

Graph Colouring Is Hard on Average for Polynomial Calculus and Nullstellensatz

Jonas Conneryd
Lund University
University of Copenhagen
jonas.conneryd@cs.lth.se

Susanna F. de Rezende
Lund University
susanna.rezende@cs.lth.se

Jakob Nordström
University of Copenhagen
Lund University
jn@di.ku.dk

Shuo Pang
University of Copenhagen
Lund University
shpa@di.ku.dk

Kilian Risse
EPFL
kilian.risse@epfl.ch

Abstract—We prove that polynomial calculus (and hence also Nullstellensatz) over any field requires linear degree to refute that sparse random regular graphs, as well as sparse Erdős-Rényi random graphs, are 3-colourable. Using the known relation between size and degree for polynomial calculus proofs, this implies strongly exponential lower bounds on proof size.

I. INTRODUCTION

Determining the *chromatic number* of a graph G , i.e., how many colours are needed for the vertices of G if no two vertices connected by an edge should have the same colour, is one of the classic 21 problems shown NP-complete in the seminal work of Karp [35]. This *graph colouring problem*, as it is also referred to, has been extensively studied since then, but there are still major gaps in our understanding.

The currently best known approximation algorithm computes a graph colouring within at most a factor $O(n(\log \log n)^2/(\log n)^3)$ of the chromatic number [31], and it is known that approximating this number to within a factor $n^{1-\varepsilon}$ is NP-hard [55]. Even under the promise that the graph is 3-colourable, the most parsimonious algorithm with guaranteed polynomial running time needs $O(n^{0.19996})$ colours [36]. This is very far from the lower bounds that are known—it is NP-hard to $(2k-1)$ -colour a k -colourable graph [10], but the question of whether colouring a 3-colourable graph with 6 colours is NP-hard remains open [39]. It is widely believed that any algorithm that colours graphs optimally has to run in exponential time in the worst case, and the currently fastest algorithm for 3-colouring has time complexity $O(1.3289^n)$ [15]. A survey on various algorithms and techniques for so-called exact algorithms for graph colouring can be found in [33].

Graph colouring instances of practical interest might not exhibit such exponential-time behaviour, however, and in such a context it is relevant to study algorithms without worst-case guarantees and examine how they perform in practice. To understand such algorithms from a computational complexity viewpoint, it is natural to investigate bounded models of computation that are strong enough to describe the reasoning

performed by the algorithms and to prove unconditional lower bounds that hold in these models.

A. Previous Work

Focusing on random graphs, McDiarmid [46] developed a method for determining k -colourability that captures a range of algorithmic approaches. Beame et al. [12] showed that this method could in turn be simulated by the resolution proof system [17], [25], [24], [52], and established average-case exponential lower bounds for resolution proofs of non- k -colourability for random graph instances sampled so as not to be k -colourable with exceedingly high probability.

Different algebraic approaches for k -colourability have been considered in [4], [43], [44], [45]. Bayer [11] seems to have been the first to use Hilbert's Nullstellensatz to attack graph colouring. Informally, the idea is to write the problem as a set of polynomial equations $\{p_i(x_1, \dots, x_n) = 0 \mid i \in [m]\}$ in such a way that legal k -colourings correspond to common roots for these polynomials. Finding polynomials q_1, \dots, q_m such that $\sum_{i=1}^m q_i p_i = 1$ then proves that the graph is not k -colourable. This latter equality is referred to as a *Nullstellensatz certificate* of non-colourability, and the *degree* of this certificate is the largest degree of any polynomial $q_i p_i$ in the sum. Later papers based on Nullstellensatz and Gröbner bases, such as [26], [48], [32], culminated in an award-winning sequence of works [28], [30], [29], [27] presenting algorithms with surprisingly good practical performance.

For quite some time, no strong lower bounds were known for these algebraic methods or the corresponding proof systems *Nullstellensatz* [13] and *polynomial calculus* [21], [2]. On the contrary, the authors of [30] reported that essentially all benchmarks they studied turned out to have Nullstellensatz certificates of small constant degree. The degree lower bound $k+1$ for k colours in [27] remained the best known until optimal, linear, degree lower bounds for polynomial calculus were established in [42] using a reduction from lower bounds for so-called functional pigeonhole principle formulas [47]. A more general reduction framework was devised in [7] to

obtain optimal degree lower bounds also for the proof systems *Sherali-Adams* [54] and *sums-of-squares* [40], [49], as well as weakly exponential size lower bounds for *Frege proofs* [23], [51] of bounded depth.

The lower bounds discussed in the previous paragraph are not quite satisfactory, in that it is not clear how much they actually tell us about the graph colouring problem, as opposed to the hardness of the problems being reduced from. In order to improve our understanding for a wider range of graph instances, it seems both natural and desirable to establish average-case lower bounds for random graphs, just as for resolution in [12]. However, this goal has remained elusive for almost two decades, as pointed out, e.g., in [47], [42], [41], [20]. For sparse random graphs, where the number of edges is linear in the number of vertices, no superconstant degree lower bounds at all have been established for algebraic or semialgebraic proof systems. On the contrary, it was shown in [9], improving on [22], that degree-2 sums-of-squares refutes k -colourability on random d -regular graphs asymptotically almost surely whenever $d \geq 4k^2$. For dense random graphs, the strongest lower bound seems to be the recent logarithmic degree bound in the sums-of-squares proof system for Erdős-Rényi random graphs with edge probability $1/2$ and $k = n^{1/2+c}$ colours [38]. Since this result is for a problem encoding using inequalities, however, it is not clear whether this has any implications for Nullstellensatz or polynomial calculus over the reals (which are known to be polynomially simulated by sums-of-squares). And for other fields nothing has been known for the latter two proof systems—not even logarithmic lower bounds.

B. Our Contribution

In this work, we establish optimal, linear, degree lower bounds and exponential size lower bounds for polynomial calculus proofs of non-colourability of random graphs.

Theorem 1 (informal). *For any $d \geq 6$, polynomial calculus (and hence also Nullstellensatz) requires asymptotically almost surely linear degree to refute that random d -regular graphs $\mathbb{G}_{n,d}$, as well as Erdős-Rényi random graphs $\mathbb{G}(n, d/n)$, are 3-colourable. These degree lower bounds hold over any field, and also imply exponential lower bounds on proof size.*

We prove our lower bound for the standard encoding in proof complexity, where binary variables $x_{v,i}$ indicate whether vertex v is coloured with colour i or not. It should be pointed out that, just as the results in [42], our degree lower bounds also apply to the k -colourability encoding introduced in [11] and used in computational algebra papers such as [28], [30], [29], [27], where a primitive k th root of unity is adjoined to the field and different colours of a vertex v are encoded by a variable x_v taking different powers of this root of unity.

Our lower bound proofs crucially use a new idea for proving degree lower bounds for colouring graphs with large girth [53]. After adapting this approach from the root-of-unity encoding to the Boolean indicator variable encoding, and replacing the proof in terms of girth with a strengthened argument using

carefully chosen properties of random graphs, we obtain a remarkably clean and simple solution to the long-standing open problem of showing average-case polynomial calculus degree lower bounds for graph colouring. We elaborate on our techniques in slightly more detail next.

C. Discussion of Proof Techniques

In most works on algebraic and semialgebraic proof systems such as Nullstellensatz, polynomial calculus, Sherali-Adams, and sums-of-squares, the focus has been on proving upper and lower bounds on the degree of proofs. Even when proof size is the measure of interest, almost all size lower bounds have been established via degree lower bounds combined with general results saying that for all of the above proof systems except Nullstellensatz strong enough lower bounds on degree imply lower bounds on size [34], [6].

At a high level, the techniques for proving degree lower bounds for the different proof systems have a fairly similar flavour. For the static proof systems, i.e., Nullstellensatz, Sherali-Adams, and sums-of-squares, it is enough to show that the dual problem is feasible and thus rule out low-degree proofs. In more detail, for Nullstellensatz, one constructs a *design* [19], which is a linear functional mapping low-degree monomials to elements in the underlying field. This functional should map low-degree monomials multiplied by any input polynomial p_i to 0, but should map 1 to a non-zero field element. If such a functional can be found, it is clear that there cannot exist any low-degree Nullstellensatz certificate $\sum_{i=1}^m q_i p_i = 1$ of unsatisfiability, as the design would map the left-hand side of the equation to zero but the right-hand side to non-zero. For Sherali-Adams, the analogous functional furthermore has to map any low-degree monomials to non-negative numbers, and for sums-of-squares this should also hold for squares of low-degree polynomials. Such a *pseudo-expectation* can be viewed as a fake probability distribution over satisfying assignments to the problem, which is indistinguishable from a true distribution for an adversary using only low-degree polynomials.

Polynomial calculus is different from these proof systems in that it does not present the certificate of unsatisfiability as a static object, but instead, given a set of polynomials \mathcal{P} , dynamically derives new polynomials in the ideal generated by \mathcal{P} . The derivation ends when it reaches the polynomial 1, i.e., the multiplicative identity in the field, showing that there is no solution. To prove degree lower bounds one designs a *pseudo-reduction operator* or *R-operator* [50], which maps all low-degree polynomials derived from \mathcal{P} to 0 but sends 1 to 1, and which is indistinguishable from a true ideal reduction operator if one is limited to reasoning with low-degree polynomials. This means that for bounded-degree polynomial calculus derivations it seems like the set of input polynomials are consistent.

Following the method in [3], a pseudo-reduction operator R can be constructed by defining it on low-degree monomials and extending it to polynomials by linearity. For every monomial m , we identify a set of related input polynomials $S(m)$, let $\langle S(m) \rangle$ be the ideal generated by these polynomials, and define $R(m) = R_{\langle S(m) \rangle}(m)$ to be the reduction of m modulo the

ideal $\langle S(m) \rangle$. Intuitively, we think of $S(m)$ as the (satisfiable) subset of polynomials that might possibly have been used in a low-degree derivation of m , but since the constant monomial 1 is not derivable in low degree it gets an empty associated set of polynomials, meaning that $R(1) = R_{\langle S(1) \rangle}(1) = 1$. In order for R to look like a real reduction operator, we need to show that for polynomials p and p' of not too high degree it holds that $R(p + p') = R(p) + R(p')$ and $R(p \cdot R(p')) = R(p \cdot p')$. The first equality is immediate, since R is defined to be a linear operator, but the second equality is more problematic. Since the polynomials p and p' will be reduced modulo different ideals—in fact, this will be the case even for different monomials within the same polynomial—a priori there is no reason why R should interact nicely with multiplication.

Proving that an R -operator behaves like an actual reduction operator for low-degree polynomials is typically the most challenging technical step in the lower bound proof. Very roughly, the proof method in [3] goes as follows. Suppose that m and m' are monomials with associated polynomial sets $S(m)$ and $S(m')$, respectively. Using expansion properties of the constraint-variable incidence graph for the input polynomials, we argue that the true reduction operator will not change if we reduce both monomials modulo the larger ideal $\langle S(m) \cup S(m') \rangle$ generated by the union of their associated sets of polynomials. This implies that we have $R(m') = R_{\langle S(m') \rangle}(m') = R_{\langle S(m) \cup S(m') \rangle}(m')$ and $R(m \cdot m') = R_{\langle S(m) \cup S(m') \rangle}(m \cdot m')$, from which it follows that $R(m \cdot R(m')) = R(m \cdot m')$ holds, just like for reduction modulo an actual ideal. To prove that expanding the ideals does not change the reduction operator is a delicate balancing act, though, since the ideals will need to be large enough to guarantee non-trivial reduction, but at the same time small enough so that different ideals can be “patched together” with only local adjustments.

All previous attempts to apply this lower bound strategy to the graph colouring problem have failed. For other polynomial calculus lower bounds it has been possible to limit the interaction between different polynomials in the input. For graph colouring, however, applying the reduction operator intuitively corresponds to partial colourings of subsets of vertices, and it has not been known how to avoid that locally assigned colours propagate new colouring constraints through the rest of the graph. In technical language, what is needed is a way to order the vertices in the graph so that there will be no long ordered paths of vertices along which colouring constraints can spread. It has seemed far from obvious how to construct such an ordering, or even whether it should exist, and due to this technical problem it has not been possible to join local ideal reduction operators into a globally consistent R -operator.

This technical problem was addressed in a recent paper [53] by an ingenious, and in hindsight surprisingly simple, idea. The main insight is to consider a proper colouring of the graph G with $\chi(G)$ colours, and then order the vertices in each colour class consecutively. In this way, order-decreasing paths are of length at most $\chi(G)$, and one can guarantee some form

of locality. Once this order is in place, the final challenge is to ensure that small cycles do not interfere when “patching together” reductions. In [53], such conflicts are avoided by ensuring that the graph should have high girth, which results in a degree lower bound linear in the girth of the graph. In terms of graph size, this cannot give better than logarithmic lower bounds, however, since the girth is at most logarithmic in the number of vertices for any graph with chromatic number larger than 3 [18].

In our work, we employ the same ordering as in [53], but instead of girth use the fact that random graphs are locally very sparse. Once the necessary technical concepts are in place, the proof becomes quite clean and elegant, which we view as an additional strength of our result.

D. Outline of This Paper

The rest of this paper is organized as follows. In Section II we present some preliminaries. In Section III we introduce our techniques and provide a proof overview. In Section IV we prove the main result of the paper. For simplicity of exposition, in this conference version we only present our results for 4-colourability on random regular graphs. For a full statement of our results with complete proofs, we refer the reader to the upcoming full-length version of this paper. We conclude with some final remarks and open problems in Section V.

II. PRELIMINARIES

Let us start by briefly reviewing the necessary preliminaries from proof complexity, graph theory, and algebra. We use standard asymptotic notation, and all logarithms in this paper have base 2. Given a natural number n , we let $[n]$ denote the set $\{1, 2, \dots, n\}$ and, for a natural number i and a set S , let $\binom{S}{i}$ denote the family of subsets of S of size i .

For a field \mathbb{F} we let $\mathbb{F}[x_1, \dots, x_n]$ denote the polynomial ring over \mathbb{F} in n variables. A *monomial* is a product of variables and a *term* is a monomial multiplied by an element of \mathbb{F} .

A. Proof Complexity

Polynomial calculus (PC) [21] is a proof system that uses algebraic reasoning to deduce that a system \mathcal{P} of polynomials over a field \mathbb{F} involving the variables x_1, \dots, x_n is infeasible, i.e., that the polynomials in \mathcal{P} have no common root. Polynomial calculus interprets \mathcal{P} as a set of generators of an ideal and derives new polynomials in this ideal through two derivation rules:

$$\text{Linear combination: } \frac{p}{ap + bq}, \quad a, b \in \mathbb{F} \quad (1a)$$

$$\text{Multiplication: } \frac{p}{x_i p}, \quad x_i \text{ any variable.} \quad (1b)$$

A *polynomial calculus derivation* π of a polynomial p starting from the set \mathcal{P} is a sequence of polynomials (p_1, \dots, p_τ) , where $p_\tau = p$ and each polynomial p_i either is in \mathcal{P} or is obtained by applying one of the derivation rules (1a)-(1b) to polynomials p_j with $j < i$. A *polynomial calculus refutation* of \mathcal{P} is a derivation of the constant polynomial 1 from \mathcal{P} . We are interested in systems of polynomial equations over

Boolean variables and therefore include the Boolean axioms $\{x_1^2 - x_1, \dots, x_n^2 - x_n\}$ in \mathcal{P} . It is well-known that polynomial calculus is sound and complete when the Boolean axioms are present.

The most common complexity measures of polynomial calculus refutations are *size* and *degree*. The *size* of a polynomial p is its number of monomials when expanded into a linear combination of distinct monomials, and the *degree* of p is the maximum degree among all its monomials. The size of a polynomial calculus refutation π is the sum of the sizes of the polynomials in π , and the degree of π is the maximum degree among all polynomials in π . We follow the convention of not counting applications of the Boolean axioms toward degree or size by tacitly working over $\mathbb{F}[x_1, \dots, x_n]/\langle x_1^2 - x_1, \dots, x_n^2 - x_n \rangle$, which only strengthens a lower bound on either measure. Polynomial calculus size and degree are connected through the *size-degree relation* [34]: if \mathcal{P} consists of polynomials with constant degree and D is the minimal degree among all polynomial calculus refutations of \mathcal{P} , then every refutation of \mathcal{P} must have size $\exp(\Omega(D^2/n))$.

The size-degree relation also applies to the stronger proof system *polynomial calculus resolution (PCR)* [2], which is polynomial calculus where additionally each variable x_i appearing in \mathcal{P} has a formal negation \bar{x}_i , enforced by adding polynomials $x_i + \bar{x}_i - 1$ to \mathcal{P} . Polynomial calculus and PCR are equivalent with respect to degree, since the map $\bar{x}_i \mapsto 1 - x_i$ sends any PCR proof to a valid polynomial calculus proof of the same degree. Therefore, to prove a lower bound on PCR size it suffices to prove a lower bound on polynomial calculus degree, and in particular all size lower bounds in this paper also apply to PCR. Finally, we remark that lower bounds on polynomial calculus degree or size also apply to the weaker *Nullstellensatz* proof system mentioned in Section I-A and Section I-B.

B. Graph Colouring and Polynomial Calculus

Given a graph G , we study the polynomial calculus degree required to refute the system $\text{Col}(G, k)$ of polynomials

$$\sum_{i=1}^k x_{v,i} - 1, \quad v \in V(G) \quad (2a)$$

$$x_{v,i}x_{v,i'}, \quad v \in V(G), \quad i \neq i' \quad (2b)$$

$$x_{u,i}x_{v,i}, \quad (u, v) \in E(G), \quad i \in [k] \quad (2c)$$

$$x_{v,i}^2 - x_{v,i}, \quad v \in V(G), \quad i \in [k] \quad (2d)$$

whose common roots correspond precisely to proper k -colourings of G . We refer to axioms in (2a) and (2b) as *vertex axioms* and to (2c) as *edge axioms*. It is known [42, Proposition 2.2] that a polynomial calculus degree lower bound for $\text{Col}(G, k)$ also applies to *Bayer's formulation* [11] of k -colourability, where each colour corresponds to a k th root of unity. This encoding has received considerable attention in computational algebra [28], [30], [29], [27], [53].

Our proof of Theorem 1 is based on the following lemma due to Razborov [50].

Lemma 2 ([50]). *Let \mathcal{P} be a set of multilinear polynomials over $\mathbb{F}[x_1, \dots, x_n]$ and let D be a positive integer. Suppose there exists an \mathbb{F} -linear operator R over multilinear polynomials with the following properties:*

- 1) $R(1) = 1$.
- 2) $R(p) = 0$ for every polynomial p in \mathcal{P} .
- 3) For every term t of degree at most $D - 1$ and every variable x_i it holds that $R(x_i t) = R(x_i R(t))$.

Then any polynomial calculus refutation of \mathcal{P} over \mathbb{F} requires degree strictly greater than D .

The proof of Lemma 2 is straightforward: apply R to all polynomials in a purported polynomial calculus refutation of \mathcal{P} and conclude by induction that it is impossible to reach contradiction in degree at most D . We call a linear operator satisfying properties 1-3 as stated in Lemma 2 an *R -operator with respect to \mathbb{F}* . Because the operators we consider in this paper satisfy these properties over any field, we will refer to them simply as *R -operators* going forward.

C. Algebra Background

The definition of our R -operator requires some standard notions from algebra, phrased for our setting in, e.g., [47]. A total ordering \prec on the multilinear monomials in $\mathbb{F}[x_1, \dots, x_n]$ is *admissible* if the following properties hold:

- 1) If $\text{Deg}(m_1) < \text{Deg}(m_2)$, then $m_1 \prec m_2$.
- 2) For any monomials m_1, m_2 and m such that $m_1 \prec m_2$ and m shares no variables with m_1 or m_2 , it holds that $mm_1 \prec mm_2$.

We identify the ordering of a term with the ordering of the corresponding monomial and the *leading term* (respectively, *leading monomial*) of a polynomial p is the largest term (respectively, monomial) in p according to \prec . For an ideal I over $\mathbb{F}[x_1, \dots, x_n]/\langle x_1^2 - x_1, \dots, x_n^2 - x_n \rangle$, a term t is *reducible modulo I* if t is the leading term of a polynomial $q \in I$; otherwise t is *irreducible modulo I* . Under a total monomial ordering it is well-known that for any ideal I and any polynomial p there exists a unique representation $p = q + r$ such that $q \in I$ and r is a linear combination of irreducible terms modulo I . We call r the *reduction of p modulo I* , and denote by R_I the *reduction operator* which maps polynomials p to the reduction of p modulo I .

We rely on the following observation whose proof can be found in [47].

Observation 3. *If I_1 and I_2 are ideals and $I_1 \subseteq I_2$, then for every variable x_i and every term t it holds that $R_{I_2}(x_i R_{I_1}(t)) = R_{I_2}(x_i t)$.*

Let us record one final observation. We prove this fact in the full version of this article.

Observation 4. *If g and $\{q_1, \dots, q_m\} = Q$ are polynomials in $\mathbb{F}[x_1, \dots, x_n]/\langle x_1^2 - x_1, \dots, x_n^2 - x_n \rangle$, and g vanishes on all the Boolean common roots of the polynomials in Q , then it holds that $g \in \langle Q \rangle$.*

D. Graph Theory

We only consider graphs $G = (V, E)$ that are finite, undirected and contain no self-loops or multi-edges. Given a subset $U \subseteq V$, the *neighbourhood* of U in G is $N(U) = \{v \in V \setminus U \mid \exists u \in U: (u, v) \in E\}$ and, for a set $W \subseteq V$, we let the *neighbourhood of U in W* be denoted by $N_W(U) = N(U) \cap W$. The set of edges between vertices in U is denoted by $E(U)$, while the set of edges with one endpoint in U and the other in $W \setminus U$ is denoted by $E(U, W)$. We let $G[U]$ denote the subgraph induced by U in G , that is, $G[U] = (U, E(U))$, and a graph is said to be d -regular if all vertices are of degree d . Note that a graph on n vertices can be d -regular only if $d < n$ and dn is even.

A *simple path of length τ* is a tuple of distinct vertices $(v_1, \dots, v_{\tau+1})$ satisfying $(v_i, v_{i+1}) \in E$ for all $i \in [\tau]$ and a *simple cycle of length τ* is a simple path of length τ with the additional requirement that $v_1 = v_{\tau+1}$. A simple path or simple cycle $(v_1, \dots, v_{\tau+1})$ of length $\tau \geq 2$ is a τ -hop with respect to U if the endpoints v_1 and $v_{\tau+1}$ are both contained in U and all other vertices are not contained in U . For a subgraph $H \subseteq G$, we refer to a τ -hop with respect to $V(H)$ as a τ -hop with respect to H . For a set $S \subseteq \mathbb{N}^+$, we say that a τ -hop is an S -hop if $\tau \in S$.

A graph is said to be k -colourable if there is a mapping $\chi: V \rightarrow [k]$ satisfying $\chi(u) \neq \chi(v)$ for all edges $(u, v) \in E$ and we refer to χ as a k -colouring of G . Finally we have the *chromatic number* of G , denoted by $\chi(G)$, which is the smallest integer k such that G is k -colourable.

The graph G is (ℓ, ε) -sparse if every subset $U \subseteq V$ of size at most ℓ satisfies $|E(U)| \leq (1 + \varepsilon)|U|$ and we say that G is an (ℓ, γ) -expander if all subsets $U \subseteq V$ of size at most ℓ satisfy $|E(U, V)| \geq \gamma|U|$. If G is d -regular, then it is not hard to see that sparsity is equivalent to expansion: a d -regular graph is (ℓ, ε) -sparse if and only if it is an $(\ell, d - 2(1 + \varepsilon))$ -expander.

We frequently use that large subsets of (ℓ, ε) -sparse graphs are 3-colourable.

Lemma 5. *If $G = (V, E)$ is (ℓ, ε) -sparse for some $\varepsilon < 1/2$, then it holds for every subset $U \subseteq V$ of size at most ℓ that $G[U]$ is 3-colourable.*

Proof. By induction on $|U|$. The base case $|U| = 1$ is immediate. For the inductive step we may assume that the claim holds for sets of size at most $s - 1$. Consider a set $U \subseteq V$ of size $s \leq \ell$. The average degree of a vertex in $G[U]$ is $2|E(U)|/s$, which is at most $2(1 + \varepsilon) < 3$ by the assumption on sparsity. Hence, since graph degrees are integers, there exists a vertex $v \in U$ with degree at most 2 in $G[U]$. The graph $G[U \setminus \{v\}]$ is 3-colourable by the inductive hypothesis, and every 3-colouring witnessing this will leave at least one colour available to properly colour v . Hence every 3-colouring of $G[U \setminus \{v\}]$ can be extended to $G[U]$, which concludes the proof. \square

We consider two models of random graphs: the *Erdős-Rényi random graph model* $\mathbb{G}(n, p)$ which is the distribution over graphs on n vertices where each edge is independently included

with probability p ; and the *random d -regular graph model* $\mathbb{G}_{n,d}$ which is the uniform distribution over the set of d -regular graphs on n vertices. A graph property P holds *asymptotically almost surely* for a random graph model $\mathbb{G} = \{\mathbb{G}_n\}_{n=1}^\infty$ if $\lim_{n \rightarrow \infty} \Pr_{G \sim \mathbb{G}_n}[G \text{ has property } P] = 1$.

Random graphs are (ℓ, ε) -sparse with excellent parameters, as stated in the following lemma. We provide a proof of this folklore result in the upcoming full version of this paper.

Lemma 6 (Sparsity lemma). *There is a constant $\delta > 0$ such that for integers $n, d \geq 3$ and all $\varepsilon > \delta d^2 / \log n$ satisfying $\varepsilon = \omega(1/\log n)$, the following holds. If G is a graph sampled from $\mathbb{G}_{n,d}$ or $\mathbb{G}(n, d/n)$, then asymptotically almost surely it is $(d^{-30(1+\varepsilon)/\varepsilon} n, \varepsilon)$ -sparse.*

Let us stress that ε may be any (small) function on n and in our applications indeed will depend on the degree d . Finally, we need some bounds on the chromatic number of graphs sampled from $\mathbb{G}(n, d/n)$ or $\mathbb{G}_{n,d}$. The upper bound is used for estimates and the lower bound ensures that $\text{Col}(G, k)$ is unsatisfiable for large enough d .

Lemma 7 ([37], [1]). *For a graph G sampled from either $\mathbb{G}(n, d/n)$ or $\mathbb{G}_{n,d}$ the following holds asymptotically almost surely. The chromatic number $\chi(G)$ is at most $2d/\log d$ and, if $d \geq 10$, then $\chi(G) \geq 5$.*

III. TECHNIQUES AND PROOF OVERVIEW

In this section we introduce the technical tools needed for our lower bounds and provide a proof overview.

A. Closure and Ordering by Colouring

Let $G = (V, E)$ be a graph with a linear ordering \prec on V . An *increasing (respectively, decreasing) path* in G is a simple path (v_1, \dots, v_τ) where $v_i \prec v_{i+1}$ (respectively, $v_{i+1} \prec v_i$) for all $i \in [\tau - 1]$. For vertices u, v in V we say that u is a *descendant* of v if there exists a decreasing path from v to u , and for a set of vertices $U \subseteq V$ we let D_U be the set of all the vertices which are a descendant of some vertex in U . We denote by $\text{Desc}(U)$ the *descendant graph* of U which is the subgraph induced by $U \cup D_U$, that is, $\text{Desc}(U) = G[U \cup D_U]$.

Recall that a τ -hop with respect to a set $U \subseteq V$ is a simple path or a simple cycle of length τ with the property that the two endpoints are both contained in U while all other vertices are not in U . The absence of τ -hops with respect to U makes the neighbourhood of U highly structured: if there are no 2-hops with respect to U , then every vertex in $N(U)$ has a single neighbour in U , and if there are no 3-hops with respect to U , then the neighbourhood of U is an independent set. The absence of longer τ -hops imply similar, more technical properties for sets of vertices that have a short path to U . Jumping ahead a bit, our lower bounds rely on the absence of $\{2, 3\}$ -hops for 4-colourability and $\{2, 3, 4\}$ -hops as well as some additional small shapes for 3-colourability (see the full version).

We now define a process that constructs, given a set $U \subseteq V$, a set of vertices that contains U , has no $\{2, 3\}$ -hops and, as we shall see, is not much larger than U as long as U itself is small.

Definition 8 (Closure [53]). Let $G = (V, E)$ be a graph with a linear ordering on V and let $U \subseteq V$. Set $H_0 = \text{Desc}(U)$. While there exists a $\{2, 3\}$ -hop Q_i with respect to H_{i-1} set $H_i = \text{Desc}(V(H_{i-1}) \cup V(Q_i))$; otherwise, stop and set $H_{\text{end}} = H_{i-1}$. The *closure* of U , denoted by $\text{Cl}(U)$, is the set of vertices $V(H_{\text{end}})$.

We collect some properties of the closure in the next lemma.

Lemma 9. For any graph $G = (V, E)$ with a linear ordering on V and any set $U \subseteq V$ the following hold.

- 1) $\text{Cl}(U)$ is uniquely defined.
- 2) U is a subset of $\text{Cl}(U)$.
- 3) The closure is monotone: for every set $U' \supseteq U$, it holds that $\text{Cl}(U') \supseteq \text{Cl}(U)$.
- 4) The closure is idempotent, that is, $\text{Cl}(\text{Cl}(U)) = \text{Cl}(U)$.

Proof. Items 2 and 4 follow immediately from the definition. We prove items 1 and 3 simultaneously by induction. Let $(H_0, \dots, H_{\text{end}})$ and $(H'_0, \dots, H'_{\text{end}})$ be two sequences of the construction of the closure in Definition 8, where H_0 is $\text{Desc}(U)$ and H'_0 is either $\text{Desc}(U)$ for item 1 or $\text{Desc}(U')$ for item 3. We intend to show that $V(H_{\text{end}}) \subseteq V(H'_{\text{end}})$, from where item 3 follows immediately while item 1 follows by symmetry. The proof is by induction on i . The base case $i = 0$ is immediate for both items since $V(H_0) \subseteq V(H'_0)$ and $V(H'_0) \subseteq V(H'_{\text{end}})$ by construction. For the induction step, we may assume that $V(H_{i-1}) \subseteq V(H'_{\text{end}})$. If no 2-hop or 3-hop with respect to $V(H_{i-1})$ exists, there is nothing to prove. Hence, let Q_i be the hop added to H_{i-1} in the i th iteration. If Q_i is a $\{2, 3\}$ -hop with respect to H_{i-1} , then either Q_i is contained in H'_{end} or, if not, some subgraph of Q_i is a hop with respect to H'_{end} . But the latter cannot be, as it contradicts that $V(H'_{\text{end}}) = \text{Cl}(U')$. It follows by induction that $V(H_{\text{end}}) \subseteq V(H'_{\text{end}})$. \square

Recall that, for a graph $G = (V, E)$, the formula $\text{Col}(G, k)$ is defined over the set of variables $\{x_{v,i}\}_{v \in V, i \in [k]}$. In order to define our R -operator we require an admissible ordering of monomials over the above mentioned variables. We induce this ordering by a linear ordering \prec on V as follows: first, we order the variables in an arbitrary but fixed manner such that for any colours $i, j \in [k]$ it holds that $x_{u,i} \prec x_{v,j}$ whenever $u \prec v$. With this order fixed we then obtain the admissible ordering on monomials by first ordering the monomials by degree and then lexicographically according to the ordering on the variables.

Next we define the linear ordering on V from which we then obtain the induced admissible monomial ordering as explained in the above paragraph.

Definition 10 (χ -ordering [53]). Let $G = (V, E)$ be a graph and let $\chi: V \rightarrow [c]$ be a proper c -colouring of G . A linear ordering \prec on V is a χ -ordering on V if $u \prec v$ whenever $\chi(u) < \chi(v)$.

We conclude this section with two simple observations that are frequently used throughout the remainder of the paper.

Observation 11. If $G = (V, E)$ is a graph that is χ -ordered by a proper colouring $\chi: V \rightarrow [c]$, then every decreasing or increasing path in G has length at most $c - 1$. If, moreover, G has maximum degree d , then for every subset U of V it holds that $|V(\text{Desc}(U))| \leq 2d^{c-1}|U|$.

B. Proof Overview

The construction of our R -operator follows the general paradigm introduced by Aleknovich and Razborov [3], which, disregarding reductions, is used in essentially all subsequent polynomial calculus degree lower bounds. First, for every monomial m we identify a subset $S(m)$ of axioms in $\text{Col}(G, k)$ that are in some sense relevant to m . Then, we define R on the monomial m as the reduction modulo $\langle S(m) \rangle$ and extend R linearly to arbitrary polynomials. The goal is to satisfy properties 1-3 in Lemma 2, which typically requires two technical lemmas. The first is the *size lemma*, which states that $S(m)$ is not much larger than the degree of m . The second is the *reduction lemma*, which states that for every ideal I that is generated by a set of polynomials that contains $S(m)$ and is furthermore not too large (in particular, I should be small enough to not contain 1), it holds that $R_I(m) = R_{\langle S(m) \rangle}(m)$. Once these lemmas are in place, the lower bound follows by standard arguments.

In a recent preprint [53], Romero and Tunçel use the notions of closure and χ -ordering to define the set $S(m)$ for each monomial m . Their proofs of the size lemma and reduction lemma assume that the underlying graph has large girth and hence is locally a tree. We use their notion of closure for our lower bound for 4-colourability on random regular graphs (Theorem 12), but the size lemma and reduction lemma cannot be proved in the same way since random graphs asymptotically almost surely contain short cycles. Instead of girth, we rely on the notion of *sparsity* from Section II-D. For the improvement to 3-colourability on random regular graphs and on Erdős-Rényi random graphs, we need new notions of closures and, in the proof of the reduction lemma, a more refined argument based on graph contraction. We defer these improvements to the full version of this paper.

IV. A LOWER BOUND FOR 4-COLOURABILITY ON RANDOM REGULAR GRAPHS

In this section we prove a linear degree polynomial calculus refutation lower bound for refuting the claim that there exists a 4-colouring of a constant degree random regular graph.

Theorem 12. There is a constant $\delta > 0$ such that the following holds for all integers n and d satisfying $\delta d^3 / \log d < \log n$. If G is a graph sampled from $\mathbb{G}_{n,d}$ and $k \geq 4$ is an integer, then asymptotically almost surely every polynomial calculus refutation of $\text{Col}(G, k)$, over any field, requires degree $2^{-O(d)}n$.

When d is constant, Theorem 12 implies polynomial calculus resolution size lower bounds for $\text{Col}(G, k)$ of the form $\exp(\Omega(n))$ through the size-degree relation mentioned in Section II-A.

Fix a graph $G = (V, E)$ and write $\mathbb{F}[\mathbf{x}]$ to denote $\mathbb{F}[x_{v,i} \mid v \in V, i \in [k]]$ modulo the Boolean axioms. As outlined, we first use the closure to associate each monomial m in $\mathbb{F}[\mathbf{x}]$ with a subset $S(m)$ of $\text{Col}(G, k)$. Let c be the chromatic number of G and let χ be a c -colouring of G . Fix a χ -ordering on V and an admissible ordering on $\mathbb{F}[\mathbf{x}]$ induced by χ . Given a polynomial p we let $V(p)$ denote the set of vertices mentioned by the variables in p .

Definition 13 (Term closure). The *closure* of a term t , denoted by $\text{Cl}(t)$, is the vertex set $\text{Cl}(V(t))$.

The ‘‘relevant’’ axiom set $S(m)$ associated with a monomial m is the set of k -colourability axioms defined over $G[\text{Cl}(m)]$, that is, the axioms in $\text{Col}(G[\text{Cl}(m)], k)$. For brevity, given a subset U of V we denote the ideal $\langle \text{Col}(G[U], k) \rangle$ by $\langle U \rangle$ and refer to the polynomials in $\text{Col}(G[U], k)$ as the generators of $\langle U \rangle$.


We let our R -operator on a monomial m be defined by $R(m) = R_{\langle \text{Cl}(m) \rangle}(m)$, where the reduction is done with respect to the admissible ordering induced by the χ -ordering, and extend R linearly to arbitrary multilinear polynomials in $\mathbb{F}[\mathbf{x}]$. To verify items 1-3 in Lemma 2 we need to establish two properties of the closure of a monomial m . First, we show that the size of the closure is closely related to the degree of the monomial m . We show that this property indeed holds if m is of low degree, G is locally sparse, and the closure is taken with respect to any admissible ordering induced by a χ -ordering on V .

Lemma 14 (Size lemma). *If $G = (V, E)$ has maximum degree d , chromatic number c and is $(\ell, 1/4c)$ -sparse, then, for every set $U \subseteq V$ of size $D \leq \ell/20c$, it holds that $|\text{Cl}(U)| \leq 40d^{c-1}D$.*

Proof. Recall from Definition 8 that the closure of a set $U \subseteq V$ is defined to be the vertex set of the final element of a sequence $(H_0, H_1, \dots, H_{\text{end}})$, where $H_0 = \text{Desc}(U)$ and H_i is obtained from H_{i-1} by appending a $\{2, 3\}$ -hop with respect to H_{i-1} and then taking the descendant graph of the set of vertices of the resulting graph. The key observation of the proof is that adding a $\{2, 3\}$ -hop to H_i adds more edges than vertices, thus increasing the edge density. As the graph G is locally sparse we may conclude that the sequence $(H_0, H_1, \dots, H_{\text{end}})$ needs to be rather short which allows us to argue the size upper bound on the closure of U .

In a bit more detail, for each H_i we identify a vertex set U_i such that the edge density of the graph $G[U_i]$ increases with i . Since the sets U_i grow very slowly and thus the local sparsity always applies, we may conclude that the number of iterations is bounded. Finally, as the graphs H_i are the descendant graphs of the sets U_i , it holds that H_{end} is the descendant graph of a set which is not much larger than the initial set U , whereby the theorem follows from the descendant graph size bound in Observation 11.

We inductively define U_i as follows. Let $U_0 = U$ and let Q_i be the hop added to H_{i-1} at iteration $i \geq 1$. If we denote by u and v the endpoints of Q_i (possibly, $u = v$) and let P_u and



■ U_{i-1} ■ U_i
□ $H_{i-1} = \text{Desc}(U_{i-1})$ □ $H_i = \text{Desc}(U_i)$

Fig. 1: Construction of U_i .

P_v be two shortest decreasing paths from U_{i-1} to u and v , respectively, then $U_i = U_{i-1} \cup V(P_u \cup P_v \cup Q_i)$.

For this definition to be meaningful, we need to establish that the paths P_u and P_v always exist.

Claim 15. For every vertex v in H_i , there exists a decreasing path in H_i from some vertex in U_i to v .

Proof. The proof is by induction on i . The base case $i = 0$ holds because $H_0 = \text{Desc}(U_0)$. For the induction step, suppose that the claim holds for $i - 1$. By definition, the vertices in $H_i \setminus H_{i-1}$ are descendants of a vertex in Q_i , and all vertices in Q_i are contained in U_i . □

Next, we show that $|U_i|$ grows slowly with i and that the edge density $|E(U_i)|/|U_i|$ exceeds the sparsity threshold $(1 + 1/4c)$ after a small number of iterations. We are deliberately loose with constants below in order for the same estimates to hold in a more general setting that we consider in the full version of this paper.

Claim 16. It holds that $|U_i \setminus U_{i-1}| \leq 2c + |V(Q_i)| - 4$ and $|E(U_i)| \geq |E(U_{i-1})| + |U_i \setminus U_{i-1}| + 1$.

Proof. Denote the graph $P_u \cup P_v \cup Q_i$ by F . By Observation 11, P_u and P_v contain at most c vertices each, so $3 \leq |V(F)| \leq 2c + |V(Q_i)| - 2$. Moreover, the endpoints of F are contained in U_{i-1} and all other vertices in F are outside of U_{i-1} . By our choice of P_u and P_v there are two cases, depending on whether F contains a cycle or not.

Case 1: If there is no cycle in F , then $|V(F) \cap U_{i-1}| = 2$ so $|U_i \setminus U_{i-1}| = |V(F)| - 2$. Moreover $|E(F)| \geq |V(F)| - 1$ since F is connected.

Case 2: If F contains a cycle, then $|V(F) \cap U_{i-1}| = 1$, hence $|U_i \setminus U_{i-1}| = |V(F)| - 1$. In addition, $|E(F)| \geq |V(F)|$ since F is connected and contains a cycle. Moreover, it holds that $|V(F)| \leq 2c + |V(Q_i)| - 3$.

In both cases, we have that $|U_i \setminus U_{i-1}| \leq 2c + |V(Q_i)| - 4$ and $|E(U_i)| \geq |E(U_{i-1})| + |U_i \setminus U_{i-1}| + 1$. □

Towards contradiction, suppose that $i \geq 4D + 1$. Note that $|V(Q_i)| \leq 4$, so by Claim 16 we have

$$\frac{|E(U_{4D+1})|}{|U_{4D+1}|} = \frac{|E(U)| + \sum_{i=1}^{4D+1} |E(U_i) \setminus E(U_{i-1})|}{|U| + \sum_{i=1}^{4D+1} |U_i \setminus U_{i-1}|} \quad (3a)$$

$$\geq \frac{|E(U)| + \sum_{i=1}^{4D+1} (|U_i \setminus U_{i-1}| + 1)}{|U| + \sum_{i=1}^{4D+1} |U_i \setminus U_{i-1}|} \quad (3b)$$

$$\geq \frac{0 + \sum_{i=1}^{4D+1} (2c + 1)}{D + \sum_{i=1}^{4D+1} 2c} \quad (3c)$$

$$> 1 + \frac{1}{4c}, \quad (3d)$$

which contradicts that G is $(\ell, 1/4c)$ -sparse since $|U_{4D+1}| < 15cD < \ell$. Therefore, it follows that $i \leq 4D$.

At iteration i in the construction, H_i is the descendant graph of $U \cup \bigcup_{j \leq i} V(Q_j)$, and the hop Q_i added to $V(H_{i-1})$ contains at most 2 vertices not already in H_i . Therefore, since $i \leq 4D$ in the last iteration, a rather loose upper bound on $|U \cup \bigcup_{j \leq i} V(Q_j)|$ is $20D$. With this estimate, it follows from Observation 11 that $|\text{Cl}(U)| \leq 40d^{c-1}D$. \square

We now prove the reduction lemma, which states that there is no difference between reducing a monomial m modulo $\langle \text{Cl}(m) \rangle$ and reducing modulo some slightly larger ideal that contains $\langle \text{Cl}(m) \rangle$.

Lemma 17 (Reduction lemma). *If $G = (V, E)$ is (ℓ, ε) -sparse for some $\varepsilon < 1/2$, then, for every monomial m and every subset $U \subseteq V$ of size at most ℓ satisfying $U \supseteq \text{Cl}(m)$, it holds that m is reducible modulo $\langle U \rangle$ if and only if m is reducible modulo $\langle \text{Cl}(m) \rangle$.*

The proof idea is to construct a function ρ mapping variables associated with vertices in $U \setminus \text{Cl}(m)$ to either constants or polynomials of smaller order such that all axioms in $\langle U \rangle \setminus \langle \text{Cl}(m) \rangle$ are either satisfied or mapped to a polynomial in $\langle \text{Cl}(m) \rangle$. It is not hard to see that such a mapping turns any polynomial in $\langle U \rangle$ with leading monomial m into a smaller polynomial in $\langle \text{Cl}(m) \rangle$ whose leading monomial is also m . It then follows that a monomial m is reducible modulo $\langle U \rangle$ if m is reducible modulo $\langle \text{Cl}(m) \rangle$. The other direction is immediate, so this suffices to prove the lemma.

We now outline the construction of ρ . Variables far from $\text{Cl}(m)$, which here means variables associated with a vertex in $U \setminus (\text{Cl}(U) \cup N_U(\text{Cl}(m)))$, are mapped according to a 3-colouring χ of the subgraph $G[U \setminus \text{Cl}(m)]$. It remains to define ρ on variables associated with vertices in $N_U(\text{Cl}(m))$. By our choice of closure, $N_U(\text{Cl}(m))$ forms an independent set, and furthermore each vertex u in $N_U(\text{Cl}(m))$ has precisely one adjacent vertex v in $\text{Cl}(m)$. Hence, as $k \geq 4$ and χ is a 3-colouring of $G[U \setminus \text{Cl}(m)]$, no matter how this vertex v is coloured there is always a colour c_u available to properly 4-colour u . We may think of c_u as a function that, given χ and the colour of v , outputs a (proper) 4-colouring of u . Variables associated with such a vertex u are mapped according to c_u by ρ .

To implement the above proof outline we need some further notation: for a polynomial p and a partial function ρ mapping variables to polynomials we let $p \upharpoonright_\rho$ denote the polynomial obtained from p by substituting every occurrence of a variable x_i in the domain of ρ by $\rho(x_i)$.

Proof of Lemma 17. Since $\text{Cl}(m)$ is a subset of U , it follows that if m is reducible modulo $\langle \text{Cl}(m) \rangle$, then m is also reducible modulo $\langle U \rangle$. For the reverse direction, define a mapping ρ as follows. Since G is (ℓ, ε) -sparse for $\varepsilon < 1/2$, and since $|U| \leq \ell$, it follows from Lemma 5 that there exists a 3-colouring χ of the subgraph $G[U \setminus \text{Cl}(m)]$. Variables associated with a vertex u in $U \setminus (\text{Cl}(m) \cup N_U(\text{Cl}(m)))$ are then mapped accordingly: $\rho(x_{u,i}) = 1$ if $\chi(u) = i$ and $\rho(x_{u,i}) = 0$ otherwise.

Next, for each vertex $u \in N_U(\text{Cl}(m))$ we define ρ on the variables associated to u . Since χ is a 3-colouring of $G[U \setminus \text{Cl}(m)]$, the neighbourhood of u is coloured by at most 2 colours. As $k \geq 4$ there are two distinct colours $c_1, c_2 \in [k]$ not appearing in the neighbourhood of u . Furthermore, since there are no 2-hops in U with respect to $\text{Cl}(m)$, the vertex u has a single neighbour $v \in \text{Cl}(m)$. Define ρ on u by

$$x_{u,c_1} \mapsto x_{v,c_2}; \quad (4a)$$

$$x_{u,c_2} \mapsto \sum_{i \in [k], i \neq c_2} x_{v,i}; \quad (4b)$$

$$x_{u,i} \mapsto 0, \text{ for all } i \neq c_1, c_2. \quad (4c)$$

This completes the definition of ρ .

Let f be a polynomial in $\langle U \rangle$ with leading monomial m . We claim that $f \upharpoonright_\rho \in \langle \text{Cl}(m) \rangle$, that all monomials m' satisfy $m' \upharpoonright_\rho \preceq m'$, and that $m = m \upharpoonright_\rho$. If so, we are done, since then m is the leading monomial of the polynomial $f \upharpoonright_\rho \in \langle \text{Cl}(m) \rangle$ and we may hence conclude that if m is reducible modulo $\langle U \rangle$, then m is also reducible modulo $\langle \text{Cl}(m) \rangle$.

We now argue that the three properties hold. The latter two are almost immediate: since ρ does not map variables associated with $\text{Cl}(m)$ (of which $V(m)$ is a subset) we have $m = m \upharpoonright_\rho$. Furthermore, since $V(\text{Desc}(\text{Cl}(m))) = \text{Cl}(m)$ it holds for every variable x that $x \upharpoonright_\rho \preceq x$, and hence every monomial m' satisfies $m' \upharpoonright_\rho \preceq m'$.

It remains to prove that $f \upharpoonright_\rho \in \langle \text{Cl}(m) \rangle$. Since there are no 3-hops in U with respect to $\text{Cl}(m)$, the vertices in $N_U(\text{Cl}(m))$ form an independent set. Therefore, the mapping ρ extends any proper k -colouring of $\text{Cl}(m)$ to a proper k -colouring of U , which in turn implies that every axiom p of $\text{Col}(G[U], k)$ satisfies $p \upharpoonright_\rho \in \langle \text{Cl}(m) \rangle$ by Observation 4. Since $f \in \langle U \rangle$ we may write $f = \sum_i a_i p_i$ for polynomials a_i in $\mathbb{F}[\mathbf{x}]$ and axioms p_i of $\text{Col}(G[U], k)$. As we noted, it holds for each of the axioms p_i that $p_i \upharpoonright_\rho \in \langle \text{Cl}(m) \rangle$ and therefore the polynomial $f \upharpoonright_\rho = \sum_i a_i \upharpoonright_\rho \cdot p_i \upharpoonright_\rho$ is in $\langle \text{Cl}(m) \rangle$ as claimed. \square

Proof of Theorem 12. Let c be the chromatic number of G and let $\chi: V \rightarrow [c]$ be a colouring of G . Fix a χ -ordering on V and an admissible ordering on $\mathbb{F}[\mathbf{x}]$ induced by χ . We assume that $c \leq 2d/\log d$ and that G is $(\ell, 1/4c)$ -sparse for $\ell = 2^{-300d}n$; both properties hold asymptotically almost

surely by Lemma 7 and Lemma 6, respectively. For the latter we use the bound $\delta d^3 / \log d \leq \log n$ to argue that the second parameter $1/4c$ is large enough.

Fix $D = 2^{-300d_n} / (40d^{c-1}) > 2^{-350d_n} = 2^{-O(d)}n$. Recall that $R(m) = R_{\langle \text{Cl}(m) \rangle}(m)$. The goal is to show that R satisfies items 1-3 in Lemma 2 for this choice of D , that is, we need to show that $R(1) = 1$, that R maps all axioms of $\text{Col}(G, k)$ to 0 and that for every term t of degree at most $D - 1$ and every variable x , it holds that $R(xt) = R(xR(t))$.

Note that $R(1) = 1$ is immediate since the closure of a constant polynomial is empty by definition. To see that R maps each axiom in $\text{Col}(G, k)$ to 0, let p be one such axiom and let m_p be the product of the variables in p . As m_p mentions at most two vertices, it follows by Lemma 14 that $|\text{Cl}(m_p)| \leq 80d^{c-1} < \ell$. Therefore, for each axiom $p = \sum_j t_j$ it holds that

$$R(p) = \sum_j R(t_j) \quad (5a)$$

$$= \sum_j R_{\langle \text{Cl}(t_j) \rangle}(t_j) \quad [\text{by definition}] \quad (5b)$$

$$= \sum_j R_{\langle \text{Cl}(m_p) \rangle}(t_j) \quad [\text{by Lemmas 9 and 17}] \quad (5c)$$

$$= R_{\langle \text{Cl}(m_p) \rangle} \left(\sum_j t_j \right) \quad [\text{by the linearity of } R] \quad (5d)$$

$$= R_{\langle \text{Cl}(m_p) \rangle}(p) \quad [\text{by definition}] \quad (5e)$$

$$= 0, \quad (5f)$$

where the last equality holds because p is an element of $\langle \text{Cl}(m_p) \rangle$ (in fact, one of its generators).

Finally, we need to show that for every term t of degree at most $D - 1$ and every variable x , it holds that $R(xt) = R(xR(t))$. By definition we have that

$$R(xR(t)) = \sum_{t' \in R(t)} R(xt') = \sum_{t' \in R(t)} R_{\langle \text{Cl}(xt') \rangle}(xt'), \quad (6)$$

where the sum is over terms t' in the polynomial $R(t)$. The next step is the main technical challenge. We want to show that if t is a term of degree at most $D - 1$, then reducing modulo $\langle \text{Cl}(xt') \rangle$ or $\langle \text{Cl}(xt) \rangle$ results in the same polynomial. More formally, we claim that

$$\sum_{t' \in R(t)} R_{\langle \text{Cl}(xt') \rangle}(xt') = \sum_{t' \in R(t)} R_{\langle \text{Cl}(xt) \rangle}(xt'). \quad (7)$$

Before proving (7) let us finish the proof of Theorem 12. By noting that

$$\begin{aligned} & R(xR(t)) \\ &= \sum_{t' \in R(t)} R_{\langle \text{Cl}(xt) \rangle}(xt') \quad [\text{by (6) and (7)}] \end{aligned} \quad (8a)$$

$$= R_{\langle \text{Cl}(xt) \rangle} \left(\sum_{t' \in R(t)} xt' \right) \quad [\text{by the linearity of } R] \quad (8b)$$

$$= R_{\langle \text{Cl}(xt) \rangle}(xR_{\langle \text{Cl}(t) \rangle}(t)) \quad [\text{by the definition of } t'] \quad (8c)$$

$$= R_{\langle \text{Cl}(xt) \rangle}(xt) \quad [\text{by Observation 3}] \quad (8d)$$

$$= R(xt) \quad (8e)$$

we establish the final property of Lemma 2 which concludes the proof of Theorem 12 modulo (7).

Let us now show that (7) holds. The choice of D and Lemma 14 together imply that $|\text{Cl}(xt)| \leq 40d^{c-1}D = \ell$, so if we can show that $\text{Cl}(xt') \subseteq \text{Cl}(xt)$ for each term t' in $R(t)$, we can apply Lemma 17, from which the claim follows immediately.

We first argue that $V(R(t)) \subseteq \text{Cl}(t)$. Suppose this is not the case. Then in particular $R(t) \neq 0$. Assign all variables that mention a vertex outside of $\text{Cl}(t)$ to 0, and denote this assignment by ρ . Then, the terms in $R(t)|_\rho$ are still irreducible modulo $\langle \text{Cl}(t) \rangle$. By definition, $R(t)$ is the *unique* sum of irreducible terms modulo $\langle \text{Cl}(t) \rangle$ such that $t = q + R(t)$ for some polynomial q in $\langle \text{Cl}(t) \rangle$. No variable in t nor in any of the generators of $\langle \text{Cl}(t) \rangle$ is assigned by ρ , so $t|_\rho = t$ and $q|_\rho \in \langle \text{Cl}(t) \rangle$. Furthermore, by our assumption that $V(R(t)) \not\subseteq \text{Cl}(t)$, it must hold that $R(t)|_\rho \neq R(t)$. But this contradicts the fact that the decomposition $t = q + R(t)$ is unique, and thus it follows that $V(R(t)) \subseteq \text{Cl}(t)$.

Note that clearly $V(x) \subseteq \text{Cl}(xt)$, and since the closure is monotone, it holds that $\text{Cl}(t) \subseteq \text{Cl}(xt)$. Hence, for any term t' in $R(t)$ we have $V(x) \cup V(t') \subseteq \text{Cl}(xt)$. Again by monotonicity, it holds that $\text{Cl}(xt') = \text{Cl}(V(x) \cup V(t')) \subseteq \text{Cl}(xt)$. Finally, since the closure is idempotent, it follows that $\text{Cl}(\text{Cl}(xt)) = \text{Cl}(xt)$. Therefore, Lemma 17 implies that $R_{\langle \text{Cl}(xt') \rangle}(xt') = R_{\langle \text{Cl}(xt) \rangle}(xt')$, which establishes (7). \square

We conclude this section with a technical remark about our result. Our degree lower bound is of the form $n/f(d)$, where f is at least exponential in d . This stands in contrast with previous results [12], [42] where f is at most polynomial in d . While the precise dependence on d is immaterial for sparse random graphs, it would be interesting to see if this can be improved. It is not immediately clear, however, what the correct dependence on d should be. For the sums-of-squares proof system, which simulates polynomial calculus over the reals [16], there exist strong upper bounds for k -colourability on random graphs and random regular graphs in some parameter regimes: the paper [9] showed that asymptotically almost surely, degree 2 sums-of-squares refutes k -colourability on d -regular random graphs if $d \geq 4k^2$. These results rule out a polynomial dependence on d in any linear sums-of-squares degree lower bound for k -colouring whenever k is fixed. However, similar upper bounds are not known to hold for polynomial calculus.

V. CONCLUDING REMARKS

In this work, we show that polynomial calculus over any field requires linear degree to refute that a sparse random regular graph or Erdős-Rényi random graph is 3-colourable. Our lower bound is optimal up to constant factors, and implies strongly exponential size lower bounds for the same problem by the well-known size-degree relation for polynomial calculus [34].

It would be interesting to investigate whether the ideas and concepts underlying this work could be extended to prove lower

bounds for colouring principles in other proof systems, the most obvious candidates being Sherali-Adams and sums-of-squares. Regarding polynomial calculus, it is worth noting that the closure operation defined in [53] and generalized in this work is not, per se, restricted to graph colouring. It is natural to ask whether similar techniques could be useful for proving degree lower bounds for other graph problems. One open problem is to improve the degree lower bound for matching on random graphs in [8] to linear in the graph size, and to make it hold for graphs of small constant degree. Another problem is to establish polynomial calculus size lower bounds for independent set and vertex cover, analogously to what was done for the resolution proof system in [14]. Finally, an intriguing technical challenge is to prove degree lower bounds for variants of the dense linear ordering principle [5] for graphs of bounded degree.

ACKNOWLEDGEMENTS

The authors would like to thank Albert Atserias and Gaia Carenini for helpful discussions during the course of this work, and we also thank Albert for making us aware of some relevant references. In addition, we benefitted from feedback from participants of the participants of the Dagstuhl Seminar 23111 “Computational Complexity of Discrete Problems”. Finally, we are grateful to the anonymous *FOCS* reviewers for comments that helped us improve the exposition in the paper considerably.

Part of this work was carried out while the authors were taking part in the semester programme *Meta-Complexity* and the extended reunion of the programme *Satisfiability: Theory, Practice, and Beyond* at the Simons Institute for the Theory of Computing at UC Berkeley in the spring of 2023.

Susanna F. de Rezende received funding from ELLIIT, from the Knut and Alice Wallenberg grant KAW 2021.0307, and from the Swedish Research Council grant 2021-05104. Kilian Risse was supported by the Swiss National Science Foundation project 200021-184656 “Randomness in Problem Instances and Randomized Algorithms”. Jonas Conneryd and Jakob Nordström were funded by the Swedish Research Council grant 2016-00782, and in addition Jonas Conneryd was also partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, whereas Jakob Nordström together with Shuo Pang were supported by the Independent Research Fund Denmark grant 9040-00389B.

REFERENCES

- [1] D. Achlioptas and A. Naor, “The two possible values of the chromatic number of a random graph,” *Annals of Mathematics*, vol. 162, no. 3, pp. 1335–1351, Nov. 2005.
- [2] M. Alekhnovich, E. Ben-Sasson, A. A. Razborov, and A. Wigderson, “Space complexity in propositional calculus,” *SIAM Journal on Computing*, vol. 31, no. 4, pp. 1184–1211, Apr. 2002, preliminary version in *STOC ’00*.
- [3] M. Alekhnovich and A. A. Razborov, “Lower bounds for polynomial calculus: Non-binomial case,” *Proceedings of the Steklov Institute of Mathematics*, vol. 242, pp. 18–35, 2003, available at <http://people.cs.uchicago.edu/~razborov/files/misha.pdf>. Preliminary version in *FOCS ’01*.
- [4] N. Alon and M. Tarsi, “Colorings and orientations of graphs,” *Combinatorica*, vol. 12, no. 2, pp. 125–134, Jun. 1992.

- [5] A. Atserias and V. Dalmau, “A combinatorial characterization of resolution width,” *Journal of Computer and System Sciences*, vol. 74, no. 3, pp. 323–334, May 2008, preliminary version in *CCC ’03*.
- [6] A. Atserias and T. Hakoniemi, “Size-degree trade-offs for Sums-of-Squares and Positivstellensatz proofs,” in *Proceedings of the 34th Annual Computational Complexity Conference (CCC ’19)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 137, Jul. 2019, pp. 24:1–24:20.
- [7] A. Atserias and J. Ochremiak, “Proof complexity meets algebra,” *ACM Transactions on Computational Logic*, vol. 20, pp. 1:1–1:46, Feb. 2019, preliminary version in *ICALP ’17*.
- [8] P. Austrin and K. Risse, “Perfect matching in random graphs is as hard as Tseitin,” *TheoretCS*, vol. 1, pp. Art. 2, 47, 2022.
- [9] J. Banks, R. Kleinberg, and C. Moore, “The Lovász theta function for random regular graphs and community detection in the hard regime,” *SIAM Journal on Computing*, vol. 48, no. 3, pp. 1098–1119, 2019.
- [10] L. Barto, J. Bulín, A. Krokhin, and J. Opršal, “Algebraic approach to promise constraint satisfaction,” *Journal of the ACM*, vol. 68, no. 4, pp. 28:1–28:66, Aug. 2021.
- [11] D. A. Bayer, “The division algorithm and the Hilbert scheme,” Ph.D. dissertation, Harvard University, Jun. 1982, available at <https://www.math.columbia.edu/~bayer/papers/Bayer-thesis.pdf>.
- [12] P. Beame, J. C. Culberson, D. G. Mitchell, and C. Moore, “The resolution complexity of random graph k -colorability,” *Discrete Applied Mathematics*, vol. 153, no. 1-3, pp. 25–47, Dec. 2005.
- [13] P. Beame, R. Impagliazzo, J. Krajíček, T. Pitassi, and P. Pudlák, “Lower bounds on Hilbert’s Nullstellensatz and propositional proofs,” in *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science (FOCS ’94)*, Nov. 1994, pp. 794–806.
- [14] P. Beame, R. Impagliazzo, and A. Sabharwal, “The resolution complexity of independent sets and vertex covers in random graphs,” *Computational Complexity*, vol. 16, no. 3, pp. 245–297, Oct. 2007, preliminary version in *CCC ’01*.
- [15] R. Beigel and D. Eppstein, “3-coloring in time $O(1.3289^n)$,” *Journal of Algorithms*, vol. 54, no. 2, pp. 168–204, Feb. 2005.
- [16] C. Berkholz, “The relation between polynomial calculus, Sherali-Adams, and sum-of-squares proofs,” in *Proceedings of the 35th Symposium on Theoretical Aspects of Computer Science (STACS ’18)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 96, Feb. 2018, pp. 11:1–11:14.
- [17] A. Blake, “Canonical expressions in Boolean algebra,” Ph.D. dissertation, University of Chicago, 1937.
- [18] B. Bollobás, “Chromatic number, girth and maximal degree,” *Discrete Mathematics*, vol. 24, no. 3, pp. 311–314, 1978.
- [19] S. R. Buss, “Lower bounds on Nullstellensatz proofs via designs,” in *Proof Complexity and Feasible Arithmetics*, ser. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 39. American Mathematical Society, 1998, pp. 59–71, available at <http://www.math.ucsd.edu/~sbuss/ResearchWeb/designs/>.
- [20] S. R. Buss and J. Nordström, “Proof complexity and SAT solving,” in *Handbook of Satisfiability*, 2nd ed., ser. Frontiers in Artificial Intelligence and Applications, A. Biere, M. J. H. Heule, H. van Maaren, and T. Walsh, Eds. IOS Press, Feb. 2021, vol. 336, ch. 7, pp. 233–350.
- [21] M. Clegg, J. Edmonds, and R. Impagliazzo, “Using the Groebner basis algorithm to find proofs of unsatisfiability,” in *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC ’96)*, May 1996, pp. 174–183.
- [22] A. Coja-Oghlan, “The Lovász number of random graphs,” *Combinatorics, Probability and Computing*, vol. 14, no. 4, pp. 439–465, 2005.
- [23] S. A. Cook and R. A. Reckhow, “The relative efficiency of propositional proof systems,” *Journal of Symbolic Logic*, vol. 44, no. 1, pp. 36–50, Mar. 1979, preliminary version in *STOC ’74*.
- [24] M. Davis, G. Logemann, and D. Loveland, “A machine program for theorem proving,” *Communications of the ACM*, vol. 5, no. 7, pp. 394–397, Jul. 1962.
- [25] M. Davis and H. Putnam, “A computing procedure for quantification theory,” *Journal of the ACM*, vol. 7, no. 3, pp. 201–215, 1960.
- [26] J. A. De Loera, “Gröbner bases and graph colorings,” *Beiträge zur Algebra und Geometrie*, vol. 36, no. 1, pp. 89–96, Jan. 1995, available at <https://www.emis.de/journals/BAG/vol.36/no.1/>.
- [27] J. A. De Loera, S. Margulies, M. Pernpeintner, E. Riedl, D. Rolnick, G. Spencer, D. Stasi, and J. Swenson, “Graph-coloring ideals: Nullstellensatz certificates, Gröbner bases for chordal graphs, and hardness of Gröbner bases,” in *Proceedings of the 40th International Symposium*

- on *Symbolic and Algebraic Computation (ISSAC '15)*, Jul. 2015, pp. 133–140.
- [28] J. A. De Loera, J. Lee, P. N. Malkin, and S. Margulies, “Hilbert’s Nullstellensatz and an algorithm for proving combinatorial infeasibility,” in *Proceedings of the 21st International Symposium on Symbolic and Algebraic Computation (ISSAC '08)*, Jul. 2008, pp. 197–206.
- [29] —, “Computing infeasibility certificates for combinatorial problems through Hilbert’s Nullstellensatz,” *Journal of Symbolic Computation*, vol. 46, no. 11, pp. 1260–1283, Nov. 2011.
- [30] J. A. De Loera, J. Lee, S. Margulies, and S. Onn, “Expressing combinatorial problems by systems of polynomial equations and Hilbert’s Nullstellensatz,” *Combinatorics, Probability and Computing*, vol. 18, no. 4, pp. 551–582, Jul. 2009.
- [31] M. M. Halldórsson, “A still better performance guarantee for approximate graph coloring,” *Information Processing Letters*, vol. 45, no. 1, pp. 19–23, Jan. 1993.
- [32] C. J. Hillar and T. Windfeldt, “Algebraic characterization of uniquely vertex colorable graphs,” *Journal of Combinatorial Theory, Series B*, vol. 98, no. 2, pp. 400–414, Mar. 2008.
- [33] T. Husfeldt, “Graph colouring algorithms,” in *Topics in Chromatic Graph Theory*, ser. Encyclopedia of Mathematics and its Applications, L. W. Beineke and R. J. Wilson, Eds. Cambridge University Press, May 2015, ch. 13, pp. 277–303.
- [34] R. Impagliazzo, P. Pudlák, and J. Sgall, “Lower bounds for the polynomial calculus and the Gröbner basis algorithm,” *Computational Complexity*, vol. 8, no. 2, pp. 127–144, 1999.
- [35] R. M. Karp, “Reducibility among combinatorial problems,” in *Complexity of Computer Computations*, ser. The IBM Research Symposia Series. Springer, 1972, pp. 85–103.
- [36] K.-I. Kawarabayashi and M. Thorup, “Coloring 3-colorable graphs with less than $n^{1/5}$ colors,” *Journal of the ACM*, vol. 64, no. 1, Mar. 2017.
- [37] G. Kemkes, X. Pérez-Giménez, and N. Wormald, “On the chromatic number of random d -regular graphs,” *Advances in Mathematics*, vol. 223, no. 1, pp. 300–328, Jan. 2010.
- [38] P. K. Kothari and P. Manohar, “A stress-free sum-of-squares lower bound for coloring,” in *Proceedings of the 36th Annual IEEE Conference on Computational Complexity (CCC '21)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 200, Jul. 2021, pp. 23:1–23:21.
- [39] A. Krokhnin and J. Opršal, “An invitation to the promise constraint satisfaction problem,” *ACM SIGLOG News*, vol. 9, no. 3, pp. 30–59, 2022.
- [40] J. B. Lasserre, “An explicit exact SDP relaxation for nonlinear 0-1 programs,” in *Proceedings of the 8th International Conference on Integer Programming and Combinatorial Optimization (IPCO '01)*, ser. Lecture Notes in Computer Science, vol. 2081. Springer, Jun. 2001, pp. 293–303.
- [41] M. Lauria, “Algorithm analysis through proof complexity,” in *Proceedings of the 14th Conference on Computability in Europe (CiE '18), Sailing Routes in the World of Computation*, ser. Lecture Notes in Computer Science. Springer International Publishing, Jul. 2018, vol. 10936, pp. 254–263.
- [42] M. Lauria and J. Nordström, “Graph colouring is hard for algorithms based on Hilbert’s Nullstellensatz and Gröbner bases,” in *Proceedings of the 32nd Annual Computational Complexity Conference (CCC '17)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 79, Jul. 2017, pp. 2:1–2:20.
- [43] L. Lovász, “Stable sets and polynomials,” *Discrete Mathematics*, vol. 124, no. 1–3, pp. 137–153, Jan. 1994.
- [44] Y. V. Matiyasevich, “A criterion for vertex colorability of a graph stated in terms of edge orientations,” *Diskretnyi Analiz*, vol. 26, pp. 65–71, 1974, English translation of the Russian original. Available at http://logic.pdmi.ras.ru/~yumat/papers/22_paper/.
- [45] —, “Some algebraic methods for calculating the number of colorings of a graph,” *Journal of Mathematical Sciences*, vol. 121, no. 3, pp. 2401–2408, May 2004.
- [46] C. McDiarmid, “Colouring random graphs,” *Annals of Operations Research*, vol. 1, no. 3, pp. 183–200, Oct. 1984.
- [47] M. Mikša and J. Nordström, “A generalized method for proving polynomial calculus degree lower bounds,” in *Proceedings of the 30th Annual Computational Complexity Conference (CCC '15)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 33, Jun. 2015, pp. 467–487.
- [48] M. Mnuk, “Representing graph properties by polynomial ideals,” in *Proceedings of the 4th International Workshop on Computer Algebra in Scientific Computing (CASC '01)*, Sep. 2001, pp. 431–444.
- [49] P. A. Parrilo, “Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization,” Ph.D. dissertation, California Institute of Technology, May 2000, available at <http://resolver.caltech.edu/CaltechETD:etd-05062004-055516>.
- [50] A. A. Razborov, “Lower bounds for the polynomial calculus,” *Computational Complexity*, vol. 7, no. 4, pp. 291–324, Dec. 1998.
- [51] R. A. Reckhow, “On the lengths of proofs in the propositional calculus,” Ph.D. dissertation, University of Toronto, 1975, available at https://www.cs.toronto.edu/~sacook/homepage/reckhow_thesis.pdf.
- [52] J. A. Robinson, “A machine-oriented logic based on the resolution principle,” *Journal of the ACM*, vol. 12, no. 1, pp. 23–41, Jan. 1965.
- [53] J. A. Romero Barbosa and L. Tunçel, “Graphs with large girth and chromatic number are hard for Nullstellensatz,” arXiv.org, Tech. Rep. 2212.05365, Dec. 2022.
- [54] H. D. Sherali and W. P. Adams, “A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems,” *SIAM Journal on Discrete Mathematics*, vol. 3, pp. 411–430, 1990.
- [55] D. Zuckerman, “Linear degree extractors and the inapproximability of max clique and chromatic number,” *Theory of Computing*, vol. 3, no. 6, pp. 103–128, Aug. 2007, preliminary version in *STOC '06*.