set of all counter-realizable cycle covers of graph H_N . We also argue that the sign of every cycle cover of graph H_N is same (w.l.o.g., we assume it to be positive). Since, α is a vertex in H_N without any self loop, any cycle cover $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \dots, \mathcal{C}_k \rangle$ of H_N must cover α with some cycle, w.l.o.g., we call it \mathcal{C}_1 which have both the edges (α, s_1^w) and (t_m^r, α) , and all other cycles in $\widehat{\mathcal{C}}$ are self-loops on all the vertices which are not covered in cycle \mathcal{C}_1 . It is easy to observe that the length of any cycle in H_N which uses vertex α is always equal to $6m + 2$. Therefore, the total number of vertices of graph H_N which are not covered in this long cycle and which will get covered by self loops in any cycle cover is $N - 6m - 2$. It immediately follows that the sign of every cycle cover of graph H_N is same.

We now show a bijection between the set of all counter-realizable cycle covers of graph H_N and the set of counter-realizable paths between s_1^w to t_m^r in graph H'_N . It is easy to see that a cycle cover $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \dots, \mathcal{C}_k \rangle$ of H_N is counter-realizable iff the long cycle which uses the vertex α is counter-realizable, w.l.o.g., we call the long cycle \mathcal{C}_1 . It is easy to note that, for any counter-realizable cycle cover $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \dots, \mathcal{C}_k \rangle$, the cycle \mathcal{C}_1 must be formed by an edge (α, s_1^w) , followed by a unique counter-realizable directed path \mathcal{P} between s_1^w to t_m^r (of graph H'_N), followed by an edge (t_m^r, α) . Also, for every counter-realizable directed path \mathcal{P} from s_1^w to t_m^r (of graph H'_N), one can form a unique counter-realizable long cycle (and therefore a counter-realizable cycle cover) in H_N where the long cycle is (α, s_1^w) , followed by directed path \mathcal{P} between s_1^w to t_m^r , followed by (t_m^r, α) . This finishes the first part of our argument.

We now argue that the signed sum of all counter-realizable cflow sequences of graph H_N which are not cycle covers is equal to 0. We now argue that there exist no cflow sequence in graph H_N which does not contain the vertex α in any of its cflow. Suppose there exist some cflow which does not contain α then we consider the graph formed by deleting the vertex α in H_N , in such a graph, the only closed walks possible are single-loops on each vertex of such a graph. But the total degree of such a cflow sequence can never be equal to N , therefore, such a cflow sequence is not a valid cflow sequence. Let us assume that α is the least numbered vertex in H_N , say, numbered with 1. Since, α is a vertex without a self-loop, any cflow involving α must have edges (α, s_1^w) and (t_m^r, α) . It is easy to see that in H_N , any counter-realizable cflow sequence, say, $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k \rangle$ satisfies the property that except the first cflow \mathcal{C}_1 (which involves the vertex α), all other cflows in the cflow sequence $\widehat{\mathcal{C}}$ are self-loops. α_1 will be the head of \mathcal{C}_1 . It is crucial to note that since all self-loops in graph H_N are labelled with **No-op**, the cflow sequence $\widehat{\mathcal{C}}$ is counter-realizable iff the first cflow \mathcal{C}_1 is counter-realizable.

We can now use similar ideas discussed in part 2 of the previous claim to show that there always exists an involution ψ on the set of counter-realizable cflow sequences of graph H_N such that ψ will map any counter-realizable cycle cover to itself and for any counter-realizable cflow sequence which is not a cycle cover, say $\widehat{\mathcal{C}}$, there exist another counter-realizable cflow sequence which is not a cycle cover, $\widehat{\mathcal{C}}'$ such that $\psi(\widehat{\mathcal{C}}) = \widehat{\mathcal{C}}'$ and $\psi(\widehat{\mathcal{C}}') = \widehat{\mathcal{C}}$ and the monomials associated with both $\widehat{\mathcal{C}}$ and $\widehat{\mathcal{C}}'$ are same but their signatures are opposite.

Obtaining G_n from H_N . To obtain G_n from this H_N , we need to give an ordering on the vertices that is consistent with Definition 10 and ensure that G_n is a complete graph. Ensuring the latter is easy. We add all the missing edges and set \mathcal{L} value for them to 0.

To describe the ordering, let us first assume that m is odd. When m is even, the ordering can be worked out similarly. We first introduce some notation. For $1 \leq i \leq m-1$ let us denote the vertex obtained by fusing t_i^w with s_{i+1}^w by θ_i . Let us denote the vertex obtained by fusing t_m^w with s_1^r by θ_m . Also for $1 \leq i \leq m-1$, let us denote the vertex obtained by fusing t_i^r with s_{i+1}^r by θ'_i .

The ordering can now be described as follows. Vertex α is set to 1 and the vertex t_m^r is set to 2. The vertices θ'_1 to θ'_{m-1} are numbered in reverse order, i.e. θ'_{m-1} is set to 3, θ'_{m-2} to 4 and so on up to θ'_1 is set to $m + 1$. We also number the vertices θ_1 to θ_m in reverse order starting from $2m + 1$ down to $m + 2$. We number the vertex s_1^w as $2m + 2$ ¹³

Now, let us assume that the symbols in Σ are ordered in some arbitrary order, say a_1, \dots, a_{m^2} . In H_N , let \mathcal{E} be defined as $\{e \mid \phi(e) \neq \text{No-op}\}$. From our construction of H_N , no two edges in \mathcal{E} share any endpoints. Also for each $a_i \in \Sigma$, there is a unique edge with $\text{Write}(a_i)$ on it and a unique edge with $\text{Read}(a_i)$ on it. We now fix the following ordering: the tail of the edge with $\text{Write}(a_i)$ on it is assigned $4(i - 1) + 1 + (2m + 2)$ and its head is assigned $4(i - 1) + 2 + (2m + 2)$, the tail of the edge with $\text{Read}(a_i)$ on it is assigned $4(i - 1) + 3 + (2m + 2)$ and finally, its head is assigned $4(i - 1) + 4 + (2m + 2)$.

5.2 Details regarding VNP-hardness of $\text{CountDet}_n(X)$ when $\text{char} = 2$

Over characteristic 2, Perm_m is known to be easy. Therefore, to prove VNP-hardness over characteristic 2, we use a different VNP-hard polynomial, which was shown to be VNP-hard over characteristic 2 fields in a work of Hrubes [4]. The polynomial is based on the algebraic variant of the well-known Edge Cover problem. We start by defining the polynomial.

Definition 16. Let $m = \binom{\tau}{2}$ for some parameter τ . Let $G = (V, E)$ be a complete undirected graph on τ vertices, i.e. $V = [\tau]$ and $E = \{(i, j) \mid 1 \leq i < j \leq \tau\}$ and let edge $e = (i, j)$ be labelled with $y_{i,j}$.

$$\text{EC}_m^*(Y) = \sum_{E' \subseteq E, E' \text{ is an edge cover}} \prod_{(i,j) \in E', i < j} y_{i,j}$$

In [4], the above polynomial was shown to be VNP-hard. We will show that we can write $\text{EC}_m^*(Y)$ as a projection of $\text{CountDet}_n(X)$, where $n = \text{poly}(m)$.

For this, we will define $m + 1$ counter graphs, W, R_1, \dots, R_m , which when interconnected appropriately will give us another counter graph H_N , where $N = \text{poly}(m)$ and it has the following two properties.

- All the counter-realizable cflow sequences in H_N have the same sign.
- Moreover, the sum of the weights of the counter-realizable cflow sequences which are cycle covers equals EC_m^* and the sum of the weights of the counter-realizable cflow sequences which are not cycle covers = 0.

Construction of W . For each edge $(i, j) \in E$ such that $1 \leq i < j \leq \tau$, we add a directed path $\rho_{i,j} = \langle (s_{i,j}, i), (i, j), (j, t_{i,j}) \rangle$ in W . We will call $s_{i,j}$ the source of $\rho_{i,j}$ and $t_{i,j}$ the sink of $\rho_{i,j}$. We arrange these paths in a linear order $\rho_{1,2}, \dots, \rho_{1,\tau}, \rho_{2,3}, \dots, \rho_{2,\tau}, \dots, \rho_{\tau-1,\tau}$ one after the other. Additionally, we do the following. We rename $s_{1,2}$ as s_1 and $t_{\tau-1,\tau}$ as t_τ

- Add edges from s_1 to all the other sources, i.e. $\forall 1 \leq i < j \leq \tau$, add edge $(s_1, s_{i,j})$.

¹³ For an odd m , note that $2m + 2$ is always a multiple of 4.

- Add edges from the sink of all the paths to t_τ , i.e. $\forall 1 \leq i < j \leq \tau$, add $(t_{i,j}, t_\tau)$.
- Also add edges from sink of a path to the sources of all the paths that come after it in the above order.
- We define $\phi((s_{i,j}, i)) = \mathbf{Write}(\alpha_{i,j})$ and $\phi((j, t_{i,j})) = \mathbf{Write}(\alpha_{j,i})$ for all $i \leq 1 < j \leq \tau$. We also define $\phi((s_1, 1)) = \mathbf{Write}(\alpha_{1,2})$ and $\phi((\tau, t_\tau)) = \mathbf{Write}(\alpha_{\tau, \tau-1})$. For all the other edges we set ϕ to be **No-op**. We also assign $\mathcal{L}((i, j)) = y_{i,j}$.

Let π be any path from s_1 to t_τ in W . We will say that an edge (i, j) of G is traversed in π if $\rho_{i,j}$ is in π , i.e. all the three edges in $\rho_{i,j}$ are traversed in π . It is easy to see the following property holds.

Observation 2 *Let $S \subseteq E$, then there is a path π_S in W such that it traverses exactly the set of edges in S . Moreover, if S is an edge cover then for each vertex $i \in [\tau]$ we would have at least one edge $(i, j) \in S$ such that upon traversing the path π_S we would have done $\mathbf{Write}(\alpha_{i,j})$ and $\mathbf{Write}(\alpha_{j,i})$ for that edge.*

Construction of R_1, \dots, R_τ . We have one graph R_i for each vertex $i \in [\tau]$. This graph will allow for reading the symbols $\alpha_{i,j}$ for all $j \neq i$. For each $i \in [\tau]$, we describe $R_i = (V_i, E_i, \Sigma_i, \phi_i)$. Here, $V_i = \bigcup_{j \in [\tau] \setminus \{i\}} \{a_{i,j}\} \cup \bigcup_{j \in [\tau] \setminus \{i\}} \{b_{i,j}\}$. Let min and max denotes the minimum and maximum number in $[\tau] \setminus \{i\}$. $E_i = \{(a_{i,j}, b_{i,j}) \mid j \in [\tau] \setminus \{i\}\} \cup \{(a_{i,min}, a_{i,j'}) \mid j' \in [\tau] \setminus \{i\} \text{ and } j' > min\} \cup \{(b_{i,j}, b_{i,max}) \mid j \in [\tau] \setminus \{i\} \text{ and } j < max\} \cup \bigcup_{k \in [\tau] \setminus \{i\}} \{(b_{i,k}, a_{i,j}) \mid j > k\}$. $\phi((a_{i,j}, b_{i,j})) = \mathbf{Read}(\alpha_{i,j})$, where $j \neq i$ and ϕ for all the other edges is **No-op**.

We relabel $a_{i,min}$ as a_i^* and we relabel $b_{i,max}$ as b_i^* . We observe the following about R_i .

- Observation 3** – *Let π be any path from a_i^* to b_i^* . There exists at least one $j \in [\tau], j \neq i$ such that we encounter $\mathbf{Read}(\alpha_{i,j})$ along π .*
- *Let $S \subseteq [\tau] \setminus \{i\}$, there exists a path from a_i^* to b_i^* , say π_S , that encounters exactly the set $\{\mathbf{Read}(\alpha_{i,j}) \mid j \in S\}$ along it.*

Construction of H_N . We now interconnect W and R_1, \dots, R_τ to create the graph H_N as follows. We add an edge from t_τ to a_1^* . We also add edges from b_i^* to a_{i+1}^* for $1 \leq i \leq \tau - 1$. Finally, we add an edge from b_τ^* to s_1 and self-loops on all nodes other than b_τ^* and s_1 . We first observe the following properties about H_N .

Claim. Let Π be any counter-realizable path from s_1 to b_τ^* in H_N . The product of the Y variables along Π , corresponds to a unique monomial in $\mathbf{EC}_m^*(Y)$.

Conversely, if \mathcal{M} is a monomial in \mathbf{EC}_m^* then there is a unique counter-realizable path Π in H_N from s_1 to b_τ^* such that the product of Y variables along Π equals \mathcal{M} .

Proof. Let Π be a counter-realizable path from s_1 to b_τ^* in H_N . Clearly, Π is obtained by concatenating the following paths and edges in this order: $\pi_0 \cdot$

$(t_\tau, a_1^*) \cdot \pi_1 \cdot (b_1^*, a_2^*) \pi_2 \dots \cdot (b_{\tau-1}^*, a_\tau^*) \cdot \pi_\tau$, where π_0 is a directed path from s_1 to t_τ in W and π_i is a directed path from a_i^* to b_i^* in R_i for $1 \leq i \leq \tau$.

By part 1 of Observation 3, we know that for each $i \in [\tau]$, π_i must encounter $\text{Read}(\alpha_{i,j})$ for at least one $j \neq i$. As the path is counter-realizable, there will not be any read operation that does not have a corresponding write operation before it. Let S be a subset of edges of G that are traversed in π_0 .

As all reads must find a corresponding write along π_0 , we have that for each vertex i , there must be at least one edge (i, j) in S , which results into $\text{Write}(\alpha_{i,j})$ and $\text{Write}(\alpha_{j,i})$ along π_0 . Therefore, S must be an edge cover. Hence the monomial obtained by taking product of the Y variables along the path gives rise to a monomial EC_m^* .

Conversely, let \mathcal{M} be a monomial in EC_m^* . Then \mathcal{M} corresponds to a subset of edges S of E that forms an edge cover of G . By Observation 2 we know that there exists a unique path in W that traverses exactly the set of edges in S . Let us call this path π_S . As S is an edge cover, we know that for each i in the vertex set of G , there exists at least one $j \neq i$ such that $\text{Write}(\alpha_{i,j})$ occurs in π_S . For $i \in [\tau]$, let $U_i = \{\alpha_{i,j} \mid j \in [\tau] \setminus \{i\} \text{ and } \text{Write}(\alpha_{i,j}) \text{ occurred along } \pi_S\}$. We know that $|U_i| \geq 1$ for $i \in [\tau]$. From the second part of Observation 3 we know that we can uniquely append π_S with paths $\pi_{U_1}, \pi_{U_2}, \dots, \pi_{U_\tau}$, where π_{U_i} is the unique path between a_i^* and b_i^* that traverses the set U_i . Therefore the path $\pi_S \cdot (t_\tau, a_1^*) \cdot \pi_{U_1} \cdot (b_1^*, a_2^*) \pi_{U_2} \dots \cdot (b_{\tau-1}^*, a_\tau^*) \cdot \pi_{U_\tau}$ is a counter-realizable path and the product of the Y variables along it gives rise to the monomial \mathcal{M} . \square

Consider the counter graph H_N as defined in Section 5.2. We will argue that the sum of weights of counter-realizable cycle covers in H_N equals the $\text{EC}_m^*(Y)$ and the sum of signed weights of counter-realizable clow sequences which are not cycle covers equals 0. We already know from Claim 5.2, that there exists a bijection between the set of monomials in $\text{EC}_m^*(Y)$ and the set of counter-realizable paths between s_1 to b_τ^* in graph H_N . It is therefore sufficient to show that there exists a bijection between the set of counter-realizable paths between s_1 to b_τ^* in graph H_N and the set consisting of all counter-realizable cycle covers of graph H_N . Since, we are working on fields of characteristic 2, sign of every cycle cover of graph H_N is positive.

We now show a bijection between the set of all counter-realizable cycle covers of graph H_N and the set of counter-realizable paths between s_1 to b_τ^* in graph H_N . It is easy to note that any cycle cover $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \dots, \mathcal{C}_k \rangle$ of H_N must use the edge (b_τ^*, s_1) , this is because, if it does not use this edge, then there is no way to cover vertices s_1 and b_τ^* by any other cycle in cycle cover $\widehat{\mathcal{C}}$ in graph H_N . We assume that \mathcal{C}_1 is the cycle in cycle cover $\widehat{\mathcal{C}}$ which uses the edge (b_τ^*, s_1) . It is also easy to note that all the vertices which are not covered in \mathcal{C}_1 must be covered by self-loops on each of them in cycle cover $\widehat{\mathcal{C}}$, that is, in other words, all cycles in the cycle cover $\widehat{\mathcal{C}}$, except \mathcal{C}_1 are all self-loops. This is because, after deleting vertices b_τ^* and s_1 , the only cycles left in graph H_N are self-loops. It is easy to see that for any counter-realizable cycle cover $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \dots, \mathcal{C}_k \rangle$, the cycle \mathcal{C}_1 must be formed by an edge (b_τ^*, s_1) , followed by a unique counter-realizable directed path \mathcal{P} between s_1 to b_τ^* of graph H_N . Also, for every counter-realizable

directed path \mathcal{P} between s_1 to b_τ^* , one can form a unique counter-realizable long cycle (and therefore a counter-realizable cycle cover) in H_N where the long cycle is formed by an edge (b_τ^*, s_1) , followed by a unique counter-realizable directed path \mathcal{P} between s_1 to b_τ^* of graph H_N . This finishes the first part of our argument.

We now prove that the sum of signed weights of all counter-realizable cflow sequences of graph H_N which are not cycle covers is equal to 0. We first show that there exist no cflow sequence in graph H_N without using the vertex s_1 in any of its cflow. For the sake of contradiction, let us assume that there exists a cflow sequence $\widehat{\mathcal{C}}' = \langle \mathcal{C}'_1, \mathcal{C}'_2, \mathcal{C}'_3, \dots, \mathcal{C}'_k \rangle$ which does not use vertex s_1 . We now consider the graph formed by deleting the vertex s_1 in H_N , in such a graph, the only closed walks possible are single-loops on each vertex of such a graph. It is easy to see that the total degree of $\widehat{\mathcal{C}}'$ is always less than N , therefore, $\widehat{\mathcal{C}}'$ is not a valid cflow sequence. Let us assume that s_1 is the least numbered vertex in H_N , say, numbered with 1.

Since, s_1 is a vertex without a self-loop, any cflow which uses s_1 must have the edge (b_τ^*, s_1) . It is easy to see that any counter-realizable cflow sequence, say, $\widehat{\mathcal{C}} = \langle \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k \rangle$ consists of the big cycle \mathcal{C}_1 (which uses the vertex s_1) and all other cflows in $\widehat{\mathcal{C}}$ are self-loops. s_1 will be the head of \mathcal{C}_1 . It is crucial to note that since all self-loops in graph H_N are labelled with `No-op`, the cflow sequence $\widehat{\mathcal{C}}$ is counter-realizable iff the cflow \mathcal{C}_1 is counter-realizable.

Using similar ideas discussed above, it can be shown that there exists an involution ψ defined on the set of all counter-realizable cflow sequences of graph H_N , such that the function ψ maps every counter-realizable cflow sequence which is also a cycle cover to itself and for any counter-realizable cflow sequence $\widehat{\mathcal{C}}$, which is not a cycle cover, there exists a counter-realizable cflow sequence $\widehat{\mathcal{C}}'$ which is also not a cycle cover such that the monomials associated with both $\widehat{\mathcal{C}}$ and $\widehat{\mathcal{C}}'$ are same but with opposite signatures, also, $\psi(\widehat{\mathcal{C}}) = \widehat{\mathcal{C}}'$ and $\psi(\widehat{\mathcal{C}}') = \widehat{\mathcal{C}}$.

Ordering of the vertices of H_N To finish the proof we also show that we can give a complete ordering of the vertices of H_N consistent with Definition 10 and turn it into a complete graph to obtain G_n to fit the description of G_n as in Definition 10. To make it a complete graph, simply add all the missing edges and assign \mathcal{L} for the newly added edges to 0. To get the ordering, fix any arbitrary ordering for the set Σ , the stack alphabet of H_N , say $a_1, a_2, \dots, a_{|\Sigma|}$. In H_N , let \mathcal{E} be defined as $\{e \mid \phi(e) \neq \text{No-op}\}$. From our construction of H_N , no two edges in \mathcal{E} share any endpoints. Also for each $a_i \in \Sigma$, there is a unique edge with `Write`(a_i) on it and a unique edge with `Read`(a_i) on it. We now fix the following ordering: the tail of the edge with `Write`(a_i) on it is assigned $4(i-1) + 1$ and its head is assigned $4(i-1) + 2$, the tail of the edge with `Read`(a_i) on it is assigned $4(i-1) + 3$ and finally, its head is assigned $4(i-1) + 4$.

Observation 4 *It is worth noting that another plausible (and probably more natural) way to define the polynomial family in Definition 8 (and Definition 10, respectively) is to consider only the stack-realizable cycle covers (counter-realizable cycle covers, respectively) instead of the stack-realizable cflow sequences (counter-realizable cflow sequences, respectively) in the overall summation. In case of such*

a variant of $\text{StackDet}_n(X)$, the VP-hardness immediately follows whereas the VP-upperbound is not known. However, in case of such a variant of $\text{CountDet}_n(X)$, both VNP-upperbound and VNP-hardness follows.

References

1. Chaugule, P., Limaye, N., Varre, A.: Variants of homomorphism polynomials complete for algebraic complexity classes. In: Computing and Combinatorics - 25th International Conference, COCOON. LNCS, vol. 11653, pp. 90–102. Springer (2019)
2. Durand, A., Mahajan, M., Malod, G., de Rugy-Altherre, N., Saurabh, N.: Homomorphism polynomials complete for vp. In: LIPIcs-Leibniz International Proceedings in Informatics. vol. 29 (2014)
3. Engels, C.: Dichotomy theorems for homomorphism polynomials of graph classes. *Journal of Graph Algorithms and Applications* **20**(1), 3–22 (2016)
4. Hrubes, P.: On hardness of multilinearization, and VNP completeness in characteristics two. *Electronic Colloquium on Computational Complexity (ECCC)* **22**, 67 (2015), <http://eccc.hpi-web.de/report/2015/067>
5. Mahajan, M., Saurabh, N.: Some complete and intermediate polynomials in algebraic complexity theory. *Theory of Computing Systems* **62**(3), 622–652 (2018)
6. Mahajan, M., Vinay, V.: Determinant: Combinatorics, algorithms, and complexity. Tech. rep. (1997)
7. Mengel, S.: Characterizing arithmetic circuit classes by constraint satisfaction problems. In: ICALP. pp. 700–711. Springer (2011)
8. Mengel, S.: Arithmetic branching programs with memory. In: International Symposium on Mathematical Foundations of Computer Science. pp. 667–678. Springer (2013)
9. Prasad Chaugule, Nutan Limaye, S.P.: Variants of the determinant polynomial and the VP-completeness (October 2020), <https://eccc.weizmann.ac.il/report/2020/152/>, [Online; posted 07-October-2020]
10. Raz, R.: Elusive functions and lower bounds for arithmetic circuits. In: Proceedings of the fortieth annual ACM symposium on Theory of computing. pp. 711–720. ACM (2008)
11. Shpilka, A., Yehudayoff, A.: Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends® in Theoretical Computer Science* **5**(3–4), 207–388 (2010)
12. Valiant, L.G.: Completeness classes in algebra. In: Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing. pp. 249–261. STOC '79 (1979)