

Graph Colouring Is Hard on Average for Polynomial Calculus and Nullstellensatz*

Jonas Conneryd
Lund University
University of Copenhagen

Susanna F. de Rezende
Lund University

Jakob Nordström
University of Copenhagen
Lund University

Shuo Pang
University of Copenhagen
Lund University

Kilian Risse
EPFL

March 21, 2025

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.

1 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 [Kar72]. 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 [Hal93], and it is known that approximating this number to within a factor $n^{1-\varepsilon}$ is NP-hard [Zuc07]. 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 [KT17]. This is very far from the lower bounds that are known—it is NP-hard to $(2k - 1)$ -colour a k -colourable graph [BBKO21], but the question of whether colouring a 3-colourable graph with 6 colours is NP-hard remains open [KO22]. 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 $O(1.3289^n)$ [BE05]. A survey on various algorithms and techniques for so-called exact algorithms for graph colouring can be found in [Hus15].

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

*This is the full-length version of a paper with the same title that appeared in the *Proceedings of the 64th Annual IEEE Symposium on Foundations of Computer Science (FOCS '23)*.

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.

1.1 Previous Work

Focusing on random graphs, McDiarmid [McD84] developed a method for determining k -colourability that captures a range of algorithmic approaches. Beame et al. [BCMM05] showed that this method could in turn be simulated by the resolution proof system [Bla37, DP60, DLL62, Rob65], 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 [AT92, Lov94, Mat74, Mat04]. Bayer [Bay82] 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 of 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 [DL95, Mnu01, HW08], culminated in an award-winning sequence of works [DLMM08, DLMO09, DLMM11, DMP⁺15] 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* [BIK⁺94] and *polynomial calculus* [CEI96, ABRW02]. On the contrary, the authors of [DLMO09] 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 [DMP⁺15] remained the best known until optimal, linear, degree lower bounds for polynomial calculus were established in [LN17] using a reduction from so-called functional pigeonhole principle formulas [MN15]. A more general reduction framework was devised in [AO19] to obtain optimal degree lower bounds also for the proof systems *Sherali-Adams* [SA90] and *sums-of-squares* [Las01, Par00], as well as weakly exponential size lower bounds for *Frege proofs* [CR79, Rec75] 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 [BCMM05]. However, this goal has remained elusive for almost two decades, as pointed out, e.g., in [MN15, LN17, Lau18, BN21]. 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 [BKM19], improving on [Coj05], 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 $k = n^{1/2+\epsilon}$ colours [KM21]. 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.

1.2 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 1.1 (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, as well as Erdős-Rényi random graphs, 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 [LN17], our degree lower bounds also apply to the k -colourability encoding introduced in [Bay82] and used in computational algebra papers such as [DLMM08, DLMO09, DLMM11, DMP⁺15], 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 [RT22]. 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 more detail next.

1.3 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 [IPS99, AH19].

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* [Bus98], 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 an expectation over a fake probability distribution supported on 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. The most common way to prove degree lower bounds is to design a *pseudo-reduction operator* [Raz98], 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 [AR03], a pseudo-reduction operator \tilde{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 $\tilde{R}(m) = R_{\langle S(m) \rangle}(m)$ 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 $\tilde{R}(1) = R_{\langle S(1) \rangle}(1) = 1$. In order for \tilde{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 $\tilde{R}(p + p') = \tilde{R}(p) + \tilde{R}(p')$ and $\tilde{R}(p \cdot p') = \tilde{R}(p) \cdot \tilde{R}(p')$. The first equality is immediate, since \tilde{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 \tilde{R} should commute with multiplication.

Proving that a pseudo-reduction operator \tilde{R} 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 [AR03] 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 $\tilde{R}(m') = R_{\langle S(m') \rangle}(m') = R_{\langle S(m) \cup S(m') \rangle}(m')$ and $\tilde{R}(m \cdot m') = R_{\langle S(m) \cup S(m') \rangle}(m \cdot m')$, from which it follows that $\tilde{R}(m \cdot \tilde{R}(m')) = \tilde{R}(m \cdot m')$ 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 such an ordering should exist, and due to this technical problem it has not been possible to join local ideal reduction operators into a globally consistent pseudo-reduction operator.

This technical problem was addressed in a recent paper [RT22] 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 [RT22], 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 [Bol78].

In our work, we employ the same ordering as in [RT22], 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 simple and elegant, which we view as an additional strength of our result.

1.4 Outline of This Paper

The rest of this paper is organized as follows. In Section 2 we present some preliminaries and then, as a warm-up, reprove the resolution width lower bound for the colourability formula in Section 3. After this, we proceed to revisit the general framework to obtain polynomial calculus lower bounds in Section 4. In Section 5 we introduce the important notion of a closure and in Section 6 prove our main theorem. We conclude with some final remarks and open problems in Section 7.

2 Preliminaries

Let us start by briefly reviewing the necessary preliminaries from proof complexity, graph theory, and algebra. We use standard asymptotic notation. In this paper \log denotes the logarithm base 2, while \ln denotes the natural logarithm.

For a field \mathbb{F} we let $\mathbb{F}[x_1, \dots, x_n]$ denote the polynomial ring over \mathbb{F} in n variables and let a *monomial* denote a product of variables. We denote by $\text{Vars}(m)$ the *variables of a monomial* m , that is, if $m = \prod_{i \in I} x_i$, then $\text{Vars}(m) = \cup_{i \in I} x_i$ and extend this notation to polynomials $p = \sum_m a_m m$ by $\text{Vars}(p) = \cup_m \text{Vars}(m)$. For polynomials $p_1, \dots, p_m \in \mathbb{F}[x_1, \dots, x_n]$ we let $\langle p_1, \dots, p_m \rangle$ denote the ideal generated by these polynomials: $\langle p_1, \dots, p_m \rangle$ contains all polynomials of the form $\sum_{i=1}^m q_i p_i$, for $q_i \in \mathbb{F}[x_1, \dots, x_n]$. For a polynomial q and a partial function ρ mapping variables to polynomials we let $q \upharpoonright_\rho$ denote the polynomial obtained from q by substituting every occurrence of a variable x_i in the domain of ρ by $\rho(x_i)$.

2.1 Proof Complexity

Polynomial calculus (PC) [CEI96] 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:

$$\text{Linear combination: } \frac{p}{ap + bq}, \quad a, b \in \mathbb{F}; \quad (2.1a)$$

$$\text{Multiplication: } \frac{p}{x_i p}, \quad x_i \text{ any variable.} \quad (2.1b)$$

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 (2.1a)-(2.1b) 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} . It is well-known that polynomial calculus is sound and complete when the system \mathcal{P} contains all Boolean axioms $\{x_1^2 - x_1, \dots, x_n^2 - x_n\}$. We often refer to \mathcal{P} as the set of *axioms*, and we say that a subset of axioms $\mathcal{Q} \subseteq \mathcal{P}$ is *satisfiable* if the polynomials in \mathcal{Q} have a common root.

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 of 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* [IPS99]: if \mathcal{P} consists of polynomials with initial degree d and contains all Boolean axioms, and if D is the minimal degree among all polynomial calculus refutations of \mathcal{P} , then every refutation of \mathcal{P} must have size $\exp(\Omega((D - d)^2/n))$.

The size-degree relation also applies to the stronger proof system *polynomial calculus resolution* (PCR) [ABRW02], 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 1.1 and Section 1.2.

2.2 Graph Colouring and Pseudo-Reductions

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

$$\sum_{i=1}^k x_{v,i} - 1 \quad v \in V(G) \quad [\text{every vertex is assigned a colour}] \quad (2.2a)$$

$$x_{v,i}x_{v,i'} \quad v \in V(G), i \neq i' \in [k] \quad [\text{no vertex gets more than one colour}] \quad (2.2b)$$

$$x_{u,i}x_{v,i} \quad (u, v) \in E(G), i \in [k] \quad [\text{no two adjacent vertices get the same colour}] \quad (2.2c)$$

$$x_{v,i}^2 - x_{v,i} \quad v \in V(G), i \in [k] \quad [\text{Boolean axioms}] \quad (2.2d)$$

whose common roots correspond precisely to proper k -colourings of G . We refer to axioms in (2.2a) and (2.2b) as *vertex axioms* and to (2.2c) as *edge axioms*. Note that the system $\text{Col}(G, k)$ of polynomials is not the standard polynomial translation of the usual CNF formula (defined in Section 3.1) as (2.2a) does not correspond to a single clause. We work with the above formulation for the sake of exposition and our lower bound also applies to the standard translation of the CNF formula.

If the field \mathbb{F} we are working over contains, or can be extended to contain, a primitive k th root of unity, then it is known [LN17, Proposition 2.2] that a polynomial calculus degree lower bound for $\text{Col}(G, k)$ also applies to *Bayer's formulation* [Bay82] of k -colourability, where each colour corresponds to a k th root of unity. This encoding has received considerable attention in computational algebra [DLMM08, DLMO09, DLMM11, DMP⁺15, RT22].

Our proof of Theorem 1.1 is based on the notion of a *pseudo-reduction* operator or *R-operator*. The following lemma is due to Razborov [Raz98].

Definition 2.1 (Pseudo-reduction). Let $D \in \mathbb{N}^+$ and \mathcal{P} be a set of polynomials over $\mathbb{F}[x_1, \dots, x_n]$. An \mathbb{F} -linear operator $\tilde{R}: \mathbb{F}[x_1, \dots, x_n] \rightarrow \mathbb{F}[x_1, \dots, x_n]$ is a *degree- D pseudo-reduction* for \mathcal{P} if

1. $\tilde{R}(1) = 1$,
2. $\tilde{R}(p) = 0$ for every polynomial $p \in \mathcal{P}$, and
3. $\tilde{R}(x_i m) = \tilde{R}(x_i \tilde{R}(m))$ for any monomial m of degree at most $D - 1$ and any variable x_i .

Lemma 2.2 ([Raz98]). *Let $D \in \mathbb{N}^+$ and \mathcal{P} be a set of polynomials over $\mathbb{F}[x_1, \dots, x_n]$. If there is a degree- D pseudo-reduction for \mathcal{P} , then any polynomial calculus refutation of \mathcal{P} over \mathbb{F} requires degree strictly greater than D .*

The proof of Lemma 2.2 is straightforward: apply \tilde{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 .

2.3 Algebra Background

The definition of our pseudo-reduction operator requires some standard notions from algebra. A total well-order $<$ on the monomials in $\mathbb{F}[x_1, \dots, x_n]$ is *admissible* if the following properties hold:

1. For any monomial m it holds that $1 < m$.
2. For any monomials m_1, m_2 and m such that $m_1 < m_2$, it holds that $mm_1 < mm_2$.

Note that this definition is the more standard definition of admissible order (also known as monomial order) used in algebra, and differs from that introduced in [Raz98], and used subsequently in [AR01, MN15], since it is defined over any monomial in $\mathbb{F}[x_1, \dots, x_n]$ and not only multilinear monomials, and it does not necessarily have to respect the total degree of monomials. We use the above definition due to its simplicity. Recall that we only concern ourselves with multilinear polynomials in this article by tacitly working over $\mathbb{F}[x_1, \dots, x_n] / \langle x_1^2 - x_1, \dots, x_n^2 - x_n \rangle$.

We write $m_1 \leq m_2$ to denote that $m_1 < m_2$ or $m_1 = m_2$. The *leading monomial* of a polynomial p is the largest monomial according to $<$ appearing in p with a non-zero coefficient. For an ideal I over $\mathbb{F}[x_1, \dots, x_n]$, a monomial m is *reducible modulo I* if m is the leading monomial of some polynomial $q \in I$; otherwise m is *irreducible modulo I* . Under a total monomial order 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 monomials 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 . It is straightforward to verify that the reduction operator is linear over the vector space of polynomials.

We will use the following standard fact.

Observation 2.3. *If I_1 and I_2 are ideals over $\mathbb{F}[x_1, \dots, x_n]$ and $I_1 \subseteq I_2$, then for any monomials m and m' it holds that $R_{I_2}(m'R_{I_1}(m)) = R_{I_2}(m'm)$.*

Proof. Write $m = q_1 + r_1$ where r_1 is the reduction of m modulo I_1 and hence $q_1 \in I_1$. Similarly, let $m' \cdot R_{I_1}(m) = m' \cdot r_1 = q_2 + r_2$ for $q_2 \in I_2$ and r_2 the reduction of $m' \cdot r_1$ modulo I_2 . We have that $R_{I_2}(m'R_{I_1}(m)) = r_2$ and we now argue that $R_{I_2}(m'm) = r_2$. Note that $m'(m - q_1) = r_2 + q_2$ and hence $m' \cdot m = r_2 + q_2 + m' \cdot q_1$. Since $I_1 \subseteq I_2$ and ideals are closed under multiplication and addition it holds that $q_2 + m_2 q_1 \in I_2$. Moreover, since r_2 is irreducible modulo I_2 , it follows that $R_{I_2}(m'm) = r_2$ by uniqueness of a reduction modulo an ideal. \square

Finally, let us record the following form of Hilbert's Nullstellensatz on the Boolean cube.

Lemma 2.4. *Let g be a polynomial and Q be a set of polynomials in $\mathbb{F}[x_1, \dots, x_n]$, and suppose that Q contains all the Boolean axioms. Then it holds that g vanishes on all common roots of Q if and only if $g \in \langle Q \rangle$.*

Note that the interesting direction of Lemma 2.4 immediately follows from the implicational completeness of the Nullstellensatz proof system. For the convenience of the reader we provide a self-contained proof in Appendix A.

2.4 Graph Theory

We consider graphs $G = (V, E)$ that are finite and undirected, and contain no self-loops or multi-edges. Given a vertex set $U \subseteq V$, the *neighbourhood of U in G* is $N(U) = \{v \in V \mid \exists u \in U : (u, v) \in E\}$, and for a second 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 $E(U)$ and we let $G[U]$ denote the subgraph induced by U in G , that is, $G[U] = (U, E(U))$. 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.

For an edge $e = (u, v) \in E$ the *contraction of G with respect to e* is the graph obtained from G by identifying the vertices u and v as a single, new, vertex v_e and adding an edge between a vertex $w \in V \setminus \{u, v\}$ and v_e if and only if there is an edge between w and at least one of u or v in G . For a set of edges $S \subseteq E$ we let the *contraction of G with respect to S* be the graph obtained from G by contracting the edges in S one at a time, in any order.

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 $(u, v) \in E$, in which case we refer to χ as a *proper k -colouring of G* . The *chromatic number* of G , denoted by $\chi(G)$, is the smallest integer k such that G is k -colourable.

Definition 2.5 (Sparsity). A graph $G = (V, E)$ is *(ℓ, ε) -sparse* if every vertex set $U \subseteq V$ of size at most ℓ satisfies $|E(U)| \leq (1 + \varepsilon)|U|$.

The essential property of sparse graphs we use to obtain our main result is that large subsets of vertices in such graphs are 3-colourable.

Lemma 2.6. *If a graph $G = (V, E)$ 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. One is the *Erdős-Rényi random graph model* $\mathbb{G}(n, p)$, which is the distribution over graphs on n vertices where each edge is included independently with probability p . The other is the *random d -regular graph model* $\mathbb{G}_{n,d}$, which is the uniform distribution over 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 sparse with excellent parameters, as stated in the following lemma, which is essentially due to [Raz17]. (See Appendix B for a proof.)

Lemma 2.7 (Sparsity lemma). For $n, d \in \mathbb{N}^+$ and $\varepsilon, \delta \in \mathbb{R}^+$ such that $\varepsilon\delta = \omega(1/\log n)$, the following holds asymptotically almost surely.

1. If G is a graph sampled from $\mathbb{G}(n, d/n)$, then it is $((4d)^{-(1+\delta)(1+\varepsilon)/\varepsilon} n, \varepsilon)$ -sparse.
2. For $d^2 \leq \varepsilon\delta \log n$, if G is a graph sampled from $\mathbb{G}_{n,d}$, then it is $((8d)^{-(1+\delta)(1+\varepsilon)/\varepsilon} n, \varepsilon)$ -sparse.

Finally, we need some bounds on the chromatic number of graphs sampled from $\mathbb{G}(n, d/n)$ or $\mathbb{G}_{n,d}$, where, in particular, the lower bounds ensure that $\text{Col}(G, k)$ is unsatisfiable for large enough d .

Lemma 2.8 ([KPGW10, CFRR02, AN05, Luc91]). For $n \in \mathbb{N}$, $d \leq n^{0.1}$ and a graph G sampled from either $\mathbb{G}(n, d/n)$ or $\mathbb{G}_{n,d}$ it holds asymptotically almost surely that the chromatic number $\chi(G)$ is at most $2d/\log d$ and, if $d \geq 6$, then $\chi(G) \geq 4$.

3 A Simple Resolution Lower Bound for 4-Colourability on Sparse Graphs

In this section we reprove the main result of Beame et al. [BCMM05], namely that resolution requires large width to refute the claim that random graphs have small chromatic number. We hope that the exposition below may serve as a gentle introduction to the polynomial calculus lower bounds that will follow later, which build on similar albeit slightly more complicated concepts. Readers only interested in the polynomial calculus lower bounds may safely skip this section.

For expert readers let us remark that while Beame et al. [BCMM05] build on the width lower bound methodology introduced by Ben-Sasson and Wigderson [BW01], we rely on the game characterization due to Pudlák [Pud00]. We recover the bounds of Beame et al. precisely for $k \geq 4$, but for $k = 3$ we incur a slight loss in parameters.

In the following we recollect some standard notions to then prove the resolution width lower bound in Section 3.2.

3.1 Graph Colouring in CNF and the Resolution Proof System

A *clause* is a disjunction over a set of *literals* $\bigvee_{i \in S} \ell_i$, where every literal ℓ_i is either a Boolean variable x or the negation \bar{x} thereof. The *width* of a clause is the number of literals in it. A *CNF formula* $F = C_1 \wedge \dots \wedge C_m$ is a conjunction of clauses, and the *width* of F is the maximum width of any clause C_i in F .

A *resolution refutation* of a CNF formula F is as an ordered sequence of clauses $\pi = (D_1, \dots, D_\tau)$ such that $D_\tau = \perp$ is the empty clause and each D_i either occurs in F or is derived from clauses D_{j_1} and D_{j_2} , with $j_1 < i$ and $j_2 < i$, by the *resolution rule*

$$\frac{B \vee x \quad C \vee \bar{x}}{B \vee C}. \quad (3.1)$$

The *width* of a resolution refutation $\pi = (D_1, \dots, D_\tau)$ is the maximum width of any clause D_i in π and the *length* of π is τ . The *size-width* relation [BW01] relates the minimum width W required to refute a CNF formula F to the minimum length of any resolution refutation of F : if F is of initial width w and defined over n variables, then any resolution refutation of F requires size $\exp(\Omega((W - w)^2/n))$.

As resolution operates over clauses, in contrast to polynomial calculus that operates over polynomials, we need to define the colourability formula as a CNF formula. In this section, for a

graph G and integer k , we let $\text{Col}(G, k)$ denote the CNF formula consisting of the clauses

$$\bigvee_{i=1}^k x_{v,i}, \quad v \in V(G) \quad \text{[every vertex is assigned a colour]} \quad (3.2a)$$

$$\bar{x}_{v,i} \vee \bar{x}_{v,i'}, \quad v \in V(G), i \neq i' \quad \text{[no vertex gets more than one colour]} \quad (3.2b)$$

$$\bar{x}_{u,i} \vee \bar{x}_{v,i}, \quad (u, v) \in E(G), i \in [k]. \quad \text{[no two adjacent vertices get the same colour]} \quad (3.2c)$$

Clearly, the CNF formula $\text{Col}(G, k)$ is satisfiable if and only if G is k -colourable.

3.2 Resolution Lower Bounds

For the sake of simplicity we prove the theorem below for the 4-colourability formula only. The theorem can be extended to 3-colourability using concepts from Section 5.

Theorem 3.1. *Let G be a graph, and suppose that $\ell \in \mathbb{N}^+$ and $\varepsilon > 0$ are such that G is $(\ell, 1/2 - \varepsilon)$ -sparse. Then every resolution refutation of $\text{Col}(G, 4)$ requires width at least $\ell/4$.*

Combining Lemma 2.7 with the above theorem we obtain the following corollaries for random graphs.

Corollary 3.2 ([BCMM05]). *For any $\varepsilon > 0$, if G is a graph sampled from $\mathbb{G}(n, d/n)$, then asymptotically almost surely every resolution refutation of $\text{Col}(G, 4)$ requires width $n/4(4d)^{3+\varepsilon}$.*

While Beame et al. [BCMM05] do not consider random regular graphs, it is not hard to see that their techniques can be used to prove the following statement.

Corollary 3.3. *For any $\varepsilon > 0$, if G is a graph sampled from $\mathbb{G}_{n,d}$ and $d^2 = o(\log n)$, then asymptotically almost surely every resolution refutation of $\text{Col}(G, 4)$ requires width $n/4(8d)^{3+\varepsilon}$.*

We prove Theorem 3.1 using the prover-adversary game as introduced by Pudlák [Pud00], which we describe here adapted to the colouring formula. The width- w prover-adversary game for colouring proceeds in rounds. In each round the prover either queries or forgets the colouring of a vertex. In response to the former, the adversary has to respond with a colouring of the queried vertex. The prover has limited memory and may only remember the partial colouring of up to w vertices. The prover wins whenever the remembered partial colouring is improper. This game characterises resolution refutation width while in our setting a more precise statement is that the prover has a winning strategy in the width- w prover-adversary game for colouring if and only if there is a resolution refutation of the colourability formula where every clause in the refutation mentions at most w vertices.

Ultimately we want to design a strategy for the adversary so that the prover cannot win with limited memory. In order not to reach a partial colouring that is impossible to extend to the newly queried vertex, whenever the prover remembers a partial colouring defined on a subset U of the vertices, the adversary maintains a consistent partial colouring defined on a *closure* of U , as defined below. Intuitively, a closure of U should be a slightly larger set that contains U and other vertices that, if not taken into account when colouring U , might not be possible to properly colour later on.

Definition 3.4 (Closure). Let $G = (V, E)$ be a graph and let $U \subseteq V$. We say that U is *closed* if all vertices $v \in V \setminus U$ satisfy $|N_U(v)| \leq 1$. A *closure* of U is any minimal closed set that contains U .

Proposition 3.5. *Every set of vertices has a unique closure.*

Proof. Suppose there are two distinct closures W_1 and W_2 of a set U . As both W_1 and W_2 contain U it holds that $U \subseteq W_\cap = W_1 \cap W_2$. By minimality of a closure, the set W_\cap is not a closure of U and hence there is a vertex v not in W_\cap such that $|N_{W_\cap}(v)| \geq 2$. But this implies that the set W_i satisfying $v \notin W_i$ is not closed, which contradicts the initial assumption. \square

In light of Proposition 3.5, for a set U we write $\text{Cl}(U)$ to denote the unique closure of U . We now show that if the underlying graph G is sparse, then the closure of a set U is not much larger than U itself.

Lemma 3.6. *Let G be an $(\ell, 1/2 - \varepsilon)$ -sparse graph for $\ell \in \mathbb{N}^+$ and $\varepsilon > 0$. Then any set $U \subseteq V$ of size at most $\ell/4$ satisfies $|\text{Cl}(U)| \leq 4|U| \leq \ell$.*

Proof. Let us consider the following process: set $U_0 = U$ and $i = 0$. While there is a vertex $v \in V \setminus U_i$ satisfying $|N_{U_i}(v)| \geq 2$, set $U_{i+1} = U_i \cup \{v\}$ and increment i .

Suppose this process terminates after t iterations. Clearly the final set U_t contains U and is closed. Furthermore it is a minimal set with these properties and is hence the closure of U . As we add at least 2 edges to $E(U_i)$ in every iteration it holds that $|E(U_i)| \geq 2i$ and, as we add a single vertex in each iteration, we have $|U_i| = |U| + i$. Suppose $t \geq 3|U|$. In iteration $i = 3|U|$ it holds that $|E(U_i)| \geq 6|U|$, while $|U_i| = 4|U| \leq \ell$. Hence $|E(U_i)| \geq 3|U_i|/2$, which contradicts the assumption on sparsity. We may thus conclude that the closure of a set U is of size at most $4|U|$, as claimed. \square

In order to design a strategy for the adversary, we prove that it is always possible to extend a colouring on a small closed set of vertices to a slightly larger set.

Lemma 3.7. *Let $G = (V, E)$ be an $(\ell, 1/2 - \varepsilon)$ -sparse graph, let $U \subseteq V$ be a closed set of size at most ℓ and suppose χ is a proper 4-colouring of $G[U]$. Then χ can be extended to $G[W]$ for any set $W \supseteq U$ of size at most $|W| \leq \ell$.*

Proof. By induction on $s = |W \setminus U|$. The statement clearly holds for $s = 0$. For the inductive step, let $v \in W \setminus U$ be such that $|N_{W \setminus U}(v)| \leq 2$, as guaranteed to exist by the assumption on sparsity since it holds that $|E(W \setminus U)| \leq (3/2 - \varepsilon)|W \setminus U|$. By the inductive hypothesis there is an extension of χ to the set $W \setminus \{v\}$.

As U is closed, the vertex v has at most one neighbour in U . Hence it holds that $|N_W(v)| \leq 3$ and we may thus extend the colouring to v . \square

With Lemmas 3.6 and 3.7 at hand we are ready to prove Theorem 3.1.

Proof of Theorem 3.1. Let us describe a strategy for the adversary. At any point in the game the adversary maintains a partial 4-colouring χ to the closure of the vertices U the prover remembers. Whenever the prover queries a vertex in the closure the adversary responds accordingly. If a vertex v outside the closure of U is queried, then the adversary extends χ to $\text{Cl}(U \cup \{v\})$ by virtue of Lemma 3.7 and responds accordingly. Finally, if the prover forgets the colouring of a vertex $u \in U$, then we shrink our closure to the closure of $U \setminus \{u\}$. Here we use the fact that by minimality of the closure it holds that $\text{Cl}(U) \supseteq \text{Cl}(U \setminus \{u\})$.

As, by Lemma 3.6, the closure of a set U is at most a factor 4 larger than U and, by Lemma 3.7, we may maintain a valid 4-colouring to the closure as long as the size of the closure is bounded by ℓ , the prover cannot win the game if $w \leq \ell/4$. Therefore, every resolution refutation of the 4-colourability formula defined over an $(\ell, 1/2 - \varepsilon)$ -sparse graph contains a clause that mentions at least $\ell/4$ vertices and hence has width at least $\ell/4$. \square

Combining Theorems 3.2 and 3.3 with the mentioned size-width relation we obtain optimal $\exp(\Omega(n))$ resolution size lower bounds for constant d .

4 Polynomial Calculus Lower Bounds: The General Framework

In this section we provide a proof overview of Theorem 1.1 and then revisit the general framework, as introduced by Alekhovich and Razborov [AR03], for obtaining polynomial calculus lower bounds.

4.1 Proof Overview

As outlined in the introduction, the construction of our pseudo-reduction \tilde{R} for the colouring formula follows the general paradigm introduced by [AR03] which has been used in several subsequent papers [GL10a, GL10b, MN15]. The idea is that given an initial set of polynomials \mathcal{P} , for every monomial m we identify a subset $S(m)$ of polynomials in \mathcal{P} that are in some sense relevant to m . Then we define \tilde{R} on the monomial m as the reduction modulo the ideal $\langle S(m) \rangle$ generated by the polynomials, and extend \tilde{R} linearly to arbitrary polynomials. The goal is to show that \tilde{R} satisfies properties 1-3 in Lemma 2.2, which typically requires showing that S satisfies two main technical properties. The first property is captured by what we call a *satisfiability lemma*, which states that if the degree of m is at most some parameter D , then the associated set $S(m)$ is satisfiable and is in some sense well-structured. The second property is described by the *reducibility lemma*, which states that for every ideal I that is generated by a well-structured set of polynomials that contains $S(m)$ and is satisfiable, it holds that $R_I(m) = R_{\langle S(m) \rangle}(m)$. Once these lemmas are in place, and as long as S satisfies some other simple conditions, a degree lower bound of D follows by an argument presented in [AR03].

We formalise these arguments in Lemma 4.6 by extracting the essential parts of the Alekhovich–Razborov [AR03] framework and making explicit the properties the map S must satisfy in order for the proof of the lower bound to go through. Apart from the properties in Definition 4.1 capturing some type of multilinearity, monotonicity and closedness of S and ensuring that axioms with leading monomial m are in $S(m)$, we introduce in Lemma 4.6 two sufficient conditions for the lower bound to follow: the first corresponds to the satisfiability lemma and the second to the reducibility lemma. With this in hand, the ensuing sections can focus on defining the map S and proving that it satisfies the conditions, without having to refer to reduction operators.

4.2 Revisiting the Alekhovich–Razborov Framework

We now revisit the general framework introduced by Alekhovich and Razborov [AR03] for proving polynomial calculus lower bounds. We extract from their proofs a formal statement that if the function S , mapping monomials to relevant subsets of axioms, satisfies two conditions, namely the *satisfiability* and the *reducibility conditions*, along with some natural conditions as discussed in the following, then this implies polynomial calculus degree lower bounds.

Before discussing the additional conditions that the mapping S needs to satisfy, we need to fix an *admissible* order $<$ on the monomials in $\mathbb{F}[x_1, \dots, x_n]$ as defined in Section 2.3. Recall that an admissible order is a total order that respects multiplication, that is, if $m_1 < m_2$ then it holds that $mm_1 < mm_2$. The *leading monomial* of a polynomial p is the largest monomial according to $<$ that appears in p . For the remainder of this section we implicitly assume that monomials in $\mathbb{F}[x_1, \dots, x_n]$ are ordered according to an admissible order $<$. Let us further adopt the convention that whenever we write a polynomial p as a sum of monomials $p = \sum_i a_i m_i$, then $a_i \in \mathbb{F}$ are field elements and the $m_i \in \mathbb{F}[x_1, \dots, x_n]$ are monomials ordered such that $m_i < m_j$ for all $j < i$. In particular, m_1 is the leading monomial of p .

Our first additional requirement on S is that it maps monomials according to the variables in the monomial, so that if two monomials m and m' both contain the same set of variables, then $S(m) = S(m')$. We further require that S is in some sense monotone, namely, for any variable x , we require that if $S(m') \subseteq S(m)$ for monomials $m' < m$, then it also holds that $S(xm') \subseteq S(xm)$. Moreover, the image of S should consist of sets that are “closed” in the sense that for any m , the set of axioms $S(m)$ contains the union of all $S(m')$ where $m' < m$ and $\text{Vars}(m') \subseteq \text{Vars}(S(m))$. Finally, we require that if m is the leading monomial of an axiom p , then $p \in S(m)$. These properties are formalised in the following definition.

Definition 4.1 (Support). Let \mathcal{P} be a set of polynomials over $\mathbb{F}[x_1, \dots, x_n]$, let $<$ be an admissible order on the monomials in $\mathbb{F}[x_1, \dots, x_n]$, and let $S: 2^{\{x_1, \dots, x_n\}} \rightarrow 2^{\mathcal{P}}$ be a function that maps subsets $Y \subseteq \{x_1, \dots, x_n\}$ of variables to subsets $S(Y) \subseteq \mathcal{P}$ of polynomials. For a monomial m we write $S(m)$ for $S(\text{Vars}(m))$ and we say that S is a \mathcal{P} -support if the following holds.

1. For every variable x and for all monomials m and m' such that $m' < m$, it holds that if $S(m') \subseteq S(m)$, then $S(xm') \subseteq S(xm)$.
2. For all monomials m and m' such that $m' < m$, it holds that if $\text{Vars}(m') \subseteq \text{Vars}(S(m))$, then $S(m') \subseteq S(m)$.
3. For all $p \in \mathcal{P}$, it holds that $p \in S(m)$, where m is the leading monomial in p .

If the set of polynomials \mathcal{P} is not essential we call a \mathcal{P} -support simply a *support*. Let us record four technical claims that follow from the above properties of a support and that we will later use to prove that an appropriate map S implies a polynomial calculus degree lower bound.

The first claim is that if m_1 is the leading monomial of an axiom p , not only does $S(m_1)$ contain p but it also contains $S(m_i)$ for all monomials m_i that appear in p .

Claim 4.2. If S is a \mathcal{P} -support, then for all $p \in \mathcal{P}$, if $p = \sum_i a_i m_i$, then $\{p\} \cup \bigcup_i (S(m_i)) \subseteq S(m_1)$.

Proof. By property 3 of Definition 4.1 it holds that $p \in S(m_1)$. This implies that $\text{Vars}(m_i) \subseteq \text{Vars}(p)$ is contained in $\text{Vars}(S(m_1))$ and thus for $i \neq 1$ since $m_i < m_1$, property 2 of Definition 4.1 implies that $S(m_i) \subseteq S(m_1)$. \square

The next claim states that if the variables of m' is a subset of the variables of m , then $S(m') \subseteq S(m)$.

Claim 4.3. If S is a support, then for all monomials m and m' such that $\text{Vars}(m') \subseteq \text{Vars}(m)$ it holds that $S(m') \subseteq S(m)$.

Proof. Note that $S(1) \subseteq S(m)$ by property 2 of Definition 4.1, where we use that $1 < m$. Thus, by inductively applying property 1 of Definition 4.1 and using the fact that $<$ respects multiplication, we obtain that $S(m') \subseteq S(m' \cdot m)$. Finally, since $\text{Vars}(m') \subseteq \text{Vars}(m)$, it holds that $\text{Vars}(m') = \text{Vars}(m' \cdot m)$ and hence $S(m) = S(m' \cdot m) \supseteq S(m')$. \square

We also prove that monomials m' that appear when m is reduced modulo $S(m)$ must be such that $S(m') \subseteq S(m)$ and that $S(xm') \subseteq S(xm)$.

Claim 4.4. If S is a support, then for all monomials m and m' such that m' appears in $R_{\langle S(m) \rangle}(m)$ it holds that $S(m') \subseteq S(m)$ and $S(xm') \subseteq S(xm)$ for any variable x .

Proof. We assume $m' \neq m$, and thus $m' < m$, otherwise the claim follows trivially. Moreover, we observe that it suffices to show that $S(m') \subseteq S(m)$, since we can then use property 1 of Definition 4.1 to conclude that $S(xm') \subseteq S(xm)$.

We now argue that $\text{Vars}(m') \subseteq \text{Vars}(m) \cup \text{Vars}(S(m))$. Towards contradiction, suppose this is not the case. Let ρ denote the assignment that assigns all variables in m' that are not in $\text{Vars}(S(m)) \cup \text{Vars}(m)$ to 0. Recall that $R_{\langle S(m) \rangle}(m)$ denotes the reduction of m modulo $\langle S(m) \rangle$ and that there is a unique representation $m = q + R_{\langle S(m) \rangle}(m)$ with $q \in \langle S(m) \rangle$. No variable in m nor in any of the generators of $\langle S(m) \rangle$ is assigned by ρ , so $m \upharpoonright_\rho = m$ and $q \upharpoonright_\rho \in \langle S(m) \rangle$. Moreover, we have $m' \upharpoonright_\rho = 0$ so $R_{\langle S(m) \rangle}(m) \upharpoonright_\rho \neq R_{\langle S(m) \rangle}(m)$. But this contradicts that the representation $m = q + R_{\langle S(m) \rangle}(m)$ is unique, and thus $\text{Vars}(m') \subseteq \text{Vars}(m) \cup \text{Vars}(S(m))$.

Write $m' = m'_1 \cdot m'_2$ such that $\text{Vars}(m'_1) \subseteq \text{Vars}(m)$ and $\text{Vars}(m'_2) \subseteq \text{Vars}(S(m))$. Observe that $m'_2 < m$ since $1 < m' < m$ and the order $<$ respects multiplication as it is admissible. Hence by property 2 of Definition 4.1 applied to m'_2 and m it follows that $S(m'_2) \subseteq S(m)$. By iteratively applying property 1 of Definition 4.1 for each $x \in \text{Vars}(m'_1)$, and again using that the order $<$ respects multiplication, we can conclude that $S(m') = S(m'_1 \cdot m'_2) \subseteq S(m'_1 \cdot m)$. Finally, since $\text{Vars}(m'_1) \subseteq \text{Vars}(m)$, we have that $S(m'_1 \cdot m) = S(m)$, and therefore $S(m') \subseteq S(m)$. \square

From this claim we can deduce that if irreducibility is preserved modulo some larger ideal, then the reduced polynomials are the same as done in the following.

Claim 4.5. Let S be a \mathcal{P} -support and let $\mathcal{Q} \subseteq \mathcal{P}$. Suppose that every monomial m irreducible modulo $\langle S(m) \rangle$ that satisfies $S(m) \subseteq \mathcal{Q}$ is also irreducible modulo $\langle \mathcal{Q} \rangle$. Then for all monomials m' such that $S(m') \subseteq \mathcal{Q}$ it holds that $R_{\langle \mathcal{Q} \rangle}(m') = R_{\langle S(m') \rangle}(m')$.

Proof. Let m' be an arbitrary monomial satisfying $S(m') \subseteq \mathcal{Q}$ and suppose that any irreducible monomial m modulo $\langle S(m) \rangle$ such that $S(m) \subseteq \mathcal{Q}$ is also irreducible modulo $\langle \mathcal{Q} \rangle$. We want to prove that $R_{\langle \mathcal{Q} \rangle}(m') = R_{\langle S(m') \rangle}(m')$.

Let us write $R_{\langle S(m') \rangle}(m') = \sum_i a_i m_i$. Note that by definition each such m_i is irreducible modulo $\langle S(m') \rangle$. As by Claim 4.4 it holds that $S(m_i) \subseteq S(m')$ we may conclude that each m_i is irreducible modulo $\langle S(m_i) \rangle$. Thus, by assumption along with the inclusions $S(m_i) \subseteq S(m') \subseteq \mathcal{Q}$, each m_i is irreducible modulo $\langle \mathcal{Q} \rangle$. This implies that $R_{\langle S(m') \rangle}(m') = R_{\langle \mathcal{Q} \rangle}(R_{\langle S(m') \rangle}(m')) = R_{\langle \mathcal{Q} \rangle}(m')$, using once more that $S(m') \subseteq \mathcal{Q}$. \square

We can now make the formal claim that polynomial calculus degree lower bounds follow from the existence of a support S satisfying the two conditions corresponding to the satisfiability lemma and the reducibility lemma. This fact is implicit in [AR03] and we make it explicit below. It is worth remarking that Filmus [Fil19] establishes a lemma of similar flavour. However, in contrast to Filmus, we do not introduce another layer of abstraction but rather state the lemma directly in terms of S .

Lemma 4.6. Let $D \in \mathbb{N}^+$, let \mathcal{P} be a set of polynomials over $\mathbb{F}[x_1, \dots, x_n]$ of degree at most D , and denote by S a \mathcal{P} -support. If all monomials m of degree at most D satisfy that

1. Satisfiability condition: the set of axioms $S(m)$ is satisfiable; and
2. Reducibility condition: for all monomials m' , if $S(m') \subseteq S(m)$, then m' is reducible modulo $\langle S(m') \rangle$ if and only if m' is reducible modulo $\langle S(m) \rangle$;

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

Proof. Let \tilde{R} be the operator defined on monomials m by $\tilde{R}(m) = R_{\langle S(m) \rangle}(m)$ and extended by linearity to polynomials. Our goal is to show that \tilde{R} is a degree- D pseudo-reduction for \mathcal{P} . That

is, according to Definition 2.1, we need to show that $\tilde{R}(1) = 1$, that \tilde{R} maps all polynomials in \mathcal{P} to 0, and that for every monomial m of degree at most $D - 1$ and every variable x , it holds that $\tilde{R}(xm) = \tilde{R}(x\tilde{R}(m))$. If we manage to show these properties, then we can appeal to Lemma 2.2 to reach the desired conclusion.

By the satisfiability condition we have that $S(1)$ is satisfiable, that is, the set of polynomials $S(1)$ has a common root. This implies that the entire ideal $\langle S(1) \rangle$ has a common root and, therefore, 1 is not in $\langle S(1) \rangle$ and hence is not reducible to 0 modulo $\langle S(1) \rangle$. As 1 is the smallest monomial in the order we conclude that $\tilde{R}(1) = R_{\langle S(1) \rangle}(1) = 1$.

To see that \tilde{R} maps each polynomial $p \in \mathcal{P}$ to 0, let $p = \sum_j a_j m_j$. By Claim 4.2 it follows that $p \in S(m_1)$ and that $S(m_j) \subseteq S(m_1)$ for all j . It therefore holds that

$$\tilde{R}(p) = \sum_j a_j R_{\langle S(m_j) \rangle}(m_j) \quad [\text{by definition of } \tilde{R} \text{ and } p] \quad (4.1a)$$

$$= \sum_j a_j R_{\langle S(m_1) \rangle}(m_j) \quad [\text{by Claim 4.5 and the reducibility condition}] \quad (4.1b)$$

$$= R_{\langle S(m_1) \rangle} \left(\sum_j a_j m_j \right) \quad [\text{by linearity of } R_{\langle S(m_1) \rangle}] \quad (4.1c)$$

$$= R_{\langle S(m_1) \rangle}(p) \quad [\text{by definition of } p] \quad (4.1d)$$

$$= 0, \quad [\text{since } p \in S(m_1)] \quad (4.1e)$$

where we note that in order to apply the reducibility condition we use the assumption that p has degree at most D .

Finally, we need to show that for every monomial m of degree at most $D - 1$ and every variable x , it holds that $\tilde{R}(xm) = \tilde{R}(x\tilde{R}(m))$. Let $\tilde{R}(m) = R_{\langle S(m) \rangle}(m) = \sum_j a_j m_j$. By definition of \tilde{R} we have that

$$\tilde{R}(x\tilde{R}(m)) = \sum_j a_j \tilde{R}(x m_j) = \sum_j a_j R_{\langle S(x m_j) \rangle}(x m_j). \quad (4.2)$$

We now argue that since m is a monomial of degree at most $D - 1$, reducing modulo $\langle S(x m_j) \rangle$ or $\langle S(x m) \rangle$ results in the same polynomial. More formally, we claim that

$$R_{\langle S(x m_j) \rangle}(x m_j) = R_{\langle S(x m) \rangle}(x m_j) \quad (4.3)$$

follows from Claim 4.5 with $\mathcal{Q} = S(x m)$ and $m' = x m_j$. Let us check that the conditions of Claim 4.5 are satisfied. We need to check that

1. any monomial \tilde{m} irreducible modulo $\langle S(\tilde{m}) \rangle$ satisfying that $S(\tilde{m}) \subseteq S(x m)$ is also irreducible modulo $\langle S(x m) \rangle$, and
2. $S(x m_j) \subseteq S(x m)$.

The latter follows immediately from Claim 4.4 and item 1 follows from the reducibility condition: since the monomial $x m$ has degree at most D , it holds for all monomials \tilde{m} , if $S(\tilde{m}) \subseteq S(x m)$, then \tilde{m} is irreducible modulo $\langle S(\tilde{m}) \rangle$ if and only if \tilde{m} is irreducible modulo $\langle S(x m) \rangle$. We may thus conclude (4.3) from Claim 4.5 applied with $\mathcal{Q} = S(x m)$ and $m' = x m_j$.




Figure 1: A set U with a 3-hop to the left, a 2-hop on top and a lasso to the right.

We finish the proof of Lemma 4.6 by noting that

$$\tilde{R}(x\tilde{R}(m)) = \sum_j a_j R_{\langle S(xm) \rangle}(xm_j) \quad [\text{by (4.2) and (4.3)}] \quad (4.4a)$$

$$= R_{\langle S(xm) \rangle} \left(\sum_j a_j xm_j \right) \quad [\text{by linearity of } R_{\langle S(xm) \rangle}] \quad (4.4b)$$

$$= R_{\langle S(xm) \rangle} (xR_{\langle S(m) \rangle}(m)) \quad [\text{as } R_{\langle S(m) \rangle}(m) = \sum_j a_j m_j] \quad (4.4c)$$

$$= R_{\langle S(xm) \rangle}(xm) \quad [\text{by Observation 2.3, using Claim 4.3}] \quad (4.4d)$$

$$= \tilde{R}(xm). \quad (4.4e)$$

This establishes that \tilde{R} is a degree- D pseudo-reduction for \mathcal{P} as defined in Definition 2.1. We may thus appeal to Lemma 2.2 to obtain the desired polynomial calculus refutation degree lower bound for \mathcal{P} . This completes the proof. \square

5 Closure for Graph Colouring

While the previous section was fairly generic and we made few assumptions on the polynomial system \mathcal{P} to be refuted, we now turn our attention to the colouring formula and make the first steps towards defining a support S for it. The goal of this section is to introduce the *closure* of a vertex set, a crucial notion in the definition of a support S for the colouring formula.

In order to state the definition of our closure we need to introduce some terminology. Let $G = (V, E)$ be a graph. A *walk of length τ* in G is a tuple of vertices $(v_1, \dots, v_{\tau+1})$ satisfying $(v_i, v_{i+1}) \in E$ for all $i \in [\tau]$. A *simple path* is a walk where all vertices are distinct and a *simple cycle* is a walk $(v_1, \dots, v_{\tau+1})$ of length $\tau \geq 3$ where $v_1 = v_{\tau+1}$ and all other vertices are distinct.

Suppose the set of vertices V has a linear order $<$ on V . An *increasing (decreasing) path* in G is a simple path (v_1, \dots, v_τ) where $v_i < v_{i+1}$ ($v_i > v_{i+1}$) for all $i \in [\tau - 1]$. For vertices $u, v \in V$ we say that v is a *descendant of u* if there exists a decreasing path from u to v , and for a set of vertices $U \subseteq V$ we say that v is a *descendant of U* if it is a descendant of some vertex in U . We write $\text{Desc}(U)$ to denote the set of all descendants of U . We define every vertex to be a descendant of itself so that $U \subseteq \text{Desc}(U)$.

The definition of our closure, besides the notion of descendants, also requires the notions of a *hop* and a *lasso*: 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 (in the case of cycles, the two endpoints coincide), while all other vertices are not in U . Similarly a *lasso with respect to U* is a paw graph, that is, a walk $(v_1, v_2, v_3, v_4, v_5)$ with $v_2 = v_5$ and all other vertices being distinct, such that only v_1 is in U . See Figure 1 for an example of a set with some hops and a lasso.

Algorithm 1 A procedure to obtain the closure of a given set U .

```

1: procedure CLOSURE( $U$ )
2:    $W_0 \leftarrow \text{Desc}(U)$ 
3:    $i \leftarrow 0$ 
4:   while exists 2-, 3-, 4-hop or lasso  $Q_{i+1}$  with respect to  $W_i$  do
5:      $W_{i+1} \leftarrow \text{Desc}(W_i \cup V(Q_{i+1}))$ 
6:      $i \leftarrow i + 1$ 
7:    $W_{\text{end}} \leftarrow W_i$ 
8:   return  $W_{\text{end}}$ 

```

With these notions in place we can now define a closure of a set $U \subseteq V$. After showing the closure is unique, we define a process that given U constructs its closure and then prove that in sparse graphs the closure of U satisfies two important properties:

1. it is the set of descendants of a set that is not much larger than U , as long as U itself is not too large, and,
2. if we remove the closure from the graph, then any small enough vertex set has a specially structured proper 3-colouring.

These two properties will then be used to prove that the satisfiability condition and the reducibility condition which will allow us to apply Lemma 4.6.

Definition 5.1 (Closure). Let $G = (V, E)$ be a graph and let $U \subseteq V$. We say that U is *closed* if $U = \text{Desc}(U)$ and there are no 2-, 3-, 4-hops or lassos with respect to U . A *closure* of U is any minimal closed set that contains U .

Proposition 5.2. *Every set of vertices has a unique closure.*

Proof. Suppose there are two distinct closures W_1 and W_2 of U . As W_1 as well as W_2 contains U it holds that $U \subseteq W_\cap = W_1 \cap W_2$. However, by minimality, W_\cap is not a closure of U and hence either there is a 2-, 3-, 4-hop or lasso Q with respect to W_\cap or there is a vertex v in $\text{Desc}(W_\cap)$ that is not in W_\cap .

In the first case this implies that some subgraph of Q is either a 2-, 3-, 4-hop or a lasso with respect to W_1 or W_2 since a subgraph of a hop or lasso is again hop or a lasso. In the second case this implies that $W_1 \not\subseteq \text{Desc}(W_1)$ or $W_2 \not\subseteq \text{Desc}(W_2)$, since $v \in \text{Desc}(W_1) \cap \text{Desc}(W_2)$. Either way, this contradicts the assumption that both W_1 and W_2 are closed. \square

In light of Proposition 5.2, for a set U we write $\text{Cl}(U)$ to denote the unique closure of U . In order to establish a bound on the size of the closure the algorithmic, but otherwise equivalent, description in Algorithm 1 will be useful.

Clearly the set of vertices W_{end} returned by Algorithm 1 is closed and minimal by construction: only adding part of a hop or a lasso results in a smaller hop. Hence $\text{Cl}(U) = W_{\text{end}}$ by Proposition 5.2.

The main property we use of our notion of closure is that the neighbourhood of a closed set W is very structured: since there are no 2-hops with respect to W , every vertex in $N(W) \setminus W$ has a single neighbour in W , and since there are no 3-hops with respect to W the neighbourhood of W is an independent set. The absence of longer 4-hops and lassos imply similar, more technical properties for sets of vertices that are connected to W via short paths. We leverage this structure to argue that after removing a closed set W from a sparse graph, all small vertex sets have a proper 3-colouring

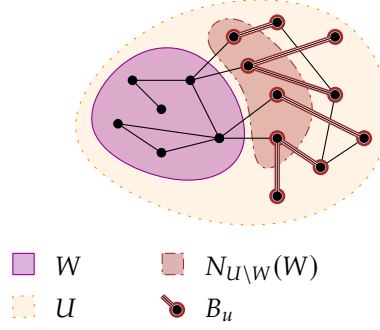


Figure 2: A depiction of the graphs B_u as defined in the proof of Lemma 5.3.

with the additional property that the neighbours of each vertex at distance 1 from W are coloured with a single colour.

Lemma 5.3. *Suppose that $G = (V, E)$ is $(\ell, 1/3\Delta)$ -sparse. Let $U, W \subseteq V$ be vertex sets such that $|U| \leq \ell$, W is closed, and every vertex in $N_{U \setminus W}(W)$ has degree at most Δ in $G[U \setminus W]$. Then there exists a proper 3-colouring of the subgraph $G[U \setminus (W \cup N_U(W))]$ such that for every $u \in N_{U \setminus W}(W)$, the set $N_{U \setminus W}(u)$ is monochromatic.*

Proof. We start the proof by establishing the following two properties.

1. First, since there are no lassos in U with respect to W , for every vertex $u \in N_{U \setminus W}(W)$, it holds that the neighbours $N_{U \setminus W}(u)$ of u are an independent set. Hence the graph $B_u = G[\{u\} \cup N_{U \setminus W}(u)]$ is a star with center u and leaves $N_{U \setminus W}(u)$.
2. Second, as there are no 3-, 4-hops in U with respect to W , the graphs $\{B_u \mid u \in N_{U \setminus W}(W)\}$ as just defined are pairwise vertex disjoint.

See Figure 2 for an illustration.

Let G' be the graph obtained from $G[U \setminus W]$ by contracting each star B_u to a vertex \tilde{u} . Note that from a 3-colouring $\tilde{\chi}$ of G' we can define a 3-colouring χ of $G[U \setminus (W \cup N_U(W))]$ with the desired property: we claim that the 3-colouring $\chi: U \setminus (W \cup N_U(W)) \rightarrow \{1, 2, 3\}$ defined by

$$\chi(w) = \begin{cases} \tilde{\chi}(\tilde{u}) & \text{if } w \text{ is a leaf of a star } B_u, \\ \tilde{\chi}(w) & \text{otherwise} \end{cases} \quad (5.1)$$

is a proper 3-colouring of $G[U \setminus (W \cup N_U(W))]$. Note that since the graphs $\{B_u \mid u \in N_{U \setminus W}(W)\}$ are pairwise disjoint, the colouring χ assigns precisely one colour per vertex and, as each of the induced graphs B_u is a star, the colouring is proper. Finally, for all vertices $u \in N_{U \setminus W}(W)$ the set $N_{U \setminus W}(u)$ is monochromatic by construction.

It remains to argue that G' is 3-colourable. Here we want to use the extreme sparsity of G : we enforce such extreme sparsity on G that even though G' is obtained from G by contractions and hence has higher edge density than G , the graph G' is still $(|V(G')|, 1/3)$ -sparse. Thus by Lemma 2.6 it follows that G' is 3-colourable.

Let us verify that G' is indeed $(|V(G')|, 1/3)$ -sparse, that is, we need to argue that every subset $T' \subseteq V(G')$ satisfies $|E(T')| \leq (1 + 1/3)|T'|$. Fix such a subset $T' \subseteq V(G')$ and let T be the preimage of T' in the contraction. We estimate $|E(T')|$ in terms of $|E(T)|$. Let $\{u_1, \dots, u_t\} = T \cap N_{U \setminus W}(W)$. Observe that in T' each star B_{u_i} is contracted to a vertex \tilde{u}_i . Let $s = \sum_{i=1}^t (|V(B_{u_i})| - 1)$. It holds that

$$s \leq (\Delta - 1)t \leq (\Delta - 1)|T'|, \quad (5.2)$$

where we use the assumption that the degree of every vertex u_i in $G[U \setminus W]$ is bounded by Δ . Furthermore, it holds that $|T| = |T'| + s$ and, because all the edges in the stars B_{u_i} are contracted, that $|E(T)| \geq |E(T')| + s$.

By assumption, G is $(\ell, 1/3\Delta)$ -sparse and $|T| \leq |U| \leq \ell$. This implies that $|E(T)| \leq (1 + 1/3\Delta)|T|$ which, if combined with the above, gives $|E(T')| + s \leq (1 + 1/3\Delta)(|T'| + s)$. Using (5.2) we may conclude that

$$|E(T')| \leq \left(1 + \frac{1}{3\Delta}\right)|T'| + \frac{s}{3\Delta} \leq \left(1 + \frac{1}{3\Delta}\right)|T'| + \frac{(\Delta - 1)|T'|}{3\Delta} = \left(1 + \frac{1}{3}\right)|T'|. \quad (5.3)$$

Therefore, G' is $(|G'|, 1/3)$ -sparse and thus, by Lemma 2.6, it follows that G' is 3-colourable. \square

The second property of the closure that we need is that we have some control on how large it is, which is necessary to show that the satisfiability condition of Lemma 4.6 holds. We establish an upper bound on the size of the closure by proving that for sparse graphs the closure of a not too large set U is the set of descendants of a set Z that is not much larger than U , as stated in the lemma below. Combining this with Lemma 2.6, it follows that in order to establish that the set $\text{Cl}(U)$ is 3-colourable—and thus $\text{Col}(G[\text{Cl}(U)], 3)$ is satisfiable—it is sufficient to prove an upper bound on the size of the set of descendants of Z . Jumping ahead a bit, we will be able to establish such an upper bound in the next section by choosing an appropriate ordering of the vertices.

Lemma 5.4. *Suppose that $G = (V, E)$ has a linear order on V and is $(\ell, 1/3a)$ -sparse for $a \geq 2$. Let $U \subseteq V$ be a set of size $|U| \leq \ell/25a$ such that any decreasing path in $G[V \setminus U]$ has at most a vertices. Then there exists a set $Z \subseteq V$ such that $Z \supseteq U$, $|Z| \leq 25|U|$ and $\text{Cl}(U) = \text{Desc}(Z)$.*

Proof. Recall from Algorithm 1 that the closure of a set $U \subseteq V$ can be defined to be the final set in a sequence $(W_0, W_1, \dots, W_{\text{end}})$, where $W_0 = \text{Desc}(U)$ and W_{i+1} is obtained from W_i by appending a 2-, 3-, 4-hop or a lasso with respect to W_i and then taking the descendants of the resulting set of vertices. A key observation is that adding such hops or lassos to W_i adds more edges to the induced graph $G[W_i]$ than vertices, thus increasing the edge density. As the graph G is locally sparse we can conclude that the sequence $(W_0, W_1, \dots, W_{\text{end}})$ needs to be rather short, which allows us to argue the size upper bound on Z .

In the following, for each W_i we identify a vertex set U_i such that the edge density of the graph $G[U_i]$ increases with i . The idea is as follows. Since the sets U_i grow very slowly and thus the local sparsity always applies, we will be able to conclude that the number of iterations in the construction of $\text{Cl}(U)$ is bounded. As the vertices in W_i are the descendants of the set U_i , it holds that the vertices in W_{end} are the descendants of a set which is not much larger than the initial set U , whereby the lemma follows.

Let us now implement this plan. We define U_i inductively as follows. Let $U_0 = U$ and let Q_i be the hop or lasso added to W_{i-1} at iteration $i \geq 1$. If we denote by u and v the endpoints of Q_i (where we could have $u = v$) and let P_u and P_v be two shortest decreasing paths from U_{i-1} to u and v , respectively, then it holds that $U_i = U_{i-1} \cup V(P_u \cup P_v \cup Q_i)$. See Figure 3 for an illustration.

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

Claim 5.5. For every vertex v in W_i , there exists a decreasing path in W_i from some vertex in U_i to v .

Proof. The proof is by induction on i . The base case $i = 0$ holds because $W_0 = \text{Desc}(U_0)$. For the induction step, suppose that the claim holds for $i - 1$. By definition, the vertices in $W_i \setminus W_{i-1}$ are descendants of a vertex in Q_i , and all vertices in Q_i are contained in U_i . The claim follows. \square




Figure 3: A depiction of the construction of U_i as defined in the proof of Lemma 5.4.

Next, we show that $|U_i|$ grows slowly with i and that the edge density $|E(U_i)|/|U_i|$ exceeds the sparsity threshold $(1 + 1/3a)$ after a small number of iterations.

Claim 5.6. For U_i as commented above it holds that $|U_i \setminus U_{i-1}| \leq 2a + |V(Q_i)| - 4$ and $|E(U_i)| \geq |E(U_{i-1})| + |U_i \setminus U_{i-1}| + 1$.

Proof. Let F be the graph defined by the union of the edges in P_u, P_v and Q_i . Since the paths P_u and P_v are decreasing paths, according to the statement of Lemma 5.4, they contain at most a vertices each, so $3 \leq |V(F)| \leq 2a + |V(Q_i)| - 2$. 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 $|V(F)| \leq 2a + |V(Q_i)| - 3$.

Since in both cases the number of added vertices is bounded $|U_i \setminus U_{i-1}| \leq 2a + |V(Q_i)| - 4$ and the number of edges in the subgraph induced by U_i can be lower bounded $|E(U_i)| \geq |E(U_{i-1})| + |U_i \setminus U_{i-1}| + 1$ the statement follows. \square

Recall that we want to show that the edge density of the induced subgraph $G[U_i]$ increases with i . Since G is sparse it thus follows that the number of iterations in the construction of $\text{Cl}(U)$ is bounded.

Let $s = |U|$. Towards contradiction, suppose that $i \geq 5s + 1$. Note that $|V(Q_i)| \leq 5$, so by Claim 5.6 we have $|U_i \setminus U_{i-1}| \leq 2a + 1$. This implies that

$$|U_{5s+1}| = s + \sum_{i=1}^{5s+1} |U_i \setminus U_{i-1}| \leq s + (5s + 1)(2a + 1) < 16a|U| < \ell, \quad (5.4)$$

and therefore, since G is $(\ell, 1/3a)$ -sparse, it holds that $|E(U_{5s+1})|/|U_{5s+1}| \leq 1 + 1/3a$. However, by

Claim 5.6, we have that

$$\frac{|E(U_{5s+1})|}{|U_{5s+1}|} = \frac{|E(U)| + \sum_{i=1}^{5s+1} |E(U_i) \setminus E(U_{i-1})|}{s + \sum_{i=1}^{5s+1} |U_i \setminus U_{i-1}|} \quad (5.5a)$$

$$\geq \frac{|E(U)| + \sum_{i=1}^{5s+1} (|U_i \setminus U_{i-1}| + 1)}{s + \sum_{i=1}^{5s+1} |U_i \setminus U_{i-1}|} \quad (5.5b)$$

$$\geq \frac{0 + \sum_{i=1}^{5s+1} (2a + 2)}{s + \sum_{i=1}^{5s+1} (2a + 1)} \quad (5.5c)$$

$$> 1 + \frac{1}{3a}, \quad (5.5d)$$

where for (5.5b) we use that $|E(U_i)| \geq |E(U_{i-1})| + |U_i \setminus U_{i-1}| + 1$, and then for (5.5c), we observe that the fraction decreases as $\sum_{i=1}^{5|U|+1} |U_i \setminus U_{i-1}|$ increases and thus, along with the bound $|U_i \setminus U_{i-1}| \leq 2a + 1$, we obtain the claimed inequality. For the final inequality (5.5d) we use $a \geq 2$. This contradicts the assumption that G is $(\ell, 1/3a)$ -sparse and it hence follows that $i \leq 5|U|$.

Let i_{end} be the last iteration and let $Z = U \cup \bigcup_{j \leq i_{\text{end}}} V(Q_j)$. We claim that $\text{Desc}(Z) = \text{Cl}(U)$ and $|Z| \leq 25|U|$. Indeed, at iteration i in the construction, W_i is the set of descendants of $U \cup \bigcup_{j \leq i} V(Q_j)$, and the hop or lasso Q_i added to W_{i-1} contains at most 4 vertices not already in W_i . Therefore, $\text{Desc}(Z) = \text{Cl}(U)$ and since $i_{\text{end}} \leq 5|U|$ it follows that $|Z| \leq |U| + 4 \cdot 5|U| \leq 25|U|$. \square

6 A Lower Bound for 3-Colourability on Sparse Random Graphs

We are now ready to prove a linear degree lower bound for polynomial calculus refutations of the claim that sparse random graphs are 3-colourable. In fact, we prove something slightly stronger: we show that if a graph is locally very sparse and has only few vertices of high degree, then it is hard for polynomial calculus to refute the claim that the graph is 3-colourable.

Theorem 6.1 (Main theorem). *Let c, Δ, k and ℓ be integers such that $c > k \geq 3$ and let $G = (V, E)$ be a c -colourable, $(\ell, 1/3\Delta)$ -sparse graph. For any set T_Δ such that $G[V \setminus T_\Delta]$ has maximum degree at most Δ , any polynomial calculus refutation of $\text{Col}(G, k)$ over any field requires degree $\ell/50\Delta^{c-1} - |T_\Delta|$.*

We defer the proof of Theorem 6.1 to Section 6.1, and first show how our results follow from this theorem. Intuitively, the above theorem holds because sparse graphs are locally 3-colourable (see Lemma 2.6) and hence the colouring formula defined over sparse graphs is locally satisfiable. Before diving into the proof of Theorem 6.1, let us see how we can use Lemma 2.7 to obtain degree lower bounds for random graphs. Note that the condition on the graph G in Theorem 6.1 is not inherently random. However, to the best of our knowledge, there are no explicit constructions of graphs that are this sparse. In fact, it is stated as an open problem in the survey by Hoory et al. [HLW06, Open problem 10.8] to explicitly construct such graphs.

Recall from Lemma 2.8 that if $d \geq 6$, then graphs sampled from the Erdős-Rényi random graph distribution $\mathbb{G}(n, d/n)$ or the random d -regular graph distribution $\mathbb{G}_{n,d}$ are asymptotically almost surely not 3-colourable. In light of this, the below statements are interesting in the parameter regime $d \geq 6$.

Corollary 6.2 (Colouring lower bound for random regular graphs). *There exists an absolute constant C such that for positive integers n and $d \geq 2$ satisfying $6d^3 \leq \log n$ the following holds. If G is a graph sampled from $\mathbb{G}_{n,d}$ and $k \geq 3$ is an integer, then asymptotically almost surely every polynomial calculus refutation of $\text{Col}(G, k)$ over any field requires degree $d^{-Cd} \cdot n$.*

Proof. Let $\Delta = d$ and $T_\Delta = \emptyset$. By Lemma 2.8 we have that asymptotically almost surely G is c -colourable for $c \leq 2d/\log d$, and by Lemma 2.7 that asymptotically almost surely G is $(\ell, 1/3d)$ -sparse for $\ell = (8d)^{-6d}n$. Note that to apply Lemma 2.7, we use $\varepsilon = 1/3d$ and $\delta = 1/2$. We can now apply Theorem 6.1 and conclude that any polynomial calculus refutation of $\text{Col}(G, k)$, over any field, requires degree $\ell/50\Delta^{c-1} - |T_\Delta| = (8d)^{-6d}n/50d^{2d/\log d-1} \geq d^{-Cd} \cdot n$, using that C is a large enough constant. \square

To prove the result for Erdős-Rényi random graphs we need the additional property that asymptotically almost surely there exists a small set T_Δ of vertices such that $G[V \setminus T_\Delta]$ has maximum degree at most Δ .

Lemma 6.3. *Let $G = (V, E)$ be a graph sampled from $\mathbb{G}(n, d/n)$ where $d = O(\log n)$. If $\Delta \geq d$ is such that $(\Delta/d)^\Delta = o(n)$, then asymptotically almost surely there exists a set T_Δ of size at most $(ed/\Delta)^\Delta \cdot 2en$ such that the maximum degree in $G[V \setminus T_\Delta]$ is at most $\Delta - 1$.*

The proof of Lemma 6.3 is mostly a standard calculation and we present it in Appendix C for completeness. We are now ready to prove our lower bound for Erdős-Rényi random graphs.

Corollary 6.4 (Colouring lower bound for Erdős-Rényi random graphs). *There exists an absolute constant C such that for $n \in \mathbb{N}^+$ and $d \in \mathbb{R}^+$ satisfying that $d > 1$ and $d^5 = o(\log n)$ the following holds. If G is a graph sampled from $\mathbb{G}(n, d/n)$ and $k \geq 3$ is an integer, then asymptotically almost surely every polynomial calculus refutation of $\text{Col}(G, k)$ over any field requires degree $d^{-Cd^5} \cdot n$.*

Proof. Fix $\Delta = (5d)^5$. By Lemma 2.8 we have that asymptotically almost surely G is c -colourable for $c \leq 2d/\log d$, and by Lemma 2.7 it holds asymptotically almost surely that G is (ℓ, ε) -sparse for $\ell = (4d)^{-4(5d)^5}n$ and $\varepsilon = 1/3\Delta$. Note that to apply Lemma 2.7, we can use $\delta = 1/4$, so that $(1 + \varepsilon)(1 + \delta)/\varepsilon \leq 4\Delta$.

Let $T_\Delta \subseteq V$ be a minimum size set such that $G[V \setminus T_\Delta]$ has maximum degree at most Δ . By Lemma 6.3, we have that

$$|T_\Delta| \leq (ed/\Delta)^\Delta \cdot 2en = \frac{2en}{(5^5 d^4/e)^{(5d)^5}} < \frac{n}{2 \cdot (5d)^{4(5d)^5}}, \quad (6.1)$$

where the last inequality follows since $(5/e)^{(5d)^5} \geq (5/e)^{5^5} > 4e$.

We can now apply Theorem 6.1 and conclude that any polynomial calculus refutation of $\text{Col}(G, k)$, over any field, requires degree

$$\frac{\ell}{50\Delta^{c-1}} - |T_\Delta| \geq \frac{n}{50 \cdot (5d)^{10d/\log d} \cdot (4d)^{4(5d)^5}} - \frac{n}{2 \cdot (5d)^{4(5d)^5}} \quad (6.2a)$$

$$\geq \frac{n}{(5d)^{4(5d)^5}} - \frac{n}{2 \cdot (5d)^{4(5d)^5}} \quad (6.2b)$$

$$\geq d^{-Cd^5} \cdot n, \quad (6.2c)$$

where for the second inequality we use that $50 \cdot (5d)^{10d/\log d} < 2^{10} \cdot (10)^{10d} \leq 2^{(5d)^5}$ and for the last inequality we use that C is a large enough constant. \square

6.1 Proof of Main Theorem

Fix T_Δ such that $G[V \setminus T_\Delta]$ has maximum degree at most Δ and let $X = \{x_{v,i} \mid v \in V(G), i \in [k]\}$.

To prove Theorem 6.1, our goal is to define a $\text{Col}(G, k)$ -support S that maps monomials in the polynomial ring $\mathbb{F}[X]$ to subsets of $\text{Col}(G, k)$ such that Lemma 4.6 holds. For brevity, given a set $U \subseteq V$ 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 $\langle U \rangle$.

Given a monomial m , we let $V(m)$ denote the set of vertices mentioned by the variables in m . Moreover, given a linear order $<_v$ on the vertices, we say an admissible order $<$ of the monomials over the variables X of $\text{Col}(G, k)$ respects $<_v$ if for any colours $i, j \in [k]$ it holds that $x_{u,i} < x_{v,j}$ whenever $u <_v v$.

The main technical lemma we need in order to prove Theorem 6.1 is the reducibility lemma below, from which the reducibility condition of Lemma 4.6 will follow. The reducibility lemma implies, in particular, that reducing a monomial m modulo $\langle W \rangle$ for any closed set $W \supseteq V(m) \cup T_\Delta$ is the same as reducing modulo $\langle U \rangle$ for any superset $U \supseteq W$ that is not too large.

Lemma 6.5 (Reducibility lemma). *Let $G = (V, E)$ be a $(\ell, 1/3\Delta)$ -sparse graph with a linear order $<_v$ on V and consider an admissible order that respects $<_v$. If the vertex sets $W \subseteq U$ satisfy that W is closed, that the size of U is $|U| \leq \ell$, and that every vertex in $N_{U \setminus W}(W)$ has degree at most Δ in $G[U \setminus W]$, then for every monomial m such that $V(m) \subseteq W$, it holds that m is reducible modulo $\langle U \rangle$ if and only if m is reducible modulo $\langle W \rangle$.*

We postpone the proof of this lemma to the end of this section. In order to ensure that we are always reducing a monomial m by $\langle W \rangle$ for some closed set W that contains $V(m)$ and is such that every vertex in $N_{U \setminus W}(W)$ has degree at most Δ in $G[U \setminus W]$ for all $U \supseteq W$, we define the closure of a monomial to include the set T_Δ .

Definition 6.6 (Monomial closure). The *monomial closure* of a monomial m , denoted by $\text{Cl}_\Delta(m)$, is the vertex set $\text{Cl}(V(m) \cup T_\Delta)$.

Looking ahead, we will prove Theorem 6.1 by showing that the map S defined by mapping monomials m to $\text{Col}(G[\text{Cl}_\Delta(m)], k) \subseteq \text{Col}(G, k)$ is a $\text{Col}(G, k)$ -support and appealing to Lemma 4.6. To establish the satisfiability condition in Lemma 4.6, we must prove that $\text{Col}(G[\text{Cl}_\Delta(m)], k)$ is satisfiable whenever m is of low degree. Since G is sparse, by Lemma 2.6, it suffices to show that the monomial closure of m is not too large. This, in turn, will follow from the next lemma. Lemma 6.7 is an almost direct consequence of Lemma 5.4 and states that under suitable technical assumptions, the size of the monomial closure of m is closely related to the degree of m and the size of T_Δ .

Lemma 6.7 (Satisfiability lemma). *Let $a, \Delta, \ell \in \mathbb{N}^+$ such that $a \geq 2$ and let $G = (V, E)$ be a $(\ell, 1/3a)$ -sparse graph with a linear order on V . Let $T_\Delta \subseteq V$ be such that $G[V \setminus T_\Delta]$ has maximum degree at most Δ , $\text{Desc}(T_\Delta) \subseteq T_\Delta$, and that any decreasing path in $G[V \setminus T_\Delta]$ has at most a vertices. Then it holds for any monomial m that if $\text{Deg}(m) + |T_\Delta| \leq \ell/25a$, then $|\text{Cl}_\Delta(m)| \leq 50\Delta^{a-1} \cdot (\text{Deg}(m) + |T_\Delta|)$.*

Proof. Let $U = V(m) \cup T_\Delta$ and note that $\text{Cl}_\Delta(m) = \text{Cl}(U)$. Note that $|U| = |V(m) \cup T_\Delta| \leq \text{Deg}(m) + |T_\Delta| \leq \ell/25a$ and that any decreasing path in $G[V \setminus U]$ has at most a vertices. We can therefore apply Lemma 5.4 to deduce that there exists a set $Z \subseteq V$ such that $U \subseteq Z$, $\text{Cl}(U) = \text{Desc}(Z)$ and $|Z| \leq 25|U|$.

Note that since $Z \supseteq T_\Delta$ and all the descendants of vertices in T_Δ are in T_Δ , we have that $\text{Desc}(Z) = (\text{Desc}(Z \setminus T_\Delta) \setminus T_\Delta) \cup T_\Delta$. Moreover, since any vertex $v \in Z \setminus T_\Delta$ has degree at most Δ in $G[V \setminus T_\Delta]$, and since any decreasing path in $G[V \setminus T_\Delta]$ has at most a vertices, it follows that v has at most $2\Delta^{a-1}$ descendants in $V \setminus T_\Delta$. We thus have the upper bound $|(\text{Desc}(Z \setminus T_\Delta) \setminus T_\Delta)| \leq 2\Delta^{a-1} \cdot |Z \setminus T_\Delta|$ from which we conclude that

$$|\text{Desc}(Z)| \leq 2\Delta^{a-1} \cdot |Z \setminus T_\Delta| + |T_\Delta| \leq 2\Delta^{a-1} \cdot |Z| \leq 50\Delta^{a-1} \cdot (\text{Deg}(m) + |T_\Delta|),$$

as claimed in the lemma. \square

We are now ready to prove our main theorem, which we restate here for convenience.

Theorem 6.1 (Main theorem, restated). *Let c, Δ, k and ℓ be integers such that $c > k \geq 3$ and let $G = (V, E)$ be a c -colourable, $(\ell, 1/3\Delta)$ -sparse graph. For any set T_Δ such that $G[V \setminus T_\Delta]$ has maximum degree at most Δ , then any polynomial calculus refutation of $\text{Col}(G, k)$ over any field requires degree $\ell/50\Delta^{c-1} - |T_\Delta|$.*

Proof of Theorem 6.1. We start by defining a linear order on V . Let $\chi_c: V \setminus T_\Delta \rightarrow [c]$ be a proper c -colouring of $G[V \setminus T_\Delta]$. We let the order $<_v$ on V be any linear order that satisfies $u <_v v$ whenever $u \in T_\Delta$ and $v \in V \setminus T_\Delta$ and whenever $u, v \in V \setminus T_\Delta$ and $\chi_c(u) < \chi_c(v)$. Observe that any decreasing path in $G[V \setminus T_\Delta]$ has at most c vertices and that $\text{Desc}(T_\Delta) \subseteq T_\Delta$.

We can define an admissible ordering $<$ of the monomials over the variables $\{x_{v,i}\}_{v \in V, i \in [k]}$ of $\text{Col}(G, k)$ that respects $<_v$ as follows: for distinct vertices u, v let $x_{u,i} < x_{v,j}$ whenever $u <_v v$ and for variables associated with the same vertex u let $x_{u,i} < x_{u,j}$ whenever $i < j$. 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.

To prove Theorem 6.1 we use Lemma 4.6. For this, we show that the map $S: m \mapsto \text{Col}(G[\text{Cl}_\Delta(m)], k)$ is a $\text{Col}(G, k)$ -support in the sense of Definition 4.1 and that it satisfies the satisfiability condition and the reducibility condition in Lemma 4.6 for $D = \ell/(50\Delta^{c-1}) - |T_\Delta|$.

We start by proving, via the properties of the closure, that the map S is a $\text{Col}(G, k)$ -support. We actually show something slightly stronger, namely that S satisfies the following four properties, which are the same as those in Definition 4.1 except that we do not require that $m' < m$ in item 2 and item 3.

1. For all monomials m and m' such that $\text{Vars}(m) = \text{Vars}(m')$, it holds that $S(m) = S(m')$.
2. For every variable x and for all monomials m and m' , if $S(m') \subseteq S(m)$, then $S(xm') \subseteq S(xm)$.
3. For all monomials m and m' , if $\text{Vars}(m') \subseteq \text{Vars}(S(m))$, then $S(m') \subseteq S(m)$.
4. For all $p \in \mathcal{P}$, it holds that $p \in S(m)$, where m is the leading monomial in p .

Item 1 follows immediately from the definition of S , since the closure of a monomial m only depends on $V(m)$. Note that since $S(m) = \text{Col}(G[\text{Cl}_\Delta(m)], k)$ it holds that $S(m') \subseteq S(m)$ if and only if $\text{Cl}_\Delta(m') \subseteq \text{Cl}_\Delta(m)$ and, therefore, item 2 is equivalent to showing that if $\text{Cl}_\Delta(m') \subseteq \text{Cl}_\Delta(m)$, then $\text{Cl}_\Delta(xm') \subseteq \text{Cl}_\Delta(xm)$. Recall that $\text{Cl}_\Delta(m) = \text{Cl}(V(m) \cup T_\Delta)$ and hence, by minimality of closure, we obtain that

$$V(x) \cup V(m') \cup T_\Delta \subseteq \text{Cl}(V(x)) \cup \text{Cl}(V(m') \cup T_\Delta) \quad [\text{since } U \subseteq \text{Cl}(U)] \quad (6.3a)$$

$$\subseteq \text{Cl}(V(x)) \cup \text{Cl}(V(m) \cup T_\Delta) \quad [\text{as } \text{Cl}_\Delta(m') \subseteq \text{Cl}_\Delta(m) \text{ by assumption}] \quad (6.3b)$$

$$\subseteq \text{Cl}(V(xm) \cup T_\Delta), \quad (6.3c)$$

where the final equation relies on the fact that $\text{Cl}(A) \cup \text{Cl}(B) \subseteq \text{Cl}(A \cup B)$ which follows by minimality of the sets $\text{Cl}(A)$ and $\text{Cl}(B)$. This allows us to derive that

$$\text{Cl}_\Delta(xm') = \text{Cl}(V(x) \cup V(m') \cup T_\Delta) \quad [\text{by definition of monomial closure}] \quad (6.4a)$$

$$\subseteq \text{Cl}(\text{Cl}(V(xm) \cup T_\Delta)) \quad [\text{by (6.3)}] \quad (6.4b)$$

$$= \text{Cl}(V(xm) \cup T_\Delta) \quad [\text{since closure is idempotent by minimality}] \quad (6.4c)$$

$$= \text{Cl}_\Delta(xm), \quad [\text{by definition of monomial closure}] \quad (6.4d)$$

and thus item 2 holds. Observe furthermore that if $\text{Vars}(m') \subseteq \text{Vars}(S(m))$ then $V(m') \subseteq \text{Cl}_\Delta(m)$ by the definition of monomial closure and since $S(m) = \text{Col}(G[\text{Cl}_\Delta(m)], k)$. Thus, using again the observation that $S(m') \subseteq S(m)$ if and only if $\text{Cl}_\Delta(m') \subseteq \text{Cl}_\Delta(m)$, in order to conclude that item 3 holds it suffices to show that if $V(m') \subseteq \text{Cl}_\Delta(m)$, then $\text{Cl}_\Delta(m') \subseteq \text{Cl}_\Delta(m)$. By again using the minimality of closure and the fact that $U \subseteq \text{Cl}(U)$, we can conclude that

$$\text{Cl}_\Delta(m') = \text{Cl}(V(m') \cup T_\Delta) \subseteq \text{Cl}(\text{Cl}_\Delta(m) \cup T_\Delta) = \text{Cl}_\Delta(m). \quad (6.5)$$

Finally, item 4 follows easily from the definition of S . Indeed, if $p \in \text{Col}(G, k)$ is an edge axiom, say $x_{u,i}x_{v,i}$, then it holds that $S(x_{u,i}x_{v,i}) = \text{Col}(G[\text{Cl}(\{u, v\} \cup T_\Delta)], k) \ni p$; and if $p \in \text{Col}(G, k)$ is a vertex axiom ($\sum_{i=1}^k x_{v,i} - 1$ or $x_{v,i}x_{v,i'}$) or a Boolean axiom ($x_{v,i}^2 - x_{v,i}$) mentioning a vertex v and m is the leading monomial in p it holds that $S(m) = \text{Col}(G[\text{Cl}(\{v\} \cup T_\Delta)], k) \ni p$. Thus, the map S is a $\text{Col}(G, k)$ -support.

We now show that the satisfiability condition and the reducibility condition in Lemma 4.6 hold for the map S and for $D = \ell/(50\Delta^{c-1}) - |T_\Delta|$. We can assume $D \geq 2$, since otherwise the theorem is trivially true. Observe that the polynomials in $\text{Col}(G, k)$ have degree at most D .

To see that the satisfiability condition holds, note that by Lemma 6.7 every monomial m of degree at most D satisfies $|\text{Cl}_\Delta(m)| \leq 50\Delta^{c-1}(D + |T_\Delta|) = \ell$, where we use that any decreasing path in $G[V \setminus T_\Delta]$ has at most c vertices, that $\text{Desc}(T_\Delta) \subseteq T_\Delta$ holds, and that $G[V \setminus T_\Delta]$ has maximum degree at most Δ . Since G is $(\ell, 1/3\Delta)$ -sparse and $|\text{Cl}_\Delta(m)| \leq \ell$, it follows from Lemma 2.6 that the graph $G[\text{Cl}_\Delta(m)]$ is 3-colourable and so $\text{Col}(G[\text{Cl}_\Delta(m)], k)$ is satisfiable.

To establish the reducibility condition, let m and m' be monomials of degree at most D such that $S(m') \subseteq S(m)$, i.e., such that $\text{Cl}_\Delta(m') \subseteq \text{Cl}_\Delta(m)$. Note that as argued above, it holds that $|\text{Cl}_\Delta(m)| \leq \ell$ since m has degree at most D . Therefore, we can apply Lemma 6.5 with $W = \text{Cl}_\Delta(m')$ and $U = \text{Cl}_\Delta(m)$ to conclude that if m' is reducible modulo $\langle W \rangle = \langle S(m') \rangle$ if and only if it is also reducible modulo $\langle U \rangle = \langle S(m) \rangle$. Note that we use the fact that all vertices in $G[V \setminus T_\Delta]$ —and hence also all vertices in $G[U \setminus W]$ —have degree at most Δ .

With the satisfiability and reducibility conditions of Lemma 4.6 in hand, we conclude that every polynomial calculus refutation of $\text{Col}(G, k)$ requires degree strictly greater than D , as desired. \square

6.2 The Reducibility Lemma

It remains to prove the reducibility lemma which we restate here for convenience.

Lemma 6.5 (Reducibility Lemma, restated). *Let $G = (V, E)$ be a $(\ell, 1/3\Delta)$ -sparse graph with a linear order $<_v$ on V and consider an admissible order that respects $<_v$. If the vertex sets $W \subseteq U$ satisfy that W is closed, that the size of U is $|U| \leq \ell$, and that every vertex in $N_{U \setminus W}(W)$ has degree at most Δ in $G[U \setminus W]$, then for every monomial m such that $V(m) \subseteq W$, it holds that m is reducible modulo $\langle U \rangle$ if and only if m is reducible modulo $\langle W \rangle$.*

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

Let us first outline the construction of ρ . Using the definition of closure, we show in Lemma 5.3 that there exists a proper 3-colouring χ of the subgraph $G[U \setminus (W \cup N_U(W))]$ that uses a *single* colour for each set $N_{U \setminus W}(u)$, where $u \in N_{U \setminus W}(W)$. Variables far from W , which here means

variables associated with a vertex in $U \setminus (W \cup N_U(W))$, are mapped according to the 3-colouring χ . It remains to define ρ on variables associated with each vertex $u \in N_{U \setminus W}(W)$. Since u has precisely one adjacent vertex v in W , and since furthermore the set $N_{U \setminus W}(u)$ is coloured with a single colour, no matter how the vertex v is coloured there is always a colour c_u available to properly colour u . We may think of c_u as a function that, given χ and the colour of v , assigns a colour to u that is consistent with the colouring of $N_U(u)$. Variables associated with the vertex u are mapped according to c_u by ρ .

Proof of Lemma 6.5. Since W is a subset of U , it follows that if m is reducible modulo $\langle W \rangle$, then m is also reducible modulo $\langle U \rangle$. For the reverse direction, we define a mapping ρ on variables as outlined above.

To this end, let χ be a proper 3-colouring of the subgraph $G[U \setminus (W \cup N_U(W))]$ that uses a single colour for each set $N_{U \setminus W}(u)$, where $u \in N_{U \setminus W}(W)$. Such a colouring exists by Lemma 5.3. Variables associated with a vertex $u \in U \setminus (W \cup N_U(W))$ are mapped according to χ : if $\chi(u) = i$, then $\rho(x_{u,i}) = 1$ and $\rho(x_{u,i'}) = 0$ for all $i' \neq i$.

Next, for each vertex $u \in N_{U \setminus W}(W)$, we define ρ on the variables associated with u . Since χ assigns, for each $u \in N_{U \setminus W}(W)$, a single colour to each set $N_{U \setminus W}(u)$ there are at least two distinct colours $c_1, c_2 \in [k]$ that are not assigned to any vertex in $N_{U \setminus W}(u)$. Since there are no 2-hops in U with respect to W the vertex u has a single neighbour $v \in W$. Furthermore, as there are no 3-hops in U with respect to W , it holds that $N_{U \setminus W}(u) \cap N_{U \setminus W}(W) = \emptyset$, which implies that the vertex v is the only neighbour of u that is not coloured by χ . Hence no matter how v is coloured by χ , either c_1 or c_2 can be used to properly colour u . Let us make this choice explicit by defining ρ on u by

$$\rho(x_{u,c}) = \begin{cases} x_{v,c_2} & \text{if } c = c_1, \\ \sum_{i \in [k], i \neq c_2} x_{v,i} & \text{if } c = c_2, \text{ and} \\ 0 & \text{otherwise, that is, if } c \notin \{c_1, c_2\}. \end{cases} \quad (6.6)$$

This completes the definition of ρ . Note that the mapping ρ extends any proper k -colouring of W to a proper k -colouring of U .

Let q be a polynomial in $\langle U \rangle$ with leading monomial m . We claim that

1. $q \upharpoonright_\rho \in \langle W \rangle$,
2. that all monomials m' satisfy $m' \upharpoonright_\rho \leq m'$, and
3. that $m = m \upharpoonright_\rho$.

If we can show this, then we are done, since m is then the leading monomial of the polynomial $q \upharpoonright_\rho \in \langle W \rangle$ and we may thus conclude that if m is reducible modulo $\langle U \rangle$, then m is also reducible modulo $\langle W \rangle$.

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

It remains to prove that $q \upharpoonright_\rho \in \langle W \rangle$. Since $q \in \langle U \rangle$ we may write $q = \sum_i q_i p_i$ for polynomials $q_i \in \mathbb{F}[X]$ and axioms $p_i \in \text{Col}(G[U], k)$. Note that the mapping ρ extends any proper k -colouring of W to a proper k -colouring of U , and thus it follows by Lemma 2.4 that every axiom $p_i \in \text{Col}(G[U], k)$ satisfies $p_i \upharpoonright_\rho \in \langle W \rangle$. We can therefore conclude that the polynomial $q \upharpoonright_\rho = \sum_i q_i \upharpoonright_\rho \cdot p_i \upharpoonright_\rho$ is in $\langle W \rangle$ as claimed. \square

With the proof of Lemma 6.5 completed we have shown the last missing piece of the proof of Theorem 6.1. This thus establishes our polynomial calculus degree lower bounds for the colouring formula over sparse graphs.

7 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 by the well-known size-degree relation for polynomial calculus [IPS99].

Our overall proof technique is the same as that of earlier papers such as [AR03, GL10a, GL10b, MN15], but our proofs have a different flavour. A central technical concept in [AR03, MN15] is (variations of) the so-called *constraint-variable incidence graph*: this graph consists of a vertex per constraint and variable, and has an edge between a constraint C and a variable x if and only if C depends on x . This graph is commonly used to argue that by expansion small sets of constraints are satisfiable, even after the removal of a closed set of vertices. By contrast, we never need to make any (explicit) use of this graph. This raises the question of whether it is possible to rephrase our proofs in language closer to that of [AR03, MN15], or are the two approaches inherently different?

The lower bound techniques in this paper, as well as those in [MN15], work over any field. For NP-hard problems such as k -colourability, we expect a polynomial calculus lower bound to hold regardless of which field is used for the derivations. However, other formulas such as the Tseitin contradictions are easy for polynomial calculus over a field of characteristic 2 and hard in other characteristics. The techniques in, for instance, [BI99, BGIP01, AR03] capture this fact, while those in [MN15] and this paper cannot. Another interesting question is therefore whether the techniques in these papers can be unified into a general approach that works both for field-dependent and field-independent lower bounds.

Our degree lower bounds for 3-colourability are of the form $n/f(d)$, where d is either degree or average degree of the graph depending on the random graph model. In our work, f is at least exponential in d , but in previous results [BCMM05, LN17], 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 the parameters in our result can be improved. We remark that it is far from clear what the correct dependence on d should be. For the sums-of-squares proof system, which simulates polynomial calculus over the reals [Ber18], there exist strong upper bounds for k -colourability on random graphs and random regular graphs in some parameter regimes: the paper [BKM19] 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 -colourability whenever k is fixed. However, similar upper bounds are not known to hold for polynomial calculus, and it should be pointed out that the latter proof system is incomparable to sum-of-squares when considered over fields of finite characteristic.

More broadly 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 [RT22] 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 [AR22] 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 [BIS07]. Finally, an intriguing technical challenge is to prove degree lower bounds for variants of the dense linear ordering principle [AD08] for graphs of bounded degree.

Acknowledgements

The authors would like to thank Albert Atserias, Gaia Carenini and Amir Yehudayoff 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 of the participants of the Dagstuhl Seminar 23111 “Computational Complexity of Discrete Problems” and Oberwolfach workshop 2413 “Proof Complexity and Beyond”. Finally, we are grateful to Maryia Kapytka for feedback on preliminary versions of this manuscript and also to the anonymous FOCS reviewers—all these comments 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. 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. Kilian Risse was supported by the Swiss National Science Foundation project 200021-184656 “Randomness in Problem Instances and Randomized Algorithms”.

A On Boolean Implication and Ideal Membership

In this appendix, we provide a proof of the folklore result that being implied by a set of polynomials and being in the ideal generated by these polynomials is the same in the Boolean setting.

Lemma 2.4 (restated). *Let g be a polynomial and Q be a set of polynomials in $\mathbb{F}[x_1, \dots, x_n]$, and suppose that Q contains all the Boolean axioms. Then it holds that g vanishes on all common roots of Q if and only if $g \in \langle Q \rangle$.*

Proof. If $g \in \langle Q \rangle$, then we can write $g = \sum_i f_i q_i$ for $f_i \in \mathbb{F}[x_1, \dots, x_n]$ and $q_i \in Q$. Observe that $\sum_i f_i q_i|_{\xi} = 0$ for any common root $\xi \in \{0, 1\}^n$ of Q . Hence g vanishes on all common roots of Q .

For the other direction, let $\xi = (\xi_1, \xi_2, \dots, \xi_n)$ be an element of $\{0, 1\}^n$ and write $\mathbf{1}_{\xi}$ for the multilinear polynomial that evaluates to 1 on ξ and to 0 on all other elements of $\{0, 1\}^n$, that is,

$$\mathbf{1}_{\xi}(x) = \prod_{i:\xi_i=1} x_i \prod_{j:\xi_j=0} (1 - x_j). \quad (\text{A.1})$$

Clearly, every function $f: \{0, 1\}^n \rightarrow \mathbb{F}$ can be expressed as a multilinear polynomial through the identity

$$f = \sum_{\xi \in \{0,1\}^n} f(\xi) \cdot \mathbf{1}_{\xi}, \quad (\text{A.2})$$

and as the polynomials $\{\mathbf{1}_\xi \mid \xi \in \{0, 1\}^n\}$ form a basis of the vector space of multilinear polynomials over $\mathbb{F}[x_1, \dots, x_n]$, this representation is unique. Let $S \subseteq \{0, 1\}^n$ be the set of common roots of the polynomials in Q . Since g vanishes on S , we may write

$$g = \sum_{\xi \in \{0, 1\}^n \setminus S} g(\xi) \cdot \mathbf{1}_\xi. \quad (\text{A.3})$$

We now show that if $\xi \in \{0, 1\}^n \setminus S$, then it holds that $\mathbf{1}_\xi \in \langle Q \rangle$. This suffices to prove the desired result, as g is then a linear combination of polynomials in $\langle Q \rangle$ and hence is also in $\langle Q \rangle$.

Let ξ be an element in $\{0, 1\}^n \setminus S$. Because S is the set of common roots of the polynomials in Q , there exists a polynomial $q \in Q$ such that $q(\xi) \neq 0$. The polynomial $\mathbf{1}_\xi \cdot (q(\xi))^{-1}q$ coincides with $\mathbf{1}_\xi$ on all of $\{0, 1\}^n$, so $\mathbf{1}_\xi \cdot (q(\xi))^{-1}q = \mathbf{1}_\xi$ modulo the Boolean axioms. Since $\langle Q \rangle$ contains both $\mathbf{1}_\xi \cdot (q(\xi))^{-1}q$ and the Boolean axioms, it follows that $\mathbf{1}_\xi \in \langle Q \rangle$. \square

B Random Graphs Are Sparse

In this appendix we prove Lemma 2.7, which is a quantitative version of Lemma 4.15 in [Raz17]. We make no claim of originality, but present this result here to make the paper self-contained.

We start by proving the sparsity lemma for graphs sampled from the Erdős-Rényi random graph distribution $\mathbb{G}(n, d/n)$ in Lemma B.1 and then establish the analogous result for random regular graphs sampled according to $\mathbb{G}_{n,d}$ in Lemma B.4.

Lemma B.1 (Sparsity lemma for Erdős-Rényi random graphs). *Let $n, d \in \mathbb{N}^+$ and $\varepsilon, \delta \in \mathbb{R}^+$ be such that $\varepsilon\delta = \omega(1/\log n)$. If G is a graph sampled from $\mathbb{G}(n, d/n)$, then asymptotically almost surely it is $((4d)^{-(1+\delta)(1+\varepsilon)/\varepsilon} n, \varepsilon)$ -sparse.*

Proof. Let $\alpha = (4d)^{-(1+\delta)(1+\varepsilon)/\varepsilon}$ and denote by \mathcal{A} the event “ G is $(\alpha n, \varepsilon)$ -sparse”. For a set $U \subseteq V$ of size s , the random variable $|E(G[U])|$ is a sum of $s(s-1)/2$ random indicator variables for the edges that are 1 with probability d/n and 0 otherwise. We apply a union bound over sets of size $s \leq \alpha n$ to conclude that

$$\Pr[\neg \mathcal{A}] \leq \sum_{\substack{U \subseteq V \\ |U| \leq \alpha n}} \Pr[|E(G[U])| \geq (1+\varepsilon)|U|] \quad (\text{B.1a})$$

$$\leq \sum_{s=1}^{\alpha n} \binom{n}{s} \binom{\frac{s(s-1)}{2}}{(1+\varepsilon)s} \cdot \left(\frac{d}{n}\right)^{(1+\varepsilon)s} \quad (\text{B.1b})$$

$$\leq \sum_{s=1}^{\alpha n} \left(\frac{en}{s}\right)^s \left(\frac{e(s-1)}{2(1+\varepsilon)}\right)^{(1+\varepsilon)s} \cdot \left(\frac{d}{n}\right)^{(1+\varepsilon)s} \quad (\text{B.1c})$$

$$\leq \sum_{s=1}^{\alpha n} \exp\left(-\varepsilon s \ln\left(\frac{n}{s}\right) + (1+\varepsilon)s \left(\ln\left(\frac{e^2 d}{2(1+\varepsilon)}\right)\right)\right) \quad (\text{B.1d})$$

$$\leq \sum_{s=1}^{\alpha n} \exp(-\delta \varepsilon s \ln(n/s)) \quad (\text{B.1e})$$

$$\leq o(1), \quad (\text{B.1f})$$

where for (B.1e) we use that $n/s \geq 1/\alpha = (4d)^{(1+\delta)(1+\varepsilon)/\varepsilon}$ to estimate that

$$(1+\varepsilon) \left(\ln\left(\frac{e^2 d}{2(1+\varepsilon)}\right)\right) \leq (1+\varepsilon) \ln(4d) = \frac{\varepsilon \ln(1/\alpha)}{1+\delta} \leq \frac{\varepsilon \ln(n/s)}{1+\delta}, \quad (\text{B.2})$$

and for (B.1f) we use that $\varepsilon\delta = \omega(1/\log n)$ and that $s \ln(n/s) \geq \ln n + s - 1$ (for $1 \leq s \leq n/e^2$). \square

We next prove the sparsity lemma for random regular graphs. In order to sample G from $\mathbb{G}_{n,d}$ we use the *configuration model*, which is defined as follows. Given n and d such that dn is even, we have a vertex set V of size n and for each vertex $v \in V$ there is a cell C_v with d elements. We sample a perfect matching M uniformly from the set \mathcal{M}_{dn} of all possible perfect matchings of the dn elements and consider the corresponding multi-graph $G_M = (V, E)$, possibly with parallel edges and loops, where $(u, v) \in E$ if and only if there exists $(x, y) \in M$ such that $x \in C_u$ and $y \in C_v$. Since all simple graphs (without parallel edges or loops) are sampled with the same probability, we can sample G from $\mathbb{G}_{n,d}$ by sampling according to the configuration model repeatedly until we sample a simple graph.

Theorem B.2 ([MW91, Wor99]). *For $d = o(n^{1/2})$, the probability that G_M is simple when M is sampled uniformly from \mathcal{M}_{dn} is equal to*

$$\exp\left(-\frac{d^2 - 1}{4} - \frac{d^3}{12n} + O(d^2/n)\right).$$

Let $S_{\ell,q}$ denote the sum of ℓ random variables that are 1 with probability q and 0 otherwise. We argue that we can bound the probability that a set of vertices $U \subseteq V$ witnesses that the graph is not sparse by bounding the probability that $S_{\ell,q}$ is too large.

Claim B.3. For any $s \leq n/2$ and $B \in \mathbb{R}^+$, if $U \subseteq V$ is of size s and $q = s/(n-s)$, it holds that

$$\Pr_{M \sim \mathcal{M}_{dn}} [|E(G_M[U])| \geq B] \leq \Pr[S_{ds,q} \geq B].$$

Proof. To see why this is true, consider the random process in the configuration model that matches one by one the elements in the cells C_v for all $v \in U$. At each step, there are at most ds elements in cells C_v where $v \in U$ that are not yet matched, and at least $d(n-s)$ elements in cells C_v where $v \notin U$ that are not yet matched. This implies that at each step the probability that we obtain an edge between cells in U is at most $ds/d(n-s) = q$. Since at least one element in a cell of U is matched at every step, there are at most ds steps in total. Hence the claim follows. \square

Lemma B.4 (Sparsity lemma for random regular graphs). *Let $n, d \in \mathbb{N}^+$ and $\varepsilon, \delta \in \mathbb{R}^+$ be such that $\varepsilon\delta = \omega(1/\log n)$ and $d^2 \leq \varepsilon\delta \log n$. If G is a graph sampled from $\mathbb{G}_{n,d}$, then asymptotically almost surely it is $((8d)^{-(1+\delta)(1+\varepsilon)/\varepsilon} n, \varepsilon)$ -sparse.*

Proof. Fix $\varepsilon > 0$, let $\alpha = (8d)^{-2(1+1/\varepsilon)}$, and denote by \mathcal{A} the event “ G is $(\alpha n, \varepsilon)$ -sparse”. We want to prove that $G \sim \mathbb{G}_{n,d}$ is $(\alpha n, \varepsilon)$ -sparse with probability that goes to 1 as n goes to infinity. To this end, we prove that if we sample G from the configuration model, the probability that it is not $(\alpha n, \varepsilon)$ -sparse is much smaller than the probability that G is a random regular graph. More formally, our goal is to prove that

$$\Pr[\neg \mathcal{A}] \cdot \exp\left(\frac{d^2 - 1}{4}\right) \leq o(1), \tag{B.3}$$

where we recall that the probability here and in what follows is taken over sampling G in the

configuration model. By union bound and using Claim B.3 we have that

$$\Pr[\neg \mathcal{A}] \cdot \exp\left(\frac{d^2 - 1}{4}\right) \leq \sum_{\substack{U \subseteq V \\ |U| \leq \alpha n}} \Pr[|E(G[U])| \geq (1 + \varepsilon)|U|] \cdot \exp\left(\frac{d^2}{4}\right) \quad (\text{B.4a})$$

$$\leq \sum_{s=1}^{\alpha n} \binom{n}{s} \Pr[S_{ds,q} \geq (1 + \varepsilon)s] \cdot \exp\left(\frac{d^2}{4}\right). \quad (\text{B.4b})$$

Our goal is to show that $\Pr[S_{ds,q} \geq (1 + \varepsilon)s]$ is so small that it compensates for the other factors. More concretely, we wish to show that for $s \leq \alpha n$ it holds that

$$\binom{n}{s} \Pr[S_{ds,q} \geq (1 + \varepsilon)s] \leq \exp(-\varepsilon \delta s \ln(n/s)), \quad (\text{B.5})$$

from which it follows, since $d^2 \leq \varepsilon \delta \log n \leq 2\varepsilon \delta \ln n$, that

$$\Pr[\neg \mathcal{A}] \cdot \exp\left(\frac{d^2 - 1}{4}\right) \leq \sum_{s=1}^{\alpha n} \exp\left(-\varepsilon \delta s \ln(n/s) + \frac{d^2}{4}\right) \quad (\text{B.6a})$$

$$\leq \sum_{s=1}^{\alpha n} \exp\left(-\frac{\varepsilon \delta s \ln(n/s)}{2}\right) \quad (\text{B.6b})$$

$$= o(1), \quad (\text{B.6c})$$

where we use the fact that $\varepsilon \delta = \omega(1/\log n)$ and that $s \ln(n/s) \geq \ln n + s - 1$ (for $1 \leq s \leq n/e^2$).

It remains to show that (B.5) holds. This follows from the sequence of calculations

$$\binom{n}{s} \Pr[S_{ds,q} \geq (1 + \varepsilon)s] \leq \binom{n}{s} \binom{ds}{(1 + \varepsilon)s} \cdot \left(\frac{s}{n - s}\right)^{(1 + \varepsilon)s} \quad (\text{B.7a})$$

$$\leq \left(\frac{en}{s}\right)^s \left(\frac{eds}{(1 + \varepsilon)s}\right)^{(1 + \varepsilon)s} \cdot \left(\frac{s}{n - s}\right)^{(1 + \varepsilon)s} \quad (\text{B.7b})$$

$$\leq \exp\left(-\varepsilon s \ln\left(\frac{n - s}{s}\right) + s + (1 + \varepsilon)s \ln\left(\frac{ed}{(1 + \varepsilon)}\right)\right) \quad (\text{B.7c})$$

$$\leq \exp\left(-\varepsilon s \ln\left(\frac{n}{s}\right) + (1 + \varepsilon)s \ln\left(\frac{e^2 d}{(1 + \varepsilon)}\right)\right) \quad (\text{B.7d})$$

$$\leq \exp(-\varepsilon \delta s \ln(n/s)), \quad (\text{B.7e})$$

where for (B.7d) we use that $\ln(n/s - 1) \geq \ln(n/s) - 1$ (for $s \leq n/2$), and for (B.7e) we use that $n/s \geq 1/\alpha = (8d)^{-(1 + \delta)(1 + \varepsilon)/\varepsilon}$ to bound

$$(1 + \varepsilon) \left(\ln\left(\frac{e^2 d}{(1 + \varepsilon)}\right) \right) \leq (1 + \varepsilon) \ln(8d) = \frac{\varepsilon \ln(1/\alpha)}{1 + \delta} \leq \frac{\varepsilon \ln(n/s)}{1 + \delta}. \quad (\text{B.8})$$

This concludes the proof of Lemma 2.7. \square

C Erdős-Rényi Graphs Almost Have Small Maximum Degree

In this section we provide a proof of Lemma 6.3 stating that Erdős-Rényi random graphs have small degree except for a small set of vertices.

Lemma 6.3 (restated). *Let $G = (V, E)$ be a graph sampled from $\mathbb{G}(n, d/n)$ where $d = O(\log n)$. If $\Delta \geq d$ is such that $(\Delta/ed)^\Delta = o(n)$, then asymptotically almost surely there exists a set T_Δ of size at most $(ed/\Delta)^\Delta \cdot 2en$ such that the maximum degree in $G[V \setminus T_\Delta]$ is at most $\Delta - 1$.*

Proof. Let \mathcal{A} denote the event that there exists a set $T_\Delta \subseteq V$ of size $\ell = (ed/\Delta)^\Delta \cdot 2en$ such that the maximum degree in $G[V \setminus T_\Delta]$ is at most $\Delta - 1$. We prove that $\Pr[\neg \mathcal{A}] = o(1)$.

Note that if \mathcal{A} does not hold, then, in particular, if we go over the vertices of G in any given fixed order, and remove from G any vertex of degree at least Δ that we encounter, we will end up removing at least ℓ vertices. Observe further that after the removal of i vertices, the probability that a vertex has degree at least Δ is at most $\binom{n-i}{\Delta} (d/n)^\Delta$ and is independent of the fact that the removed vertices had degree at least Δ . Therefore, by taking a union bound over all sets of size ℓ , we can bound the probability of the event \mathcal{A} not holding by

$$\Pr[\neg \mathcal{A}] \leq \binom{n}{\ell} \prod_{i=1}^{\ell} \binom{n-i}{\Delta} \left(\frac{d}{n}\right)^\Delta \quad (\text{C.1a})$$

$$\leq \binom{n}{\ell} \left(\binom{n}{\Delta} \left(\frac{d}{n}\right)^\Delta \right)^\ell \quad (\text{C.1b})$$

$$\leq \left(\frac{en}{\ell}\right)^\ell \left(\frac{ed}{\Delta}\right)^{\Delta \ell} \quad (\text{C.1c})$$

$$= o(1), \quad (\text{C.1d})$$

where for (C.1d) we use that $\ell = 2en (ed/\Delta)^\Delta$ to derive that $(en/\ell)^\ell (ed/\Delta)^{\Delta \ell} = 2^{-\ell}$, and then use that $\ell = \omega(1)$, which holds since $(\Delta/ed)^\Delta = o(n)$, to conclude $2^{-\ell} = o(1)$. The lemma follows. \square

References

- [ABRW02] Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Space complexity in propositional calculus. *SIAM Journal on Computing*, 31(4):1184–1211, April 2002. Preliminary version in *STOC '00*.
- [AD08] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, May 2008. Preliminary version in *CCC '03*.
- [AH19] Albert Atserias and Tuomas Hakoniemi. Size-degree trade-offs for Sums-of-Squares and Positivstellensatz proofs. In *Proceedings of the 34th Annual Computational Complexity Conference (CCC '19)*, volume 137 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 24:1–24:20, July 2019.
- [AN05] Dimitris Achlioptas and Assaf Naor. The two possible values of the chromatic number of a random graph. *Annals of Mathematics*, 162(3):1335–1351, November 2005.
- [AO19] Albert Atserias and Joanna Ochremiak. Proof complexity meets algebra. *ACM Transactions on Computational Logic*, 20:1:1–1:46, February 2019. Preliminary version in *ICALP '17*.

References

- [AR01] Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS '01)*, pages 190–199, October 2001.
- [AR03] Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. *Proceedings of the Steklov Institute of Mathematics*, 242:18–35, 2003. Available at <http://people.cs.uchicago.edu/~razborov/files/misha.pdf>. Preliminary version in FOCS '01.
- [AR22] Per Austrin and Kilian Risse. Perfect matching in random graphs is as hard as Tseitin. *TheoretCS*, 1:Art. 2, 47, 2022. Preliminary version in SODA '22.
- [AT92] Noga Alon and Michael Tarsi. Colorings and orientations of graphs. *Combinatorica*, 12(2):125–134, June 1992.
- [Bay82] David Allen Bayer. *The Division Algorithm and the Hilbert Scheme*. PhD thesis, Harvard University, June 1982. Available at <https://www.math.columbia.edu/~bayer/papers/Bayer-thesis.pdf>.
- [BBKO21] Libor Barto, Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *Journal of the ACM*, 68(4):28:1–28:66, August 2021. Preliminary version in STOC '19.
- [BCMM05] Paul Beame, Joseph C. Culberson, David G. Mitchell, and Cristopher Moore. The resolution complexity of random graph k -colorability. *Discrete Applied Mathematics*, 153(1-3):25–47, December 2005.
- [BE05] Richard Beigel and David Eppstein. 3-coloring in time $O(1.3289^n)$. *Journal of Algorithms*, 54(2):168–204, February 2005.
- [Ber18] Christoph 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)*, volume 96 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:14, February 2018.
- [BGIP01] Samuel R. Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. Linear gaps between degrees for the polynomial calculus modulo distinct primes. *Journal of Computer and System Sciences*, 62(2):267–289, March 2001. Preliminary version in CCC '99.
- [BI99] Eli Ben-Sasson and Russell Impagliazzo. Random CNF's are hard for the polynomial calculus. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS '99)*, pages 415–421, October 1999. Journal version in [BI10].
- [BI10] Eli Ben-Sasson and Russell Impagliazzo. Random CNF's are hard for the polynomial calculus. *Computational Complexity*, 19(4):501–519, 2010. Preliminary version in FOCS '99.
- [BIK⁺94] Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel 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)*, pages 794–806, November 1994.

- [BIS07] Paul Beame, Russell Impagliazzo, and Ashish Sabharwal. The resolution complexity of independent sets and vertex covers in random graphs. *Computational Complexity*, 16(3):245–297, October 2007. Preliminary version in CCC '01.
- [BKM19] Jess Banks, Robert Kleinberg, and Cristopher Moore. The Lovász theta function for random regular graphs and community detection in the hard regime. *SIAM Journal on Computing*, 48(3):1098–1119, 2019.
- [Bla37] Archie Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, University of Chicago, 1937.
- [BN21] Samuel R. Buss and Jakob Nordström. Proof complexity and SAT solving. In Armin Biere, Marijn J. H. Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, chapter 7, pages 233–350. IOS Press, 2nd edition, February 2021.
- [Bol78] Béla Bollobás. Chromatic number, girth and maximal degree. *Discrete Mathematics*, 24(3):311–314, 1978.
- [Bus98] Samuel R. Buss. Lower bounds on Nullstellensatz proofs via designs. In *Proof Complexity and Feasible Arithmetics*, volume 39 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 59–71. American Mathematical Society, 1998. Available at <http://www.math.ucsd.edu/~sbuss/ResearchWeb/designs/>.
- [BW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, March 2001. Preliminary version in STOC '99.
- [CEI96] Matthew Clegg, Jeffery Edmonds, and Russell 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)*, pages 174–183, May 1996.
- [CFRR02] Colin Cooper, Alan Frieze, Bruce Reed, and Oliver Riordan. Random regular graphs of non-constant degree: independence and chromatic number. *Combinatorics, Probability and Computing*, 11(4):323–341, 2002.
- [Coj05] Amin Coja-Oghlan. The Lovász number of random graphs. *Combinatorics, Probability and Computing*, 14(4):439–465, 2005.
- [CR79] Stephen A. Cook and Robert A. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, March 1979. Preliminary version in STOC '74.
- [DL95] Jesús A. De Loera. Gröbner bases and graph colorings. *Beiträge zur Algebra und Geometrie*, 36(1):89–96, January 1995.
- [DLL62] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem proving. *Communications of the ACM*, 5(7):394–397, July 1962.
- [DLMM08] Jesús A. De Loera, Jon Lee, Peter N. Malkin, and Susan 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)*, pages 197–206, July 2008.

References

- [DLMM11] Jesús A. De Loera, Jon Lee, Peter N. Malkin, and Susan Margulies. Computing infeasibility certificates for combinatorial problems through Hilbert’s Nullstellensatz. *Journal of Symbolic Computation*, 46(11):1260–1283, November 2011.
- [DLMO09] Jesús A. De Loera, Jon Lee, Susan Margulies, and Shmuel Onn. Expressing combinatorial problems by systems of polynomial equations and Hilbert’s Nullstellensatz. *Combinatorics, Probability and Computing*, 18(4):551–582, July 2009.
- [DMP⁺15] Jesús A. De Loera, Susan Margulies, Michael Pernpeintner, Eric Riedl, David Rolnick, Gwen Spencer, Despina Stasi, and Jon 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)*, pages 133–140, July 2015.
- [DP60] Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, 1960.
- [Fil19] Yuval Filmus. Another look at degree lower bounds for polynomial calculus. *Theoretical Computer Science*, 796:286–293, December 2019.
- [GL10a] Nicola Galesi and Massimo Lauria. On the automatizability of polynomial calculus. *Theory of Computing Systems*, 47(2):491–506, August 2010.
- [GL10b] Nicola Galesi and Massimo Lauria. Optimality of size-degree trade-offs for polynomial calculus. *ACM Transactions on Computational Logic*, 12(1):4:1–4:22, November 2010.
- [Hal93] Magnús M. Halldórsson. A still better performance guarantee for approximate graph coloring. *Information Processing Letters*, 45(1):19–23, January 1993.
- [HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, October 2006.
- [Hus15] Thore Husfeldt. Graph colouring algorithms. In Lowell W. Beineke and Robin J. Wilson, editors, *Topics in Chromatic Graph Theory*, Encyclopedia of Mathematics and its Applications, chapter 13, pages 277–303. Cambridge University Press, May 2015.
- [HW08] Christopher J. Hillar and Troels Windfeldt. Algebraic characterization of uniquely vertex colorable graphs. *Journal of Combinatorial Theory, Series B*, 98(2):400–414, March 2008.
- [IPS99] Russell Impagliazzo, Pavel Pudlák, and Jiří Sgall. Lower bounds for the polynomial calculus and the Gröbner basis algorithm. *Computational Complexity*, 8(2):127–144, 1999.
- [Kar72] Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, The IBM Research Symposia Series, pages 85–103. Springer, 1972.
- [KM21] Pravesh K. Kothari and Peter Manohar. A stress-free sum-of-squares lower bound for coloring. In *Proceedings of the 36th Annual IEEE Conference on Computational Complexity (CCC ’21)*, volume 200 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 23:1–23:21, July 2021.
- [KO22] Andrei Krokhin and Jakub Opršal. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News*, 9(3):30–59, 2022.

- [KPGW10] Graeme Kemkes, Xavier Pérez-Giménez, and Nicholas Wormald. On the chromatic number of random d -regular graphs. *Advances in Mathematics*, 223(1):300–328, January 2010.
- [KT17] Ken-Ichi Kawarabayashi and Mikkel Thorup. Coloring 3-colorable graphs with less than $n^{1/5}$ colors. *Journal of the ACM*, 64(1), March 2017. Preliminary version in *STACS '14*.
- [Las01] Jean 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)*, volume 2081 of *Lecture Notes in Computer Science*, pages 293–303. Springer, June 2001.
- [Lau18] Massimo 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*, volume 10936 of *Lecture Notes in Computer Science*, pages 254–263. Springer International Publishing, July 2018.
- [LN17] Massimo Lauria and Jakob 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)*, volume 79 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:20, July 2017.
- [Lov94] László Lovász. Stable sets and polynomials. *Discrete Mathematics*, 124(1–3):137–153, January 1994.
- [Łuc91] Tomasz Łuczak. The chromatic number of random graphs. *Combinatorica*, 11(1):45–54, 1991.
- [Mat74] Yuri V. Matiyasevich. A criterion for vertex colorability of a graph stated in terms of edge orientations. *Diskretnyi Analiz*, 26:65–71, 1974. English translation of the Russian original. Available at http://logic.pdmi.ras.ru/~yumat/papers/22_paper/.
- [Mat04] Yuri V. Matiyasevich. Some algebraic methods for calculating the number of colorings of a graph. *Journal of Mathematical Sciences*, 121(3):2401–2408, May 2004.
- [McD84] Colin McDiarmid. Colouring random graphs. *Annals of Operations Research*, 1(3):183–200, October 1984.
- [MN15] Mladen Mikša and Jakob Nordström. A generalized method for proving polynomial calculus degree lower bounds. In *Proceedings of the 30th Annual Computational Complexity Conference (CCC '15)*, volume 33 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 467–487, June 2015.
- [Mnu01] Michal Mnuć. Representing graph properties by polynomial ideals. In *Proceedings of the 4th International Workshop on Computer Algebra in Scientific Computing (CASC '01)*, pages 431–444, September 2001.
- [MW91] Brendan D. McKay and Nicholas C. Wormald. Asymptotic enumeration by degree sequence of graphs with degrees $o(n^{1/2})$. *Combinatorica*, 11(4):369–382, December 1991.
- [Par00] Pablo A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, May 2000. Available at <http://resolver.caltech.edu/CaltechETD:etd-05062004-055516>.

References

- [Pud00] Pavel Pudlák. Proofs as games. *American Mathematical Monthly*, pages 541–550, 2000.
- [Raz98] Alexander A. Razborov. Lower bounds for the polynomial calculus. *Computational Complexity*, 7(4):291–324, December 1998.
- [Raz17] Alexander A. Razborov. On the width of semialgebraic proofs and algorithms. *Mathematics of Operations Research*, 42(4):1106–1134, May 2017.
- [Rec75] Robert A. Reckhow. *On the Lengths of Proofs in the Propositional Calculus*. PhD thesis, University of Toronto, 1975. Available at https://www.cs.toronto.edu/~sacook/homepage/reckhow_thesis.pdf.
- [Rob65] John Alan Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, January 1965.
- [RT22] Julián Ariel Romero Barbosa and Levent Tunçel. Graphs with large girth and chromatic number are hard for Nullstellensatz. Technical Report 2212.05365, arXiv.org, December 2022.
- [SA90] Hanif D. Sherali and Warren P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3:411–430, 1990.
- [Wor99] Nicholas C. Wormald. Models of random regular graphs. In J. D. Lamb and D. A. Preece, editors, *Surveys in Combinatorics, 1999*, London Mathematical Society Lecture Note Series, pages 239—298. Cambridge University Press, 1999.
- [Zuc07] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(6):103–128, August 2007. Preliminary version in *STOC '06*.