Electronic Colloquium on Computational Complexity, Revision 6 of Report No. 147 (2023)

Black-Box Identity Testing of Noncommutative Rational Formulas in Deterministic Quasipolynomial Time*

V. Arvind† Abhranil Chatterjee‡ Partha Mukhopadhyay§

Abstract

Rational Identity Testing (RIT) is the decision problem of determining whether or not a noncommutative rational formula computes zero in the free skew field. It admits a deterministic polynomial-time white-box algorithm [GGdOW16, IQS18, HH21], and a randomized polynomial-time algorithm [DM17] in the black-box setting, via singularity testing of linear matrices over the free skew field. Indeed, a randomized NC algorithm for RIT in the white-box setting follows from the result of Derksen and Makam [DM17].

Designing an efficient deterministic black-box algorithm for RIT and understanding the parallel complexity of RIT are major open problems in this area. Despite being open since the work of Garg, Gurvits, Oliveira, and Wigderson [GGdOW16], these questions have seen limited progress. In fact, the only known result in this direction is the construction of a quasipolynomial-size hitting set for rational formulas of only *inversion height* two [ACM22].

In this paper, we significantly improve the black-box complexity of this problem and obtain the first quasipolynomial-size hitting set for *all* rational formulas of polynomial size. Our construction also yields the first deterministic quasi-NC upper bound for RIT in the white-box setting.

*A preliminary version of this work appeared in the proceedings of the *56th Annual ACM Symposium on Theory of Computing (STOC, 2024)*.

†Institute of Mathematical Sciences (HBNI), and Chennai Mathematical Institute, Chennai, India. Email: arvind@imsc.res.in.

‡Indian Institute of Technology, Kanpur, India. Email: abhneil@gmail.com.

§Chennai Mathematical Institute, Chennai, India. Email: partham@cmi.ac.in.

ISSN 1433-8092

Contents

1 Introduction

The goal of algebraic circuit complexity is to understand the complexity of computing multivariate polynomials and rational expressions using basic arithmetic operations, such as additions, multiplications, and inverses. Algebraic formulas and algebraic circuits are some of the well-studied computational models.

In the commutative setting, the role of inverses is well understood, but in noncommutative computation it is quite subtle. To elaborate, it is known that *any* commutative rational expression can be expressed as fg^{-1} where f and g are two commutative polynomials [Str73]. However, noncommutative rational expression such as $x^{-1} + y^{-1}$ or $xy^{-1}x$ cannot be represented as fg^{-1} or $f^{-1}g$ for any noncommutative polynomials f and g. Therefore, the presence of *nested inverses* makes a rational expression more complicated, for example $(z + xy^{-1}x)^{-1} - z^{-1}$.

A noncommutative rational expression is not always defined on a matrix substitution. For a noncommutative rational expression Φ, its *domain of definition* is the set of matrix tuples (of any dimension) where Φ is defined.

Two rational expressions Φ_1 and Φ_2 are *equivalent* if they agree on every matrix substitution in the intersection of their domains of definition. This induces an equivalence relation on the set of all noncommutative rational expressions (with nonempty domains of definition). Interestingly, this computational definition was used by Amitsur in the characterization of the *universal* free skew field [Ami66]. The free skew field consists of these equivalence classes, called *noncommutative rational functions*. One can think of the free skew field $\mathbb{F}\langle\!\langle x_1, \ldots, x_n \rangle\!\rangle$ as the smallest field that contains the noncommutative polynomial ring $\mathbb{F}\langle x_1, \ldots, x_n\rangle$. The free skew field been extensively studied in mathematics [Ami66, Coh71, Coh95, FR04].

The complexity-theoretic study of noncommutative rational functions was initiated by Hrubeš and Wigderson [HW15]. Computationally (and in this paper), noncommutative rational functions are represented by algebraic formulas using addition, multiplication, and inverse gates over a set of noncommuting variables, and they are called noncommutative rational formulas. Hrubeš and Wigderson [HW15] also addressed the *rational identity testing* problem (RIT): decide efficiently whether a given noncommutative rational formula Φ computes the zero function in the free skew field. Equivalently, the problem is to decide whether Φ is zero on its domain of definition, follows from Amitsur's characterization [Ami66].

For example, the rational expression $(x + xy^{-1}x)^{-1} + (x + y)^{-1} - x^{-1}$ is a rational identity, known as Hua's identity [Hua49]. Rational expressions exhibit peculiar properties which seem to make the RIT problem quite different from the noncommutative polynomial identity testing. For example, Bergman has constructed an explicit rational formula, of inversion height two, which is an identity for 3×3 matrices but not an identity for 2×2 matrices [Ber76]. Also, the apparent lack of *canonical representations*, like a sum of monomials representation for polynomials, and the use of nested inverses in noncommutative rational expressions complicate the problem. This motivates the definition of *inversion height* of a rational formula which is the maximum number of inverse gates in a path from an input gate to the output gate. The *inversion height* of a rational function is the minimum over all the inversion heights of formulas representing the function. For example, consider the rational expression $(x + xy^{-1}x)^{-1}$. Athough it has a nested inverse, by Hua's identity it represents a rational function of inversion height one. In fact, Hrubeš and Wigderson obtain the following interesting bound on the inversion height of any rational function [HW15]. This is obtained by adapting Brent's depth reduction for the commutative formulas [Bre74].

Fact 1. *For any noncommutative n-variate rational formula Φ_1 of size s, one can construct a rational formula Φ_2 of size s with the following properties:*

1. *Both Φ_1 and Φ_2 compute the same rational function.*
2. *The domains of definition of Φ_1 and Φ_2 are identical.*
3. *The inversion height of Φ_2 is at most $O(\log s)$.*

Consequently, to design a black-box RIT algorithm for rational formulas of size at most s, it suffices to construct a hitting set for rational formulas of inversion height at most $O(\log s)$. This upper bound is crucial for our proof.

Hrubeš and Wigderson have given an efficient reduction from the RIT problem to the singularity testing problem of linear matrices in noncommuting variables over the free skew field (NSingular). Equivalently, given a linear matrix $T = A_1x_1 + \ldots + A_nx_n$ over noncommuting variables $\{x_1, x_2, \ldots, x_n\}$, the problem NSingular is to decide whether or not there is a matrix substitution $(p_1, \ldots, p_n)$ such that $\det\left(\sum_{i=1}^{n} A_i \otimes p_i\right) \neq 0$ [IQS18]. It is the noncommutative analogue of Edmonds' problem of symbolic determinant identity testing (Singular). While Singular can be easily solved in randomized polynomial time using Polynomial Identity Lemma [DL78, Zip79, Sch80], finding a deterministic polynomial-time algorithm remains completely elusive [KI04].

Remarkably, NSingular $\in$ P thanks to two independent breakthrough results [GGdOW16, IQS18]. In particular, the algorithm of Garg, Gurvits, Oliveira, and Wigderson [GGdOW16] is analytic in nature and based on operator scaling which works over $\mathbb{Q}$. The algorithm of Ivanyos, Qiao, and Subrahmanyam [IQS18] is purely algebraic. Moreover, the algorithm in their paper [IQS18] works over both $\mathbb{Q}$ and fields of positive characteristic. Subsequently, a third algorithm based on convex optimization was developed by Hamada and Hirai [HH21]. Not only are these beautiful results, but also they have enriched the field of computational invariant theory greatly [BFG+19, DM20, MW19]. As an immediate consequence, RIT can also be solved in deterministic polynomial time in the *white-box* setting. Both problems admit a randomized polynomial-time black-box algorithm due to Derksen and Makam [DM17]. Essentially, the result of [DM17] shows that to test whether a rational formula of size s is zero or not (more generally, whether a linear matrix of size $2s$ is invertible or not over the free skew field), it suffices to evaluate the formula (resp. the linear matrix) on random $2s \times 2s$ matrices.

Two central open problems in this area are to design faster deterministic algorithms for the NSingular problem and RIT problem in the black-box setting, raised in [GGdOW16, GGdOW20]. The algorithms in [GGdOW16] and [IQS18] are inherently sequential and seem unlikely to be helpful for designing a subexponential-time black-box algorithm. Even for the RIT problem (which could be easier than the NSingular problem), there is limited progress towards designing an efficient deterministic black-box algorithm. A deterministic quasipolynomial-time black-box algorithm is known for identity testing of rational formulas of inversion height two [ACM22]. More recently, it is known that certain ABP (algebraic branching program)-hardness of polynomial identities (PI) for matrix algebras will yield a black-box subexponential-time derandomization of RIT in almost general setting [ACG+23]. However, the result is conditional as such a hardness result is not proven. It is interesting to note that in the literature of identity testing, the NSingular problem and the RIT problem are rare examples of problems with deterministic polynomial-time white-box algorithms but no known deterministic subexponential-time black-box algorithms.

It is well-known [GGdOW16] that an efficient black-box algorithm (via a hitting set construction) for NSingular would generalize the celebrated quasi-NC algorithm for bipartite perfect matching significantly [FGT16]. This motivates the study of the parallel complexity of NSingular and RIT. From the result of Derksen and Makam [DM17], one can observe that RIT in the white-box setting can be solved in randomized NC which involve formula evaluation, and matrix operations (addition, multiplication, and inverse computation) [Bre74, Csa76, Ber84, HW15].[1] Designing a hitting set in quasi-NC for this problem would therefore yield a deterministic quasi-NC algorithm for this problem.

1.1 Our Results

In this paper, we focus on the RIT problem and improve the black-box complexity significantly by showing the following result.

Theorem 2. *For n-variate noncommutative rational formulas of size s and inversion height θ, we can construct a hitting set $\mathcal{H}_{n,s,\theta} \subseteq \mathrm{Mat}_{\ell_\theta}(\mathbb{Q})^n$ of size $(ns)^{O(\theta^5 \log s)}$ in deterministic $(ns)^{O(\theta^5 \log s)}$ time where $\ell_\theta = s^{O(\theta)}$.*

Here $\mathrm{Mat}_{\ell_\theta}(\mathbb{Q})$ denotes the ℓ_θ-dimensional matrix algebra over $\mathbb{Q}$. An immediate corollary of Theorem 2 and Fact 1 is the following.

Corollary 3 (black-box RIT)**.** *In the black-box setting, RIT can be solved in deterministic quasipolynomial time via an explicit hitting set construction.*

Note that even for noncommutative formulas i.e. when the inversion height $\theta = 0$, the best known hitting set is of quasipolynomial size and improving it to a polynomial-size hitting set is a long standing open problem [FS13]. In this light, Theorem 2 is nearly the best result one can hope for, apart from further improving the logarithmic factors on the exponent.

We further show that our hitting set construction can in fact be performed in quasi-NC. In the white-box setting, we can evaluate a given rational formula on the hitting set points in parallel. This involves the evaluation of the formula in parallel, and supporting matrix addition, multiplication, and inverse computation. It is already observed that Brent's formula evaluation [Bre74] can be adapted to the setting of noncommuatative rational formulas [HW15], and such matrix operations can be performed in NC [Csa76, Ber84]. Combining these results with the quasi-NC construction of the hitting set, we obtain the following corollary.

Corollary 4 (white-box RIT)**.** *In the white-box setting RIT is in deterministic quasi-*NC*.*

1.2 Proof Idea

The main idea is to construct a hitting set for rational formulas of every inversion height inductively. Our goal is now to construct a hitting set for rational formulas of inversion height θ given a hitting set for formulas of height $\theta - 1$. To design a black-box RIT algorithm for rational formulas of size

[1] Similarly NSingular is also in randomized NC via the determinant computation [Ber84, DM17].

at most s, it suffices to construct a hitting set for rational formulas of inversion height at most $O(\log s)$, we can stop the induction at that stage.

As already mentioned, a noncommutative rational formula is nonzero in the free skew field if it evaluates to nonzero for some matrix substitution. However, the difficulty is that a rational formula may be undefined for a matrix substitution. For instance, this will happen if there is an inversion gate to which the input subformula evaluated to a singular matrix. Informally speaking, it is somewhat easier to maintain that nonzero subformulas evaluate to nonzero matrices, but it is much harder to ensure that they evaluate to *non-singular* matrices. Therefore, a rational formula of inversion height θ may not even be defined on any matrix tuple in the hitting set of rational formulas of inversion height $\theta - 1$.

One way to handle this difficulty is to evaluate rational formulas on elements of a division algebra. Finite dimensional division algebras are associative algebras in which every nonzero elements are invertible.

This idea of embedding inside a division algebra proves to be very useful. We formally define the notion of hitting set for rational formulas (of any arbitrary inversion height) inside a division algebra.

Definition 5 (Division algebra hitting set)**.** For a class of n-variate rational formulas, a *division algebra hitting set* is a hitting set for that class of formulas that is contained in D^n for some finite-dimensional division algebra D.

The advantage of a division algebra hitting set is that, whenever a rational formula evaluates to some nonzero value (at an n-tuple in the hitting set), the output is invertible.[2]

Therefore every n-variate rational formula of inversion height θ is defined on some n-tuple in any given division algebra hitting set for rational formulas of inversion height $\theta - 1$. Can we build on this to efficiently construct a division algebra hitting set for rational formulas of inversion height θ? In that case, we could inductively construct a division algebra hitting set for rational formulas of every inversion height. For the base case of the induction, we want to construct a division algebra hitting set for noncommutative formulas. An earlier work [ACM22] shows this by essentially embedding the hitting set obtained by Forbes and Shpilka [FS13] for noncommutative *polynomials* computed by algebraic branching programs (ABPs) in a cyclic division algebra (see Section 2.3 for the definition of a cyclic division algebra) of suitably small index i.e. the dimension of its matrix representation. The inductive construction of the division algebra hitting set is the main technical step we implement here using several conceptual and technical ideas.

At this point, we take a detour and examine the connection between the RIT and NSingular problems. It is known that RIT is polynomial-time reducible to NSingular [HW15]. But do we need the full power of NSingular to solve the RIT problem for rational formulas of height θ given a hitting set for formulas of height $\theta - 1$? Consider the following promise version of NSingular. The input is a linear matrix $T(x_1, \ldots, x_n)$ and a matrix tuple $(p_1, \ldots, p_n) \in D_1^n$ for a cyclic division algebra D_1. The promise is that there is a submatrix T' of size $s - 1$ (obtained by removing the i^{th}

[2]Division algebras are used in [IQS18] as part of the regularity lemma and their noncommutative rank algorithm. Roughly speaking, after each rank-increment step, if the matrix substitutions are d-dimensional these are replaced by d-dimensional division algebra elements. This is a *rank rounding* step which ensures that the rank is always a multiple of d after each increment. This is a white-box process as it crucially uses the structure of the linear matrix. In contrast, our paper deals with the black-box setting where we do not have an explicit description of the rational formula (and the corresponding linear matrix).

row and j^{th} column, for some $i, j \in [s]$) such that $T'(p_1, \ldots, p_n)$ is invertible. It is easier to think such a tuple $(p_1, \ldots, p_n)$ as a *witness*. The question is now to check the singularity of T over the free skew field. We show that the construction of a hitting set for rational formulas of inversion height θ inductively reduces to this special case where the witness is some tuple in the hitting set for height $\theta - 1$. We then consider the shifted matrix $T(x_1 + p_1, \ldots, x_n + p_n)$. Applying Gaussian elimination, we can convert the shifted matrix of form:

$$U \cdot T(x_1 + p_1, \ldots, x_n + p_n) \cdot V = \left[\begin{array}{c|c} I_{s-1} - L & A_j \\ \hline B_i & C_{ij} \end{array}\right], \tag{1}$$

where the entries of L, A_j, B_i, C_{ij} are homogeneous D_1-linear forms. Here B_i is a row vector and A_j is a column vector. At a high level, this is conceptually similar to an idea useful for approximating commutative rank [BBJP19].

It is not too difficult to prove that T is invertible, if and only if, $C_{ij} - B_i(I_{s-1} - L)^{-1}A_j = C_{ij} - B_i(\sum_{k \geqslant 0} L^k)A_j$ is a nonzero series. By a standard result on noncommutative formal series [Eil74, Corollary 8.3], this is equivalent to saying that the truncated polynomial $C_{ij} - B_i(\sum_{k \leqslant (s-1)\ell} L^k)A_j$ is nonzero where ℓ is the index of D_1. However, this series and the polynomial obtained will have division algebra elements interleaved with variables. Such a series (resp. polynomial) is called a generalized series (resp. generalized polynomial) and is studied extensively, for example, in the work of Volčič [Vol18]. We define similar notion of generalized ABPs and show that the truncated generalized polynomial of our interest is indeed computable by a polynomial-size generalized ABP. Finally, (up to a certain scaling by scalars) the upshot is that the division algebra hitting set construction for rational formulas is inductively reducible to the division algebra hitting set construction for such generalized ABPs.

We now consider such generalized ABPs where the coefficients lie inside a cyclic division algebra D of index ℓ, and call such ABPs as D-ABPs. Our goal is to construct a division algebra hitting set for such ABPs. To do so, a key conceptual idea is to introduce new noncommuting indeterminates for every variable and use the following mapping:

$$x_i \mapsto \sum_{j,k=1}^{\ell_1} C_{jk} \otimes y_{ijk},$$

where the matrices $\{C_{jk}\}$ form a basis for the division algebra D. The idea is to overcome the problem of interleaving division algebra elements using the property of tensor products. This substitution reduces the problem to identity testing of a noncommutative ABP in the $\{y_{ijk}\}$ variables. Luckily, a division algebra hitting set construction for noncommutative ABPs is already known [ACM22]. However, we need the hitting set to be inside a division algebra that contains D as a subalgebra. A natural thought could be to take the tensor product of D and the division algebra, say D', that defines the division algebra hitting set for the noncommutative ABP in the $\{y_{ijk}\}$ variables. However, in general, the tensor product of two division algebras is not a division algebra. At this point, we use a result [Pie82, Proposition, pg 292] that the tensor product of two cyclic division algebras of relatively prime indices ℓ_1 and ℓ_2, respectivly, is a cyclic division algebra of index $\ell_1\ell_2$. However, the division algebra hitting set construction [ACM22] for noncommutative ABPs was only for division algebras whose index is only a power of two. Thus, in order to apply the above tensor product construction [Pie82] in several stages recursively, we need a division algebra hitting set construction for a division algebra whose index is a power of any *arbitrary* prime $\mathfrak{p}$.

We now informally describe how to construct such a division algebra hitting set for noncommutative formulas (more generally for noncommutative ABPs). For simplicity, suppose the prime is p and the ABP degree is p^d. In [FS13], it is assumed that the degree of the ABP is 2^d and the construction has a recursive structure. In fact, itt is essentially a reduction to the hitting set construction for ROABPs (read-once oblivious algebraic branching programs) in commuting variables $u_1, u_2, \ldots, u_{2^d}$. The recursive step in their construction is by combining hitting sets (via hitting set generator $\mathcal{G}_{d-1}$) for two halves of degree 2^{d-1} [FS13] with a rank preserving step of matrix products to obtain the generator $\mathcal{G}_d$ at the d^{th} step. More precisely, $\mathcal{G}_d$ is a map from $\mathbb{F}^{d+1} \to \mathbb{F}^{2^d}$ that stretches the seed $(\alpha_1, \ldots, \alpha_{d+1})$ to a 2^d tuple for the read-once variables.

For our case, the high-level idea is to decompose the ABP of degree p^d in p consecutive parts each of length p^{d-1}. We can adapt the rank preserving step for *two matrix products* in [FS13] to *p many matrix products*. However, the main task is to ensure that the hitting set points lie inside a division algebra. For our purpose, we take a classical construction of cyclic division algebras [Lam01, Chapter 5]. The cyclic division algebra $D = (K/F, \sigma, z)$ is defined using a indeterminate x as the ℓ-dimensional vector space:

$$D = K \oplus Kx \oplus \cdots \oplus Kx^{\ell-1},$$

where the (noncommutative) multiplication for D is defined by $x^\ell = z$ and $xb = \sigma(b)x$ for all $b \in K$. Here $\sigma : K \to K$ is an automorphism of K that fixes F, such that σ generates the Galois group $\mathrm{Gal}(K/F)$. The fields we choose are $F = \mathbb{Q}(\omega_0, z)$ and $K = F(\omega)$, where z is an indeterminate, ω_0 is a root of unity whose order is relatively prime to p, and ω is an ℓ^{th} primitive root of unity where ℓ is a suitable power of p. The matrix representation of the division algebra element xb for any $b \in K$ is of particular interest in this paper. It has the following circulant form:

$$\begin{bmatrix} 0 & b & 0 & \cdots & 0 \\ 0 & 0 & \sigma(b) & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & \sigma^{\ell-2}(b) \\ z\sigma^{\ell-1}(b) & 0 & \cdots & 0 & 0 \end{bmatrix}.$$

Now, the broad idea to get the hitting set points inside a division algebra of index ℓ that is a power of p is as follows. Inductively assume that the construction follows the σ-automorphism (as in the matrix structure above) for each part of the generator output of length p^{d-1}. Then we need to satisfy the σ-action at the $p-1$ boundaries between the successive p^{d-1} length parts. In order to follow this inductive structure, we think of the division algebra construction as a tower of extension fields of F, with a higher-order root of unity at each stage.

Specifically, let $\omega_i = \omega^{p^{a_i}}$ for $a_1 > a_2 > \cdots > a_d > a_{d+1} > 0$, where a_i are positive integers suitably chosen. Let $K_i = F(\omega_i)$ be the cyclic Galois extension for $1 \leqslant i \leqslant d+1$ giving a tower of extension fields

$$F \subset F(\omega_1) \subset F(\omega_2) \subset \cdots \subset F(\omega_d) \subset F(\omega_{d+1}) \subset F(\omega).$$

We require two properties of $\omega_i, 1 \leqslant i \leqslant d+1$. Firstly, for the hitting set generator $\mathcal{G}_i$ we will choose the root of unity as ω_i and the variable α_i will take values only in the set $W_i = \{\omega_0^{j'}\omega_i^{j} \mid 1 \leqslant j \leqslant p^{L-a_i} 0 \leqslant j' \leqslant \operatorname{ord} \omega_0 - 1\}$. We also require that for $1 \leqslant i \leqslant d+1$ the map σ^{p^i} has $F(\omega_i)$ as its fixed field.

The construction of matrix tuples in D satisfying the above properties is the main technical step in Theorem 33. It turns out that choosing ω_0 suitably plays an important role in the analysis of the final hitting set size (see Section 5 for details).

We now conclude this section with a summary of the key steps involved in the hitting set construction.

Informal summary.

1. Division algebra hitting set construction for generalized ABPs defined over cyclic division algebras (Section 3).
   (a) Given any prime $\wp$, we build hitting set for noncommutative ABPs inside a cyclic division algebra whose index is a power of $\wp$ (Theorem 33).
   (b) We reduce the hitting set problem for generalized ABPs to that of noncommutative ABPs using the map $x_i \mapsto \sum C_{jk} \otimes y_{ijk}$ (the key idea in Theorem 41).
2. We construct a hitting set conditioned on a witness point for NSingular problem (see Section 4 for details). This builds on the construction of division algebra hitting set for generalized ABPs in Step 1 (Theorem 46).
3. We construct the hitting set for rational formulas by induction on the inversion height θ. The final hitting set construction and analysis are presented in Section 5.1.

1.3 Related results

The basic result underlying the hitting set construction in this paper is the Forbse-Shpilka hitting set construction [FS13]. At a high level, our approach builds on the framework introduced in [ACM22, ACG+23]. In [ACM22], the authors present a hitting set construction for rational formulas of inversion height two. The main ingredient in their construction is a division algebra hitting set construction for noncommutative formulas (more generally, for noncommutative ABPs). Additionally, they proposed an approach to constructing a hitting set by induction on inversion height as a possible approach to derandomize RIT in the black-box setting. Unfortunately they could not obtain a division algebra hitting set even for rational formulas of inversion height one. In [ACG+23], the authors have shown a conditional result. Assuming a conjecture on hardness of polynomial identities [BW05] they obtain a hitting set for every inversion height by induction. A limitation of that construction is that it is based on an unproven conjecture, and even so it only yields a quasipolynomial-size hitting set for rational formulas that is barely more than constant inversion height.

In this paper, we are able to overcome both the difficulties as we unconditionally build the quasipolynomial-size hitting set for *all* polynomial-size rational formulas.

As already mentioned, the results of [GGdOW20, IQS18, HH21] solve the more general NSingular problem in order to solve the RIT problem in the white-box setting.

In contrast, our hitting set construction crucially uses the inversion height of the input rational formula inductively. Furthermore, the hitting set construction for the NSingular problem is known for the following special cases only: when the input matrix is a symbolic matrix (for which bipartite perfect matching is a special case) [FGT16], or more generally, when the input matrix consists of rank-1 coefficient matrices (for which linear matroid intersection is a special case) [GT20], and more recently, when the input matrix consists of rank-2 skew-symmetric coefficient matrices (fractional linear matroid matching is a special instance of it) [GOR24].[3]

An exponential lower bound on the size of the rational formula computed as an entry of the inverse of a symbolic matrix is known [HW15]. Therefore, our hitting set construction does not subsume these results. Similarly, it is quite unlikely to reduce the RIT problem (in the general setting) to any of these special cases of NSingular problem. Thus it seems that these results are incomparable.

1.4 Organization

In Section 2, we provide a background on algebraic complexity theory, cyclic division algebras, and noncommutative formal power series. hitting set for noncommutative ABPs over cyclic division algebras whose index is any prime power. In Section 3 we give, in a self-contained presentation, the construction of a hitting set for generalized ABPs defined over cyclic division algebras. In Section 4, we obtain for the NSingular problem a hitting set conditioned on a witness point. The main result and analysis are in Section 5. In Section 5.1 we present the proof of our main result : a hitting set for arbitrary rational formulas (Theorem 2). Section 5.2 contains the proof of Corollary 4. Finally, we raise a few questions for further research in Section 6.

2 Preliminaries

2.1 Notation

Throughout the paper, we use $\mathbb{F}, F, K$ to denote fields, and $\mathrm{Mat}_m(\mathbb{F})$ (resp. $\mathrm{Mat}_m(F), \mathrm{Mat}_m(K)$) to denote m-dimensional matrix algebra over $\mathbb{F}$ (resp. over F, K). Similarly, $\mathrm{Mat}_m(\mathbb{F})^n$ (resp. $\mathrm{Mat}_m(F)^n, \mathrm{Mat}_m(K)^n$) denote the set of n-tuples over $\mathrm{Mat}_m(\mathbb{F})$ (resp. $\mathrm{Mat}_m(F), \mathrm{Mat}_m(K)$), respectively. D is used to denote finite-dimensional division algebras. We use $\wp$ to denote an arbitrary prime number. Let $\underline{x}$ denote the set of variables $\{x_1, \ldots, x_n\}$. Sometimes we use $\underline{p} = (p_1, \ldots, p_n)$ and $\underline{q} = (q_1, \ldots, q_n)$ to denote the matrix tuples in suitable matrix algebras where n is clear from the context. The free noncommutative ring of polynomials over a field $\mathbb{F}$ is denoted by $\mathbb{F}\langle \underline{x} \rangle$. For matrices A and B, their usual tensor product is denoted by $A \otimes B$. For a polynomial f and a monomial m, we use $[m]f$ to denote the coefficient of m in f.

2.2 Algebraic Complexity Theory

Definition 6 (Algebraic Branching Program)**.** An *algebraic branching program* (ABP) is a layered directed acyclic graph. The vertex set is partitioned into layers $0, 1, \ldots, d$, with directed edges only between adjacent layers (i to $i+1$). There is a *source* vertex of in-degree 0 in the layer 0, and one out-degree 0 *sink* vertex in layer d. Each edge is labeled by an affine $\mathbb{F}$-linear form in variables, say, $x_1, x_2, \ldots, x_n$. The polynomial computed by the ABP is the sum over all source-to-sink directed paths of the ordered product of affine forms labeling the path edges.

[3] For the first two cases, invertibility over the (commutative) function field and invertibility over the (noncommutative) free skew field coincide.

The *size* of the ABP is defined as the total number of nodes and the *width* is the maximum number of nodes in a layer, and the depth or length is the number of layers in the ABP. An ABP can compute a commutative or a noncommutative polynomial, depending on whether the variables $x_1, x_2, \ldots, x_n$ occurring in the $\mathbb{F}$-linear forms are commuting or noncommuting. ABPs of width w can also be defined as an iterated matrix multiplication $\underline{u}^t \cdot M_1 M_2 \cdots M_\ell \cdot \underline{v}$, where $\underline{u}, \underline{v} \in \mathbb{F}^n$ and each M_i is of form $\sum_{i=1}^{n} A_i x_i$ for matrices $A_i \in \mathrm{Mat}_w(\mathbb{F})$, assuming without loss of generality that all matrices $M_j, 1 \leqslant j \leqslant \ell$ are $w \times w$. Here, $\underline{u}^t$ is the transpose of $\underline{u}$.

We say a set $\mathcal{H} \subseteq \mathbb{F}^n$ is a hitting set for a (commutative) algebraic circuit class $\mathcal{C}$ if for every n-variate polynomial f in $\mathcal{C}$, $f \not\equiv 0$ if and only if $f(\underline{a}) \neq 0$ for some $\underline{a} \in \mathcal{H}$.

A special class of ABPs in commuting variables are the *read-once oblivious* ABPs (in short, ROABPs). In ROABPs a different variable is used for each layer, and the edge labels are univariate polynomials over that variable. For the class of ROABPs, Forbes and Shpilka [FS13] obtained the first quasipolynomial-time black-box algorithm by constructing a hitting set of quasipolynomial size.

Theorem 7. [FS13] *For the class of polynomials computable by a width r, depth d, individual degree $< n$ ROABPs of known order, if $|\mathbb{F}| \geqslant (2dnr^3)^2$, there is a* $\mathrm{poly}(d, n, r)$*-explicit hitting set of size at most* $(2dn^2r^4)^{\lceil \log d+1 \rceil}$.

Indeed, they proved a more general result.

Definition 8 (Hitting Set Generator)**.** A polynomial map $\mathcal{G} : \mathbb{F}^t \longrightarrow \mathbb{F}^n$ is a generator for a circuit class $\mathcal{C}$ if for every n-variate polynomial f in $\mathcal{C}$, $f \equiv 0$ if and only if $f \circ \mathcal{G} \equiv 0$.

Theorem 9. [FS13, Construction 3.13, Lemma 3.21] *For the class of polynomials computable by a width r, depth d, individual degree $< n$ ROABPs of known order, one can construct a hitting set generator $\mathcal{G} : \mathbb{F}^{\lceil \log d+1 \rceil} \longrightarrow \mathbb{F}^d$ of degree dnr^4 efficiently.*

As a consequence, Forbes and Shpilka [FS13], obtain an efficient construction of quasipolynomial-size hitting set for noncommutative ABPs as well. Consider the class of noncommutative ABPs of width r, and depth d computing polynomials in $\mathbb{F}\langle \underline{x} \rangle$. The result of Forbes and Shpilka provide an explicit construction (in quasipolynomial-time) of a set $\mathrm{Mat}_{d+1}(\mathbb{F})$, such that for any ABP (with parameters r and d) computing a nonzero polynomial f, there always exists $(p_1, \ldots, p_n) \in \mathcal{H}_{n,r,d}$, $f(\underline{p}) \neq 0$.

Theorem 10 (Forbes and Shpilka [FS13])**.** *For all $n, r, d \in \mathbb{N}$, if $|\mathbb{F}| \geqslant \mathrm{poly}(d, n, r)$, then there is a hitting set $\mathcal{H}_{n,r,d} \subset \mathrm{Mat}_{d+1}(\mathbb{F})$ for noncommutative ABPs of parameters $|\mathcal{H}_{n,r,d}| \leqslant (rdn)^{O(\log d)}$ and there is a deterministic algorithm to output the set $\mathcal{H}_{n,r,d}$ in time $(rdn)^{O(\log d)}$.*

2.3 Cyclic Division Algebras

A division algebra D is an associative algebra over a (commutative) field $\mathbb{F}$ such that all nonzero elements in D are units (they have a multiplicative inverse). In this paper, we are interested in finite-dimensional division algebras. Specifically, we focus on cyclic division algebras and their construction [Lam01, Chapter 5]. Let $F = \mathbb{Q}(z)$, where z is a commuting indeterminate. Let ω be an ℓ^{th} primitive root of unity. To be specific, let $\omega = e^{2\pi\iota/\ell}$. Let $K = F(\omega) = \mathbb{Q}(\omega, z)$ be the cyclic Galois extension of F obtained by adjoining ω. So, $[K : F] = \ell$ is the degree of the extension. The elements of K are polynomials in ω (of degree at most $\ell - 1$) with coefficients from F.

Define $\sigma : K \to K$ by letting $\sigma(\omega) = \omega^k$ for some k relatively prime to ℓ and stipulating that $\sigma(a) = a$ for all $a \in F$. Then σ is an automorphism of K with F as fixed field and it generates the Galois group $\mathrm{Gal}(K/F)$.

The division algebra $D = (K/F, \sigma, z)$ is defined using a new indeterminate x as the ℓ-dimensional vector space:

$$D = K \oplus Kx \oplus \cdots \oplus Kx^{\ell-1},$$

where the (noncommutative) multiplication for D is defined by $x^\ell = z$ and $xb = \sigma(b)x$ for all $b \in K$. The parameter ℓ is called the *index* of D [Lam01, Theorem 14.9].

The elements of D has matrix representation in $K^{\ell\times\ell}$ from its action on the basis $\mathcal{X} = \{1, x, \ldots, x^{\ell-1}\}$. I.e., for $a \in D$ and $x^j \in \mathcal{X}$, the j^{th} row of the matrix representation is obtained by writing $x^j a$ in the $\mathcal{X}$-basis.

For example, the matrix representation $M(x)$ of x is:

$$M(x)[i,j] = \begin{cases} 1 & \text{if } j = i+1, i \leqslant \ell - 1 \\ z & \text{if } i = \ell, j = 1 \\ 0 & \text{otherwise.} \end{cases}$$

$$M(x) = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 1 \\ z & 0 & \cdots & 0 & 0 \end{bmatrix}.$$

For each $b \in K$ its matrix representation $M(b)$ is the diagonal matrix:

$$M(b)[i,j] = \begin{cases} b & \text{if } i = j = 1 \\ \sigma^{i-1}(b) & \text{if } i = j, i \geqslant 2 \\ 0 & \text{otherwise.} \end{cases}$$

$$M(b) = \begin{bmatrix} b & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma(b) & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma^2(b) & 0 & 0 & 0 \\ 0 & 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma^{\ell-2}(b) & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma^{\ell-1}(b) \end{bmatrix}$$

Proposition 11. *For all $b \in K$, $M(bx) = M(b) \cdot M(x)$*

Also, the matrix representation of $xb = \sigma(b)x$ is easy to see in the basis $\{1, x, \ldots, x^{\ell-1}\}$:

$$M(\sigma(b)x) = \begin{bmatrix} 0 & \sigma(b) & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma^2(b) & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma^3(b) & 0 & 0 \\ 0 & 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma^{\ell-1}(b) \\ \sigma^\ell(b)z & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Define $C_{ij} = M(\omega^{j-1})\cdot M(x^{i-1})$ for $1 \leqslant i, j \leqslant \ell$. Observe that $\mathfrak{B} = \{C_{ij}, i, j \in [\ell]\}$ is a F-generating set for the division algebra D.

Fact 12. *The F-linear span of $\mathfrak{B}$ is the cyclic division algebra D in the matrix algebra* $\mathrm{Mat}_\ell(K)$.

The following proposition is a standard fact.

Proposition 13. [Lam01, Section 14(14.13)] *The K-linear span of $\mathfrak{B}$ is the entire matrix algebra* $\mathrm{Mat}_\ell(K)$.

The following theorem gives us a way of constructing new division algebras using tensor products. This construction plays an important role in our main result.

Theorem 14. [Pie82, Proposition, Page 292] *Let K, L be cyclic extensions of the field F such that their extension degrees, $[K : F]$ and $[L : F]$, are relatively prime. Let $D_1 = (K/F, \sigma_1, z)$, and $D_2 = (L/F, \sigma_2, z)$ be the corresponding cyclic division algebras as defined above. Then their tensor product $D_1 \otimes D_2$ is also a cyclic division algebra.*

2.4 Noncommutative Rational Series

Let D be a division algebra and P be a series over the noncommuting variables $x_1, x_2, \ldots, x_n$ defined as follows:

$$P = c - B\left(\sum_{k\geqslant 0} L^k\right)A,$$

where c is a D-linear form (over $x_1, \ldots, x_n$), B (resp. A) is a $1 \times s$ (resp. $s \times 1$) dimensional vector, and L is a $s \times s$ matrix. The entries of B, L, A are D-linear forms over $x_1, \ldots, x_n$. Furthermore, the variables $x_i : 1 \leqslant i \leqslant n$ commute with the elements in D. Define the truncated polynomial $\widetilde{P}$ as follows:

$$\widetilde{P} = c - B\left(\sum_{k\leqslant s-1} L^k\right)A. \tag{2}$$

The next statement shows that the infinite series $P \neq 0$ is equivalent in saying that $\widetilde{P}$ is nonzero. The proof of the fact is standard when D is a (commutative) field [Eil74, Corollary 8.3, Page 145]. For the case of division algebras, the proof can be found in [DK21, Example 8.2, Page 23]. However, we include a self-contained proof.

Fact 15. *The infinite series $P \neq 0$ if and only if its truncation $\widetilde{P} \neq 0$.*

Proof. If $P = 0$, then obviously $\widetilde{P} = 0$, since the degrees in different homogeneous components do not match. Now, suppose $\widetilde{P} = 0$. Notice that the terms in c are linear forms and the degree of any term in $B\left(\sum_{k\geqslant 0} L^k\right)A$ is at least two. Hence, c must be zero. Write the row and column vectors B and A as $B = \sum_\ell B_\ell x_\ell$, $A = \sum_\ell A_\ell x_\ell$. Similarly, write $L = \sum_\ell L_\ell x_\ell$.

Suppose BL^sA contributes a nonzero monomial (word) $w = x_{i_1}x_{i_2}\ldots x_{i_{s+2}}$. Clearly the coefficient of w is $B_{i_1}L_{i_2}\ldots L_{i_{s+1}}A_{i_{s+2}}$. Consider the vectors $v_1 = B_{i_1}, v_2 = B_{i_1}L_{i_2}, \ldots, v_{s+1} = B_{i_1}L_{i_2}\ldots L_{i_{s+1}}$ corresponding to the prefixes $w_1 = x_{i_1}, w_2 = x_{i_1}x_{i_2}, \ldots, w_{s+1} = x_{i_1}\ldots x_{i_{s+1}}$. These vectors $v_i, 1 \leqslant i \leqslant s+1$ all lie in the (left) D-module D^s which has rank s. As D is a division algebra, these vectors cannot all be D-linearly independent. Hence, there are elements $\lambda_1, \ldots, \lambda_{s+1}$ in

D, not all zero, such that the linear combination $\lambda_1 v_1 + \ldots + \lambda_{s+1} v_{s+1} = 0$. However, $v_{s+1} A_{i_{s+2}} \neq 0$ by the assumption. Hence, there is at least one vector $v_\ell : 1 \leqslant \ell \leqslant s$ such that $v_\ell A_{i_{s+2}} \neq 0$. This means that the coefficient of the word $w_\ell x_{i_{s+2}}$, which is of length at most $s+1$, is nonzero in $\widetilde{P}$, which is not possible by assumption.

Now, with $k = s$ as the base case, we can inductively apply the above argument to show that $BL^k A$ is zero for each $k \geqslant s$. ∎

2.5 Generalized Formal Power Series

We now define the notion of generalized series first introduced by Volčič. For a detailed exposition, see [Vol18].

A *generalized word* or a *generalized monomial* in $x_1, \ldots, x_n$ over the matrix algebra $\mathrm{Mat}_m(\mathbb{F})$ allows the matrices to interleave between variables. That is to say, a generalized monomial is of the form: $a_0 x_{k_1} a_2 \cdots a_{d-1} x_{k_d} a_d$, where $a_i \in \mathrm{Mat}_m(\mathbb{F})$, and its degree is the number of variables d occurring in it. A finite sum of generalized monomials is a *generalized polynomial* in the ring $\mathrm{Mat}_m(\mathbb{F})\langle \underline{x} \rangle$. A *generalized formal power series* over $\mathrm{Mat}_m(\mathbb{F})$ is an infinite sum of generalized monomials such that the sum has finitely many generalized monomials of degree d for any $d \in \mathbb{N}$. The ring of generalized series over $\mathrm{Mat}_m(\mathbb{F})$ is denoted $\mathrm{Mat}_m(\mathbb{F})\langle\langle \underline{x} \rangle\rangle$.

A generalized series (resp. polynomial) S over $\mathrm{Mat}_m(\mathbb{F})$ admits the following canonical description. Let $E = \{e_{i,j}, 1 \leqslant i, j \leqslant m\}$ be the set of elementary matrices. Express each coefficient matrix a in S in the E basis by a $\mathbb{F}$-linear combination and then expand S. Naturally each monomial of degree-d in the expansion looks like $e_{i_0,j_0} x_{k_1} e_{i_1,j_1} x_{k_2} \cdots e_{i_{d-1},j_{d-1}} x_{k_d} e_{i_d,j_d}$ where $e_{i_l,j_l} \in E$ and $x_{k_l} \in \underline{x}$. We say the series S (resp. polynomial) is identically zero if and only if it is zero under such expansion i.e. the coefficient associated with each generalized monomial is zero.

The evaluation of a generalized series over $\mathrm{Mat}_m(\mathbb{F})$ is defined on any $k'm \times k'm$ matrix algebra for some integer $k' \geqslant 1$ [Vol18]. To match the dimension of the coefficient matrices with the matrix substitution, we use an inclusion map $\iota : \mathrm{Mat}_m(\mathbb{F}) \to \mathrm{Mat}_{k'm}(\mathbb{F})$, for example, ι can be defined as $\iota(a) = a \otimes I_{k'}$ or $\iota(a) = I_{k'} \otimes a$. Now, a generalized monomial $a_0 x_{k_1} a_1 \cdots a_{d-1} x_{k_d} a_d$ over $\mathrm{Mat}_m(\mathbb{F})$ on matrix substitution $(p_1, \ldots, p_n) \in \mathrm{Mat}_{k'm}(\mathbb{F})^n$ evaluates to

$$\iota(a_0) p_{k_1} \iota(a_1) \cdots \iota(a_{d-1}) p_{k_d} \iota(a_d)$$

under some inclusion map $\iota : \mathrm{Mat}_m(\mathbb{F}) \to \mathrm{Mat}_{k'm}(\mathbb{F})$. All such inclusion maps are known to be compatible by the Skolem-Noether theorem [Row80, Theorem 3.1.2]. Therefore, if a series S is zero with respect to some inclusion map $\iota : \mathrm{Mat}_m(\mathbb{F}) \to \mathrm{Mat}_{k'm}(\mathbb{F})$, then it is zero w.r.t. any such inclusion map. Henceforth, we define the inclusion map as $\iota(a) = a \otimes I$ w.l.o.g. to evaluate a generalized series.

We naturally extend the definition of usual ABPs (Definition 6) to the generalized ABPs.

Definition 16 (Generalized Algebraic Branching Program)**.** A *generalized algebraic branching program* is a layered directed acyclic graph. The vertex set is partitioned into layers $0, 1, \ldots, d$, with directed edges only between adjacent layers (i to $i+1$). There is a *source* vertex of in-degree 0 in the layer 0, and one out-degree 0 *sink* vertex in layer d. Each edge is labeled by a generalized linear form of $\sum_{i=1}^{n} a_i x_i b_i$ where $a_i, b_i \in \mathrm{Mat}_m(\mathbb{F})$ for some integer m. As usual, *width* is the maximum number of vertices in a layer. The generalized polynomial computed by the ABP is the sum over all source-to-sink directed paths of the ordered product of generalized linear forms labeling the path edges.

Remark 17. It is clear from the definition above that such generalized ABPs with d layers compute homogeneous generalized polynomials of degree d.

3 Division Algebra Hitting Set for Generalized ABPs over Cyclic Division Algebras

In this section, we consider generalized ABPs where the coefficients are from a cyclic division algebra. We will construct a hitting set for such ABPs inside another cyclic division algebra.

Definition 18 (D-ABP)**.** Let $D = (K/F, \sigma, z)$ be a cyclic division algebra of index ℓ. We define a D-ABP as a generalized ABP $\mathcal{A}$ in $\{x_1, x_2, \ldots, x_n\}$ variables (as defined in Definition 16) where each edge is labeled by $\sum_{i=1}^{n} a_i x_i b_i : a_i, b_i \in D$. The ABP $\mathcal{A}$ computes a generalized polynomial over D.

The main result of this section is a hitting set construction inside a cyclic division algebra for such division algebra ABPs. A key ingredient is the construction of a hitting set for noncommutative ABPs in a cyclic division algebra whose index is a power of any arbitrary prime $\wp$. We first present this construction in the next subsection before addressing the general case.

3.1 Division algebra hitting set for noncommutative ABPs

Fix a prime number $\wp$. In particular, $\wp$ is independent of the input ABP. In this section we show that the quasipolynomial-size hitting set construction for noncommutative ABPs by Forbes and Shpilka [FS13] can be adapted to a more general setting where the hitting set points consist of matrices that lie in a finite-dimensional cyclic division algebra whose index is a power of $\wp$. The construction here requires some more detail, especially as we will choose parameters keeping in mind that the rational formula hitting set problem.

The Forbes -Shpilka construction gives a quasipolynomial size hitting set for ROABPs and hence obtains a similar size hitting set for noncommutative ABPs. In order to obtain a cyclic division algebra hitting set, we will take the Forbes-Shpilka hitting set construction for ROABPs and ensure additional conditions. We now present our ROABP hitting set construction. It is essentially based on [FS13]. However, we present all the details, keeping it largely self-contained.

3.1.1 Properties of the span of an ROABP

Let F be a characteristic zero field, which is a finite extension of the function field $\mathbb{Q}(z)$ in indeterminate z. Let $\{u_i\}_{0 \leqslant i \leqslant \wp-1}$ be $\wp$ commuting indeterminates for prime $\wp$. The ring $\mathrm{Mat}_r(F[u_i])$ consists of $r \times r$ matrices whose entries are univariate polynomials in u_i over F. Equivalently, any element $M \in \mathrm{Mat}_r(F[u_i])$ can be written as a univariate polynomial $M = \sum_{j=0}^{d} M_j u_i^j$ in u_i of degree $\deg(M) = d$ with matrix coefficients $M_j \in \mathrm{Mat}_r(F)$ where $M_d \neq 0$. Let $\overline{F}$ denote the algebraic closure of F. The following lemmas are useful generalizations of the results presented in [FS13, Subsection 3.1].

Let $\xi \in \overline{F}$ be a root of unity and let $K = F(\xi)$ denote the field extension of F by ξ. For matrices $M_i \in \mathrm{Mat}_r(K), 1 \leqslant i \leqslant t$ over a field K, let $\mathrm{span}_K\{M_i\}$ denote the linear space $\{\sum_i \alpha_i M_i \mid \alpha_i \in K\}$. The next two lemmas are a straightforward generalization of [FS13, Lemma 3.2, Lemma 3.3], we skip their proofs.

Lemma 19. [FS13, Lemma 3.2] *Let K be any field. For $1 \leqslant i \leqslant t$ and $1 \leqslant \ell \leqslant t'$ let $M_{i\ell}, M'_{i\ell} \in \mathrm{Mat}_r(K)$ be matrices such that $\mathrm{span}_K\{M_{i\ell}\}_\ell = \mathrm{span}_K\{M'_{i\ell}\}_\ell$ for each i. Then,*

$$\mathrm{span}_K\left\{\prod_{i=0}^{\mathscr{p}-1} M_{i\ell}\right\}_\ell = \mathrm{span}_K\left\{\prod_{i=0}^{\mathscr{p}-1} M'_{i\ell}\right\}_\ell .$$

Lemma 20.

$$\mathrm{span}_K\left\{\left[\prod_{i=0}^{\mathscr{p}-1} u_i^{j_i}\right]\prod_{i=0}^{\mathscr{p}-1} M_i(u_i)\right\}_{\text{each } j_i\in\{0,1,\ldots,n-1\}} = \mathrm{span}_K\left\{\prod_{i=0}^{\mathscr{p}-1} M_i(u_i)\right\}_{\text{each } u_i\in F} .$$

Lemma 21. *[GR08][FS13, Lemma 3.4] Let $\xi \in \overline{F}$ be a root of unity whose (finite) order $\mathrm{ord}(\xi) > \mathfrak{n}$ and $K = F(\xi)$ be the field extension by ξ. Let $M \in \mathrm{Mat}_{\mathfrak{n},r}(F)$ For $\alpha \in K$, define $A_\alpha \in \mathrm{Mat}_{r,\mathfrak{n}}(K)$ by $(A_\alpha)_{i,j} = (\xi^i\alpha)^j$. Then there are $< \mathfrak{n}r$ values of $\alpha \in K$ such that the first r rows of $\mathrm{rank}(A_\alpha M) < r$.*

Lemma 22. *For each $i \in [\mathscr{p}]$, let $M_i(u_i) \in \mathrm{Mat}_r(F[u_i])$ be matrix polynomials of degree less than $\mathfrak{n}$. Let $\xi \in \overline{F}$ be a root of unity such that $\mathrm{ord}(\xi) > \mathfrak{n}^{\mathscr{p}}$ and let $K = F(\xi)$. Then for any $\alpha \in K$ and any $\mu \geqslant \mathfrak{n}$,*

$$\mathrm{span}_K\left\{\left[\prod_{i=0}^{\mathscr{p}-1} u_i^{j_i}\right]\prod_{i=0}^{\mathscr{p}-1} M_i(u_i)\right\}_{\text{each } j\in\{0,1,\ldots,n-1\}} \supseteq \mathrm{span}_K\left\{\prod_{i=0}^{\mathscr{p}-1} M_i((\xi^\ell\alpha)^{\mu^i})\right\}_{0\leqslant\ell\leqslant\mathrm{ord}(\xi)-1} .$$

Moreover, for all but $\mathfrak{n}^p r^2$ many values of α in K,

$$\mathrm{span}_K\left\{\left[\prod_{i=0}^{\mathscr{p}-1} u_i^{j_i}\right]\prod_{i=0}^{\mathscr{p}-1} M_i(u_i)\right\}_{\text{each } j\in\{0,1,\ldots,n-1\}} = \mathrm{span}_K\left\{\prod_{i=0}^{\mathscr{p}-1} M_i((\xi^\ell\alpha)^{\mu^i})\right\}_{0\leqslant\ell\leqslant r^2-1} .$$

Proof. By definition we have

$$\prod_{i=1}^{\mathscr{p}} M_i(u_i) = \sum_{j_i}\left(\left[\prod_{i=0}^{\mathscr{p}-1} u_i^{j_i}\right]\prod_{i=1}^{\mathscr{p}} M_i(u_i)\right)\cdot\prod_{i=0}^{\mathscr{p}-1} u_i^{j_i}.$$

Therefore, by substitution we have

$$\prod_{i=1}^{\mathscr{p}} M_i((\xi^\ell\alpha)^{\mu^{i-1}}) = \sum_{j_i}\left(\left[\prod_{i=0}^{\mathscr{p}-1} u_i^{j_i}\right]\prod_{i=1}^{\mathscr{p}} M_i(u_i)\right)\cdot(\xi^\ell\alpha)^{j_0+j_1\mu+\ldots+j_{\mathscr{p}-1}\mu^{\mathscr{p}-1}}, \tag{3}$$

which implies the first part of the lemma.

Now, we define a rectangular matrix $C \in \mathrm{Mat}_{\mathfrak{n}^{\mathscr{p}},r^2}(F)$ as follows. Each row of C is indexed by a tuple $(j_0, j_1, \ldots, j_{\mathscr{p}-1}) \in \{0, 1, \ldots, \mathfrak{n}-1\}^{\mathscr{p}}$. For each such tuple $(j_0, j_1, \ldots, j_{\mathscr{p}-1})$, treating the $r \times r$

matrix $\left[\prod_{i=0}^{\mathfrak{p}-1} u_i^{j_i}\right] \prod_{i=0}^{\mathfrak{p}-1} M_i(u_i)$ as an r^2-dimensional vector, we define it as the corresponding row $C_{(j_0,j_1,\ldots,j_{\mathfrak{p}-1})}$. By definition,

$$\text{row-span}(C) = \text{span}\left\{\left[\prod_{i=0}^{\mathfrak{p}-1} u_i^{j_i}\right] \prod_{i=0}^{\mathfrak{p}-1} M_i(u_i)\right\}_{\text{each } j \in \{0,1,\ldots,\mathfrak{n}-1\}}.$$

Next, consider the rectangular matrix $A_\alpha \in \text{Mat}_{r^2,\mathfrak{n}^{\mathfrak{p}}}(K)$ whose columns are indexed by tuples $(j_0, j_1, \ldots, j_{\mathfrak{p}-1}) \in \{0, 1, \ldots, \mathfrak{n}-1\}^{\mathfrak{p}}$ with entries defined as

$$(A_\alpha)_{\ell,(j_0,j_1,\ldots,j_{\mathfrak{p}-1})} = (\xi^\ell \alpha)^{j_0+j_1\mu+\ldots+j_{\mathfrak{p}-1}\mu^{\mathfrak{p}-1}}.$$

By Lemma 21, for all but $\mathfrak{n}^{\mathfrak{p}} r^2$ values of α, we have $\text{rank}(A_\alpha C) = \text{rank}(C)$. Multiplying the ℓ^{th} row of A_α with C we get

$$(A_\alpha)_\ell C = \sum_{j_i} \left(\left[\prod_{i=0}^{\mathfrak{p}-1} u_i^{j_i}\right] \prod_{i=0}^{\mathfrak{p}-1} M_i(u_i)\right) \cdot (\xi^\ell \alpha)^{j_0+j_1\mu+\ldots+j_{\mathfrak{p}-1}\mu^{\mathfrak{p}-1}} = \prod_{i=0}^{\mathfrak{p}-1} M_i((\xi^\ell \alpha)^{\mu^i}).$$

Therefore,

$$\text{row-span}(A_\alpha C) = \text{span}\left\{M_0(\xi^\ell \alpha) M_1((\xi^\ell \alpha)^\mu) \cdots M_{\mathfrak{p}-1}((\xi^\ell \alpha)^{\mu^{\mathfrak{p}-1}})\right\}_{0 \leqslant \ell \leqslant r^2-1\}}.$$

As row-span(C) contains row-span$(A_\alpha C)$, if rank(C) = rank$(A_\alpha C)$ then we have row-span(C) = row-span$(A_\alpha C)$. Therefore, for all but $\mathfrak{n}^{\mathfrak{p}} r^2$ values of α, row-span(C) = row-span$(A_\alpha C)$. ∎

We inductively assume that the ROABP is transformed into an iterated matrix product of $\mathfrak{p}^d$ matrices with the following property. We can group this iterated matrix product into $\mathfrak{p}$ consecutive sections, where the i^{th} section is a product of $\mathfrak{p}^{d-1}$ consecutive matrices the entries of which are polynomial in the variable $u_i, 0 \leqslant i \leqslant \mathfrak{p}-1$. Furthermore, this transformation of the original ROABP is identity preserving, that is, the polynomial computed by the original ROABP is nonzero if and only if a designated entry of this matrix product is nonzero.

Then, broadly speaking, the next lemma gives a method to show that the span of the full matrix product can be captured by span of the matrix products over a *single* variable. This will enable us to transform the given iterated matrix product over variables $\{u_i\}$ into an iterated matrix product over a single variable. This lemma, based on Lagrangian interpolation, is a variant of [FS13, Lemma 3.7].

We recall the definition of Lagrange interpolation polynomials. Let χ be a positive integer and $S = \{\beta_0, \beta_1, \ldots, \beta_{\chi-1}\}$ be a set of χ distinct scalars. The Lagrange interpolation polynomials with respect to S is the set of χ many univariate polynomials $q_{\ell,S}(v), 0 \leqslant \ell \leqslant \chi-1$ of degree $\chi-1$ in variable v defined as:

$$q_{\ell,S}(v) = \prod_{k \neq \ell} \frac{v-\beta_k}{\beta_\ell - \beta_k} \quad \text{such that} \quad q_{\ell,S}(\beta_k) = \begin{cases} 1 & \text{if } \ell = k, \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 23. *Consider p many families of $r \times r$ matrices $\{M_{ij}\}_{0 \leqslant j \leqslant p, 0 \leqslant i \leqslant p^{d-1}}$ where for the j^{th} family the entries are univariate polynomials over $F[u_j]$ of degree less than n. Let $(f_0(u_j), f_1(u_j), \ldots, f_{p^{d-1}-1}(u_j)) \in F[u_j]$ be polynomials of degree at most m for each $0 \leqslant j \leqslant p-1$. Let ω be a primitive root of unity of order p^L for prime $p > 2$, where $K = F(\omega)$. Let $\omega_0 \in F$ be a root of unity of order $\Lambda = q^\tau > (p^d nm)^p$ where $q \neq p$ is a prime number. Let $1 \leqslant \gamma \leqslant p^L$ such that γ is relatively prime to p^L and $\gamma \pmod{\Lambda} = 1$. Define polynomials in indeterminate v:*

$$f'_{ij}(v) \quad = \quad \sum_{\ell=0}^{r^2-1} \sum_{\ell'=0}^{\mathrm{ord}(\omega)-1} f_i((\omega_0{}^{\ell} \omega^{\ell'} \alpha)^{\mu_{i,j}}) q_{\ell\ell',S}(v)$$

where $\mu_{i,j} = \mu^{j-1}\gamma^{i-1}$, and $q_{\ell\ell'}(v)$ denotes the Lagrange interpolation polynomials with respect to a set $S = \{\beta_{\ell\ell'}\}_{0 \leqslant \ell \leqslant r^2-1, 0 \leqslant \ell' \leqslant \mathrm{ord}(w)-1}$ where each $\beta_{\ell\ell'} \in F$.

Then, for all but $(p^{d-1}nm)^p r^2$ many values of α,

$$\mathrm{span}_K \left\{ \prod_{j=0}^{p-1} \prod_{i=0}^{p^{d-1}-1} M_{ij}(f_i(u_j)) \right\}_{\text{each } u_j \in F} \subseteq \mathrm{span}_K \left\{ \prod_{j=0}^{p-1} \prod_{i=0}^{p^{d-1}-1} M_{ij}(f'_i(v)) \right\}_{v \in F}.$$

Proof. As before, all spans are K-linear spans. For each j, let $R_j(u_j) = \prod_{i=0}^{p^{d-1}-1} M_{ij}(f_i(u_j))$. Note that $R_j(u_j)$ is a matrix of univariate polynomials in u_j of degree less than $p^{d-1}nm$. By definition,

$$\prod_{j=1}^{p} \prod_{i=0}^{p^{d-1}-1} M_{ij}(f_i(u_j)) = \prod_{j=1}^{p} R_j(u_j).$$

Lemma 22 and Lemma 20 imply that for $\mu > p^d nm$, except for $< (p^d nm)^p r^2$ many values of α in K,

$$\mathrm{span}_K \left\{ \prod_{j=0}^{p-1} R_j(u_j) \right\}_{\text{each } u_j \in K} = \mathrm{span}_K \left\{ \prod_{j=0}^{p-1} R_j(((\omega_0\omega)^{\ell}\alpha)^{\mu^j}) \right\}_{0 \leqslant \ell \leqslant r^2-1}. \tag{4}$$

Claim 24.

$$\mathrm{span} \left\{ \prod_{j=0}^{p-1} \prod_{i=0}^{p^{d-1}-1} M_{ij}(f_i(((\omega_0\omega)^{\ell}\alpha)^{\mu^j})) \right\}_{0 \leqslant \ell \leqslant r^2-1} \subseteq \mathrm{span} \left\{ \prod_{j=0}^{p-1} \prod_{i=0}^{p^{d-1}-1} M_{ij}(f_i((\omega_0^{\ell}\omega^{\ell'}\alpha)^{\mu^j\gamma^i})) \right\}_{\substack{0 \leqslant \ell \leqslant r^2-1, \\ 0 \leqslant \ell' \leqslant \mathrm{ord}(\omega)-1}}.$$

Proof. It suffices to show $\mathrm{span}\left\{M_{ij}(f_i((\omega^{\ell}\alpha)^{\mu^j}))\right\} = \mathrm{span}\left\{M_{ij}(f_i((\omega_0^{\ell}\omega^{\ell'}\alpha)^{\mu^j\gamma^i}))\right\}$ (we can then apply Lemma 19). Let $\alpha = \omega_0^{j_1}\omega^{j_2}$. Therefore, $(\omega_0\omega)^{\ell}\alpha = \omega_0^{\ell+j_1}\omega^{\ell+j_2}$. Now, for each i the map $a \mapsto a^{\gamma^i}$ for all $a \in \{\omega^{\ell} \mid 0 \leqslant \ell \leqslant \mathrm{ord}(\omega)-1\}$ is a bijection as γ is relatively prime to $\mathrm{ord}(\omega)$. Hence,

$$((\omega_0\omega)^{\ell}\alpha)^{\gamma^{-i}} = (\omega_0^{\ell+j_1}\omega^{\ell+j_2})^{\gamma^{-i}} = \omega_0^{\ell+j_1}(\omega^{\ell+j_2})^{\gamma^{-i}} = \omega_0^{\ell+j_1}\omega^{\ell'+j_2} = \omega_0^{\ell}\omega^{\ell'}\alpha,$$

for a unique $0 \leqslant \ell' \leqslant \mathrm{ord}(\omega)-1$, as γ is relatively prime to the order of ω. ∎

We now set $\mu_{i,j} = \mu^j \gamma^i$ in Claim 24 to get,

$$\operatorname{span}\left\{\prod_{j=0}^{p-1}\prod_{i=0}^{p^{d-1}-1} M_{ij}(f_i((\omega^\ell \alpha)^{\mu^j}))\right\}_{0\leqslant \ell\leqslant r^2-1} \subseteq \operatorname{span}\left\{\prod_{j=0}^{p-1}\prod_{i=0}^{p^{d-1}-1} M_{j,i}(f_i((\omega_0^\ell \omega^{\ell'} \alpha)^{\mu_{i,j}}))\right\}_{\substack{0\leqslant \ell\leqslant r^2-1,\\ 0\leqslant \ell'\leqslant \operatorname{ord}(\omega)-1}}. \tag{5}$$

For each j, let $T_j(v) = \prod_{i=0}^{p^{d-1}-1} M_{ij}(f'_{i+jp^{d-1}}(v))$. By the definition of the Lagrange interpolation polynomials, letting $q_\ell(\beta_k) = \delta_{\ell k}$ where each β_k is distinct, we have

$$T_j(\beta_{\ell\ell'}) = \prod_{i=0}^{p^{d-1}-1} M_{ij}(f_i((\omega_0^\ell \omega^{\ell'} \alpha)^{\mu_{i,j}})), \tag{6}$$

where $(\omega_0\omega)^\ell \alpha = (\omega_0^\ell \omega^{\ell'} \alpha)^{\gamma^i}$. Combining Equation (4), Equation (5), and Equation (6), we obtain:

$$\operatorname{span}\left\{\prod_{j=0}^{p-1} R_j(u_j)\right\}_{\text{each } u_j\in F} \subseteq \operatorname{span}\left\{\prod_{j=0}^{p-1} T_j(v)\right\}_{v\in F}. \qquad \blacksquare$$

Remark 25. Lemma 23 can be seen as a generalization of [FS13, Lemma 3.7]. However, here we use two different roots of unity ω_0 and ω unlike [FS13]. The order of ω_0 is sufficiently large which allows us to use Lemma 22 in order to prove the span containment, similar to [FS13]. Whereas ω plays an important role to embed the hitting set points for the noncommutaitve ABPs inside a cyclic division algebra as shown in Section 3.1.3.

3.1.2 Construction of the ROABP hitting set generator

We will now proceed to describe the construction of a hitting set generator for ROABPs. This in turn will yield a quasipolynomial size hitting set for noncommutative ABPs, where the matrix substitutions will be from a cyclic division algebra. Our construction and choice of parameters here is tailored towards the main aim of the paper, which is obtaining a quasipolynomial size hitting set for noncommutative rational formulas. That we shall describe in subsequent sections.

Fix a prime $p > 2$ and integers n, r, d. Let $\omega = e^{\frac{2\pi\iota}{p^L}}$, a root of unity of order p^L for some even number $L > d$. Let $F = \mathbb{Q}(z, \omega_0)$ where ω_0 is a root of unity of order $\Lambda = q^\tau > (p^d mn)^p \cdot r^2$ where $q \neq p$ is a prime number and $K = F(\omega)$ is its (finite) extension by ω. Consider the K-automorphism $\sigma : a \mapsto a^{1+\Lambda p^\kappa}$ where the positive integer $\kappa = L/2$. It is easy to check that $\sigma(\omega_0) = \omega_0$. As $\sigma(\omega) = \omega^{1+\Lambda p^\kappa}$ and $1 + \Lambda p^\kappa$ is relatively prime to p^L it follows that σ generates all the automorphisms of K that fix F.

Let $\omega_i = \omega^{p^{\kappa-i}}$ for each $1 \leqslant i \leqslant d$. We denote by K_i the cyclic Galois extension $K_i = F(\omega_i)$ of F by ω_i, for $1 \leqslant i \leqslant d$. This gives a tower of field extensions

$$F \subset F(\omega_1) \subset F(\omega_2) \subset \cdots \subset F(\omega_d) \subset F(\omega) = K.$$

Observation 26. *For each $1 \leqslant i \leqslant d$, σ^{p^i} fixes $F(\omega_i)$.*

Proof. Define $a_i = \kappa - i$. As $\sigma(\omega) = \omega^{\Lambda p^{\kappa}+1}$, we have $\sigma(\omega_i) = \omega^{p^{a_i}(\Lambda p^{\kappa}+1)}$. Therefore,

$$\sigma^{p^i}(\omega_i) = \omega^{p^{a_i}(\Lambda p^{\kappa}+1)^{p^i}}.$$

Now, $(\Lambda p^{\kappa}+1)^{p^i} = \sum_{j=0}^{p^i}\binom{p^i}{j}\Lambda^j p^{\kappa j}$. As $\kappa = L/2$, we have $\omega^{p^{\kappa j}} = 1$ for $j \geqslant 2$. Therefore,

$$\sigma^{p^i}(\omega_i) = \omega^{p^{a_i}(\Lambda p^{i+\kappa}+1)} = \omega_i \cdot \omega^{\Lambda p^{a_i+i+\kappa}}.$$

The choice of a_i, $a_i + i + \kappa = L$ for $1 \leqslant i \leqslant d$, implies that $\omega^{\Lambda p^{a_i+i+\kappa}} = 1$ ensuring that σ^{p^i} fixes ω_i. ∎

For each $1 \leqslant i \leqslant d$, define the set

$$W_i = \{\omega_0^{j_1}\omega_i^{j_2} \mid 1 \leqslant j_1 \leqslant \Lambda,\ 1 \leqslant j_2 \leqslant p^{\kappa+i}\}.$$

Hitting Set Generator. We now proceed to define the hitting set generator as a mapping $\mathcal{G}_d : K^{d+1} \to K^{p^d}$. The map $\mathcal{G}_d$ will be recursively defined in terms of $\mathcal{G}_{d-1}$ which is for matrix products of length p^{d-1}. More generally, we will define $\mathcal{G}_i : K^{i+1} \to K^{p^i}$ in terms of $\mathcal{G}_{i-1}$, $1 \leqslant i \leqslant d$.

For $0 \leqslant k \leqslant p^i - 1$, we will denote the k^{th} coordinate of the output of $\mathcal{G}_i$ as the function $\mathcal{G}_{i,k} : K^{i+1} \to K$. Lemma 23 will play a role in the construction. Specifically, the choice of the interpolating scalar set $S_i \subseteq F$ used in the lemma and the effect of the automorphism σ on the Lagrange interpolation polynomials will be important. In the definition of $\mathcal{G}_i$, we will choose a subset $S_i \subset \mathcal{S} = \{\omega_0^{\ell} \mid 0 \leqslant \ell \leqslant \Lambda - 1\}$ of size $|S_i| = r^2 \cdot \mathrm{ord}(\omega_i) = r^2 \cdot p^{\kappa+i}$. By choice $\Lambda = q^{\tau} > r^2 \cdot \mathrm{ord}(\omega_i)$ for $1 \leqslant i \leqslant d$. Notice that $S_i \subset W_i$ and each element of S_i is fixed by σ.

For $0 \leqslant k \leqslant p^i - 1$, we write $k = jp^{i-1} + k'$ for some $0 \leqslant j \leqslant p - 1$ and $0 \leqslant k' \leqslant p^{i-1}$. Define $\mathcal{G}_{0,k}(\alpha_1) = \alpha_1$ and for $i > 0$ define the mapping $\mathcal{G}_{i,k}$ recursively as:

$$\mathcal{G}_{i,k}(\alpha_1,\ldots,\alpha_i,\alpha_{i+1}) = \sum_{\ell'=0}^{\mathrm{ord}(\omega_i)-1}\sum_{\ell=0}^{r^2-1}\mathcal{G}_{i-1,k'}(\alpha_1,\ldots,\alpha_{i-1},(\omega_0^{\ell}\omega_i^{\ell'}\alpha_i)^{\mu_i^j\gamma^{k'}})q_{\ell\ell',S_i}(\alpha_{i+1}), \tag{7}$$

where $\mu_i = 1 + \Lambda p^{i-1+\kappa}$ and $\gamma = 1 + \Lambda p^{\kappa}$.

Remark 27. It is useful to note that for all $\alpha_i \in W_i$, from the choice of μ and γ, we have $\alpha_i^{\gamma^k} = \alpha_i^{\mu_i^j\gamma^{k'}}$. Hence in Lemma 23, we can equivalently write $(\omega_0^{\ell}\omega_i^{\ell'}\alpha_i)^{\gamma^k}$ instead of $(\omega_0^{\ell}\omega_i^{\ell'}\alpha_i)^{\mu_i^j\gamma^{k'}}$.

Remark 28. We briefly explain the parameters μ_i^j and γ^k. The parameter μ_i^j is from setting μ in Lemma 23 as $\mu_i = 1 + \Lambda p^{i-1+\kappa}$. Then $\mu_i^j = (1 + \Lambda p^{i-1+\kappa})^j$. Notice that if $\alpha_i \in W_i$ then $\mathrm{ord}(\alpha_i) \leqslant |W_i| = \Lambda \cdot \mathrm{ord}(\omega_i)$ and, as observed in Remark 27, $\alpha_i^{\mu_i^j\gamma^{k'}} = \alpha_i^{\gamma^k}$. Letting $e_{ik} = \gamma^k (\mathrm{mod}\ \Lambda\,\mathrm{ord}(\omega_i))$ we can write $\alpha_i^{\gamma^k} = \alpha_i^{e_{ik}}$ for each i, where $e_{ik} \leqslant \Lambda\,\mathrm{ord}(\omega_i) - 1$. This implies that the degree of α_i for $i \leqslant d$ in the polynomial $\mathcal{G}_d$ can be bounded by $\Lambda\,\mathrm{ord}(\omega_i) - 1$.

Lemma 29. *For a positive integer d let $M_k \in F[u_k]^{r\times r}$ for $0 \leqslant k \leqslant \wp^d - 1$, be matrices whose entries are polynomials of degree at most n. Then the following containment holds for the mapping $\mathcal{G}_d : K^{d+1} \to K^{\wp^d}$ defined in Equation (7):*

$$\operatorname{span}_K \left\{ \prod_{k=0}^{\wp^d-1} M_k(u_k) \right\}_{\text{each } u_k\in F} \subseteq \operatorname{span}_K \left\{ \prod_{k=0}^{\wp^d-1} M_k(\mathcal{G}_{d,k}(\alpha_1,\ldots,\alpha_d,\alpha_{d+1})) \right\}_{\substack{\alpha_{d+1}\in\mathcal{S},\\ \text{each } \alpha_i\in W_i, 1\leqslant i\leqslant d}}.$$

Proof. We will prove this by a suitable induction using Lemma 23. In order to set up the induction we observe that the $\wp^d$ product of the matrices $M_k(u_k)$ (in $\wp^d$ distinct commuting variables u_k) can be grouped into $\wp$-products of matrices at a time forming a full $\wp$-ary tree of depth $d+1$. Counting from the bottom layer numbered as layer 0, at the i^{th} layer of this $\wp$-ary tree of matrix products we have $\wp^{d-i}$ nodes, where each node represents a $\wp^i$ matrix product. For $0 \leqslant k \leqslant \wp^d - 1$, we can uniquely express $k = j\wp^i + k'$ for $0 \leqslant k' \leqslant p^i - 1$ and $0 \leqslant j \leqslant \wp^{d-i} - 1$. Denoting, for convenience, the matrix M_k as $M_{jk'}$ we can write

$$\prod_{k=0}^{\wp^d-1} M_k(u_k) = \prod_{j=0}^{\wp^{d-i}-1} \prod_{k'=0}^{\wp^i-1} M_{jk'}(u_k).$$

Now, we can state the claim which we prove by induction and obtain the lemma. We want to show for $0 \leqslant i \leqslant d$ that

$$\operatorname{span}_K \left\{ \prod_{k=0}^{\wp^d-1} M_k(u_k) \right\}_{\text{each } u_k\in F} \subseteq \operatorname{span}_K \left\{ \prod_{j=0}^{\wp^{d-i}-1} \prod_{k'=0}^{\wp^i-1} M_{jk'}(\mathcal{G}_{i,k}(\alpha_1,\ldots,\alpha_i,\hat{u}_j)) \right\}_{\substack{\alpha_\ell\in W_\ell,\ 1\leqslant\ell\leqslant i,\\ \hat{u}_j\in\mathcal{S}}}. \tag{8}$$

Base Case: $i = 0$**.** For the base case $i = 0$ we note that the right hand side coincides with the left because for $i = 0$ we have $j = k$ and $\mathcal{G}_{0,k}(\hat{u}_j) = \hat{u}_k$ for $\hat{u}_k \in F$. Then as the entries of M_k are polynomials of degree at most n in u_k and $|\mathcal{S}| = \Lambda$ is sufficiently larger than n by choice, we can apply Lemma 22 (choosing $\xi = \omega_0$ and $\alpha \in \mathcal{S}$ in that lemma)

$$\operatorname{span}_K \left\{ \prod_{k=0}^{\wp^d-1} M_k(u_k) \right\}_{\text{each } u_k\in F} \subseteq \operatorname{span}_K \left\{ \prod_{k=0}^{\wp^d-1} M_k(u_k) \right\}_{\text{each } u_k\in \mathcal{S}}.$$

Induction Step Now suppose the containment of Equation 8 holds for some $i < d$. Our aim is to prove it for $i + 1$. From the induction hypothesis, it suffices to show the containment of the right hand side of Equation 8 in the span

$$\operatorname{span}_K \left\{ \prod_{j'=0}^{\wp^{d-i-1}-1} \prod_{k''=0}^{\wp^{i+1}-1} M_{j'k''}(\mathcal{G}_{i+1,k}(\alpha_1,\ldots,\alpha_{i+1},\hat{u}_j)) \right\}_{\substack{\alpha_\ell\in W_\ell,\ 1\leqslant\ell\leqslant i,\\ \hat{u}_j\in F}}.$$

Here, as we did for layer i, we are writing $k = j'\wp^{i+1} + k''$ where $0 \leqslant k'' \leqslant p^{i+1} - 1$ and $0 \leqslant j' \leqslant \wp^{d-i-1} - 1$. Also, we are denoting the matrix M_k as $M_{j'k''}$. We are going to prove this containment by applying Lemma 23. It is helpful to write the RHS of Equation 8 as the three-fold product

$$\operatorname{span}_K \left\{ \prod_{j'=0}^{\mathcal{p}^{d-i-1}-1} \prod_{j_1=0}^{\mathcal{p}-1} \prod_{k'=0}^{\mathcal{p}^{i}-1} M_{jk'}(\mathcal{G}_{i,k}(\alpha_1, \ldots, \alpha_i, \hat{u}_j))) \right\}_{\substack{\alpha_\ell \in W_\ell,\ 1 \leqslant \ell \leqslant i, \\ \hat{u}_j \in F}},$$

where the index j used in RHS of Equation 8 is essentially split into j' and j_1 where $j = j'\mathcal{p}^{d-i-1} + j_1$. Likewise, note that we can also write the index k'' as $k'' = j_1\mathcal{p}^i + k'$ for $0 \leqslant k'' \leqslant p^{i+1} - 1$. Then, by the short-hand we are using, the matrix M_k is the same as $M_{jk'}$ and $M_{j'k''}$.

By Lemma 19 it suffices to show for each $0 \leqslant j' \leqslant \mathcal{p}^{d-i-1}$ and each fixed value of $(\alpha_1, \alpha_2, \ldots, \alpha_i) \in W_1 \times W_2 \times \cdots \times W_i$: $\operatorname{span}_K \left\{ \prod_{j_1=0}^{\mathcal{p}-1} \prod_{k'=0}^{\mathcal{p}^i-1} M_{j'k''}(\mathcal{G}_{i,k}(\alpha_1, \ldots, \alpha_i, \hat{u}_{j_1})) \right\}_{\hat{u}_{j_1} \in F}$ is contained in $\operatorname{span}_K \left\{ \prod_{k''=0}^{\mathcal{p}^{i+1}-1} M_{j'k''}(\mathcal{G}_{i+1,k}(\alpha_1, \ldots, \alpha_{i+1}, v)) \right\}_{\alpha_{i+1} \in W_{i+1}, v \in \mathcal{S}}$.

Each matrix entry in the first span above is a univariate polynomial over $\hat{u}_{j_1}$ only. Therefore, by Lemma 19 and Lemma 23 we obtain,

$$\begin{aligned}
&\operatorname{span}_K \left\{ \prod_{j_1=0}^{\mathcal{p}-1} \prod_{k'=0}^{\mathcal{p}^i-1} M_{j'k''}(\mathcal{G}_{i,k}(\alpha_1, \ldots, \alpha_i, \hat{u}_{j_1})) \right\}_{\hat{u}_{j_1} \in F} \\
&\subseteq \operatorname{span}_K \left\{ \prod_{j_1=0}^{\mathcal{p}-1} \prod_{k'=0}^{\mathcal{p}^i-1} M_{j'k''} \left(\left(\sum_{\ell'=0}^{\operatorname{ord}(\omega_{i+1})-1} \sum_{\ell=0}^{r^2-1} \mathcal{G}_{i,k}(\alpha_1, \ldots, \alpha_i, ((\omega_0^{\ell} \omega_{i+1}^{\ell'} \alpha_{i+1})^{\mu_i^{j_1} \gamma^{k'}}) \right) q_{\ell\ell', S_i}(v) \right) \right\}_{\alpha_{i+1} \in W_{i+1}, v \in \mathcal{S}} \\
&\subseteq \operatorname{span}_K \left\{ \prod_{k''=0}^{\mathcal{p}^{i+1}-1} M_{j'k''}(\mathcal{G}_{i+1,k}(\alpha_1, \ldots, \alpha_{i+1}, v)) \right\}_{\alpha_{i+1} \in W_{i+1}, v \in \mathcal{S}}.
\end{aligned}$$

Importantly, by Lemma 23 the number of bad choices for $\alpha_{i+1} \in W_{i+1}$ is only a small fraction of $|W_{i+1}|$ by our choice of Λ.

■

Lemma 30. *The mapping $\mathcal{G}_d : K^{d+1} \to K^{\mathcal{p}^d}$ as defined in Equation (7) has the property that for $(\alpha_1, \alpha_2, \ldots, \alpha_{d+1}) \in W_1 \times W_2 \times \cdots \times W_d \times \mathcal{S}$ and for any $1 \leqslant k < \mathcal{p}^d$*

$$\sigma(\mathcal{G}_{d,k}(\alpha_1, \ldots, \alpha_d, \alpha_{d+1})) = \mathcal{G}_{d,k+1}(\alpha_1, \ldots, \alpha_d, \alpha_{d+1}).$$

We need to show for $\alpha_i \in W_i$, $1 \leqslant i \leqslant d$ and $\alpha_{d+1} \in \mathcal{S}$, for any $1 \leqslant k < \mathcal{p}^d$

$$\sigma(\mathcal{G}_{d,k}(\alpha_1, \ldots, \alpha_d, \alpha_{d+1})) = \mathcal{G}_{d,k+1}(\alpha_1, \ldots, \alpha_d, \alpha_{d+1}).$$

The proof is by induction with some tedious details. Recall that $\sigma(\omega) = \omega^\gamma$. For readability we shall use $\underline{\alpha}$ for the $(d-2)$-tuple $(\alpha_1, \alpha_2, \ldots, \alpha_{d-2})$. Inductively, suppose for any $\alpha_i \in W_i$ for $1 \leqslant i \leqslant d-1$ and $\alpha_d \in \mathcal{S}$, and for $1 \leqslant k' < \mathcal{p}^{d-1}$ we have $\sigma(\mathcal{G}_{d-1,k}(\underline{\alpha}, \alpha_{d-1}, \alpha_d)) = \mathcal{G}_{d-1,k+1}(\underline{\alpha}, \alpha_{d-1}, \alpha_d)$. Then by definition of $\mathcal{G}_{d-1}$ (Equation 7), writing $k = j\mathcal{p}^{d-1} + k'$ we have

$$\sigma\left(\sum_{\ell,\ell'} \mathcal{G}_{d-2,k}(\underline{\alpha}, (\omega_0^{\ell} \omega_{d-1}^{\ell'} \alpha_{d-1})^{\mu_{d-1}^{j} \gamma^{k'}}) q_{\ell\ell', S_{d-2}}(\alpha_d) \right) = \sum_{\ell,\ell'} \mathcal{G}_{d-2,k+1}(\underline{\alpha}, (\omega_0^{\ell} \omega_{d-1}^{\ell'} \alpha_{d-1})^{\mu_{d-1}^{j} \gamma^{k'+1}}) q_{\ell\ell', S_{d-2}}(\alpha_d).$$

For $\alpha_d \in \mathcal{S}$ we have $\sigma(\alpha_d) = \alpha_d$. Hence, we have for each ℓ and ℓ',

$$\sigma(\mathcal{G}_{d-2,k}(\alpha_1, \dots, (\omega_{d-1}^{\ell}\alpha_{d-1})^{\gamma^k}) = \mathcal{G}_{d-2,k+1}(\alpha_1, \dots, (\omega_{d-1}^{\ell}\alpha_{d-1})^{\gamma^{k+1}}).$$

At this point, it is useful to introduce some notation for ease of reading. We will assume $(\alpha_1, \alpha_2, \dots, \alpha_{d-2})$ is a fixed $(d-2)$-tuple. Let $f'_k(\alpha_{d-1}, \alpha_d, \alpha_{d+1})$, $g_k(\alpha_{d-1}, \alpha_d)$ and $h_k(\alpha_{d-1})$ denote the maps $\mathcal{G}_{d,k}(\underline{\alpha}, \alpha_{d-1}, \alpha_d, \alpha_{d+1})$, $\mathcal{G}_{d-1,k}((\underline{\alpha}, \alpha_{d-1}, \alpha_d)$, and $\mathcal{G}_{d-2,k}((\underline{\alpha}, \alpha_{d-1})$ respectively. We now define $G_{d,k}$ in terms of $\mathcal{G}_{d-2,k}$ using the above notation using Remark 27 that $a^{\gamma^k} = a^{\mu_i^j \gamma^{k'}}$ for $a \in W_i$ where $k = j\wp^{i-1} + k'$ and $k' = j'\wp^{d-2} + k''$.

$$\begin{aligned}
\mathcal{G}_{d,k}(\alpha_1, \dots, \alpha_d, \alpha_{d+1}) &= f'_k(\alpha_{d-1}, \alpha_d, \alpha_{d+1}) \\
&= \sum_{\ell,\ell'} g_{k'}(\alpha_{d-1}, (\omega_0^{\ell}\omega_d^{\ell'}\alpha_d)^{\gamma^k}) q_{\ell\ell',S_d}(\alpha_{d+1}) \\
&= \sum_{\ell,\ell'} \left(\sum_{\ell_1,\ell_2} h_{k''}((\omega_0^{\ell_1}\omega_{d-1}^{\ell_2}\alpha_{d-1})^{\gamma^{k'}}) q_{\ell_1\ell_2,S_{d-1}}((\omega_0^{\ell_1}\omega_d^{\ell_2}\alpha_d)^{\gamma^k}) \right) q_{\ell\ell',S_d}(\alpha_{d+1}).
\end{aligned}$$

Therefore, $\sigma(\mathcal{G}_{d,k}(\alpha_1, \dots, \alpha_d, \alpha_{d+1})) = \sigma(f'_k(\alpha_{d-1}, \alpha_d, \alpha_{d+1}))$

$$\begin{aligned}
&= \sum_{\ell,\ell'} \left(\sum_{\ell_1,\ell_2} \sigma(h_{k''}((\omega_0^{\ell_1}\omega_{d-1}^{\ell_2}\alpha_{d-1})^{\gamma^{k'}}) q_{\ell_1\ell_2,S_{d-1}}(\sigma(\omega_0^{\ell}\omega_d^{\ell'}\alpha_d^{\gamma^k})) \right) q_{\ell\ell',S_d}(\sigma(\alpha_{d+1})) \\
&= \sum_{\ell,\ell'} \left(\sum_{\ell_1,\ell_2} h_{k_1}((\omega_0^{\ell_1}\omega_{d-1}^{\ell_2}\alpha_{d-1})^{\gamma^{k'+1}}) q_{\ell_1\ell_2,S_{d-1}}((\omega_0^{\ell}\omega_d^{\ell'}\alpha_d)^{\gamma^{k+1}})) \right) q_{\ell\ell',S_d}(\alpha_{d+1}) \\
&= \sum_{\ell,\ell'} \left(\sum_{\ell_1,\ell_2} \mathcal{G}_{d-2,k_1}(\underline{\alpha}, (\omega_0^{\ell_1}\omega_{d-1}^{\ell_2}\alpha_{d-1})^{\mu_{d-1}^{j'}\gamma^{k''+1}}) q_{\ell_1\ell_2,S_{d-1}}((\omega_0^{\ell}\omega_d^{\ell'}\alpha_d)^{\mu_d^j\gamma^{k'+1}}) \right) q_{\ell\ell',S_d}(\alpha_{d+1}) \\
&= \mathcal{G}_{d,k+1}(\alpha_1, \dots, \alpha_d, \alpha_{d+1})
\end{aligned}$$

for $\alpha_i \in W_i, 1 \leqslant i \leqslant d$ and $\alpha_{d+1} \in \mathcal{S}$. The index k_1 requires an explanation. Notice that k'' is determined by k' by the equation $k' = j'\wp^{d-2} + k''$. Thus, if k' increases to $k'+1$ then we have a new equation $k'+1 = j_1\wp^{d-2} + k_1$ and, depending on whether $k' < j\wp^{d-1} - 2$ or not, $k_1 = k''+1, j_1 = j'$ or $k_1 = 0, j_1 = j'+1$. Either case is covered by the above argument proving the lemma. ■

Degree reduction of the generator $\mathcal{G}_d$ By Remark 28 the variable α_i has degree e_{ik} bounded by $\Lambda \operatorname{ord}(\omega_i) - 1$ for $1 \leqslant i \leqslant d$, and α_{d+1} has degree $r^2 \operatorname{ord}(\omega_d)$, the degree of the interpolating polynomial. Therefore, for each d and $0 \leqslant k \leqslant \wp^d - 1$, there is a polynomial $\mathcal{G}'_{d,k}$ over F in variables $\alpha_i, 1 \leqslant i \leqslant d+1$ such that $\deg(\alpha_i) \leqslant \Lambda\wp^{\kappa+i} - 1$ and $\deg(\alpha_d) = r^2\wp^{\kappa+d}$. Furthermore, $\mathcal{G}'_{d,k}$ agrees with the hitting set generator $\mathcal{G}_{d,k}$ on all the points in the set $W_1 \times W_2 \times \cdots \times W_d \times \mathcal{S}$ for all k. This defines the reduced degree modified generator $\mathcal{G}'_d$. Following Lemma 29, the following is immediate:

$$\operatorname{span}_K \left\{ \prod_{k=0}^{\wp^d-1} M_k(u_k) \right\}_{\text{each } u_k \in F} \subseteq \operatorname{span}_K \left\{ \prod_{k=0}^{\wp^d-1} M_k(\mathcal{G}'_{d,k}(\alpha_1, \dots, \alpha_d, \alpha_{d+1})) \right\}_{\substack{\alpha_{d+1}\in\mathcal{S}, \\ \text{each } \alpha_i \in W_i, 1\leqslant i\leqslant d}} . \quad (9)$$

Remark 31 (Construction of the modified generator)**.** We note that the polynomial $\mathcal{G}'_{d,k}$ has at most $O(\Lambda \cdot p^{2\kappa})^{d+1}$ many monomials. Moreover, its values at any point in $W_1 \times \cdots \times W_d \times \mathcal{S}$ agrees with the polynomial $\mathcal{G}_{d,k}$ which can be evaluated at these points in parallel time $O((d+1) \cdot \log \Lambda)$. Thus, finding the coefficients of $\mathcal{G}'_{d,k}$ can also be solved in parallel time $O((d+1) \cdot \log \Lambda)^{O(1)}$ as solving a system of linear equations is in NC. Thus, if Λ is quasipolynomial, we have a quasi-NC algorithm for computing the hitting set generator polynomial $\mathcal{G}'_d$ in sparse representation.

Theorem 32. *For prime p, let $M_k \in F[u_k]^{r \times r}$ for $1 \leqslant k \leqslant p^d$, where each entry of M_k is of degree at most n. Then the generator $\mathcal{G}'_d : K^{d+1} \to K^{p^d}$ (whose k^{th} component $\mathcal{G}'_{d,k}$ is described above) is a hitting set generator for the ROABP defined by the matrix product $\prod_{k=0}^{p^d-1} M_k(u_k)$. More precisely, if $f(\underline{u})$ is the polynomial computed at the $(1,1)^{th}$ entry of the matrix product $\prod_{k=0}^{p^d-1} M_k(u_k)$ then $f \equiv 0$ if and only if the $(1,1)^{th}$ entry of the matrix product $\prod_{k=0}^{p^d-1} M_k(\mathcal{G}'_{d,k}(\alpha_1, \alpha_2, \ldots, \alpha_{d+1})) = 0$ for all $(\alpha_1, \alpha_2, \ldots, \alpha_{d+1}) \in W_1 \times \cdots \times W_d \times \mathcal{S}$.*

Proof. If $f(\underline{u})$ is nonzero then, as F is an infinite field, we can set $u_i = a_i \in F$ such that $f(\underline{a}) \neq 0$. By Equation (9) there is some $(\alpha_1, \alpha_2, \ldots, \alpha_{d+1}) \in W_1 \times W_2 \times \cdots \times \mathcal{S}$ such that $\prod_{k=0}^{p^d-1} M_k(\mathcal{G}'_{d,k}(\alpha_1, \alpha_2, \ldots, \alpha_{d+1}) \neq 0$. Conversely, suppose $f \equiv 0$. Then for all $a_i \in K, 0 \leqslant i \leqslant p^d - 1$ we have $f(\underline{a}) = 0$. This implies the $(1,1)^{th}$ entry of the matrix product $\prod_{k=0}^{p^d-1} M_k(\mathcal{G}'_{d,k}(\alpha_1, \alpha_2, \ldots, \alpha_{d+1})) = 0$ for all $(\alpha_1, \alpha_2, \ldots, \alpha_{d+1}) \in W_1 \times W_2 \times \cdots \times \mathcal{S}$. ■

3.1.3 Division algebra hitting set construction

Now we are ready to prove the main theorem of the section, a division algebra hitting set construction for noncommutative ABPs.

Theorem 33. *Let $p > 2$ be a prime. For n-variate degree-$\tilde{d}$ noncommutative polynomials computed by homogeneous ABPs of width r over rationals $\mathbb{Q}$, we construct a hitting set $\widehat{\mathcal{H}}_{n,r,\tilde{d}} \subseteq D^n$ of size $(nr\tilde{d})^{O(p \log_p \tilde{d})}$ in $(nr\tilde{d})^{O(p \log_p \tilde{d})}$ time, where D is a cyclic division algebra of index $\ell = p^L$ where $L = O(\log_p \tilde{d})$.*

Proof. We will set $\ell = p^L$ as the index of the division algebra D, where L will be determined in the analysis below. A necessary condition is $p^L \geqslant \tilde{d}$. We assume without loss of generality that $\tilde{d} = p^d$ (and we may increase the degree by a factor of at most p).

A key idea in [FS13] is to convert the given ABP into a set-multilinear form and eventually an ROABP. Specifically, they replace the noncommutative variable x_i by the matrix $M(x_i)$:

$$M(x_i) = \begin{bmatrix} 0 & x_{i1} & 0 & \cdots & 0 \\ 0 & 0 & x_{i2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & x_{i\tilde{d}} \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix}.$$

and the (commuting) variables $x_{i1}, x_{i2}, \ldots, x_{i\tilde{d}}$ will in turn be replaced by powers $u_1^i, u_2^i, \ldots, u_{\tilde{d}}^i$ of commuting variables u_j. The variables u_j will be finally substituted in the [FS13] construction by the output of a generator $\mathcal{G}_d$ (where $\mathcal{p}^{d-1} < \tilde{d} \leqslant \mathcal{p}^d$) that stretches a seed $(\alpha_1, \alpha_2, \ldots, \alpha_{d+1})$ to $(\mathcal{G}_{d,0}(\underline{\alpha}), \mathcal{G}_{d,1}(\underline{\alpha}), \ldots, \mathcal{G}_{d,\mathcal{p}^d-1}(\underline{\alpha}))$.

Obviously, the above matrices are nilpotent and they cannot be elements of a division algebra. Here, our plan will be to replace x_i by the following matrix $M(x_i)$:

$$M(x_i) = \left[\begin{array}{ccccc|ccc} 0 & f_0^i(\underline{\alpha}) & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 0 & f_1^i(\underline{\alpha}) & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & f_{\mathcal{p}^d-1}^i(\underline{\alpha}) & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 & f_{\mathcal{p}^d}^i(\underline{\alpha}) & \cdots & 0 \\ \hline \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & f_{\mathcal{p}^L-2}^i(\underline{\alpha}) \\ z f_{\mathcal{p}^L-1}^i(\underline{\alpha}) & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \end{array}\right],$$

where for each j, $f_j(\underline{\alpha}) = \mathcal{G}_{d,j}(\underline{\alpha})$ is an element from an extension field $K = F(\omega)$ as described. Furthermore, we will ensure that the automorphism σ of K with F as its fixed field has the property that $\sigma(f_j(\underline{\alpha})) = f_{j+1}(\underline{\alpha})$ and $\sigma^{\mathcal{p}^L}$ is the identity automorphism. This will ensure that the matrix substitutions $M(x_i)$ for all choices of $\underline{\alpha}$ from the hitting set will live in the cyclic division algebra defined by (K, F, σ). But this is precisely achieved by the construction of the degree-reduced hitting set generator $\mathcal{G}'_d$ in Theorem 32.

Size of the hitting set

We first analyze the size of the ROABP hitting set of Theorem 32. It is just $|\mathcal{S}| \cdot \prod_{i=1}^{d} |W_i|$. Now $|W_i| = \Lambda \cdot p^{\kappa+i}$ and $|\mathcal{S}| = \Lambda \mathcal{p}^{k+d}$. Thus

$$|\mathcal{S}| \cdot \prod_{i=1}^{d} |W_i| \leqslant \Lambda^{d+1} \cdot \mathcal{p}^{\kappa(d+1)} \prod_{i=1}^{d+1} \mathcal{p}^i.$$

The construction requires $\kappa \geqslant d+1$. Setting $\kappa = d+1$ gives us an upper bound of $\Lambda^{d+1} \cdot \mathcal{p}^{2(d+1)^2}$. By the above construction this is also the size of the hitting set for the noncommutative ABP. Now, by Lemma 23, it suffices to choose $\Lambda > (\mathcal{p}^d nm)^{\mathcal{p}} r^2$, where we can see that $m \leqslant r^2 \mathcal{p}^{2d+1}$. Recall that Λ is a power of some prime number $\mathcal{q}$ (as described in Section 3.1.2). Therefore, we can choose $\Lambda \leqslant \mathcal{q}(\mathcal{p}^d nm)^{\mathcal{p}} r^2$. We choose $\mathcal{p}$ and $\mathcal{q}$ to be two distinct successive primes, and hence $\mathcal{q} \leqslant 2\mathcal{p}$. With this choice of Λ and $\mathcal{q}$, the hitting set size is bounded by $2p(\mathcal{p}^{3d+1} n r^2)^{\mathcal{p}} r^2 \cdot \mathcal{p}^{2(d+1)^2}$. Now, noting that the actual degree satisfies $\mathcal{p}^{d-1} \leqslant \tilde{d} \leqslant \mathcal{p}^d$ we obtain the size bound $((\mathcal{p}\tilde{d})^4 r^2 n)^{\mathcal{p}(\log_{\mathcal{p}} \tilde{d}+1)}$ which is $(\tilde{d} r n)^{O(\mathcal{p} \log_{\mathcal{p}} \tilde{d})}$. ■

3.2 Hitting set for generalized ABPs

In order to apply Theorem 33 to generalized ABPs, we basically describe an efficient reduction that computes a hitting set for generalized ABPs given as input a hitting set for noncommutative ABPs. This is an important conceptual part of the proof.

Informally, the observations contained in the following three lemmas, Lemma 34, Lemma 35, and Lemma 36, show that a nonzero polynomial computed by a D-ABP can be evaluated to nonzero at a point in a cyclic division algebra $D \otimes D'$ where the index of the cyclic division algebra D' is relatively prime to the index of D.

Lemma 34. *For any nonzero n-variate degree-d D-ABP $\mathcal{A}$, for every $d' \geqslant \ell_1 d$, there is a $d' \times d'$ matrix tuple such that the D-ABP is nonzero evaluated on that tuple. Here ℓ is the index of D.*

Proof. Fix an edge of $\mathcal{A}$ and let its label be $\sum_{i=1}^{n} a_i x_i b_i$, for $a_i, b_i \in D_1$. Replace each $a_i, b_i \in D$ by its matrix representation in $\text{Mat}_\ell(K)$ and the variable x_i by Z_i, an $\ell_1 \times \ell_1$ matrix whose $(j,k)^{th}$ entry is a new noncommuting indeterminate z_{ijk}. Therefore, each edge is now labeled by an $\ell \times \ell$ matrix whose entries are K-linear terms in $\{z_{ijk}\}$ variables. After the substitution, $\mathcal{A}$ is now computing a matrix M of degree-d noncommutative polynomials. Clearly, it is an identity-preserving substitution. I.e., $\mathcal{A}$ is nonzero if and only if M is nonzero. Therefore, if $\mathcal{A}$ is nonzero, there exists a $d \times d$ matrix substitution for the $\{z_{ijk}\}$ variables such that M evaluated on that substitution is nonzero.[4] Hence, we obtain an $\ell d \times \ell d$ matrix tuple for the $\underline{x}$ variables such that $\mathcal{A}$ is nonzero on that substitution. ∎

Lemma 35. *Suppose for a nonzero n-variate degree-d D-ABP $\mathcal{A}$,there is a matrix tuple $(p_1, \ldots, p_n) \in \text{Mat}_{d'}(K)^n$ such that the ABP is nonzero evaluated on that tuple. Let $\widetilde{D} = \left(\widetilde{K}/F, \widetilde{\sigma}, z\right)$ be a cyclic division algebra of index d', where K is a subfield of $\widetilde{K}$. Then there is a tuple in $\widetilde{D}^n$ such that the D-ABP $\mathcal{A}$ is nonzero evaluated on that tuple as well.*

Proof. Let $\left\{\widetilde{C}_{j,k}\right\}_{1 \leqslant j,k \leqslant d'}$ be the basis of the division algebra $\widetilde{D}$ as defined in Section 2.3. By Proposition 13, we can write each matrix $p_i = \sum_{j,k} \lambda_{ijk} \widetilde{C}_{jk}$ where each $\lambda_{ijk} \in \widetilde{K}$. Define new commuting indeterminates $\{u_{ijk}\}$ and let $\widetilde{p}_i = \sum u_{ijk} \widetilde{C}_{jk}$. Evaluating $\mathcal{A}$ on $(\widetilde{p}_1, \ldots, \widetilde{p}_n)$ then gives a nonzero matrix of commutative polynomials, as it is nonzero if $u_{ijk} \leftarrow \lambda_{ijk}$. There is a substitution for each $u_{ijk} \leftarrow \gamma_{ijk} \in \mathbb{Q}$ such that such a nonzero polynomial evaluates to nonzero. Hence, we can define a tuple $(q_1, \ldots, q_n)$ where each $q_i = \sum \gamma_{ijk} \widetilde{C}_{jk}$ such that $\mathcal{A}$ is nonzero on $(q_1, \ldots, q_n)$. Now the proof follows since each $q_i \in \widetilde{D}$. ∎

Lemma 36. *For any nonzero polynomial computed by an n-variate degree-d D-ABP $\mathcal{A}$,there is a cyclic division algebra $\widetilde{D}$ of index $\ell\ell'$ (where $\ell' \geqslant d$ and ℓ' is relatively prime to ℓ) and a tuple in $\widetilde{D}^n$ such that $\mathcal{A}$ is nonzero evaluated on that tuple.*

Proof. Consider a cyclic division algebra D' of index ℓ'. Define $\widetilde{D} = D \otimes D'$. By assumption, $\ell'(\geqslant d)$ is relatively prime to ℓ. More precisely, suppose $D = (K/F, \sigma, z)$ and $D' = (K'/F', \sigma', z')$, where $F = \mathbb{Q}(z, \omega_0)$, $K = F(\omega_1)$, $F' = \mathbb{Q}(z', \omega'_0)$ and $K' = F'(\omega'_1)$, where ω_0, ω'_0, ω_1, and ω'_1 are roots of unity whose orders are relatively prime to each other. Then we define the field $F_1 = \mathbb{Q}(z, \omega_0\omega'_0)$, and the fields $K_1 = F_1(\omega_1)$ and $K'_1 = F_1(\omega'_1)$ and note that the cyclic division algebra D is isomorphic to $(K_1/F_1, \hat{\sigma}, z)$ and D' is isomorphic to $(K'_1/F_1, \hat{\sigma}', z)$ for suitably defined automorphisms $\hat{\sigma}$ and $\hat{\sigma}'$. Now by Theorem 14 it follows that $\widetilde{D}$ is also a cyclic division algebra of index $\ell\ell'$ defined using the field extension $F_1(\omega_1\omega_2)$ of F_1. The proof follows from Lemma 34 and Lemma 35. ∎

[4]In fact, $\lceil d/2 \rceil + 1$-dimensional matrix substitutions will suffice [AL50].

The following lemma now reduces the problem of zero testing of a division algebra ABP to a corresponding noncommutaitve ABP.

Lemma 37. *Let D be a cyclic division algebra of index ℓ and $\mathcal{A}$ be an n-variate degree-d D-ABP of width r over $\{x_1, x_2, \ldots, x_n\}$. Then, there is an $\ell^2 n$-variate degree-d noncommutative ABP $\mathcal{B}$ of width ℓr such that the following holds:*

- *The polynomial f computed by the D-ABP $\mathcal{A}$ is nonzero if and only if the polynomial g computed by the noncommutative ABP $\mathcal{B}$ is nonzero.*
- *For any matrix tuple $(q_{111}, \ldots, q_{n\ell\ell})$, if the polynomial g computed by $\mathcal{B}$ is nonzero on that tuple, then the polynomial f computed by the D-ABP $\mathcal{A}$ evaluated on $(q_1, \ldots, q_n)$ is also nonzero where for each i, $q_i = \sum_{j,k} C_{jk} \otimes q_{ijk}$ and $\{C_{jk}\}_{1 \leqslant j,k \leqslant \ell}$ is a basis of D.*

Proof. Introduce a set of noncommuting indeterminates $\{y_{ijk}\}_{i\in[n],j,k\in[\ell_1]}$. Consider the following mapping:

$$x_i \mapsto \sum_{j,k} C_{jk} \otimes y_{ijk}.$$

Equivalently, each x_i is substituted by an $\ell \times \ell$ matrix. Fix a D-ABP $\mathcal{A}$. Consider each edge of $\mathcal{A}$ labeled as $\sum_{i=1}^{n} a_i x_i b_i$ where $a_i, b_i \in D$. Replace each $a_i, b_i \in D$ by its matrix representation in $\mathrm{Mat}_\ell(K)$ and x_i by the $\ell \times \ell$ matrix $\sum_{j,k} C_{jk} \otimes y_{ijk}$. Therefore, each edge is now labeled by an $\ell \times \ell$ matrix whose entries are K-linear terms in $\{y_{ijk}\}$ variables. After the substitution, $\mathcal{A}$ is now computing a matrix M of degree-d noncommutative polynomials in $\{y_{ijk}\}$ variables.

Claim 38. *If the D-ABP $\mathcal{A}(\underline{x})$ is nonzero then the matrix $M \in \mathrm{Mat}_\ell(\mathbb{F}\langle \underline{y} \rangle)$ is nonzero.*

Proof. If $\mathcal{A}(\underline{x})$ is nonzero, then it is nonzero evaluated at some $(p_1, \ldots, p_n) \in \widetilde{D}^n$ where $\widetilde{D} = D \otimes D'$ (Lemma 36). We can therefore expand the D component in the $\{C_{jk}\}$ basis and write each $p_i = \sum C_{jk} \otimes q_{ijk}$ for some $q_{ijk} \in D'$. Therefore M is nonzero under the substitution each $y_{ijk} \leftarrow q_{ijk}$. ∎

We now claim that each entry of M is computable by a small ABP.

Claim 39. *For each $1 \leqslant j, k \leqslant \ell_1$, the $(j,k)^{th}$ entry of the matrix $M \in \mathrm{Mat}_\ell(\mathbb{F}\langle \underline{y} \rangle)$ is computable by an $\ell^2 n$-variate degree-d noncommutative homogeneous ABP of width ℓr.*

Proof. For each vertex v in the D-ABP $\mathcal{A}$, make ℓ copies of v (including the source S and sink T), let us call it $(v,1), \ldots, (v,\ell_1)$. For any two vertices u and v, suppose the edge is labeled by $\sum_{i=1}^{n} a_i x_i b_i$ and $M_{u,v}$ be the corresponding $\ell \times \ell$ matrix after substitution. Then for each $1 \leqslant \hat{j}, \hat{k} \leqslant \ell$, we add an edge $((u,\hat{j}),(v,\hat{k}))$ labeled by the $(\hat{j},\hat{k})^{th}$ entry of $M_{u,v}$. Note that product of the edge labels of a path exactly captures the corresponding matrix product. Therefore, if we consider the ABP with source $(S,\hat{j})$ and sink $(T,\hat{k})$, it is computing the $(\hat{j},\hat{k})^{th}$ entry of the matrix M. Note that the width of the new ABP is ℓr. ∎

We now consider a nonzero entry of the matrix M which is computable by an $\ell^2 n$-variate degree-d noncommutative homogeneous ABP of width ℓr. ∎

Remark 40. The map $x_i \mapsto \sum C_{jk} \otimes y_{ijk}$ transforms a D-ABP into an element of $\mathrm{Mat}_\ell(\mathbb{F}\langle \underline{y} \rangle)$ where ℓ is the index of D. Note that, the inclusion map $\iota : \mathrm{Mat}_\ell(\mathbb{F}) \to \mathrm{Mat}_{d\ell}(\mathbb{F})$ defined by $\iota(a) = a \otimes I_d$ is implicitly assumed here (see Section 2). If a different inclusion map is applied for evaluating the D-ABP, the mapping must be modified accordingly. For instance, if the inclusion map is defined by $\iota(a) = I_d \otimes a$, then the mapping becomes $x_i \mapsto \sum y_{ijk} \otimes C_{jk}$.

We are now ready to prove the main result of this section.

Theorem 41 (Division algebra hitting set for D-ABPs)**.** *Fix integers n, r, d. Let $\widehat{\mathcal{H}}_{n,r,d}$ be a division algebra hitting set for the class of n-variate width-r noncommutative ABPs of degree d. Then, for any cyclic division algebra D of index ℓ, we can design a hitting set $\widehat{\mathcal{H}}^D_{n,r,d} \subseteq \widetilde{D}^n$ for the class of n-variate width-r D-ABPs of degree d as follows:*

$$\widehat{\mathcal{H}}^D_{n,r,d} = \left\{ (q_1, \ldots, q_n) : q_i = \sum_{j,k} C_{jk} \otimes q_{ijk} \quad \textit{where} \quad (q_{111}, \ldots, q_{n\ell\ell}) \in \widehat{\mathcal{H}}_{\ell^2 n, \ell r, d} \right\}, \tag{10}$$

where $\widehat{\mathcal{H}}_{\ell^2 n,\ell r,d} \subseteq D'^n$ such that D' is a cyclic division algebra of index ℓ' which is relatively prime to ℓ and $\widetilde{D} = D \otimes D'$ is a cyclic division algebra of index $\ell\ell'$.

Proof. From the assumption, $\widetilde{D} = D \otimes D'$. Note that, it is a cyclic division algebra of index $\ell\ell'$ by Theorem 14. Our is now to output a division algebra hitting set for a D-ABP in the cyclic division algebra $\widetilde{D}^n$. which immediately follows from Lemma 37. ∎

The above theorem combined with Theorem 33 immediately yields the following corollary.

Corollary 42. *Fix two distinct prime numbers $\mathfrak{p}$ and $\mathfrak{p}'$. For the class of polynomials computed by n-variate D-ABPs of degree-d and width r, where D is a cyclic division algebra of index $\ell = \mathfrak{p}^L$, there is an algorithm that outputs a hitting set $\widehat{\mathcal{H}}_{\ell^2 n,\ell r,d} \subseteq \widetilde{D}^n$ of size $(\ell_1 n r d)^{O(\mathfrak{p}' \log d)}$ and $\widehat{D}$ is a cyclic division algebra of index $\ell\ell'$, for $\ell' = \mathfrak{p}'^{L'}$ where L' is $O(\log_{\mathfrak{p}'} d)$.*

4 Hitting Set for NSingular Conditioned on a Matrix Tuple

Fix a matrix tuple $\underline{p} = (p_1, \ldots, p_n)$. For any integer s, consider the class of linear matrices such that a submatrix of size $s - 1$ is invertible on $\underline{p}$. We say *the class of linear matrices is conditioned on the matrix tuple $\underline{p}$*. In this section, we construct a hitting set for the NSingular problem for the class of linear matrices conditioned on $\underline{p}$.

The result of this section is crucial for the hitting set construction for rational formulas in Section 5. More precisely, we construct the hitting set for rational formulas inductively on the inversion height. To construct a hitting set for inversion height θ from inversion height $\theta - 1$, we consider the union of all hitting sets conditioned on a matrix tuple in the hitting set of height $\theta - 1$.

Lemma 43. *Let D be a cyclic division algebra of index ℓ and for any n, $(p_1, \ldots, p_n) \in D^n$ be a matrix tuple. Fix an integer s. Consider an n-variate linear matrix T of size s, conditioned on $(p_1, \ldots, p_n)$. Then, there is an n-variate degree-$(s-1)$ D-ABP $\mathcal{A}$ of width $s - 1$ such that the following holds:*

- *The linear matrix T is invertible over the free skew field if and only if the D-ABP $\mathcal{A}$ is nonzero.*
- *For any integer ℓ' and a matrix tuple $(q_1, \ldots, q_n) \in \mathrm{Mat}_{\ell\ell'}(\mathbb{Q})^n$, if $\mathcal{A}$ is invertible on that tuple, then for some indeterminate t, T evaluated on $(tq_1 + p_1 \otimes I_{\ell'}, \ldots, tq_n + p_n \otimes I_{\ell'})$ is invertible.*

Proof. Let $T(\underline{p})$ is not invertible. We can then find two invertible transformations U, V in $\mathrm{Mat}_s(D)$ such that

$$U \cdot T(p_1, p_2, \ldots, p_n) \cdot V = \left[\begin{array}{c|c} I_{s-1} & 0 \\ \hline 0 & 0 \end{array}\right],$$

where I_{s-1} is the identity matrix whose diagonal elements are the identity element of D. This is possible since one can do Gaussian elimination over division algebras.

Notice that $T(\underline{x} + \underline{p}) = T(\underline{p}) + T(\underline{x})$. Hence, we can write

$$T(\underline{x} + \underline{p}) = U^{-1} \cdot \left(\left[\begin{array}{c|c} I_{s-1} & 0 \\ \hline 0 & 0 \end{array}\right] + U \cdot T(\underline{x}) \cdot V\right) \cdot V^{-1}.$$

Let the invertible submatrix T' of T of size $s-1$ is obtained by removing the i^{th} row and j^{th} column, for some $i, j \in [s]$. We can therefore write,

$$T(\underline{x} + \underline{p}) = U^{-1} \cdot \left[\begin{array}{c|c} I_{s-1} - \widetilde{T} & A_j \\ \hline B_i & c_{ij} \end{array}\right] \cdot V^{-1},$$

where each entry of $\widetilde{T}, A_j, B_i, c_{ij}$ are D-linear forms in $\underline{x}$ variables with no constant term. We can simplify it further by multiplying both sides by invertible matrices and writing,

$$T(\underline{x} + \underline{p}) = U^{-1}U' \left[\begin{array}{c|c} I_{s-1} - \widetilde{T} & 0 \\ \hline 0 & c_{ij} - B_i(I_{s-1} - \widetilde{T})^{-1}A_j \end{array}\right] V'V^{-1}. \tag{11}$$

$$\text{where, } U' = \left[\begin{array}{c|c} I_{s-1} & 0 \\ \hline B_i(I_{s-1} - \widetilde{T})^{-1} & 1 \end{array}\right], \qquad V' = \left[\begin{array}{c|c} I_{s-1} & (I_{s-1} - \widetilde{T})^{-1}A_j \\ \hline 0 & 1 \end{array}\right].$$

Note that, here 1 denotes the unit in the division algebra D.

$$\text{Let,} \quad P_{ij}(\underline{x}) = c_{ij} - B_i(I_{s-1} - \widetilde{T})^{-1}A_j.$$

We can also represent P_{ij} as a formal series:

$$P_{ij}(\underline{x}) = c_{ij} - B_i \left(\sum_{k \geqslant 0} \widetilde{T}^k\right) A_j. \tag{12}$$

This is a generalized series (in $\underline{x}$ variables) over the division algebra D where the division algebra elements can interleave in between the variables. The following claim truncates this infinite series to a finite value.

Claim 44. *Consider a generalized D-series P as defined in Equation (12).*

$$P(\underline{x}) = c - B\left(\sum_{k\geqslant 0}\widetilde{T}^k\right)A,$$

$$\textit{Define its truncation:}\quad \widetilde{P}(\underline{x}) = c - B\left(\sum_{0\leqslant k\leqslant s-1}\widetilde{T}^k\right)A.$$

$$\textit{Then } P(\underline{x}) = 0 \iff \widetilde{P}(\underline{x}) = 0.$$

Proof. Suppose P is nonzero. Substitute each $\{x_i : 1 \leqslant i \leqslant n\}$ by the following map used in the proof of Theorem 41:

$$x_i \mapsto \sum_{j,k} C_{jk} \otimes y_{ijk}.$$

Consider an entry of $\widetilde{T}$ which is of form $\sum_{i=1}^{n} a_i x_i b_i$ for some $a_i, b_i \in D$. Since $C_{jk}, 1 \leqslant j, k \leqslant \ell$ is a basis for the division algebra D, we can write each entry of $\widetilde{T}$ as $\sum \beta_{ijk}\beta'_{ij'k'} C_{jk} x_i C_{j'k'}$ for some $\beta_{ijk}, \beta_{ij'k'} \in F$. Substituting each x_i as above and identifying each C_{jk} with $C_{jk} \otimes 1$, it follows that each entry of $\widetilde{T}$ can be expressed as $\sum_{j,k}(C_{jk} \otimes \sum_i \alpha_{ijk} y_{ijk})$, where each $\alpha_{ijk} \in F$. Therefore, it now computes a series $\sum C_{jk} \otimes f_{jk} \in D \otimes_F F\langle\langle \underline{y} \rangle\rangle$. We first observe the following claim. Its proof is omitted as it is a straightforward generalization of the proof of Claim 38.

Claim 45. $P(\underline{x}) = 0 \iff \sum C_{jk} \otimes f_{jk} = 0.$

Recall that, $D\langle\langle y \rangle\rangle$ denotes the formal power series in noncommuting $\underline{y}$ variables where the coefficients are in D and $\underline{y}$ variables commute with the elements in D. We now define the following map:

$$\psi : D \otimes_F F\langle\langle y \rangle\rangle \to D\langle\langle y \rangle\rangle,$$
$$C_{jk} \otimes y_{ijk} \mapsto C_{jk} y_{ijk}.$$

Note that, ψ is an isomorphism. Each entry of the matrix L is now of form $\sum_{i,j,k} \gamma_{ijk} y_{ijk}$ (where $\gamma_{ijk} \in D$). Therefore, substituting each $x_i \mapsto \sum_{j,k} C_{jk} \otimes y_{ijk}$ and then applying ψ-map on P computes a series in $D\langle\langle y \rangle\rangle$. We can now apply Fact 15 and truncate it to degree $s-1$ preserving the nonzeroness.

Clearly, applying the substitution $x_i \mapsto \sum_{j,k} C_{jk} \otimes y_{ijk}$ and then the ψ-map on $\widetilde{P}$ will have the same effect. Therefore, $\widetilde{P}$ is also nonzero. ∎

We can now write, $\quad P_{ij} = 0 \iff \left(c_{ij} = 0 \quad \text{and} \quad \text{for each } 0 \leqslant k \leqslant (s-1)\ell_1, \quad B_i\widetilde{T}^k A_j = 0\right),$

where each $B_i\widetilde{T}^k A_j$ is a generalized polynomial over D, indeed it is a D-ABP.

Let k_0 be the minimum k such that $B_i L^k A_j \neq 0$. By definition, $k_0 \leqslant s-1$. Define, a generalized ABP $\widetilde{P}_{ij} = B_i L^{k_0} A_j$. We now reduce the singularity testing of T to identity testing of this D-ABP.[5]

[5]In a recent work [CM23], a similar idea is used to show a polynomial-time reduction form NSingular to identity testing of noncommutative ABPs in the white-box setting.

Let $\widetilde{P}_{ij}$ be zero, therefore P_{ij} is also zero by Claim 44. Following Equation (11), $T(\underline{x}+\underline{p})$, and hence, $T(\underline{x})$ is not invertible.

For the other direction, if $\widetilde{P}_{ij} \neq 0$, by assumption there exists a matrix tuple $(q_1, \ldots, q_n) \in \mathrm{Mat}_{\ell\ell'}(\mathbb{Q})$ such that $\widetilde{P}_{ij}(\underline{q})$ is invertible. We then evaluate $T(\underline{x}+\underline{p})$ on $(tq_1, \ldots, tq_n)$ where t is a commutative variable. Clearly, the infinite series P_{ij} is nonzero at $t\underline{q}$ since the different degree-t parts do not cancel each other. Also, $(I_{s-1} - \widetilde{T})(t\underline{q})$ is invertible. Therefore, by definition $T(\underline{x}+\underline{p})$ is also invertible.

It completes the proof. ∎

Theorem 46. *Fix integers n, r, d. Let D be a cyclic division algebra of index ℓ and $\widehat{\mathcal{H}}^{D}_{n,r,d} \subseteq \widehat{D}^n$ be a division algebra hitting set for the class of n-variate width-r D-ABP of degree d where $\widehat{D}$ is of index $\ell\ell'$. Then, given a matrix tuple $\underline{p} = (p_1, \ldots, p_n) \in D^n$, we can design a hitting set $\widetilde{\mathcal{H}}^{\underline{p}}_{n,s,\ell}$ for the class of n-variate linear matrices of size s conditioned on tuple $\underline{p}$ as follows:*

$$\widetilde{\mathcal{H}}^{\underline{p}}_{n,s,\ell} = \left\{(aq_1 + p_1 \otimes I_{\ell'}, \ldots, aq_n + p_n \otimes I_{\ell'}) : \underline{q} \in \widehat{\mathcal{H}}^{D}_{n,s-1,s-1} \quad \text{and} \quad a \in \Gamma, |\Gamma| = (s\ell\ell')^{O(1)}\right\}. \tag{13}$$

Proof. The proof immediately follows from Lemma 43 except the bound on Γ. Note that, For any linear matrix T of size s, consider the matrix T evaluated on $(tq_1 + p_1 \otimes I_{\ell'}, \ldots, tq_n + p_n \otimes I_{\ell'})$. Clearly, it is a matrix of size $s\ell\ell'$ where each entry is a linear polynomial in t. Consider the rational expressions in the inverse of this matrix. Since the degrees of the polynomials in the rational expressions and the determinant are bounded by a polynomial in the size of the matrix, and we need to only avoid the roots of the numerator and the denominator polynomials present in each entry, we can vary the parameter t over a set $\Gamma \subset \mathbb{Q}$ such that $|\Gamma| = (s\ell\ell')^{O(1)}$. ∎

Corollary 47. *Let D be a cyclic division algebra of index ℓ and $\wp$ be any prime that is not a divisor of ℓ. Given a tuple $(p_1, \ldots, p_n) \in D^n$, consider the class of n-variate linear matrix of size s conditioned on $(p_1, \ldots, p_n) \in D^n$. We can construct a hitting set $\widetilde{\mathcal{H}}^{\underline{p}}_{n,s,\ell} \subseteq \widetilde{D}^n$ of size $(\ell n s)^{O(\wp \log s)}$ in deterministic $(\ell n s)^{O(\wp \log s)}$-time for this class where $\widetilde{D}$ is a cyclic division algebra of index $\ell s^{O(1)}$.*

Proof. The proof follows from the combination of Theorem 41 and Theorem 46.

∎

Remark 48. As previously discussed in Remark 40, the witness matrix tuple is also influenced by the choice of the inclusion map used to evaluate the D division algebra ABP. More precisely, if the inclusion map $\iota : \mathrm{Mat}_{\ell_1}(\mathbb{F}) \to \mathrm{Mat}_{\ell_1\ell_2}(\mathbb{F})$ is applied for evaluating the D-ABP, then the hitting set is redefined as follows:

$$\widetilde{\mathcal{H}}^{\underline{p}}_{n,s,\ell_1} = \left\{(tq_1 + \iota(p_1), \ldots, tq_n + \iota(p_n)) : \underline{q} \in \widehat{\mathcal{H}}^{D}_{n,s-1,s-1} \quad \text{and} \quad a \in \Gamma\right\}.$$

5 Derandomizing Black-box RIT

In this section we will first prove Theorem 2 by explaining the quasipolynomial-size hitting set construction for rational formulas along with analyzing its size. Then we will outline the quasi-NC white-box RIT algorithm for rational formulas.

5.1 Hitting set construction for rational formulas

The hitting set construction is by induction on the inversion height of a rational formula. We will show that for every inversion height θ we can construct a hitting set $\mathcal{H}_{n,s,\theta} \subseteq D_\theta^n$ as claimed, where D_θ is a cyclic division algebra of index ℓ_θ. The base case $\theta = 0$ is for noncommutative formulas (which have inversion height 0). By Theorem 33 we have such a hitting set construction of size $(ns)^{O(\log s)}$ for noncommutative ABPs (and hence for noncommutative formulas without inversion gates).

Inductively assume that we have such a construction for the class of n-variate rational formulas of size s and inversion height $\theta - 1$. Let $\Phi(\underline{x})$ be any rational formula of inversion height θ in $\mathbb{Q}\langle\!\langle \underline{x} \rangle\!\rangle$ of size s. We first show the following.

Lemma 49. *For every rational formula Φ of inversion height θ in $\mathbb{Q}\langle\!\langle \underline{x} \rangle\!\rangle$ of size s, there exists a hitting set point $\underline{p} \in \mathcal{H}_{n,s,\theta-1}$ at which $\Phi(\underline{p})$ is defined.*

Proof. Let $\mathcal{F} = \{\mathfrak{g}_i \mid 1 \leqslant i \leqslant k\}$ be the subset of all inverse gates $\mathfrak{g}_i$ in the formula Φ such that there are no other inverse gates on the path from the output gate to the gate $\mathfrak{g}_i$. For each $\mathfrak{g}_i \in \mathcal{F}$, let ϕ_i be the subformula that is input to gate $\mathfrak{g}_i$. Now, consider the product formula $\phi = \phi_1\phi_2 \cdots \phi_k$ (where $k = |\mathcal{F}|$). Notice that the formula ϕ is size at most s since for each $i \neq j$, the subformulas ϕ_i and ϕ_j are disjoint (and we can account for the $k - 1$ new product gates in ϕ for the k gates $\mathfrak{g}_i$ of Φ). Furthermore, we note that the formula ϕ is of inversion height $\theta - 1$. Therefore, as $\mathcal{H}_{n,s,\theta-1}$ is a division algebra hitting set for ϕ, for some $\underline{p} \in \mathcal{H}_{n,s,\theta-1}$, $\phi(\underline{p})$ is nonzero and hence invertible. Consequently, each ϕ_i is also invertible at $\underline{p}$. As the path from the output gate of Φ to each $\mathfrak{g}_i$ has no other inverse gate, it follows that $\Phi(\underline{x})$ is defined at $\underline{p}$. ∎

If the rational formula Φ is of size s then, as shown in [HW15, Theorem 2.6], the formula Φ can be represented as the top-right corner of the inverse of a linear matrix of size at most $2s$. More precisely, $\Phi(\underline{x}) = u^t \mathcal{L}^{-1} v$ where $\mathcal{L}$ is a linear matrix of size at most $2s$ and $u, v \in \mathbb{Q}^{2s}$ are $2s$-dimensional column vectors whose first (resp. last) entry is 1 and others are zero. Here u^t denotes the transpose of u. Therefore, Φ^{-1} can be written as[HW15, Equation 6.3]:

$$\Phi^{-1}(\underline{x}) = [1\ 0\ \ldots\ 0] \cdot \widehat{\mathcal{L}}^{-1} \cdot \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix} \qquad \text{where} \quad \widehat{\mathcal{L}} = \left[\begin{array}{c|c} v & \mathcal{L} \\ \hline 0 & -u^t \end{array}\right].$$

By Lemma 49 the formula Φ is defined for some $\underline{p} \in \mathcal{H}_{n,s,\theta-1}$. Therefore, $\mathcal{L}$ is invertible at $\underline{p}$ (see [HW15, Proposition 7.1]).

Our goal is to find a point (that it, a tuple over a suitable division algebra) such that Φ evaluates to a nonzero and hence invertible value in the division algebra. Equivalently, Φ^{-1} needs to be invertible at that tuple, and therefore $\widehat{\mathcal{L}}$ is invertible when evaluated at that tuple [HW15, Proposition 7.1].

Note that $\widehat{\mathcal{L}}$ is of size at most $2s + 1$. Moreover, we know a tuple $\underline{p} \in \mathcal{H}_{n,s,\theta-1}$ such that a submatrix $\mathcal{L}$ of $\widehat{\mathcal{L}}$ of size $2s$ is invertible on $\underline{p}$. We can now use the construction of $\widetilde{\mathcal{H}}^{p}_{n,2s+1,\ell_{\theta-1}}$ (where $\ell_{\theta-1}$ is the index of the cyclic division algebra $D_{\theta-1}$), as described in Theorem 46, to find a tuple $\underline{q}$ inside a division algebra of index ℓ_θ such that $\widehat{\mathcal{L}}(\underline{q})$ is invertible, therefore $\Phi(\underline{q})$ is nonzero.

We now obtain the following hitting set $\mathcal{H}_{n,s,\theta}$ for the class of n-variate noncommutative rational formulas of height θ and size s defined by the three equations below:

$$\mathcal{H}_{n,s,\theta} = \bigcup_{\underline{p} \in \mathcal{H}_{n,s,\theta-1} \subseteq D_{\theta-1}^n} \widetilde{\mathcal{H}}^{\underline{p}}_{n,2s+1,\ell_{\theta-1}}, \tag{14}$$

$$\widetilde{\mathcal{H}}^{\underline{p}}_{n,2s+1,\ell_{\theta-1}} = \left\{(aq_1 + p_1 \otimes I_\ell, \ldots, aq_n + p_n \otimes I_\ell) : \underline{q} \in \widehat{\mathcal{H}}^{D_{\theta-1}}_{n,2s,2s+1},\ \ell_\theta = \ell_{\theta-1}\ell',\ \text{and } a \in \Gamma\right\}, \tag{15}$$

$$\widehat{\mathcal{H}}^{D_{\theta-1}}_{n,2s,2s+1} = \left\{(q_1, \ldots, q_n) : q_i = \sum_{j,k} C_{jk} \otimes q_{ijk}\ :\ (q_{111}, \ldots, q_{n\ell_{\theta-1}\ell_{\theta-1}}) \in \widehat{\mathcal{H}}_{\ell^2_{\theta-1}n,2\ell_{\theta-1}s,2s+1}\right\}, \tag{16}$$

Here $\widetilde{\mathcal{H}}^{\underline{p}}_{n,2s+1,\ell_{\theta-1}}$ is the hitting set for the class of linear matrices of size $2s+1$ conditioned on a matrix tuple $\underline{p} \in D^n_{\theta-1}$ and $\Gamma \subseteq Q$ is a set of size $(s\ell_{\theta-1}\ell)^{O(1)}$ (Theorem 46), $\widehat{\mathcal{H}}^{D_{\theta-1}}_{n,2s,2s+1} \subseteq D^n_\theta$ is a division algebra hitting set for an n-variate width-$2s$ and degree-$2s+1$ $D_{\theta-1}$-ABP (Theorem 41) and finally $\widehat{\mathcal{H}}_{\ell^2_{\theta-1}n,2\ell_{\theta-1}s,2s+1}$ is a division algebra hitting set for $\ell^2_{\theta-1}n$-variate width-$2\ell_{\theta-1}$ and degree-$2s+1$ noncommutative ABP (Theorem 33).

Let ℓ be the dimension of the matrices in $\widehat{\mathcal{H}}_{\ell^2_{\theta-1}n,2\ell_{\theta-1}s,2s+1}$. By Theorem 33 we obtain $\ell \leqslant s^{O(1)}$.

Noting that $\ell_\theta = \ell \cdot \ell_{\theta-1}$, we now have $\ell_\theta = s^{O(\theta)}$.

We also have from Equation (14), Equation (15), Equation (16) that,

$$|\mathcal{H}_{n,s,\theta}| = |\mathcal{H}_{n,s,\theta-1}| \cdot |\widehat{\mathcal{H}}_{\ell^2_{\theta-1}n,2\ell_{\theta-1}s,2s+1}| \cdot |\Gamma|.$$

We now apply Theorem 46 to bound $|\Gamma|$ by $(s\ell\ell_{\theta-1})^{O(1)}$. Therefore, $|\Gamma| = s^{O(\theta)}$. Let p_θ be the θ^{th} prime number. In order to apply Theorem 33, and consequently Theorem 41, we choose $p_{\theta+1}$ and $p_{\theta+2}$ in the construction of $\widehat{\mathcal{H}}_{\ell^2_{\theta-1}n,2\ell_{\theta-1}s,2s+1}$. As evident in the proof of Theorem 33, the size of the hitting set is bounded by $(ns^2\ell^3_{\theta-1})^{O(\theta^2 \log s)}$ as $p_\theta \leqslant \theta^2$.

$$|\widehat{\mathcal{H}}_{\ell^2_{\theta-1}n,2\ell_{\theta-1}s,2s+1}| \cdot |\Gamma| \leqslant (ns^2\ell^3_{\theta-1})^{O(\theta^2 \log s)} \cdot s^{O(\theta)} \leqslant (ns\ell_{\theta-1})^{O(\theta^2 \log s)}.$$

As shown above $\ell_{\theta-1} \leqslant s^{O(\theta)}$. Hence, we now obtain

$$|\mathcal{H}_{n,s,\theta}| = |\mathcal{H}_{n,s,\theta-1}| \cdot (ns\ell_{\theta-1})^{O(\theta^2 \log s)} \leqslant |\mathcal{H}_{n,s,\theta-1}|(ns)^{O(\theta^3 \log s)}.$$

Therefore, $|\mathcal{H}_{n,s,\theta}| \leqslant (ns)^{O(\theta^5 \log s)}$.

Final steps Note that $\mathcal{H}_{n,s,\theta} \subseteq D^n_\theta$. From our construction, the entry of each matrix in the hitting set is in $\mathbb{Q}(z, \underline{\omega}_0, \underline{\omega})$ where z is a commuting indeterminate; $\underline{\omega}_0 = (\omega_{0,1}, \ldots, \omega_{0,\theta})$ where each $\omega_{0,i}$ is a root of unity of order $\Lambda_i = q_i^\tau$; and $\underline{\omega} = (\omega_1, \ldots, \omega_\theta)$ where each ω_i is a root of unity of order p_i^τ, and all $\{p_i\}_{0\leqslant i\leqslant\theta}$ and $\{q_i\}_{0\leqslant i\leqslant\theta}$ are distinct primes.

Firstly, we need to bound the parameter Λ_θ for each inversion height θ. As described in the construction of Theorem 33, $\Lambda_\theta = (ns\ell_{\theta-1})^{O(\theta^2)}$. From the bound on ℓ_θ, therefore, $\Lambda_\theta = (ns)^{O(\theta^3)}$. Note that, the degree of each ω_i and of z is of degree $\leqslant s^{O(\theta)}$.

We now discuss how to obtain a hitting set over $\mathbb{Q}$ itself. In the hitting set points we replace (as discussed in Section 1.2) $\underline{\omega}$ and $\underline{\omega_0}$ by commuting indeterminates $\underline{t}_1, \underline{t}_2$. Then, for any nonzero rational formula Φ of size s there is a matrix tuple in the hitting set on which Φ evaluates to a nonzero matrix $M(z, \underline{t}_1, \underline{t}_2)$ of dimension $s^{O(\theta)}$ over the commutative function field $\mathbb{Q}(z, \underline{t}_1, \underline{t}_2)$. Each entry of $M(z, \underline{t}_1, \underline{t}_2)$ is a commutative rational expression of the form a/b, where a and b are polynomials in z, $\underline{t}_1$ and $\underline{t}_2$ and the degrees of both a and b are $(ns)^{O(\theta^3)}$. We can now vary the variables $z, \underline{t}_1, \underline{t}_2$ over a $(ns)^{c\theta^3}$-size large set $\widetilde{T} \subseteq \mathbb{Q}$ for some sufficiently large constant c such that it avoids the roots of all the numerator and denominator polynomials involved in the computation. Therefore, finally we obtain the hitting set $\mathcal{H}_{n,s,\theta} \subseteq \mathrm{Mat}_{\ell_\theta}(\mathbb{Q})$ (with a slight abuse of notation) where

$$\ell_\theta \leqslant s^{O(\theta)} \qquad \text{and,} \quad |\mathcal{H}_{n,s,\theta}| \leqslant (ns)^{O(\theta^4 \log s)}.$$

It completes the proof of Theorem 2.

5.2 RIT is in quasi-NC

Recall that, NC is the class of problems which can be solved in poly-logarithmic time using polynomially many processors in parallel. Similarly, quasi-NC is the class of problems which can be solved in poly-logarithmic time using quasipolynomially many processors in parallel. We now prove Corollary 4, by outlining a quasi-NC RIT algorithm for rational formulas in the white-box setting. The proof consists of two steps. Firstly we show that the hitting set presented in the last section can be constructed in quasi-NC. We then show that given a matrix tuple, a noncommutative rational formula can be evaluated in NC. We now describe each step.

Step 1. Quasi-NC hitting set construction: Firstly, note that the matrix operations like additions, and tensor products are routinely computable in NC. The hitting set generator (see Equation (7)) requires repeated application of σ to a root of unity, say b, in order to compute the sequence $\sigma^i(b), 0 \leqslant i \leqslant \wp^L - 1$. However, as the automorphism σ is defined by exponentiation, we can compute all these $\sigma^i(b)$ in parallel by exponentation. Now, for each i we will need to compute a power b^N for some N which can be done in $\log N$ parallel rounds by repeated squaring. Moreover, for the construction described in Equation (7), the size of N in binary is bounded by quasipolynomial in the input size. Hence this computation too is in quasi-NC. Putting it together, as argued in Remark 31, the division algebra hitting set construction for the noncommutative ABPs is in quasi-NC.

Now, from the description of the hitting set $\mathcal{H}_{n,s,\theta}$ given in Section 5.1 via Equation (14), Equation (15), Equation (16), it is easy build the hitting set in quasi-NC by induction on θ. It is directly based on θ many constructions of division algebra hitting sets for noncommutative ABPs (with different size and width parameters at each stage). Each of these θ many computations is in quasi-NC as we outlined above and θ itself is bounded by $\log s$.

Step 2. Parallel evaluation of rational formulas:

Let $\Phi(x_1, x_2, \ldots, x_n)$ be a rational formula of size s in the noncommuting variables x_i. Given a matrix tuple $(p_1, p_2, \ldots, p_n)$ of $\ell \times \ell$ matrices over $\mathbb{Q}$, our aim is to give an NC algorithm for evaluating $\Phi(p_1, p_2, \ldots, p_n)$ if Φ is defined at this matrix tuple and otherwise detecting that it is undefined.

We first note that if the formula Φ has depth $O(\log s)$ then it is amenable to parallel evaluation on the input $(p_1, p_2, \ldots, p_n)$ using the formula structure. Matrix multiplication, addition, and matrix inversion are all in NC^2 [Csa76, Ber84]. Hence evaluation of $\Phi(p_1, p_2, \ldots, p_n)$ is in NC^3 in this case.

In general, the formula Φ may have depth $O(s)$. Hrubes and Wigderson, in [HW15, Proposition 4.1], have shown that any noncommutative rational formula of size s can be transformed into an equivalent rational formula of depth $O(\log s)$ (and hence size $s^{O(1)}$). Their transformation, though algorithmic, uses RIT for rational formulas as a subroutine.[6] Now, as RIT for rational formulas is in deterministic polynomial time [IQS18, GGdOW16], it follows that rational formula depth-reduction is in deterministic polynomial time.

However, our aim is only *parallel evaluation* of a rational formula of size s (without any depth constraint) at a matrix tuple $(p_1, p_2, \ldots, p_n)$. Examining the Hrubes and Wigderson, in [HW15, Proposition 4.1] depth reduction of rational formulas, we note that that their algorithm is essentially based on Brent's classical result [Bre74] on depth reduction for commutative arithmetic formulas.[7] However, there are some new aspects. It turns out that if Ψ is a rational formula in noncommuting variables $z, y_1, y_2, \ldots, y_m$ with z occurring exactly once as input then Ψ has a z-normal form expression:

$$\Psi = (Az + B)(Cz + D)^{-1}$$

where A, B, C, D are small rational formulas with no occurrence of z. They exploit this structure in their divide and conquer algorithm for constructing the equivalent formula $\widehat{\Phi}$ in polynomial time.

Using this structure, there is a simple NC algorithm [Jog23] for the evaluation of $\Phi(p_1, p_2, \ldots, p_n)$, given as input a rational formula Φ and a matrix tuple $(p_1, p_2, \ldots, p_n)$ which we briefly describe below.

1. The input rational formula Φ is a binary tree. Let r denote its root. By standard NC computation we can find a gate v in Φ such that the size of the subformula Φ_v rooted at v has size between $s/3$ and $2s/3$.

2. We compute the path $P = (v, v_1, v_2, \ldots, v_t = r)$ of all gates from v to r in Φ. Then we find all the gates u in Φ such that $u \notin P$ and u is input to some gate $v_i \in P$. Notice that for $v_i \in P$ such that $v_i \in \{+, \times\}$ has exactly one such input u. The inversion gates v_i are unary.

3. Recursively evaluate Φ_v and each such Φ_u on the input $(p_1, p_2, \ldots, p_n)$.

4. We are left with the problem of evaluation a *skew* rational formula Φ' consisting of the gates along path P with the already computed Φ_u and Φ_v as inputs. Using z-normal forms [HW15] (defined above) it is easy to obtain a simple divide-and-conquer parallel algorithm for evaluating skew rational formulas (in particular, evaluating Φ') (as described in the claim below).

More precisely, a *skew rational formula* Φ' is a rational formula whose internal gates $v_1, v_2, \ldots, v_t = r$, where r is the output gate, form a path. Every gate $v_i, i > 1$ that is a $+$ or

[6] As observed in [AJ24], depth-reduction of noncommutative rational formulas is NC-reducible to RIT for rational formulas.

[7] We note that Brent actually describes a detailed parallel algorithm for carrying out the depth reduction.

$\times$ has one input that is a formula input. That is, in the binary tree underlying the skew formula if $v_i, i > 1$ has two children then exactly one of them is a leaf node. If the gate v_1 is an inversion then it has a single input (which is leaf). Otherwise, $v_1 \in \{+, \times\}$ will have two leaves as children.

Claim 50. *Let Φ' be a skew rational formula with internal $v_1, v_2, \ldots, v_t = r$, where r is the output gate. Then there is an NC^3 algorithm to evaluate Φ' on a given matrix assignment to the input variables.*

Proof of Claim. Every non-inversion gate v_i for $i > 1$ in Φ' is either a + or a $\times$ and has one input that is a matrix (over F).

We will solve a more general problem where the last gate v_1 in Φ' is variable z. Then, using the z-normal form of [HW15], for a matrix assignment to all input variables the formula Φ' evaluates to $(Az + B)(Cz + D)^{-1}$, where A, B, C, D are matrices over F. We will show that this z-normal form is computable in $O(\log t)$ many parallel rounds of matrix operations (which are matrix multplications and inversions which can be performed in NC^2 using standard algorithms).

We split Φ' into two skew formulas: formula Φ'_1 defined by the path $z = v_1, v_2, \ldots, v_{t/2}$ and Φ'_2 defined by the path $v_{t/2} = z_1, \ldots, v_t = r$. Recursively suppose now that we have computed for the given matrix input $\Phi'_1 = (A_1 z_1 + B_1)(C_1 z_1 + D_1)^{-1}$ and $\Phi'_2 = (A_2 z + B_2)(C_2 z + D_1)^{-1}$, where $A_1, B_1, C_1, D_1, A_2, B_2, C_2, D_2$ are matrices over F. Then, exactly as shown in the proof of [HW15, Proposition 4.1], we can write $\Phi' = (A_1 h_1 + B_1 h_2)(C_1 h_1 + D_1 h_2)^{-1}$, where $h_1 = A_2 z + B_2$ and $h_2 = C_2 z + D_2$. This means we can recover the z-normal form for Φ' in NC^2 with some matrix operations. This sets up a divide-and-conquer algorithm on the path length t which means we have an NC algorithm that has at most $\log t \leqslant \log s$ parallel rounds of NC^2 computations giving us an NC^3 upper bound. ■

We analyze the running time of the above parallel algorithm. Let $T(s)$ bound the number of rounds of parallel matrix multiplications, additions, inversions required to evaluate a size s rational formula. Then, the above algorithm yields the bound

$$T(s) \leqslant T(2s/3) + O(\log s),$$

which implies $T(s) \leqslant O(\log^2 s)$. Notice that the term $T(2s/3)$ bounds the running time for recursive evaluation of Φ_v and each Φ_u, all in parallel, because each of these subformulas have size at most $2s/3$. The term $O(\log s)$ is the bound[8] for the separate parallel algorithm, mentioned above, for evaluating the *skew* rational formula Φ'.[9]

As each matrix operation can be performed in NC^2, it follows that rational formula evaluation is in NC^4. We obtain the following.

Lemma 51. *There is an NC^4 algorithm for evaluating a noncommutative rational formula Φ on a given matrix input $(p_1, p_2, \ldots, p_n)$.*

Remark 52. Notice that the NC algorithm described above, with minor changes, will yield an $O(\log^2 s)$ depth, poly(s) size rational formula equivalent to Φ.

The size of our final hitting set is $(ns)^{O(\theta^5 \log s)} = (ns)^{O(\log^6 s)}$ and the dimension of the matrices in the hitting set is $s^{O(\theta)} = s^{O(\log s)}$. Using the rational formula evaluation procedure, on each such matrix tuple, it can be evaluated within quasi-NC. This is in parallel repeated for $(ns)^{O(\log^6 s)}$ points in the hitting set. This completes the proof of Corollary 4.

[8]Here again we mean $O(\log s)$ rounds of parallel matrix inversions, multiplications, and additions.

[9]A skew rational formula is a rational formula in which at least one input to each binary gate is a formula input.

6 Conclusion

In this paper, we nearly settle the black-box complexity of the RIT problem. However, designing a black-box algorithm for the NSingular problem remains wide open. The connection of this problem to the parallel algorithm for bipartite matching [FGT16] is already discussed in Section 1.

We believe that the techniques introduced in this paper might be useful in designing efficient hitting sets for the NSingular problem.

Recall that, the result of Derksen and Makam [DM17] implies that for a nonzero rational formula of size s, there is a $2s \times 2s$ matrix tuple such that the evaluation is nonzero. Therefore, the quasipolynomial bound on the dimension of the hitting set point obtained in Theorem 2 is far from the optimal bound known. An interesting open problem is to construct a hitting set where the dimension is polynomially bounded in the size of the formula.

Another interesting problem is to show that, in the white-box setting RIT can be solved in NC. Recall that, the identity testing of noncommutative formulas can be performed in NC in the white-box setting [AJS09, For14].

Acknowledgments. We are grateful to the referees for their several thoughtful comments and suggestions which have helped us organize the paper and results. We thank Pushkar S. Joglekar for discussions concerning depth reduction of rational formulas in parallel, and for sharing his observation [Jog23]. We also thank the reviewers of the conference version of this work for their helpful comments.

References

[ACG+23] Vikraman Arvind, Abhranil Chatterjee, Utsab Ghosal, Partha Mukhopadhyay, and C. Ramya. On identity testing and noncommutative rank computation over the free skew field. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPIcs*, pages 6:1–6:23, MIT, Cambridge, Massachusetts, USA, 2023. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

[ACM22] Vikraman Arvind, Abhranil Chatterjee, and Partha Mukhopadhyay. Black-box identity testing of noncommutative rational formulas of inversion height two in deterministic quasipolynomial time. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2022*, volume 245 of *LIPIcs*, pages 23:1–23:22, online, 2022. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

[AJ24] Vikraman Arvind and Pushkar S. Joglekar. A multivariate to bivariate reduction for noncommutative rank and related results. In *51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, July 8-12, 2024, Tallinn, Estonia*, volume 297 of *LIPIcs*, pages 14:1–14:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

[AJS09] Vikraman Arvind, Pushkar S. Joglekar, and Srikanth Srinivasan. Arithmetic circuits and the Hadamard product of polynomials. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2009, December 15-17, 2009, IIT Kanpur, India*, pages 25–36, IIT Kanpur, India, 2009. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

[AL50] A. S. Amitsur and J. Levitzki. Minimal identities for algebras. *Proceedings of the American Mathematical Society*, 1(4):449–463, 1950.

[Ami66] S.A Amitsur. Rational identities and applications to algebra and geometry. *Journal of Algebra*, 3(3):304 – 359, 1966.

[BBJP19] Vishwas Bhargava, Markus Bläser, Gorav Jindal, and Anurag Pandey. A deterministic PTAS for the algebraic rank of bounded degree polynomials. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 647–661, San Diego, California, USA, 2019. SIAM.

[Ber76] George M Bergman. Rational relations and rational identities in division rings. *Journal of Algebra*, 43(1):252–266, 1976.

[Ber84] Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Inf. Process. Lett.*, 18(3):147–150, 1984.

[BFG$^+$19] Peter Bürgisser, Cole Franks, Ankit Garg, Rafael Mendes de Oliveira, Michael Walter, and Avi Wigderson. Towards a theory of non-commutative optimization: Geodesic 1st and 2nd order methods for moment maps and polytopes. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 845–861, Baltimore, Maryland, USA, 2019. IEEE Computer Society.

[Bre74] Richard P. Brent. The parallel evaluation of general arithmetic expressions. *J. ACM*, 21(2):201–206, 1974.

[BW05] Andrej Bogdanov and Hoeteck Wee. More on noncommutative polynomial identity testing. In *20th Annual IEEE Conference on Computational Complexity (CCC 2005), 11-15 June 2005, San Jose, CA, USA*, pages 92–99, San Jose, CA, USA, 2005. IEEE Computer Society.

[CM23] Abhranil Chatterjee and Partha Mukhopadhyay. The noncommutative Edmonds' problem re-visited, 2023.

[Coh71] P. M. Cohn. The Embedding of Firs in Skew Fields. *Proceedings of the London Mathematical Society*, s3-23(2):193–213, 10 1971.

[Coh95] P. M. Cohn. *Skew Fields: Theory of General Division Rings*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1995.

[Csa76] L. Csanky. Fast parallel matrix inversion algorithms. *SIAM J. Comput.*, 5(4):618–623, 1976.

[DK21] Manfred Droste and Dietrich Kuske. Weighted automata. In Jean-Éric Pin, editor, *Handbook of Automata Theory*, pages 113–150. European Mathematical Society Publishing House, Zürich, Switzerland, 2021.

[DL78] Richard A. Demillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193 – 195, 1978.

[DM17] Harm Derksen and Visu Makam. Polynomial degree bounds for matrix semi-invariants. *Advances in Mathematics*, 310:44–63, 2017.

[DM20] Harm Derksen and Visu Makam. Algorithms for orbit closure separation for invariants and semi-invariants of matrices. *Algebra & Number Theory*, 14(10):2791–2813, 2020.

[Eil74] Samuel Eilenberg. *Automata, Languages, and Machines (Vol A)*. Pure and Applied Mathematics. Academic Press, 1974.

[FGT16] Stephen Fenner, Rohit Gurjar, and Thomas Thierauf. Bipartite perfect matching is in quasi-nc. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, page 754–763, New York, NY, USA, 2016. Association for Computing Machinery.

[For14] Michael Andrew Forbes. *Polynomial Identity Testing of Read-Once Oblivious Algebraic Branching Programs*. PhD thesis, MIT, 2014.

[FR04] Marc Fortin and Christophe Reutenauer. Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank. *Séminaire Lotharingien de Combinatoire [electronic only]*, 52, 01 2004.

[FS13] Michael A. Forbes and Amir Shpilka. Quasipolynomial-time identity testing of non-commutative and read-once oblivious algebraic branching programs. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 243–252, Berkeley, CA, USA, 2013. IEEE Computer Society.

[GGdOW16] Ankit Garg, Leonid Gurvits, Rafael Mendes de Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing. In Irit Dinur, editor, *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 109–117, Hyatt Regency, New Brunswick, New Jersey,USA, 2016. IEEE Computer Society.

[GGdOW20] Ankit Garg, Leonid Gurvits, Rafael Mendes de Oliveira, and Avi Wigderson. Operator scaling: Theory and applications. *Found. Comput. Math.*, 20(2):223–290, 2020.

[GOR24] Rohit Gurjar, Taihei Oki, and Roshan Raj. Fractional linear matroid matching is in quasi-NC. *arXiv preprint arXiv:2402.18276*, 2024.

[GR08] Ariel Gabizon and Ran Raz. Deterministic extractors for affine sources over large fields. *Comb.*, 28(4):415–440, 2008.

[GT20] Rohit Gurjar and Thomas Thierauf. Linear matroid intersection is in quasi-nc. *Comput. Complex.*, 29(2):9, 2020.

[HH21] Masaki Hamada and Hiroshi Hirai. Computing the nc-rank via discrete convex optimization on CAT(0) spaces. *SIAM J. Appl. Algebra Geom.*, 5(3):455–478, 2021.

[Hua49] Loo-Keng Hua. Some properties of a sfield. *Proceedings of the National Academy of Sciences of the United States of America*, 35(9):533–537, 1949.

[HW15] Pavel Hrubeš and Avi Wigderson. Non-commutative arithmetic circuits with division. *Theory of Computing*, 11(14):357–393, 2015.

[IQS18] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Constructive non-commutative rank computation is in deterministic polynomial time. *Computational Complexity*, 27(4):561–593, 2018.

[Jog23] Pushkar S. Joglekar. Personal communication, 2023.

[KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Comput. Complex.*, 13(1-2):1–46, 2004.

[Lam01] T.Y. Lam. A first course in noncommutative rings (second edition), 2001.

[MW19] Visu Makam and Avi Wigderson. Singular tuples of matrices is not a null cone (and, the symmetries of algebraic varieties). *Electron. Colloquium Comput. Complex.*, TR19-114, 2019.

[Pie82] Richard S. Pierce. *Associative Algebras*. Springer-Verlag, 1982.

[Row80] Louis Halle Rowen. *Polynomial identities in ring theory*. Pure and Applied Mathematics. Academic Press, 1980.

[Sch80] Jacob T. Schwartz. Fast probabilistic algorithm for verification of polynomial identities. *J. ACM.*, 27(4):701–717, 1980.

[Str73] Volker Strassen. Vermeidung von divisionen. *Journal für die reine und angewandte Mathematik*, 264:184–202, 1973.

[Vol18] Jurij Volčič. Matrix coefficient realization theory of noncommutative rational functions. *Journal of Algebra*, 499:397–437, 04 2018.

[Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation, EUROSAM '79, An International Symposiumon Symbolic and Algebraic Computation, Marseille, France, June 1979, Proceedings*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226, Marseille, France, 1979. Springer.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il