

QUASI-POLYNOMIAL HITTING-SET FOR SET-DEPTH- Δ FORMULAS

MANINDRA AGRAWAL, CHANDAN SAHA, AND NITIN SAXENA

ABSTRACT. We call a depth-4 formula C *set-depth-4* if there exists a (unknown) partition $X_1 \sqcup \dots \sqcup X_d$ of the variable indices $[n]$ that the top product layer respects, i.e. $C(\mathbf{x}) = \sum_{i=1}^k \prod_{j=1}^d f_{i,j}(\mathbf{x}_{X_j})$, where $f_{i,j}$ is a *sparse* polynomial in $\mathbb{F}[\mathbf{x}_{X_j}]$. Extending this definition to any depth - we call a depth- Δ formula C (consisting of alternating layers of Σ and Π gates, with a Σ -gate on top) a *set-depth- Δ* formula if every Π -layer in C respects a (unknown) partition on the variables; if Δ is even then the product gates of the bottom-most Π -layer are allowed to compute arbitrary monomials.

In this work, we give a hitting-set generator for set-depth- Δ formulas (over *any* field) with running time polynomial in $\exp((\Delta^2 \log s)^{\Delta-1})$, where s is the size bound on the input set-depth- Δ formula. In other words, we give a *quasi*-polynomial time *blackbox* polynomial identity test for such constant-depth formulas. Previously, the very special case of $\Delta = 3$ (also known as *set-multilinear* depth-3 circuits) had no known sub-exponential time hitting-set generator. This was declared as an open problem by Shpilka & Yehudayoff (FnT-TCS 2010); the model being first studied by Nisan & Wigderson (FOCS 1995). Our work settles this question, not only for depth-3 but, up to depth $\epsilon \log s / \log \log s$, for a fixed constant $\epsilon < 1$.

The technique is to investigate depth- Δ formulas via depth- $(\Delta - 1)$ formulas over a *Hadamard algebra*, after applying a ‘shift’ on the variables. We propose a new algebraic conjecture about the *low-support rank-concentration* in the latter formulas, and manage to prove it in the case of set-depth- Δ formulas.

1. INTRODUCTION

Polynomial identity testing (PIT) - the algorithmic question of examining if a given arithmetic circuit computes an identically zero polynomial - has received some attention in the recent times, primarily due to its close connection to circuit lower bounds. It is now known that a complete (blackbox) derandomization of PIT for depth-4 formulas, via a particular kind of pseudorandom generators, implies $\text{VP} \neq \text{VNP}$ (an algebraic analogue of the much coveted result: $\text{P} \neq \text{NP}$). It is also known that $\text{VP} \neq \text{VNP}$, which amounts to proving exponential circuit lower bounds, must necessarily be shown before proving $\text{P} \neq \text{NP}$ ([Val79, SV85]). Blackbox identity testing (equivalently, the problem of designing hitting-set generators), being a promising approach to proving lower bounds, naturally calls for a closer examination. Towards this, some progress has been made in the form of polynomial time hitting set generators for the following models:

- depth-2 formulas [KS01],
- depth-3 formulas with bounded top fanin [ASSS12, SS11],
- depth-4 (bounded depth) constant-occur formulas [ASSS12],

and a quasi-polynomial time hitting-set generator for

- multilinear constant-read formulas [AvMV11],

Key words and phrases. polynomial identity testing, hitting-set, set-multilinear formula, Hadamard algebra, shift, low-support rank concentration.

among some others (refer to the surveys [SY10, Sax09, AS09]). The hope is, by studying these special but interesting models we might develop a deeper understanding of the nature of hitting sets and thereby get a clue as to what techniques can be lifted to solve PIT in general (i.e. for depth-4 formulas). One such potentially effective technique is the study of *partial derivatives* of formulas.

Despite the apparent difference between the approaches of [ASSS12] and [AvMV11], at a finer level they share a common ingredient - the use of partial derivatives. The partial derivative based method was introduced in the seminal paper by Nisan and Wigderson [NW97] for proving circuit lower bounds, and since then it has been successfully applied (with more sophistications) to prove various interesting results on lower bounds, identity testing and reconstruction of circuits [ASSS12, AvMV11, GKQ12, GKKS12] (refer to the surveys [SY10, CKW11] for much more).

Partial derivatives & shifting - the intuition: In a way, partial derivatives *shift* the variables by some amount - for e.g., if $f(x_1, x_2, \dots, x_n)$ is a multilinear polynomial then its partial derivative with respect to x_1 is $f(x_1 + 1, x_2, \dots, x_n) - f(x_1, \dots, x_n)$. Out of curiosity, one might ask what happens if we shift the polynomial by arbitrary field constants? If we shift a monomial $f(\mathbf{x}) = x_1 x_2 \dots x_n$ by $\mathbf{c} = (c_1, \dots, c_n) \in \mathbb{F}^n$, $c_i \neq 0$, we get the polynomial $f(\mathbf{x} + \mathbf{c}) = (x_1 + c_1)(x_2 + c_2) \dots (x_n + c_n)$. Something interesting has happened here: The polynomial $f(\mathbf{x} + \mathbf{c})$ has many *low-support* monomials. By a low-support monomial, we mean that the number of variables involved in the monomial is less than a predefined small quantity, say ℓ .

Is it possible that shifting has a similar effect on a more general polynomial $f(\mathbf{x})$, i.e. $f(\mathbf{x} + \mathbf{c})$ has low-support monomials with nonzero coefficients, if $f \neq 0$? Surely, this is true if \mathbf{c} is chosen randomly from \mathbb{F}^n (by Schwartz-Zippel [Sch80, Zip79]). But, f is not just any arbitrary polynomial, it is a polynomial computed by a formula (say, depth-3 or depth-4 formula). This makes it an interesting proposition to investigate the following derandomization question: Let $f \neq 0$ be a polynomial computed by a formula. Is it possible to efficiently compute a *small* collection of points $\mathcal{T} \subset \mathbb{F}^n$, such that there exists a $\mathbf{c} \in \mathcal{T}$ for which $f(\mathbf{x} + \mathbf{c})$ has a low-support monomial with nonzero coefficient?

If the answer to the above question is yes, then it is fairly straightforward to do an efficient blackbox identity test on f : For the right choice of $\mathbf{c} \in \mathcal{T}$, $g(\mathbf{x}) = f(\mathbf{x} + \mathbf{c}) \neq 0$ has a low-support monomial. To witness that $g(\mathbf{x}) \neq 0$, it suffices to keep a set of ℓ variables intact and set the remaining $n - \ell$ variables to zero in g ; running over all possible choices of ℓ variables whom we choose to keep intact, we can witness the fact that $g \neq 0$. Since ℓ is presumably small, $g(\mathbf{x})$ restricted to ℓ variables is a sparse polynomial which can be efficiently tested for nonzeroness in a blackbox fashion [KS01].

Indeed, we prove that the above intuition is true for the class of set-depth- Δ formulas (precisely defined in Section 1.1) - a highly interesting class capturing many other previously studied models (see Section 1.1), including *set-multilinear* depth-3 circuits.

Set-multilinear depth-3 circuits: A circuit $C = \sum_{i=1}^k \prod_{j=1}^d f_{i,j}(\mathbf{x}_{X_j})$ is called a set-multilinear depth-3 circuit if $X_1 \sqcup \dots \sqcup X_d$ is a partition of the variable indices $[n]$ and $f_{i,j}(\mathbf{x}_{X_j})$ is a linear polynomial in the variables \mathbf{x}_{X_j} i.e. the set of variables corresponding to the partition X_j . The set-multilinear depth-3 model, first defined by [NW97], kicked off a flurry of activity. Though innocent-looking, it has led researchers to various arithmetic inventions - the *partial derivative* method for circuit lower bounds [NW97], noncommutative whitebox PIT [RS05], the relationship between *tensor-rank* and super-polynomial

circuit lower bounds [Raz10], hitting-set for tensors, low-rank recovery of matrices, rank-metric codes [FS12], and reconstruction (or learnability) of circuits [KS06]. Although, an exponential lower bound for set-multilinear depth-3 circuits is known [NW97, RY09], the closely associated problem of efficient blackbox identity testing on this model remained an open question, until this work.

Our contribution: Hitting set for set-depth- Δ formulas - A whitebox deterministic polynomial time identity test for set-depth- Δ follows from the noncommutative PIT results [RS05]. We are interested in *blackbox* PIT and, naturally, we cannot see inside C and the underlying partitions of $[n]$. The only information we have is the circuit-size bound, s . To our knowledge, there was no sub-exponential time hitting-set known for the set-depth- Δ model. Our work improves this situation to quasi-polynomial for *any* underlying field (refer Theorem 1). We remark that even the very special case of set-multilinear depth-3 circuits had no sub-exponential hitting-set known (see [SY10, Problem 27]); closest being the recent result of [FS12] where they give a quasi-polynomial hitting-set for *tensors*, i.e. the *knowledge of the sets* X_1, \dots, X_d is required.

Furthermore, set-depth-4 covers other well-studied models - diagonal circuits [Sax08] & semi-diagonal circuits [SSS12] - that had whitebox identity tests but no blackbox sub-exponential PIT were known. For these (and set-multilinear depth-3), our hitting-set has time complexity $s^{O(\log s)}$, although, for general set-depth-4 it requires $s^{O(\log^2 s)}$.

Depth-4 formulas being the ultimate frontier for PIT (and lower bounds) [AV08], one might wonder about the utility of our result on hitting-set for set-depth- Δ formulas beyond $\Delta = 4$. It turns out that there is an interesting connection: We show that a quasi-polynomial hitting set generator for set-depth-6 formulas implies a quasi-polynomial hitting set generator for depth-3 formulas of the form $C = \sum_{i=1}^k \prod_{j=1}^d f_{i,j}(\mathbf{x}_{X_j})^{e_{i,j}}$, where $X_1 \sqcup \dots \sqcup X_d$ defines a partition on $[n]$ and $f_{i,j}$ are linear polynomials. Since arbitrary powers $e_{i,j} \geq 0$ are allowed, the above depth-3 model is stronger than set-multilinear depth-3 formulas (as there is no restriction of multilinearity). This appears to be temptingly close to the general depth-3 model modulo the partition on variables, and provides us with a good motivation to understand the strength of our approach against depth-3 formulas.

Technical novelty of our approach - As mentioned before, many works have looked at the partial derivatives of a formula and related matrices, e.g. the Jacobian [ASSS12, BMS11]. From a geometric viewpoint, the study via derivatives shifts the variables by an *infinitesimal* amount and hopes to discover interesting structure. We take a more radical approach; we shift the circuit by *formal* variables and look at how the circuit changes by considering a *transfer* matrix T . The transfer matrix originates from the study of a formula with field coefficients via a simpler one having *Hadamard algebra* coefficients. This makes the transfer process more amenable to an attack using matrices and linear algebra; proving properties that are vaguely reminiscent of the case of top-fanin $k = 1$.

The main technicality lies in proving the invertibility of a transfer matrix, which is an exponential-sized matrix. Some of the arguments here are combinatorial in nature involving greedy and binary-search paradigms.

Although, Hadamard algebra is implicit in the whitebox identity test of [RS05] and the study of PIT over commutative algebras of [SSS09] (Theorem 6 in [SSS09]), the novelty of our approach lies in understanding the effect of shift by viewing it through the lens of Hadamard algebra, and thereby observing the remarkable phenomenon of *low-support rank concentration*, which in turn implies that a low-support monomial survives after shifting.

We state our results more precisely now.

1.1. Our results. *Set-depth & set-height formulas* - Let C be an arithmetic formula over a field \mathbb{F} in n variables \mathbf{x} , consisting of alternating layers of addition (Σ) and multiplication (Π) gates, with a Σ -gate on top. The number of layers of Π -gates in C is called the *product-depth* (or simply *height*) of C and will be denoted by H . Naturally, the depth of C - which is the number of layers of gates in C - is either $\Delta = 2H$ or $2H + 1$. Counting the Π -layers from the top, we label these layers by numbers in the range $[H]$ and will be referring to a layer as the h -th Π -layer in C , for $h \in [H]$.

We say that C is a *set-depth- Δ* formula if for every h -th Π -layer in C , there exists a partition $X_{h,1} \sqcup \dots \sqcup X_{h,d_h}$ of variable indices $[n]$ that the product gates of the h -th Π -layer respect. In other words, for every $h \in [H]$ the i -th product gate in the h -th Π -layer computes a polynomial of the form $\prod_{j=1}^{d_h} f_{i,j}(\mathbf{x}_{X_{h,j}})$, where each $f_{i,j}(\mathbf{x}_{X_{h,j}})$ is a set-depth- $(\Delta - 2h)$ formula of height $H - h$ on the variable set $\mathbf{x}_{X_{h,j}}$. If $\Delta = 2H$ then the product gates of the H -th Π -layer are allowed to compute arbitrary monomials, i.e. here the H -th Π -layer *need not* respect any partition of the variables.

We will also refer to C as a *set-height- H* formula. *Size* of C , denoted by s or $|C|$, is the number of gates (including the input gates) in C .

Theorem 1 (Main). *There is a hitting-set generator for set-height- H formulas, of size s , that runs in time polynomial in $\exp((2H^2 \log s)^{H+1})$, over any field \mathbb{F} .*

Remarks. 1. For blackbox PIT of set-multilinear depth-3 formulas this gives a *quasi-polynomial* time complexity of $s^{O(\log s)}$ - this is the first sub-exponential time algorithm. 2. For constants $H > 1$ the formula may *not* be multilinear, though the hitting-set remains quasi-polynomial. The time complexity remains *sub-exponential* up to $H = \epsilon \log s / \log \log s$, for a fixed constant $\epsilon < 1$.

An interesting model that is not set-depth- Δ but still Theorem 1 could be applied is - semi-diagonal formula. The reason being the *duality* transformation [Sax08, SSS12] that helps us view it as a set-depth-4 formula. We recall - a depth-4 ($\Sigma\Pi\Sigma\Pi$) formula C is *semi-diagonal* if, for all i , its i -th (top) product-gate computes a polynomial of the form $m_i \cdot \prod_{j=1}^b f_{i,j}^{e_{i,j}}$, where m_i is a monomial, $f_{i,j}$ is a sum of univariate polynomials, and b is a constant. We give two applications, with similar proofs but, for different looking formulas.

Corollary 2 (Semi-diagonal depth-4). *There is a hitting-set generator for semi-diagonal depth-4 formulas, of size s , that runs in time $s^{O(\log s)}$ (assuming $\text{char}(\mathbb{F})$ zero or large).*

Corollary 3 (Set-depth-3 with powers). *Consider a depth-3 formula $C = \sum_{i=1}^k \prod_{j=1}^d f_{i,j}(\mathbf{x}_{X_j})^{e_{i,j}}$, where $f_{i,j}$ is a linear polynomial in $\mathbb{F}[\mathbf{x}_{X_j}]$, $e_{i,j} \in \mathbb{N}$, and $X_1 \sqcup \dots \sqcup X_d$ partitions $[n]$. There is a hitting-set generator for such formulas, of size s , that runs in time $s^{O(\log^2 s)}$ (assuming $\text{char}(\mathbb{F})$ zero or large). The result continues to hold even if $f_{i,j}$ is a sum of univariates.*

Remarks - The restriction on $\text{char}(\mathbb{F})$ in the above two corollaries comes from the use of the duality trick. We think this restriction can be lifted by using *Galois rings* ([Sax08, SSS12]), and defining *rank* for a Hadamard algebra over a Galois ring appropriately. We avoid working out the details here just to keep the focus on the main contributions of this work.

1.2. Organization. We develop an extensive terminology in Section 2, which would be useful later. This section also shows the proof idea at work for the example case of diagonal circuits. Section 3 proves the first structural property - a small shift ensures *low-block-support rank-concentration* in a product of polynomials, that have disjoint variables and

only low-weight monomials. Starting with this as a base case, Section 4 proves the second structural property - a small shift ensures low-support rank-concentration in set-depth- Δ formulas (thus, achieving the presence of a low-support monomial). Finally, the proofs of our main results (or hitting-sets) are completed in Section 5.

2. THE BASICS

2.1. Polynomials. Let $\mathbb{N} := \mathbb{Z}_{\geq 0}$ and $[n] := \{1, \dots, n\}$. Let R be a commutative ring. In the motivating cases R will be a field \mathbb{F} , which we implicitly assume to be large enough. This we can do as the required field extensions are constructible in deterministic polynomial time [AL86], further, as in blackbox PIT we are allowed to evaluate the circuit over any ‘small’ field extension.

Not always will we use bold notation for a vector, hopefully the context will avoid the confusion. For a vector $e \in \mathbb{Z}^n$ we define $|e| := \sum_i e_i$. Also, let the *support* be $S(e) := \{i \mid e_i \neq 0\}$ and the *weight* $s(e)$ be its size. For an *exponent vector* $e \in \mathbb{N}^n$, we define a *coefficient* operator $\text{Coef}(e) : R[\mathbf{x}] \rightarrow R$ that on a polynomial $f \in R[\mathbf{x}]$ equals the coefficient of x^e in f . Clearly, it is an R -module homomorphism but is not multiplicative. Define the *support of f* as $S(f) := \{e \in \mathbb{N}^n \mid \text{Coef}(e)(f) \neq 0\}$ and the *sparsity* $s(f)$ be its size. The *monomial-weight of f* is $\mu(f) := \max_{e \in S(f)} s(e)$. Further, define the *cone of f* as $\mathcal{S}(f) := \{e' \in \mathbb{N}^n \mid \exists e \in S(f), e' \leq e\}$, where the inequality is coordinate-wise, and its size as $\mathfrak{s}(f)$. Note that for a sparse polynomial f , $s(f)$ is small but $\mathfrak{s}(f)$ is usually exponential.

Lemma 4 (Cone). *For an n -variate polynomial f , of degree bound d and monomial-weight μ , we have $\mathfrak{s}(f) \leq \binom{n+1}{\mu} \cdot \binom{d+\mu}{\mu}$.*

For $u, v, a \in \mathbb{N}^n$ define $v! := \prod_{i \in [n]} v_i!$, $\binom{v}{u} := \prod_{i \in [n]} \binom{v_i}{u_i} = \frac{v!}{u! \cdot (v-u)!}$, and $a^{v-u} := \prod_{i \in [n]} a_i^{v_i - u_i}$. We keep in mind the conventions: For all $a < b \in \mathbb{N}$, $\binom{a}{b} = 0$ and $\binom{a}{0} = 1$.

Lemma 5 (Shift on monomials). *Let $u, v \in \mathbb{N}^n$, $a_1, \dots, a_n \in R$ and $f = \prod_{i \in [n]} (x_i + a_i)^{v_i}$. Then, $\text{Coef}(u)(f) = \binom{v}{u} \cdot a^{v-u}$.*

For a polynomial f a shift does not change $\mu(f)$ but, might blow up $s(f)$ exponentially.

2.2. Hadamard algebras. For a commutative ring R and $\kappa \in \mathbb{N}$, we define the *Hadamard algebra* $H_\kappa(R) := (R^\kappa, +, \star)$, on the free R -module R^κ , by defining: $u \star v := (u_i \cdot v_i)_{i \in [\kappa]}$, where \cdot is the multiplication in R . $H_\kappa(R)$ is an R -algebra (it is closed, associative, distributive and commutative) with the zero vector as *zero* and the all-one vector as *unity*.

We can now naturally define the *polynomial ring over $H_\kappa(R)$* , $H_\kappa(R)[\mathbf{x}]$. It inherits the operations $+$, \star , and all the elements of $H_\kappa(R)$. Also there is an obvious isomorphism between the algebras $H_\kappa(R)[\mathbf{x}]$ and $H_\kappa(R[\mathbf{x}])$. (View the elements of $H_\kappa(R)$ and $H_\kappa(R)[\mathbf{x}]$ as ‘column vectors’ with entries from R and $R[\mathbf{x}]$, respectively.)

For an $e \in \mathbb{N}^n$ and $f \in H_\kappa(R)[\mathbf{x}]$, we have the natural notions – *coefficient* operator $\text{Coef}(e) : H_\kappa(R)[\mathbf{x}] \rightarrow H_\kappa(R)$, *support* $S(f) \subset \mathbb{N}^n$, and *sparsity* $s(f)$.

Low-support coefficient-space - For any polynomial f over a Hadamard algebra $H_\kappa(R)$, where R is a field, and $\ell \in \mathbb{N}_{>0}$, define $V_\ell(f) := \text{sp}_R\{\text{Coef}(e)(f) \mid e \in \mathbb{N}^n, s(e) < \ell\} \subseteq H_\kappa(R)$. We call f *ℓ -concentrated over $H_\kappa(R)$* if $V_\ell(f) = \text{sp}_R\{\text{Coef}(e)(f) \mid e \in \mathbb{N}^n\}$.

We can extend the above definition also to the case when R is an integral domain, as we can then work with the associated field of fractions.

We demonstrate the usefulness of Hadamard algebra & ‘shifting’ in achieving low-support rank concentration, using the example case of diagonal circuits (see Section A).

2.3. Proof ideas. With the spirit of the argument (as in Section A) in mind, let us state the proof ideas. Let $C(\mathbf{x}) = \sum_{i=1}^k \prod_{j=1}^d f_{i,j}(\mathbf{x}_{X_j})$, where $f_{i,j}$ is a *sparse* polynomial in $\mathbb{F}[\mathbf{x}_{X_j}]$, be a set-depth-4 formula. Consider a $\Pi\Sigma\Pi$ formula

$$D(\mathbf{x}) := f_1(\mathbf{x}_{X_1}) \star \cdots \star f_d(\mathbf{x}_{X_d}) \quad \text{over } H_k(\mathbb{F}),$$

where the i -th coordinate of $f_j(\mathbf{x}_{X_j})$ is $f_{i,j}(\mathbf{x}_{X_j})$. Note that $C(\mathbf{x})$ can be expressed as $(1, 1, \dots, 1) \cdot D(\mathbf{x})$, where \cdot is the usual matrix product. Denote $(1, 1, \dots, 1)$ by $\mathbf{1}$.

For a subspace $V \subseteq \mathbb{F}^k$ and polynomials $D_1, D_2 \in H_k(\mathbb{F})[\mathbf{x}]$, we say $D_1 \equiv D_2 \pmod{V}$ if each coefficient of $D_1 - D_2$ is in V . Somewhat wishfully, we would like to propose a *low-support rank-concentration* property:

Conjecture 6 (Wishful!). *If $\ell > \log |D|$ then $D(\mathbf{x}) \equiv 0 \pmod{V_\ell(D)}$.*

If this is true then the coefficient of \mathbf{x}^e , in D , is in the \mathbb{F} -span of those coefficients that correspond to low support, i.e. $O(\log |D|)$. Suppose we verify the zeroness of $\pi_S \circ C(\mathbf{x}) = \mathbf{1} \cdot D(\pi_S \mathbf{x})$, for $S \in \binom{[n]}{\ell-1}$ and $\pi_S : x_i \mapsto (x_i \text{ if } i \in S, \text{ else } 0)$. This means that $\forall e \in \mathbb{N}^n$ with $s(e) < \ell$ we have $\mathbf{1} \cdot \text{Coef}(e)(D) = 0$. Now the conjecture implies that also $\forall e \in \mathbb{N}^n$ with $s(e) \geq \ell$ we have $\mathbf{1} \cdot \text{Coef}(e)(D) = 0$, clearly implying, $C(\mathbf{x}) = \mathbf{1} \cdot D(\mathbf{x}) = 0$. In other words, we have a blackbox PIT for set-depth-4 in time $\text{poly}(n^{\log |C|})$.

Unfortunately, Conjecture 6 is easily false! For example, let $D(\mathbf{x}) = x_1 \cdots x_n$ and $1 < \ell \leq n$. Then obviously $D(\mathbf{x}) \not\equiv 0 \pmod{V_\ell(D)}$.

Here is where ‘shifting’ enters the picture. The goal in this paper is to prove that after a ‘small’ shift of the variables, D begins to satisfy something like Conjecture 6. This requires a rather elaborate study of how a formula changes when shifted; the meat is expressed through certain *transfer* equations. Looking ahead, we conjecture (without proof) that the phenomena continue to hold in *general* constant-depth formulas.

2.4. Set-height formulas over Hadamard algebra. Just as we have defined set-height formulas over a field \mathbb{F} - meaning, the underlying constants come from \mathbb{F} , we can also define set-height formula in a natural way over any Hadamard algebra $H_\kappa(R)$. The reason we can extend the definition to arbitrary $H_\kappa(R)$ is that the defining property of set-height formulas is the existence of a partition of variables for every Π -layer (irrespective of where the constants of the formula come from). *Size* of a formula C over $H_\kappa(R)$ is defined as κ times the number of gates in C .

Let C be a set-height- H formula (over \mathbb{F}) of depth Δ - we will count depth of C from the top, i.e. the top Σ -gate is at depth 1. If Δ is even (resp. odd) then the gates of the bottom-most Σ -layer compute sparse polynomials (resp. linear polynomials) in the variables. Let k be the maximum among the fanin of the Σ -gates of C (barring the gates of the bottom-most Σ -layer), and d the maximum among the fanin of the Π -gates in C .

Uniform fanin of Σ and Π -gates - With the definitions of k and d as above, we can assume that the fanin of every Σ -gate in C (barring the gates of the bottom-most Σ -layer) is k , and fanin of every Π -gate is d . This can be achieved by introducing ‘dummy’ gates: The ‘dummy’ Σ -gates introduced as children of a Π -gate compute the field constant 1, and the ‘dummy’ Π -gates introduced as children of a Σ -gate also compute 1 except that some of the field constants on the wires are set to zeroes. This process keeps C a set-height- H formula but might bloat up the size from s to s^Δ , although it does not change k and d (according to the way we have defined them). Of course, formula C is not modified physically as it is presented as a blackbox. But the point is, even in the blackbox setting we can treat C as a set-height- H formula with *uniform fanin* of Σ and Π -gates. We will call this

uniform fanin of the Σ and Π -gates as the Σ -fanin and Π -fanin, respectively. Note that the definition of Σ -fanin excludes the gates of the bottom-most Σ -layer - they are handled next.

Fanin bound on bottom-most Σ -gates - If Δ is even, denote the set of monomials computed by the H -th Π -layer by M ; if Δ is odd then $M := \mathbf{x} \cup \{1\}$. The fanin of every gate of the bottom-most Σ -layer is bounded by $\lambda := |M| + 1$. Refer to λ as the *sparsity parameter*.

Henceforth, we will assume uniform Σ and Π -fanin of C (k and d respectively), keeping in mind that the fanin of every gate of the bottom-most Σ -layer is bounded by λ . All of k, d and λ are in turn bounded by s . Denote this class of formulas over \mathbb{F} by $\mathcal{C}_0(k, d, \lambda, \mathbf{x})$.

Recursive structure of set-height formulas over Hadamard algebras - Let $\mathcal{C}_h(k, d, \lambda, \mathbf{x})$ be the class of set-height- $(H - h)$ formulas, of depth $(\Delta - 2h)$, in the variables \mathbf{x} with Σ -fanin k , Π -fanin d and sparsity parameter λ , over the Hadamard algebra $\mathcal{R}_h := \mathbf{H}_{k^h}(\mathbb{F})$. (Eg., to begin with $h = 0$ and the input formula $C \in \mathcal{C}_0(k, d, \lambda, \mathbf{x})$.) Assume that k, d and λ are less than s , which is the size of the input formula C . Let C_h be a formula in $\mathcal{C}_h(k, d, \lambda, \mathbf{x})$.

$$(1) \quad C_h(\mathbf{x}) = \sum_{i \in [k]} c_i \cdot \prod_{j \in [d]} f_{i,j}(\mathbf{x}_{X_j}),$$

$c_i \in \mathcal{R}_h$, $f_{i,j}(\mathbf{x}_{X_j})$ is a set-height- $(H - h - 1)$ formula over \mathcal{R}_h on the variables \mathbf{x}_{X_j} , and $X_1 \sqcup \dots \sqcup X_d$ is the partition of $[n]$ that the first Π -layer of $C_h(\mathbf{x})$ respects. Let $\mathcal{R}_{h+1} := \mathbf{H}_k(\mathcal{R}_h) = \mathbf{H}_{k^{h+1}}(\mathbb{F})$. Define $f_j(\mathbf{x}_{X_j}) := (f_{1,j}(\mathbf{x}_{X_j}), \dots, f_{k,j}(\mathbf{x}_{X_j}))^T \in \mathcal{R}_{h+1}[\mathbf{x}_{X_j}]$. Let

$$D_h(\mathbf{x}) := f_1(\mathbf{x}_{X_1}) \star \dots \star f_d(\mathbf{x}_{X_d}) = \prod_{j \in [d]} f_j(\mathbf{x}_{X_j}) \quad \text{over } \mathcal{R}_{h+1},$$

where \star denotes the Hadamard product in the algebra \mathcal{R}_{h+1} (extended naturally to the polynomial ring over \mathcal{R}_{h+1}). Evidently,

$$(2) \quad C_h(\mathbf{x}) = (c_1, \dots, c_k) \cdot D_h(\mathbf{x}) = \mathbf{c}^T \cdot D_h(\mathbf{x}),$$

where \cdot is the product for matrices over $\mathcal{R}_h[\mathbf{x}]$. We intend to understand the nature of the circuit $C_h(\mathbf{x})$ by studying the properties of the circuit $D_h(\mathbf{x})$ - it is here that the recursive structure reveals itself as in Lemma 7. Let $\mathcal{P}_h(h') := \{X_{h',1}, \dots, X_{h',d}\}$ be the partition of $[n]$ that the h' -th Π -layer of C_h respects. (Recall that when the depth of C_h is even then the bottom-most Π -layer need not respect any partition - this attribute would always remain implicit in our discussions.) Define the partition $\mathcal{P}_h(h', X_j) := \{X_{h',1} \cap X_j, \dots, X_{h',d} \cap X_j\}$ (ignore here the empty sets), for every $1 \leq j \leq d$.

Lemma 7. *For every $j \in [d]$, $f_j(\mathbf{x}_{X_j})$ is a set-height- $(H - h - 1)$ formula in $\mathcal{R}_{h+1}[\mathbf{x}_{X_j}]$ with Σ -fanin k , Π -fanin d and sparsity parameter λ , i.e. $f_j(\mathbf{x}_{X_j}) \in \mathcal{C}_{h+1}(k, d, \lambda, \mathbf{x}_{X_j})$, such that every h' -th Π -layer of $f_j(\mathbf{x}_{X_j})$ respects the partition $\mathcal{P}_h(h' + 1, X_j)$. (Pf. in App. B)*

2.5. Matrices. A matrix M with coefficients in ring R , and the rows (resp. columns) indexed by \mathcal{I} (resp. \mathcal{J}) is compactly denoted as: $M \in (\mathcal{I} \times \mathcal{J} \rightarrow R)$. M is simultaneously a map from $\mathcal{I} \times \mathcal{J}$ to R , and a R -linear transformation from $R^{|\mathcal{J}|}$ to $R^{|\mathcal{I}|}$. When R is an integral domain, we denote the *rank* by $\text{rk}_R M$. Note that the row-rank and column-rank are equal for a matrix. We call a matrix $M \in (\mathcal{I} \times \mathcal{J} \rightarrow R)$, $|\mathcal{I}| = |\mathcal{J}| - 1$, *strongly full* if for all $u \in \mathcal{J}$, $M_{\mathcal{I}, \mathcal{J} \setminus \{u\}}$ is invertible. For two matrices M_1, M_2 and a R -module V , we write $M_1 \equiv M_2 \pmod{V}$ to mean that each column of $M_1 - M_2$ is in V . For two matrices $M_1 \in (\mathcal{I}_1 \times \mathcal{J}_1 \rightarrow R)$ and $M_2 \in (\mathcal{I}_2 \times \mathcal{J}_2 \rightarrow R)$, the matrices M_1^{-1} (when $|\mathcal{I}_1| = |\mathcal{J}_1|$), $M_1 M_2$ (when $\mathcal{J}_1 = \mathcal{I}_2$) and $M_1 \otimes M_2$ are in $(\mathcal{J}_1 \times \mathcal{I}_1 \rightarrow R)$, $(\mathcal{I}_1 \times \mathcal{J}_2 \rightarrow R)$ and $((\mathcal{I}_1 \times \mathcal{I}_2) \times (\mathcal{J}_1 \times \mathcal{J}_2) \rightarrow R)$ respectively. For a matrix $M \in R^{\kappa \times a}$ and an element

$v \in H_\kappa(R)$, $v \star M$ is the matrix obtained after taking the Hadamard product of each column with v . For two matrices $M_1 \in (\mathcal{I} \times \mathcal{J}_1 \rightarrow R)$, $M_2 \in (\mathcal{I} \times \mathcal{J}_2 \rightarrow R)$ the *Hadamard-tensor* matrix $M_1 \otimes M_2 \in (\mathcal{I} \times (\mathcal{J}_1 \times \mathcal{J}_2) \rightarrow R)$ is defined as: Its (j_1, j_2) -th column is $(M_1)_{\mathcal{I}, j_1} \star (M_2)_{\mathcal{I}, j_2}$. We list some intuitive formulas.

Lemma 8 (Matrices). *For any column-vector v and matrices E_i, M_i, Z_i , with suitable assumptions on the sizes and invertibility, we have:*

- (1) $(\otimes_i E_i) \cdot (\otimes_i M_i) = \otimes_i (E_i M_i)$.
- (2) $\otimes_i M_i^{-1} = (\otimes_i M_i)^{-1}$.
- (3) $(v \star M_1) \cdot M_2 = v \star (M_1 M_2)$.
- (4) $(Z_1 M_1) \otimes (Z_2 M_2) = (Z_1 \otimes Z_2) \cdot (M_1 \otimes M_2)$.

3. LOW-BLOCK-SUPPORT RANK-CONCENTRATION

For $i \in [\ell]$, let $f_i \in H_\kappa(\mathbb{F})[\mathbf{x}_{X_i}]$ be a polynomial of degree at most δ , where the X_i 's are disjoint subsets of $[n]$. Define $\mu := \max_i \{\mu(f_i)\}$. By Lemma 4, the sparsity parameter $\lambda := \max_i \{\mathfrak{s}(f_i)\}$ of the f_i 's is bounded by $(\delta + n + \mu)^{O(\mu)}$. Define $\ell := 2 \lceil \log_2 \kappa \rceil + 1$.

Consider the depth-3 $(\Pi\Sigma\Pi)$ formula over $H_\kappa(\mathbb{F})$,

$$D := f_1(\mathbf{x}_{X_1}) \star \cdots \star f_\ell(\mathbf{x}_{X_\ell}) \text{ in } H_\kappa(\mathbb{F})[\mathbf{x}].$$

We *shift* it by formal variables \mathbf{t} to get $D(\mathbf{x} + \mathbf{t}) = f_1(\mathbf{x}_{X_1} + \mathbf{t}_{X_1}) \star \cdots \star f_\ell(\mathbf{x}_{X_\ell} + \mathbf{t}_{X_\ell})$ in $H_\kappa(\mathbb{F}[\mathbf{t}])[\mathbf{x}]$. Wlog we can assume that, $\forall i \in [\ell]$, $f_i(\mathbf{t})$ is a unit in $H_\kappa(\mathbb{F}(\mathbf{t}))$. This is because not being a unit only means that the vector $f_i \in \mathbb{F}(\mathbf{t})^\kappa$ has a zero coordinate, say at place $j \in [\kappa]$. Then the j -th coordinate of $D(\mathbf{t})$ is zero, and we can forget this position altogether; project the setting to the simpler algebra $H_{\kappa-1}(\mathbb{F})$. We *normalize* f_i to $f'_i(\mathbf{x}) := f_i(\mathbf{t})^{-1} \star f_i(\mathbf{x} + \mathbf{t})$. Define $D'(\mathbf{x}) := f'_1(\mathbf{x}_{X_1}) \star \cdots \star f'_\ell(\mathbf{x}_{X_\ell})$ in $H_\kappa(\mathbb{F}(\mathbf{t}))[\mathbf{x}]$.

$$(3) \quad D(\mathbf{x} + \mathbf{t}) = D(\mathbf{t}) \star D'(\mathbf{x}).$$

Any exponent $e \in \mathbb{N}^n$, possibly appearing in D' , can be written uniquely as $e = \sum_{i \in [\ell]} e_i$, where $e_i \in \mathcal{S}(f_i)$, because f_i 's are on disjoint set of variables. We will frequently use this identification. We define the *block-support* of e , $\text{bs}(e) := \{i \in [\ell] \mid e_i \neq 0\}$, and let the *block-weight* $\text{bs}(e)$ be its size. Based on this we define a relevant vector space, for $l \in \mathbb{N}_{>0}$,

$$\mathcal{V}_l(D') := \text{sp}_{\mathbb{F}(\mathbf{t})} \{ \text{Coef}(e)(D') \mid e \in \mathbb{N}^n, \text{bs}(e) < l \}.$$

Ordering & Kronecker-based map - We define a *term ordering* on the monomials t^e , $e \in \mathbb{N}^n$, and their inverses. For a $w \in \mathbb{N}^n$ we denote the ordering as $t^e \preceq_w t^{e'}$, or equivalently $1/t^{e'} \preceq_w 1/t^e$, if $\sum_{i \in [n]} w_i e_i \leq \sum_{i \in [n]} w_i e'_i$. Note that the ordering is *multiplicative* on the monomials, equivalently, the induced ordering on the exponents is additive.

For reasons of efficiency, useful later but skippable for now, we assume: \prec_w keeps the monomials $\left\{ \prod_{i \in [\ell]} t^{e_i} \mid \forall i \in [\ell], e_i \in \mathcal{S}(f_i) \right\}$ distinct. If we fix such a $w \in \mathbb{N}_{>0}$ (note: it could be found in time $\lambda^{O(\ell)}$), then the Kronecker-like homomorphism $\tau : t_i \mapsto y^{w_i}$ ($\forall i \in [n]$) will obviously also map the aforementioned monomials to distinct *univariate* ones. We extend τ to a homomorphism from $H_\kappa(\mathbb{F}[\mathbf{t}])[\mathbf{x}]$ to $H_\kappa(\mathbb{F}[y])[\mathbf{x}]$, by keeping \mathbf{x} unchanged. Its domain can be further extended to a subset of $H_\kappa(\mathbb{F}(\mathbf{t}))[\mathbf{x}]$ (i.e. as long as τ does not cause a *division by zero*).

We would like to prove something like Conjecture 6 for $D(\mathbf{x} + \mathbf{t})$. Note that it suffices to focus on $D'(\mathbf{x})$ as its coefficients are all scaled-up by the same nonzero ‘constant’ $D(\mathbf{t})$. The rest of the section is devoted to proving the following theorem.

Theorem 9 (Low block-support suffices). $D'(\mathbf{x}) \equiv 0 \pmod{\mathcal{V}_\ell(D')}$. Further, it remains true under the map τ .

3.1. Shift-&-normalizing D . We investigate the effect of shift-&-normalizing on f_i . Write, for $i \in [\ell]$, $f_i(\mathbf{x}_{X_i}) =: \sum_{v_i \in \mathcal{S}(f_i)} z_{i,v_i} x^{v_i}$. (Note: $v_i \in \mathbb{N}^n$ and we will denote its j -th coordinate by $v_{i,j} \in \mathbb{N}$.) This yields, after shift-&-normalize (division by *units* is allowed in $H_\kappa(\mathbb{F}(\mathbf{t}))$),

$$f'_i(\mathbf{x}) := f_i(\mathbf{x} + \mathbf{t})/f_i(\mathbf{t}) =: \sum_{u_i \in \mathcal{S}(f_i)} z'_{i,u_i} x^{u_i} \in H_\kappa(\mathbb{F}(\mathbf{t}_{X_i}))[\mathbf{x}_{X_i}].$$

The last step defines

$$(4) \quad z'_{i,u_i} = \text{Coef}(u_i)(f'_i) = f_i(\mathbf{t})^{-1} \star \sum_{v_i \in \mathcal{S}(f_i)} z_{i,v_i} \binom{v_i}{u_i} t^{v_i - u_i}$$

for all exponent vectors $u_i \in \mathcal{S}(f'_i) \subseteq \mathcal{S}(f'_i) = \mathcal{S}(f_i)$. The constant coefficient of f'_i , $z'_{i,0} = 1$.

3.2. Transfer equation of a single polynomial. Let f be one of the polynomials f_1, \dots, f_ℓ over $H_\kappa(\mathbb{F})$. Let $S := \mathcal{S}(f)$ and $\mathcal{S} := \mathcal{S}(f)$. For $v \in \mathcal{S}$ define $z_v := \text{Coef}(v)(f)$, and $z'_v := \text{Coef}(v)(f')$. Since f is a unit, obviously, $S \neq \emptyset$ and $\mathcal{S} \neq \emptyset$. Let $Z \in ([\kappa] \times \mathcal{S} \rightarrow \mathbb{F})$ be such that: Its v -th column is the vector z_v . Note that exactly $s(f)$ of these columns are nonzero. Let $Z' \in ([\kappa] \times \mathcal{S} \rightarrow \mathbb{F}(\mathbf{t}))$ be such that: Its u -th column is the vector z'_u . For any $\mathcal{C} \subseteq \mathcal{S}(f)$ we define a diagonal matrix $N_{\mathcal{C}} \in (\mathcal{C} \times \mathcal{C} \rightarrow \mathbb{F}[\mathbf{t}])$ as: Its u -th diagonal element is t^u . Let the *transfer matrix* (of $\Sigma\Pi$ formulas) $T \in (\mathcal{S} \times \mathcal{S} \rightarrow \mathbb{F})$ be such that: Its (v, u) -th entry is $\binom{v}{u}$. We are ready to state the promised *transfer equation*.

Lemma 10 (Transfer equation - primal). $Z' = f(\mathbf{t})^{-1} \star Z N_{\mathcal{S}} T N_{\mathcal{S}}^{-1}$. (Pf. in Appendix C)

For later use, we need a ‘modulo’ version of this transfer equation. As shorthand denote $Z'_{[\kappa], \mathcal{C}}$ by $Z'_{\mathcal{C}}$, for any $\mathcal{C} \subseteq \mathcal{S}$. Note that the transfer matrix captures a transformation, from Z to Z' , which is clearly invertible. Thus, T is an invertible matrix. Define $T' := (T_{\mathcal{S}, \mathcal{S}})^{-1} \in (\mathcal{S} \times \mathcal{S} \rightarrow \mathbb{F})$ and $\mathcal{S}^* := \mathcal{S} \setminus \{0\}$. If $\mathcal{S}^* = \emptyset$ then it only means that $f \in H_\kappa(\mathbb{F})$, and is invertible. Such an f could be dropped from D right in the beginning. From now on we assume $\mathcal{S}^* \neq \emptyset$. We deduce a modulo version now.

Lemma 11 (Transfer equation - mod). We have $f(\mathbf{t})^{-1} \star Z \equiv Z'_{\mathcal{S}^*} N_{\mathcal{S}^*} T'_{\mathcal{S}^*, \mathcal{S}} N_{\mathcal{S}}^{-1} \pmod{z'_0}$. Further, $T'_{\mathcal{S}^*, \mathcal{S}}$ is strongly full. (Pf. in Appendix C)

3.3. Transfer equation of D : Hadamard tensoring. For two subsets $B_1, B_2 \subset \mathbb{N}^n$ we define $B_1 + B_2 := \{b_1 + b_2 \mid b_1 \in B_1, b_2 \in B_2\}$, where the sum is coordinate-wise. For $i \in [\ell]$, let $\mathcal{S}_i := \mathcal{S}(f_i)$ and $\mathcal{S}_i^* := \mathcal{S}_i \setminus \{0\}$. Define $\mathcal{S} := \sum_{i \in [\ell]} \mathcal{S}_i$ and $\mathcal{S}' := \sum_{i \in [\ell]} \mathcal{S}_i^*$. Note that there is a natural identification between \mathcal{S}' and $\times_{i \in [\ell]} \mathcal{S}_i^*$. We will be implicitly using this. For $i \in [\ell]$, define $Z_i \in ([\kappa] \times \mathcal{S}_i \rightarrow \mathbb{F})$ such that: Its u_i -th column is the vector $z_{i,u_i} := \text{Coef}(u_i)(f_i)$. Let $Z \in ([\kappa] \times \mathcal{S} \rightarrow \mathbb{F})$ such that: Its u -th column is the vector $z_u := \text{Coef}(u)(D)$. Note that $Z = \otimes_{i \in [\ell]} Z_i$. For $i \in [\ell]$, define $Z'_i \in ([\kappa] \times \mathcal{S}_i^* \rightarrow \mathbb{F})$ such that: Its v_i -th column is the vector $z'_{i,v_i} := \text{Coef}(v_i)(f'_i)$. (Note that Z'_i has fewer columns than Z_i .) Let $Z' \in ([\kappa] \times \mathcal{S}' \rightarrow \mathbb{F})$ such that: Its v -th column is the vector $z'_v := \text{Coef}(v)(D')$. Note that $Z' = \otimes_{i \in [\ell]} Z'_i$. For any $\mathcal{C} \subseteq \mathcal{S}$ we define a diagonal matrix $N_{\mathcal{C}} \in (\mathcal{C} \times \mathcal{C} \rightarrow \mathbb{F}[\mathbf{t}])$ as: Its u -th diagonal element is t^u . For $i \in [\ell]$, define $T'_i := T'_{\mathcal{S}_i^*, \mathcal{S}_i}$.

Let the *transfer matrix* (of $\Pi\Sigma\Pi$ formulas) $T' \in (\mathcal{S}' \times \mathcal{S} \rightarrow \mathbb{F})$ be $\otimes_{i \in [\ell]} T'_i$.

Lemma 12 (Tf. eqn. depth-3). $D(\mathbf{t})^{-1} \star Z \equiv Z' N_{\mathcal{S}'} T' N_{\mathcal{S}}^{-1} \pmod{\mathcal{V}_\ell(D')}$. (Pf. App. C)

3.4. Combinatorial juggernaut: To select columns of T' . Recall that T' has rows (resp. columns) indexed by \mathcal{S}' (resp. \mathcal{S}) and has entries in \mathbb{F} . Let \mathcal{M} be some $\kappa > 0$ columns that we intend to remove from T' ; we call them *marked* and the others $\mathcal{S} \setminus \mathcal{M}$ are *unmarked*. We make the following claim about the submatrices of T' not involving \mathcal{M} .

Theorem 13 (Invertible minor). *There exist unmarked columns $\mathcal{C} \subseteq \mathcal{S}$, $|\mathcal{C}| = |\mathcal{S}'|$, such that $|T'_{\mathcal{S}', \mathcal{C}}| \neq 0$. (Proof in Appendix C)*

3.5. T' on the nullspace of Z : Finishing Theorem 9. Recall that the columns of Z are indexed by \mathcal{S} . Think of these *ordered* by the weight vector w , as discussed in the beginning of this section. Pick a basis \mathcal{M} , size at most κ , of the column vectors of Z by *starting from the largest column*. Formally, \mathcal{M} gives the unique (once \prec is fixed) basis such that for each u -th, $u \in \mathcal{S} \setminus \mathcal{M}$, column of Z there exist columns $u_1, \dots, u_r \in \mathcal{M}$ spanning the u -th column, and $u \prec u_r \prec \dots \prec u_1$. We think of the columns \mathcal{M} of T' *marked*, and invoke Theorem 13 to get the $\mathcal{C} \subseteq \mathcal{S}$. We define an $A \in (\mathcal{S} \times \mathcal{C} \rightarrow \mathbb{F})$: If a is the v -th column of A then $Z \cdot a = 0$ expresses the \mathbb{F} -linear dependence of z_v on $\{z_{v'} \mid v' \in \mathcal{M}, v \prec v'\}$; in particular, the *least* row where a is nonzero is the v -th, the entry being 1. Recall the transfer equation, Lemma 12, for the following.

Lemma 14 (T' on nullspace of Z). $|T' N_{\mathcal{S}}^{-1} A| \neq 0$. *Further, the leading nonzero inverse-monomial in the determinant has the coefficient $|T'_{\mathcal{S}', \mathcal{C}}|$. (Proof in Appendix C)*

Finally, we use A to finish the proof of our main structure theorem.

Proof of Theorem 9. From the transfer equation, Lemma 12, we recall

$$D(\mathbf{t})^{-1} \star Z \equiv Z' N_{\mathcal{S}'} T' N_{\mathcal{S}}^{-1} \pmod{\mathcal{V}_\ell(D')}.$$

Right-multiplying by A , we get

$$(5) \quad 0 = D(\mathbf{t})^{-1} \star (ZA) \equiv Z' N_{\mathcal{S}'} T' N_{\mathcal{S}}^{-1} A \pmod{\mathcal{V}_\ell(D')}.$$

Since $T' N_{\mathcal{S}}^{-1} A$ is invertible from Lemma 14 and $N_{\mathcal{S}'}$ is obviously invertible, we get

$$Z' \equiv 0 \pmod{\mathcal{V}_\ell(D')}.$$

(Here we do use that the matrices are over $\mathbb{F}(\mathbf{t})$ and that $\mathcal{V}_\ell(D')$ is an $\mathbb{F}(\mathbf{t})$ -vector space.) This immediately implies the first part of Theorem 9, as Z' collected exactly those coefficients of D' that we a priori did not know in $\mathcal{V}_\ell(D')$. The second part of the theorem follows easily as: (1) τ keeps $D(\mathbf{t})$ a unit, and (2) τ corresponds to the correct term ordering \preceq_w . These two properties allow the above proof also work after applying τ . \square

4. LOW-SUPPORT RANK-CONCENTRATION

We will prove that a set-height- H formula, after a ‘small’ shift, begins to have ‘low’-support rank-concentration. The proof is by induction on the height of the formulas over Hadamard algebras. For this, we would need the following concepts.

For $H > h \in \mathbb{N}$, let $\mathbf{t}_h := \{t_{H-1}, \dots, t_{h+1}, t_h\}$ be a set of formal variables and $\mathbb{F}(\mathbf{t}_h)$ be the function field. These \mathbf{t}_h -variables are different from the variables \mathbf{x} involved in the formula C . Let $\mathcal{R}'_h := H_{k^h}(\mathbb{F}(\mathbf{t}_h))$ be a Hadamard algebra over $\mathbb{F}(\mathbf{t}_h)$; $k^h = \dim_{\mathbb{F}(\mathbf{t}_h)} \mathcal{R}'_h$. Further, $\mathcal{R}'_{h+1}[t_h]$ denotes the (univariate) polynomial ring over \mathcal{R}'_{h+1} , and $\mathcal{R}'_{h+1}(t_h)$ is the corresponding *ring of fractions*. ($\mathcal{R}'_{h+1}(t_h)$ is basically $H_{k^{h+1}}(\mathbb{F}(\mathbf{t}_h))$.)

Low-support shift for $\mathcal{C}_h(k, d, \lambda, \mathbf{x})$ - Let τ_h be a map from $\mathbb{F}[\mathbf{x}]$ to $\mathbb{F}(\mathbf{t}_h)[\mathbf{x}]$ defined as,

$$\tau_h : x_i \mapsto x_i + \alpha_{H-1,i} t_{H-1}^{a_{H-1,i}} + \dots + \alpha_{h,i} t_h^{a_{h,i}}, \quad \text{for } x_i \in \mathbf{x},$$

$a_{H-1,i}, \dots, a_{h,i} \in \mathbb{Z}^+$ and $\alpha_{H-1,i}, \dots, \alpha_{h,i} \in \mathbb{F}$. (τ_h fixes \mathbb{F} , i.e. $\tau_h(c) = c$ for $c \in \mathbb{F}$.) In short, we will write $\tau_h : \mathbf{x} \mapsto \mathbf{x} + \boldsymbol{\alpha}_h \mathbf{t}_h^{\mathbf{a}_h}$. For $\ell_h \in \mathbb{N}$, the map τ_h (as above) is called an ℓ_h -support shift for the class of formulas $\mathcal{C}_h(k, d, \lambda, \mathbf{x})$ if for every formula $C_h \in \mathcal{C}_h(k, d, \lambda, \mathbf{x})$, the polynomial $\tau_h(C_h(\mathbf{x})) = C_h(\mathbf{x} + \boldsymbol{\alpha}_h \mathbf{t}_h^{\mathbf{a}_h})$ is ℓ_h -concentrated over \mathcal{R}'_h .

For the rest of our discussion, we will fix ℓ_h as follows, for $H > h \geq 0$:

$$\ell_h := \begin{cases} (2H \lceil H \log_2 k \rceil)^{H-h-1} \cdot 2 \lceil H \log_2(k\lambda) \rceil + 1, & \text{if } \Delta \text{ is even,} \\ (2H \lceil H \log_2 k \rceil)^{H-h} + 1, & \text{if } \Delta \text{ is odd (\& for } h = H, \ell_H := 2). \end{cases}$$

The above setting satisfies the relation $\ell_h = (\ell_{h+1} - 1)H(\ell - 1) + 1$, where $\ell := 2 \lceil H \log_2 k \rceil + 1$, for every $H - 1 > h \geq 0$ (and also for $h = H - 1$ when Δ is odd).

Recall Equation 2 that says - for each $h \in \{0, \dots, H-1\}$ and C_h , there exists $\mathbf{c} \in \mathbb{H}_k(\mathcal{R}_h)$ such that $C_h = \mathbf{c}^T \cdot D_h$. This section is dedicated to proving the following theorem.

Theorem 15 (Low support suffices). *We can construct τ_0 such that $\tau_0 \circ D_0$ is ℓ_0 -concentrated over $\mathcal{R}'_1[t_0]$, in time polynomial in $(d + n + \ell_0)^{\ell_0}$, where $n := |\mathbf{x}|$.*

Proof strategy ahead - The idea is to construct the map τ_h by applying induction on height $H - h$ of the class $\mathcal{C}_h(k, d, \lambda, \mathbf{x})$. By Equation 2,

$$C_h(\mathbf{x}) = \mathbf{c}^T \cdot (f_1(\mathbf{x}_{X_1}) \star \dots \star f_d(\mathbf{x}_{X_d})).$$

From Lemma 7, $f_j(\mathbf{x}_{X_j}) \in \mathcal{C}_{h+1}(k, d, \lambda, \mathbf{x}_{X_j})$. By definition, $\tau_{h+1} : x_i \mapsto x_i + \alpha_{H-1,i} t_{H-1}^{a_{H-1,i}} + \dots + \alpha_{h+1,i} t_{h+1}^{a_{h+1,i}}$ is an ℓ_{h+1} -support shift for $\mathcal{C}_{h+1}(k, d, \lambda, \mathbf{x}_{X_j})$ for every $1 \leq j \leq d$. Here is where we use induction on height $H - h$: We will build the map τ_h from the inductive knowledge of τ_{h+1} . Basically, we will show that it is possible to efficiently compute $a_{h,1}, \dots, a_{h,n} \in \mathbb{Z}^+$ and $\alpha_{h,1}, \dots, \alpha_{h,n} \in \mathbb{F}$ such that $\tau_h : x_i \mapsto \tau_{h+1}(x_i) + \alpha_{h,i} t_h^{a_{h,i}}$ is an ℓ_h -support shift for $\mathcal{C}_h(k, d, \lambda, \mathbf{x})$.

The proof of Theorem 15. The proof proceeds by induction on height $H - h$ of the class $\mathcal{C}_h(k, d, \lambda, \mathbf{x})$ (in other words, *reverse* induction on h). The induction hypothesis is that τ_{h+1} , an ℓ_{h+1} -support shift for the class $\mathcal{C}_{h+1}(k, d, \lambda, \mathbf{x})$, can be constructed in time polynomial in $(d + n + \ell_{h+1})^{\ell_{h+1}}$, where $n := |\mathbf{x}|$. Overall this means, by varying $h \in [0, \dots, H-1]$, we get a hitting-set of size polynomial in $\prod_{h=0}^{H-1} (d + n + \ell_h)^{\ell_h} \leq (d + n + \ell_0)^{\sum_h \ell_h} < (d + n + \ell_0)^{2\ell_0}$. We discuss the base case and the inductive step in separate detail. Keep in mind that $f_j(\mathbf{x}_{X_j}) \in \mathcal{C}_{h+1}(k, d, \lambda, \mathbf{x}_{X_j})$.

4.1. Base case ($h+1 \geq H-1$). The base case is when $H - h - 1 = 1$ or 0 , i.e. $f_j(\mathbf{x}_{X_j})$'s are sparse polynomials or linear polynomials over \mathcal{R}_{h+1} , depending on whether Δ is even or odd, respectively. These two base cases have varying level of difficulty. If $H - h - 1 = 0$ then $\ell_{h+1} = \ell_H = 2$, hence taking τ_H as the identity map suffices (since $f_j(\mathbf{x}_{X_j})$'s are linear polynomials) as an ℓ_H -support shift for the class $\mathcal{C}_H(k, d, \lambda, \mathbf{x})$. If $H - h - 1 = 1$ then $f_j(\mathbf{x}_{X_j})$'s are sparse polynomials. We first prove an, independently interesting, property.

Lemma 16 (Sparse polynomial). *Let $f \in \mathbb{H}_\kappa(\mathbb{F})[\mathbf{x}]$ be a polynomial with degree bound δ . Let $\ell' := 1 + \min\{2 \lceil \log_2(\kappa \cdot s(f)) \rceil, \mu(f)\}$. We can construct a map $\sigma : x_i \mapsto x_i + t^{b_i}$, in time polynomial in $(\delta + n + \ell')^{\ell'}$, such that $\sigma(f)$ is ℓ' -concentrated over $\mathbb{H}_\kappa(\mathbb{F}(t))$. (Ap. D)*

Now we apply the lemma to the sparse polynomial $f_j(\mathbf{x}_{X_j})$, which has the sparsity parameter λ . Hence we define $\tau_{h+1} = \tau_{H-1} : x_i \mapsto x_i + t_{H-1}^{b_i}$ (in other words, $a_{H-1,i} := b_i$). This, by Lemma 16, ensures that the concentration parameter is $2 \lceil \log_2(k^{H-1} \cdot \lambda) \rceil + 1 \leq$

$2 \lceil H \log_2(k\lambda) \rceil + 1 = \ell_{H-1} = \ell_{h+1}$. Finally, τ_{H-1} is an ℓ_{H-1} -support shift for the class $\mathcal{C}_{H-1}(k, d, \lambda, \mathbf{x})$, and it can be constructed in time polynomial in $(d + n + \ell_{H-1})^{\ell_{H-1}}$.

4.2. Induction ($h+1$ to h). Let $\widehat{f}_j(\mathbf{x}_{X_j}) := \tau_{h+1}(f_j(\mathbf{x}_{X_j}))$. Then,

$$\widehat{D}_h(\mathbf{x}) := \tau_{h+1}(D_h(\mathbf{x})) = \widehat{f}_1(\mathbf{x}_{X_1}) \star \cdots \star \widehat{f}_d(\mathbf{x}_{X_d}),$$

where every \widehat{f}_j is ℓ_{h+1} -concentrated over \mathcal{R}'_{h+1} (by induction hypothesis). Let $\mathbf{t} := \{t_{h,1}, \dots, t_{h,n}\}$ be a set of ‘fresh’ formal variables. (We will keep in mind that the \mathbf{t} -variables would be eventually set as univariates in a variable t_h .) As before in Eqn. 3,

$$\widehat{D}_h(\mathbf{x} + \mathbf{t}) = \prod_{j \in [d]} \widehat{f}_j(\mathbf{x}_{X_j} + \mathbf{t}_{X_j}) = \prod_{j \in [d]} \widehat{f}_j(\mathbf{t}_{X_j}) \star \widehat{f}'_j(\mathbf{x}_{X_j}) = \widehat{D}_h(\mathbf{t}) \star \widehat{D}'_h(\mathbf{x}).$$

In the same spirit as Theorem 9, we would like to show that $\widehat{D}'_h(\mathbf{x}) \equiv 0 \pmod{\mathcal{V}_\ell(\widehat{D}'_h)}$, where $\mathcal{V}_\ell(\widehat{D}'_h) := \text{sp}_{\mathbb{F}(\mathbf{t}_{h+1}, \mathbf{t})} \left\{ \text{Coef}(e)(\widehat{D}'_h) \mid e \in \mathbb{N}^n, \text{bs}(e) < \ell \right\}$, and $\ell = 2 \lceil H \log_2 k \rceil + 1$. As before (see ‘key argument’ in Lemma 16), it is sufficient to prove the typical case (i.e. product of the first ℓ polynomials), $\widehat{D}'_{h,\ell}(\mathbf{x}) := \prod_{j \in [\ell]} \widehat{f}'_j(\mathbf{x}_{X_j}) \equiv 0 \pmod{\mathcal{V}_\ell(\widehat{D}'_{h,\ell})}$. Towards this, we define the *truncated* polynomials, $\widehat{g}_j(\mathbf{x}_{X_j}) := \sum_{e: s(e) < \ell_{h+1}} \text{Coef}(e)(\widehat{f}_j) \mathbf{x}_{X_j}^e$ and let the corresponding product be $\widehat{E}_h(\mathbf{x}) := \prod_{j \in [d]} \widehat{g}_j(\mathbf{x}_{X_j})$. Sparsity of $\widehat{g}_j(\mathbf{x}_{X_j})$ over \mathcal{R}'_{h+1} is bounded by $(d^{H-h-1} + n + \ell_{h+1})^{\ell_{h+1}} =: \lambda_h$. Mimicking the notations on \widehat{D}_h let,

$$\widehat{E}_h(\mathbf{x} + \mathbf{t}) = \prod_{j \in [d]} \widehat{g}_j(\mathbf{x}_{X_j} + \mathbf{t}_{X_j}) = \widehat{E}_h(\mathbf{t}) \star \widehat{E}'_h(\mathbf{x}) \quad \text{and} \quad \widehat{E}'_{h,\ell}(\mathbf{x}) := \prod_{j \in [\ell]} \widehat{g}'_j(\mathbf{x}_{X_j}).$$

By Theorem 9, we can find $a_{h,1}, \dots, a_{h,n} \in \mathbb{Z}^+$ in time $(d\lambda_h)^{O(\ell)} = (d + n + \ell_h)^{O(\ell_h)}$ such that by setting $t_{h,i} = \alpha_{h,i} t_h^{a_{h,i}}$ (any $\alpha_{h,i} \in \mathbb{F} \setminus \{0\}$ works), where t_h is a ‘fresh’ formal variable, we can ensure that the following is satisfied:

$$(6) \quad \widehat{E}'_{h,\ell}(\mathbf{x}) \equiv 0 \pmod{\mathcal{V}_\ell(\widehat{E}'_{h,\ell})}.$$

The claim is that the same setting $t_{h,i} = \alpha_{h,i} t_h^{a_{h,i}}$ (now with carefully chosen $\alpha_{h,i}$ ’s) also ensures that $\widehat{D}'_{h,\ell}(\mathbf{x}) \equiv 0 \pmod{\mathcal{V}_\ell(\widehat{D}'_{h,\ell})}$. Consequently, \widehat{D}'_h is $(\ell - 1)(\ell_{h+1} - 1) + 1 < \ell_h$ concentrated over $\mathcal{R}'_{h+1}(t_h)$. This is what we argue next. Equation 6 implies

$$(7) \quad \widehat{E}_{h,\ell}(\mathbf{x} + \boldsymbol{\alpha} \mathbf{t}) = \prod_{j \in [\ell]} \widehat{g}_j(\mathbf{x}_{X_j} + \boldsymbol{\alpha}_{X_j} \mathbf{t}_{X_j}) = \widehat{E}_{h,\ell}(\boldsymbol{\alpha} \mathbf{t}) \star \widehat{E}'_{h,\ell}(\mathbf{x}) \equiv 0 \pmod{\mathcal{V}_\ell(\widehat{E}_{h,\ell}(\mathbf{x} + \boldsymbol{\alpha} \mathbf{t}))},$$

where (reusing symbol) $\mathbf{t} := (t_h^{a_{h,1}}, \dots, t_h^{a_{h,n}})$ and $\boldsymbol{\alpha} := (\alpha_{h,1}, \dots, \alpha_{h,n})$. Define, $\widehat{D}_{h,\ell}(\mathbf{x}) := \prod_{j=1}^\ell \widehat{f}_j(\mathbf{x}_{X_j})$. We need to take a closer look at how the coefficients of $\widehat{D}_{h,\ell}(\mathbf{x})$, $\widehat{D}'_{h,\ell}(\mathbf{x} + \boldsymbol{\alpha} \mathbf{t})$, $\widehat{E}_{h,\ell}(\mathbf{x})$ and $\widehat{E}_{h,\ell}(\mathbf{x} + \boldsymbol{\alpha} \mathbf{t})$ are related to each other. Towards this, define:

$$\begin{aligned} \widehat{z}_{j,u_j} &:= \text{Coef}(u_j)(\widehat{f}_j(\mathbf{x}_{X_j})) \in \mathcal{R}'_{h+1}, \\ \widehat{z}'_{j,u_j} &:= \text{Coef}(u_j)(\widehat{f}_j(\mathbf{x}_{X_j} + \boldsymbol{\alpha}_{X_j} \mathbf{t}_{X_j})) \in \mathcal{R}'_{h+1}[t_h], \\ \widetilde{z}_{j,u_j} &:= \text{Coef}(u_j)(\widehat{g}_j(\mathbf{x}_{X_j})) \in \mathcal{R}'_{h+1}; \text{ equals } \widehat{z}_{j,u_j} \text{ if } u_j \in S(\widehat{g}_j), \\ \widetilde{z}'_{j,u_j} &:= \text{Coef}(u_j)(\widehat{g}_j(\mathbf{x}_{X_j} + \boldsymbol{\alpha}_{X_j} \mathbf{t}_{X_j})) \in \mathcal{R}'_{h+1}[t_h]. \end{aligned}$$

Let,

$$\begin{aligned}\widehat{B}_j &:= \{u_j : \widehat{z}_{j,u_j} \text{ is in the } \mathbb{F}(\mathbf{t}_{h+1})\text{-basis of the coefficients of } \widehat{f}_j\} \text{ and} \\ \widetilde{B}_j &:= \{u_j : \widetilde{z}_{j,u_j} \text{ is in the } \mathbb{F}(\mathbf{t}_{h+1})\text{-basis of the coefficients of } \widehat{g}_j\}\end{aligned}$$

with respect to some fixed basis that comprises coefficients of monomials of as low support as possible. Note that $\widehat{B}_j = \widetilde{B}_j =: B_j$, as \widehat{f}_j is ℓ_{h+1} -concentrated over \mathcal{R}'_{h+1} .

The crucial observation is that, for any $v_j \in B_j$, \widehat{z}'_{j,v_j} gets a t_h -free contribution only from the monomial x^{v_j} , thus, its basis representation looks like:

$$\widehat{z}'_{j,v_j} = (1 + a(v_j, v_j)) \cdot \widehat{z}_{j,v_j} + \sum_{u_j \in B_j \setminus \{v_j\}} a(u_j, v_j) \cdot \widehat{z}_{j,u_j},$$

where a 's are in $\mathbb{F}(\mathbf{t}_{h+1})[t_h]$ and t_h divides each $a(\cdot, v_j)$. Similarly,

$$\widetilde{z}'_{j,v_j} = (1 + b(v_j, v_j)) \cdot \widehat{z}_{j,v_j} + \sum_{u_j \in B_j \setminus \{v_j\}} b(u_j, v_j) \cdot \widehat{z}_{j,u_j},$$

where b 's are in $\mathbb{F}(\mathbf{t}_{h+1})[t_h]$ and t_h divides each $b(\cdot, v_j)$. Now define the following matrices:

$$\begin{aligned}\widehat{Z}_j &\in ([k^{h+1}] \times B_j \rightarrow \mathbb{F}(\mathbf{t}_{h+1})) \quad ; \quad \text{with } u_j\text{-th column } \widehat{z}_{j,u_j}, \\ \widehat{Z}'_j &\in ([k^{h+1}] \times B_j \rightarrow \mathbb{F}(\mathbf{t}_h)) \quad ; \quad \text{with } u_j\text{-th column } \widehat{z}'_{j,u_j}, \\ \widetilde{Z}_j &\in ([k^{h+1}] \times B_j \rightarrow \mathbb{F}(\mathbf{t}_h)) \quad ; \quad \text{with } u_j\text{-th column } \widetilde{z}'_{j,u_j}.\end{aligned}$$

From the above crucial observation,

$$(8) \quad \widehat{Z}'_j = \widehat{Z}_j \cdot \widehat{M}' \quad \text{and} \quad \widetilde{Z}'_j = \widehat{Z}_j \cdot \widetilde{M}',$$

where $\widehat{M}', \widetilde{M}' \in (B_j \times B_j \rightarrow \mathbb{F}(\mathbf{t}_{h+1})[t_h])$ with rows indexed by $u_j \in B_j$ and columns indexed by $v_j \in B_j$. The (u_j, v_j) -th entry of \widehat{M}' contains $a(u_j, v_j)$ if $u_j \neq v_j$, otherwise $1 + a(u_j, v_j)$ if $u_j = v_j$. Similarly, the (u_j, v_j) -th entry of \widetilde{M}' contains $b(u_j, v_j)$ if $u_j \neq v_j$, otherwise $1 + b(u_j, v_j)$ if $u_j = v_j$. Note that both \widehat{M}' and \widetilde{M}' are invertible over $\mathbb{F}(\mathbf{t}_{h+1})(t_h)$ as $\det(\widehat{M}') \equiv \det(\widetilde{M}') \equiv 1 \pmod{t_h}$. Therefore,

$$(9) \quad \widehat{Z}'_j = \widetilde{Z}'_j \cdot (\widetilde{M}'^{-1} \widehat{M}') \quad \text{and} \quad \widetilde{Z}'_j = \widehat{Z}'_j \cdot (\widehat{M}'^{-1} \widetilde{M}')^{-1}.$$

Now observe that any coefficient of $\widehat{D}_{h,\ell}(\mathbf{x} + \boldsymbol{\alpha} \mathbf{t})$ is an $\mathbb{F}(\mathbf{t}_h)$ -linear combination of the columns of $\otimes_{j \in [\ell]} \widehat{Z}_j$ (by the definition of B_j), which by Equation 8 (& Lemma 8-(4)) is an $\mathbb{F}(\mathbf{t}_h)$ -linear combination of the columns of $\otimes_{j \in [\ell]} \widehat{Z}'_j$ - this in turn is an $\mathbb{F}(\mathbf{t}_h)$ -linear combination of the columns of $\otimes_{j \in [\ell]} \widetilde{Z}'_j$ (by Equation 9). By Equation 7, any $\mathbb{F}(\mathbf{t}_h)$ -linear combination of the columns of $\otimes_{j \in [\ell]} \widetilde{Z}'_j$ can be expressed as an $\mathbb{F}(\mathbf{t}_h)$ -linear combination of those columns u of $\otimes_{j \in [\ell]} \widetilde{Z}'_j$ for which $\text{bs}(u) < \ell$, which in turn can be expressed as an $\mathbb{F}(\mathbf{t}_h)$ -linear combination of those columns u of $\otimes_{j \in [\ell]} \widehat{Z}'_j$ for which $\text{bs}(u) < \ell$ (by Equation 9 again). In other words, we have shown the following: $\widehat{D}_{h,\ell}(\mathbf{x} + \boldsymbol{\alpha} \mathbf{t}) \equiv 0 \pmod{\mathcal{V}_\ell(\widehat{D}_{h,\ell}(\mathbf{x} + \boldsymbol{\alpha} \mathbf{t}))}$. This would imply that $\widehat{D}'_{h,\ell}(\mathbf{x}) \equiv 0 \pmod{\mathcal{V}_\ell(\widehat{D}'_{h,\ell})}$, if we choose $\boldsymbol{\alpha}$ so that the map $t_{h,i} \mapsto \alpha_{h,i} t_h^{a_{h,i}}$ ensures that $\widehat{f}_j(\boldsymbol{\alpha}_{X_j} \mathbf{t}_{X_j})^{-1}$ is *well-defined* in $\mathcal{R}'_{h+1}(t_h)$. Such an $\boldsymbol{\alpha}$ can be constructed, by Lemma 17, in time polynomial in $\lambda_h = (d^{H-h-1} + n + \ell_{h+1})^{\ell_{h+1}}$. Therefore, $\tau_h : x_i \mapsto \tau_{h+1}(x_i) + \alpha_{h,i} t_h^{a_{h,i}}$ is such that $\tau_h(D_h(\mathbf{x}))$ is ℓ_h -concentrated over $\mathcal{R}'_{h+1}[t_h]$. Since $C_h(\mathbf{x}) = c^T \cdot D_h(\mathbf{x})$, hence $\tau_h(C_h(\mathbf{x}))$ is ℓ_h -concentrated over \mathcal{R}'_h . This finishes the construction of τ_h , given τ_{h+1} , in time $(d + n + \ell_h)^{O(\ell_h)}$. \square

Lemma 17 (Preserve invertibility). *Let $f \in H_\kappa(\mathbb{F})[\mathbf{x}]$ be a polynomial with degree bound δ . Assume that f is ℓ' -concentrated over $H_\kappa(\mathbb{F})$, and that $f^{-1} \in H_\kappa(\mathbb{F}(\mathbf{x}))$. Then, we can construct an $\alpha \in \mathbb{F}^n$, in time polynomial in $\kappa(\delta + n + \ell')^{\ell'}$, such that $f(\alpha)^{-1} \in H_\kappa(\mathbb{F})$.*

(Proof in Appendix D.)

5. READING OFF THE HITTING-SET

5.1. Proof of Theorem 1. Suppose we are given a blackbox access to a set-height- H nonzero formula C of size s , more so we can think of $C = C_0 \in \mathcal{C}_0(k, d, \lambda, \mathbf{x})$. Using Theorem 15 we can construct a map $\tau_0 : \mathbb{F}[\mathbf{x}] \mapsto \mathbb{F}[\mathbf{t}_0][\mathbf{x}]$ such that $\widehat{D} := \tau_0 \circ D_0$ is ℓ_0 -concentrated over $\mathcal{R}'_1[t_0]$, in time $(d + n + \ell_0)^{O(\ell_0)}$. Clearly, $\widehat{D} \in H_\kappa(\mathbb{F}[\mathbf{t}_0])[\mathbf{x}]$ and $C' := \tau_0 \circ C = \mathbf{c}^T \cdot \widehat{D}$. For $X \subseteq [n]$ of size at most ℓ_0 , define $\sigma_X : x_j \mapsto (x_j \text{ if } j \in X, \text{ else } 0)$ for all $j \in [n]$. Clearly, $\sigma_X \circ C'$ is only ℓ_0 -variate, thus it has sparsity $(d^H + \ell_0)^{O(\ell_0)}$. By the assumption on \widehat{D} we know that there exists such an X for which $\sigma_X \circ C' \neq 0$. Thus, using standard sparse PIT methods (see [BHLV09]) we can construct a hitting-set for C' , in time $(d^H + n + \ell_0)^{O(\ell_0)} = 2^{O(\ell_0 H \log(s + \ell_0))} = \exp(O(\ell_0 H^2 \log s))$, which is time polynomial in $\exp((2H^2 \log s)^{H+1})$. \square

5.2. Proof of Corollary 2. Suppose we are given a blackbox access to a semi-diagonal formula $C = \sum_{i=1}^k m_i \cdot \prod_{j=1}^b f_{i,j}^{e_{i,j}}$ over field \mathbb{F} , where m_i is a monomial, $f_{i,j}$ is a sum of univariate polynomials, and b is a constant. Call its size s .

Assume $p := \text{char}(\mathbb{F})$ is zero (or larger than $\max_{i,j} \{e_{i,j}\}$). Using the *duality* trick (see [SSS12, Theorem 2.1]), there exists another representation of C as $C' := \sum_{i=1}^{k'} \prod_{j=1}^n g_{i,j}(x_j)$ of size $s^{O(b)}$. Rewrite this, using the obvious Hadamard algebra $H_{k'}(\mathbb{F})$, as $C' = \mathbf{c}^T \cdot D$, where $D = G_1(x_1) \star \cdots \star G_n(x_n) \in H_{k'}(\mathbb{F})[\mathbf{x}]$. Trivially, the monomial-weight of each G_j is bounded by 1. Thus, by invoking Theorem 9 (& the ‘key argument’ in Lemma 16) we can shift D , in time $s^{O(\log k')}$, such that it becomes $O(\log k')$ -concentrated. On top of the shift, the usual sparse PIT gives a hitting-set for C in time $s^{O(\log s)}$. \square

5.3. Proof of Corollary 3. Suppose we are given a blackbox access to the formula $C = \sum_{i=1}^k \prod_{j=1}^d f_{i,j}(\mathbf{x}_{X_j})^{e_{i,j}}$, where $f_{i,j}$ is a sum of univariate polynomials in $\mathbb{F}[\mathbf{x}_{X_j}]$, $e_{i,j} \in \mathbb{N}$, and $X_1 \sqcup \cdots \sqcup X_d$ partitions $[n]$. Let the formula size be s .

Assume $\text{char}(\mathbb{F})$ is zero (or larger than $\max_{i,j} \{e_{i,j}\}$). Using the *duality* trick (see [SSS12, Theorem 2.1]), there exists another representation of $f_{i,j}(\mathbf{x}_{X_j})^{e_{i,j}}$ as $F_{i,j} := \sum_{p=1}^{k_{i,j}} \prod_{q \in X_j} g_{i,j,p,q}(x_q)$ of size $s^{O(1)}$. Trivially, the monomial-weight of each $g_{i,j,p,q}$ is bounded by 1. Overall, we can represent C now as $C' := \sum_{i=1}^k \prod_{j=1}^d F_{i,j}$, which is a set-depth-6 formula. Recall the inductive proof of Theorem 15 on C' . It will have $H = 3$ inductive steps. The crucial observation is that in the base case (dealing with sparse polynomials) we can use a better bound $\ell' = 2$ in Lemma 16, as $\mu(g_{i,j,p,q}) \leq 1$. This leads us to an improvement on Theorem 15 - we construct τ_0 such that $\tau_0 \circ D_0$ is $O(\log^2 s)$ -concentrated over $\mathcal{R}'_1[t_0]$, in time polynomial in $s^{\log^2 s}$. Again, on top of the shift, the usual sparse PIT gives a hitting-set for C in time $s^{O(\log^2 s)}$. \square

6. CONCLUSION

We have identified a natural phenomena - low-support rank-concentration - in constant-depth formulas, that is directly useful in their blackbox PIT (up to quasi-polynomial time). In this work we gave a proof for the interesting special case of set-depth- Δ formulas.

More work is needed to prove such rank-concentration in full generality. Next, it would be interesting to prove rank-concentration for depth-3 formulas. Another direction is to improve this proof technique to give polynomial-time hitting-sets for set-depth- Δ formulas.

ACKNOWLEDGMENTS

This work was initiated when MA and NS visited Max Planck Institute for Informatics, and would like to thank the institute for its generous hospitality. The travel of MA was funded by Humboldt Forschungspreis, and that of NS by MPII. CS and NS would like to thank Hausdorff Center for Mathematics (Bonn) for the generous support during the research work. Additionally, CS is supported by the IMPECS fellowship.

REFERENCES

- [AL86] L. M. Adleman and H. W. Lenstra, *Finding irreducible polynomials over finite fields*, Proceedings of the 18th Annual ACM Symposium on Theory of Computing (STOC), 1986, pp. 350–355.
- [AS09] M. Agrawal and R. Saptharishi, *Classifying polynomials and identity testing*, Current Trends in Science, Indian Academy of Sciences, 2009, pp. 149–162.
- [ASSS12] Manindra Agrawal, Chandan Saha, Ramprasad Saptharishi, and Nitin Saxena, *Jacobian hits circuits: hitting-sets, lower bounds for depth- D occur- k formulas & depth-3 transcendence degree- k circuits*, STOC, 2012, pp. 599–614.
- [AV08] M. Agrawal and V. Vinay, *Arithmetic circuits: A chasm at depth four*, Proceedings of the 49th Annual Symposium on Foundations of Computer Science (FOCS), 2008, pp. 67–75.
- [AvMV11] M. Anderson, D. van Melkebeek, and I. Volkovich, *Derandomizing polynomial identity testing for multilinear constant-read formulae*, Proceedings of the 26th Annual IEEE Conference on Computational Complexity (CCC), 2011, pp. 273–282.
- [BHLV09] M. Bläser, M. Hardt, R. J. Lipton, and N. K. Vishnoi, *Deterministically testing sparse polynomial identities of unbounded degree*, Information Processing Letters **109** (2009), no. 3, 187–192.
- [BMS11] M. Beecken, J. Mittmann, and N. Saxena, *Algebraic independence and blackbox identity testing*, Proceedings of the 38th Annual International Colloquium on Automata, Languages and Programming (ICALP), 2011, pp. 137–148.
- [CKW11] Xi Chen, Neeraj Kayal, and Avi Wigderson, *Partial Derivatives in Arithmetic Complexity and Beyond*, Foundations and Trends in Theoretical Computer Science **6** (2011), no. 1-2, 1–138.
- [Ell69] W.J. Ellison, *A ‘Waring’s problem’ for homogeneous forms*, Proceedings of the Cambridge Philosophical Society **65** (1969), 663–672.
- [FS12] M. A. Forbes and A. Shpilka, *On identity testing of tensors, low-rank recovery and compressed sensing*, Proceedings of the Annual 44th ACM Symposium on Theory of Computing (STOC), 2012, pp. 163–172.
- [GKKS12] Ankit Gupta, Pritish Kamath, Neeraj Kayal, and Ramprasad Saptharishi, *An exponential lower bound for homogeneous depth four arithmetic circuits with bounded bottom fanin*, Electronic Colloquium on Computational Complexity (ECCC) (2012), 98.
- [GKQ12] Ankit Gupta, Neeraj Kayal, and Youming Qiao, *Random Arithmetic Formulas can be Reconstructed Efficiently*, Electronic Colloquium on Computational Complexity (ECCC) **19** (2012), 33.
- [KS01] A. Klivans and D. Spielman, *Randomness efficient identity testing of multivariate polynomials*, Proceedings of the 33rd Symposium on Theory of Computing (STOC), 2001, pp. 216–223.
- [KS06] Adam R. Klivans and Amir Shpilka, *Learning restricted models of arithmetic circuits*, Theory of Computing **2** (2006), no. 1, 185–206.
- [NW97] N. Nisan and A. Wigderson, *Lower bounds on arithmetic circuits via partial derivatives*, Computational Complexity **6** (1997), no. 3, 217–234, (Conference version in FOCS 1995).
- [Raz10] R. Raz, *Tensor-rank and lower bounds for arithmetic formulas*, Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC), 2010, pp. 659–666.
- [RS05] R. Raz and A. Shpilka, *Deterministic polynomial identity testing in non-commutative models*, Computational Complexity **14** (2005), no. 1, 1–19, (Conference version in CCC 2004).
- [RY09] Ran Raz and Amir Yehudayoff, *Lower bounds and separations for constant depth multilinear circuits*, Computational Complexity **18** (2009), no. 2, 171–207.

- [Sax08] N. Saxena, *Diagonal circuit identity testing and lower bounds*, Proceedings of the 35th Annual International Colloquium on Automata, Languages and Programming (ICALP), 2008, pp. 60–71.
- [Sax09] ———, *Progress on polynomial identity testing*, Bulletin of the European Association for Theoretical Computer Science (EATCS)- Computational Complexity Column (2009), no. 99, 49–79.
- [Sch80] J. T. Schwartz, *Fast probabilistic algorithms for verification of polynomial identities*, J. ACM **27** (1980), no. 4, 701–717.
- [SS11] N. Saxena and C. Seshadhri, *Blackbox identity testing for bounded top fanin depth-3 circuits: the field doesn't matter*, Proceedings of the 43rd annual ACM Symposium on Theory of Computing (STOC), 2011, pp. 431–440.
- [SSS09] C. Saha, R. Satharishi, and N. Saxena, *The power of depth 2 circuits over algebras*, Proceedings of the 29th Foundations of Software Technology and Theoretical Computer Science (FSTTCS), 2009, pp. 371–382.
- [SSS12] ———, *A case of depth-3 identity testing, sparse factorization and duality*, Computational Complexity (2012), <http://eccc.hpi-web.de/report/2011/021/>.
- [SV85] Sven Skyum and Leslie G. Valiant, *A Complexity Theory Based on Boolean Algebra*, J. ACM **32** (1985), no. 2, 484–502.
- [SY10] A. Shpilka and A. Yehudayoff, *Arithmetic circuits: A survey of recent results and open questions*, Foundations and Trends in Theoretical Computer Science **5** (2010), no. 3–4, 207–388.
- [Val79] L. G. Valiant, *Completeness classes in algebra*, Proceedings of the 11th Annual ACM Symposium on Theory of Computing (STOC), 1979, pp. 249–261.
- [Zip79] R. Zippel, *Probabilistic algorithms for sparse polynomials*, Proceedings of the International Symposium on Symbolic and Algebraic Manipulation (EUROSAM), 1979, pp. 216–226.

APPENDIX A. DIAGONAL CIRCUITS: THE SPIRIT OF THE ARGUMENT

A circuit $C = \sum_{i=1}^k f_i^d$ is a diagonal circuit if f_i is a linear polynomial in n variables, \mathbf{x} .¹ We can associate a formula over a Hadamard algebra with C , namely

$$D(\mathbf{x}) := F^d \quad \text{over } H_k(\mathbb{F}),$$

where $F = z_0 + z_1x_1 + \dots + z_nx_n$, every $z_j \in \mathbb{F}^k$ and F restricted to the i -th coordinate of the vectors z_0, \dots, z_n is the linear polynomial f_i . Clearly, $C = (1, 1, \dots, 1) \cdot D(\mathbf{x})$, where \cdot is the usual matrix product. Assume that $\text{char}(\mathbb{F}) = 0$ or $> d$.

Consider shifting every x_j by a formal variable t_j , i.e. $x_j \mapsto x_j + t_j$. Then,

$$D(\mathbf{x} + \mathbf{t}) = F(\mathbf{x} + \mathbf{t})^d = D(\mathbf{t}) \star (1 + z'_1x_1 + \dots + z'_nx_n)^d =: D(\mathbf{t}) \star D'(\mathbf{x}),$$

where $z'_j = D(\mathbf{t})^{-1}z_j$. We have stated before (in Section 1) that variables would be ultimately shifted by field constants. Here is a way to set t_j a field constant: To ensure that $D(\mathbf{t})^{-1}$ makes sense when t_j 's are set to constants, we map $t_j \mapsto y^j$ where y is a fresh variable and then set y to an $\alpha \in \mathbb{F}$ such that α is not a root of any of the polynomials $f_i(y, y^2, \dots, y^n)$, $1 \leq i \leq k$. With this setting, we can safely assume that $D(\mathbf{t})$ and $z'_1, \dots, z'_n \in H_k(\mathbb{F})$.

Clearly, $C(\mathbf{x} + \mathbf{t}) = (1, 1, \dots, 1) \cdot D(\mathbf{x} + \mathbf{t})$ is zero if and only if $C = 0$. We would like to show that for $\ell = \lceil \log k \rceil$, $C(\mathbf{x} + \mathbf{t})$ is ℓ -concentrated over \mathbb{F} . The coefficient of a monomial $\mathbf{x}^e = \prod_{j \in [n]} x_j^{e_j}$ in $D(\mathbf{x} + \mathbf{t})$ is $D(\mathbf{t}) \star \text{Coef}(e)(D') = \binom{d}{e} D(\mathbf{t}) \star \prod_{j \in [n]} z_j'^{e_j} = \binom{d}{e} D(\mathbf{t}) \star \mathbf{z}'^e$, where $\binom{d}{e} = \binom{d}{e_1, \dots, e_n}$. For a moment, treat \mathbf{z}'^e as a ‘monomial’ in z_1, \dots, z_n . List down all monomials in z_1, \dots, z_n with degree bounded by d in degree-lexicographic order. The idea is to form a basis of $\text{sp}_{\mathbb{F}}\{\text{Coef}(e)(D') \mid e \in \mathbb{N}^n\}$ by picking terms \mathbf{z}'^e , the coefficient of \mathbf{x}^e in D' (upto scaling by $\binom{d}{e}$), from the ordered list. We pick a term $z_{j_1}^{e_1} \dots z_{j_m}^{e_m}$ ($e_j > 0$) from the ordered list if it is not in the span of the already picked terms. The

¹A lemma by Ellison [Ell69] states that every n -variate polynomial of degree d over \mathbb{C} has a diagonal circuit representation although k can be exponentially large.

claim is, if $z_{j_1}^{e_1} \dots z_{j_m}^{e_m}$ ($e_j > 0$) is picked then so are the terms $\prod_{r \in S} z_{j_r}$, for every set $S \subseteq [m]$ - this follows easily from the degree-lexicographic ordering of the list. This implies that $m < \lceil \log k \rceil = \ell$, as dimension of $\text{sp}_{\mathbb{F}}\{\text{Coef}(e)(D') \mid e \in \mathbb{N}^n\}$ is bounded by k and there are 2^m such terms $\prod_{r \in S} z_{j_r}$. Therefore, $D'(\mathbf{x})$ is ℓ -concentrated over $\mathbb{H}_k(\mathbb{F})$ which implies that $D(\mathbf{x} + \mathbf{t}) = D(\mathbf{t}) \star D'(\mathbf{x})$ is ℓ -concentrated over $\mathbb{H}_k(\mathbb{F})$. Since, $C(\mathbf{x} + \mathbf{t}) = (1, 1, \dots, 1) \cdot D(\mathbf{x} + \mathbf{t})$, $C(\mathbf{x} + \mathbf{t})$ is also ℓ -concentrated over \mathbb{F} .

Thus, by shifting $x_j \mapsto x_j + \alpha^j$, where $\alpha \in \mathbb{F}$ is such that none of the $f_i(\alpha, \alpha^2, \dots, \alpha^n)$ is zero, we are guaranteed that the shifted diagonal circuit satisfies $\lceil \log k \rceil$ -concentration. Such an α is always present among a set of $kn + 1$ distinct elements of \mathbb{F} . A quasi-polynomial hitting set generator for $C(\mathbf{x})$ ensues immediately (as sketched in Section 1).

APPENDIX B. MISSING PROOFS OF SECTION 2

B.1. Proof of Lemma 7.

Proof. Recall that $f_j(\mathbf{x}_{X_j}) = (f_{1,j}(\mathbf{x}_{X_j}), \dots, f_{k,j}(\mathbf{x}_{X_j}))^T$, where every $f_{i,j}(\mathbf{x}_{X_j})$ is a set-height- $(H - h - 1)$ formula over \mathcal{R}_h . The proof is by induction on height $(H - h - 1)$ of $f_j(\mathbf{x}_{X_j})$ (in other words, *reverse* induction on h).

Base case ($h + 1 \geq H - 1$): The base case is when $H - h - 1 = 1$ or 0 , i.e. $f_{i,j}(\mathbf{x}_{X_j})$'s are sparse polynomials or linear polynomials depending on whether Δ is even or odd, respectively. In this case, $f_j(\mathbf{x}_{X_j})$ is a set-height- $(H - h - 1)$ formula over \mathcal{R}_{h+1} . Also, the sparsity parameter λ remains the same by its definition. Hence, $f_j(\mathbf{x}_{X_j}) \in \mathcal{C}_{h+1}(k, d, \lambda, \mathbf{x}_{X_j})$. (Here we do not care about the partition.)

Inductive step ($h + 2$ to $h + 1$): The crucial property to note here is that the formulas $f_{i,j}(\mathbf{x}_{X_j})$'s appear as sub-formulas of C_h at depth-3 (Equation 1). Therefore, the *corresponding* Π -layers of $f_{1,j}(\mathbf{x}_{X_j}), \dots, f_{k,j}(\mathbf{x}_{X_j})$ respect the *same* partitions of \mathbf{x}_{X_j} . In particular, we can express every $f_{i,j}(\mathbf{x}_{X_j})$ as,

$$f_{i,j}(\mathbf{x}_{X_j}) = \sum_{p=1}^k b_{i,j,p} \cdot \prod_{q=1}^d g_{i,j,p,q}(\mathbf{x}_{Y_{j,q}}),$$

where $b_{i,j,p} \in \mathcal{R}_h$, $g_{i,j,p,q}(\mathbf{x}_{Y_{j,q}})$ is a set-height- $(H - h - 2)$ formula over \mathcal{R}_h , and the first Π -layer of all $f_{i,j}(\mathbf{x}_{X_j})$, for $1 \leq i \leq k$, respect the same partition $\mathcal{P}_h(2, X_j)$. In other words, $Y_{j,q}$'s partition X_j as do $X_{2,q} \cap X_j$. (Note: With j fixed, here $X_{2,q} \cap X_j$ are the only relevant variable indices.) Hence,

$$(10) \quad f_j(\mathbf{x}_{X_j}) = \sum_{p=1}^k b_{j,p} \cdot \prod_{q=1}^d g_{j,p,q}(\mathbf{x}_{Y_{j,q}}),$$

where $b_{j,p} = (b_{1,j,p}, \dots, b_{k,j,p})^T \in \mathcal{R}_{h+1}$ and $g_{j,p,q}(\mathbf{x}_{Y_{j,q}}) = (g_{1,j,p,q}(\mathbf{x}_{Y_{j,q}}), \dots, g_{k,j,p,q}(\mathbf{x}_{Y_{j,q}}))^T \in \mathcal{R}_{h+1}[\mathbf{x}_{Y_{j,q}}]$.

In order to apply induction, we make a comparison between $f_{i,j}(\mathbf{x}_{X_j})$ and $g_{i,j,p,q}(\mathbf{x}_{Y_{j,q}})$ (and between $f_j(\mathbf{x}_{X_j})$ and $g_{j,p,q}(\mathbf{x}_{Y_{j,q}})$). Just like $f_{i,j}(\mathbf{x}_{X_j})$ is a set-height- $(H - h - 1)$ formula over \mathcal{R}_h occurring as a sub-formula at depth-3 of the formula C_h , $g_{i,j,p,q}(\mathbf{x}_{Y_{j,q}})$ is a set-height- $(H - h - 2)$ formula over \mathcal{R}_h occurring as a sub-formula at depth-5 of the formula C_h . Hence, by induction, $g_{j,p,q}(\mathbf{x}_{Y_{j,q}})$ is a set-height- $(H - h - 2)$ formula in $\mathcal{R}_{h+1}[\mathbf{x}_{Y_{j,q}}]$ with Σ -fanin k , Π -fanin d and sparsity parameter λ i.e., $g_{j,p,q}(\mathbf{x}_{Y_{j,q}}) \in \mathcal{C}_{h+2}(k, d, \lambda, \mathbf{x}_{Y_{j,q}})$, such that every h' -th Π -layer of $g_{j,p,q}(\mathbf{x}_{Y_{j,q}})$ respects the partition $\mathcal{P}_h(h' + 2, Y_{j,q})$. Since

$g_{j,p,q}(\mathbf{x}_{Y_{j,q}})$ has only variables $\mathbf{x}_{Y_{j,q}}$ and $Y_{j,q} \subseteq X_j$, we can also say that every h' -th Π -layer of $g_{j,p,q}(\mathbf{x}_{Y_{j,q}})$ respects the partition $\mathcal{P}_h(h'+2, X_j)$. The h' -th Π -layers of the $g_{j,p,q}(\mathbf{x}_{Y_{j,q}})$'s (for $1 \leq q \leq d$) correspond to the $(h'+1)$ -th Π -layer of $f_j(\mathbf{x}_{X_j})$. Hence, by Equation 10, we infer that every h' -th Π -layer of $f_j(\mathbf{x}_{X_j})$ respects the partition $\mathcal{P}_h(h'+1, X_j)$. Note that the Σ -fanin, Π -fanin and the sparsity parameter remain k, d and λ , respectively. This proves the claim. \square

APPENDIX C. MISSING PROOFS OF SECTION 3

C.1. Proof of Lemma 10.

Proof. Consider a column $u \in \mathcal{S}$ of Z' ; it is z'_u . Now

$$\begin{aligned} z'_u &= f(\mathbf{t})^{-1} \star \sum_{v \in \mathcal{S}} z_v \binom{v}{u} t^{v-u} \quad [\text{by Equation 4}] \\ &= f(\mathbf{t})^{-1} \star \sum_{v \in \mathcal{S}} z_v \cdot t^v \cdot \binom{v}{u} \cdot t^{-u} \\ &= f(\mathbf{t})^{-1} \star Z \cdot (u\text{-th column of } N_{\mathcal{S}} T N_{\mathcal{S}}^{-1}). \end{aligned}$$

Running over all $u \in \mathcal{S}$ gives us the result. \square

C.2. Proof of Lemma 11.

Proof. Lemma 10 gives $Z'_{\mathcal{S}} = f(\mathbf{t})^{-1} \star Z N_{\mathcal{S}} T_{\mathcal{S}, \mathcal{S}} N_{\mathcal{S}}^{-1}$. Rewrite it as,

$$f(\mathbf{t})^{-1} \star Z = Z'_{\mathcal{S}} N_{\mathcal{S}} T' N_{\mathcal{S}}^{-1}.$$

Going modulo the subspace $\text{sp}_{\mathbb{F}(\mathbf{t})}\{z'_0\}$ kills the 0-th column of $Z'_{\mathcal{S}}$ and yields,

$$f(\mathbf{t})^{-1} \star Z \equiv Z'_{\mathcal{S}^*} N_{\mathcal{S}^*} T'_{\mathcal{S}^*, \mathcal{S}} N_{\mathcal{S}}^{-1} \pmod{z'_0}.$$

For the second part we exploit the independence of $T'_{\mathcal{S}^*, \mathcal{S}}$ from Z and the Hadamard algebra. Formally, fix a large enough $\tilde{\kappa}$, say $|\mathcal{S}|$, and the Hadamard algebra $H_{\tilde{\kappa}}(\mathbb{F})$. Let $e \in \mathcal{S}$. Fix \tilde{Z} as: Its e -th column is 0 and the rest are linearly independent modulo 1 (note: $1 = \tilde{z}'_0$). For this ‘generic’ setting we still have the equation, $\tilde{f}(\mathbf{t})^{-1} \star \tilde{Z} \equiv \tilde{Z}'_{\mathcal{S}^*} N_{\mathcal{S}^*} T'_{\mathcal{S}^*, \mathcal{S}} N_{\mathcal{S}}^{-1} \pmod{\tilde{z}'_0}$. Implying,

$$\tilde{f}(\mathbf{t})^{-1} \star \tilde{Z}_{\mathcal{S} \setminus \{e\}} \equiv \tilde{Z}'_{\mathcal{S}^*} N_{\mathcal{S}^*} T'_{\mathcal{S}^*, \mathcal{S} \setminus \{e\}} N_{\mathcal{S} \setminus \{e\}}^{-1} \pmod{\tilde{z}'_0}.$$

Since the LHS is a matrix of rank $|\mathcal{S}| - 1$, we deduce that $T'_{\mathcal{S}^*, \mathcal{S} \setminus \{e\}}$ is invertible. In other words, $T'_{\mathcal{S}^*, \mathcal{S}}$ is strongly full. \square

C.3. Proof of Lemma 12.

Proof. For $i \in [\ell]$, we can apply Lemma 11 to f_i and get,

$$(11) \quad f_i(\mathbf{t})^{-1} \star Z_i \equiv Z'_i N_{\mathcal{S}_i^*} T'_i N_{\mathcal{S}_i}^{-1} \pmod{1}$$

where the 1 is the unity, the all one vector, in $H_{\tilde{\kappa}}(\mathbb{F})$. Denote the u_i -th column of the matrix on the RHS, of the above congruence, by C_{i, u_i} .

Consider a column $u \in \mathcal{S}$ of Z ; it is z_u . Now

$$\begin{aligned}
D(\mathbf{t})^{-1} \star z_u &= \prod_{i \in [\ell]} f_i(\mathbf{t})^{-1} \star z_{i,u_i} \\
&= \prod_{i \in [\ell]} (\alpha_i + C_{i,u_i}) \quad [\text{for some } \alpha_i \in \mathbb{F}(\mathbf{t}) \text{ by Equation 11}] \\
&\equiv \prod_{i \in [\ell]} C_{i,u_i} \pmod{\mathcal{V}_\ell(D')} \quad [\because \text{the product of } \ell \text{ or less } C_{i,u_i} \text{ vanishes}]
\end{aligned}$$

Running over all $u \in \mathcal{S}$ gives us,

$$\begin{aligned}
D(\mathbf{t})^{-1} \star Z &\equiv \otimes_{i \in [\ell]} \left(Z'_i N_{\mathcal{S}_i^*} T'_i N_{\mathcal{S}_i}^{-1} \right) \\
&\equiv \left(\otimes_{i \in [\ell]} Z'_i \right) \cdot \otimes_{i \in [\ell]} \left(N_{\mathcal{S}_i^*} T'_i N_{\mathcal{S}_i}^{-1} \right) \quad [\text{by Lemma 8-(4)}] \\
&\equiv Z' \cdot N_{\mathcal{S}'} \cdot T' \cdot N_{\mathcal{S}}^{-1} \pmod{\mathcal{V}_\ell(D')} \quad [\text{by Lemma 8-(1)}]
\end{aligned}$$

□

C.4. Proof of Theorem 13.

Proof. We know that $T' = \otimes_{i \in [\ell]} T'_i$, where each $T'_i \in (\mathcal{S}_i^* \times \mathcal{S}_i \rightarrow \mathbb{F})$ is strongly full (Lemma 11 for f_i). Thus, we can apply invertible row operations $E_i \in (\mathcal{S}_i^* \times \mathcal{S}_i^* \rightarrow \mathbb{F})$ such that $E_i T'_i$ has a $|\mathcal{S}_i^*|$ -sized identity submatrix, and another column that has only nonzero entries.

Since, from now on, we are not going to use the properties of the index sets $\mathcal{S}_i^*, \mathcal{S}_i$, we replace them by a more readable identification: Define, for $i \in [\ell]$, $n_i := |\mathcal{S}_i^*| > 0$ and identify \mathcal{S}_i^* (resp. \mathcal{S}_i) with $U_i := [n_i]$ (resp. $W_i := [0..n_i]$). Let $U := \times_{i \in [\ell]} U_i$ and $W := \times_{i \in [\ell]} W_i$. Wlog we keep the following setting: For all $i \in [\ell]$,

- (1) $(T'_i)_{U_i, U_i} = I_{n_i}$ [by Lemma 8-(1), and taking $E_i T'_i$ to be our new T'_i], and
- (2) the column $(T'_i)_{U_i, 0}$ is zero free.

Define an *indicator* function (note: $\delta(\cdot)$ equals 1, if the boolean condition is true, else 0)

$$\varepsilon : \mathbb{N}_{>0} \times \mathbb{N} \rightarrow \{0, 1\}; (u, w) \mapsto \delta((w = 0) \vee (w \neq 0 \wedge w = u)).$$

Extend it to $\mathbb{N}_{>0}^\ell \times \mathbb{N}^\ell$ by defining $\varepsilon : (u, w) \mapsto \prod_{r \in [\ell]} \varepsilon(u_r, w_r)$.

Note that the (u, w) -th entry in T'_i is nonzero iff $\varepsilon(u, w) = 1$. Thus, ε *exactly indicates the non-zeroneess in T'_i* .

Similarly, by tensoring, the (u, w) -th entry in $T' \in (U \times W \rightarrow \mathbb{F})$ is nonzero iff $\varepsilon(u, w) = 1$. Thus, ε *exactly indicates the non-zeroneess in T'* .

We will build \mathcal{C} incrementally, starting with $\mathcal{C} = \emptyset$. During this build up we might apply row permutations R on T' .

Consider a column u , $u \in U \subset W$, of T' . This column has exactly one nonzero entry; appearing at the row indexed by $u \in U$. Put all these unmarked columns u in \mathcal{C} , and collect the marked ones in \mathcal{M}_1 .

If $\mathcal{M}_1 = \emptyset$ then we already have $|\mathcal{C}| = |U|$ and we are done (infact, $T'_{U, \mathcal{C}}$ is identity). So assume $|\mathcal{M}_1| =: m_1 \in [\kappa]$ and define $m_2 := \kappa - m_1 < \kappa$. Let the other marked columns be $\mathcal{M}_2 := \mathcal{M} \setminus \mathcal{M}_1$; they lie in $W \setminus U$ and are m_2 many.

Consider the unmarked columns in $W \setminus U$; collect them in $\mathcal{L} := W \setminus (U \cup \mathcal{M}_2)$. We will now focus on the submatrix $T'_{\mathcal{M}_1, W \setminus U} =: T'_1$. Note that its column-indices are ℓ -tuples with at least one zero.

Claim 18. *There exists a row-permutation $R_1 \in \mathbb{F}^{m_1 \times m_1}$, and m_1 unmarked columns $\mathcal{C}_1 \subseteq \mathcal{L}$ such that: $(R_1 T'_1)_{\mathcal{M}_1, \mathcal{C}_1}$ is a lower-triangular $m_1 \times m_1$ matrix with w -th ($w \in \mathcal{C}_1$) diagonal entry being nonzero.*

Proof of Claim 18. We will again build \mathcal{C}_1 incrementally, starting from \emptyset .

Recall that each row of T'_1 is indexed by an ℓ -tuple u in U . For $i \in [\ell]$ we denote the i -th coordinate in u by $u(i)$, and for an $I \subseteq [\ell]$, $u(I)$ denotes the ordered set $\{u(i) | i \in I\}$. For $w \in W$, define the *support* $S(w) := \{i \in [\ell] | w(i) \neq 0\}$. We want to permute the rows so that the coordinates of the row-indices appear in a *decreasing order of frequency*. Formally, pick $R_1 \in \mathbb{F}^{m_1 \times m_1}$ to reorder the rows of T'_1 as $\mathcal{M}_1 = (u_1, \dots, u_{m_1})$ such that:

- The ordered list $u_1(1), \dots, u_{m_1}(1)$ has repetitions only in contiguous locations and the frequencies are non-increasing. In equation terms: The list has some r distinct elements $\alpha_1, \dots, \alpha_r \in U_1$ with respective frequencies $i_1 \geq \dots \geq i_r$ (summing to m_1), and they appear as $\alpha_1(i_1 \text{ times}), \dots, \alpha_r(i_r \text{ times})$.
- The ordered list $(u_1(1), u_1(2)), \dots, (u_{m_1}(1), u_{m_1}(2))$ has repetitions only in contiguous locations and the frequencies are non-increasing.
- The same as above holds for 3-tuples, 4-tuples, \dots, ℓ -tuples.

We now describe an iterative process to build \mathcal{C}_1 one element at a time. In the i -th iteration, $i \in [m_1]$, we will add an unmarked, *unpicked* column $w_i \in \mathcal{L}$ to \mathcal{C}_1 . The process maintains the invariant: $(R_1 T'_1)_{\mathcal{M}_1, \mathcal{C}_1}$ is a *lower-triangular* matrix.

Iteration $i = 1$ - The row u_1 of T'_1 has exactly $2^\ell - 1$ nonzero columns. (Why? Zero-out at least one coordinate of u_1 .) Since $2^\ell - 1 \geq \kappa > |\mathcal{M}_2|$ we can pick a column $w_1 \in \mathcal{L}$ such that $\varepsilon(u_1, w_1) \neq 0$, thus $(T'_1)_{u_1, w_1} \neq 0$. Add w_1 to \mathcal{C}_1 .

Iteration $i \geq 2$ - Consider the list u_1, \dots, u_i . We claim that there are positions $I \subset [\ell]$, $|I| \leq \lceil \lg i \rceil$, such that $u_i(I)$ is not contained in any of the previous sets in the list. The proof is by *binary-search* in the list. Start with $I = \emptyset$. Pick the least $j_1 \in [\ell]$ such that $u_1(j_1), \dots, u_i(j_1)$ are not all the same; add j_1 to I . By the ordering on u 's the frequency μ_1 of $u_i(j_1)$ is at most $i/2$. If it is one then we stop with this I , otherwise we zoom-in on the 'halved' list $u_{i-\mu_1+1}, \dots, u_i$. Again we pick the least $j_2 \in [j_1 + 1, \ell]$ such that $u_{i-\mu_1+1}(j_2), \dots, u_i(j_2)$ are not all the same; add j_2 to I . This leads to a further halving of the list, and so on. Finally, we do have our positions I , $|I| \leq \lceil \lg i \rceil$, such that $u_i(I)$ appears for the first time in u_i .

We deduce that each column w of T'_1 , with $I \subseteq S(w) \subsetneq [\ell]$ and $w(S(w)) = u_i(S(w))$, has the first nonzero entry at the u_i -th row. (Why? Consider $\varepsilon(u_j, w) = \varepsilon(u_j(S(w)), w(S(w))) = \varepsilon(u_j(S(w)), u_i(S(w)))$.) The number of such columns w , that are unmarked and unpicked, is at least $(2^{\ell-|I|} - 1) - m_2 - (i - 1) \geq 2^{\ell-|I|} - \kappa \geq 2^{\ell-\lceil \lg i \rceil} - \kappa \geq 2^{\ell-\lceil \lg \kappa \rceil} - \kappa = 2^{\lceil \lg \kappa \rceil + 1} - \kappa > 0$. So we can pick such a column, say, $w_i \in \mathcal{L} \setminus \mathcal{C}_1$ and add to \mathcal{C}_1 .

Note that the square submatrix of T'_1 thus far, $(R_1 T'_1)_{\{u_1, \dots, u_i\}, \mathcal{C}_1}$ is lower-triangular with a nonzero diagonal.

After the iteration $i = m_1$ - The square matrix $(R_1 T'_1)_{\mathcal{M}_1, \mathcal{C}_1}$ is lower-triangular with a nonzero diagonal.

This finishes the claim. \square

Since R_1 permutes the rows of T'_1 , its action can be lifted to the rows of T' ; call this action R . Also, append \mathcal{C}_1 to the current \mathcal{C} (making its size $|U|$). Define $\overline{\mathcal{M}}_1 := U \setminus \mathcal{M}_1$

and $\bar{\mathcal{C}}_1 := \mathcal{C} \setminus \mathcal{C}_1$. Consider the square matrix $(RT')_{U, \mathcal{C}}$. It looks like,

$$\left[\begin{array}{c|c} (RT')_{\bar{\mathcal{M}}_1, \bar{\mathcal{C}}_1} & (RT')_{\bar{\mathcal{M}}_1, \mathcal{C}_1} \\ \hline (RT')_{\mathcal{M}_1, \bar{\mathcal{C}}_1} & (RT')_{\mathcal{M}_1, \mathcal{C}_1} \end{array} \right] = \left[\begin{array}{c|c} I_{\bar{\mathcal{M}}_1, \bar{\mathcal{C}}_1} & (RT')_{\bar{\mathcal{M}}_1, \mathcal{C}_1} \\ \hline 0_{\mathcal{M}_1, \bar{\mathcal{C}}_1} & (R_1 T'_1)_{\mathcal{M}_1, \mathcal{C}_1} \end{array} \right].$$

Clearly, its determinant equals $|(R_1 T'_1)_{\mathcal{M}_1, \mathcal{C}_1}| \neq 0$. Thus, $|T'_{U, \mathcal{C}}| \neq 0$ and we are done. \square

C.5. Proof of Lemma 14.

Proof. Let a be the v -th column of A . Let $a' \in \mathbb{F}^{|\mathcal{M}|}$ be the vector having the entries of a appearing at the rows \mathcal{M} . Consider $(T' N_S^{-1}) \cdot a$. By the property of a we can write,

$$\begin{aligned} (T' N_S^{-1})a &= (T' N_S^{-1})_{S', v} + (T' N_S^{-1})_{S', \mathcal{M}} \cdot a' \\ &= T'_{S', v} \cdot t^{-v} + (T' N_S^{-1})_{S', \mathcal{M}} \cdot a'. \end{aligned}$$

Thus, the v -th column of A has the leading monomial t^{-v} which ‘contributes’ the vector $T'_{S', v}$. Going over the columns a , running $v \in \mathcal{C}$, by the column-linearity of determinant and the multiplicativity of the inverse-monomial ordering, we deduce that the largest possible (inverse-monomial) term in the expression $|T' N_S^{-1} A|$ is:

$$|T'_{S', \mathcal{C}}| \cdot t^{-\sum_{v \in \mathcal{C}} v}.$$

We know this is nonzero, by the property of \mathcal{C} , thus it is *indeed* the leading term. In particular, $|T' N_S^{-1} A| \neq 0$. \square

APPENDIX D. MISSING PROOFS OF SECTION 4

D.1. Proof of Lemma 16.

Proof. If $2 \lceil \log_2(\kappa \cdot s(f)) \rceil \geq \mu(f)$ then $\ell' = 1 + \mu(f)$. In this case trivially, for any shift σ , $\sigma(f)$ is ℓ' -concentrated over $H_\kappa(\mathbb{F}(t))$. So, from now on we assume $2 \lceil \log_2(\kappa \cdot s(f)) \rceil < \mu(f)$, thus $\ell' = 1 + 2 \lceil \log_2(\kappa \cdot s(f)) \rceil$.

Define $\mathcal{R} := H_{s(f)}(H_\kappa(\mathbb{F}))$. Let $f =: \sum_{e \in S(f)} z_e x^e$. Define a column vector $D \in (S(f) \times [1] \rightarrow H_\kappa(\mathbb{F}[\mathbf{x}]))$ with e -th entry being $z_e x^e$; D can be seen as a polynomial over \mathcal{R} . Rewrite D as a product of univariate polynomials over \mathcal{R} as:

$$D(\mathbf{x}) = g_1(x_1) \star \cdots \star g_n(x_n).$$

Clearly, each g_i has degree, hence sparsity, bounded by δ , and can be seen as an element in $H_{\kappa \cdot s(f)}(\mathbb{F})[x_i]$.

For any $X \subseteq [n]$ of size ℓ' , define $D_X(\mathbf{x}) := \prod_{i \in X} g_i(x_i)$. Recalling Theorem 9 we can construct a shift σ for D_X , such that $\sigma \circ D_X$ is ℓ' -concentrated, in time polynomial in $(\delta + n + \ell')^{\ell'}$. Using induction on the number of variables, it is easy to see that if $\sigma \circ D_X$ is ℓ' -concentrated ($\forall X \in \binom{[n]}{\ell'}$) then so is $\sigma \circ D$. The *key argument* is: Since the constant coefficient in each g'_i (i.e. shift-&-normalized g_i) is one, deduce that the coefficient of any term in D' (i.e. shift-&-normalized D) of block-weight $\leq \ell'$ is produced by the product of some $\leq \ell'$ g'_i 's, so this case is covered by some $X \in \binom{[n]}{\ell'}$. Also, deduce that the coefficient of any term in D' of block-weight $> \ell'$ can be inductively written down as a linear combination of $\{\text{Coef}(e)(D') \mid e \in \mathbb{N}^n, s(e) < \ell'\}$. Finally, $\sigma \circ D$ inherits this concentration property from D' .

Recall $f = 1^T \cdot D$, where 1 is the unity in $\mathcal{R} = H_{s(f)}(H_\kappa(\mathbb{F}))$. Thus, from the ℓ' -concentration of $\sigma \circ D$ (over \mathcal{R}), we can deduce the ℓ' -concentration of $\sigma \circ f$ (over $H_\kappa(\mathbb{F})$). This completes the construction of σ . \square

D.2. Proof of Lemma 17.

Proof. View f as a vector with κ coordinates; each entry is in $\mathbb{F}[\mathbf{x}] \setminus \{0\}$. Call the i -th entry f_i . Clearly, f_i has variables (resp. degree) at most n (resp. δ). Also, by the concentration property there exists $e_i \in \mathbb{N}^n$, with $s(e_i) \leq \ell'$, such that $\text{Coef}(e_i)(f_i) \neq 0$.

For $X \subseteq [n]$ of size at most ℓ' , define $\sigma_X : x_j \mapsto (x_j \text{ if } j \in X, \text{ else } 0)$ for all $j \in [n]$. Clearly, $\sigma_X \circ f_i$ is only ℓ' variate, thus it has sparsity $(\delta + \ell')^{O(\ell')}$. By the assumption on f_i we know that $X_i := S(e_i)$ is of size at most ℓ' , and $\sigma_{X_i} \circ f_i \neq 0$. Using standard sparse PIT methods (see [BHLV09]), we can construct a hitting-set for $\sigma_{X_i} \circ f_i$ in time $(\delta + \ell')^{O(\ell')}$. Varying over all subsets $X \subseteq [n]$ of size at most ℓ' , we get a hitting-set for f_i in time $(\delta + n + \ell')^{O(\ell')}$. For convenience, denote this hitting-set as a set of evaluation-maps $\{\sigma_{i,1}, \dots, \sigma_{i,r}\}$; each map is from \mathbf{x} to \mathbb{F} and we write $\sigma_{i,j} \circ f_i$ to mean $f_i(\sigma_{i,j}(\mathbf{x}))$. Overall we are ensured the existence of a j , for a given i , such that $\sigma_{i,j} \circ f_i \neq 0$. We will now show how to combine all these into a single map.

Pick distinct κr elements $\beta_{1,1}, \dots, \beta_{\kappa,r} \in \mathbb{F}$. Consider the univariate polynomial $g(u) := \prod_{i \in [\kappa], j \in [r]} (u - \beta_{i,j})$. Define $g_{i,j}(u) := g(u)/(u - \beta_{i,j})$, for all i, j . Consider an evaluation map from $\mathbb{F}[\mathbf{x}]$ to $\mathbb{F}[u, v] - \sigma := v \cdot \sum_{i \in [\kappa], j \in [r]} g_{i,j}(u) \cdot \sigma_{i,j}$. We claim that, for all $i \in [\kappa]$, $\sigma \circ f_i \neq 0$. To see this, note that there is some $j \in [r]$ for which $\sigma_{i,j} \circ f_i \neq 0$. Further, let f'_i be a homogeneous part of f_i , say of degree δ_i , such that $\sigma_{i,j} \circ f'_i \neq 0$. Consider the partial evaluation $(\sigma \circ f_i)(\beta_{i,j}, v) = f_i(v \cdot g_{i,j}(\beta_{i,j}) \cdot \sigma_{i,j}(\mathbf{x}))$. Here the coefficient of the monomial v^{δ_i} is $g_{i,j}(\beta_{i,j})^{\delta_i} \cdot (\sigma_{i,j} \circ f'_i) \neq 0$. Consequently, $\sigma \circ f_i \neq 0$.

Thus, for all $i \in [\kappa]$, $\sigma \circ f_i$ is a nonzero bivariate polynomial in $\mathbb{F}[u, v]$. Since its degree remains bounded by $\delta \cdot \kappa r$, we can again apply [BHLV09] to replace u, v by a hitting-set. Finally, we hit an $\alpha \in \mathbb{F}^n$, in time polynomial in $\kappa(\delta + n + \ell')^{\ell'}$, such that for all $i \in [\kappa]$, $f_i(\alpha) \neq 0$. This finishes the proof. \square

INDIAN INSTITUTE OF TECHNOLOGY, KANPUR, INDIA.
E-mail address: manindra@iitk.ac.in

MAX PLANCK INSTITUT FÜR INFORMATIK, SAARBRÜCKEN, GERMANY.
E-mail address: csaha@mpi-inf.mpg.de

HAUSDORFF CENTER FOR MATHEMATICS, BONN, GERMANY.
E-mail address: ns@hcm.uni-bonn.de