Discovering the roots: Uniform closure results for algebraic classes under factoring[*]

Pranjal Dutta [†]   Nitin Saxena [‡]   Amit Sinhababu [§]

Abstract

Newton iteration (NI) is an almost 350 years old recursive formula that approximates a simple root of a polynomial quite rapidly. We generalize it to a matrix recurrence (allRootsNI) that approximates *all* the roots simultaneously. In this form, the process yields a better circuit complexity in the case when the number of roots r is small but the multiplicities are exponentially large. Our method sets up a linear system in r unknowns and iteratively builds the roots as formal power series. For an algebraic circuit $f(x_1, \ldots, x_n)$ of size s we prove that *each* factor has size at most a polynomial in: s and the degree of the squarefree part of f. Consequently, if f_1 is a $2^{\Omega(n)}$-hard polynomial then any nonzero multiple $\prod_i f_i^{e_i}$ is equally hard for *arbitrary* positive e_i's, assuming that $\sum_i \deg(f_i)$ is at most $2^{O(n)}$.

It is an old open question whether the class of poly(n)-sized formulas (resp. algebraic branching programs) is closed under factoring. We show that given a polynomial f of degree $n^{O(1)}$ and formula (resp. ABP) size $n^{O(\log n)}$ we can find a similar size formula (resp. ABP) factor in randomized poly$(n^{\log n})$-time. Consequently, if determinant requires $n^{\Omega(\log n)}$ size formula, then the same can be said about any of its nonzero multiples.

As part of our proofs, we identify a new property of multivariate polynomial factorization. We show that under a random linear transformation τ, $f(\tau \overline{x})$ *completely* factors via power series roots. Moreover, the factorization adapts well to circuit complexity analysis. This with allRootsNI are the techniques that help us make progress towards the old open problems; supplementing the large body of classical results and concepts in algebraic circuit factorization (eg. Zassenhaus, J.NT 1969; Kaltofen, STOC 1985-7 & Bürgisser, FOCS 2001).

2012 ACM CCS concept: Theory of computation– Algebraic complexity theory, Problems, reductions and completeness; Computing methodologies– Algebraic algorithms, Hybrid symbolic-numeric methods; Mathematics of computing– Combinatoric problems.
Keywords: circuit factoring, formula, ABP, randomized, hard, VF, VBP, VP, VNP, quasipoly.

Contents

1 Introduction

Algebraic circuits provide a way, alternate to Turing machines, to study computation. Here, the complexity classes contain (multivariate) polynomial families instead of languages. It is a natural question whether an algebraic complexity class is closed under factors. This is also a useful, and hence, a very well studied question both from the point of view of practice and theory. We study the following two questions related to multivariate polynomial factorization: **(1)** Let $\{f_n(x_1, \ldots, x_n)\}_n$ be a polynomial family in an algebraic complexity class $\mathcal{C}$ (egs. VP, VF, VBP, VNP or $\overline{\text{VP}}$ etc.). Let g_n be an arbitrary factor of f_n. Can we say that $\{g_n\}_n \in \mathcal{C}$? Equivalently, is the class $\mathcal{C}$ *closed under factoring*? **(2)** Can we design an *efficient*, i.e. randomized poly(n)-time, algorithm to output the factor g_n with a representation in $\mathcal{C}$? (*Uniformity*)

Different classes give rise to new challenges for the closure questions. Before discussing further, we give a brief overview of the algebraic complexity classes relevant for our paper. For more details, see [Mah14, SY10, BCS13].

Algebraic circuit is a natural model to represent a polynomial compactly. An *algebraic circuit* has the structure of a layered directed acyclic graph. It has leaf nodes labelled as input variables $x_1, \ldots, x_n$ and constants from the underlying field $\mathbb{F}$. All the other nodes are labelled as addition and multiplication gates. It has a root node that outputs the polynomial computed by the circuit. Some of the complexity parameters of a circuit are *size* (number of edges and nodes), *depth* (number of layers), syntactic *degree* (the maximum degree polynomial computed by any node), *fan-in* (maximum number of inputs to a node) and *fan-out*. An *algebraic formula* is a circuit whose underlying graph is a *directed tree*. In a formula, the fan-out of the nodes is at most one, i.e. 'reuse' of intermediate computation is not allowed.

The class VP (resp. VF) contains the families of n-variate polynomials of degree $n^{O(1)}$ over $\mathbb{F}$, computed by $n^{O(1)}$-sized circuits (resp. formulas). The class VF is sometimes denoted as VP_e, for it collects 'expressions' which is another name for formulas. Similarly, one can define VQP (resp. VQF) which contains the families of n-variate polynomials of degree $n^{O(1)}$ over $\mathbb{F}$, computed by $2^{\mathrm{poly}(\log n)}$-sized circuits (resp. formulas). If we relax the condition on the degree in the definition of VP, by allowing the degree to be possibly exponential, then we define the class VP_{nb}. Such circuits can compute constants of exponential bit-size (unlike VP).

Algebraic branching program (ABP) is another model for computing polynomials which we define in Sec.A. The class VBP contains the families of polynomials computed by $n^{O(1)}$-sized ABPs. We have the easy containments: $\mathrm{VF} \subseteq \mathrm{VBP} \subseteq \mathrm{VP} \subseteq \mathrm{VQP} = \mathrm{VQF}$ [BOC92, VSBR83].

Finally, we give an overview of the class VNP, which can be seen as a non-deterministic analog of the class VP. A family of polynomials $\{f_n\}_n$ over $\mathbb{F}$ is in VNP if there exist polynomials $t(n), s(n)$ and a family $\{g_n\}_n$ in VP such that for every n, $f_n(\overline{x}) = \sum_{w \in \{0,1\}^{t(n)}} g_n(\overline{x}, w_1, \ldots, w_{t(n)})$. Here, *witness* size is $t(n)$ and *verifier* circuit g_n has size $s(n)$. VP is contained in VNP and it is believed that this containment is strict (Valiant's Hypothesis [Val79]).

Newton iteration based numerical methods are very popular in engineering [OR00, GMS$^+$86, BSR$^+$05]. This work introduces a new process to approximate all the roots of a circuit assuming that they are few and their multiplicites are known. This is based on a matrix recurrence, which in turn is derived from a new identity (Claim 6). Based on the process (called *allRootsNI* in Section 1.3) we get several consequences in very high-degree circuit factoring (eg. Theorem 1):

Every factor of a given circuit C has size polynomial in: size(C) and the degree of the squarefree part of C (note that it may be exponentially smaller than deg(C)).

and in factoring other poly-degree algebraic models (eg. Theorems 3, 14 & 15):

Every factor, of a degree-d polynomial with VF (respectively VBP, VNP) complexity s, has VF (respectively VBP, VNP) complexity poly$(s, d^{\log d})$. The latter is poly(s) if degree $d = 2^{O(\sqrt{\log s})}$.

Now, we briefly discuss the state of the art on the closure questions for various algebraic complexity classes. (Those interested in numerical analysis could first read Section 1.2.1.) To cover more depth and breadth, see [Kal90, Kal92, FS15].

1.1 Previously known closure results

Famously, Kaltofen [Kal85, Kal86, Kal87, Kal89] showed that VP is *uniformly* closed under factoring, i.e. for a given d degree n variate polynomial f of circuit size s, there exists a randomized poly(snd)-time algorithm that outputs its factor as a circuit whose size is bounded by poly(snd). This fundamental result has several applications such as 'hardness versus randomness' in algebraic complexity [KI03, AV08, DSY09, AGS18], derandomization of Noether Normalization Lemma [Mul17], in the problem of circuit reconstruction [KS09, Sin16], and polynomial equivalence testing [Kay11]. In general, multivariate polynomial factoring has several applications including decoding of Reed-Solomon, Reed-Muller codes [GS98, Sud97], integer factoring [LLMP90], primary decomposition of polynomial ideals [GTZ88] and algebra isomorphism [KS06, IKRS12].

It is natural to ask whether Kaltofen's VP factoring result can be extended to VP_{nb} which allows degree of the polynomials to be exponentially high. It is known that *not every* factor of a high degree polynomial has a small sized circuit. For example, the polynomial $x^{2^s} - 1$ can be computed in size s, but it has factors over $\mathbb{C}$ that require circuit size $\Omega\left(2^{s/2}/\sqrt{s}\right)$ [LS78, Sch77]. It is conjectured [Bür13, Conj.8.3] that *low* degree factors of high degree small-sized circuits have *small* circuits. Partial results towards it are known. It was shown in [Kal87] that if polynomial f given by a circuit of size s factors as $g^e h$, where g and h are coprime, then g can be computed by a circuit of size poly$(e, \deg(g), s)$. The question left open is to remove the dependency on e. In the special case where $f = g^e$, it was established that g has circuit size poly$(\deg(g), \mathrm{size}(f))$.

On the other hand, several algorithmic problems are NP-hard, eg. computing the degree of the squarefree part, gcd, or lcm; even in the case of supersparse univariate polynomials [Pla77b].

Now, we discuss the closure results for classes more restrictive than VP (such as VF, VBP etc.). Unfortunately, Kaltofen's technique [Kal89] for VF will give a superpolynomial-sized factor formula; as it heavily *reuses* intermediate computations while working with linear algebra and Euclid gcd. The same holds for the class VBP. In contrast, extending the idea of [DSY09], Oliveira [Oli16] showed that an n-variate polynomial with *bounded individual degree* and computed by a formula of size s, has factors of formula size poly(n, s). Furthermore, it was established that for a given n-variate individual-degree-r polynomial, computed by a circuit (resp. formula) of size s and depth Δ, there exists a poly(n^r, s)-time randomized algorithm that outputs any factor of f computed by a circuit (resp. formula) of depth $\Delta + 5$ and size poly(n^r, s). We are not aware of any work specifically on VBP factoring, except a special case in [KK08]—it dealt with the elimination of a *single* division gate from skew circuits (also see Section A.1 & Lemma 20)—and another special case result in [Jan11] that was weakened later owing to proof errors.

Going beyond VP we can ask about the closure of VNP. Bürgisser conjectured [Bür13, Conj.2.1] that VNP is closed under factoring. Kaltofen's technique [Kal89] for factoring VP circuits does not yield the closure of VNP. After our paper was sent to review, Chou, Kumar and Solomon [CKS18] have confirmed that VNP is indeed closed under factors.

Looking at the Border. Recently, *approximative* algebraic complexity classes like $\overline{\text{VP}}$ [GMQ16] have become objects of interest, especially in the context of the geometric complexity program [Mul12a, Mul12b, Gro15], but also in the framework that yields the fastest matrix multiplication algorithms ([LG14] surveys the recent developments). Interestingly, [Mul17, Thm.4.9] shows that the following three fundamental concepts are tightly related mainly due to circuit factoring results: **1)** efficient blackbox polynomial identity testing (PIT) for $\overline{\text{VP}}$, **2)** strong lower bounds against $\overline{\text{VP}}$, and **3)** efficiently computing an 'explicit system of parameters' for the invariant ring of an explicit variety with a given group action.

The related algorithmic questions, including factorization of $\overline{\text{VP}}$ circuits, have been recently put in PSPACE [FS17]. This is the best uniform derandomization result known currently.

$\overline{\text{VP}}$ contains families of polynomials of degree poly(n) that can be approximated (infinitesimally closely) by poly(n)-sized circuits. Bürgisser [Bür04, Bür01a, Bür01b] discusses approximative complexity of factors, proving that low degree factors of high degree circuits have small approximative complexity. In particular, $\overline{\text{VP}}$ is closed under factoring [Bür01a, Thm.4.1]. Like the standard versions, closure of $\overline{\text{VF}}$ resp. $\overline{\text{VBP}}$ is an open question. Recently, it has been shown that $\overline{\text{VF}} = \text{width-2-}\overline{\text{VBP}}$ [BIZ17] while classically it is false [AW11]. The new methods that we present extend nicely to approximative classes because of their analytic nature (Theorem 14).

We conclude by stating a few reasons why closure results under factoring are interesting and non-trivial. First, there are classes that are *not* closed under factors. For example, the class of sparse polynomials; as a factor's sparsity may blowup super-polynomially [vzGK85]. Closure under factoring indicates the robustness of an algebraic complexity class, as, it proves that all nonzero multiples of a *hard* polynomial remain hard. For this reason, closure results are also important for proving lower bounds on the power of some algebraic proof systems [FSTW16].

Finally, factoring is the key reason why PIT, for VP, can be reduced to very special cases, and gets tightly related to circuit lower bound questions (like VP$\neq$VNP?). See [KI03, Thm.4.1] for whitebox PIT connection and [AGS18, KST18] for blackbox PIT. One of the central reasons is: Suppose a polynomial $f(\overline{y})$ is such that for a nonzero size-s circuit C, $C(f(\overline{y})) = 0$. Then, using factoring results for low degree C, one deduces that f also has circuit size poly(s). This gives us the connection: *If we picked a "hard" polynomial f then $f(\overline{y})$ would be a hitting-set generator (hsg) for C* [KI03, Thm.7.7]. Our work is strongly motivated by the open question of proving such a result for size-s circuits C that have high degree (i.e. $s^{\omega(1)}$). Our first factoring

result (Theorem 1) implies such a 'hardness to hitting-set' connection for arbitrarily high degree circuits C assuming that: the squarefree part C_{sqfree} of C has low degree. In such a case we only have to find a hitting-set for C_{sqfree} which, as our result proves, has low algebraic circuit complexity.

1.2 Our results

Before stating the results, we describe some of the assumptions and notations used throughout the paper. Set $[n]$ refers to $\{1, 2, \ldots, n\}$. Logarithms are wrt base 2. For a polynomial f, $\text{size}(f)$ refers to the smallest size of circuits computing f; it is the *algebraic circuit complexity* of f.

Field. We denote the underlying field as $\mathbb{F}$ and assume that it is of characteristic 0 and algebraically closed. For eg. complex $\mathbb{C}$, algebraic numbers $\overline{\mathbb{Q}}$ or algebraic p-adics $\overline{\mathbb{Q}}_p$. All the results partially hold for other fields (such as $\mathbb{R}, \mathbb{Q}, \mathbb{Q}_p$ or finite fields of characteristic >degree of the input polynomial). For a brief discussion on this issue, see Section 5.

Ideal. We denote the variables $(x_1, \ldots, x_n)$ as $\overline{x}$. The *ideal* $I := \langle \overline{x} \rangle$ of the polynomial ring will be of special interest, and its power ideal I^d, whose generators are all degree d monomials in n variables. Often we will reduce the polynomial ring modulo I^d (inspired from *Taylor series of an analytic function around* $\overline{0}$ [Tay15]).

Radical. For a polynomial $f = \prod_i f_i^{e_i}$, with f_i's coprime irreducible nonconstant polynomials and multiplicity $e_i > 0$, we define the squarefree part as the *radical* $\text{rad}(f) := \prod_i f_i$.

What can we say about these f_i's if f has a circuit of size s? Our main result gives a good circuit size bound when $\text{rad}(f)$ has small degree. A more general formulation (with u_0) is:

Theorem 1. *If $f = u_0 u_1$ is a nonzero product in the polynomial ring $\mathbb{F}[\overline{x}]$, with $\text{size}(f) + \text{size}(u_0) \leq s$, then every factor of u_1 has a circuit of size $\text{poly}(s + \deg(\text{rad}(u_1)))$.*

Note that Kaltofen's proof technique in the VP factoring paper [Kal89] does not extend to the *exponential* degree regime (even when degree of $\text{rad}(f)$ is small) because it requires solving equations with $\deg_{x_i}(f)$ many unknowns for some x_i, where $\deg_{x_i}(f)$ denotes *individual degree* of x_i in f, which can be very high. Also, basic operations like 'determining the coefficient of a univariate monomial' become #P-hard in the exponential-degree regime [Val82]. The proof technique in Kaltofen's single factor Hensel lifting paper [Kal87, Thm.2] works only in the perfect-power case of $f = g^e$. It can be seen that $\text{rad}(f)$ "almost" equals $f / \gcd(f, \partial_{x_i}(f))$, but the gcd itself can be of exponential-degree and so one cannot hope to use [Kal87, Thm.4] to compute the gcd either. Univariate high-degree gcd computation is NP-hard [Pla77a, Pla77b].

Interestingly, our result when combined with [Kal87, Thm.3] implies that every factor g of f has a circuit of size polynomial in: $\text{size}(f)$, $\deg(g)$ and $\min\{\deg(\text{rad}(f)), \text{size}(\text{rad}(f))\}$. We leave it as an open question whether the latter expression is polynomially related to $\text{size}(f)$.

Theorem 1 shows an interesting way to create *hard* polynomials. In the theorem statement let the size concluded be $(s + \deg(\text{rad}(u_1)))^e$, for some constant e. If one has a polynomial $f_1(x_1, \ldots, x_n)$ that is 2^{cn}-hard, then any nonzero $f := \prod_i f_i^{e_i}$ is also $2^{\Omega(n)}$-hard for *arbitrary* positive e_i's, as long as $\sum_i \deg(f_i) \leq 2^{\frac{cn}{e}-1}$.

1.2.1 A detour into numerical analysis (via arithmetic circuits)

Root approximation of univariate polynomials has been an interesting problem in mathematics & engineering. Interestingly, it has also found applications in various other problems such as computing the largest eigenvalue, and checking whether a matrix is approximately PSD [LV16].

We can quantify the same question about 'approximating roots' of a univariate polynomial and analyze the bit-complexity of the root; where the measure is the *bitsize* of the best circuit. For an $f(x) \in \mathbb{R}[x]$ we define $\text{bitsize}(f) = s$ if we can compute $f(x)$ by a circuit using $\{+, -, \times, \div\}$

gates and of overall bitsize s. Note that the degree of f could be exponential (eg. 2^s) and the coefficients may have exponential bit-length (eg. 2^s bits, or values 2^{-2^s} to 2^{2^s}).

For this complexity notion, we can show a surprising fact for the roots.

Claim: *If $a \in (0,1)$ is 2^m-bit approximation of a root of $f(x)$ and $\epsilon \in (0,1)$ lower bounds the gap between a and the other roots of $f(x)$, then $bitsize(a) \leq O(s + m \log(\frac{1}{\epsilon}))$, where $bitsize(f) =: s$.*

Note that the estimate is highly nontrivial, as it is essentially expressing 2^m bits of the root using 'only' bitsize m; so, roots of small circuits are rather special, in contrast to generic strings that are incompressible [VL97]. It is relevant here to recall Shub-Smale's tau conjecture that states that 'small' circuits have 'few' *integral* roots! A proof of tau conjecture would imply P$\neq$NP over $\mathbb{C}$ [BCSS98]. This motivates well the study of roots of circuits.

Our Theorem 1 is an algebraic analogue of the above; there $m \log(\frac{1}{\epsilon})$ gets replaced by the degree of the radical (or, the number of roots). Also, algebraic result is better in the sense that we do not need $\div$ gates.

Proof sketch– Now we give the main steps in the numerical result (the algebraic analogue Thm.1 would be more involved to prove as it is stronger). We use general Newton iteration (or, NI with multiplicity). Observe that, we are interested in an existential statement; so, assume that we know the multiplicity γ of the root a before-hand. Suppose, $f(x) = (x-a)^\gamma g(x)$ where $g(a) \neq 0$. Assume that $g(x) = \prod_{i=1}^r (x - b_i)^{\gamma_i}$.

Hypothesis says that $|a - b_i| \geq \epsilon$, for all i. We will use general NI to approximate 2^m bits of a (in this case after the decimal-point only). We will start with y_0 such that $|y_0 - a| \leq \epsilon \cdot 2^{-3m-1}$, and wlog assume $m \geq s \geq 2$. Thus, to start with, y_0 has a trivial circuit of bitsize $O(m \log(\frac{1}{\epsilon}))$.

Next, we use the general NI formula [DB08, Eqn.6.3.13], i.e.

$$ y_{t+1} \;=\; y_t \;-\; \gamma \cdot \left. \frac{f}{f'} \right|_{y_t} . $$

Finally, we need to show that the process has *quadratic* convergence. Inductively, we want to show that for all $t \geq 0$,

$$ |y_t - a| \;\leq\; \epsilon \cdot 2^{-3m} \cdot 2^{-2^t} . \tag{1} $$

Note that the above inequality implies that y_m is a 2^m-bit approximation of a. Moreover, computing y_{t+1} as a circuit– given the value y_t and circuits for $f(x), f'(x)$ –requires $O(1)$ additional bitsize (as we already have γ, f, f' and $\div$ is allowed in the circuit). So, y_m will have a circuit of bitsize $O(s + m \log(\frac{1}{\epsilon}))$.

We are only left to prove Equation 1. We have,

$$ \frac{f(y_t)}{f'(y_t)} \;=\; \frac{(y_t - a)^\gamma g(y_t)}{(y_t - a)^{\gamma-1} \cdot (\gamma g(y_t) + (y_t - a)g'(y_t))} $$
$$ =\; \frac{(y_t - a)g(y_t)}{\gamma g(y_t) + (y_t - a)g'(y_t)} . $$

Hence,

$$ |y_{t+1} - a| \;=\; \left| y_t - a - \gamma \frac{f(y_t)}{f'(y_t)} \right| $$
$$ =\; \left| (y_t - a)\left(1 - \frac{\gamma g(y_t)}{\gamma g(y_t) + (y_t - a)g'(y_t)} \right) \right| $$
$$ =\; |y_t - a|^2 \cdot \left| \frac{g'(y_t)}{\gamma g(y_t) + (y_t - a)g'(y_t)} \right| $$

$$= |y_t - a|^2 \cdot \left| \gamma \frac{g(y_t)}{g'(y_t)} + (y_t - a) \right|^{-1}$$

$$\leq |y_t - a|^2 \cdot \left(\left| \frac{g(y_t)}{g'(y_t)} \right| - |y_t - a| \right)^{-1}.$$

Observe that by Leibniz rule:

$$\frac{g'(y_t)}{g(y_t)} = \sum_{i=1}^{r} \frac{\gamma_i}{y_t - b_i}.$$

By the induction hypothesis $|y_t - a| \leq \epsilon 2^{-3m} \cdot 2^{-2^t}$; so, we have for all $i \in [r]$:

$$|y_t - b_i| \geq |a - b_i| - |y_t - a| \geq \epsilon \cdot (1 - 2^{-2^t} 2^{-3m}).$$

Since $\gamma_i, r \leq 2^s$ ($\because$ bitsize-s circuit can have degree at most 2^s), we can upper bound as:

$$\left| \frac{g'(y_t)}{g(y_t)} \right| \leq \sum_{i \in [r]} \frac{\gamma_i}{|y_t - b_i|} \leq \frac{2^{2s}}{\epsilon(1 - 2^{-2^t} 2^{-3m})}.$$

This means, $\left| \frac{g(y_t)}{g'(y_t)} \right| \geq 2^{-2s-1} \epsilon$. Consequently,

$$\left| \frac{g(y_t)}{g'(y_t)} \right| - |y_t - a| \geq 2^{-2s-1} \epsilon - \epsilon 2^{-3m} \cdot 2^{-2^t}$$

$$= \epsilon 2^{-3s} \cdot \left(2^{s-1} - 2^{3s-3m} \cdot 2^{-2^t} \right)$$

Hence,

$$|y_{t+1} - a| \leq |y_t - a|^2 \cdot \left(\left| \frac{g(y_t)}{g'(y_t)} \right| - |y_t - a| \right)^{-1}$$

$$\leq 2^{-2^{t+1}} (\epsilon \cdot 2^{-3m})^2 \cdot \left(\epsilon 2^{-3s} \right)^{-1}$$

$$\leq 2^{-2^{t+1}} \cdot \epsilon \cdot 2^{-3m}$$

This finishes the inductive step and we are done. $\qquad \square$

We leave some interesting questions open: Whether bitsize(a) could be improved to poly($s +$ m)? Whether a bound for bitsize(a) could be proved without requiring $\div$ gates?

1.2.2 Back to multivariate algebraic models

In general, for a high degree circuit f, rad(f) can be of high degree (exponential in size of the circuit). Ideally, we would like to show that every degree d factor of f has poly(size(f), d)-size circuit. The next theorem reduces the above question to a special kind of modular division, where the denominator polynomial may *not* be invertible but the quotient is well-defined (eg. x^2/x mod x). All that remains is to somehow eliminate this kind of *non-unit division* operator (which we leave as an open question). Consider 'random' elements $\alpha_i, \beta_i \in_r \mathbb{F}$ and the corresponding random linear map $\tau : x_i \mapsto \alpha_i y + x_i + \beta_i$, $i \in [n]$, where y is a new variable apart from $x_1, \ldots, x_n$.

Theorem 2. *If nonzero $f \in \mathbb{F}[\overline{x}]$ can be computed by a circuit of size s, then any degree d factor of $f(\tau \overline{x})$ is of the form A/B mod $\langle \overline{x} \rangle^{d+1}$ where polynomials A, B have circuits of size poly(sd).*

Note that in Theorem 2, B may be non-invertible in $\mathbb{F}[\overline{x}]/\langle\overline{x}\rangle^{d+1}$ and may have a high degree (eg. 2^s). So, we cannot use the famous trick of Strassen to do division elimination here [Str73].

We prove uniform closure results, under factoring, for the algebraic complexity classes defined below. Let $s : \mathbb{N} \longrightarrow \mathbb{N}$ be a function. Define the class $\mathrm{VF}(s)$ to contain families $\{f_n\}_n$ such that n-variate f_n can be computed by an algebraic formula of size $\mathrm{poly}(s(n))$ and has degree $\mathrm{poly}(n)$. Similarly, $\mathrm{VBP}(s)$ contains families $\{f_n\}_n$ such that f_n can be computed by an ABP of size $\mathrm{poly}(s(n))$ and has degree $\mathrm{poly}(n)$. Finally, $\mathrm{VNP}(s)$ denotes the class of families $\{f_n\}_n$ such that f_n has witness size $\mathrm{poly}(s(n))$, verifier circuit size $\mathrm{poly}(s(n))$, and has degree $\mathrm{poly}(n)$.

Theorem 3. *The classes $VF(n^{\log n})$, $VBP(n^{\log n})$, $VNP(n^{\log n})$ are all closed under factoring.*

Moreover, there exists a randomized $\mathrm{poly}(n^{\log n})$-time algorithm that: for a given $n^{O(\log n)}$ sized formula (resp. ABP) f of $\mathrm{poly}(n)$-degree, outputs $n^{O(\log n)}$ sized formula (resp. ABP) of a nontrivial factor of f (if one exists).

Remark. The "time-complexity" in the algorithmic part makes sense only in certain cases. For example, when $\mathbb{F} \in \{\mathbb{Q}, \mathbb{Q}_p, \mathbb{F}_q\}$, or when one allows computation in the BSS-model [BSS89]. In the former case our algorithm takes $\mathrm{poly}(n^{\log n})$ bit operations (assuming that the characteristic is zero or larger than the degree; see Theorem 15 in Section 5.2).

It is important to note that Theorem 3 does not follow by invoking Kaltofen circuit factoring [Kal89] and VSBR transformation [VSBR83] from circuit to log-depth formula. Formally, if we are given a formula (resp. ABP) of size $n^{O(\log n)}$ and degree $\mathrm{poly}(n)$, then it has factors which can be computed by a circuit of size $n^{O(\log n)}$ and depth $O(\log n)$. If one converts the factor circuit to a formula (resp. ABP), one would get the size upper bound of the factor formula to be a much larger $(n^{O(\log n)})^{\log n} = n^{O(\log^2 n)}$. Moreover, Kaltofen's methods crucially rely on the circuit representation to do linear algebra, division with remainder, and Euclid gcd in an efficient way; a nice overview of the implementation level details to keep in mind is [KSS15, Sec.3].

Our proof methods extend to the approximative versions $\mathcal{C}(n^{\log n})$ for $\mathcal{C} \in \{\overline{\mathrm{VF}}, \overline{\mathrm{VBP}}, \overline{\mathrm{VNP}}\}$ as well (Theorem 14).

As before, Theorem 3 has an interesting lower bound consequence: If f has VF (resp. VBP resp. VNP) complexity $n^{\omega(\log n)}$ then any nonzero fg has similar hardness (for $\deg(g) \leq \mathrm{poly}(n)$).

In fact, the method of Theorem 3 yields a formula factor of size $s^e d^{2\log d}$ for a given degree-d size-s formula (e is a constant). This means— If determinant $\det_n$ requires $n^{a\log n}$ size formula, for $a > 2$, then *any* nonzero degree-$O(n)$ multiple of $\det_n$ requires $n^{\Omega(\log n)}$ size formula.

Similarly, if we conjecture that a VP-complete polynomial f_n (say the homomorphism polynomial in [DMM+14, Thm.19]) has $n^{a\log n}$ ABP complexity, for $a > 4$, then *any* nonzero degree-$O(n)$ multiple of f_n has $n^{\Omega(\log n)}$ ABP complexity.

1.3   Proof techniques

We begin by describing the new techniques that we have developed. Since they also give a new viewpoint on classic properties, they may be of independent interest. The techniques are *analytic* at heart ([KP12] has a good historical perspective). The way they appear in algebra is through the *formal power series* ring $\mathbb{F}[[x_1, \ldots, x_n]]$. The elements of this ring are multivariate formal power series, with degree as precision. So, an element f is written as $f = \sum_{i=0}^{\infty} f^{=i}$, where $f^{=i}$ is the *homogeneous part* of degree i of f. In algebra texts it is also called the *completion* of $\mathbb{F}[x_1, \ldots, x_n]$ wrt the ideal $\langle x_1, \ldots, x_n \rangle$ (see [Kem10, Chap.13]). The *truncation* $f^{\leq d}$, i.e. homogeneous parts up to degree d, can be obtained by reducing modulo the ideal $\langle \overline{x} \rangle^{d+1}$. Here d is seen as the *precision* parameter of the respective approximation of f.

The advantages of the ring $\mathbb{F}[[\overline{x}]]$ are many. They usually emerge because of the *inverse* identity $(1-x_1)^{-1} = \sum_{i\geq 0} x_1^i$, which would not have made sense in $\mathbb{F}[\overline{x}]$ but is available now.

First, we introduce a factorization pattern of a polynomial f, over the power series ring, under a random linear transformation. Next, we discuss how this factorization helps us to bound the size of factors of the original polynomial.

Power series complete split: We are interested in the *complete* factorization pattern of a polynomial $f(x_1, \ldots, x_n)$. We can view f as a univariate polynomial in one variable, say x_n, with coefficients coming from $\mathbb{F}[x_1, \ldots, x_{n-1}]$. It is easy to connect linear factors with the roots: $x_n - g$ is a factor of f iff $f(x_1, \ldots, x_{n-1}, g(x_1, \ldots, x_{n-1})) = 0$.

Of course, one should not expect that a polynomial always has a factor which is linear in one variable. But, if one works with an algebraically closed field, then a univariate polynomial completely splits into linear factors (also see the *fundamental theorem of algebra* [CRS96, §2.5.4]). So, if we go to the algebraic closure of $\mathbb{F}(x_1, \ldots, x_{n-1})$, any multivariate polynomial which is monic in x_n will split into factors all linear in x_n. A representation of the elements of $\overline{\mathbb{F}(x_1, \ldots, x_{n-1})}$ as a finite circuit is impossible (eg. $\sqrt{x_1}$). On the other hand, we show in this work that *all* the roots (wrt a new variable y) are actually elements from $\mathbb{F}[[x_1, \ldots, x_n]]$, after a *random* linear transformation on the variables, $\tau : \overline{x} \mapsto \overline{x} + \overline{\alpha}y + \overline{\beta}$, is applied (Theorem 4). Note– By a random choice $\alpha \in_r \mathbb{F}$ we will mean that choose randomly from a fixed finite set $S \subseteq \mathbb{F}$ of appropriate size (namely $> \deg(f)$). This will be in the spirit of [Sch80].

Our proof of the existence of power series roots is *constructive*, as it also gives an algorithm to find approximation of the roots up to any precision, using formal power series version of the Newton iteration method (see [BCS13, Thm.2.31]). We try to explain the above idea using the following example. Consider $f = (y^2 - x^3) \in \mathbb{F}[x, y]$. Does it have a factor of the form $y - g$ where $g \in \mathbb{F}[[x]]$? The answer is clearly 'no' as $x^{3/2}$ does not have any power series representation in $\mathbb{F}[[x]]$. But, what if we shift x randomly? For example, if we use the shift $y \mapsto y, x \mapsto x + 1$. Then, by Taylor series around 1, we see that $(x+1)^{3/2}$ has a power series expansion, namely $1 + \frac{3}{2}x + \frac{3/2 \times 1/2}{2!}x^2 + \ldots$.

Formally, Theorem 4 shows that under a random $\tau : \overline{x} \mapsto \overline{x} + \overline{\alpha}y + \overline{\beta}$ where $\overline{\alpha}, \overline{\beta} \in_r \mathbb{F}^n$, polynomial f can be factored as $f(\tau\overline{x}) = \prod_{i=1}^{d_0}(y - g_i)^{\gamma_i}$, where $g_i \in \mathbb{F}[[\overline{x}]]$ with the constant terms $g_i(\overline{0})$ being distinct, $d_0 := \deg(\text{rad}(f))$ and $\gamma_i > 0$.

Reducing factoring to computing power series root approximations: Using the split Theorem 4, we show that multivariate polynomial factoring reduces to power series root finding up to certain precision. Following the above notation f splits as $f(\tau\overline{x}) = \prod_{i=1}^{d_0}(y - g_i)^{\gamma_i}$. For all $t \geq 0$, it is easy to see that $f(\tau\overline{x}) \equiv \prod_{i=1}^{d_0}(y - g_i^{\leq t})^{\gamma_i} \mod I^{t+1}$, where $I := \langle x_1, \ldots, x_n \rangle$. Note that there is a one-one correspondence, induced by τ, between the polynomial factors of f and $f(\tau\overline{x})$ ($\because \tau$ is invertible and f is y-free). We remark that the leading-coefficient of $f(\tau\overline{x})$ wrt y is a nonzero element in $\mathbb{F}$; so, we call it *monic* (Lemma 28). Next, we show case by case how to find a *polynomial* factor of $f(\tau\overline{x})$ from the approximate power series roots.

Case 1- Computing a linear factor of the form $y - g(\overline{x})$: If the degree of the input polynomial is d, all the non-trivial factors have degree $\leq (d-1)$. So, if we compute the approximations of all the power series roots (wrt y) up to precision of degree $t = d - 1$, then we can recover all the factors of the form $y - g(x_1, \ldots, x_n)$. Technically, this is supported by the uniqueness of power series factorization (Proposition 1).

Case 2- Computing a monic non-linear factor: Assume that a factor g of total degree t is of the form $y^k + c_{k-1}y^{k-1} + \ldots + c_1 y + c_0$, where for all i, $c_i \in \mathbb{F}[\overline{x}]$. Now this factor g also splits into linear (in y) factors above $\mathbb{F}[[\overline{x}]]$ and obviously these linear factors are also linear factors of the original polynomial $f(\tau\overline{x})$. So we have to take the right combination of some k power series roots, with their approximations (up to the degree t wrt $\overline{x}$), and take the product mod I^{t+1}. Note that if we only want to give an existential proof of the size bound of the factors, we need not find the combination of the power series roots forming a factor algorithmically. Doing it

through brute-force search takes exponential time ($\binom{d}{k}$ choices). Interestingly, using a classical (linear algebra) idea due to Kaltofen, it can be done in randomized polynomial time. We will spell out the ideas later, while discussing the algorithm part of Theorem 3.

Once we are convinced that looking at approximate (power series) roots is enough, we need to investigate methods to compute them. We will now sketch two methods. The first one approximates all the roots *simultaneously* up to precision δ. The next ones approximate the roots *one at a time.* In the latter, multiplicity of the root plays an important role.

Recursive root finding via matrices (allRootsNI): We *simultaneously* find the approximations of all the power series roots g_i of $f(\tau\overline{x})$. At each recursive step we get a better precision wrt degree. We show that knowing approximations $g_i^{<\delta}$, of g_i up to degree $\delta - 1$, is enough to (simultaneously for all i) calculate approximations of g_i up to degree δ. This new technique, of finding approximations of the power series roots, is at the core of Theorem 1.

First, let us introduce a nice identity. From now on we assume $f(\overline{x}, y) = \prod_i (y - g_i)^{\gamma_i}$ (i.e. relabel $f(\tau\overline{x})$). By applying the derivative operator ∂_y, we get a classic identity (which we call *logarithmic derivative identity*): $(\partial_y f)/f = \sum_i \gamma_i/(y - g_i)$. Reduce the above identity modulo $I^{\delta+1}$ and define $\mu_i := g_i(\overline{0}) \equiv g_i \mod I$. This gives us (see Claim 6):

$$\frac{\partial_y f}{f} = \sum_{i=1}^{d_0} \frac{\gamma_i}{y - g_i} \equiv \sum_{i=1}^{d_0} \frac{\gamma_i}{y - g_i^{<\delta}} + \sum_{i=1}^{d_0} \frac{\gamma_i \cdot g_i^{=\delta}}{(y - \mu_i)^2} \mod I^{\delta+1}.$$

In terms of the d_0 unknowns $g_i^{=\delta}$, the above is a linear equation. (Note- We treat γ_i, μ_i's as known.) As y is a free variable above, we can fix it to d_0 "random" elements c_i in $\mathbb{F}$, $i \in [d_0]$. One would expect these fixings to give a linear system with a unique solution for the unknowns. We can express the system of linear equations succinctly in the following matrix representation: $M \cdot v_\delta = W_\delta \mod I^{\delta+1}$. Here M is a $d_0 \times d_0$ matrix; each entry is denoted by $M(i, j) := \frac{\gamma_i}{(c_i - \mu_j)^2}$. Vector v_δ resp. W_δ is a $d_0 \times 1$ matrix where each entry is denoted by $v_\delta(i) := g_i^{=\delta}$ resp. $W_\delta(i) := \frac{\partial_y f}{f}\big|_{y=c_i} - G_{i,\delta}$, where $G_{i,\delta} := \sum_{k=1}^{d_0} \gamma_k/(c_i - g_k^{<\delta})$. We ensure that $\{c_i, \mu_i \mid i \in [d_0]\}$ are distinct, and show that the determinant of M is non-zero (Lemma 29). So, by knowing approximations up to $\delta - 1$, we can recover δ-th part by solving the above system as $v_\delta = M^{-1}W_\delta \mod I^{\delta+1}$. An important point is that the random c_i's will ensure: all the reciprocals involved in the calculation above do exist mod $I^{\delta+1}$.

Self-correction property: Does the above recursive step need an exact $g_i^{<\delta}$? We show the self correcting behavior of this process of root finding, i.e. in this iterative process there is no need to filter out the "garbage" terms of degree $\geq \delta$ in each step. If one has recovered g_i correct up to degree $\delta - 1$, i.e. say we have calculated $g'_{i,\delta-1} \in \mathbb{F}(\overline{x})$ such that $g'_{i,\delta-1} \equiv g_i^{<\delta} \mod I^\delta$, and say we solve $M\tilde{v}_\delta = \widetilde{W}_\delta$ exactly, where $\widetilde{W}_\delta(i) := \frac{\partial_y f}{f}\big|_{y=c_i} - \widetilde{G}_{i,\delta}$, and $\widetilde{G}_{i,\delta} := \sum_{k=1}^{d_0} \gamma_k/(c_i - g'_{k,\delta-1})$. Still, we can show that $g'_{i,\delta} := g'_{i,\delta-1} + \tilde{v}_\delta(i) \equiv g_i^{\leq\delta} \mod I^{\delta+1}$ (Claim 7). So, we made progress in terms of the precision (wrt degree).

Rapid Newton Iteration with multiplicity: We show that from allRootsNI, we can derive a formula that finds $g_1^{<2^{t+1}}$ using *only* $g_1^{<2^t}$, i.e. the process has quadratic convergence and it does not involve roots other than g_1. Rewrite $\partial_y f/f = \sum_{i=1}^{d_0} \gamma_i/(y - g_i) = (1 + L_1)\gamma_1/(y - g_1)$, where $L_1 := \sum_{1 < i \leq d_0} \frac{\gamma_i}{y - g_i} \cdot \frac{y - g_1}{\gamma_1}$. This implies $f/\partial_y f = (1 + L_1)^{-1} \cdot (y - g_1)/\gamma_1$. Now, if we put $y = y_t := g_1^{<2^t}$, then $y_t - g_i = g_1^{<2^t} - g_i$ is a unit in $\mathbb{F}[[\overline{x}]]$ for $i \neq 1$ ($\because$ it is a nonzero constant mod I). Also, $y_t - g_1 = g_1^{<2^t} - g_1 \equiv 0 \mod I^{2^t}$, implying $L_1|_{y=y_t} \equiv 0 \mod I^{2^t}$. Thus, $(L_1 \cdot (y - g_1))\big|_{y=y_t} \equiv 0 \mod I^{2^{t+1}}$.

Hence, $f/\partial_y f\big|_{y=y_t} = (y_t - g_1)/\gamma_1 \mod I^{2^{t+1}}$.

This shows that, if $f(\overline{x}, y) = (y - g)^e h$, where $h|_{y=g} \neq 0 \mod I$ and $e > 0$, then the power series for g can be approximated by the recurrence:

$$y_{t+1} := y_t - e \cdot \frac{f}{\partial_y f}\bigg|_{y=y_t} \tag{2}$$

where $y_t \equiv g \mod I^{2^t}$. This we call a *generalized Newton Iteration* formula, as it works with any multiplicity $e > 0$.

In fact, when $e = 1$, g is called a *simple root* of f; the above is an alternate proof of the classical Newton Iteration (NI) [New69] that finds a simple root in a recursive way (see Lemma 27). When *all* the roots are simple there are numerical methods to simultaneously approximate them [Dur60, Ker66, Ehr67, Abe73]. However, it is well known that NI *fails* to approximate the roots that repeat (see [Lec02]). In that case either NI is used on the function $f/\partial_y f$ or, though less frequently, the generalized NI is used in numerical methods (see [DB08, Eqn.6.3.13]).

There is a technical point about our Eqn.2 when $e \geq 2$. The denominator $\partial_y f|_{y=y_t}$ is zero mod I, thus, its reciprocal does not exist! However, the ratio $(f/\partial_y f)|_{y=y_t}$ does exist in $\mathbb{F}[[\overline{x}]]$. On the other hand, if $e = 1$ then the denominator $\partial_y f|_{y=y_t}$ is nonzero mod I, thus, it is invertible in $\mathbb{F}[[\overline{x}]]$ and that is necessary for fast algebraic circuit computation (esp. division elimination).

We can compare the NI formula with the recurrence formula (which we call *slow* Newton Iteration) used in [DSY09, Eqn.5], [Oli16, Lem.4.1] for root finding. The slow NI formula is $Y_{t+1} = Y_t - \frac{f(\overline{x}, Y_t)}{\partial_y f(\overline{0}, Y_1)}$, where $Y_t \equiv g \mod I^t$. The rate of convergence of this iteration is linear, as it takes δ many steps (instead of $\log \delta$) to get precision up to degree δ. One can also compare NI with other widespread processes like multifactor Hensel lifting [vzGG13, Sec.15.5], [Zas69] and the implicit function theorem paradigm [KP12, Sec.1.3], [KS16, PSS16]; however, we would not like to digress too much here as the latter concept covers a whole lot of ground in mathematics.

1.4 Proof overview

In all our proofs, we use the reduction of factoring to power series root approximation, and then find the latter using various techniques described before.

Proof idea of Theorem 1: We use the technique of allRootsNI to find the approximations of all the power series roots of $f(\tau \overline{x})$. As we already discussed how to find a polynomial factor g of u_1 (that divides f) from the roots of $f(\tau \overline{x})$, what remains is to analyze the size bound for power series roots that we get from allRootsNI process. We note a few crucial points that help to prove the size bound.

Let d_0 be the degree of $\text{rad}(u_1)$. The number of distinct power series roots, of $u_1(\tau \overline{x})$ wrt y, is d_0. It suffices to approximate the power series roots up to degree d_0, as any nontrivial polynomial factor of $\text{rad}(u_1(\tau \overline{x}))$ has degree less than d_0. Also, a size bound on these factors of the radical directly gives a size bound on the polynomial factor g.

The logarithmic derivative satisfies: $\partial_y \log f(\tau \overline{x}) = \partial_y \log u_0(\tau \overline{x}) + \partial_y \log u_1(\tau \overline{x})$. Since we have size s circuits for both f and u_0, and y is later fixed to random c_i's in $\mathbb{F}$, we can approximate the first two logarithmic derivative circuits modulo I^{d_0+1}. This approximates $\partial_y u_1(\tau \overline{x})/u_1(\tau \overline{x})$.

On this, allRootsNI process is used to approximate the power series roots of $u_1(\tau \overline{x})$ up to degree d_0. The self correcting behavior of the allRootsNI is crucial in the size analysis. If one had to truncate modulo I^{d_0+1} at each recursive step, there would have been a multiplicative blowup (by d_0) in each step, which would end up with an exponential blow up in the size of the roots. The self correcting property allows to complete allRootsNI process, with division gates and partially correct roots $g'_{i,\delta}$, to get a circuit of size $\text{poly}(sd_0)$. The truncation modulo I^{d_0+1}, to get a root of degree $\leq d_0$, is performed only once in the end. See Section 4.1.

11

The steps in the proof of Theorem 1 are constructive. However, to claim that we have an efficient algorithm we will need, in advance, the multiplicity of each of the d_0 roots. It is not clear how to find them efficiently, even in the univariate case $n = 1$, as the multiplicity could be exponentially large.

Proof idea of Theorem 2: The main technique used is NI with multiplicity. The main barrier in resolving high degree case is handling roots with high multiplicities (i.e. super-polynomial in size s). If all the roots of the polynomial have multiplicity equal to one, then we can use classical Newton iteration. If the multiplicity of a root is low (up to poly(s)), we can differentiate and bring down the multiplicity to one. In Theorem 1, we handled the case of high multiplicity by assuming that the radical has small degree.

So, the only remaining case is when both the number of roots, and their multiplicities, are high. Newton iteration with multiplicity helps here. Note that we need to know the multiplicity of the root *exactly* to apply NI with multiplicity; here, we will simply guess them non-uniformly. In the end, the process gives a circuit of size poly(sd) with division gates, giving the root mod I^{d+1}. By using a standard method the division gates can all be pushed "out" to the root. See Section 4.2.

Proof idea of Theorem 3: Here, we show the closure under factoring for the algebraic complexity classes $VF(n^{\log n}), VBP(n^{\log n}), VNP(n^{\log n})$. In fact, we also give randomized $n^{O(\log n)}$-time algorithm to output the factors as formula (resp. algebraic branching program). The key technique here is the classical Newton Iteration. The crucial advantage of NI over other approaches of power series root finding is that NI requires only $\log d$ steps to get precision up to degree d, whereas allRootsNI, [DSY09, Eqn.5] or [Oli16, Lem.4.1] require d steps. This leads to a slower size blow up in the case of restricted models like formula or ABP.

In a formula resp. ABP, we cannot reuse intermediate computations. So each recursive step of NI incurs a blow up by d^2, as one needs to substitute y_t in a degree d polynomial $f(y)$ which may require that many copies of y_t-powers. But, as the NI process has only $\log d$ steps, ultimately, we get $d^{2\log d}$ blow up in the size bound. This is the main idea of the existential results in Theorem 3. Moreover, an interesting by-product is that VF, VBP and VNP are closed under factors if we only consider polynomials with individual degree *constant* (also see [Oli16]).

All the steps in the proof of the existential result are algorithmically efficient except for one. We are recovering all the power series roots and multiplying a few of them to get a non-trivial factor. How do we choose the right combination of the roots which gives a non-trivial factor? If we search for the right combination in a brute-force way, it would need exponential (like 2^d) time complexity. Here, linear algebra saves us; the idea dates back to Kaltofen's algorithm for bivariate factoring. Our contribution lies in the careful analysis of the different steps, coming up with a new algorithm for computing gcd, and making sure that everything works with formulas resp. ABPs.

Consider the transformed polynomial $f(\tau\overline{x})$ that is monic and degree d in y. It will help us if we think of this polynomial as a bivariate (i.e. in y and a new degree-counter T). This somewhat reduces the problem to a two-dimensional case and makes the modular computations feasible (see [KSS15, Sec.1.2.2]). So, we need to apply the map $\overline{x} \mapsto T\overline{x}$, where T is a new formal variable; call the resulting polynomial $\tilde{f}(\overline{x}, T, y)$. This map preserves the power series roots; in fact, we can get the roots of $f(\tau\overline{x})$ by putting $T = 1$. Now comes the most important idea in the algorithm. Approximate a root g_i up to large enough precision (say $k := 2d^2$). Solve the system of linear equations $u = (y - g_i^{\leq k}(Tx)) \cdot v \mod T^{k+1}$ for monic polynomials u, v. Then, u will give a non-trivial factor when we compute $\gcd_y(u, \tilde{f})$. Intuitively, the gcd gives us the irreducible polynomial factor whose root is the power series g_i that we had earlier computed by NI.

Note that a modified gcd computation is needed to actually get a factor as a formula resp. ABP. If one uses the classical Euclidean algorithm, there are d recursive steps to execute; at each step there would be a blow up of d (as for formula or ABP, we cannot reuse any intermediate computation). So, in this approach (eg. the one used in [KSS15]), gcd of the two formulas will be of exponential size. The way we achieve a better bound is by first using NI to approximate *all* the power series roots of u and $\tilde{f}$. Subsequently, we filter the ones that appear in both to learn the gcd. There is an alternate way as well based on our Claim 11. See Section 4.3.

2   Preliminaries

In our proofs we will need some basic results about formulas, ABPs and circuits. In particular, we can efficiently eliminate a division gate, we can extract a homogeneous part, and we can compute a (first-order) derivative. Also, see [KSS15, Sec.2].

Determinant is in VBP and is computable by a $n^{O(\log n)}$ size formula.

We will use properties of $\gcd(f, g)$ and a related determinant polynomial called *resultant*.

To save space we have moved the well known details to Section A.

3   Power series factorization of polynomials

Instead of looking into the factorization over $\mathbb{F}[\overline{x}]$, we look into the more analytic factorization pattern of a polynomial over $\mathbb{F}[[x_1, \ldots, x_n]]$, namely, formal power series of n-variables over field $\mathbb{F}$. To talk about factorization, we need the notion of *uniqueness* which the following proposition ensures.

Proposition 1. *[ZS75, Chap.VII] Power series ring $\mathbb{F}[[x_1, \ldots, x_n]]$ is a unique factorization domain (UFD), and so is $\mathbb{F}[[\overline{x}]][y]$.*

As discussed before, we need to first apply a random linear map, that will make sure that the resulting polynomial splits completely over the ring $\mathbb{F}[[\overline{x}]]$. (Recall: $\mathbb{F}$ is algebraically closed.)

Theorem 4 (Power Series Complete Split). *Let $f \in \mathbb{F}[\overline{x}]$ with $deg(rad(f)) =: d_0 > 0$. Consider $\alpha_i, \beta_i \in_r \mathbb{F}$ and the map $\tau : x_i \mapsto \alpha_i y + x_i + \beta_i$, $i \in [n]$, where y is a new variable.*

Then, over $\mathbb{F}[[\overline{x}]]$, $f(\tau\overline{x}) = k \cdot \prod_{i \in [d_0]} (y - g_i)^{\gamma_i}$, where $k \in \mathbb{F}^*$, $\gamma_i > 0$, and $g_i(\overline{0}) := \mu_i$. Moreover, μ_i's are distinct nonzero field elements.

Proof. Let the irreducible factorization of f be $\prod_{i \in [m]} f_i^{e_i}$. We apply a random τ so that f, thus all its factors, become monic in y (Lemma 28). The monic factors $\tilde{f}_i := f_i(\tau\overline{x})$ remain irreducible ($\because \tau$ is invertible). Also, $\tilde{f}_i(\overline{0}, y) = f_i(\overline{\alpha}y + \overline{\beta})$ and $\partial_y \tilde{f}_i(\overline{0}, y)$ remain coprime ($\because \overline{\beta}$ is random, apply Lemma 26). In other words, $\tilde{f}_i(\overline{0}, y)$ is square free (Lemma 25).

In particular, one can write $\tilde{f}_1(\overline{0}, y)$ as $\prod_{i=1}^{\deg(f_1)} (y - \mu_{1,i})$ for distinct nonzero field elements $\mu_{1,i}$ (ignoring the constant which is the coefficient of the highest degree of y in $\tilde{f}_1$). Using classical Newton Iteration (see Lemma 27 or [BCS13, Thm.2.31]), one can write $\tilde{f}_1(\overline{x}, y)$ as a product of power series $\prod_{i=1}^{\deg(f_1)} (y - g_{1,i})$, with $g_{1,i}(\overline{0}) := \mu_{1,i}$. Thus, each $f_i(\tau\overline{x})$ can be factored into linear factors in $\mathbb{F}[[\overline{x}]][y]$.

As f_i's are irreducible coprime polynomials, by Lemma 26, it is clear that $\tilde{f}_i(\overline{0}, y)$, $i \in [m]$, are mutually coprime. In other words, $\mu_{j,i}$ are distinct and they are $\sum_i \deg(f_i) = d_0$ many. Hence, $f(\tau\overline{x})$ can be completely factored as $\prod_{i \in [m]} f_i(\tau\overline{x})^{e_i} = \prod_{i \in [d_0]} (y - g_i)^{\gamma_i}$, with $\gamma_i > 0$ and the field constants $g_i(\overline{0})$ being distinct. $\qquad\square$

Corollary 5. *Suppose g is a polynomial factor of f. As before let $f(\tau\overline{x}) = \prod_{i \in [m]} f_i(\tau\overline{x})^{e_i}$ $= k \cdot \prod_{i \in [d_0]} (y - g_i)^{\gamma_i}$. As $g(\tau x) \mid f(\tau\overline{x})$ we deduce that $g(\tau x) = k' \prod (y - g_i)^{c_i}$ with $0 \le c_i \le \gamma_i$.*
Moreover, we can get back g by applying τ^{-1} on the resulting polynomial $g(\tau\overline{x})$.

4 Main Results

This section proves Theorems 1–3. The proofs are self contained and we assume for the sake of simplicity that the underlying field $\mathbb{F}$ is algebraically closed and has characteristic 0. When this is not the case, we discuss the corresponding theorems in Section 5.

4.1 Factors of a circuit with low-degree radical: Proof of Theorem 1

In this section, we use Theorem 4 and allRootsNI to partially solve the case of circuits with exponential degree (stated in [Kal86] and studied in [Kal87, Bür04]).

Proof of Theorem 1. From the hypothesis $f = u_0 u_1$. Define $\deg(f) =: d$. Suppose $u_1 = h_1^{e_1} \ldots h_m^{e_m}$, where h_i's are coprime irreducible polynomials. Let d_0 be the degree of $\mathrm{rad}(u_1) = \prod_i h_i$. Note that $\deg(h_i), m \le d_0$ and the multiplicity $e_i \le d \le s^{O(s)}$, where s is the size bound of the input circuit. Thus, to get the size bound of any factor of u_1, it is enough to show that for each i, h_i has a circuit of size $\mathrm{poly}(sd_0)$.

Using Theorem 4, we have $\tilde{f}(\overline{x}, y) := f(\tau\overline{x}) = k \cdot u_0(\tau\overline{x}) \cdot \prod_{i \in [d_0]} (y - g_i)^{\gamma_i}$, with $g_i(\overline{0}) := \mu_i$ being distinct. From Corollary 5 we deduce that $h_i(\tau\overline{x}) = k_i \prod_{i \in [d_0]} (y - g_i^{\le d_0})^{\delta_i} \bmod I^{d_0+1}$, with ideal $I := \langle x_1, \ldots, x_n \rangle$, exponent $\delta_i \in \{0, 1\}$ and nonzero $k_i \in \mathbb{F}$. We can get h_i by applying τ^{-1}. Hence, it is enough to bound the size of $g_i^{\le d_0}$.

Let $\tilde{u}_0 := u_0(\tau\overline{x})$. From the repeated applications of Leibniz rule of the derivative ∂_y, we deduce, $\partial_y \tilde{f} / \tilde{f} = \partial_y \tilde{u}_0 / \tilde{u}_0 + \sum_{i=1}^{d_0} \gamma_i / (y - g_i)$. (Recall: $\partial_y(FG) = (\partial_y F)G + F(\partial_y G)$.)

At this point we move to the formal power series, so that the reciprocals can be approximated as polynomials. Note that $y - g_i$ is invertible in $\mathbb{F}[[\overline{x}]]$ when y is assigned any value $c_i \in \mathbb{F}$ which is not equal to μ_i. We intend to find $g_i \bmod I^\delta$ inductively, for all $\delta \ge 1$. We assume that μ_i's and γ_i's are known. Suppose, we have recovered up to $g_i \bmod I^\delta$ and we want to recover $g_i \bmod I^{\delta+1}$. The relevant recurrence, for $\delta \ge 1$, is:

Claim 6 (Recurrence). $\sum_{i=1}^{d_0} \gamma_i \cdot g_i^{=\delta} / (y - \mu_i)^2 \equiv \partial_y \tilde{f} / \tilde{f} - \partial_y \tilde{u}_0 / \tilde{u}_0 - \sum_i \gamma_i / (y - g_i^{<\delta}) \bmod I^{\delta+1}$.

Proof of Claim 6. Using a power series calculation (Lemma 31), we have $\frac{1}{y - g_i} \equiv \frac{1}{y - (g_i^{<\delta} + g_i^{=\delta})} \equiv$ $\frac{1}{y - g_i^{<\delta}} + \frac{g_i^{=\delta}}{(y - \mu_i)^2} \bmod I^{\delta+1}$. Multiplying by γ_i and summing over $i \in [d_0]$, the claim follows. $\square$

By knowing approximation up to the $\delta - 1$ homogeneous parts of g_i, we want to find the δ-th part by solving a linear system. For concreteness, assume that we have a rational function $g'_{i,\delta-1} := C_{i,\delta-1}/D_{i,\delta-1}$ such that $g'_{i,\delta-1} \equiv g_i^{<\delta} \bmod I^\delta$. Next, we show how to compute $g_i^{\le\delta}$.

We recall the process as outlined in allRootsNI (Section 1.3). In the free variable y, we plug-in d_0 random field value c_i's and get the following system of linear equations: $M \cdot v_\delta = W_\delta$, where M is a $d_0 \times d_0$ matrix with (i, j)-th entry, $M(i, j) := \gamma_j / (c_i - \mu_j)^2$. Column v_δ resp. W_δ is a $d_0 \times 1$ matrix whose i-th entry is denoted $v_\delta(i)$ resp. $(\partial_y \tilde{f} / \tilde{f} - \partial_y \tilde{u}_0 / \tilde{u}_0)|_{y=c_i} - \tilde{G}_{i,\delta}$, where $\tilde{G}_{i,\delta} := \sum_{j=1}^{d_0} \gamma_j / (c_i - g'_{j,\delta-1})$. Think of the solution v_δ as being both in $\mathbb{F}(\overline{x})^{d_0}$ and in $\mathbb{F}[[\overline{x}]]^{d_0}$; both the views help.

Now we will prove two interesting facts. First, M is invertible (Lemma 29). Second, define $g'_{i,0} := \mu_i$ and, for $\delta \ge 1$, $g'_{i,\delta} := g'_{i,\delta-1} + v_\delta(i)$. Then, $g'_{i,\delta}$ approximates g_i well:

Claim 7 (Self-correction). *Let $i \in [d_0]$ and $\delta \geq 0$. Then, $g'_{i,\delta} \equiv g_i^{\leq \delta} \bmod I^{\delta+1}$.*

Proof of Claim 7. We prove this by induction on δ. It is true for $\delta = 0$ by definition. Suppose it is true for $\delta - 1$. This means we have $g'_{i,\delta-1} \equiv g_i^{<\delta} \bmod I^{\delta}$ for all i. Let us write $g'_{i,\delta-1} =: g_i^{<\delta} + A_{i,\delta} + A'_{i,\delta}$, where $A'_{i,\delta} \equiv 0 \bmod I^{\delta+1}$ and $A_{i,\delta}$ is homogeneous of degree δ. Hence, for $i \in [d_0]$, the linear constraint is: $\sum_{j=1}^{d_0} \gamma_j \cdot v_\delta(j)/(c_i - \mu_j)^2 \equiv \partial_y \tilde{f}/\tilde{f} - \partial_y \tilde{u}_0/\tilde{u}_0 - \sum_j \gamma_j/(c_i - g'_{j,\delta-1}) \bmod I^{\delta+1}$.

The "garbage" term $A_{j,\delta}$ in RHS can be isolated using Lemma 31 as: $1/(c_i - g'_{j,\delta-1}) \equiv \frac{1}{c_i - (g_j^{<\delta} + A_{j,\delta})} \equiv 1/(c_i - g_j^{<\delta}) + A_{j,\delta}/(c_i - \mu_j)^2 \bmod I^{\delta+1}$. So, we get:

$$\sum_{j=1}^{d_0} \frac{\gamma_j \cdot v_\delta(j)}{(c_i - \mu_j)^2} \equiv \frac{\partial_y \tilde{f}}{\tilde{f}} - \frac{\partial_y \tilde{u}_0}{\tilde{u}_0} - \sum_{j=1}^{d_0} \frac{\gamma_j}{c_i - g_j^{<\delta}} - \sum_{j=1}^{d_0} \frac{\gamma_j \cdot A_{j,\delta}}{(c_i - \mu_j)^2} \bmod I^{\delta+1}.$$

Rewriting this, using Claim 6, we get:

$$\sum_{j=1}^{d_0} \frac{\gamma_j}{(c_i - \mu_j)^2} \left(v_\delta(j) + A_{j,\delta} \right) \equiv \sum_{j=1}^{d_0} \frac{\gamma_j}{(c_i - \mu_j)^2} \cdot g_j^{=\delta} \bmod I^{\delta+1}.$$

Thus, $\sum_{j=1}^{d_0} \gamma_j \cdot (v_\delta(j) + A_{j,\delta} - g_j^{=\delta})/(c_i - \mu_j)^2 \equiv 0 \bmod I^{\delta+1}$. As we vary $i \in [d_0]$ we deduce, by Lemma 29, that $v_\delta(j) + A_{j,\delta} - g_j^{=\delta} \equiv 0 \bmod I^{\delta+1}$. Hence, $g'_{j,\delta} = g'_{j,\delta-1} + v_\delta(j) \equiv (g_j^{<\delta} + A_{j,\delta}) + (g_j^{=\delta} - A_{j,\delta}) = g_j^{\leq \delta} \bmod I^{\delta+1}$. This proves it for all $j \in [d_0]$. $\qquad\square$

Size analysis: Here we give the overall process of finding factors using allRootsNI technique and analyze the circuit size needed at each step to establish the size bound of the factors. As discussed before, we need to analyze only the power series root approximation $g_i^{\leq \delta}$ or $g'_{i,\delta}$.

At the $(\delta - 1)$-th step of allRootsNI process, we have a multi-output circuit (with division gates) computing $g'_{i,\delta-1}$ as a rational function, for all $i \in [d_0]$. Specifically, let us assume that $g'_{i,\delta-1} =: C_{i,\delta-1}/D_{i,\delta-1}$, where $D_{i,\delta-1}$ is invertible in $\mathbb{F}[[\overline{x}]]$. So, the circuit computing $g'_{i,\delta-1}$ has a division gate at the top that outputs $C_{i,\delta-1}/D_{i,\delta-1}$. We would eliminate this division gate only in the end (see the standard Lemma 21). Now we show how to construct the circuit for $g'_{i,\delta}$, given the circuits for $g'_{i,\delta-1}$.

From $v_\delta = M^{-1} W_\delta$, it is clear that there exist field elements β_{ij} such that $v_\delta(i) = \sum_{j=1}^{d_0} \beta_{ij} W_\delta(j) = \sum_{j=1}^{d_0} \beta_{ij} \left((\partial_y \tilde{f}/\tilde{f} - \partial_y \tilde{u}_0/\tilde{u}_0)|_{y=c_j} - \tilde{G}_{j,\delta} \right)$.

Initially we precompute, for all $j \in [d_0]$, $(\partial_y \tilde{f}/\tilde{f} - \partial_y \tilde{u}_0/\tilde{u}_0)|_{y=c_j}$: Note that $\partial_y \tilde{f}$ has poly(s) size circuit (high degree of the circuit does not matter, see Lemma 22). Invertibility of $\tilde{f}|_{y=c_j}$ and $\tilde{u}_0|_{y=c_j}$ follows from the fact that we chose c_j's randomly. In particular, $\tilde{f}(\overline{0}, y)$, and so $\tilde{u}_0(\overline{0}, y)$, have roots in $\mathbb{F}$ which are distinct from c_j, $j \in [d_0]$. Thus, $\tilde{f}(\overline{x}, c_j)$ and $\tilde{u}_0(\overline{x}, c_j)$ have non-zero constants and so are invertible in $\mathbb{F}[[\overline{x}]]$. Similarly, $\gamma_\ell/(c_j - g'_{\ell,\delta-1})$ exists in $\mathbb{F}[[\overline{x}]]$.

Thus, the matrix recurrence allows us to calculate the polynomials $C_{i,\delta}$ and $D_{i,\delta}$, given their $\delta - 1$ analogues, by adding poly(d_0) many wires and nodes. The precomputations costed us size poly(s, δ). Hence, both $C_{i,\delta}$ and $D_{i,\delta}$ has poly(s, δ, d_0) sized circuit.

We can assume we have only one division gate at the top, as for each gate G we can keep track of numerator and denominator of the rational function computed at G, and simulate all the algebraic operations easily in this representation. When we reach precision $\delta = d_0$, we can eliminate the division gate at the top. As D_{i,d_0} is a unit, we can compute its inverse using the power series inverse formula and approximate only up to degree d_0 (Lemma 20). Finally, the circuit for the polynomial $g_i^{\leq d_0} \equiv C_{i,d_0}/D_{i,d_0} \bmod I^{d_0+1}$, for all $i \in [d_0]$, has size poly(s, d_0).

Altogether, it implies that any factor of u_1 has a circuit of size poly(s, d_0). $\qquad\square$

4.2 Low degree factors of general circuits: Proof of Theorem 2

Here, we introduce an approach to handle the general case when $\mathrm{rad}(f)$ has exponential degree. We show that allowing a special kind of modular division gate gives a small circuit for any low degree factor of f.

The *modular division* problem is to show that if f/g has a representative in $\mathbb{F}[[\overline{x}]]$, where polynomials f and g can be computed by a circuit of size s, then $f/g \bmod \langle \overline{x}^d \rangle$ can be computed by a circuit of size $\mathrm{poly}(sd)$. Note that if g is invertible in $\mathbb{F}[[\overline{x}]]$, then the question of modular division can be solved using Strassen's trick of division elimination [Str73]. But, in our case g is not invertible in $\mathbb{F}[[\overline{x}]]$ (though f/g is well-defined).

Proof of Theorem 2. As discussed before, to show size bound for an arbitrary factor (with low degree) of f, it is enough to show the size bound for the approximations of power series roots. From Theorem 4, $\tilde{f}(\overline{x}, y) = f(\tau \overline{x}) = k \cdot \prod_{i=1}^{d_0} (y - g_i)^{\gamma_i}$, with $g_i(\overline{0}) := \mu_i$ being distinct.

Fix an i from now on. To calculate $g_i^{\leq \delta}$, we iteratively use Newton iteration with multiplicity (as described in Section 1.3) for $\log \delta + 1$ many times. We know that there are rational functions $\hat{g}_{i,t}$ such that $\hat{g}_{i,t+1} := \hat{g}_{i,t} - \gamma_i \cdot \frac{\tilde{f}}{\partial_y \tilde{f}}\big|_{y = \hat{g}_{i,t}}$ and $\hat{g}_{i,t} \equiv g_i \bmod \langle \overline{x} \rangle^{2^t}$. We compute $\hat{g}_{i,t}$'s incrementally, $0 \leq t \leq \log \delta + 1$, by a circuit with division gates. As before, $\tilde{f}$ and $\partial_y \tilde{f}$ have $\mathrm{poly}(s)$ size circuits.

If $\hat{g}_{i,t}$ has S_t size circuit with division, then $S_{t+1} = S_t + O(1)$. Hence, $\hat{g}_{i,\lg \delta+1}$ has $\mathrm{poly}(s, \log \delta)$ size circuit with division.

By keeping track of numerator and denominator of the rational function computed at each gate, we can assume that the only division gate is at the top. As the size of $\hat{g}_{i,\log \delta+1}$ was initially $\mathrm{poly}(s, \log \delta)$ with intermediate division gates, it is easy to see that when division gates are pushed at the top, it computes A/B with size of both A and B still $\mathrm{poly}(s, \log \delta)$.

Finally, a degree δ polynomial factor $h \mid f$ will require us to estimate $g_i^{\leq \delta}$ for that many i's. Thus, such a factor has $\mathrm{poly}(s\delta)$ size circuit, using a single modular division. $\qquad \square$

4.3 Closure of restricted complexity classes: Proof of Theorem 3

This subsection is dedicated towards proving closure results for certain algebraic complexity classes. In fact, for "practical" fields like $\mathbb{Q}, \mathbb{Q}_p$, or $\mathbb{F}_q$ for prime-power q, we give efficient randomized algorithm to output the complete factorization of polynomials belonging to that class (stated as Theorem 15). We use the notation $g \,\|\, f$ to denote that g divides f but g^2 does not divide f. Again, we denote $I := \langle x_1, \ldots, x_n \rangle$

Proof of Theorem 3. There are essentially two parts in the proof. The first part talks only about the existential closure results. In the second part, we discuss the algorithm.

Proof of closure: Given f of degree d, we randomly shift by $\tau : x_i \mapsto x_i + y\alpha_i + \beta_i$. From Theorem 4 we have that $\tilde{f}(\overline{x}, y) := f(\tau \overline{x})$ splits like $\tilde{f} = \prod_{i=1}^{d_0} (y - g_i)^{\gamma_i}$, with $g_i(\overline{0}) =: \mu_i$ being distinct. Here is the detailed size analysis of the factors of polynomials represented by various models of our interest.

Size analysis for formula: Suppose f has a formula of size $n^{O(\log n)}$. To show size bound for all the factors, it is enough to show that the approximations of the power series roots, i.e. $g_i^{\leq d}$ has size $n^{O(\log n)}$ size formula. This follows from the reduction of factoring to approximations of power series roots.

We differentiate $\tilde{f}$ wrt y, $(\gamma_i - 1)$ many times, so that the multiplicity of the root we want to recover becomes exactly one. The differentiation would keep the size $\mathrm{poly}(n^{\log n})$ (Lemma 22). Now, we have $(y - g_i) \,\|\, \tilde{f}^{(\gamma_i - 1)}$ and we can apply classical Newton iteration formula (Section 1.3). For all $0 \leq t \leq \log d + 1$, we compute A_t and B_t such that $A_t/B_t \equiv g_i \bmod I^{2^t}$. Moreover, B_t is invertible in $\mathbb{F}[[\overline{x}]]$ ($\because g_i$ is a simple root of $\tilde{f}^{(\gamma_i - 1)}$).

To implement this iteration using the formula model, each time there would be a blow up of d^2. Note that in a formula, there can be many copies of the same variable in the leaf nodes and if we want to feed something in that variable, we have to make equally many copies. That means we may need to make s ($= \text{size}(f)$) many copies at each step. We claim that it can be reduced to only d^2 many copies.

We can pre-compute (with blow up at most $\text{poly}(sd)$) all the coefficients $C_0, \ldots, C_d$ wrt y, given the formula of $\tilde{f} =: C_0 + C_1 y + \ldots + C_d y^d$ using interpolation. We can do the same for the derivative formula. For details on this interpolation trick, see [Sap16, Lem.5.3]. Using interpolation, we can convert the formula of $\tilde{f}$ and its derivative to the form $C_0 + C_1 y + \ldots + C_d y^d$. In this modified formula, there are $O(d^2)$ many leaves labelled as y. So in the modified formula of the polynomial $\tilde{f}$ and in its derivative, we are computing and plugging in (for y) d^2 copies of $g_i^{<2^t}$ to get $g_i^{<2^{t+1}}$. This leads to d^2 blow up at each step of the iteration.

As B_t's are invertible, we can keep track of the division gates across iterations and, in the end, eliminate them causing a one-time size blow up of $\text{poly}(sd)$ (Lemma 21).

Now, assume that $\text{size}(A_t, B_t) \le S_t$. Then we have $S_{t+1} \le O(d^2 S_t) + \text{poly}(sd)$. Finally, we have $S_{\log d+1} = \text{poly}(sd) \cdot d^{2 \log d} = \text{poly}(n^{\log n})$.

Hence, $g_i^{\le d} \equiv A_{\log d+1}/B_{\log d+1} \bmod I^{d+1}$ has $\text{poly}(n^{\log n})$ size formula, and so does every polynomial factor of f after applying τ^{-1}.

Size analysis for ABP: This analysis is similar to that of the formula model, as the size blow up in each NI iteration for differentiation, division, and truncation (to degree $\le d$) is the same as that for formulas. A noteworthy difference is that we need to eliminate division in *every* iteration (Lemma 20) and we cannot postpone it. This leads to a blow up of d^4 in each step. Hence, $S_{\lg d+1} = \text{poly}(sd) \cdot d^{4 \log d} = \text{poly}(n^{\log n})$.

Size analysis for VNP: Suppose f can be computed by a verifier circuit of size, and witness size, $n^{O(\log n)}$. We call both the verifier circuit size and witness size as size parameter. Now, our given polynomial $\tilde{f}$ has $n^{O(\log n)}$ size parameters. As before, it is enough to show that $g_i^{\le d}$ has $n^{O(\log n)}$ size parameters.

For the preprocessing (taking $\gamma_i - 1$-th derivative of $\tilde{f}$ wrt y), the blow up in the size parameters is only $\text{poly}(n^{\log n})$. Now we analyze the blow up due to classical Newton iteration. We compute A_t and B_t such that $A_t/B_t \equiv g_i \bmod I^{2^t}$. Using the closure properties of VNP (discussed in Section C.1), we see that each time there is a blow up of d^4. The main reason for this blow up is due to the *composition* operation, as we are feeding a polynomial into another polynomial.

Assume that the verifier circuit $\text{size}(A_t, B_t) \le S_t$ and witness size $\le W_t$. Then we have $S_{t+1} \le O(d^4 S_t) + \text{poly}(n^{\log n})$. So, finally we have $S_{\log d+1} = \text{poly}(sd) \cdot d^{4 \log d} = \text{poly}(n^{\log n})$. It is clear that $g_i^{\le d} \equiv A_{\log d+1}/B_{\log d+1} \bmod I^{d+1}$ has $\text{poly}(n^{\log n})$ size verifer circuit. Same analysis works for W_t and witness size remains $n^{O(\log n)}$. Moreover, we get the corresponding bounds for every polynomial factor of f after applying τ^{-1}.

Remark. Recently in a follow-up paper, Chou, Kumar and Solomon [CKS18] have improved our result on VNP, showing that VNP is closed under factors. Their proof uses the reduction of polynomial factoring to power series root approximation. To avoid division gates, they use the slow variant of Newton iteration (as done in [DSY09, Eqn.5], [Oli16, Lem.4.1]) and use it to compute the circuit of an *approximator* polynomial. An approximator polynomial is a polynomial function of the coefficients (w.r.t one variable, say y) of the circuit that gives the power series roots (w.r.t y) approximated up to certain degree. It can be proved that the approximator polynomial has a small sized circuit. To get the approximate power series roots, one has to compose this circuit with the coefficients of the given polynomial. To finish the proof, use Valiant's lemma [Val82] showing VNP is closed under composition.

The same idea does not solve the VF (resp. VBP) closure under factoring questions as it is not clear if there is an approximator polynomial that has a small sized formula (resp. ABP). If one wants to use the slow Newton iteration iteratively, in each step there would be a multiplicative blow-up, as in the formula (resp. ABP) model we have to make copies of the same computation.

Before moving to the constructive part, we discuss a new method for computing gcd of two polynomials, which not only fits well in the algorithm but is also of independent interest. We recall the definition of gcd of two polynomials f, g in the ring $\mathbb{F}[\overline{x}]$: $\gcd(f, g) =: h \iff h|f, h|g$ and $(h'|f, h'|g \Rightarrow h'|h)$. It is unique up to constant multiples.

Claim 8 (Computing formula gcd)**.** *Given two polynomials $f, g \in \mathbb{F}[\overline{x}]$ of degree d and computed by a formula (resp. ABP) of size s. One can compute a formula (resp. ABP) for $\gcd(f, g)$, of size $\mathrm{poly}(s, d^{\log d})$, in randomized $\mathrm{poly}(s, d^{\log d})$ time.*

Proof of Claim 8. The idea is the following. Suppose, $\gcd(f, g) =: h$ is of degree $d > 0$, then we will compute $h(\tau \overline{x})$ for a random map τ as in Theorem 4. We know wlog that $\tilde{f} := f(\tau \overline{x}) = \prod_i (y - A_i)^{a_i}$ and $\tilde{g} := g(\tau \overline{x}) = \prod_i (y - B_i)^{b_i}$, where $A_i, B_i \in \mathbb{F}[[\overline{x}]]$. Since $\mathbb{F}[\overline{x}] \subset \mathbb{F}[[\overline{x}]]$ are UFDs (Proposition 1), we could say wlog that $h(\tau \overline{x}) = \prod_{i \in S}(y - A_i)^{\min(a_i, b_i)}$, where $S = \{i \mid A_i = B_i\}$ after possible rearrangement. Now, as τ is a random invertible map, we can assume that, for $i \neq j$, $A_i \neq B_j$ and that $A_i(\overline{0}) \neq B_j(\overline{0})$ (Lemma 26). So, it is enough to compute $A_i^{\leq d}$ and $B_j^{\leq d}$ and compare them using evaluation at $\overline{0}$. If indeed $A_i = B_i$, then $A_i^{\leq d} = B_i^{\leq d}$. If they are not, they mismatch at the constant term itself! Hence, we know the set S and so we are done once we have the power series roots with repetition.

Using univariate factoring, wrt y, we get all the multiplicities, of the roots, a_i and b_i's, additionally we get the corresponding starting points of classical Newton iteration, i.e. $A_i(\overline{0})$ and $B_i(\overline{0})$'s. Using NI, one can compute $A_i^{\leq d}$ and $B_i^{\leq d}$, for all i. Suppose, after rearrangement of A_i and B_i's (if necessary), we have $A_i = B_i$ for $i \in [s] =: S$ and $A_i \neq B_j$ for $i \in [s+1, d], j \in [s+1, d]$. Lemma 26 can be used to deduce that $A_i(\overline{0}) \neq B_j(\overline{0})$ for $i, j \in [1, d] - S$. So, we have in $\gcd(\tilde{f}, \tilde{g}) = \prod_{i \in S}(y - A_i)^{\min(a_i, b_i)}$: the index set S, the exponents and $A_i(\overline{0})$'s computed.

Size analysis: We compute $A_i^{\leq d}$ and $B_i^{\leq d}$ by NI, (possibly) after making the corresponding multiplicity one by differentiation. It is clear that at each NI step there will be a multiplicative d^2 blow up (due to interpolation, division and truncation). There are $\log d$ iterations in NI. Altogether the truncated roots have $\mathrm{poly}(s, d^{\log d})$ size formula (resp. ABP). This directly implies that $\gcd(\tilde{f}, \tilde{g})$ has $\mathrm{poly}(s, d^{\log d})$ size formula (resp. ABP). By taking the product of the linear factors, truncating to degree d, and applying τ^{-1}, we can compute the polynomial $\gcd(f, g)$.

Randomization is needed for τ and possibly for the univariate factoring over $\mathbb{F}$. Also, it is important to note that $\mathbb{F}$ may not be algebraically closed. Then one has to go to an extension, do the algebraic operations and return back to $\mathbb{F}$. For details, see Section 5.2. □

Randomized Algorithm. We give the broad steps of our algorithm below. We are given $f \in \mathbb{F}[\overline{x}]$, of degree $d > 0$, as input.

1. Choose $\overline{\alpha}, \overline{\beta} \in_r \mathbb{F}^n$ and apply $\tau : x_i \to x_i + \alpha_i y + \beta_i$. Denote the transformed polynomial $f(\tau \overline{x})$ by $\tilde{f}(\overline{x}, y)$. Wlog, from Theorem 4, $\tilde{f}$ has factorization of the form $\prod_{i=1}^{d_0}(y - g_i)^{\gamma_i}$, where $\mu_i := g_i(\overline{0})$ are distinct.

2. Factorize $\tilde{f}(\overline{0}, y)$ over $\mathbb{F}[y]$. This will give γ_i and μ_i's.

3. Fix $i = i_0$. Differentiate $\tilde{f}$, wrt y, $(\gamma_{i_0} - 1)$ many times to make g_{i_0} a simple root.

4. Apply Newton iteration (NI), on the differentiated polynomial, for $k := \lceil \log(2d^2 + 1) \rceil$ iterations; starting with the approximation $\mu_{i_0} \pmod{I}$. We get $g_{i_0}^{<2^k}$ at the end of the process $(\bmod\ I^{2^k})$.

5. Apply the transformation $x_i \mapsto Tx_i$ (T acts as a degree-counter). Consider $\tilde{g}_{i_0} := g_{i_0}^{<2^k}(T\overline{x})$. Solve the following homogeneous linear system of equations, over $\mathbb{F}[\overline{x}]$, in the unknowns u_{ij} and v_{ij}'s,

$$\sum_{0 \le i+j < d} u_{ij} \cdot y^i T^j \;=\; (y - \tilde{g}_{i_0}) \cdot \sum_{\substack{0 \le i < d \\ 0 \le j < 2^k}} v_{ij} \cdot y^i T^j \;\bmod\; T^{2^k} .$$

Solve this system, using Lemma 19, to get a nonzero polynomial (if one exists) $u := \sum_{0 \le i+j < d} u_{ij} \cdot y^i T^j$.

6. If there is no solution, return "f is irreducible".

7. Otherwise, find the minimal solution wrt $\deg_y(u)$ by brute force (try all possible degrees wrt y; it is in $[d-1]$).

8. Compute $G(\overline{x}, y, T) := \gcd_y(u(\overline{x}, y, T), \tilde{f}(T\overline{x}, y))$ using Claim 8.

9. Compute $G(\overline{x}, y, 1)$ and transform it by $\tau^{-1} : x_i \mapsto x_i - \alpha_i y - \beta_i$, $i \in [n]$, and $y \mapsto y$. Output this as an irreducible polynomial factor of f.

Claim 9 (Existence). *If f is reducible, then the linear system (Step 5) has a non-trivial solution.*

Proof of Claim 9. If f is reducible, then let $f = \prod f_i^{e_i}$ be its prime factorization. Assume wlog that $y - g_{i_0} \mid \tilde{f}_1 := f_1(\tau \overline{x})$. Of course $0 < \deg_y(\tilde{f}_1) = \deg(f_1) < d$.

Observe that we are done by picking u to be $\tilde{f}_1(T\overline{x}, y)$. For, total degree of f_1 is $< d$, and so that of $\tilde{f}_1(T\overline{x}, y)$ wrt the variables y, T is $< d$.

Moreover, $y - g_{i_0} \mid \tilde{f}_1 \implies \tilde{f}_1 = (y - g_{i_0})v$, for some $v \in \mathbb{F}[[\overline{x}]][y]$ with $\deg_y < d$. Hence, $\tilde{f}_1 \equiv (y - g_{i_0}^{<2^k}) \cdot v \bmod I^{2^k} \implies u \equiv (y - \tilde{g}_{i_0}) \cdot v(T\overline{x}, y) \bmod T^{2^k}$. This shows the existence of a nontrivial solution of the linear system (Step 5). $\square$

Now, we show that if the linear system has a solution, then the solution corresponds to a non-trivial polynomial factor of f.

Claim 10 (Step 8's success). *If the linear system (Step 5) has a non-trivial solution, then $0 < \deg_y G \le \deg_y u < d$.*

Proof of Claim 10. Suppose (u, v) is the solution provided by the algorithm in Lemma 19 (u being in the unknown LHS and v being the unknown RHS). Consider $G = \gcd_y(u, \tilde{f}(Tx, y))$. We know that there are polynomials a and b such that $au + b\tilde{f}(Tx, y) = \mathrm{Res}_y(u, \tilde{f}(Tx, y))$ (Section A.4). Consider $\deg_T(\mathrm{Res}_y(u, \tilde{f}(Tx, y)))$. As degree of T in u and $\tilde{f}(Tx, y)$ can be at most d, hence degree of T in Resultant can be atmost $2d^2$ (Section A.4). Clearly, $\deg_y G \le \deg_y u < d$. If $\deg_y G = 0$ then the resultant of $u, \tilde{f}(T\overline{x}, y)$ wrt y will be nonzero (Proposition 2). Suppose the latter happens.

Now, we have $u = (y - \tilde{g}_{i_0})v \bmod T^{2^k}$. Since $y - g_{i_0} \mid \tilde{f}$ we get that $y - g_{i_0}(T\overline{x}) \mid \tilde{f}(T\overline{x}, y)$. Assume that $\tilde{f}(Tx, y) =: (y - g_{i_0}(T\overline{x})) \cdot w$.

Thus, we can rewrite the previous equation as: $au + b\tilde{f}(T\overline{x}, y) \equiv (y - \tilde{g}_{i_0})(av + bw) \equiv \mathrm{Res}_y(u, \tilde{f}(Tx, y)) \bmod T^{2^k}$. Note that the latter is nonzero mod T^{2^k} because the resultant is a

nonzero polynomial of $\deg_T < 2^k$. Putting $y = \tilde{g}_{i_0}$ the LHS vanishes, but RHS does not ($\because$ it is independent of y). This gives a contradiction.

Thus, $\mathrm{Res}_y(u, \tilde{f}(Tx, y)) = 0$. This implies that $0 < \deg_y G < d$. $\qquad \square$

Next we show that if one takes the minimal solution u (wrt degree of y), then it will correspond to an irreducible factor of f. We will use the same notation as above.

Claim 11 (Irred. factor). *Suppose $y - g_{i_0} \mid \tilde{f}_1$ and f_1 is an irreducible factor of f. Then, $G = c \cdot \tilde{f}_1(Tx, y)$, for $c \in \mathbb{F}^*$, and $\deg_y(G) = \deg_y(u) = \deg_y(f_1)$ in Step 8.*

Proof of Claim 11. Suppose f is reducible, hence as shown above, G is a non-trivial factor of $\tilde{f}(T\overline{x}, y)$. Recall that $\tilde{f}(T\overline{x}, y) = \prod_i (y - g_i(T\overline{x}))^{\gamma_i}$ is a factorization over $\mathbb{F}[[\overline{x}, T]]$. We have that $y - \tilde{g}_{i_0} \mid G \bmod T^{2^k}$. Thus, $y - g_{i_0}(T\overline{x}) \mid G$ absolutely ($\because$ the power series ring is a UFD and use Theorem 4). So, $y - g_{i_0}(T\overline{x}) \mid \gcd_y(G, \tilde{f}_1(T\overline{x}, y))$ over the power series ring. Since, $\tilde{f}_1(T\overline{x}, y)$ is an irreducible polynomial, we can deduce that $\tilde{f}_1(T\overline{x}, y) \mid G$ in the polynomial ring. So, $\deg_y(f_1) \le \deg_y(G)$.

We have $\deg_y(\tilde{f}_1(T\overline{x}, y)) = \deg(f_1) =: d_1$. By the above discussion, the linear system in Step 7 will not have a solution of $\deg_y(u)$ below d_1. Let us consider the linear system in Step 7 that wants to find u of $\deg_y = d_1$. This system has a solution, namely the one with $u := \tilde{f}_1(T\overline{x}, y) \bmod T^{2^k}$. Then, by the above claim, we will get the G as well in the subsequent Step 8. This gives $\deg_y(G) \le \deg_y(u) = d_1$. With the previous inequality we get $\deg_y(G) = \deg_y(u) = \deg_y(f_1)$. In particular, G and $\tilde{f}_1(Tx, y)$ are the same up to a nonzero constant multiple. $\qquad \square$

Alternative to Claim 8: The above proof (Claim 11) suggests that the gcd question of Step 8 is rather special: One can just write u as $\sum_{0 \le i \le d_1} c_i(\overline{x}, T) y^i$ and then compute the polynomial $G = \sum_{0 \le i \le d_1} (c_i / c_{d_1}) \cdot y^i$ as a formula (resp. ABP), by eliminating division (Lemma 20).

Once we have the polynomial G we can fix $T = 1$ and apply τ^{-1} to get back the irreducible polynomial factor f_1 (with power series root g_{i_0}).

The running time analysis of the algorithm is by now routine. If we start with an f computed by a formula (resp. ABP) of size $n^{O(\log n)}$, then as observed before, one can compute $\tilde{g}_{i_0}$ which has $n^{O(\log n)}$ size formula (resp. ABP). This takes care of Steps 1-4.

Now, solve the linear system in Steps 5-7 of the algorithm. Each entry of the matrix is a formula (resp. ABP) size $n^{O(\log n)}$. The time complexity is similar by invoking Lemma 19.

Steps 8 is to compute gcd of two $n^{O(\log n)}$ size formulas (resp. ABPs) which again can be done in $n^{O(\log n)}$ time giving a size $n^{O(\log n)}$ formula (resp. ABP) as discussed above.

This completes the randomized poly($n^{\log n}$)-time algorithm that outputs $n^{O(\log n)}$ sized factors. $\qquad \square$

Remarks.

1. The above results hold true for the classes $VBP(s), VF(s), VNP(s)$ for any size function $s = n^{\Omega(\log n)}$.

2. By using a *reversal* technique [Oli16, Sec.1.1.2] and a modified τ, our size bound can be shown to be poly($s, d^{\log r}$), where r (resp. d) is the individual-degree (resp. degree) bound of f. So, when r is constant, we get a factor as a poly(s)-size formula (resp. ABP). Oliveira [Oli16] proved the same result for formulas. But, [Oli16] used *slow* Newton iteration and in each iteration the method was different, owing to which the size was poly(s, d^r).

5 Extensions

5.1 Closure of approximative complexity classes

In this section, we show that all our closure results, under factoring, can be naturally generalized to corresponding approximative algebraic complexity classes.

In computer science, the notion of approximative algebraic complexity emerged in early works on matrix multiplication (the notion of border rank, see [BCS13]). It is also an important concept in the geometric complexity theory program (see [GMQ16]). The notion of approximative complexity can be motivated through two ways, *topological* and *algebraic* and both the perspectives are known to be equivalent. Both allow us to talk about the *convergence* $\epsilon \to 0$.

In what follows, we can see ϵ as a formal variable and $\mathbb{F}(\epsilon)$ as the function field. For an algebraic complexity class C, the approximation is defined as follows [BIZ17, Defn.2.1].

Definition 12 (Approximative closure of a class [BIZ17])**.** *Let C be an algebraic complexity class over field $\mathbb{F}$. A family (f_n) of polynomials from $\mathbb{F}[\overline{x}]$ is in the* class $\overline{C}(\mathbb{F})$ *if there are polynomials $f_{n;i}$ and a function $t : \mathbb{N} \mapsto \mathbb{N}$ such that g_n is in the class C over the field $\mathbb{F}(\epsilon)$ with*
$$g_n(\overline{x}) = f_n(\overline{x}) + \epsilon f_{n;1}(\overline{x}) + \epsilon^2 f_{n;2}(x) + \ldots + \epsilon^{t(n)} f_{n;t(n)}(\overline{x}).$$

The above definition can be used to define closures of classes like VF, VBP, VP, VNP which are denoted as $\overline{\text{VF}}, \overline{\text{VBP}}, \overline{\text{VP}}, \overline{\text{VNP}}$ respectively. In these cases one can assume wlog that the degrees of g_n and $f_{n;i}$ are poly(n).

Following Bürgisser [Bür01a]:- Let $K := \mathbb{F}(\epsilon)$ be the rational function field in variable ϵ over the field $\mathbb{F}$. Let R denote the subring of K that consists of rational functions defined in $\epsilon = 0$. Eg. $1/\epsilon \notin R$ but $1/(1+\epsilon) \in R$.

Definition 13. *[Bür01a, Defn.3.1] Let $f \in \mathbb{F}[x_1, \ldots, x_n]$. The* approximative complexity $\overline{size}(f)$ *is the smallest number r, such that there exists F in $R[x_1, \ldots, x_n]$ satisfying $F|_{\epsilon=0} = f$ and circuit size of F over* constants K *is $\leq r$.*

Note that the circuit of F may be using division by ϵ implicitly in an intermediate step. So, we cannot simply assign $\epsilon = 0$ and get a circuit free of ϵ. Also, the degree involved can be arbitrarily large wrt ϵ. Thus, potentially $\overline{size}(f)$ can be smaller than size(f).

Using this new notion of size one can define the analogous class $\overline{\text{VP}}$. It is known to be closed under factors [Bür01a, Thm.4.1]. The idea is to work over $\mathbb{F}(\epsilon)$, instead of working over $\mathbb{F}$, and use Newton iteration to approximate power series roots. Note that in the case of $\overline{\text{VF}}, \overline{\text{VBP}}, \overline{\text{VP}}$ and $\overline{\text{VNP}}$ the polynomials have poly(n) degree. So, by using repeated differentiation, we can assume the power series root (of $\tilde{f} := f(\tau\overline{x})$) to be simple (i.e. multiplicity$= 1$) and apply classical NI. We need to carefully analyze the implementation of this idea.

Root finding using NI over K. For degree-d $f \in \mathbb{F}[\overline{x}]$ if $\overline{size}(f) = s$ then: $\exists F \in R[\overline{x}]$ with a size s circuit satisfying $F|_{\epsilon=0} = f$. The degree of F wrt $\overline{x}$ may be greater than d. In that case we can extract the part up to degree d and truncate the rest [Bür04, Prop.3.1]. So wlog $\deg_{\overline{x}}(F) = \deg(f)$.

By applying a random τ (using constants $\mathbb{F}$) we can assume that $\tilde{F} := F(\tau\overline{x}) \in R[\overline{x}, y]$ is *monic* (i.e. leading-coefficient, wrt y in $\tilde{F}$, is invertible in R). Otherwise, $\deg_y(\tilde{F}) = \deg_y(\tilde{f}) = \deg_{\overline{x}}(f)$ will decrease on substituting $\epsilon = 0$ contradicting $F|_{\epsilon=0} = f$. Wlog, we can assume that the leading-coefficient of $\tilde{F}$ wrt y is 1 and the y-monomial's degree is d. From now on we have $\tilde{F}|_{\epsilon=0} = \tilde{f}$ and both have their leading-coefficients 1 wrt y.

Let μ be a root of $\tilde{f}(\overline{0}, y)$ of multiplicity one (as discussed before). Since $\tilde{F}(\overline{0}, y) \equiv \tilde{f}(\overline{0}, y)$ mod ϵ, we can build a power series root $\mu(\epsilon) \in \mathbb{F}[[\epsilon]]$ of $\tilde{F}(\overline{0}, y)$ using NI, with μ as the starting point. But $\mu(\epsilon)$ may not converge in K. To overcome this obstruction [Bür01a] devised a clever trick.

Define $\hat{F} := \tilde{F}(\overline{x}, y + \mu + \epsilon) - \tilde{F}(\overline{0}, \mu + \epsilon)$. Note that $(\overline{0}, 0)$ is a simple root of $\hat{F}(\overline{x}, y)$ [Bür04, Eqn.5]. So, a power series root y_∞ of $\hat{F}$ can be built iteratively by classic NI (Lemma 27):

$$y_{t+1} := y_t - \left.\frac{\hat{F}}{\partial_y \hat{F}}\right|_{y = y_t}.$$

Where, $y_\infty \equiv y_t \mod \langle \overline{x} \rangle^{2^t}$. One can easily prove that y_t is defined over the coefficient field K, using induction on t.

Note that $\hat{F}|_{\epsilon=0} = \tilde{f}(\overline{x}, y + \mu) - \tilde{f}(\overline{0}, \mu) = \tilde{f}(\overline{x}, y + \mu)$. So, y_∞ is associated with a root of $\tilde{f}$ as well. This implies that by using several such roots y_∞, we can get an appropriate product $\hat{G} \in R[\overline{x}, y]$, such that an actual polynomial factor of $\tilde{f}$ (over field $\mathbb{F}$) equals $\hat{G}|_{\epsilon=0}$.

The above process, when combined with the first part of the proof of Theorem 3, does imply:

Theorem 14 (Approximative factors). *The approximative complexity classes $\overline{VF}(n^{\log n})$, $\overline{VBP}(n^{\log n})$ and $\overline{VNP}(n^{\log n})$ are closed under factors.*

The same question for the classes $\overline{VF}$, $\overline{VBP}$ and $\overline{VNP}$ we leave as an open question. (Though, for the respective bounded individual-degree polynomials we have the result as before.)

5.2 When field $\mathbb{F}$ is not algebraically closed

We show that all our results "partially" hold true for fields $\mathbb{F}$ which are not algebraically closed. The common technique used in all the proofs is the structural result (Theorem 4) which talks about power series roots with respect to y. Recall that we use a random linear map $\tau : x_i \mapsto x_i + \alpha_i y + \beta_i$, where $\alpha_i, \beta_i \in_r \mathbb{F}$, to make the input polynomial f monic in y and the individual degree of y equal to $d := \deg(f)$. If we set all the variables to zero except y, we get a univariate polynomial $\tilde{f}(\overline{0}, y)$ whose roots we are interested in finding explicitly.

The other common technique in our proofs is the classical NI, which starts with just one field root, say μ_1 of $\tilde{f}(\overline{0}, y)$, and builds the full power series on it. Let $E \subsetneq \overline{\mathbb{F}}$ be the smallest field where a root μ_1 can be found. Say, $g | \tilde{f}_1(\overline{0}, y)$ is the minimal polynomial for μ_1. The degree of the extension $E := \mathbb{F}[z]/(g(z))$ is at most d. So, computations over E can be done efficiently. The key idea is to view $E/\mathbb{F}$ as a vector space and simulate the arithmetic operations over E by operations over $\mathbb{F}$. The details of this kind of simulation can be seen in [vzGG13]. In circuits it means that we make $\deg(E/\mathbb{F})$ copies of each gate and simulate the algebraic operations on these 'tuples' following the $\mathbb{F}$-module structure of $E[\overline{x}]$.

Once we have found all the power series roots of $\tilde{f}(\overline{x}, y)$ over $E[[\overline{x}]]$, say starting from each of the conjugates $\mu_1, \ldots, \mu_i \in E$, it is easy to get a polynomial factor in $E[\overline{x}, y]$. This factor will not be in $\mathbb{F}[\overline{x}, y]$, unless E is a splitting field of $\tilde{f}_1(\overline{0}, y)$. A more practical method is: While solving the linear system over E in Steps 5-7 (Algorithm in Theorem 3) we can demand an $\mathbb{F}$-solution u. Basically, at the level of algorithm in Lemma 19, we can rewrite the linear system $Mw = (\sum_{0 \le i \le d} M_i z^i) \cdot w = 0$ as $M_i w = 0$ ($i \in [0, d]$), where the entries of the matrix M_i are given as formulas (resp. ABP) computing a poly(n) degree polynomial in $\mathbb{F}[\overline{x}]$. This way we get the desired $\mathbb{F}$-solution u. Then, Steps 8-9 will yield an irreducible polynomial factor of f in $\mathbb{F}[\overline{x}, y]$. This sketches the following more practical version of Theorem 3.

Theorem 15. *For $\mathbb{F}$ a number field, a local field, or a finite field (with characteristic $> \deg(f)$), there exists a randomized poly($sn^{\log n}$)-time algorithm that: for a given $n^{O(\log n)}$ size formula (resp. ABP) f of poly(n)-degree and bitsize s, outputs $n^{O(\log n)}$ sized formulas (resp. ABPs) corresponding to each of the nontrivial factors of f.*

Note that over these fields there are famous randomized algorithms to factor univariate polynomials in the base case, see [vzGG13, Part III] & [Pau01].

The allRootsNI method in Theorem 1 seems to require all the roots $\mu_i, i \in [d_0]$, to begin with. Let $\tilde{u}_1 := \mathrm{rad}(u_1(\tau\overline{x}))$. Since μ_i's are in the *splitting field* $E \subset \overline{\mathbb{F}}$ of $\mathrm{rad}(\tilde{u}_1(\overline{0}, y))$, we do indeed get the size bound of the power series roots $g_i^{\leq d_0}$ of $\tilde{u}_1$ assuming the constants from E. As seen in the proof, any irreducible polynomial factor $\tilde{h}_i := h_i(\tau\overline{x})$ of $\mathrm{rad}(\tilde{u}_1)$ is some product of these $(y - g_i^{\leq d_0})$'s mod I^{d_0+1}. So, for the polynomial $\tilde{h}_i$ in $\mathbb{F}[\overline{x}, y]$ we get a size upper bound over constants E. We leave it as an open question to transfer it over constants $\mathbb{F}$ (note: $E/\mathbb{F}$ can be of exponential degree).

5.3 Multiplicity issue in prime characteristic

The main obstruction in prime characteristic is when the multiplicity of a factor is a p-multiple, where $p \geq 2$ is the characteristic of $\mathbb{F}$. In this case, all versions of Newton iteration fail. This is because the derivative of a p-powered polynomial vanishes. When p is greater than the degree of the input polynomial, these problems do not occur, so all our theorems hold (also see Section 5.2).

When p is smaller than the degree of the input polynomial in Theorem 3, adapting an idea from [KSS15, Sec.3.1], we claim that we can give $n^{O(\lambda \log n)}$-sized formula (resp. ABP) for the p^{e_i}-th power of f_i, where f_i is a factor of f whose multiplicity is divisible exactly by p^{e_i}, and λ is the number of distinct p-powers that appear.

Note that presently it is an open question to show that: If a circuit (resp. formula resp. ABP) of size s computes f^p, then f has a poly(sp)-sized circuit (resp. formula resp. ABP).

Theorem 3 can be extended to all characteristic as follows.

Theorem 16. *Let $\mathbb{F}$ be of characteristic $p \geq 2$. Suppose the $\mathrm{poly}(n)$-degree polynomial given by a $n^{O(\log n)}$ size formula (resp. ABP) factors into irreducibles as $f(\overline{x}) = \prod_i f_i^{p^{e_i} j_i}$, where $p \nmid j_i$. Let $\lambda := \#\{e_i | i\}$.*

Then, there is a $\mathrm{poly}(n^{\lambda \log n})$-size formula (resp. ABP) computing $f_i^{p^{e_i}}$ over $\overline{\mathbb{F}}_p$.

Proof sketch. Note that $\lambda = O(\log_p n)$.

Let the transformed polynomial of degree d split into power series roots as follows: $\tilde{f} := f(\tau\overline{x}, y) = \prod_{i=1}^{d_0} (y - g_i)^{\gamma_i}$.

$p \nmid \gamma_i$: If g_i is such that $p \nmid \gamma_i$, then we can find the corresponding power series roots using Newton iteration and recover all such factors. After recovering all such irreducible polynomial factors, we can divide $\tilde{f}$ by their product. Let $G := \tilde{f} / \prod_{p \nmid \gamma_i} (y - g_i)^{\gamma_i}$. Clearly, G is now a p-power polynomial.

$p \mid \gamma_i$: Computing the highest power of p that divides the exponent of G (given by a formula resp. ABP) is easy. First, write the polynomial as $G = c_0 + c_1 y + ... + c_d y^d$ using interpolation. Note that it is a p^e-th power iff: $c_i = 0$ whenever $p^e \nmid i$, and p^{e+1} does not have this property. After computing the right value of p^e, we can reduce factoring to the case of a non-p-power.

Rewrite G as $\hat{G} := \sum_{p^e | i} c_i(\overline{x}) \cdot y^{i/p^e}$, i.e. replacing y^{p^e} by y. Clearly, g is an irreducible factor of G iff $\hat{g}$ is an irreducible factor of $\hat{G}$.

We can now apply NI to find the roots of $\tilde{G}$, that have multiplicity coprime to p. Divide by their product and then repeat the above.

Size analysis. If G can be computed by a size s formula (resp. ABP), $\hat{G}$ can be computed by a size $O(d^2 s)$ formula (resp. ABP). Similarly, a single division gate leads to a blow up by a factor of $O(d^2)$. The number of times we need to eliminate division is at most $\lambda \log d$. So the overall size is $n^{O(\lambda \log n)}$.

However, the splitting field E where we get all the roots of $\tilde{f}(\overline{0}, y)$ may be of degree $\Omega(d!)$. So, we leave the efficiency aspects of the algorithm as an open question. □

High degree case. Note that the above idea cannot be implemented efficiently in the case of high degree circuits. Still we can extend our Theorem 1 using allRootsNI. The key observation is that the allRootsNI formula still holds but the summands that appear are exactly the ones corresponding to g_i with $\gamma_i \neq 0 \mod p$.

This motivates the definition of a partial radical: $\mathrm{rad}_p(f) := \prod_{p \nmid e_i} f_i$, if the prime factorization of f is $\prod_i f_i^{e_i}$.

Theorem 17. *Let $\mathbb{F}$ be of characteristic $p \geq 2$. Let $f = u_0 u_1$ such that size(f)+size(u_0) $\leq s$. Any factor of $rad_p(u_1)$ has size poly($s + \deg(rad_p(u_1))$) over $\overline{\mathbb{F}}$.*

Proof idea: Observe that the roots with multiplicity divisible by p do not contribute to the allRootsNI process. So, the process works with $\mathrm{rad}_p(u_1)$ and the linear algebra complexity involved is polynomial in its degree.

6 Conclusion

The old *Factors conjecture* states that for a nonzero polynomial f: $g \mid f \implies \mathrm{size}(g) \leq \mathrm{poly}(\mathrm{size}(f), \deg(g))$. Motivated by Theorem 1, we would like to strengthen Factors conjecture to a conjecture about the squarefree part of a circuit:

Conjecture 1 (radical). *For a nonzero f: $min\{deg(rad(f)), size(rad(f))\} \leq poly(size(f))$.*

Is the Radical conjecture true if we replace size by $\overline{\mathrm{size}}$?

In low degree regime also there are many open questions. Can we identify a class "below" VP that is closed under factoring? We conclude with some interesting questions.

1. Are VF, VBP or VNP closed under factoring? We might consider Theorem 3 as a positive evidence. Additionally, note that these classes are already closed under e-th root taking. This is easy to see using the classic Taylor series of $(1 + f)^{1/e}$, where $f \in \langle \overline{x} \rangle$.

   In fact, what about the classes which are contained in $VF(n^{\log n})$ but larger than VF. For example, is $\mathrm{VF}(n^{\log \log n})$ closed under factoring?

2. Can we find a suitable analog of Strassen's (non-unit) division elimination for high degree circuits? This, by Theorem 2, will resolve Factors conjecture.

   Interestingly, if we can do Strassen's (unit) division elimination in size polynomial in $\log(\deg)$ then it will imply VP= $\overline{\mathrm{VP}}$. In particular, it will prove Factors conjecture.

3. Our results weaken when $\mathbb{F}$ is not algebraically closed or has a small prime characteristic (Sections 5.2, 5.3). Can we strengthen the methods to work for all $\mathbb{F}$?

References

[Abe73]   Oliver Aberth. Iteration methods for finding all zeros of a polynomial simultaneously. *Mathematics of computation*, 27(122):339–344, 1973. 11

[AGS18]   Manindra Agrawal, Sumanta Ghosh, and Nitin Saxena. Bootstrapping variables in algebraic circuits. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 1166–1179, 2018. 3, 4

[AV08]   Manindra Agrawal and V Vinay. Arithmetic circuits: A chasm at depth four. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 67–75. IEEE, 2008. 3

[AW11]   Eric Allender and Fengming Wang. On the power of algebraic branching programs of width two. *Automata, Languages and Programming*, pages 736–747, 2011. 4

[BCS13]   Peter Bürgisser, Michael Clausen, and Amin Shokrollahi. *Algebraic complexity theory*, volume 315. Springer Science & Business Media, 2013. 2, 9, 13, 21, 34, 36

[BCSS98]   Lenore Blum, Felipe Cucker, Michael Shub, and Steve Smale. *Complexity and Real Computation*. Springer Science & Business Media, 1998. 6

[BIZ17]   Karl Bringmann, Christian Ikenmeyer, and Jeroen Zuiddam. On algebraic branching programs of small width. In *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia*, pages 20:1–20:31, 2017. 4, 21

[BOC92]   Michael Ben-Or and Richard Cleve. Computing algebraic formulas using a constant number of registers. *SIAM Journal on Computing*, 21(1):54–58, 1992. 3

[BSR+05]   Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pages 89–96. ACM, 2005. 3

[BSS89]   Lenore Blum, Mike Shub, and Steve Smale. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bulletin (New Series) of the American Mathematical Society*, 21(1):1–46, 1989. 8

[Bür01a]   Peter Bürgisser. The complexity of factors of multivariate polynomials. In *In Proc. 42th IEEE Symp. on Foundations of Comp. Science*, 2001. 4, 21

[Bür01b]   Peter Bürgisser. On implications between P-NP-hypotheses: Decision versus computation in algebraic complexity. In *MFCS*, pages 3–17. Springer, 2001. 4

[Bür04]   Peter Bürgisser. The complexity of factors of multivariate polynomials. *Foundations of Computational Mathematics*, 4(4):369–396, 2004. (Preliminary version in FOCS 2001). 4, 14, 21, 22

[Bür13]   Peter Bürgisser. *Completeness and reduction in algebraic complexity theory*, volume 7. Springer Science & Business Media, 2013. 3, 4, 36

[CKS18]   Chi-Ning Chou, Mrinal Kumar, and Noam Solomon. Some closure results for polynomial factorization and applications. eccc report tr18-052. In *Electronic Colloquium on Computational Complexity (ECCC)*, 2018. 4, 17

[CRS96]   Richard Courant, Herbert Robbins, and Ian Stewart. *What is Mathematics?: an elementary approach to ideas and methods.* Oxford University Press, USA, 1996. 9

[DB08]   Germund Dahlquist and Åke Björck. Numerical methods in scientific computing, volume I. *Society for Industrial and Applied Mathematics*, 2008. 6, 11

[DMM+14]   Arnaud Durand, Meena Mahajan, Guillaume Malod, Nicolas de Rugy-Altherre, and Nitin Saurabh. Homomorphism polynomials complete for VP. In *34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS*, pages 493–504, 2014. 8

[DSY09]   Zeev Dvir, Amir Shpilka, and Amir Yehudayoff. Hardness-randomness tradeoffs for bounded depth arithmetic circuits. *SIAM Journal on Computing*, 39(4):1279–1293, 2009. (Preliminary version in STOC'08). 3, 4, 11, 12, 17

[Dur60]   Émile Durand. Solutions numériques des équations algébriques. Tome I, Équations du type F(x)= 0, racines d'un polynôme. 1960. 11

[Ehr67]   Louis W Ehrlich. A modified Newton method for polynomials. *Communications of the ACM*, 10(2):107–108, 1967. 11

[FS15]   Michael A Forbes and Amir Shpilka. Complexity theory column 88: Challenges in polynomial factorization. *ACM SIGACT News*, 46(4):32–49, 2015. 3

[FS17]   Michael A. Forbes and Amir Shpilka. A PSPACE construction of a hitting set for the closure of small algebraic circuits. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:163, 2017. (To appear in 50th ACM Symposium on Theory of Computing (STOC), 2018). 4

[FSTW16]   Michael A Forbes, Amir Shpilka, Iddo Tzameret, and Avi Wigderson. Proof complexity lower bounds from algebraic circuit complexity. In *Proceedings of the 31st Conference on Computational Complexity*, page 32. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016. 4

[GMQ16]   Joshua A. Grochow, Ketan D. Mulmuley, and Youming Qiao. Boundaries of VP and VNP. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55, pages 34:1–34:14, 2016. 4, 21

[GMS+86]   Philip E Gill, Walter Murray, Michael A Saunders, John A Tomlin, and Margaret H Wright. On projected Newton barrier methods for linear programming and an equivalence to Karmarkar's projective method. *Mathematical programming*, 36(2):183–209, 1986. 3

[Gro15]   Joshua A Grochow. Unifying known lower bounds via geometric complexity theory. *computational complexity*, 24(2):393–475, 2015. 4

[GS98]   Venkatesan Guruswami and Madhu Sudan. Improved decoding of reed-solomon and algebraic-geometric codes. In *Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on*, pages 28–37. IEEE, 1998. 3

[GTZ88]   Patrizia Gianni, Barry Trager, and Gail Zacharias. Gröbner bases and primary decomposition of polynomial ideals. *Journal of Symbolic Computation*, 6(2):149–167, 1988. 3

[IKRS12]   Gábor Ivanyos, Marek Karpinski, Lajos Rónyai, and Nitin Saxena. Trading grh for algebra: algorithms for factoring polynomials and related structures. *Mathematics of Computation*, 81(277):493–531, 2012. 3

[Jan11]   Maurice J Jansen. Extracting roots of arithmetic circuits by adapting numerical methods. In *2nd Symposium on Innovations in Computer Science (ICS 2011)*, pages 87–100, 2011. 4

[Kal85]   Erich Kaltofen. Computing with polynomials given by straight-line programs I: greatest common divisors. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, pages 131–142, 1985. 3

[Kal86]   Erich Kaltofen. Uniform closure properties of p-computable functions. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, May 28-30, 1986, Berkeley, California, USA*, pages 330–337, 1986. 3, 14

[Kal87]   Erich Kaltofen. Single-factor hensel lifting and its application to the straight-line complexity of certain polynomials. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 443–452. ACM, 1987. 3, 5, 14, 32

[Kal89]   Erich Kaltofen. Factorization of polynomials given by straight-line programs. *Randomness and Computation*, 5:375–412, 1989. 3, 4, 5, 8

[Kal90]   Erich Kaltofen. Polynomial factorization 1982-1986. *Dept. of Comp. Sci. Report*, pages 86–19, 1990. 3

[Kal92]   Erich Kaltofen. Polynomial factorization 1987–1991. *LATIN'92*, pages 294–313, 1992. 3

[Kay11]   Neeraj Kayal. Efficient algorithms for some special cases of the polynomial equivalence problem. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 1409–1421. Society for Industrial and Applied Mathematics, 2011. 3

[Kem10]   Gregor Kemper. *A course in Commutative Algebra*, volume 256. Springer Science & Business Media, 2010. 8

[Ker66]   Immo O Kerner. Ein Gesamtschrittverfahren zur Berechnung der Nullstellen von Polynomen. *Numerische Mathematik*, 8(3):290–294, 1966. 11

[KI03]   Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 355–364. ACM, 2003. 3, 4

[KK08]   Erich Kaltofen and Pascal Koiran. Expressing a fraction of two determinants as a determinant. In *Proceedings of the twenty-first international symposium on Symbolic and algebraic computation*, pages 141–146. ACM, 2008. 4

[KP12]   Steven G Krantz and Harold R Parks. *The implicit function theorem: history, theory, and applications.* Springer Science & Business Media, 2012. 8, 11

[KS06]   Neeraj Kayal and Nitin Saxena. Complexity of ring morphism problems. *computational complexity*, 15(4):342–390, 2006. 3

[KS09]   Zohar S Karnin and Amir Shpilka. Reconstruction of generalized depth-3 arithmetic circuits with bounded top fan-in. In *Computational Complexity, 2009. CCC'09. 24th Annual IEEE Conference on*, pages 274–285. IEEE, 2009. 3

[KS16]   Mrinal Kumar and Shubhangi Saraf. Arithmetic circuits with locally low algebraic rank. In *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, pages 34:1–34:27, 2016. 11

[KSS15]   Swastik Kopparty, Shubhangi Saraf, and Amir Shpilka. Equivalence of polynomial identity testing and polynomial factorization. *computational complexity*, 24(2):295–331, 2015. 8, 12, 13, 23, 30, 31

[KST18]   Mrinal Kumar, Ramprasad Saptharishi, and Anamay Tengse. Near-optimal bootstrapping of hitting sets for algebraic circuits. *arXiv preprint arXiv:1807.06323*, 2018. (To appear in SODA 2019). 4

[Lec02]   Grégoire Lecerf. Quadratic newton iteration for systems with multiplicity. *Foundations of Computational Mathematics*, 2(3):247–293, 2002. 11

[LG14]   François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th international symposium on symbolic and algebraic computation*, pages 296–303. ACM, 2014. 4

[LLMP90]   Arjen K Lenstra, Hendrik W Lenstra, Mark S Manasse, and John M Pollard. The number field sieve. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 564–572. ACM, 1990. 3

[LN97]   Rudolph Lidl and Harald Niederreiter. *Finite Fields*. Cambridge University Press, Cambridge, UK, 1997. 33

[LS78]   Richard J Lipton and Larry J Stockmeyer. Evaluation of polynomials with super-preconditioning. *Journal of Computer and System Sciences*, 16(2):124–139, 1978. 3

[LV16]   Anand Louis and Santosh Srinivas Vempala. Accelerated newton iteration for roots of black box polynomials. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS'16*, pages 732–740, 2016. 5

[Mah14]   Meena Mahajan. Algebraic complexity classes. In *Perspectives in Computational Complexity*, pages 51–75. Springer, 2014. 2, 30

[Mul12a]   Ketan D. Mulmuley. The GCT program toward the P vs. NP problem. *Commun. ACM*, 55(6):98–107, June 2012. 4

[Mul12b]   Ketan D. Mulmuley. Geometric complexity theory V: Equivalence between blackbox derandomization of polynomial identity testing and derandomization of Noether's normalization lemma. In *FOCS*, pages 629–638, 2012. 4

[Mul17]   Ketan Mulmuley. Geometric complexity theory V: Efficient algorithms for Noether normalization. *Journal of the American Mathematical Society*, 30(1):225–309, 2017. 3, 4

[MV97]   Meena Mahajan and V Vinay. A combinatorial algorithm for the determinant. In *SODA*, pages 730–738, 1997. 30, 31

[New69]   Isaac Newton. De analysi per aequationes numero terminorum infinitas [on analysis by infinite series] (in latin). 1669. (published in 1711 by William Jones). 11

[Oli16]   Rafael Oliveira. Factors of low individual degree polynomials. *Computational Complexity*, 2(25):507–561, 2016. (Preliminary version in CCC'15). 4, 11, 12, 17, 20

[OR00]   James M Ortega and Werner C Rheinboldt. *Iterative solution of nonlinear equations in several variables*. SIAM, 2000. 3

[Pau01]   Sebastian Pauli. Factoring polynomials over local fields. *Journal of Symbolic Computation*, 32(5):533–547, 2001. 22

[Pla77a]   David Alan Plaisted. New NP-hard and NP-complete polynomial and integer divisibility problems. In *Foundations of Computer Science, 18th Annual Symposium on*, pages 241–253. IEEE, 1977. 5

[Pla77b]   David Alan Plaisted. Sparse complex polynomials and polynomial reducibility. *Journal of Computer and System Sciences*, 14(2):210–221, 1977. 4, 5

[PSS16]   Anurag Pandey, Nitin Saxena, and Amit Sinhababu. Algebraic independence over positive characteristic: New criterion and applications to locally low algebraic rank circuits. In *41st International Symposium on Mathematical Foundations of Computer Science, MFCS 2016, August 22-26, 2016 - Kraków, Poland*, pages 74:1–74:15, 2016. (In print, Computational Complexity, 2018). 11

[Sap16]   Ramprasad Saptharishi. A survey of lower bounds in arithmetic circuit complexity. *URL https: // github. com/dasarpmar/lowerbounds-survey/releases. Version*, 3(0), 2016. 17, 32

[Sch77]   Claus-Peter Schnorr. Improved lower bounds on the number of multiplications/divisions which are necessary to evaluate polynomials. In *International Symposium on Mathematical Foundations of Computer Science*, pages 135–147. Springer, 1977. 3

[Sch80]   J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, October 1980. 9, 31, 32, 34, 35

[Sin16]   Gaurav Sinha. Reconstruction of real depth-3 circuits with top fan-in 2. In *31st Conference on Computational Complexity*, 2016. 3

[Str73]   Volker Strassen. Vermeidung von divisionen. *Journal für die reine und angewandte Mathematik*, 264:184–202, 1973. 8, 16, 31

[Sud97]   Madhu Sudan. Decoding of reed solomon codes beyond the error-correction bound. *Journal of complexity*, 13(1):180–193, 1997. 3

[SY10]   Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends® in Theoretical Computer Science*, 5(3–4):207–388, 2010. 2, 30, 31, 32

[Tay15]   Brook Taylor. Methodus incrementorum directa et inversa [direct and reverse methods of incrementation] (in latin). 1715. (Translated into English in Struik, D. J. (1969). A Source Book in Mathematics 1200–1800. Cambridge, Massachusetts: Harvard University Press. pp. 329–332.). 5

[Val79]   Leslie G. Valiant. Completeness classes in algebra. In *Proceedings of the 11h Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA*, pages 249–261, 1979. 3

[Val82]   L Valiant. Reducibility by algebraic projections in: Logic and algorithmic. In *Symposium in honour of Ernst Specker*, pages 365–380, 1982. 5, 17

[VL97]   Paul MB Vitanyi and Ming Li. An introduction to kolmogorov complexity and its applications. 34(10), 1997. 6

[VSBR83]   Leslie G. Valiant, Sven Skyum, Stuart Berkowitz, and Charles Rackoff. Fast parallel computation of polynomials using few processors. *SIAM Journal on Computing*, 12(4):641–644, 1983. 3, 8

[vzGG13]   Joachim von zur Gathen and Jürgen Gerhard. *Modern computer algebra*. Cambridge university press, 2013. 11, 22, 33

[vzGK85]   Joachim von zur Gathen and Erich Kaltofen. Factoring sparse multivariate polynomials. *Journal of Computer and System Sciences*, 31(2):265–287, 1985. 4

[Zas69]   Hans Zassenhaus. On Hensel factorization, I. *Journal of Number Theory*, 1(3):291–311, 1969. 11

[ZS75]   Oscar Zariski and Pierre Samuel. *Commutative algebra. II. Reprint of the 1960 edition*, volume 29. Graduate Texts in Mathematics, 1975. 13

A   Preliminaries

A.1   Definition of ABP

ABP is a *skew* circuit, i.e. each multiplication gate has fanin two with at least one of its inputs being a variable or a field constant. A completely different definition can be given via layered graphs or iterated matrix multiplication or symbolic determinant. Famously, they are all equivalent up to polynomial blow up [Mah14].

Definition 18 (Algebraic Branching Program). *An algebraic branching program (ABP) is a layered graph with a unique source vertex (say s) and a unique sink vertex (say t). All edges are from layer i to $i + 1$ and each edge is labelled by a linear polynomial. The polynomial computed by the ABP is defined as $f = \sum_{\gamma:s \rightsquigarrow t} wt(\gamma)$, where for every path γ from s to t, the weight $wt(\gamma)$ is defined as the product of the labels over the edges forming γ.*

Size of the ABP is defined as the total number of edges in the ABP. Width *is the maximum number of vertices in a layer.*

Equivalently, one can define f as a product of matrices (of dimension at most the width), each one having linear polynomials as entries. For more details, see [SY10].

It is a famous result that the ABP model is the same as symbolic determinant [MV97].

A.2   Randomized algorithm for linear algebra using PIT

The following lemma from [KSS15] discusses how to perform linear algebra when the coefficients of vectors are given as formula (resp. ABP). This will be crucially used in Theorem 3 when we would give an algorithm to output the factors.

Lemma 19. *(Linear algebra using PIT [KSS15, Lem.2.6]) Let $M = (M_{i,j})_{k \times n}$ be a matrix (where k is $n^{O(1)}$) with each entry being a degree $\leq n^{O(1)}$ polynomial in $\mathbb{F}[\overline{x}]$. Suppose, we have algebraic formula (resp. ABP) of size $\leq n^{O(\log n)}$ computing each entry. Then, there is a randomized poly($n^{\log n}$)-time algorithm that either:*

- *finds a formula (resp. ABP) of size poly($n^{\log n}$) computing a non-zero $u \in (\mathbb{F}[\overline{x}])^n$ such that $Mu = 0$, or*

- *outputs 0 which declares that $u = 0$ is the only solution.*

Proof. This was proved in [KSS15, Lem.2.6] for the circuit model. Since we are using a different model we repeat the details. The idea is the following. Iteratively, for every $r = 1, \ldots, n$ we shall find an $r \times r$ minor contained in the first r columns that is full rank. While continuing this process, we either reach $r = n$ in which case it means that the matrix has full column rank, hence, $u = 0$ is the only solution, or we get stuck at some value say $r = r_0$. We use the fact that r_0 is rank and using this minor we construct the required non-zero vector u.

We explain the process in a bit more detail. Using a randomized algorithm, we look for some non-zero entry in the first column. If no such entry is found we can simply take $u = (1, 0, \ldots, 0)$. So assume that such a non-zero entry is found. After permuting the rows we can assume wlog that this is $M_{1,1}$. Thus, we have found a 1×1 minor satisfying the requirements. Assume that we have found an $r \times r$ full rank minor that is composed of the first r rows and columns (we can always rearrange and hence it can be assumed wlog that they correspond to first r rows and columns). Denote this minor by M_r.

Now for every $(r + 1) \times (r + 1)$ submatrix of M contained in the first $r + 1$ columns and containing M_r, we check whether the determinant is 0 by randomized algorithm. If any of these submatrices have nonzero determinant, then we pick one of them and call it M_{r+1}. Otherwise, we have found that first $r + 1$ columns of M are linearly dependent. As M_r is full rank, there is $v \in \mathbb{F}(\overline{x})^r$ such that $M_r v = (M_{1,r+1}, \ldots, M_{r,r+1})^T$. This can be solved by applying Cramer's rule. The i-th entry of v is of the form $\det(M_r^{(i)})/\det(M_r)$, where $M_r^{(i)}$ is obtained by replacing i-th column of M_r with $(M_{1,r+1}, \ldots, M_{r,r+1})^T$. Observe that $\det(M_r)$, as well as $\det(M_r^{(i)})$, are both in $\mathbb{F}[\overline{x}]$.

Then it is immediate that $u := (\det(M_r^{(1)}), \ldots, \det(M_r^{(r)}), -\det(M_r), 0, \ldots, 0)^T$ is the desired vector.

To find M_r, each time we have to calculate the determinant and decide whether it is 0 or not. This is simply PIT for a determinant polynomial with entries of algebraic complexity $n^{O(\log n)}$ and degree $n^{O(1)}$. So, we have a comparable randomized algorithm for this. Determinant of a symbolic $n \times n$ matrix has $n^{O(\log n)}$ size formula (resp. poly(n) ABP) [MV97]. When the entries of the matrix have $n^{O(\log n)}$ size formula (resp. ABP), altogether, the determinant polynomial has the same algebraic complexity. There are $< n^2$ PIT invocations to test zeroness of the determinant. Altogether, we have a poly($n^{\log n}$)-time randomized algorithm for this [Sch80]. $\square$

A.3 Basic operations on formula, ABP and circuit

We use the following standard results on size bounds for performing some basic operations (like taking derivative) of circuits, formulas, ABPs.

Lemma 20. *(Eliminate single division [Str73], [SY10, Thm.2.1]) Let f and g be two degree-D polynomials, each computed by a circuit (resp. ABP resp. formula) of size-s with $g(\overline{0}) \neq 0$. Then $f/g \mod \langle \overline{x} \rangle^{d+1}$ can be computed by $O((s + d)d^3)$ (resp. $O(sd^2D)$ resp. $O(sd^2D^2)$) size circuit (resp. ABP resp. formula).*

Proof. Assume wlog that $g(\overline{0}) = 1$; we can ensure this by appropriate normalization. So, we have the following power series identity in $\mathbb{F}[[\overline{x}]]$:

$$f/g = f/(1 - (1 - g)) = f + f(1 - g) + f(1 - g)^2 + f(1 - g)^3 + \cdots .$$

Note that this is a valid identity as $1 - g$ is constant free. For all $d \geq 0$, LHS=RHS $\mathrm{mod} \langle \overline{x} \rangle^{d+1}$.

If we want to compute f/g mod $\langle \overline{x} \rangle^{d+1}$, we can take the RHS of the above identity up to the term $f(1 - g)^d$ and discard the remaining terms of degree greater than d. The degree$> d$ monomials can be truncated, using Strassen's *homogenization* trick, in the case of circuits and ABPs (see [Sap16, Lem.5.2]), and an *interpolation* trick in the case of formulas (which also works for ABPs and low degree circuits, [Sap16, Lem.5.4]). A careful analysis shows that the size blow up is at most $O((s + d)d^2 \cdot d)$ (resp. $O(sd \cdot D \cdot d)$ resp. $O(sd \cdot D^2 \cdot d)$) for circuits (resp. ABP resp. formula).

Using the above result, it is easy to see, that we get poly(s, d) size circuit (resp. ABP resp. formula) for computing f/g mod $\langle \overline{x} \rangle^{d+1}$. □

Remark. Note that it may happen that $g(\overline{0}) = 0$, thus $1/g$ does not exist in $\mathbb{F}[[\overline{x}]]$, yet f/g may be a polynomial of degree d. In such a case, we need to discuss a modified *normalization* that works. We can shift the polynomials f, g by some random $\overline{\alpha} \in \mathbb{F}^n$. The constant term of the shifted polynomial is non-zero with high probability [Sch80]. Now, we compute $f(\overline{x}+\overline{\alpha})/g(\overline{x}+\overline{\alpha})$ using the method described above. Finally, we recover the polynomial f/g by applying the reverse shift $\overline{x} \mapsto \overline{x} - \overline{\alpha}$.

What if our model has several division gates?

Lemma 21. *(Div. gates elimination [SY10, Thm.2.12]) Let f be a polynomial computed by a circuit (resp. formula), using division gates, of size s. Then, f mod $\langle \overline{x} \rangle^{d+1}$ can be computed by poly(sd) size circuit (resp. formula).*

Proof idea. We preprocess the circuit (resp. formula) so that the only division gate used in the modified circuit (resp. formula) is at the top. Now to remove the single division gate at the top, we use the above power series trick.

The idea of the pre-processing is the following. We can separately keep track of numerator and denominator computed at each gate and simulate addition, multiplication and division gates in the original circuit. This pre-processing incurs only poly(sd) blow up in the case of circuits. In the case of formulas one has to ensure that in any path from the leaf to the root, there are only $O(\log sd)$ division gates. □

Lemma 22 (Derivative computation)**.** *If a polynomial $f(\overline{x}, y)$ can be computed by a circuit (resp. formula resp. ABP) of size s and degree d. Then, any $\frac{\partial^k f}{\partial y^k}$ can be computed by circuit (resp. formula resp. ABP) of size poly(sk).*

Proof. The idea is simply to use the homogenization and interpolation properties [Sap16, Sec.5.1-2].

Let $f(\overline{x}, y) = c_0 + c_1 y + c_2 y^2 + \ldots + c_\delta y^\delta$, where $c_0, c_1, \ldots, c_\delta \in \mathbb{F}[\overline{x}]$. Given the circuit (resp. formula resp. ABP) computing polynomial $f(\overline{x}, y)$, we can get the circuits (resp. formula resp. ABP) computing $c_0, \ldots, c_\delta$ using homogenization and interpolation as discussed before. Given $c_0, \ldots, c_\delta$, computing $\frac{\partial^k f}{\partial y^k}$ in size poly(sd) is trivial. We use this approach of computing derivative when the polynomial is of degree $d \leq$ poly(s).

In the case of high degree circuits, we cannot use the above approach. [Kal87, Thm.1] shows that $\frac{\partial^k f}{\partial y^k}$ can be computed by a circuit of size $O(k^2 s)$, i.e. the degree of the circuit does not matter. The main idea is to inductively use the Leibniz product rule of k-th order derivative. □

A.4 Sylvester matrix & resultant

First, let us look at the notion of resultant of two univariate polynomials. Let $p(x), q(x) \in \mathbb{F}[x]$ be of degree a, b respectively. From Euclid's extended algorithm, it can be shown that there exist two polynomials $u(x), v(x) \in \mathbb{F}[x]$ such that $u(x)p(x) + v(x)q(x) = \gcd(p(x), q(x))$. This is known as Bezout's identity. If $\gcd(p(x), q(x)) = 1$, then (u, v) with $\deg(u) \leq b$ and $\deg(v) \leq a$ is unique. Let $u(x) = u_0 + u_1 x + u_2 x^2 + \ldots + u_b x^b$ and $v(x) = v_0 + v_1 x + \ldots + v_a x^a$.

Now, if we use the equation $u(x)p(x) + v(x)q(x) = \gcd(p(x), q(x))$ and compare the coefficients of x^i, for $0 \leq i \leq a + b$, we get a system of linear equations in the $a + b + 2$ many unknowns (u_i's and v_i's). The system of linear equations can be represented in the matrix form as $Mx = y$, where x consists of the unknowns. Resultant of f, g is defined as the determinant of the matrix M. It is easy to see that M is invertible if and only if the polynomials are coprime.

Now, the notion of resultant can be extended to multivariate, by defining resultant of polynomials $f(\overline{x}, y)$ and $g(\overline{x}, y)$ wrt some variable y. The idea is same as before, now we take gcd wrt the variable y and get a system of linear equations from Bezout's identity. The matrix can be explicitly written with entries being polynomial coefficients (or they could be from $\mathbb{F}[[\overline{x}]]$). This is known as Sylvester matrix, which we define next.

Definition 23. *Let $f(\overline{x}, y) = \sum_{i=0}^{l} f_i(\overline{x}) y^i$ and $g(\overline{x}, y) = \sum_{i=0}^{m} g_i(\overline{x}) y^i$. Define Sylvester matrix of f and g wrt y as the following $(m + l + 1) \times (m + l + 1)$ matrix:*

$$
Syl_y(f, g) := \begin{bmatrix}
f_l & 0 & 0 & \ldots & 0 & g_m & 0 & 0 & 0 \\
f_{l-1} & f_l & 0 & \ldots & 0 & g_{m-1} & g_m & 0 & 0 \\
f_{l-2} & f_{l-1} & f_l & \ldots & 0 & g_{m-2} & g_{m-1} & g_l & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
f_0 & f_1 & \ldots & \ldots & f_l & g_0 & g_1 & \ldots & g_m \\
0 & f_0 & \ldots & \ldots & \ldots & 0 & g_0 & \ldots & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
0 & \ldots & \ldots & \ldots & f_0 & 0 & \ldots & \ldots & g_0
\end{bmatrix}
$$

So, resultant can be formally defined as follows (for more details and alternate definitions, see [LN97, Chap.1]).

Definition 24. *Given two polynomials $f(\overline{x}, y)$ and $g(\overline{x}, y)$, define the resultant of f and g wrt y as determinant of the Sylvester matrix,*

$$
Res_y(f, g) := det(Syl_y(f, g)).
$$

From the definition, it can be seen that $Res_y(f, g)$ is a polynomial in $\mathbb{F}[x]$ with degree bounded by $2\deg(f)\deg(g)$. Now, we state the following fundamental property of the Resultant, which is crucially used.

Proposition 2 (Res vs gcd). *   1. Let $f, g \in \mathbb{F}[\overline{x}, y]$ be polynomials with positive degree in y. Then, $Res_y(f, g) = 0 \iff f$ and g have a common factor in $\mathbb{F}[\overline{x}, y]$ which has positive degree in y.*

2. There exists $u, v \in \mathbb{F}[\overline{x}]$ such that $uf + vg = Res_y(f, g)$.

The proof of this standard proposition can be found in many standard books on algebra including [vzGG13, Sec.6].

Lemma 25 (Squarefree-ness). *Let $f \in \mathbb{F}(\overline{x})[y]$ be a polynomial with $deg_y(f) \geq 1$. f is square free iff $f, f' := \partial_y f$ are coprime wrt y.*

Proof. The main idea is to show that there does not exist $g \in \mathbb{F}(\overline{x})[y]$ with positive degree in y such that $g \mid \gcd_y(f(\overline{x}, y), f'(\overline{x}, y))$. This is true because– suppose g is an irreducible polynomial with positive degree in y that divides both $f(\overline{x}, y)$ and $f'(\overline{x}, y)$. So,

$$f(\overline{x}, y) = gh \implies f'(\overline{x}, y) = gh' + g'h \implies g \mid g'h.$$

As g is irreducible and $\deg_y(g') < \deg_y(g)$ we deduce that $g \mid h$. Hence, $g^2 \mid f$. This contradicts the hypothesis that f is square free. $\square$

Now, we state another standard lemma, which is useful to us and which is proved using the property of Resultant.

Lemma 26 (Coprimality)**.** *Let $f, g \in \mathbb{F}(\overline{x})[y]$ be coprime polynomials wrt y (& nontrivial in y). Then, for $\overline{\beta} \in_r \mathbb{F}^n$, $f(\overline{\beta}, y)$ and $g(\overline{\beta}, y)$ are coprime (& nontrivial in y).*

Proof. Consider $f = \sum_{i=1}^{d} f_i y^i$ and $g = \sum_{i=1}^{e} g_i y^i$. Choose a random $\overline{\beta} \in_r \mathbb{F}^n$. Then, by Proposition 2 & [Sch80], $f_d \cdot g_e \cdot \mathrm{Res}_y(f, g)$ at $\overline{x} = \overline{\beta}$ is nonzero. This in particular implies that $\mathrm{Res}_y(f(\overline{\beta}, y), g(\overline{\beta}, y)) \neq 0$.

This implies, by Proposition 2, $f(\overline{\beta}, y)$ and $g(\overline{\beta}, y)$ are coprime. $\square$

B Useful in Section 3

Lemma 27. *(Power series root [BCS13, Thm.2.31]) Let $P(\overline{x}, y) \in \mathbb{F}(\overline{x})[y]$, $P'(\overline{x}, y) = \frac{\partial P(\overline{x}, y)}{\partial y}$ and $\mu \in \mathbb{F}$ be such that $P(\overline{0}, \mu) = 0$ but $P'(\overline{0}, \mu) \neq 0$. Then, there is a unique power series S such that $S(\overline{0}) = \mu$ and $P(\overline{x}, S) = 0$ i.e.*

$$y - S(\overline{x}) \mid P(\overline{x}, y).$$

Moreover, there exists a rational function y_t, $\forall t \geq 0$, such that

$$y_{t+1} = y_t - \frac{P(\overline{x}, y_t)}{P'(\overline{x}, y_t)} \text{ and } S \equiv y_t \bmod \langle \overline{x} \rangle^{2^t} \text{ with } y_0 = \mu.$$

Proof. We give an inductive proof of existence and uniqueness together. Suppose $P = \sum_{i=0}^{d} c_i y^i$. We show that there is y_t, a rational function $\frac{A_t}{B_t}$ such that $y_t \in \mathbb{F}[[\overline{x}]]$, For all $t \geq 0$, $P(\overline{x}, y_t) \equiv 0 \bmod \langle \overline{x} \rangle^{2^t}$ and for all $t \geq 1$, $y_t \equiv y_{t-1} \bmod \langle \overline{x} \rangle^{2^{t-1}}$. The proof is by induction. Let $y_0 := \mu$. Thus, base case is true. Now suppose such y_t exists. Define $y_{t+1} := y_t - \frac{P(\overline{x}, y_t)}{P'(\overline{x}, y_t)}$.

Now, $y_t \equiv y_{t-1} \bmod \langle \overline{x} \rangle^{2^{t-1}} \implies y_t(\overline{0}) = \mu$. Hence $P'(\overline{x}, y_t)|_{\overline{x}=\overline{0}} = P'(\overline{0}, \mu) \neq 0$ and so $P'(\overline{x}, y_t)$ is a unit in the power series ring. So, $y_{t+1} \in \mathbb{F}[[\overline{x}]]$. Let us verify that it is an improved root of P; we use Taylor expansion.

$$
\begin{aligned}
P(\overline{x}, y_{t+1}) &= P\left(\overline{x}, y_t - \frac{P(\overline{x}, y_t)}{P'(\overline{x}, y_t)}\right) \\
&= P(\overline{x}, y_t) - P'(\overline{x}, y_t)\frac{P(\overline{x}, y_t)}{P'(\overline{x}, y_t)} + \frac{P''(\overline{x}, y_t)}{2!}\left(\frac{P(\overline{x}, y_t)}{P'(\overline{x}, y_t)}\right)^2 - \ldots \\
&= 0 \bmod \langle \overline{x} \rangle^{2^{t+1}}.
\end{aligned}
$$

Thus, $P(\overline{x}, y_{t+1}) \equiv 0 \bmod \langle \overline{x} \rangle^{2^{t+1}}$ and $y_{t+1} \equiv y_t \bmod \langle \overline{x} \rangle^{2^t}$. This completes the induction step.

Moreover, using the notion of limit, we have $\lim_{t \to \infty} y_t = S$, a formal power series. It is unique as μ is a non-repeated root of $P(\overline{0}, y)$. In particular, we get that for all $t \geq 0$, $P(\overline{x}, S) = 0$ or $y - S \mid P$. $\square$

Lemma 28 (Transform to monic). *For a polynomial $f(\overline{x})$ of total degree $d \geq 0$ and random $\alpha_i \in_r \mathbb{F}$, the transformed polynomial $g(\overline{x}, y) := f(\overline{\alpha}y + \overline{x})$ has a nonzero constant as coefficient of y^d, and degree wrt y is d.*

Proof. Suppose the transformation is $x_i \mapsto x_i + \alpha_i y$ where $i \in [n]$. Write $f = \sum_{|\overline{\beta}|=d} c_{\overline{\beta}} \overline{x}^{\overline{\beta}} +$ lower degree terms . Coefficient of y^d in g is $\sum_{|\overline{\beta}|=d} c_{\overline{\beta}} \overline{\alpha}^{\overline{\beta}}$. Clearly, for a random $\overline{\alpha}$ this coefficient will not vanish [Sch80], and it is the highest degree monomial in g.

This ensures $\deg_y(g) = \deg(f) = d$ and that g is monic wrt y. $\qquad\square$

C Useful in Section 4

Lemma 29 (Matrix inverse). *Let $\mu_i, i \in [d]$, be distinct nonzero elements in $\mathbb{F}$. Define a $d \times d$ matrix A with the (i,j)-th entry $1/(y_i - \mu_j)^2$. Its entries are in the function field $\mathbb{F}(\overline{y})$. Then, $\det(A) \neq 0$.*

Proof. The idea is to consider the power series of the function $1/(y_i - \mu_j)^2$ and show that a monomial appears nontrivially in that of $\det(A)$.

We first need a claim about the coefficient operator on the determinant.

Claim 30. *Let $f_j = \sum_{i \geq 0} \beta_{j,i} x^i$ be a power series in $\mathbb{F}[[x]]$, for $j \in [d]$. Then, $\mathrm{Coeff}_{\overline{x}^{\overline{\alpha}}} \circ \det(f_j(x_i)) = \det(\beta_{j,\alpha_i})$.*

Proof of Claim 30. Observe that the rows of the matrix have disjoint variables. Thus, $x_i^{\alpha_i}$ could be produced only from the i-th row. This proves: $\mathrm{Coeff}_{\overline{x}^{\overline{\alpha}}} \circ \det(f_j(x_i)) = \det\left(\mathrm{Coeff}_{x_i^{\alpha_i}} \circ f_j(x_i)\right) = \det(\beta_{j,\alpha_i})$. $\qquad\square$

By Taylor expansion we have

$$\frac{1}{(x-\mu)^2} = \frac{1}{\mu^2} \sum_{j \geq 1} j \left(\frac{x}{\mu}\right)^{j-1} .$$

Hence, the coefficient of y_i^{i-1} in $A(i,j)$ is

$$\frac{1}{\mu_j^2} \frac{i}{\mu_j^{i-1}} = \frac{i}{\mu_j^{i+1}} .$$

By the above claim, the coefficient of $\prod_{i \in [d]} y_i^{i-1}$ in $\det(A)$ is: $\det\left(\left(\frac{i}{\mu_j^{i+1}}\right)\right)$. By cancelling i (from each row) and $1/\mu_j^2$ (from each column), we simplify it to the Vandermonde determinant:

$$\det \begin{bmatrix} \frac{1}{\mu_1^0} & \frac{1}{\mu_2^0} & \cdots & \frac{1}{\mu_d^0} \\ \frac{1}{\mu_1^1} & \frac{1}{\mu_2^1} & \cdots & \frac{1}{\mu_d^1} \\ \vdots & \vdots & \cdots & \vdots \\ \frac{1}{\mu_1^{d-1}} & \frac{1}{\mu_2^{d-1}} & \cdots & \frac{1}{\mu_d^{d-1}} \end{bmatrix} = \prod_{i<j \in [d]} \left(\frac{1}{\mu_i} - \frac{1}{\mu_j}\right) \neq 0 .$$

Hence, the determinant of A is non-zero. $\qquad\square$

Remark. If the characteristic of $\mathbb{F}$ is a prime $p \geq 2$ then the above proof needs a slight modification. One should consider the coefficient of $\prod_{i \in [d]} y_i^{s_i-1}$ in $\det(A)$ for a set $S = \{s_1, \ldots, s_d\}$ of distinct non-negative integers that are not divisible by p. Moreover, one has to consider 'random' μ_i's to deduce $\det(A) \neq 0$.

Lemma 31 (Series inverse). *Let $\delta \geq 1$. Assume that A is a polynomial of degree $< \delta$ and B is a homogeneous polynomial of degree δ, such that $A(\overline{0}) =: \mu \neq 0$. Then, we have the following identity in $\mathbb{F}[[\overline{x}]](y) \cap \mathbb{F}[[\overline{x}]][[y]]$:*

$$\frac{1}{y - (A + B)} \equiv \frac{1}{y - A} + \frac{B}{(y - \mu)^2} \mod \langle \overline{x} \rangle^{\delta+1}$$

Proof. We will use the notation $A^{[1,\delta-1]}$ to refer to the sum of the homogeneous parts of A of degrees between 1 and $\delta - 1$ (equivalently, it is $A^{<\delta} - \mu$). Note that $B \cdot A^{[1,\delta-1]}$ vanishes mod $\langle \overline{x} \rangle^{\delta+1}$. Now, in $\mathbb{F}[[\overline{x}]][[y]]$,

$$\frac{1}{y - (A + B)} \equiv \frac{1}{y - \mu - (A^{[1,\delta-1]} + B)} \mod \langle \overline{x} \rangle^{\delta+1}$$

$$\equiv \frac{1}{y - \mu} \left(\frac{1}{1 - \frac{A^{[1,\delta-1]} + B}{y - \mu}} \right) \mod \langle \overline{x} \rangle^{\delta+1}$$

$$\equiv \frac{1}{y - \mu} \left(1 + \left(\frac{A^{[1,\delta-1]} + B}{y - \mu} \right) + \left(\frac{A^{[1,\delta-1]} + B}{y - \mu} \right)^2 + \ldots \ldots \right) \mod \langle \overline{x} \rangle^{\delta+1}$$

$$\equiv \frac{1}{y - \mu} \left(1 + \left(\frac{A^{[1,\delta-1]} + B}{y - \mu} \right) + \left(\frac{A^{[1,\delta-1]}}{y - \mu} \right)^2 + \left(\frac{A^{[1,\delta-1]}}{y - \mu} \right)^3 + \ldots \ldots \right) \mod \langle \overline{x} \rangle^{\delta+1}$$

$$\equiv \frac{1}{y - \mu} \left(1 + \left(\frac{A^{[1,\delta-1]}}{y - \mu} \right) + \left(\frac{A^{[1,\delta-1]}}{y - \mu} \right)^2 + \ldots \ldots \right) + \frac{B}{(y - \mu)^2} \mod \langle \overline{x} \rangle^{\delta+1}$$

$$\equiv \frac{1}{y - \mu} \left(\frac{1}{1 - \frac{A^{[1,\delta-1]}}{y - \mu}} \right) + \frac{B}{(y - \mu)^2} \mod \langle \overline{x} \rangle^{\delta+1}$$

$$\equiv \frac{1}{y - A} + \frac{B}{(y - \mu)^2} \mod \langle \overline{x} \rangle^{\delta+1} \quad .$$

$\square$

C.1 Closure properties for VNP

VNP-size parameter (w, v) of F refers to w being the witness size and v being the size of the verifier circuit f.

Let $F(\overline{x}, y), G(\overline{x}, y), H(\overline{x})$ have verifier polynomials f, g, h and the VNP size parameters $(w_f, v_f), (w_g, v_g), (w_h, v_h)$ respectively. Let the degree of F wrt y be d. Then, the following closure properties can be shown ([BCS13] or [Bür13, Thm.2.19]):

1. Add (resp. Multiply): $F + G$ (resp. FG) has VNP-size parameter $(w_f + w_g, v_f + v_g + 3)$.

2. Coefficient: $F_i(\overline{x})$ has VNP-size parameter $(w_f, (d+1)(v_f+1))$, where $F(\overline{x}, y) =: \sum_{i=0}^{d} F_i(\overline{x}) y^i$.

3. Compose: $F(\overline{x}, H(\overline{x}))$ has VNP-size parameter $((d+1)(w_f + dw_h), (d+1)^2(v_f + v_h + 1))$.

Proof. All the above statements are easy to prove using the definition of VNP.

1. $(FG)(\overline{x}, y) = \left(\sum_{u \in \{0,1\}^{w_f}} f(\overline{x}, u_1, \ldots, u_{w_f}) \right) \cdot \left(\sum_{u \in \{0,1\}^{w_g}} g(\overline{x}, u_1, \ldots, u_{w_g}) \right) = \sum_{u \in \{0,1\}^{w_f + w_g}} A(\overline{x}, u_1, \ldots, u_{w_f + w_g})$. Where, $A(\overline{x}, u_1, \ldots, u_{w_f + w_g}) := f(\overline{x}, u_1, \ldots, u_{w_f}) \cdot$

$g(\overline{x}, u_{w_f+1}, \ldots, u_{w_f+w_g})$. Trivially, A has size $v_f + v_g + 3$ (extra: one node, two edges) and witness size is $w_f + w_g$.

Similarly, with $F + G$.

2. Interpolation gives, $f_i(\overline{x}) = \sum_{j=0}^{d} \alpha_j F(\overline{x}, \beta_j)$, for some distinct arguments $\beta_j \in \mathbb{F}$. Clearly, $F(\overline{x}, \beta_j)$ has VNP-size parameter (w_f, v_f). Using the previous addition property we get that the verifier circuit has size $(d+1)(v_f + 1)$. Witness size remains w_f as we can reuse the witness string of F.

3. Write $F(\overline{x}, y) =: \sum_{i=0}^{d} F_i(\overline{x}) y^i$. We know that F_i has VNP-size parameter $(w_f, (d+1)(v_f + 1))$. For $0 \le i \le d$, H^i has VNP-size parameter $(iw_h, (i+1)v_h)$ using i-fold product (Item 1). Substituting $y = H$ in F, we can calculate the VNP-size parameter.

Suppose F_i and H^i have corresponding verifier circuits A_i and B_i respectively. Then, $F(\overline{x}, H(\overline{x})) = \sum_{i=0}^{d} F_i(\overline{x}) H^i(\overline{x}) = \sum_{i=0}^{d} \left(\sum_{u \in \{0,1\}^{w_f}} A_i(\overline{x}, u) \right) \cdot \left(\sum_{u \in \{0,1\}^{iw_h}} B_i(\overline{x}, u) \right)$. Thus, the witness size is $< (d+1)(w_f + dw_h)$. The corresponding verifier circuit size is $< (d+1)^2 (v_f + v_h + 1)$.

$\square$