Lower Bounds for CSP Hierarchies Through Ideal Reduction*

Jonas Conneryd† Yassine Ghananne‡ Shuo Pang§

November 24, 2025

Abstract

We present a generic way to obtain level lower bounds for (promise) CSP hierarchies from degree lower bounds for algebraic proof systems. More specifically, we show that pseudo-reduction operators in the sense of Alekhnovich and Razborov [Proc. Steklov Inst. Math. 2003] can be used to fool the cohomological k-consistency algorithm. As applications, we prove optimal level lower bounds for c vs. ℓ-coloring for all $\ell \geq c \geq 3$, and give a simplified proof of the lower bounds for lax and null-constraining CSPs of Chan and Ng [STOC 2025].

1 Introduction

The class of *constraint satisfaction problems (CSPs)* is a cornerstone of research in logic, algorithms, and complexity theory, encompassing a wide range of natural computational tasks such as Boolean satisfiability, graph k-colorability, and solving systems of linear equations. Despite their generality, CSPs possess enough structure to enable a detailed understanding of which variants are tractable, which are not, and what fundamentally distinguishes the two.

Every constraint satisfaction problem can be formulated as the task of determining, given two relational structures $\mathbf{A}$ and $\mathbf{T}$, whether there is a homomorphism $\mathbf{A} \to \mathbf{T}$ [Jea98, FV98]. A widely studied variant is when the structure $\mathbf{T}$ is kept fixed; $\mathbf{T}$ is then called the *template* of the CSP. A significant amount of research on CSPs has been motivated by the *dichotomy conjecture* of Feder and Vardi [FV98]. Broadly generalizing classical results of Schaefer [Sch78] and Hell and Nešetřil [HN90], the conjecture states that every fixed-template CSP is either in P or NP-complete.

The dichotomy conjecture was resolved independently by Bulatov [Bul17] and Zhuk [Zhu20], as the culmination of two decades of efforts into the so-called *algebraic approach* to CSPs; for an overview of this area, see [BKW17]. Informally, the results of Bulatov and Zhuk show that a CSP is tractable if and only if its template admits a certain type of non-trivial symmetry, or *polymorphism*. Any CSP lacking such symmetries was previously known to be NP-hard [MM08], and Bulatov and Zhuk exploit such a symmetry to construct a polynomial-time algorithm for the CSP.

Following the resolution of the dichotomy conjecture, a new research direction has emerged around a broad generalization of CSPs, known as *promise CSPs* or *PCSPs*. An overview of this area can be found in [KO22]. Roughly speaking, the task in a PCSP is to decide whether an instance is satisfiable in a strong sense, or not even satisfiable in a weaker sense. More formally, given a pair of relational structures $(\mathbf{S}, \mathbf{T})$ such that $\mathbf{S} \to \mathbf{T}$ and an instance structure $\mathbf{A}$, the *decision version* of the PCSP with template $(\mathbf{S}, \mathbf{T})$ is the task of determining whether

*This is the publically available version of a paper with the same title to appear in *SODA '26*.

†Lund University. jonas.conneryd@cs.lth.se

‡University of Copenhagen and Lund University. yagh@di.ku.dk

§University of Bristol. s.pang@bristol.ac.uk

$\mathbf{A} \to \mathbf{S}$ or $\mathbf{A} \not\to \mathbf{T}$, promised one of the two holds. In the *search version*, the task is to find a homomorphism $\mathbf{A} \to \mathbf{T}$ promised that $\mathbf{A} \to \mathbf{S}$.

The formal study of PCSPs was initiated by Austrin, Guruswami, and Håstad [AGH17], but special cases have been investigated for much longer. For instance, perhaps the most well-known PCSP is *approximate graph coloring* [GJ76], also called *c vs. ℓ-coloring*, where $\ell \geq c \geq 3$ and the task is to decide whether an input graph G is c-colorable or not even ℓ-colorable, promised one of the two holds. There are large gaps in our understanding of even this very natural PCSP: approximate graph coloring is widely believed to be NP-hard for all constants $\ell \geq c \geq 3$, but NP-hardness is known only for c vs. $(2c-1)$-coloring for $c \geq 3$ [BBKO21], and for c vs. $\binom{c}{\lfloor c/2 \rfloor} - 1$)-coloring when $c \geq 4$ [KOWŽ23]. By contrast, the best known polynomial-time algorithm for coloring 3-colorable graphs guarantees only an $O(n^{0.19747})$-coloring for an n-vertex graph [KTY24].

In general, the existing algorithms for solving (P)CSPs, including those of Bulatov and Zhuk, can be said to consist of two main building blocks: *local consistency*, and solving systems of *linear equations*. Local consistency algorithms check whether small subproblems of the original CSP are satisfiable, and then enforce that the local solutions are consistent with each other. These checks take various forms: for instance, the well-known *k-consistency* algorithm accepts an instance if and only if a consistent collection of partial homomorphisms with domain size at most k exists. The more stringent *Sherali–Adams hierarchy* [SA90] of linear programming relaxations asks for a probability distribution on partial solutions that is consistent with the marginals of a distribution over satisfying assignments. The *sum-of-squares* or *Lasserre* hierarchy [Par00, Las01] of semidefinite programming relaxations further demands that the probabilities assigned to the local solutions satisfy certain orthogonality conditions when viewed as vectors.

Algorithms that use linear equations to solve (P)CSPs are in turn based on the observation that any fixed-template CSP can be formulated as a system of linear equations over Boolean variables. These algorithms then relax the variable domain to one where, e.g., Gaussian elimination can be used to efficiently search for a *global* solution over the relaxed domain. A prominent example is the family of *affine relaxations*, which solve (variations of) the linear system formulation of CSPs over the integers. This can be done efficiently using, e.g., the classical Smith normal form, see [Laz96]. Examples include the *affine integer programming (AIP)* [BG19, BBKO21] and *linear Diophantine equation* [BG17] relaxations.

It is well known that neither local consistency checking nor the linear equation paradigm alone solves all tractable CSPs, and both Bulatov's and Zhuk's algorithms use elaborate combinations of the two. In recent years, quite some work has gone into designing and analyzing simpler CSP algorithms than those of Bulatov and Zhuk. Besides providing a more fine-grained understanding of CSP families, the ultimate, if distant, goal of this line of work is to find a *uniform* CSP algorithm—that is, one where the template $\mathbf{T}$ is part of the input—which runs in polynomial time whenever CSP($\mathbf{T}$) is tractable. In addition to its intrinsic importance, such an algorithm might also help establish the border of tractability for more general problems, such as PCSPs. The most well-studied algorithms in this line of work are so-called *hierarchies* of relaxations, each defined with respect to a parameter k called its *level*. Larger values of k yield increasingly tighter relaxations, while the k-th level of a hierarchy runs in time $m^{O(k)}$ on an instance of size m (with some caveats in the case of semidefinite programming relaxations, see [O'D17]).

1.1 CSP Hierarchies

Several authors have introduced and analyzed (P)CSP hierarchies that combine convex and affine relaxations. Brakensiek, Guruswami, Wrochna, and Živný [BGWŽ20], generalizing [BG19], combined the first level of the SA hierarchy with AIP and showed that the resulting algorithm, which they call BLP+AIP, solves all *Boolean* tractable CSPs. A generalization of BLP+AIP to higher levels of the SA hierarchy, called the *BA^k hierarchy*, was suggested in [BGWŽ20] and analyzed by Ciardo and Živný [CŽ25a]; see also [BG17]. The *CLAP* and *C(BLP+AIP)*

algorithms due to Ciardo and Živný [CŽ23] refine BLP+AIP by iteratively pruning variable domains and solving the pruned relaxation. Both CLAP and C(BLP+AIP) can be strengthened to hierarchies of relaxations, as noted in [CŽ23] and [CN25], respectively.

Ciardo and Živný [CŽ25b] and Chan, Ng, and Peng [CNP24] independently defined and analyzed a combination of the sum-of-squares and AIP hierarchies. We refer to this algorithm as the *SDA^k hierarchy* in this paper, following [CŽ25b]. A related but distinct algorithm, which we do not discuss further here, was recently introduced by Bhangale, Khot, and Minzer [BKM25] in the context of hardness of approximation.

Another line of work blends the k-consistency algorithm with affine relaxations. Dalmau and Opršal [DO24] defined the *$\mathbb{Z}$-affine k-consistency* algorithm, which first finds a consistent set κ of partial homomorphisms through the k-consistency algorithm and then searches for a solution supported on κ to an affine relaxation. Ó Conghaile [ÓC22] introduced *cohomological k-consistency*, which is a strengthened version of $\mathbb{Z}$-affine k-consistency with additional requirements on the solutions to the affine relaxation; we describe this algorithm in more detail in Section 1.5.

Recently, Lichter and Pago [LP25] exhibited a tractable CSP that the BA^k hierarchy, $\mathbb{Z}$-affine k-consistency algorithm, and (even parameterized versions of) CLAP all fail to solve. Hence, SDA, C(BLP+AIP), and the cohomological k-consistency algorithm are the remaining candidates for a uniform CSP algorithm among those mentioned in this section.

1.2 Lower Bounds for CSP Hierarchies

An integral part of the design and analysis of (P)CSP algorithms is to understand their limitations; this is especially relevant for PCSPs, where the border of tractability is not known. For hierarchies, such results take the form of lower bounds on the level required to solve a given (P)CSP. Besides disqualifying candidates for uniform CSP algorithms through showing they fail on tractable CSPs, level lower bounds may help identify recurring features of hard instances for different families of hierarchies, which is instructive both in algorithm design and for proving stronger hardness results.

Level lower bounds for c vs. ℓ-coloring are known for all hierarchies mentioned in Section 1.1, for varying parameter regimes. The k-consistency algorithm does not solve c vs. ℓ-coloring on n-vertex graphs for any $\ell \geq c \geq 3$ and $k \leq \Omega(n)$ [AD22, CŽ24a, CN25]; this lower bound is asymptotically optimal. In a sequence of works, Ciardo and Živný [CŽ25a, CŽ25b] showed that no constant level of the AIP, BA, or SDA hierarchy solves c vs. ℓ-coloring for any $\ell \geq c \geq 3$. The results of [CŽ25b] apply also to the more general *approximate graph homomorphism* problem. Chan, Ng, and Peng [CNP24, Appendix C] showed that the k-th level of the SDA hierarchy does not solve c vs. $\binom{c}{\lfloor c/2 \rfloor}$-coloring[1] for $c \geq 4$ and $k \leq \Omega(n)$, and Chan and Ng [CN25] proved the same result for the cohomological k-consistency and C(BLP+AIP) hierarchies. The results of [CNP24, CN25] are based on hardness reductions, and therefore seem unlikely to be improved beyond the strongest parameters for which NP-hardness is known, established in [KOWŽ23].

In addition to their lower bounds for a tractable CSP mentioned above, Lichter and Pago [LP25] show that cohomological k-consistency correctly solves the CSP on which the other algorithms fail, for $k \geq 4$. They also prove superconstant lower bounds for the cohomological k-consistency algorithm for solving a certain NP-complete CSP. Chan, Ng, and Peng [CNP24] and Chan and Ng [CN25] together establish optimal level lower bounds for all hierarchies in Section 1.1. Their lower bounds are for random instances of so-called *τ-wise neutral* CSPs for SDA and weaker hierarchies, *null-constraining CSPs* for the k-consistency algorithm, *lax and pairwise uniform* CSPs for the C(BLP+AIP) hierarchy, and for *lax and null-constraining* CSPs for cohomological k-consistency. To the best of our knowledge, no known non-trivial tractable

[1] Their result is stated for c vs. $\binom{c}{\lfloor c/2 \rfloor}$-coloring but seemingly uses [KOWŽ23], which only applies for $\ell = \binom{c}{\lfloor c/2 \rfloor} - 1$.

CSP has the properties required for the lower bounds in [CNP24, CN25], except for those for the k-consistency algorithm.

1.3 CSP Hierarchies and Proof Complexity

All hierarchies for CSP($\mathbf{T}$) mentioned in Section 1.1 always accept instances $\mathbf{A}$ such that $\mathbf{A} \to \mathbf{T}$. Therefore, to prove a lower bound on the level required to solve CSP($\mathbf{T}$), we should exhibit a *fooling instance*: a structure $\mathbf{A}$ such that $\mathbf{A} \not\to \mathbf{T}$, but which is nevertheless accepted by the k-th hierarchy level. For a PCSP with template $(\mathbf{S}, \mathbf{T})$, a fooling instance is a structure $\mathbf{A}$ such that $\mathbf{A} \not\to \mathbf{T}$, but such that the k-th level of the hierarchy for CSP($\mathbf{S}$) accepts $\mathbf{A}$.

This setting closely resembles that in *proof complexity*, where the goal is to prove lower bounds on the complexity of certificates of unsatisfiability of a propositional formula; in this case, the formula would encode that $\mathbf{A} \to \mathbf{T}$. The proof complexity of CSPs has a long and extensive history, arguably going back to the work of Tseitin [Tse68]. However, it seems fair to say that the area started in earnest with the seminal paper of Chvátal and Szemerédi [CS88], who showed that, with high probability, random k-SAT requires exponential size to refute in the *resolution* proof system. Their results were formative in showing that *random* CSP instances are often hard, and by now, lower bounds for multiple families of random CSPs are known for resolution [Mit02a, Mit02b, BCMM05, BIS07, MS07, CM13], as well as for other proof systems. As two examples relevant to this work, we mention *polynomial calculus* [BI10, AR03, MN24, CdRN$^+$25] and *sum-of-squares* [Gri01, Sch08, BCK15, BHK$^+$19, KMOW17, BKM19, KM21, JPR$^+$21, Pan21, KPX24, PX25].

Multiple results and techniques from proof complexity have informed the study of lower bounds for CSP algorithms and vice versa; some examples include [BG15, BG17, LN17, AO19, CNP24, CN25]. The main result of this work provides yet another example of a fruitful connection between the two areas, by showing that a technique used for proving polynomial calculus degree lower bounds can be used to establish level lower bounds for the cohomological k-consistency algorithm (and hence also for weaker hierarchies).

Our results are perhaps most closely related to [BG15, BG17, AO19], which relate polynomial calculus degree lower bounds to hardness results for solving CSPs in various established algorithmic paradigms. However, in contrast to the algorithms in these works, lower bounds for cohomological k-consistency do not follow in a black-box fashion from polynomial calculus degree lower bounds. For instance, the tractable CSP in the work of Lichter and Pago [LP25] is solvable by constantly many levels of cohomological k-consistency, but requires linear polynomial calculus degree over every field.[2]

1.4 Our Contributions

We show that the *Alekhnovich–Razborov method* [AR03] for polynomial calculus degree lower bounds can be strengthened to prove level lower bounds for certain CSP hierarchies. Since this method has been successfully applied to many problems in proof complexity, e.g., [AR03, GL10a, GL10b, MN24, CdRN$^+$25], we are hopeful that it will find use also in this new context.

As a first application, we prove an asymptotically optimal lower bound on the level required for the cohomological k-consistency algorithm to solve approximate graph coloring, building on the corresponding polynomial calculus degree lower bound in [CdRN$^+$25].

Theorem 1.1 (Informal)**.** *With high probability as $n \to \infty$, a random d-regular n-vertex graph has chromatic number at least $d/4 \log d$ but is accepted by the cohomological k-consistency*

[2]This CSP is a certain encoding of a disjunction of two so-called *Tseitin contradictions* on expander graphs over two different fields [LP25, Sections 3-5]. Tseitin contradictions over $\mathbb{F}_q$ on expander graphs require linear polynomial calculus degree over fields of characteristic different from q [BGIP01] and can be obtained from such CSP instances through variable restrictions. Hence, this CSP requires linear degree over every field.

algorithm for 3-colorability for all $k \leq n \cdot d^{-O(d)}$. Consequently, the cohomological k-consistency algorithm does not solve 3 vs. $d/4 \log d$-coloring on n-vertex graphs for any $k \leq n \cdot d^{-O(d)}$.

For a full statement of Theorem 1.1, see Section 6. Setting d to be a constant in Theorem 1.1, it follows that cohomological k-consistency does not solve 3 vs. c-coloring on n-vertex graphs for any constant $c \geq 3$ and any $k \leq \Omega(n)$; our lower bound on k is then optimal up to constant factors. On the other extreme, we obtain a superconstant lower bound on k for all $d = o((\log n)/\log\log n)$. Our results improve those of Chan and Ng [CN25], who show through reductions that cohomological k-consistency does not solve c vs. $\binom{c}{\lfloor c/2 \rfloor}$-coloring for any $c \geq 4$ and $k \leq \Omega(n)$; recall that for these parameters, approximate graph coloring is known to be NP-hard [KOWŽ23].

As previously noted, the only remaining candidates for uniform CSP algorithms among those in Section 1.1 are cohomological k-consistency, SDA, and C(BLP+AIP). This fact makes Theorem 1.1 noteworthy in light of the results of Banks, Kleinberg, and Moore [BKM19], who show that with high probability, semidefinite programming—and hence also algorithms based on it, such as the sum-of-squares and SDA hierarchies—correctly identifies that a random d-regular graph has chromatic number at least $\sqrt{d}/2$. Therefore, with high probability, the SDA hierarchy solves $\sqrt{d}/2$ vs. $d/4 \log d$-coloring on a random d-regular graph. Thus, Theorem 1.1 yields a separation of SDA from cohomological k-consistency, albeit for a PCSP that is widely believed to be NP-hard.

As another application, we give a different proof of the lower bounds for lax and null-constraining CSPs (for definitions, see Section 7) in [CN25, Theorem 1.3].

Theorem 1.2 (Informal, cf. [CN25], Theorem 1.3)**.** *If an n-variable t-CSP $\mathcal{P}$ is non-trivial, lax, and null-constraining, then there exist $\Delta > 0$ and $\delta > 0$ such that, with probability $\delta - o_n(1)$, a randomly sampled instance of $\mathcal{P}$ with Δn constraints is unsatisfiable but is accepted by the cohomological k-consistency algorithm for all $k \leq \Omega(n)$. Consequently, the cohomological k-consistency algorithm does not solve $\mathcal{P}$ for any $k \leq \Omega(n)$.*

The full statement of Theorem 1.2 appears in Section 7. Since the proof of Theorem 1.2 goes through the connection to pseudo-reduction operators, we immediately obtain the same lower bound of $\Omega(n)$ for the polynomial calculus degree required to refute a random instance of a lax and null-constraining CSP; see Corollary 7.5.

1.5 Technical Overview

We now review the connection between polynomial calculus degree lower bounds and level lower bounds for cohomological k-consistency, formally established in Lemmas 3.1 and 4.4. To this end, we first provide a brief overview of the necessary technical background.

Proof Complexity Background: The starting point of our connection is a technique for proving lower bounds for the *polynomial calculus* proof system [CEI96, ABRW02]. Given a field $\mathbb{F}$ and a set $\mathcal{P}$ of polynomials in $\mathbb{F}[x_1, \ldots, x_n]$, polynomial calculus proves that the polynomials in $\mathcal{P}$ have no common root through deriving new polynomials in the ideal $\langle \mathcal{P} \rangle$ generated by $\mathcal{P}$ until reaching the constant polynomial 1.

The most common complexity measures of a polynomial calculus proof are *size*, defined as the number of monomials in it, and *degree*, defined as the maximum degree among all polynomials appearing in the proof. It is known that a strong enough lower bound on degree implies lower bounds on size in a black-box fashion [IPS99]; indeed, almost all known polynomial calculus size lower bounds are proved in this way. Barring some early results [BGIP01, Gri01, BI10, BCIP02], almost all polynomial calculus degree lower bounds are in turn proved by constructing a so-called *pseudo-reduction operator* or *R-operator* [Raz98], which is an $\mathbb{F}$-linear operator that sends all

low-degree polynomials derived from $\mathcal{P}$ to 0 but sends 1 to 1; thus, there can be no low-degree polynomial calculus proof that $\mathcal{P}$ is unsatisfiable.

Pseudo-reduction operators are often based on the notion of *reduction modulo ideals* from Gröbner basis theory. Given a so-called *admissible* total order $\prec$ on the monomials in $\mathbb{F}[x_1, \ldots, x_n]$ and an ideal $I \subseteq \mathbb{F}[x_1, \ldots, x_n]$, the *reduction operator* $R_I(\cdot)$ on $\mathbb{F}[x_1, \ldots, x_n]$ acts on a polynomial p by iteratively identifying a term t in p and a polynomial $q \in I$ whose largest term with respect to $\prec$ is t, subtracting q from p, and repeating until no such q exists. It then returns the resulting polynomial. As an analogy (and indeed a special case), one can think of a system of linear equations being reduced to row-echelon form.

If the polynomials in a set $\mathcal{P}$ have a common root, then $R_{\langle \mathcal{P} \rangle}$ is a pseudo-reduction operator: all polynomials derived from $\mathcal{P}$, of any degree, belong to $\langle \mathcal{P} \rangle$ and are thus reduced to 0, and since the polynomials in $\mathcal{P}$ have a common root, we have that $1 \notin \langle \mathcal{P} \rangle$ and hence $R_{\langle \mathcal{P} \rangle}(1) = 1$. A pseudo-reduction operator R should therefore behave as $R_{\langle \mathcal{P} \rangle}$ would on low-degree monomials if the polynomials in $\mathcal{P}$ had a common root. In line with this heuristic, Alekhnovich and Razborov [AR03] pioneered an approach for constructing pseudo-reduction operators based on *local* reductions: to each low-degree monomial m, associate a satisfiable subset $S(m)$ of $\mathcal{P}$, define $R(m) = R_{\langle S(m) \rangle}(m)$, and define R on arbitrary polynomials by linear extension. We refer to operators constructed in this way as *local* pseudo-reduction operators.

Cohomological k-Consistency: Given an instance $\mathbf{A}$ of CSP($\mathbf{T}$), the cohomological k-consistency algorithm maintains a collection $\mathcal{H} = \{\mathcal{H}(X) : X \in \binom{A}{\leq k}\}$ of partial homomorphisms $\mathbf{A}[X] \to \mathbf{T}$, which is iteratively refined in two steps, outlined below. When $\mathcal{H}$ stabilizes, the algorithm rejects $\mathbf{A}$ if $\mathcal{H}(X) = \varnothing$ for some $X \in \binom{A}{\leq k}$ and accepts otherwise.

A *k-consistent* collection κ of partial homomorphisms is one where (1) for each $\varphi \in \kappa$ with domain size smaller than k, there is some $\psi_S \in \kappa$ extending φ to any larger domain S of size k, and (2) all restrictions of all $\varphi \in \kappa$ to smaller domains are also in κ. In the first step of each iteration, the algorithm finds the maximal k-consistent collection κ in $\mathcal{H}$ and assigns $\mathcal{H} = \kappa$.

In the second step, the algorithm seeks a family of solutions over $\mathbb{Z}$ to a linear system formulation $L_k(\mathbf{A}, \mathbf{T})$ of CSP($\mathbf{T}$). The system $L_k(\mathbf{A}, \mathbf{T})$ has variables x_φ for all partial homomorphisms $\varphi \colon \mathbf{A} \to \mathbf{T}$ with domain size at most k. The solutions to $L_k(\mathbf{A}, \mathbf{T})$ must all be supported on κ, and for each $\varphi \in \kappa$ there must be a solution that sets x_φ to 1 and x_ψ to 0 for all $\psi \in \kappa$ with the same domain as φ. If no such solution exists, then φ is removed from $\mathcal{H}$ in the next iteration.

A Link Between the Two: Homomorphisms $\mathbf{A} \to \mathbf{T}$ can be encoded as the common roots of a set $\mathcal{P}_{\mathbf{A} \to \mathbf{T}}$ of polynomials over $\mathbb{F}$ in Boolean variables $\{x_{a,i} \mid a \in A, i \in T\}$. Any (partial) function $\varphi : A \to T$ can then be encoded as a monomial $m_\varphi = \prod_{(a,i) \in \varphi} x_{a,i}$.

Suppose that $\mathbf{A} \not\to \mathbf{T}$, that $\mathbb{F}$ has characteristic 0, and that $\mathcal{P}_{\mathbf{A} \to \mathbf{T}}$ admits a degree-D local pseudo-reduction operator as above, where for each monomial m we associate a substructure $\mathbf{A}(m)$ of $\mathbf{A}$ and define $R(m) = R_{\langle \mathcal{P}_{\mathbf{A}(m) \to \mathbf{T}} \rangle}(m)$. We use R to show that $\mathbf{A}$ is then a fooling instance for the cohomological k-consistency algorithm for all $k \leq D$. More specifically, we show that:

1. the functions $\varphi : A \to T$ such that $R(m_\varphi) \neq 0$ form a D-consistent collection of partial homomorphisms;

2. the assignment $x_\varphi \mapsto R(m_\varphi)$ is a solution to $L_k(\mathbf{A}, \mathbf{T})$, except over the polynomial ring $\mathbb{F}[\mathbf{x}_{A,T}]$ instead of $\mathbb{Z}$;

3. for each monomial m of degree at most D, the polynomial $R(m)$ has integer coefficients.

Consequently, composing the map $x_\varphi \mapsto R(m_\varphi)$ with any evaluation on a point in $\mathbb{Z}^{|A| \cdot |T|}$, or more generally with any linear map $\mathbb{Z}[\mathbf{x}_{A,T}] \to \mathbb{Z}$ such that $h(1) = 1$, yields a solution to

$L_k(\mathbf{A}, \mathbf{T})$ over $\mathbb{Z}$. As we show in Lemma 4.4, the specific properties of local pseudo-reduction operators lets us judiciously choose these maps to produce the family of solutions to $L_k(\mathbf{A}, \mathbf{T})$ sought by the cohomological k-consistency algorithm.

Items 1 and 2 are established in Lemma 3.1 and follow rather naturally from the definition of $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ and general properties of pseudo-reduction operators. The integrality of $R(m)$ in Item 3 is proved in Lemma 5.3, and the key tool we use is the *lex game* [FRR06]. The lex game over a polynomial ring $\mathbb{F}[x_1, \ldots, x_n]$ with a lexicographic monomial order $\prec$ has two players named Lea and Stan, and is defined with respect to a set $V \subseteq \mathbb{F}^n$ and monomial m.

It is shown in [FRR06] that Lea has a winning strategy in the lex game if and only if the ideal $I(V)$ of polynomials that vanish on V contains a polynomial with largest term m. The key insight for showing Item 3 is that this winning strategy can itself be encoded as a polynomial in $I(V)$, which has largest term m and has integer coefficients when $V \subseteq \{0,1\}^n$. Therefore, each reduction step performed by $R_{I(V)}$ can be done using a polynomial with integer coefficients, and hence the output is integral. Since the ideals $\langle \mathcal{P}_{\mathbf{A}(m)\to\mathbf{T}} \rangle$ are over Boolean variables, our result applies to them, and so Item 3 follows.

Given the generality of our connection and the flexibility it offers in constructing solutions to the cohomological k-consistency algorithm, we expect it to be useful for proving lower bounds for other (P)CSPs, and potentially also for other hierarchies.

Applications: To prove Theorem 1.1, we apply our connection to the polynomial calculus degree lower bound in [CdRN+25], which uses local pseudo-reduction operators. For Theorem 1.2, we construct a local pseudo-reduction operator for $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ by using the *BW closure* from [CN25] to define the subset $S(m)$ of constraints for each low-degree monomial m. As detailed in Lemma 4.3, this amounts to showing that the so-called *satisfiability* and *reducibility* conditions are satisfied for the BW closure. Roughly speaking, the satisfiability condition asks that $S(m)$ is satisfiable, and the reducibility condition states that $R_{\langle S(m)\rangle}(m) = R_{\langle S(m)\cup T\rangle}(m)$ for any small set $T \subseteq \mathcal{P}_{\mathbf{A}\to\mathbf{T}}$. We establish these conditions in Lemmas 7.18 and 7.20 respectively, using many ideas of [CN25].

1.6 Paper Organization

The rest of this paper is organized as follows. In Section 2 we present the necessary preliminaries. Sections 3 and 4 establish our connection between lower bounds for polynomial calculus degree and CSP hierarchy level. In Section 5 we introduce the lex game and prove our key results regarding integrality of reductions. We prove our lower bounds for approximate graph coloring in Section 6, and for lax and null-constraining CSPs in Section 7. Finally, we present some concluding remarks and open problems in Section 8.

2 Preliminaries

This section reviews the necessary preliminaries from algebra, constraint satisfaction, and proof complexity. For a positive integer n, we write $[n]$ to denote the set $\{1, 2, \ldots, n\}$. Tuples are denoted by boldface lowercase letters. We say a set S is *a-small* if $|S| \leq a$. We denote the family of subsets of S by 2^S, the family of subsets of S of size k by $\binom{S}{k}$, and the family of subsets of size at most k by $\binom{S}{\leq k}$.

We use standard asymptotic notation: for functions $f, g\colon \mathbb{N} \to \mathbb{R}$, we write $f = O(g)$ if there exist nonnegative constants c and n_0 such that $f(n) \leq c \cdot g(n)$ for all $n \geq n_0$, and write $f = \Omega(g)$ if $g = O(f)$. If $f = O(g)$ and $g = O(f)$, we write $f = \Theta(g)$.

2.1 Relational Structures and Constraint Satisfaction Problems

A *signature* σ is a finite set of relation symbols $\Gamma_1, \dots, \Gamma_\ell$, each with respective arity $\mathrm{ar}(\Gamma_i) \in \mathbb{N}$. A *$\sigma$-structure* $\mathbf{A}$ consists of a domain A and, for each $\Gamma \in \sigma$, a relation $\Gamma^{\mathbf{A}} \subseteq A^{\mathrm{ar}(\Gamma)}$. A *relational structure* is a σ-structure, for some signature σ. For a σ-structure $\mathbf{A}$ and a set $X \subseteq A$, the *substructure of* $\mathbf{A}$ *on* X, denoted $\mathbf{A}[X]$, is the σ-structure with domain X where, for $\Gamma \in \sigma$, the relation $\Gamma^{\mathbf{A}[X]}$ consists of the tuples in $\Gamma^{\mathbf{A}}$ that mention only variables in X. A relational structure $\mathbf{A}$ is *finite* if $|A|$ is finite. All relational structures in this paper are finite unless otherwise stated.

For two σ-structures $\mathbf{A}$ and $\mathbf{B}$, a *homomorphism* from $\mathbf{A}$ to $\mathbf{B}$ is a map $\varphi\colon A \to B$ such that for each symbol $\Gamma \in \sigma$ and each tuple $\mathbf{a} = (a_1, \dots, a_{\mathrm{ar}(\Gamma)})$, it holds that if $\mathbf{a} \in \Gamma^{\mathbf{A}}$ then $\varphi(\mathbf{a}) = (\varphi(a_1), \dots, \varphi(a_{\mathrm{ar}(\Gamma)})) \in \Gamma^{\mathbf{B}}$. If there exists a homomorphism from $\mathbf{A}$ to $\mathbf{B}$, we write $\mathbf{A} \to \mathbf{B}$; if not, we write $\mathbf{A} \not\to \mathbf{B}$. The restriction of a homomorphism $\varphi\colon \mathbf{A} \to \mathbf{T}$ to a set $X \subseteq A$ is denoted $\varphi \restriction_X$. Given sets $Y \subseteq X \subseteq A$ and a homomorphism $\varphi\colon \mathbf{A}[Y] \to \mathbf{T}$, if there exists a homomorphism $\psi\colon \mathbf{A}[X] \to \mathbf{T}$ such that $\psi \restriction_Y = \varphi$, we say that ψ *extends* φ. We denote the set of homomorphisms from $\mathbf{A}$ to $\mathbf{B}$ by $\mathrm{Hom}(\mathbf{A}, \mathbf{B})$.

For a σ-structure $\mathbf{T}$, an *instance of the constraint satisfaction problem with template* $\mathbf{T}$ is a σ-structure $\mathbf{A}$ such that, for all $\Gamma \in \sigma$, there are no tuples in $\Gamma^{\mathbf{A}}$ with repeated elements.[3] We define $\mathrm{CSP}(\mathbf{T})$ as the computational problem where the input is an instance $\mathbf{A}$ and the task is to determine whether $\mathbf{A} \to \mathbf{T}$.

For a pair $(\mathbf{S}, \mathbf{T})$ of fixed σ-structures such that $\mathbf{S} \to \mathbf{T}$, the *promise constraint satisfaction problem (PCSP) with template* $(\mathbf{S}, \mathbf{T})$, denoted by $\mathrm{PCSP}(\mathbf{S}, \mathbf{T})$, is the computational task of distinguishing whether $\mathbf{A} \to \mathbf{S}$ or $\mathbf{A} \not\to \mathbf{T}$ for a given instance $\mathbf{A}$, promised that one of the two holds.[4] Note that $\mathrm{PCSP}(\mathbf{T}, \mathbf{T}) = \mathrm{CSP}(\mathbf{T})$.

2.2 Graphs and Hypergraphs

A *hypergraph* consists of a finite vertex set V together with a collection $\mathcal{E}$ of subsets of V. The elements of $\mathcal{E}$ are called *hyperedges*. A hyperedge set $\mathcal{E}$ is *t-uniform* if all hyperedges in $\mathcal{E}$ contain exactly t vertices, and a hypergraph $H = (V, \mathcal{E})$ is t-uniform if its hyperedge set is. For a hyperedge set $\mathcal{E}$, we denote the set $\bigcup_{e \in \mathcal{E}} e$ of vertices in $\mathcal{E}$ by $V(\mathcal{E})$. For a vertex $v \in V$, the *degree* $\deg_{\mathcal{E}}(v)$ of v in $\mathcal{E}$ is the number of hyperedges in $\mathcal{E}$ containing v. The subscript $\mathcal{E}$ is usually omitted if the hyperedge set is clear from context. A hypergraph is *d-regular* if all its vertices have degree exactly d. For a set $U \subseteq V$ and a hyperedge set $\mathcal{E}$, the *neighborhood* of U in $\mathcal{E}$ is $N_{\mathcal{E}}(U) = \{v \in V(\mathcal{E}) \setminus U : \exists u \in U, \exists e \in \mathcal{E} \mid \{u, v\} \subseteq e\}$. For a hypergraph $H = (V, \mathcal{E})$ and $U \subseteq V$, the subhypergraph of H induced by U is $H[U] = (U, \mathcal{E}[U])$, where $\mathcal{E}[U] = \{e \in \mathcal{E} : e \subseteq U\}$.

A *walk of length* ℓ in a hyperedge set $\mathcal{E}$ is a sequence of vertices $v_0, \dots, v_\ell$ and a sequence of hyperedges $e_0, \dots, e_{\ell-1}$ such that $\{v_i, v_{i+1}\} \subseteq e_i$ for all $0 \le i < \ell$. A walk is a *(Berge) path* if $v_0, \dots, v_\ell$ are distinct and $e_0, \dots, e_\ell$ are also distinct; the vertices $v_0, \dots, v_\ell$ are then the *connecting vertices* of the path, and v_0 and v_ℓ are its *endpoints*. A walk is a *(Berge) cycle* if (1) $\ell \ge 2$, (2) $v_0, \dots, v_{\ell-1}$ are distinct, (3) $v_\ell = v_0$, and (4) $e_0, \dots, e_{\ell-1}$ are distinct. The *girth* of a hyperedge set $\mathcal{E}$, denoted $\mathrm{girth}(\mathcal{E})$, is the minimum length of a cycle in $\mathcal{E}$, or ∞ if there are no cycles in $\mathcal{E}$. A hyperedge e_i in a path is *simple* if every vertex in $e_i \setminus \{v_i, v_i + 1\}$ has degree 1 in $\{e_j : 0 \le j < \ell\}$, and a connecting vertex v_i in a path is simple if it has degree 2 in $\{e_j : 0 \le j < \ell\}$. A path is simple if all its hyperedges and non-endpoint connecting vertices are

[3] This technical condition corresponds to having no repeated variables in the scope of any constraint in the instance.

[4] This is the *decision* variant of PCSP. In the *search* variant, the input is an instance $\mathbf{A}$ such that $\mathbf{A} \to \mathbf{S}$, and the task is to find a homomorphism from $\mathbf{A}$ to $\mathbf{T}$. For CSPs, it is a standard fact that the search version reduces to the decision version, but no generic search-to-decision reduction is known for PCSPs. Since an efficient algorithm for the search version can be used to also solve the decision version, a lower bound for the latter also holds for the former.

simple. A cycle is simple if all its hyperedges and connecting vertices are. A path is *pendant* in a hyperedge set $\mathcal{E}$ if none of its non-endpoints belong to any hyperedge in $\mathcal{E} \setminus \{e_0, \ldots, e_{\ell-1}\}$.

A *proper c-coloring* of a hypergraph $H = (V, \mathcal{E})$ is a map $\chi \colon V \to [c]$ such that $|\chi(e)| \geq 2$ for all $e \in \mathcal{E}$. If there exists a proper c-coloring of G, we say that G is *c-colorable*. The *chromatic number* of a hypergraph G is the smallest c such that G is c-colorable.

A *graph* $G = (V, E)$ is a 2-uniform hypergraph; in particular, all hypergraph notions in this section specialize to their usual graph-theoretic meaning.

2.3 Polynomials and Reduction Modulo Ideals

For a ring K, we denote the polynomial ring over K in n variables by $K[x_1, \ldots, x_n]$, and the set of variables appearing in a polynomial p by $\mathrm{Vars}(p)$. A *monomial* is a product of variables. For a vector $\alpha = (\alpha_1, \ldots, \alpha_n) \in \mathbb{N}^n$, we write x^α to denote the monomial $\prod_{i=1}^n x_i^{\alpha_i} \in K[x_1, \ldots, x_n]$. Given a monomial m, its *multilinearization* m_{mult} is the product of the variables appearing in m. For a set of polynomials $\mathcal{P} = \{p_1, \ldots, p_m\}$ in $K[x_1, \ldots, x_n]$, the set of polynomials of the form $\sum_{i \in [m]} q_i p_i$ for $q_i \in K[x_1, \ldots, x_n]$ is called the *ideal generated by* $\mathcal{P}$ and is denoted by $\langle \mathcal{P} \rangle$. For a set $S \subseteq K^n$, we denote the ideal of polynomials vanishing on all of S by $I(S)$. If S is empty, we define $I(S)$ to be $K[x_1, \ldots, x_n]$. For a polynomial p and a (possibly partial) mapping $\rho \colon \mathrm{Vars}(p) \to K$, we denote the mapping $K[x_1, \ldots, x_n] \to K[x_1, \ldots, x_n]$ that evaluates p on ρ by eval_ρ, and the polynomial obtained by evaluating p according to ρ by $p \restriction_\rho$.

For a field $\mathbb{F}$, a total well-order $\prec$ on the monomials in $\mathbb{F}[x_1, \ldots, x_n]$ is *admissible* if (1) for any monomial m with at least one variable, it holds that $1 \prec m$, and (2) for any monomials m_1, m_2 and m such that $m_1 \prec m_2$, it holds that $m \cdot m_1 \prec m \cdot m_2$. A prominent example which will be central to this work is the *lexicographic* order, where $x_1 \succ \ldots \succ x_n$ and for any two monomials $x^{\alpha_1}, x^{\alpha_2} \in \mathbb{F}[x_1, \ldots, x_n]$, it holds that $x^{\alpha_1} \succ x^{\alpha_2}$ if and only if the leftmost entry of the vector difference $\alpha_1 - \alpha_2 \in \mathbb{Z}^n$ is positive.

Fixing an admissible order $\prec$ on the monomials in $\mathbb{F}[x_1, \ldots, x_n]$, the *leading term* of a polynomial is its term with largest monomial with respect to $\prec$. A polynomial p is *reducible modulo* an ideal I if there exists $q \in I$ with the same leading term as p; otherwise, p is *irreducible* modulo I. It is straightforward to show that for any polynomial p and any ideal I, there exists a unique representation $p = q + r$ where $q \in I$ and where r is a linear combination of irreducible monomials modulo I. The polynomial r is called the *reduction* of p modulo I, and the operator R_I on $\mathbb{F}[x_1, \ldots, x_n]$ that takes a polynomial p to its reduction modulo I is called the *reduction operator modulo* I. A straightforward argument shows that the reduction operator is $\mathbb{F}$-linear.

Finally, we recall the Boolean Nullstellensatz, omitting its standard proof (see, e.g., [CdRN+25, Appendix A]).

Lemma 2.1 (Boolean Nullstellensatz)**.** *For any finite set $\mathcal{P} \subseteq \mathbb{F}[x_1, \ldots, x_n]$ that contains the set of Boolean axioms $\{x_1^2 - x_1, \ldots, x_n^2 - x_n\}$, and any polynomial $q \in \mathbb{F}[x_1, \ldots, x_n]$, q vanishes on all common roots of the polynomials in $\mathcal{P}$ if and only if $q \in \langle \mathcal{P} \rangle$. In particular, $1 \in \langle \mathcal{P} \rangle$ if and only if the polynomials in $\mathcal{P}$ have no common root.*

2.4 Algebraic Proof Complexity

Nullstellensatz [BIK+94] is a proof system that uses algebraic reasoning to prove that a set $\mathcal{P} = \{p_1, \ldots, p_m\}$ of polynomials in variables $x_1, \ldots, x_n$ over a field $\mathbb{F}$ is unsatisfiable, that is, that the polynomials in $\mathcal{P}$ have no common root in $\mathbb{F}$. 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. A *Nullstellensatz refutation* of $\mathcal{P}$ is a certificate that 1 is in the ideal generated by $\mathcal{P}$, in the form of a polynomial identity

$$\sum_{i=1}^{m} p_i q_i = 1 \tag{1}$$

where $q_1, \ldots, q_m$ are arbitrary polynomials in $\mathbb{F}[x_1, \ldots, x_n]$.

Nullstellensatz is sound: no certificate (1) can exist if the polynomials in $\mathcal{P}$ have a common root, since assigning it to both sides of (1) results in $0 = 1$. It is well known that Nullstellensatz is complete over any field $\mathbb{F}$ when the variables $x_1, \ldots, x_n$ can only take values in $\{0, 1\}$, meaning that the *Boolean axioms* $\{x_i^2 - x_i : i \in [n]\}$ are present in $\mathcal{P}$.[5] The Boolean axioms are present in all sets $\mathcal{P}$ considered in this paper. The primary complexity measure of a Nullstellensatz refutation as in (1) is *degree*, defined as the maximum degree among the polynomials in it.

The proof system considered in this paper is *polynomial calculus* [CEI96], which can be seen as a dynamic version of Nullstellensatz. Starting from $\mathcal{P}$, polynomial calculus derives new polynomials in the ideal $\langle \mathcal{P} \rangle$ through two derivation rules:

$$\textit{Linear combination:} \ \frac{p \quad q}{ap + bq}, \ a, b \in \mathbb{F}\,; \tag{2a}$$

$$\textit{Multiplication:} \ \frac{p}{x_i p}, \ x_i \text{ any variable.} \tag{2b}$$

A *polynomial calculus derivation* of a polynomial p from the set $\mathcal{P}$ is a sequence of polynomials $(p_1, \ldots, p_\tau)$, where $p_\tau = p$ and each p_i is either in $\mathcal{P}$ or obtained by applying one of (2a)–(2b) to polynomials p_j with $j < i$. A *polynomial calculus refutation of* $\mathcal{P}$ is a derivation of 1 from $\mathcal{P}$.

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.

This paper only concerns degree lower bounds for polynomial calculus, but nevertheless we should mention that a strong enough polynomial calculus degree lower bound immediately implies a size lower bound through the *size-degree relation* [IPS99]: if all polynomials in a set $\mathcal{P} \subseteq \mathbb{F}[x_1, \ldots, x_n]$ have degree at most D_0 and any polynomial calculus refutation of $\mathcal{P}$ requires degree at least D, then any polynomial calculus refutation of $\mathcal{P}$ requires size $\exp(\Omega((D - D_0)^2/n))$.

For technical reasons, most papers about polynomial calculus in fact consider the better-behaved proof system *polynomial calculus resolution (PCR)* [ABRW02], where additionally each variable x_i appearing in $\mathcal{P}$ has a formal negation $\overline{x}_i$ enforced by adding the set of polynomials $\{x_i + \overline{x}_i - 1 : i \in [n]\}$ to $\mathcal{P}$. Polynomial calculus and PCR are equivalent with respect to degree; hence, we will be lax in differentiating between them in this paper.

The most common technique for proving polynomial calculus degree lower bounds is by constructing a *pseudo-reduction operator*, as mentioned in Section 1.5.

Definition 2.2 (Pseudo-reduction operator)**.** Let $\mathbb{F}$ be a field, let $D \in \mathbb{N}^+$ and let $\mathcal{P} \subseteq \mathbb{F}[x_1, \ldots, x_n]$ be unsatisfiable. An $\mathbb{F}$-linear operator R on $\mathbb{F}[x_1, \ldots, x_n]$ is a *degree-D pseudo-reduction operator for $\mathcal{P}$* if

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

The kernel of a pseudo-reduction operator provides an overapproximation of what can be derived from $\mathcal{P}$ in degree at most D which is fine-grained enough to not contain 1. The existence of a degree-D pseudo-reduction operator for a set $\mathcal{P}$ of polynomials implies a polynomial calculus degree lower bound of $D + 1$ for refuting $\mathcal{P}$, as stated in Lemma 2.3.[6]

[5] More generally, Nullstellensatz is complete over any field whenever the domains of $x_1, \ldots x_n$ are all finite.

[6] It is folklore that the converse direction also holds, and hence that degree-D pseudo-reduction operators for $\mathcal{P}$ characterize the existence of degree-D polynomial calculus refutations of $\mathcal{P}$. However, to the best of our knowledge, this result has never been published.

Lemma 2.3 ([Raz98], Lemma 3.2)**.** *Let $D \in \mathbb{N}^+$ and let $\mathcal{P}$ be a set of polynomials over $\mathbb{F}[x_1, \ldots, x_n]$. If there exists a degree-D pseudo-reduction operator for $\mathcal{P}$, then any polynomial calculus refutation of $\mathcal{P}$ over $\mathbb{F}$ requires degree strictly greater than D.*

Proof sketch. Apply R to all polynomials in a purported polynomial calculus refutation of $\mathcal{P}$ of degree at most D. Conclude by induction on refutation length that it is impossible to reach the constant polynomial 1. □

2.5 Polynomial Encoding of Homomorphisms

Given an instance $\mathbf{A}$ of CSP($\mathbf{T}$), we encode the statement that $\mathbf{A} \to \mathbf{T}$ as a set of polynomials $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ whose common roots correspond precisely to the homomorphisms $\mathbf{A} \to \mathbf{T}$. We use Boolean variables $\{x_{a,i} : a \in A, i \in T\}$ to indicate that $a \in A$ is mapped to $i \in T$. For a field $\mathbb{F}$, we denote the polynomial ring over $\mathbb{F}$ in these variables by $\mathbb{F}[\mathbf{x}_{A,T}]$. A (partial) function $f\colon A \to T$ is thus encoded as a monomial $m_f = \prod_{(a,f(a))} x_{a,f(a)}$. For a monomial $m \in \mathbb{F}[\mathbf{x}_{A,T}]$, we denote the set of vertices in A that appear in some variable of m by $\mathrm{Verts}_A(m)$. The set $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ consists of the polynomials in (3a)–(3d).

$$\sum_{i\in T} x_{a,i} - 1 \qquad a \in A \qquad \text{[every vertex } a \in A \text{ mapped to some } i \in T\text{]} \quad (3a)$$

$$x_{a,i}x_{a,i'} \qquad a \in A,\ i \neq i' \in T \qquad \text{[no } a \in A \text{ mapped to } > 1 \text{ vertex } i \in T\text{]} \quad (3b)$$

$$\prod_{a\in\mathbf{a}} x_{a,f(a)} \qquad \Gamma \in \sigma,\ \mathbf{a} \in \Gamma^{\mathbf{A}},\ f\colon \mathbf{a} \mapsto f(\mathbf{a}) \in T^{\mathrm{ar}(\Gamma)} \setminus \Gamma^{\mathbf{T}} \qquad \text{[no tuple in } \Gamma^{\mathbf{A}} \text{ mapped to tuple not in } \Gamma^{\mathbf{T}}\text{]} \quad (3c)$$

$$x_{a,i}^2 - x_{a,i} \qquad a \in A,\ i \in T \qquad \text{[Boolean axioms]} \quad (3d)$$

2.6 CSP Hierarchies

In this section we define the CSP hierarchy algorithms that we consider, namely the *k-consistency* algorithm and its *$\mathbb{Z}$-affine* and *cohomological* strengthenings.

The k-Consistency Algorithm: Given an instance $\mathbf{A}$ of CSP($\mathbf{T}$), the well-known *k-consistency algorithm* accepts $\mathbf{A}$ if and only if there exists a nonempty collection $\kappa = \{\kappa(X) \mid X \in \binom{A}{\leq k}\}$ which contains, for each $X \in \binom{A}{\leq k}$, a set $\kappa(X)$ of partial homomorphisms $\mathbf{A}[X] \to \mathbf{T}$ such that the following conditions are satisfied for all $Y \subseteq X \in \binom{A}{\leq k}$:

> *Extendability:* Each homomorphism $\varphi \in \kappa(Y)$ extends to some $\psi \in \kappa(X)$.[7]
>
> *Down-closure:* For every homomorphism $\varphi \in \kappa(X)$, it holds that $\varphi \restriction_Y \in \kappa(Y)$.

A collection κ satisfying the extendability and down-closure conditions is called *k-consistent*. Since a union of k-consistent collections is again k-consistent, there is a unique maximal k-consistent collection in any collection of partial homomorphisms.

$\mathbb{Z}$-Affine k-Consistency: The *$\mathbb{Z}$-affine k-consistency algorithm* [DO24, Section 4.3] is a combination of the k-consistency algorithm and an integer affine relaxation which, for an instance $\mathbf{A}$ of CSP($\mathbf{T}$), proceeds as follows. First, the k-consistency algorithm is run to obtain a k-consistent collection $\kappa = \{\kappa(X) \mid X \in \binom{A}{\leq k}\}$ of partial homomorphisms. Then, the system $L_k(\mathbf{A}, \mathbf{T}, \kappa)$ of linear equations in (4a)–(4b) is solved over $\mathbb{Z}$. The system $L_k(\mathbf{A}, \mathbf{T}, \kappa)$ is over

[7] This condition is called *forth* by some authors in constraint satisfaction, and *flasqueness* or *flabbiness* in sheaf terminology.

the variables $x_{X,\varphi}$, for all $X \in \binom{A}{\leq k}$ and all $\varphi \in \mathrm{Hom}(\mathbf{A}[X], \mathbf{T})$.

$$\sum_{\varphi \in \kappa(X)} x_{X,\varphi} = 1 \qquad \text{for all } X \in \binom{A}{\leq k} \tag{4a}$$

$$\sum_{\varphi \in \kappa(X), \varphi\restriction_Y = \psi} x_{X,\varphi} = x_{Y,\psi} \qquad \text{for all } Y \subseteq X \in \binom{A}{\leq k} \text{ and } \psi \in \kappa(Y) \tag{4b}$$

The $\mathbb{Z}$-affine k-consistency algorithm accepts $\mathbf{A}$ if and only if there exists a solution to $L_k(\mathbf{A}, \mathbf{T})$ for some k-consistent set κ (without loss of generality, the maximal such κ). Note that if k is at least the maximal arity among the relations in σ, then the equations (4a)–(4b) have a $\{0,1\}$-valued solution over the full set of partial homomorphisms of domain size at most k (which is not necessarily k-consistent) if and only if $\mathbf{A} \to \mathbf{T}$.

Cohomological k-Consistency: The *cohomological k-consistency* algorithm due to Ó Conghaile [ÓC22] further strengthens the $\mathbb{Z}$-affine k-consistency algorithm. The algorithm was originally formulated in sheaf-theoretic language but can also be defined in more elementary terms as follows.

Cohomological k-consistency algorithm for $\mathrm{CSP}(\mathbf{T})$ *on an instance* $\mathbf{A}$:

1. Maintain a collection $\mathcal{H} = \{\mathcal{H}(X) \mid X \in \binom{A}{\leq k}\}$ of partial homomorphisms $\mathbf{A} \to \mathbf{T}$. For each $X \in \binom{A}{\leq k}$, initialize $\mathcal{H}(X) = \mathrm{Hom}(\mathbf{A}[X], \mathbf{T})$.
2. Repeat until each set $\mathcal{H}(X)$ stabilizes:
   (a) Find the maximal k-consistent collection κ in $\mathcal{H}$. Update $\mathcal{H} \leftarrow \kappa$.
   (b) For each $X \in \binom{A}{\leq k}$ and $\varphi \in \mathcal{H}(X)$, check whether $L_k(\mathbf{A}, \mathbf{T}, \mathcal{H})$ has a solution satisfying $x_{X,\varphi} = 1$ and $x_{X,\psi} = 0$ for every $\psi \in \mathcal{H}(X) \setminus \{\varphi\}$. If not, remove φ from $\mathcal{H}(X)$ for the next iteration of the loop.
3. If $\mathcal{H}(X) = \varnothing$ for some $X \in \binom{A}{\leq k}$, then reject; otherwise accept.

The k-consistency algorithm always correctly decides whether $\mathbf{A} \to \mathbf{T}$ when $k = |A|$ (but then does not necessarily run in time polynomial in $\mathbf{A}$). Hence, the same is true for the $\mathbb{Z}$-affine and cohomological k-consistency algorithms, and a level lower bound of $\Omega(|A|)$ is therefore asymptotically optimal for all three hierarchies.

It is not hard to see that all algorithms for $\mathrm{CSP}(\mathbf{T})$ in this section always accept $\mathbf{A}$ if $\mathbf{A} \to \mathbf{T}$. To prove that these algorithms do not solve $\mathrm{CSP}(\mathbf{T})$, we should therefore exhibit a *fooling instance*: a structure $\mathbf{A}$ such that $\mathbf{A} \not\to \mathbf{T}$, but such that the algorithms still accept $\mathbf{A}$. To show that a hierarchy does not solve a PCSP with template $(\mathbf{S}, \mathbf{T})$, we should find an instance $\mathbf{A}$ such that $\mathbf{A} \not\to \mathbf{T}$, but such that the hierarchy for $\mathrm{CSP}(\mathbf{S})$ accepts $\mathbf{A}$. This setting is highly reminiscent of that in proof complexity, and in Sections 3 and 4 we show how techniques from proof complexity can be leveraged to construct fooling instances.

3 CSP Hierarchies and Pseudo-Reduction Operators

In this section, we show that the existence of a pseudo-reduction operator satisfying a certain integrality condition implies level lower bounds for the $\mathbb{Z}$-affine k-consistency hierarchy. In Section 4, we provide a method for constructing such pseudo-reduction operators; this construction also lets us extend our lower bounds to hold for the cohomological k-consistency algorithm.

We are particularly interested in pseudo-reduction operators R for $\mathcal{P}_{\mathbf{A} \to \mathbf{T}}$ on $\mathbb{F}[\mathbf{x}_{A,T}]$ where $\mathbb{F}$ is of characteristic zero and hence contains $\mathbb{Z}$ as a subring. If restricting the domain of R to

$\mathbb{Z}[\mathbf{x}_{A,T}]$ results in an operator $\mathbb{Z}[\mathbf{x}_{A,T}] \to \mathbb{Z}[\mathbf{x}_{A,T}]$, which by linearity is equivalent to R mapping all monomials to polynomials in $\mathbb{Z}[\mathbf{x}_{A,T}]$, we say that R is an *integral* pseudo-reduction operator, or *IPR* operator for short. For a degree-D pseudo-reduction operator R for $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$, the *$R$-support* is the collection $\kappa_R = \{\kappa_R(X) : X \in \binom{A}{\leq k}\}$ where $\kappa_R(X) = \{\varphi \in \mathrm{Hom}(\mathbf{A}[X], \mathbf{T}) \mid R(m_\varphi) \neq 0\}$.

The main result of this section is Lemma 3.1.

Lemma 3.1. *If R is a degree-D pseudo-reduction operator for $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$, then κ_R is k-consistent for all $k \leq D$. Moreover, if R is integral, then for any linear $h : \mathbb{Z}[\mathbf{x}_{A,T}] \to \mathbb{Z}$ satisfying $h(1) = 1$, the assignment to the variables $x_{X,\varphi}$ of (4a)–(4b) which sets $x_{X,\varphi}$ to $h \circ R(m_\varphi)$ for $X \in \binom{A}{\leq D}$ and $\varphi \in \mathrm{Hom}(\mathbf{A}[X], \mathbf{T})$ is a solution to (4a)–(4b) over $\mathbb{Z}$.*

Consequently, if there exists a degree-D IPR operator for $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$, then the $\mathbb{Z}$-affine k-consistency algorithm for $\mathrm{CSP}(\mathbf{T})$ accepts $\mathbf{A}$ for all $k \leq D$.

Proof. We first observe that κ_R in fact contains the whole support of R on monomials of degree at most D. To see this, note that every monomial m in the variables of $P_{\mathbf{A}\to\mathbf{T}}$ encodes a relation $\Gamma_m \subseteq A \times T$, where $(a, i) \in A \times T$ is contained in Γ_m if and only if $x_{a,i} \in m$, with domain size at most D. If Γ_m is not a partial homomorphism $\mathbf{A} \to \mathbf{T}$, then Γ_m either maps some $a \in A$ to multiple $i \in T$ or Γ_m maps a tuple in a relation $\Gamma^{\mathbf{A}}$ in $\mathbf{A}$ to a tuple which is not in $\Gamma^{\mathbf{T}}$. In both cases, we have that $R(m) = 0$ since m is then a multiple of an axiom in (3b) or (3c). Therefore, the support of R consists of monomials that encode partial homomorphisms in A.

We first show that the collection κ_R is k-consistent for all $k \leq D$ and then show that the output of R satisfies (4a)–(4b) *as polynomials*, using only the homomorphisms in κ_R. Therefore, composing R with a linear map $h\colon \mathbb{Z}[\mathbf{x}_{A,T}] \to \mathbb{Z}$ evaluating to 1 on 1 yields a solution over $\mathbb{Z}$ to (4a)–(4b).

To prove that κ_R is D-consistent, we need to show that κ_R is non-empty and satisfies the extendability and down-closure conditions for all $k \leq D$. For the down-closure condition we use the contrapositive. Recall that for a (partial) homomorphism $\varphi\colon \mathbf{A} \to \mathbf{T}$, we denote the monomial $\prod_{(a,\varphi(a))} x_{a,\varphi(a)}$ by m_φ. Note that for $Y \subsetneq X \in \binom{A}{\leq k}$ and homomorphisms $\varphi\colon \mathbf{A}[X] \to \mathbf{T}$ and $\psi\colon \mathbf{A}[Y] \to \mathbf{T}$ such that $\varphi \restriction_Y = \psi$, it holds that m_φ is a multiple of m_ψ. Therefore, if $R(m_\psi) = 0$ then by Property 3 of a pseudo-reduction operator we have

$$R(m_\varphi) = R(m' \cdot m_\psi) = R(m' \cdot R(m_\psi)) = R(0) = 0 \ , \tag{5}$$

so if $\psi \notin \kappa_R(Y)$ then for any $X \supseteq Y$, no extension of ψ is is in $\kappa_R(X)$.

To establish the extendability condition, let $Y \subsetneq X \in \binom{A}{\leq k}$ where $X = Y \cup \{a\}$. For any homomorphism $\psi\colon \mathbf{A}[Y] \to \mathbf{T}$ in $\kappa_R[Y]$, we have by Property 3 of a pseudo-reduction operator that

$$R\Big(\Big(\sum_{i\in T} x_{a,i} - 1\Big) \cdot m_\psi\Big) = R\Big(R\Big(\sum_{i\in T} x_{a,i} - 1\Big) \cdot m_\psi\Big) = 0 \ , \tag{6}$$

and by linearity we also have

$$R\Big(\Big(\sum_{i\in T} x_{a,i} - 1\Big) \cdot m_\psi\Big) = \sum_{i\in T} R(x_{a,i} \cdot m_\psi) - R(m_\psi) \ . \tag{7}$$

Since $\psi \in \kappa_R$ and hence $R(m_\psi) \neq 0$, it follows that at least one of the extensions $x_{a,i}m_\psi$ of ψ to X is also in κ_R. With the extendability and down-closure conditions established, we conclude that κ_R is k-consistent.

Now, we show that the output of R satisfies (4a)–(4b). For (4a), we use induction on $|X|$. The base case, where $|X| = 0$, follows since $R(1) = 1$. For the induction step, let $S(X)$ denote the sum $\sum_{\varphi\in k_R(X)} m_\varphi$ for $k \leq D$ and suppose that for all sets $Y \in \binom{A}{\leq k-1}$ we have that $R(S(Y)) = 1$. For any set $X = Y \cup \{a\}$ of size k, note that all summands in $S(X)$ are contained in the sum $\sum_{i\in T} x_{a,i} \cdot S(Y)$ by k-consistency, and by Property 3 of a pseudo-reduction operator we have

$$R\Big(\sum_{i\in T} x_{a,i} \cdot (S(Y) - 1)\Big) = R\Big(\sum_{i\in T} x_{a,i} \cdot R(S(Y) - 1)\Big) = 0 \ . \tag{8}$$

Moreover, we have by linearity that

$$R\Big(\sum_{i\in T} x_{a,i}\cdot(S(Y)-1)\Big) = R\Big(\sum_{i\in T} x_{a,i}\cdot S(Y)\Big) - R\Big(\sum_{i\in T} x_{a,i}\Big) \tag{9a}$$

$$= R\Big(\sum_{i\in T} x_{a,i}\cdot S(Y)\Big) - 1 \qquad \text{[induction hypothesis]} \tag{9b}$$

$$= R\Big(S(X)\Big) - 1\ , \qquad \text{[definition of } \kappa_R\text{]} \tag{9c}$$

which proves that (4a) is satisfied when composing R with any integer-valued $\mathbb{F}$-linear map h such that $h(1) = 1$. To establish (4b), first observe that for $Y \subsetneq X \in \binom{A}{\leq k}$ we have that $R(S(X \setminus Y)) = 1$. Fix a homomorphism $\psi : \mathbf{A}[Y] \to \mathbf{T}$ in κ_R. By Property 3 of a pseudo-reduction operator, we have

$$R((S(X\setminus Y)-1)\cdot m_\psi) = R(R((S(X\setminus Y)-1))\cdot m_\psi) = 0 \tag{10}$$

and moreover

$$0 = R((S(X\setminus Y)-1)\cdot m_\psi) \tag{11a}$$

$$= \Big(\sum_{\mu\in\kappa_R(X\setminus Y)} R(m_\mu\cdot m_\psi)\Big) - R(m_\psi) \tag{11b}$$

$$= \Big(\sum_{\varphi\in\kappa_R(X):\varphi\restriction_Y=\psi} R(m_\varphi)\Big) - R(m_\psi) \qquad \text{[definition of } \kappa_R\text{]} \tag{11c}$$

which establishes (4b). Hence, again composing R with any linear map $h : \mathbb{Z}[\mathbf{x}_{A,T}] \to \mathbb{Z}$ such that $h(1) = 1$ furnishes a level-D solution for the $\mathbb{Z}$-affine k-consistency algorithm. □

Remark 3.2. The notion of an IPR operator R for a set $\mathcal{P}$ is only well-defined over fields of characteristic zero. Note, however, that the existence of such an operator R over any field of characteristic zero implies a polynomial calculus degree lower bound for refuting $\mathcal{P}$ over *all* fields, even those of finite (prime) characteristic ℓ. This is because $((\cdot) \mod \ell) \circ R$ is then a pseudo-reduction operator over the finite field $\mathbb{F}_\ell$, and every field of characteristic ℓ contains $\mathbb{F}_\ell$ as a subfield.

4 From Pseudo-Reduction To Cohomological k-Consistency

In Section 3, we showed that the existence of a degree-D IPR operator for $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ implies that the $\mathbb{Z}$-affine k-consistency algorithm for CSP($\mathbf{T}$) accepts $\mathbf{A}$ for all $k \leq D$. This section introduces the *Alekhnovich–Razborov method* [Raz98, AR03] for constructing such operators. The proof that operators produced by this method are integral is deferred to Section 5. Given integrality, this section concludes by showing that a degree-D pseudo-reduction operator for $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ constructed using the Alekhnovich–Razborov method can be used to prove that $\mathbf{A}$ is a fooling instance for the cohomological k-consistency algorithm for CSP($\mathbf{T}$); this is formally established in Lemma 4.4.

The treatment of the Alekhnovich–Razborov method in this section partly follows that of [CdRN+25, Section 4.2] but is specialized to sets of polynomials $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ as in Section 2.5. For more general perspectives, see [AR03] and the more recent papers [MN24, CdRN+25].

Recall from Section 1.5 that if $\mathbf{A} \to \mathbf{T}$, then the true reduction operator $R_{\langle\mathcal{P}_{\mathbf{A}\to\mathbf{T}}\rangle}$ is a degree-D pseudo-reduction operator for all D. For structures $\mathbf{A} \not\to \mathbf{T}$, the Alekhnovich–Razborov method constructs operators that attempt to mimic the behavior of $R_{\langle\mathcal{P}_{\mathbf{A}\to\mathbf{T}}\rangle}$ as closely as possible while mapping 1 to 1. The construction is as follows: Let $\prec$ be an admissible ordering on the monomials in $\mathbb{F}[\mathbf{x}_{A,T}]$. To each monomial m, associate a *satisfiable* set $S(m) \subseteq \mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ and define $R(m)$ to be $R^{\prec}_{\langle S(m)\rangle}(m)$. Finally, define R on arbitrary polynomials by linear extension.

For a set of polynomials $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ and a monomial $m = \prod x_{a,i}$ in $\mathbb{F}[\mathbf{x}_{A,T}]$, a natural choice for $S(m)$ is a set $\mathcal{P}_{\mathbf{A}[X_m]\to\mathbf{T}} \subseteq \mathcal{P}_{\mathbf{A}\to\mathbf{T}}$, where $X_m \subseteq A$ depends on m and $\mathbf{A}[X_m]$ is a satisfiable substructure of $\mathbf{A}$. In addition, we require that X_m only depends on the vertices of A mentioned by m. To obtain lower bounds for the cohomological k-consistency algorithm, it is necessary for our proof that $S(m)$ be of this form; see Lemma 4.4. Therefore, we only consider such sets $S(m)$ and again refer the reader to [CdRN+25, Section 4.2] for a more general treatment. Given a set $X \subseteq A$, we denote the ideal $\langle \mathcal{P}_{\mathbf{A}[X]\to\mathbf{T}} \rangle$ by $\langle X \rangle$ for brevity when the relational structures $\mathbf{A}$ and $\mathbf{T}$ are clear from context.

As previously noted, the substructure $\mathbf{A}[X_m]$ should be the part of $\mathbf{A}$ that is "relevant" to m. The next definition takes the first step toward formalizing this intuition by putting some reasonable constraints on S, which are precisely those of the well-known notion of a *closure operator*. A closure operator takes each set $X \subseteq A$ to its *closure* $\mathrm{cl}(X)$, which the smallest "closed" set containing X. Here, a closed set $C \subseteq A$ with respect to $\mathbf{A}$ and $\mathbf{T}$ is one where the substructure $\mathbf{A}[C]$ contains the part of $\mathbf{A}$ that is "relevant" for reducing monomials that mention only vertices in C modulo $\langle \mathcal{P}_{\mathbf{A}\to\mathbf{T}} \rangle$, in a sense we make precise in Lemma 4.2. In general, the definition of a closed set varies with the application; see Sections 6 and 7 for two examples.

Definition 4.1 (Size-D closure operator)**.** Given a set A, a map $\mathrm{cl}\colon 2^A \to 2^A$ is a *size-D closure operator* if the following conditions hold for all sets $X, Y \subseteq A$ of size at most D.

Self-containment: $X \subseteq \mathrm{cl}(X)$.

Monotonicity: If $X \subseteq Y$, then $\mathrm{cl}(X) \subseteq \mathrm{cl}(Y)$.

Idempotence: $\mathrm{cl}(\mathrm{cl}(X)) = \mathrm{cl}(X)$.

If these conditions hold for $D = |A|$, we say that cl is a *closure operator.*

Given a closure operator $\mathrm{cl}(\cdot)$ on A, we define the closure of a monomial $m \in \mathbb{F}[\mathbf{x}_{A,T}]$ to be the set $\mathrm{cl}(m) = \mathrm{cl}(\mathrm{Verts}_A(m)) \subseteq A$. The next lemma formalizes that the closure of a monomial should be the part of A such that $\mathbf{A}[\mathrm{cl}(m)]$ contains all "relevant" information about reducing m modulo $\langle \mathcal{P}_{\mathbf{A}\to\mathbf{T}} \rangle$, in the sense that adding a small set to $\mathrm{cl}(m)$ does not change the result of reduction. We formulate this property as the so-called *reducibility condition*, following [CdRN+25].

Lemma 4.2. *Let $D \in \mathbb{N}$, let σ be a signature whose relation symbols all have arity at most D, let $\mathbf{T}$ be a σ-structure and $\mathbf{A}$ an instance of $\mathrm{CSP}(\mathbf{T})$, and let $\mathrm{cl}\colon 2^A \to 2^A$ be a size-D closure operator for A. Finally, let $\prec$ be an admissible order of the monomials in $\mathbb{F}[\mathbf{x}_{A,T}]$. Suppose that the following conditions hold for every monomial $m \in \mathbb{F}[\mathbf{x}_{A,T}]$ of degree at most D:*

Satisfiability: *It holds that $\mathbf{A}[\mathrm{cl}(m)] \to \mathbf{T}$.*

Reducibility: *For every monomial m' such that $\mathrm{cl}(m') \subseteq \mathrm{cl}(m)$, it holds that m' is reducible modulo $\langle \mathrm{cl}(m') \rangle$ if and only if m' is reducible modulo $\langle \mathrm{cl}(m) \rangle$.*

Then, the $\mathbb{F}$-linear extension R of the map $m \mapsto R^{\prec}_{\langle \mathrm{cl}(m) \rangle}(m)$ is a degree-D pseudo-reduction operator for $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$.

The proof of Lemma 4.2 can be found in Section A, and is by a now-standard argument first introduced by Alekhnovich and Razborov [AR03] and subsequently clarified by Galesi and Lauria [GL10a, GL10b] as well as Mikša and Nordström [MN24].

In Section 5, we prove that operators $\mathbb{F}[\mathbf{x}_{A,T}] \to \mathbb{F}[\mathbf{x}_{A,T}]$ defined as in Lemma 4.2 with respect to a lexicographic order restrict to operators $\mathbb{Z}[\mathbf{x}_{A,T}] \to \mathbb{Z}[\mathbf{x}_{A,T}]$. We record this fact as Lemma 4.3 and defer its proof to after we prove Lemma 5.3.

Lemma 4.3 (Integrality assumption)**.** *Let* $\mathbf{A}, \mathbf{T}, \mathrm{cl}(\cdot)$*, and* D *be defined as in Lemma 4.2. Let* $\mathbb{F}$ *be any field of characteristic* 0 *and let* $\prec_{\mathrm{lex}}$ *be any lexicographic order on the monomials in* $\mathbb{F}[\mathbf{x}_{A,T}]$*. For any monomial* m *of degree at most* D*, it holds that* $R^{\prec_{\mathrm{lex}}}_{\langle \mathrm{cl}(m)\rangle}(m) \in \mathbb{Z}[\mathbf{x}_{A,T}]$*.*

Given Lemma 4.3, we can establish our main technical lemma needed to prove lower bounds for the cohomological k-consistency algorithm.

Lemma 4.4. *Suppose that the assumptions of Lemma 4.2 are satisfied for all monomials* $m \in \mathbb{F}[\mathbf{x}_{A,T}]$ *of degree at most* D*, where* $\mathbb{F}$ *has characteristic zero and monomials are ordered by a lexicographic order* $\prec_{\mathrm{lex}}$*. Then, the cohomological* k*-consistency algorithm for* $\mathrm{CSP}(\mathbf{T})$ *accepts* $\mathbf{A}$ *for all* $k \leq s$*, where* s *is the largest integer such that* $|\mathrm{cl}(m)| \leq D$ *for all monomials* m *of degree at most* s*.*

Proof. By Lemma 4.3, the operator $R\colon \mathbb{F}[\mathbf{x}_{A,T}] \to \mathbb{F}[\mathbf{x}_{A,T}]$ in the statement of Lemma 4.2 restricts to an operator $\mathbb{Z}[\mathbf{x}_{A,T}] \to \mathbb{Z}[\mathbf{x}_{A,T}]$, and by Lemma 4.2, we have that R is a degree-D pseudo-reduction operator for $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$.

By Lemma 3.1, the collection $\kappa_R = \{\kappa_R(X) : X \in \binom{A}{\leq k}\}$ where $\kappa_R(X) = \{\varphi \in \mathrm{Hom}(\mathbf{A}[X], \mathbf{T}) \mid R(m_\varphi) \neq 0\}$ is k-consistent for all $k \leq D$. Moreover, also by Lemma 3.1 the assignment to the variables $x_{X,\varphi}$ of (4a)–(4b) which sets $x_{X,\varphi}$ to $\mathrm{eval}_y \circ R(m_\varphi)$ for $X \in \binom{A}{\leq k}$ and $\varphi \in \mathrm{Hom}(\mathbf{A}[X], \mathbf{T})$ is a solution to (4a)–(4b) over $\mathbb{Z}$ for any point $y \in \mathbb{Z}^{|A|\cdot|T|}$. Therefore, it suffices to show that for all sets $X \in \binom{A}{\leq s}$ and every homomorphism $\varphi \in \kappa_R(X)$, there exists a point $y \in \mathbb{Z}^{|A|\cdot|T|}$ such that $\mathrm{eval}_y \circ R(m_\varphi) = 1$ and $\mathrm{eval}_y \circ R(m_\psi) = 0$ for all homomorphisms $\psi \neq \varphi \in \kappa_R(X)$.

Denote the polynomial $R(m_\varphi)$ by r, and recall that $r = R^{\prec_{\mathrm{lex}}}_{\langle \mathrm{cl}(m_\varphi)\rangle}(m_\varphi)$, which by definition is the unique sum of irreducible terms modulo $\langle \mathrm{cl}(m_\varphi)\rangle = \langle \mathrm{cl}(X)\rangle$ such that $m_\varphi = r + q$, where $q \in \langle \mathrm{cl}(X)\rangle$. Note that q vanishes when evaluated on (the indicator vector of) any homomorphism $\mathbf{A}[\mathrm{cl}(X)] \to \mathbf{T}$ and that $m_\varphi = 1$ when evaluated on φ. Since $\mathrm{Verts}_A(m_\varphi) = X \subseteq \mathrm{cl}(X)$, it follows that r evaluates to 1 on any homomorphism $\mathbf{A}[\mathrm{cl}(X)] \to \mathbf{T}$ that extends φ. Moreover, for all homomorphisms $\psi \neq \varphi \in \kappa_R(X)$, the monomial m_ψ evaluates to 0 on that assignment since $\mathrm{cl}(m_\varphi) = \mathrm{cl}(m_\psi)$, and hence $R(m_\psi)$ also does. Therefore, to conclude the proof it remains to show that for all homomorphisms $\varphi\colon \mathbf{A}[X] \to \mathbf{T}$, there exists a homomorphism $\mu\colon \mathbf{A}[\mathrm{cl}(X)] \to \mathbf{T}$ that extends φ. But this follows immediately from the D-consistency of κ_R, since $|\mathrm{cl}(X)| \leq D$ by assumption.

Thus, the cohomological k-consistency algorithm stabilizes on κ_R, and hence the cohomological k-consistency algorithm for $\mathrm{CSP}(\mathbf{T})$ accepts $\mathbf{A}$ for all $k \leq s$ as desired. □

Remark 4.5 (Room of choice of solutions)**.** In the proof of Lemma 4.4, there exist many maps $h : \mathbb{F}[x] \to \mathbb{F}$ whose composition with R fools the cohomological k-consistency algorithm. Specifically, given a k-small X and a homomorphism $\varphi\colon \mathbf{A}[X] \to \mathbf{T}$ where $R(m_\varphi) \neq 0$, we can let h be an arbitrary $\mathbb{F}$-linear extension of the evaluation map $\mathrm{eval}_\mu\colon \mathbb{F}[\mathbf{x}_{\mathrm{cl}(X),T}] \to \mathbb{F}$ from the subspace $\mathbb{F}[\mathbf{x}_{\mathrm{cl}(X),T}]$ to $\mathbb{F}[\mathbf{x}_{A,T}]$, where μ is an arbitrary homomorphism on $\mathbf{A}[\mathrm{cl}(X)]$ extending φ (which exists by the D-consistency of κ_R). Then $h \circ R$ provides a $\mathbb{Z}$-affine k-consistency solution satisfying $h \circ R(\varphi) = 1$ and $h \circ R(\psi) = 0$ for all other $\psi\colon X \to T$.

Remark 4.6 (On requiring a lexicographic order)**.** Insisting on a lexicographic order in Lemmas 4.3 and 4.4 might seem like a technical artifact of our proof techniques, but there is in fact reason to believe that the same result does not hold for other orders. For instance, the pseudo-reduction operators in [AR03] can be put on the same form as in Lemma 4.2, but are defined with respect to a *degree-respecting* monomial order, which the lexicographic order is not. This requirement turns out to be crucial for their proof that the reducibility condition holds (see [AR03, Theorem 3.1]). Moreover, the operators in [AR03] cannot be integral in general since they then apply over every field (see Remark 3.2); this is impossible, since some of the lower bounds in [AR03] apply to sets of polynomials that are easy to refute over some fields.

5 Integrality of Reductions and the Lex Game

In this section we establish an integrality property of polynomial reductions (Lemma 5.3), from which Lemma 4.3 follows as an easy corollary. This property says that for any finite set $V \subseteq \{0,1\}^n$ and any field $\mathbb{F}$ of characteristic 0, the reduction of a monomial $m \in \mathbb{F}[x_1, \ldots, x_n]$ modulo the ideal $I(V)$ under the lexicographic order is always a polynomial with coefficients in $\mathbb{Z}$.

The proof of Lemma 5.3 uses a game characterization of reducibility under a lexicographic order, namely the *lex game* of Felszeghy, Ráth, and Rónyai [FRR06]; we use an equivalent formulation from [FR09].

Definition 5.1 (Lex game; [FR09], Section 3)**.** Let $\mathbb{F}$ be a field, $V \subseteq \mathbb{F}^n$ a finite set, and $\alpha = (\alpha_1, \ldots, \alpha_n) \in \mathbb{N}^n$ an n-dimensional vector of natural numbers. The *lex game* $\mathrm{Lex}(V; \alpha)$ is the following deterministic game played between two players named Lea and Stan. Lea and Stan both know V and α. If $V = \varnothing$, Lea wins immediately. For a nonempty V, the lex game proceeds in rounds. In round i, Lea picks α_{n-i+1} elements of $\mathbb{F}$, and Stan picks a value y_{n-i+1} which is different from all of Lea's choices in that round; if he cannot do so, Lea wins immediately. After n rounds, if Stan's answers $(y_1, \ldots, y_n)$ is a vector in V, he wins; otherwise, Lea wins.

The lex game is finite, deterministic, and cannot end in a draw. Therefore, for any V and α, one of the players has a winning strategy for $\mathrm{Lex}(V; \alpha)$. The main result of [FRR06] is that the existence of a winning strategy for Lea characterizes whether the monomial x^α is reducible modulo $I(V)$ under the lexicographic order.

Theorem 5.2 ([FRR06], Theorem 2)**.** *For any field $\mathbb{F}$, finite set $V \subseteq \mathbb{F}^n$, and $\alpha \in \mathbb{N}^n$, the monomial $x^\alpha \in \mathbb{F}[x_1, \ldots, x_n]$ is reducible modulo $I(V)$ in the lexicographic order induced by $x_1 \succ \ldots \succ x_n$ if and only if Lea has a winning strategy in the game* $\mathrm{Lex}(V; \alpha)$.

Following [FRR06], we show the following corollary of the game characterization. We note that similar results were also shown in [ARS02, Section 4], [Fel07, Corollary 5.1.2], and [FR09, Corollary 3] by arguments different from ours.

Lemma 5.3 (Integrality lemma)**.** *If $\mathbb{F}$ is a field of characteristic zero and V is a subset of $\{0,1\}^n \subseteq \mathbb{F}^n$, then for any monomial $m = x^\alpha \in \mathbb{F}[x_1, \ldots, x_n]$ it holds that $R^{\mathrm{lex}}_{I(V)}(m) \in \mathbb{Z}[x_1, \ldots, x_n]$.*

Intuitively, given a reducible monomial m, we can express Lea's guesses in her winning strategy as a set of polynomials $f_1(x), \ldots, f_n(x)$, where $x = (x_1, \ldots, x_n)$ represents Stan's answers and each f_i is a Boolean-valued function that depends only on $\{x_{i+1}, \ldots, x_n\}$. We will use these polynomials to rewrite m into smaller terms.

Proof. We use induction over the lexicographic order of $m = x^\alpha$. Without loss of generality, we may assume m is multilinear, as $m = m_{\mathrm{mult}}$ modulo $\langle x_1^2 - x_1, \ldots, x_n^2 - x_n \rangle$ where the latter ideal is $I(\{0,1\}^n) \subseteq I(V)$.

The base case $m = 1$ is immediate: 1 either reduces to 0 or itself, both of which are integers.

For the induction step, we denote $I = I(V)$. If m is irreducible modulo I then $R_I(m) = m$, which trivially has integer coefficients. Hence, we may suppose that m is reducible modulo I. By Theorem 5.2, in this case Lea has a winning strategy in the game $\mathrm{Lex}(V; \alpha)$. Recall that Lea makes a guess for the value of x_i for each i where $\alpha_i = 1$, her guess for x_i depends only on Stan's previous answers, and her guesses are $\{0,1\}$-valued since $V \subseteq \{0,1\}^n$. Therefore, we can write her guess for x_i as a polynomial

$$f_i(x_{i+1}, \ldots, x_n) = \sum_{a \in \{0,1\}^{n-i}} c_a m_a$$

where $m_a = \prod_{j:a_j=1} x_j \prod_{j:a_j=0}(1-x_j)$ is the indicator polynomial of "$x = a$" and $c_a \in \{0,1\}$ is her guess when Stan answered $(x_{i+1},\dots,x_n) = a$. Clearly, $f_i \in \mathbb{Z}[x_{i+1},\dots,x_n]$. As in [FRR06], we now consider the polynomial

$$p(x) = \prod_{i:\ \alpha_i=1} (x_i - f_i) \tag{12}$$

which has integer coefficients. The leading term of each factor $x_i - f_i$ is x_i as f_i depends only on $x_{i+1},\dots,x_n$, so the leading term of p is $x^\alpha = m$. Moreover, we claim that p is in the ideal I. To see this, notice that for any $a \in V$, we have that $p(a) = 0$. That is, since Stan loses the lex game instance corresponding to α, Lea can always stop him from answering a. This means that there is some $i \in [n]$ such that $f_i(a_{i+1},\dots,a_n) = a_i$. It follows that p vanishes on V and so is contained in I.

We now complete the induction by considering the reduction of $m - p$. As $p \in I$, we have $R_I(p) = 0$ so $R_I(m) = R_I(m-p)$. We have shown that $p \in \mathbb{Z}[x_1,\dots,x_n]$ and that p has leading term m, so $m - p = \sum_i d_i m_i$ where $m_i \prec m$ for all i and all d_i are integers. Therefore, $R_I(m-p) = \sum_i d_i R_I(m_i)$ where each $R_I(m_i)$ is in $\mathbb{Z}[x_1,\dots,x_n]$ by the induction hypothesis. It follows that $R_I(m) \in \mathbb{Z}[x_1,\dots,x_n]$. The proof is complete. □

Lemma 4.3 follows as a straightforward corollary of Lemma 5.3.

Lemma 4.3 (Restated)**.** *Let $\mathbf{A}, \mathbf{T}, \mathrm{cl}(\cdot)$, and D be defined as in Lemma 4.2. Let $\mathbb{F}$ be any field of characteristic 0 and let $\prec_{\mathrm{lex}}$ be any lexicographic order on the monomials in $\mathbb{F}[\mathbf{x}_{A,T}]$. For any monomial m of degree at most D, it holds that $R^{\prec_{\mathrm{lex}}}_{\langle \mathrm{cl}(m)\rangle}(m) \in \mathbb{Z}[\mathbf{x}_{A,T}]$.*

Proof. Note that $\langle \mathrm{cl}(m)\rangle = I(V_{\mathrm{Hom}(\mathbf{A}[\mathrm{cl}(m)],\mathbf{T})})$ as ideals in $\mathbb{F}[x_{\mathbf{A}[\mathrm{cl}(m)],\mathbf{T}}]$ by the Boolean Nullstellensatz (Lemma 2.1), where $V_{\mathrm{Hom}(\mathbf{A}[\mathrm{cl}(m)],\mathbf{T})} \subseteq \{0,1\}^{|\mathrm{cl}(m)|\cdot|T|}$ denotes the set of (indicator vectors of) homomorphisms $\mathbf{A}[\mathrm{cl}(m)] \to \mathbf{T}$. Then, by Lemma 5.3, $R^{\prec_{\mathrm{lex}}}_{\langle \mathrm{cl}(m)\rangle}(m) \in \mathbb{Z}[\mathbf{x}_{\mathrm{cl}(m),T}]$ for every monomial m over $\mathbf{x}_{\mathrm{cl}(m),T}$ that has degree at most D. The same remains true when m is over $\mathbf{x}_{A,T}$ and reductions are done within the larger polynomial ring $\mathbb{F}[\mathbf{x}_{A,T}]$. To see this, set all variables outside of $\mathbf{x}_{\mathrm{cl}(m),T}$ to 0; this assignment does not touch any variable of any generator of $\langle \mathrm{cl}(m)\rangle$, nor of m, so the reductions in the two polynomial rings coincide by uniqueness of the representation $m = R^{\prec_{\mathrm{lex}}}_{\langle \mathrm{cl}(m)\rangle}(m) + q$ for $q \in \langle \mathrm{cl}(m)\rangle$. □

Remark 5.4. For readers familiar with the polynomial calculus literature, we note that the lex game characterization of monomial reducibility broadly generalizes the *pigeon dance* criterion, established in [Raz98] and simplified in [IPS99], for ideals arising from the *pigeonhole principle.*

6 Lower Bounds for Approximate Graph Coloring

In this section, we use the machinery developed in Sections 3-5 to prove lower bounds for the level needed for the cohomological k-consistency algorithm to solve approximate graph coloring, also called c vs. ℓ-coloring. In this problem, the task is to decide whether the input graph G is either c-colorable or not even ℓ-colorable, for some $\ell \geq c$. Phrased as a PCSP, approximate graph coloring is $\mathrm{PCSP}(K_c, K_\ell)$, where K_a denotes the complete graph on a vertices.

The fooling instances are sampled from the *random regular graph model* $\mathbb{G}_{n,d}$, which is the uniform distribution over d-regular graphs on n vertices.[8] We write $G \sim \mathbb{G}_{n,d}$ to denote that G is a graph sampled from $\mathbb{G}_{n,d}$. More quantitatively, the main result in this section is Theorem 6.1.

Theorem 6.1. *There is an absolute constant C such that the following holds for natural numbers n and d such that $d \geq 10$ and $6d^3 \leq \log n$. With high probability as $n \to \infty$, a graph $G \sim \mathbb{G}_{n,d}$*

[8] Note that a d-regular graph on n vertices can only exist if dn is even. We ignore this technical issue in what follows.

has chromatic number strictly larger than $d/4\log d$, but the cohomological k-consistency algorithm for 3-colorability accepts G for all $k \leq n \cdot d^{-Cd}$. Consequently, the cohomological k-consistency algorithm does not solve 3 vs. $d/4\log d$-coloring on n-vertex graphs, for any $k \leq n \cdot d^{-Cd}$.

The same result holds for sparse Erdös-Rényi random graphs (where each edge in the graph appears independently with some fixed probability p), with a slighly worse dependence on d in the parameters. The proof is essentially the same as that of Theorem 6.1.

The rest of this section is devoted to the proof of Theorem 6.1. Fortunately, all the technical work has already been done in [CdRN+25, Sections 5-6] (partially building on ideas of [RT24]) for the purpose of proving degree lower bounds for the polynomial calculus proof system in Section 2.4. Hence, Theorem 6.1 is, more or less, an immediate consequence of their degree lower bound combined with Lemma 4.4, and all that is needed in this section is to reiterate and rephrase their proof for our setting.

We first recall the well-known fact that the chromatic number of a graph $G \sim \mathbb{G}_{n,d}$ is concentrated around $d/\log d$.

Lemma 6.2 ([KPGW10, CFRR02])**.** *With probability $1 - o_n(1)$, a graph G sampled from $\mathbb{G}_{n,d}$ where $d \leq n^{0.1}$ has chromatic number strictly larger than $d/4\log d$ and at most $d/\log d$.*[9]

The crucial property of random graphs that is used to establish both the satisfiability and reducibility conditions is *sparsity*.

Definition 6.3 (Sparsity)**.** A graph $G = (V, E)$ is (s, ε)-*sparse* if every vertex set $U \subseteq V$ of size at most s satisfies $|E(U)| \leq (1 + \varepsilon)|U|$.

As is well known, random regular graphs are sparse with high probability and with strong parameters.

Lemma 6.4 (See, e.g., [CdRN+25], Lemma 2.7)**.** *For $n, d \in \mathbb{N}^+$ and $\varepsilon, \delta \in \mathbb{R}^+$ satisfying $\varepsilon\delta = \omega(1/\log n)$ and $d^2 \leq \varepsilon\delta \log n$, it holds with high probability that a graph G sampled from $\mathbb{G}_{n,d}$ is (s, ε)-sparse for $s = (8d)^{-(1+\delta)(1+\varepsilon)/\varepsilon} n$.*

All not-too-large subgraphs in a sparse graph are 3-colorable, as shown in the next lemma. Since the proof of this folklore result is short and instructive, we include it with no claim of originality.

Lemma 6.5 (Folklore, see, e.g., [CdRN+25], Lemma 2.6)**.** *If a graph $G = (V, E)$ is (s, ε)-sparse for some $\varepsilon < 1/2$, then it holds for every subset $U \subseteq V$ of size at most s that $G[U]$ is 3-colorable.*

Proof. We proceed by induction on $|U|$. The base case $|U| = 1$ is immediate. For the induction step, assume that the claim holds for all sets of size at most $s - 1$. Consider a set $U \subseteq V$ of size $a \leq s$. The average degree in $G[U]$ is $2|E(U)|/a$, which is at most $2(1 + \varepsilon) < 3$ by sparsity. Hence, there exists a vertex $v \in U$ with degree at most 2 in $G[U]$. The graph $G[U \setminus \{v\}]$ is 3-colorable by the induction hypothesis, and every 3-coloring witnessing this will leave at least one color available for v. Hence, every 3-coloring of $G[U \setminus \{v\}]$ can be extended to $G[U]$. The lemma follows. □

The main technical insight needed to establish the satisfiability and reducibility conditions is an appropriate notion of *closed* set (cf. Section 4), for which we need the following definitions from [CdRN+25, Section 5]. Given a linear order $\prec$ on V, an *increasing (decreasing) path* in G is a simple path $(v_1, \ldots, v_\tau)$ where $v_i \prec v_{i+1}$ ($v_i \succ v_{i+1}$) 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 $U \subseteq V$, a vertex v is a *descendant of* U if it is a descendant of some vertex in U. We denote the

[9] The cited references establish much stronger concentration bounds than those stated in Lemma 6.2, but these bounds do not meaningfully improve the parameters in our main results.

set of all descendants of U by $\mathrm{Desc}(U)$ and define every vertex to be a descendant of itself so that $U \subseteq \mathrm{Desc}(U)$.

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 its two endpoints are both contained in U (in the case of cycles, the two endpoints coincide), while all other vertices are not in U. A *lasso with respect to U* is a walk $(v_1, v_2, v_3, v_4, v_5)$ where $v_2 = v_5$ and all other vertices are distinct, and where $v_1 \in U$.

Definition 6.6 (Closure)**.** Let $G = (V, E)$ be a graph and let $U \subseteq V$. We say that U is *closed* if $U = \mathrm{Desc}(U)$ 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.

It is shown in [CdRN+25, Proposition 5.2] that there is in fact a unique closure for every set $U \subseteq V$, which we denote by $\mathrm{cl}(U)$. The map $\mathrm{cl}(\cdot)\colon 2^V \to 2^V$ satisfies self-containment by definition, and monotonicity and idempotence are straightforward consequences of minimality. Hence, $\mathrm{cl}(\cdot)$ is a closure operator.

The following lemma shows that the closure of a small enough set $U \subseteq V$ is not much larger than U, given a bound on the sizes of descendant graphs.

Lemma 6.7 ([CdRN+25], 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 $U \subseteq Z \subseteq V$ such that $|Z| \leq 25\,|U|$ and $\mathrm{cl}(U) = \mathrm{Desc}(Z)$.*

Finally, the next lemma establishes the reducibility condition for (s, ε)-sparse graphs of bounded maximum degree. Recall that $\mathbf{x}_{V,[3]}$ denotes the variables in the encoding (3a)-(3d) of 3-colorability. Given a linear order $\prec_V$ on the vertices of a graph G, we say that an order $\prec$ on the monomials in $\mathbb{F}[\mathbf{x}_{V,[3]}]$ *respects* $\prec_V$ if for all $u, v \in V$ and all $i, j \in [3]$, it holds that $x_{u,i} \prec x_{v,j}$ whenever $u \prec_V v$.

Lemma 6.8 ([CdRN+25], Lemma 6.5)**.** *Let $G = (V, E)$ be a $(s, 1/3\Delta)$-sparse graph with a linear ordering $\prec_V$ on V. Let $\mathbb{F}$ be any field and consider an admissible order on $\mathbb{F}[\mathbf{x}_{V,[3]}]$ that respects $\prec_V$. For all sets $W \subseteq U$ satisfying that (1) W is closed, (2) $|U| \leq \ell$, and (3) every vertex in $N_{U \setminus W}(W)$ has degree at most Δ in $G[U \setminus W]$, it holds for every monomial m such that $\mathrm{Verts}_V(m) \subseteq W$ that m is reducible modulo $\langle U \rangle$ if and only if m is reducible modulo $\langle W \rangle$.*

We now have all components needed to prove Theorem 6.1.

Proof of Theorem 6.1. Let $G \sim \mathbb{G}_{n,d}$ and fix a field $\mathbb{F}$ of characteristic 0. By Lemmas 6.4 and 6.2, it holds with high probability that G is $(s, 1/3d)$-sparse for $s = (8d)^{-6d} \cdot n$ and that $\chi(G) \in [d/4 \log d, d/\log d]$. Given these properties, we establish the satisfiability and reducibility conditions for all monomials of degree at most D, where $D = s/25d^{\chi(G)} = d^{-C_1 d} \cdot n$, for some large enough constant C_1.

We first put a linear order $\prec_V$ on the vertices of G as follows. Let $\chi\colon V \to [\chi(G)]$ be a proper coloring of G. Order the vertices of G by color classes, so that $u \prec_V v$ whenever $\chi(u) < \chi(v)$ for all $u, v \in V$, and order the vertices in the same color class arbitrarily. Finally, let $\prec$ be the lexicographic order induced by any linear order on the variables of $\mathrm{Col}(G, 3)$ such that $x_{u,i} \prec x_{v,j}$ whenever $u \prec_V v$. Clearly, this lexicographic order respects $\prec_V$, and from now on the closure and reduction operators are defined with respect to this order.

Now we establish the satisfiability condition. Note that with the ordering $\prec_V$, every decreasing path in G has at most $\chi(G)$ vertices, and hence $|\mathrm{Desc}(U)| \leq d^{\chi(G)-1}\,|U|$ for all $U \subseteq V$. Therefore, since $\chi(G) d^{\chi(G)-1} < d^{\chi(G)}$ we have by Lemma 6.7 that for all sets $U \subseteq V$ of size at most $s/25d^{\chi(G)}$ that $|\mathrm{cl}(U)| \leq 25d^{\chi(G)}\,|U| \leq s$, and hence $G[\mathrm{cl}(U)]$ is 3-colorable by Lemma 6.5. The satisfiability condition is thus established for all $D \leq s/25d^{\chi(G)}$.

Finally, we turn to the reducibility condition. Let m be a monomial of degree at most D and let $m' \in \mathbb{F}[\mathbf{x}_{V,[3]}]$ satisfy that $\mathrm{cl}(m') \subseteq \mathrm{cl}(m)$. The argument in the previous paragraph shows that $|\mathrm{cl}(m)| \leq s$, and therefore the reducibility condition follows immediately from Lemma 6.8 with $W = \mathrm{cl}(m')$ and $U = \mathrm{cl}(m)$.

With the satisfiability and reducibility conditions in place, Lemma 4.4 implies that the cohomological k-consistency algorithm for 3-colorability accepts G for all $k \leq D/25d^{\chi(G)} \leq d^{-C_2 d} \cdot n$, for a large enough constant C_2. The proof is complete. □

7 Hardness of Lax and Null-Constraining CSPs

All the technical work of establishing the satisfiability and reducibility conditions for approximate graph coloring was already accomplished in [CdRN+25]. This section shows the techniques of Section 4 "in action" by providing a different proof of [CN25, Theorem 1.3]. There, Chan and Ng establish that no n-variable CSP whose template $\mathbf{T}$ satisfies two technical conditions they call *lax* and *null-constraining* is solved by the cohomological k-consistency algorithm for any $k \leq \Omega(n)$. Just as in Section 6, the fooling instances for $\mathrm{CSP}(\mathbf{T})$ are randomly sampled.

The main result of this section is Theorem 7.4. The overall proof structure is the same as in Section 6. We first recall a few definitions in order to introduce the CSP instances and closure in [CN25]. Then, we show that the closure of a small set is small (Corollary 7.13) and that the pseudo-reduction operator based on it is idempotent (Lemma 7.14). Finally, we establish the reducibility condition (Lemma 7.20).[10] Theorem 7.4 then follows via the connection to CSP hierarchies in Section 4. Our argument simplifies that in [CN25] by avoiding their use of *insular families* and *augmented closures* [CN25, Section 6], as well as their framework for combined hierarchies [CN25, Section 8].

In this section we consider only *t-uniform* relational structures, where all relations have arity t. For a t-uniform σ-structure $\mathbf{T}$, an instance $\mathbf{A}$ of $\mathrm{CSP}(\mathbf{T})$ induces a t-uniform hypergraph $H = (A, \mathcal{E})$, where $\mathcal{E} = \{\{a_1, \ldots, a_t\} : \exists \Gamma \in \sigma \mid (a_1, \ldots, a_t) \in \Gamma^{\mathbf{A}}\}$.

Definition 7.1 (Lax; [CN25], Definition 4.9)**.** A t-uniform σ-structure $\mathbf{T}$ is *lax* if for every symbol $\Gamma \in \sigma$ and every $j \in [t]$, there exists $\mathbf{b} \in T^{[t]\setminus\{j\}}$ such that $\{(b_1, \ldots, b_{j-1}, a, b_{j+1}, \ldots, b_t) \mid a \in T\} \subseteq \Gamma^{\mathbf{T}}$.

In other words, if $\mathrm{CSP}(\mathbf{T})$ is lax with input $\mathbf{A}$, then for every relation $\Gamma^{\mathbf{A}}$, every tuple $\mathbf{a} \in \Gamma^{\mathbf{A}}$, and every element $a_j \in \mathbf{a}$, there exists a map $\varphi\colon \mathrm{Vars}(\mathbf{a}) \setminus \{a_j\} \to T$ such that any extension of φ to $\mathrm{Vars}(\mathbf{a})$ maps $\mathbf{a}$ to some $\mathbf{i} \in \Gamma^{\mathbf{T}}$. Properties similar to but distinct from the lax condition, such as *robustness* [ABRW04] and *respectful boundary expansion* [MN24], have also been studied in proof complexity.

To introduce the null-constraining property, originally defined by Chan and Molloy [CM13], we need a few definitions. Given a t-uniform σ-structure $\mathbf{T}$, a *simple path instance of* $\mathrm{CSP}(\mathbf{T})$ (of length ℓ) is an instance $\mathbf{A}$ of $\mathrm{CSP}(\mathbf{T})$ such that the hypergraph of $\mathbf{A}$ is a simple path (of length ℓ). A simple path instance $\mathbf{A}$ of $\mathrm{CSP}(\mathbf{T})$ with endpoints $(a_1, a_2) \in A^2$ *permits* a pair $(i_1, i_2) \in T^2$ if there exists a homomorphism $\varphi\colon \mathbf{A} \to \mathbf{T}$ such that $\varphi(a_1) = i_1$ and $\varphi(a_2) = i_2$.

Definition 7.2 (Null-constraining; [CN25], Definition 4.2)**.** For $\ell \geq 1$ and a t-uniform σ-structure $\mathbf{T}$, we say that $\mathrm{CSP}(\mathbf{T})$ is *ℓ-null-constraining* if for every pair $(i_1, i_2) \in T^2$, it holds that every simple path instance of $\mathrm{CSP}(\mathbf{T})$ of length at least ℓ permits (i_1, i_2). We say that $\mathrm{CSP}(\mathbf{T})$ is *null-constraining* if it is ℓ-null-constraining for some ℓ. A structure $\mathbf{T}$ is null-constraining if $\mathrm{CSP}(\mathbf{T})$ is.

[10] For readers familiar with polynomial calculus, we note that the proof of the reducibility condition here is based on *variable restrictions*, which is simpler than its counterpart in Section 6 and similar to that in [MN24].

Note that graph coloring with at least 3 colors is 2-null-constraining but not lax. An example of a non-trivial lax and null-constraining CSP is t-uniform hypergraph coloring for $t \geq 3$ and at least 2 colors. For more examples, see [CN25, Section 4.3].

A *t-uniform random hypergraph* $H = (A, \mathcal{E})$ with n vertices and m edges is one chosen uniformly at random from the set of t-uniform n-vertex hypergraphs with m edges. Given a t-uniform σ-structure $\mathbf{T}$ that is lax and null-constraining, the fooling instances for CSP($\mathbf{T}$) are σ-structures $\mathbf{A}$ sampled according to the following distribution, which we denote by $\mathbb{A}_{n,m,\mathbf{T}}$. First, pick a t-uniform random n-vertex hypergraph with m edges. Then, for each $e \in \mathcal{E}$, pick a permutation $\pi_e = (v_1, \ldots, v_t)$ of the vertices in e and a relation symbol $\Gamma_e \in \sigma$, both uniformly at random. Finally, let $\mathbf{A}$ be the σ-structure over A where, for each $\Gamma \in \sigma$, we define $\Gamma^{\mathbf{A}} = \{\pi_e : \Gamma_e = \Gamma\}$.

For a σ-structure $\mathbf{T}$, we say that CSP($\mathbf{T}$) is *trivially satisfiable* if there exists some $i \in T$ such that, for any instance $\mathbf{A}$ of CSP($\mathbf{T}$), the map $A \to T$ that sends all elements in A to i is a homomorphism $\mathbf{A} \to \mathbf{T}$.[11] A random instance of such a CSP is clearly always satisfiable, so to obtain a fooling instance we must assume that CSP($\mathbf{T}$) is not trivially satisfiable. This assumption turns out to also be sufficient to guarantee that randomly sampled instances with linearly many constraints in the number of variables are unsatisfiable with high probability.

Lemma 7.3 ([AD22], Lemma 3)**.** *If $\mathbf{T}$ is a t-uniform relational structure such that* CSP($\mathbf{T}$) *is not trivially satisfiable, then there exists $\Delta > 0$ depending only on $\mathbf{T}$ such that, with probability $1 - o_n(1)$, a structure $\mathbf{A} \sim \mathbb{A}_{n,\Delta n,\mathbf{T}}$ is not homomorphic to $\mathbf{T}$.*

In [AD22], the proof is for random instances whose hypergraphs are sampled from the *Erdős-Rényi* random model. By, e.g., [MS07, Lemma 25], the same result holds for our random hypergraph model.

We are now ready to state Theorem 7.4, which is the main result of this section.

Theorem 7.4. *Let $t \geq 2$, and let $\mathbf{T}$ be a t-uniform relational structure that is lax and ℓ-null-constraining, and such that* CSP($\mathbf{T}$) *is not trivially satisfiable. There exists $\Delta > 0$ depending only on $\mathbf{T}$, and parameters $\zeta = \zeta(t, \ell, \Delta) > 0$ and $\delta = \delta(t, \ell, \Delta) > 0$ such that an instance $\mathbf{A} \sim \mathbb{A}_{n,\Delta n,\mathbf{T}}$ of* CSP($\mathbf{T}$) *is unsatisfiable with high probability as $n \to \infty$, but with probability $\delta - o_n(1)$ is accepted by the cohomological k-consistency algorithm for all $k \leq \zeta n$. Consequently, the cohomological k-consistency algorithm does not solve* CSP($\mathbf{T}$) *for any $k \leq \zeta n$.*

The proof of Theorem 7.4 goes through Lemma 4.2, and the proof of the reducibility condition (Lemma 7.20) applies to any field, not only those of characteristic zero. Hence, we also obtain the same lower bound of ζn as in Theorem 7.4 for the polynomial calculus degree required to refute $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ over any field under the same assumptions on $\mathbf{A}$ and $\mathbf{T}$.

Corollary 7.5. *For $\mathbf{A}, \mathbf{T}, \zeta, \delta$ as in Theorem 7.4, it holds with probability $\delta - o_n(1)$ that polynomial calculus requires degree at least ζn to refute $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$ over any field.*

To the best of our knowledge, Corollary 7.5 is new and in particular does not follow from [CN25, Theorem 1.3].

The rest of this section is devoted to the proof of Theorem 7.4. To begin with, we recall several technical properties of hypergraphs from [CN25]. Just as in Section 6, our results rely on a notion of *sparsity*.

Definition 7.6 (Sparsity)**.** A t-uniform hyperedge set $\mathcal{E}$ over a set V is *ε-sparse* if $|\mathcal{E}| \leq \frac{1+\varepsilon}{t-1} |V(\mathcal{E})|$ and is *(s, ε)-sparse* if all s-small subsets of $\mathcal{E}$ are ε-sparse. A hypergraph is (s, ε)-sparse if its hyperedge set is.

[11] Such CSPs are called *reflexive* in [AD22].

With constant probability as $n \to \infty$, random hypergraphs are sparse and have large girth, as stated in the next lemma. Both properties are standard. For the claim regarding sparsity, see, e.g., [MS07, Lemma 10]. For the lower bound on girth see, e.g., [JŁR00, Theorem 3.19]. The calculations there are for graphs, but generalize readily.

Lemma 7.7 (Properties of random hypergraphs)**.** *Let H be a random t-uniform hypergraph with n vertices and Δn edges. For any $\varepsilon > 0$, there exists $\mu = \mu(\Delta, t, \varepsilon)$ such that H is $(\mu n, \varepsilon)$-sparse with high probability as $n \to \infty$. Moreover, for all $\ell \in \mathbb{N}^+$ there exists $\delta = \delta(\Delta, \ell, t) > 0$ such that, with probability $\delta - o_n(1)$, the hypergraph H has no Berge cycle of length at most ℓ.*

To define closure we need the following notions, which are rephrasings of [CN25, Definition 5.20] in light of [CN25, Remark 5.16].

Definition 7.8 ((U,ℓ)-boundary, (U,ℓ)-bad)**.** Given $\ell \in \mathbb{N}^+$, a hypergraph $H = (V, \mathcal{E})$, subset $U \subseteq V$, and hyperedge subset $\mathcal{F} \subseteq \mathcal{E}$, let $B_1(U, \mathcal{F})$ be the set of edges in $\mathcal{F}$ that contain at most one vertex u such that $u \in U$ or $\deg_{\mathcal{F}}(u) > 1$, and let $B_2^\ell(U, \mathcal{F})$ be the set of all pendant paths (each viewed as a set of hyperedges) of length ℓ in $\mathcal{F} \setminus \mathcal{E}[U]$ whose non-endpoints are all outside U. The (U,ℓ)-*boundary* of $\mathcal{F}$ is $B_1(U,\mathcal{F}) \cup B_2^\ell(U,\mathcal{F})$ and is denoted $B^\ell(U,\mathcal{F})$. A set $\mathcal{F} \subseteq \mathcal{E}$ such that $B^\ell(U,\mathcal{F}) = \varnothing$ is called (U,ℓ)-*bad*.[12]

Definition 7.9 (Expanding; [CN25], Definition 5.27)**.** A t-uniform hyperedge set $\mathcal{E}$ over a set V is (ℓ, s, γ)-*expanding* if every s-small $\mathcal{F} \subseteq \mathcal{E}$ satisfies $|B^\ell(\varnothing, \mathcal{F})| \geq \gamma |\mathcal{E}|$. A hypergraph is (ℓ, s, γ)-expanding if its hyperedge set is.

If a hypergraph H is sparse and has no short Berge cycles, as is the case for a random hypergraph by Lemma 7.7, then H is expanding by the next lemma.

Lemma 7.10 ([CN25], Lemma 5.28)**.** *Let $\ell \geq 2$ and let $H = (V, \mathcal{E})$ be a t-uniform hypergraph. There exists $\varepsilon = \varepsilon(t, \ell)$ such that the following holds. If $H = (V, \mathcal{E})$ is (s, ε)-sparse and $\operatorname{girth}(\mathcal{E}) > \ell$, then H is (ℓ, s, γ)-expanding for $\gamma = \frac{t-1}{72\ell^2 t^3(1+\varepsilon)}$, provided that $\gamma \leq 1$.*

We collect some properties of bad sets in the next lemma. For their (short) proofs, see [CN25, Lemmas 5.3, 5.17, and 5.21].

Lemma 7.11. *For any hypergraph $H = (V, \mathcal{E})$, vertex sets $U \subseteq W \subseteq V$, and $\ell \in \mathbb{N}^+$, the following properties of (U,ℓ)-bad sets hold.*

1. *Any union of (U,ℓ)-bad hyperedge sets is (U,ℓ)-bad.*
2. *If a hyperedge set $\mathcal{F} \subseteq \mathcal{E}$ is (U,ℓ)-bad, then $\mathcal{F}$ is (W,ℓ)-bad.*
3. *Chaining property: Given hyperedge sets $\mathcal{F}, \mathcal{F}' \subseteq \mathcal{E}$, the set $\mathcal{F}'$ is $(U \cup V(\mathcal{F}), \ell)$-bad if and only if $\mathcal{F}' \cup \mathcal{F}$ is (U,ℓ)-bad.*

To define closure, Chan and Ng proceed differently from in Section 6: building on the work of [KMOW17], given a hypergraph $H = (V, \mathcal{E})$ they define the *s-local closure*[13] of U as

$$\operatorname{cl}_\ell^s(U) = U \cup \bigcup_{\substack{\mathcal{F} \subseteq \mathcal{E}:\ |\mathcal{F}| \leq s \\ \mathcal{F} \text{ is } (U,\ell)\text{-bad}}} V(\mathcal{F}). \tag{13}$$

Readers familiar with proof complexity may note the similarity between (13) and the notion of *support* [AR03, Fil19, MN24] defined using *unique neighbor expansion*.

By definition, the s-local closure is unique, contains U, and satisfies monotonicity by Lemma 7.11. The next two lemmas show that $\operatorname{cl}_\ell^s(U)$ is small when U is small, and that taking local closure is an idempotent operation.

[12] Chan and Ng call such sets *closed* [CN25, Remark 5.16], which unfortunately clashes with our usage of that word.

[13] This construction is called the *BW closure* in [CN25].

Lemma 7.12 ([CN25], Lemma 5.33)**.** *Let $H = (V, \mathcal{E})$ be a hypergraph, let $U \subseteq V$, and let $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2 \subseteq E$ be the union of two s-small (U, ℓ)-bad hyperedge sets $\mathcal{F}_1$ and $\mathcal{F}_2$. If $\mathcal{E}$ is $(\ell, 2s, \gamma)$-expanding and $|U| \leq r\gamma/\ell$ for some $r \geq 0$, then $\mathcal{F}$ is r-small.*

Corollary 7.13 (Closure size lemma, cf. [CN25], Lemma 6.11)**.** *If a hypergraph $H = (V, \mathcal{E})$ is t-uniform and $(\ell, 2s, \gamma)$-expanding, then for all sets $U \subseteq V$ such that $|U| \leq r\gamma/\ell$ it holds that $|\mathrm{cl}^s_\ell(U)| \leq |U|\,(1 + t\ell/\gamma)$.*

Proof. Since $\mathrm{cl}^s_\ell(U) = U \cup \bigcup\{V(\mathcal{F}) : \mathcal{F} \subseteq \mathcal{E}, |\mathcal{F}| \leq s, \mathcal{F}\ (U, \ell)\text{-bad in } \mathcal{E}\}$ and the union of two (U, ℓ)-bad sets is again (U, ℓ)-bad by Lemma 7.11, the claim follows readily from induction on the number of (U, ℓ)-bad sets coupled with Lemma 7.12. □

Lemma 7.14 (Small-set idempotence of $\mathrm{cl}^s_\ell(\cdot)$)**.** *If a hypergraph $H = (V, \mathcal{E})$ is t-uniform and $(\ell, 2s, \gamma)$-expanding, then for all sets $U \subseteq V$ such that $|\mathrm{cl}^s_\ell(U)| \leq s\gamma/\ell$, it holds that $\mathrm{cl}^s_\ell(\mathrm{cl}^s_\ell(U)) = \mathrm{cl}^s_\ell(U)$.*

Proof. Clearly $\mathrm{cl}^s_\ell(U) \subseteq \mathrm{cl}^s_\ell(\mathrm{cl}^s_\ell(U))$, so our goal is to establish the converse inclusion. Denote $\mathrm{cl}^s_\ell(U)$ by T and let $\mathcal{X}$ denote the union of s-small (U, ℓ)-bad sets. Note that $\mathcal{X}$ is (U, ℓ)-bad by Lemma 7.11 and hence (T, ℓ)-bad, again by Lemma 7.11. Moreover, $\mathcal{X}$ is s-small by Corollary 7.13 and Lemma 7.12.

Consider any s-small (T, ℓ)-bad set $\mathcal{F} \subseteq \mathcal{E}$. Since T is $s\gamma/\ell$-small, the set $\mathcal{F} \cup \mathcal{X}$ is s-small by Lemma 7.12 and is (U, ℓ)-bad by the chaining property in Lemma 7.11. Thus, $\mathcal{F} \subseteq \mathcal{F} \cup \mathcal{X} \subseteq \mathcal{X}$ by the definition of $\mathcal{X}$. It follows that the union $\mathcal{Y}$ of s-small (T, ℓ)-bad sets is contained in $\mathcal{X}$. Since $T = U \cup V(\mathcal{X})$, we thus have that $\mathrm{cl}^s_\ell(\mathrm{cl}^s_\ell(U)) = (T \cup V(\mathcal{Y})) \subseteq (U \cup V(\mathcal{X})) = \mathrm{cl}^s_\ell(U)$ as desired. □

The final ingredient needed to establish the reducibility condition is *insularity* [CN25, Section 6]. Intuitively, a vertex set U in a hypergraph $H = (V, \mathcal{E})$ is insular with respect to an edge set $\mathcal{F}$ if $\mathcal{F}$ has no additional (U, ℓ)-bad edge subsets beyond those already present in $\mathcal{E}[U]$.

Definition 7.15 (Insular, cf. [CN25], Definition 6.1 and Section 7.1)**.** Let $H = (V, \mathcal{E})$ be a hypergraph, $U \subseteq V$, and $\mathcal{F} \subseteq \mathcal{E}$. We say that U is *ℓ-insular in $\mathcal{F}$* if all (U, ℓ)-bad hyperedge sets in $\mathcal{F}$ are contained in $\mathcal{E}[U]$.

Observation 7.16. *Let $H = (V, \mathcal{E})$ be a hypergraph, $U \subseteq V$, and $\mathcal{F} \subseteq \mathcal{E}$. If $\mathrm{cl}^s_\ell(U) = U$ and $\mathcal{F}$ is s-small, then U is ℓ-insular in $\mathcal{F}$.*

Proof. If $\mathcal{X} \subseteq \mathcal{F}$ is (U, ℓ)-bad, then since $\mathcal{F}$ is s-small, we have that $\mathcal{X} \subseteq \mathcal{E}[\mathrm{cl}^s_\ell(U)]$ by the definition of $\mathrm{cl}^s_\ell(\cdot)$. Since $\mathrm{cl}^s_\ell(U) = U$, it follows that $\mathcal{X} \subseteq \mathcal{E}[U]$. □

The next lemma is used to establish the reducibility condition, and states that if a set U is ℓ-insular in $\mathcal{F}$, then the neighborhood of U in $\mathcal{F}$ is $\Omega(\ell)$-insular in $\mathcal{F}$ given some technical assumptions on $\mathcal{F}$. Lemma 7.17 is proved in [CN25, Section 7.1].

Lemma 7.17. *Let $\ell \in \mathbb{N}^+$, let $H = (V, \mathcal{E})$ be a t-uniform hypergraph, and let $U \subseteq V$ and $\mathcal{F} \subseteq \mathcal{E}$. Denote $U \cup N_{\mathcal{F}}(U)$ by W. There exists $\varepsilon = \varepsilon(t, \ell) > 0$ such that the following holds. If $\mathcal{F}$ is ε-sparse and $\mathrm{girth}(\mathcal{F}) > 3\ell$, then if U is $(3\ell + 2)$-insular in $\mathcal{F}$, it holds that W is ℓ-insular in $\mathcal{F}$.*

The following lemma is [CN25, Lemma 6.19], which is implicit in [CM13, Lemma 5.8]. It is used both for the satisfiability and reducibility conditions.

Lemma 7.18 (Extension lemma)**.** *Let $\mathbf{T}$ be a uniform, ℓ-null-constraining relational structure, let $\mathbf{A}$ be an instance of $\mathrm{CSP}(\mathbf{T})$, and let $H = (A, \mathcal{E})$ be the hypergraph of $\mathbf{A}$. For any set $U \subseteq A$, hyperedge set $\mathcal{F} \subseteq \mathcal{E}$, and hypergraph $h = (V(h), \mathcal{E}(h)) \in B^\ell(U, \mathcal{F})$, any partial homomorphism $\rho\colon \mathbf{A} \to \mathbf{T}$ with domain $U \cup V(\mathcal{F} \setminus \mathcal{E}(h))$ can be extended to a partial homomorphism $\rho'\colon \mathbf{A} \to \mathbf{T}$ with domain $U \cup V(\mathcal{F})$.*

Proof. If $h \in B_1(U, \mathcal{F})$, then h is a hyperedge that intersects $U \cup V(\mathcal{F} \setminus h)$ in at most one vertex, say u. Thinking of u as the endpoint of a simple path instance of length ℓ, the null-constraining property yields a partial homomorphism on h.

If $h \in B_2^\ell(U, \mathcal{F}) \setminus B_1(U, \mathcal{F})$, then h is a pendant path of length ℓ in $\mathcal{F}$ whose non-endpoints are outside $U \cup V(\mathcal{F} \setminus \mathcal{E}(h))$. By the null-constraining property, we obtain a partial homomorphism on h that agrees with ρ on $(U \cup V(\mathcal{F} \setminus \mathcal{E}(h))) \cap V(h)$.

In both cases, combining ρ with the map on h yields the desired ρ'. □

Corollary 7.19. *Let* $\mathbf{T}$ *be a uniform,* ℓ*-null-constraining relational structure, let* $\mathbf{A}$ *be an instance of* $\mathrm{CSP}(\mathbf{T})$*, and let* $H = (A, \mathcal{E})$ *be the hypergraph of* $\mathbf{A}$*. Let* $U \subseteq A$ *and* $\mathcal{F} \subseteq \mathcal{E}$*. If* U *is* ℓ*-insular in* $\mathcal{F}$*, then every partial homomorphism on* U *can be extended to* $U \cup V(\mathcal{F})$*.*

Proof. We proceed by induction on the size of $\mathcal{F}$. For the base case $|\mathcal{F}| = 0$ there is nothing to prove. For the induction step, suppose that the claim holds for all hyperedge sets of size at most $s-1$, and consider a set $\mathcal{F}$ of size s. Since U is ℓ-insular in $\mathcal{F}$, it holds that $B^\ell(U, \mathcal{F})$ is non-empty, and hence contains some hypergraph $h = (V_h, \mathcal{E}_h)$. Let $\mathcal{F}' = \mathcal{F} \setminus \mathcal{E}_h$. Then U is ℓ-insular in $\mathcal{F}'$. By the induction hypothesis, any partial homomorphism on U can be extended to $U \cup V(\mathcal{F}')$, and by Lemma 7.18 also to $U \cup V(\mathcal{F})$. □

Now we are finally in a position to establish the reducibility condition. The proof uses similar ideas as that of [CN25, Theorem 7.13].

Lemma 7.20 (Reducibility condition)**.** *Let* $t \geq 2$*,* $\ell \geq 1$*, and let* $\mathbf{T}$ *be a* t*-uniform, lax, and* ℓ*-null-constraining relational structure. Let* $\mathbf{A}$ *be an instance of* $\mathrm{CSP}(\mathbf{T})$ *and let* $H = (A, \mathcal{E})$ *be the hypergraph of* $\mathbf{A}$*. Suppose that* $\mathrm{girth}(\mathcal{E}) > 3\ell + 2$*, and let* $\varepsilon = \varepsilon(t, \ell)$ *be maximal such that Lemma 7.17 holds, let* s *be maximal such that* H *is* $(2s, \varepsilon)$*-sparse, let* $\ell' = 3\ell + 2$*, and let* $\gamma = \frac{t-1}{72\ell'^2 t^3(1+\varepsilon)}$*.*

Let $\mathbb{F}$ *be any field and consider any admissible order on the monomials in* $\mathbb{F}[\mathbf{x}_{A,T}]$*. For all monomials* $m \in \mathbb{F}[\mathbf{x}_{A,T}]$ *of degree at most* $D \leq s\gamma/(\ell'(1 + t\ell'/\gamma))$*, it holds for all* m' *such that* $\mathrm{cl}^s_{\ell'}(m') \subseteq \mathrm{cl}^s_{\ell'}(m)$ *that* m' *is reducible modulo* $\langle \mathrm{cl}^s_{\ell'}(m') \rangle$ *if and only if* m' *is reducible modulo* $\langle \mathrm{cl}^s_{\ell'}(m) \rangle$*.*

Proof. For brevity, denote $\mathrm{cl}^s_{\ell'}(m')$ by U and $\mathrm{cl}^s_{\ell'}(m)$ by W. Since $U \subseteq W$, if m' is reducible modulo $\langle U \rangle$ then it is reducible modulo $\langle W \rangle$. For the converse direction, suppose that m' is reducible modulo $\langle W \rangle$, which is to say that there exists a polynomial $p \in \langle W \rangle$ such that m' is the monomial in the leading term of p. The idea is to apply a variable restriction ρ such that $p \restriction_\rho$ is a polynomial in the ideal $\langle U \rangle$ with m' as the leading monomial. To do so, we first establish some technical properties of U and W.

By Lemma 7.10, the hypergraph H is $(\ell', 2s, \gamma)$-expanding. Therefore, given $|\mathrm{Verts}_V(m)| \leq D$ and $U \subseteq W$, we have by Corollary 7.13 that $|U| \leq s\gamma/\ell'$, and hence $\mathrm{cl}^s_{\ell'}(U) = U$ by Lemma 7.14. Moreover, the set $\mathcal{E}[W]$ is s-small by sparsity and Lemma 7.12, so by Observation 7.16 we have that U is ℓ'-insular in $\mathcal{E}[W]$, and hence $U \cup N_{\mathcal{E}[W]}(U)$ is ℓ-insular in $\mathcal{E}[W]$ by Lemma 7.17.

Consider the set of edges in $\mathcal{E}[W]$ that nontrivially intersect U, denoted $\mathcal{F} = \{e \in \mathcal{E}[W] \setminus \mathcal{E}[U] : e \cap U \neq \varnothing\}$. Each edge $e \in \mathcal{F}$ intersects U in exactly one vertex, since otherwise $\{e\}$ is (U, ℓ')-bad. Moreover, there is no 2- or 3-path, nor 2- or 3-cycle, in $\mathcal{E}[W] \setminus \mathcal{E}[U]$ whose endpoints are in U, as such a set of edges would be (U, ℓ')-bad given $\ell' > 3$. It follows that the edges in $\mathcal{F}$ are pairwise vertex disjoint outside of U and that each edge in $\mathcal{E}[W] \setminus \mathcal{E}[U]$ intersects at most one edge in $\mathcal{F}$. Moreover, since $\mathcal{E}$ is t-uniform, the hyperedge set $\mathcal{E}[N_{\mathcal{E}[W]}(U)]$ must be empty, as otherwise there would be an edge in $\mathcal{E}[W] \setminus \mathcal{E}[U]$ intersecting at least two edges in $\mathcal{F}$.

We are ready to define ρ. For each edge $e \in \mathcal{F}$, the lax property yields a partial homomorphism on $V(e) \setminus U$ that satisfies e. By the above, the union of these homomorphisms is a well-defined partial homomorphism τ on all of $\mathcal{F}$, whose domain is $N_{\mathcal{E}[W]}(U)$. We extend τ to all of $\mathcal{E}[W] \setminus (\mathcal{F} \cup \mathcal{E}[U])$ as follows. Since any $(N_{\mathcal{E}[W]}(U), \ell)$-bad subset of $\mathcal{E}[W] \setminus (\mathcal{F} \cup \mathcal{E}[U])$ is also

$(U \cup N_{\mathcal{E}[W]}(U), \ell)$-bad (Lemma 7.11), by the ℓ-insularity of $U \cup N_{\mathcal{E}[W]}(U)$ in $\mathcal{E}[W]$ (recorded two paragraphs above), the vertices of edges in this bad set are contained in $N_{\mathcal{E}[W]}(U)$, but such edges cannot exist as noted at the end of the previous paragraph. This means $N_{\mathcal{E}[W]}(U)$ is ℓ-insular in $\mathcal{E}[W] \setminus (\mathcal{F} \cup \mathcal{E}[U])$. Therefore, we can extend τ to vertices in $\mathcal{E}[W] \setminus \mathcal{E}[U]$ satisfying all of $\mathcal{E}[W] \setminus \mathcal{E}[U]$ by Corollary 7.19; this assignment—or more precisely, its Boolean encoding as a partial assignment to the variables in $\mathbf{x}_{A,T}$—is our ρ. Note that the domain of ρ is disjoint from $\mathbf{x}_{U,T}$.

It remains to show two facts: (1) $p \restriction_\rho \in \langle U \rangle$, and (2) m' is its leading monomial. Fact (1) holds since $p \in \langle W \rangle$ and ρ leaves the generating polynomials of $\langle U \rangle$ untouched while setting the remaining generating polynomials of $\langle W \rangle$ to 0. To see fact (2), we notice that $m' \restriction_\rho = m'$ (since $\mathrm{Verts}_A(m') \subseteq U$ and ρ assigns no variable over U), and for the remaining monomials in p, applying variable restrictions never increases their order in an admissible order. □

We now have all components needed to prove Theorem 7.4.

Proof of Theorem 7.4. The existence of $\Delta > 0$ depending only on $\mathbf{T}$ such that an instance $\mathbf{A}$ sampled from $\mathbb{A}_{n,\Delta n,\mathbf{T}}$ is unsatisfiable with high probability as $n \to \infty$ follows from Lemma 7.3.

Fix $\mathbf{A} \sim \mathbb{A}_{n,\Delta n,\mathbf{T}}$, a field $\mathbb{F}$ of characteristic 0, and any lexicographic order on the monomials of $\mathbb{F}[\mathbf{x}_{A,T}]$. Let $H = (A, \mathcal{E})$ be the hypergraph of $\mathbf{A}$. Let $\varepsilon = \varepsilon(t, \ell)$ be maximal such that Lemma 7.17 holds whenever $\mathrm{girth}\, \mathcal{E} > 3\ell + 2$, let s be maximal such that H is $(2s, \varepsilon)$-sparse, let $\ell' = 3\ell + 2$, and let $\gamma = \frac{t-1}{72\ell'^2 t^3(1+\varepsilon)}$.

By Lemma 7.7, there exists $\delta = \delta(\Delta, t, \ell) > 0$ such that, with probability $\delta - o_n(1)$, it holds that $\mathrm{girth}(H) > 3\ell + 2$ and that there exists $\mu = \mu(\Delta, t, \varepsilon) > 0$ such that $s \geq \mu n$. Then H is $(\ell', 2s, \gamma)$-expanding by Lemma 7.10. Assuming these properties hold, we show that $\mathrm{cl}^s_{\ell'}(\cdot)$ is a closure operator satisfying the satisfiability and reducibility conditions for all monomials $m \in \mathbb{F}[\mathbf{x}_{A,T}]$ of degree at most $D \leq s\gamma/(\ell'(1 + t\ell'/\gamma))$.

Let m be a monomial of degree at most D. For all D-small sets $U \subseteq A$, we have by Lemma 6.7 that $|\mathrm{cl}^s_{\ell'}(U)| \leq s\gamma/\ell'$. Hence, by Lemma 7.14, the operator $\mathrm{cl}^s_{\ell'}(\cdot)$ is idempotent on all D-small vertex sets, and is therefore a size-D closure operator.

The reducibility condition for $\mathrm{cl}^s_{\ell'}(\cdot)$ on monomials of degree at most D is precisely Lemma 7.20, so it remains to establish the satisfiability condition. Let m be a monomial of degree D. Since $|\mathrm{Verts}_V(m)| \leq D$, we have by Corollary 7.13 that $|\mathrm{cl}^s_{\ell'}(m)| \leq s\gamma/\ell'$. Moreover, the set $\mathcal{E}[\mathrm{cl}^s_{\ell'}(m)]$ is s-small by sparsity and Lemma 7.12. Since H is $(\ell', 2s, \gamma)$-expanding, it follows by Corollary 7.19 (instantiated with $U = \varnothing$, which is ℓ'-insular in $\mathcal{E}[\mathrm{cl}^s_{\ell'}(m)]$ by expansion) that there exists a satisfying assignment to $\mathcal{E}[\mathrm{cl}^s_{\ell'}(m)]$.

With the satisfiability and reducibility conditions established for all monomials of degree at most D, it follows from Corollary 7.13 and Lemma 4.4 that the cohomological k-consistency algorithm for CSP($\mathbf{T}$) accepts $\mathbf{A}$ for all $k \leq D/(1 + t\ell'/\gamma) \leq s\gamma/(\ell'(1 + t\ell'/\gamma)^2) = \zeta n$, for some $\zeta = \zeta(\Delta, t, \ell)$. The proof is complete. □

8 Concluding Remarks

We show that pseudo-reduction operators from proof complexity that are constructed via the Alekhnovich-Razborov method can be used to prove level lower bounds for (P)CSP hierarchies. Using this connection, we prove optimal level lower bounds for the cohomological k-consistency algorithm for approximate graph coloring, and give a new proof of the level lower bounds in [CN25] for lax and null-constraining CSPs.

Perhaps the most pressing open problem is whether the cohomological k-consistency algorithm solves all tractable CSPs. A closely related question is whether technical connections from proof complexity, such as the one presented in this work, can be useful to prove hierarchy lower bounds for tractable CSPs, and whether they can be used to extend our lower bounds to other hierarchies such as the C(BLP+AIP) hierarchy [CŽ23, CN25].

It is also natural to ask whether our results for approximate graph coloring can be generalized to *approximate graph homomorphism*, which is null-constraining but not lax. Lower bounds for this problem have been established for multiple hierarchies, the strongest of which is SDA [CŽ24b], but bounds with parameters comparable to those in Theorem 6.1 are only known for the k-consistency algorithm [CŽ24a] (see also [CN25, Theorem 6.24]).

Unlike our results for lax and null-constraining CSPs, it seems unclear how to prove our results for approximate graph coloring without using the connection to polynomial calculus degree lower bounds. In particular, the definition of closure there depends on the choice of a vertex order, which in turn seems motivated by the reduction operator. It would be interesting to see a proof of Theorem 6.1 that does not use pseudo-reduction operators; such a proof would likely also yield insights applicable to polynomial calculus degree lower bounds.

More technically, the use of the lex game in Section 5 limits the applicability of our results to lexicographic orders on the monomials. Can we show similar results for, say, the graded lexicographic order? This would likely require different techniques from those in this paper, or might simply be impossible; see Remark 4.6.

Acknowledgements

We would like to thank Siu-On Chan, Nicola Galesi, Tamio-Vesa Nakajima, Jakub Opršal, Per Austrin, Kilian Risse, and Benedikt Pago for insightful discussions. The initial stages of this work were carried out when the authors attended the 2025 *Complexity as a Kaleidoscope* research school at CIRM, Marseille, France. We thank the organizers for creating such an inspiring environment, and Noel Arteche, David Engström, Stefan Grosser, and Santiago Guzmán Pro for helpful conversations while there. We are also grateful to the anonymous *SODA* reviewers for their helpful comments, all of which improved the exposition of the paper.

JC was partially supported by the *Wallenberg AI and Autonomous Systems Program (WASP)* and would also like to thank Susanna de Rezende for encouraging him to apply for funding to attend the CIRM research school. YG received funding from the Independent Research Fund Denmark grant 9040-00389B.

A Proof of Lemma 4.2

In this appendix we prove Lemma 4.2, restated below for convenience. As noted in Section 4, the argument is standard and hence we make no claim of originality. We follow the presentation of Conneryd, de Rezende, Nordström, Pang, and Risse [CdRN+25, Lemma 4.6].

Lemma 4.2 (Restated)**.** *Let $D \in \mathbb{N}$, let σ be a signature whose relation symbols all have arity at most D, let $\mathbf{A}$ and $\mathbf{T}$ be σ-structures, and let $\mathrm{cl}\colon 2^A \to 2^A$ be a size-D closure operator for A. Finally, let $\prec$ be an admissible ordering of the monomials in $\mathbb{F}[\mathbf{x}_{A,T}]$. Suppose that the following conditions hold for every monomial $m \in \mathbb{F}[\mathbf{x}_{A,T}]$ of degree at most D:*

> Satisfiability: *It holds that $\mathbf{A}[\mathrm{cl}(m)] \to \mathbf{T}$.*
>
> Reducibility: *For every monomial m' such that $\mathrm{cl}(m') \subseteq \mathrm{cl}(m)$, it holds that m' is reducible modulo $\langle \mathrm{cl}(m') \rangle$ if and only if m' is reducible modulo $\langle \mathrm{cl}(m) \rangle$.*

Then, the operator R on $\mathbb{F}[\mathbf{x}_{A,T}]$ defined to be the $\mathbb{F}$-linear extension of the map that takes each monomial $m \in \mathbb{F}[\mathbf{x}_{A,T}]$ to $R^{\prec}_{\langle \mathrm{cl}(m) \rangle}(m)$ is a degree-D pseudo-reduction operator for $\mathcal{P}_{\mathbf{A} \to \mathbf{T}}$.

Before commencing the proof of Lemma 4.2, we record a basic fact of reduction operators as well as two technical claims that both concern how closure operators interact with reduction modulo ideals.

Fact A.1. Let I and J be ideals in $\mathbb{F}[x_1, \ldots, x_n]$ such that $I \subseteq J$. For all monomials m and m', it holds that $R_J(m' R_I(m)) = R_J(m'm)$.

Proof. Write $m = q + r$ where $q \in I$ and r is the reduction of m modulo I. Then $m' R_I(m) = m' r$, and also $m' R_I(m) = q' + r'$ where $q' \in J$ and r' is the reduction of $m' R_I(m)$ modulo J. Observe that $R_J(m' R_I(m)) = r'$, so the claim follows if we show that $R_J(m'm) = r'$. Note that $m' R_I(m) = m'(m - q) = q' + r'$, and hence $m'm = q' + r' - m'q$. Since $I \subseteq J$ it holds that $m'q \in J$. Therefore, we can write $m'm = q' - m'q + r'$ where $q' - m'q \in J$ and r' is irreducible modulo J. By uniqueness of this representation, it follows that $R_J(m'm) = r'$. □

Claim A.2. *If the conditions of Lemma 4.2 are satisfied, then for every monomial m of degree at most D and every monomial m' that appears in $R^{\prec}_{\langle \mathrm{cl}(m) \rangle}(m)$, it holds that $\mathrm{cl}(m') \subseteq \mathrm{cl}(m)$.*

Proof. It suffices to show that $\mathrm{Verts}_A(m') \subseteq \mathrm{cl}(m)$, since monotonicity and idempotence of $\mathrm{cl}(\cdot)$ then imply that $\mathrm{cl}(m') \subseteq \mathrm{cl}(\mathrm{cl}(m)) = \mathrm{cl}(m)$. The claim trivially follows if $m' = m$, so suppose this is not the case. Toward a contradiction, suppose that $\mathrm{Verts}_A(m') \not\subseteq \mathrm{cl}(m)$. Denote the assignment that sets $x_{a,i}$ to 0 for all $i \in \mathbf{T}$ and all $a \notin \mathrm{cl}(m)$ by ρ. Denote the polynomial $R^{\prec}_{\langle \mathrm{cl}(m) \rangle}(m)$ by r, and recall that r is the unique sum of irreducible terms such that $m = r + q$, for some $q \in \langle \mathrm{cl}(m) \rangle$. No variable of m nor of any generator of $\langle \mathrm{cl}(m) \rangle$ is assigned by ρ, so $m \restriction_\rho = m$ and $q \restriction_\rho \in \langle \mathrm{cl}(m) \rangle$. However, we have that $m' \restriction_\rho = 0$, so $r \restriction_\rho \neq r$. Moreover, all terms in $r \restriction_\rho$ are irreducible modulo $\langle \mathrm{cl}(m) \rangle$, as otherwise r would contain a reducible term. But this contradicts that the representation $m = r + q$ is unique, and thus we conclude that $\mathrm{Verts}_A(m') \subseteq \mathrm{cl}(m)$. The claim follows. □

The next claim states that if the reducibility condition in Lemma 4.2 is satisfied, then reducing a monomial m modulo $\langle \mathrm{cl}(m) \rangle$ and reducing m modulo a larger ideal $\langle X \rangle$ generated by a set $X \supseteq \mathrm{cl}(m)$ results in the same polynomial. Thus, the subinstance $\mathbf{A}[\mathrm{cl}(m)]$ can be said to be the part of $\mathbf{A}$ that is "relevant" for reducing m.

Claim A.3. *Let $\mathbf{A}, \mathbf{T}, \prec$, and $\mathrm{cl}(\cdot)$ be defined as in Lemma 4.2, and let $X \subseteq A$. Suppose that for all monomials $m \in \mathbb{F}[\mathbf{x}_{A,T}]$ such that $\mathrm{cl}(m) \subseteq X$, it holds that m is irreducible modulo $\langle X \rangle$ if m is irreducible modulo $\langle \mathrm{cl}(m) \rangle$. Then, for all monomials m' such that $\mathrm{cl}(m') \subseteq X$, it holds that $R^{\prec}_{\langle \mathrm{cl}(m') \rangle}(m') = R^{\prec}_{\langle X \rangle}(m')$.*

Proof. Let m be a monomial such that $\mathrm{cl}(m) \subseteq X$, and write $R^{\prec}_{\langle \mathrm{cl}(m) \rangle}(m) = \sum_j a_j m_j$. Recall that m_j is irreducible modulo $\langle \mathrm{cl}(m) \rangle$ for all j by definition. Moreover, we have by Claim A.2 that $\mathrm{cl}(m_j) \subseteq \mathrm{cl}(m)$, and each m_j is also irreducible modulo $\langle \mathrm{cl}(m_j) \rangle$. Therefore, it follows by assumption that each m_j is irreducible modulo $\langle X \rangle$. By definition, the polynomial $R^{\prec}_{\langle \mathrm{cl}(m) \rangle}(m)$ is the unique sum of irreducible terms modulo $\langle \mathrm{cl}(m) \rangle$ such that $m = R^{\prec}_{\langle \mathrm{cl}(m) \rangle}(m) + q$, for some $q \in \langle \mathrm{cl}(m) \rangle$. Note that $\mathrm{cl}(m) \subseteq \langle X \rangle$ and hence $q \in \langle X \rangle$. Since all terms in $R^{\prec}_{\langle \mathrm{cl}(m) \rangle}(m)$ are also irreducible modulo $\langle X \rangle$, it follows by uniqueness that $R^{\prec}_{\langle \mathrm{cl}(m) \rangle}(m) = R^{\prec}_{\langle X \rangle}(m)$. □

With Claims A.2 and A.3 established, we are ready to prove Lemma 4.2.

Proof of Lemma 4.2. Let R be the operator in the lemma statement. We need to show that R satisfies Properties 1-3 of a pseudo-reduction operator, that is, that (1) $R(1) = 1$, (2) $R(p) = 0$ for every polynomial $p \in \mathcal{P}_{\mathbf{A} \to \mathbf{T}}$, and (3) $R(x_{a,i} \cdot m) = R\big(x_{a,i} \cdot R(m)\big)$ for every monomial $m \in \mathbb{F}[\mathbf{x}_{A,T}]$ of degree at most $D - 1$ and every variable $x_{a,i}$.

To see that $R(1) = 1$, note that $\mathbf{A}[\mathrm{cl}(1)] \to \mathbf{T}$ by the satisfiability condition, and hence $1 \notin \langle \mathrm{cl}(1) \rangle$ by the Boolean Nullstellensatz (Lemma 2.1). Since 1 is minimal in $\prec$, it follows that $R(1) = 1$.

To establish Property 2, observe that $\mathrm{cl}(m)$ contains all vertices in A mentioned by m, which implies that R maps all axioms (3b)–(3d) to 0. To see that R also maps the axioms (3a) to 0, further note that $\mathrm{cl}(1) \subseteq \mathrm{cl}(\{a\})$ for all $a \in A$, and hence

$$R\Big(\sum_{i\in T} x_{a,i} - 1\Big) = \sum_{i\in T} R^{\prec}_{\langle \mathrm{cl}(x_{a,i})\rangle}(x_{a,i}) - R^{\prec}_{\langle \mathrm{cl}(1)\rangle}(1) \qquad \text{[definition of } R\text{]} \tag{14a}$$

$$= \sum_{i\in T} R^{\prec}_{\langle \mathrm{cl}(x_{a,i})\rangle}(x_{a,i}) - R^{\prec}_{\langle \mathrm{cl}(x_{a,i})\rangle}(1) \qquad \text{[reducibility and Claim A.3]} \tag{14b}$$

$$= 0 \ , \qquad \text{[definition of } R^{\prec}_{\langle \mathrm{cl}(x_{a,i})\rangle}\text{]} \tag{14c}$$

as desired.

Let us now establish Property 3, which is subtler than the first two. Let m be a monomial of degree at most $D-1$, and write the polynomial $R(m)$ as $\sum_j a_j m_j$ where $a_j \in \mathbb{F}$. We have by definition that

$$R(x_{a,i}R(m)) = \sum_j a_j R(x_{a,i}m_j) = \sum_j a_j R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m_j)\rangle}(x_{a,i}m_j) \ . \tag{15}$$

The next step is to show that $R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m_j)\rangle}(x_{a,i}m_j) = R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m)\rangle}(x_{a,i}m_j)$ for all j, for which we need Claim A.2. Indeed, note that

$$\mathrm{cl}(x_{a,i}m_j) \subseteq \mathrm{cl}(\{a\} \cup \mathrm{cl}(m_j)) \qquad \text{[monotonicity and self-containment]} \tag{16a}$$

$$\subseteq \mathrm{cl}(\{a\} \cup \mathrm{cl}(m)) \qquad \text{[monotonicity and Claim A.2]} \tag{16b}$$

$$\subseteq \mathrm{cl}(\mathrm{cl}(x_{a,i}m)) \qquad \text{[monotonicity and self-containment]} \tag{16c}$$

$$= \mathrm{cl}(x_{a,i}m) \ . \qquad \text{[idempotence]} \tag{16d}$$

By reducibility and Claim A.3, we obtain that $R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m_j)\rangle}(x_{a,i}m_j) = R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m)\rangle}(x_{a,i}m_j)$ for all j. Summing up, we have thus far established that

$$R(x_{a,i}R(m)) = \sum_j a_j R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m)\rangle}(x_{a,i}m_j) \ . \tag{17}$$

From (17), we can conclude that Property 3 holds. Indeed, we have

$$\sum_j a_j R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m)\rangle}(x_{a,i}m_j) = R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m)\rangle}\Big(x_{a,i}\sum_j a_j m_j\Big) \qquad \text{[linearity]} \tag{18a}$$

$$= R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m)\rangle}(x_{a,i}R^{\prec}_{\langle \mathrm{cl}(m)\rangle}(m)) \qquad \text{[definition of } \textstyle\sum_j a_j m_j\text{]} \tag{18b}$$

$$= R^{\prec}_{\langle \mathrm{cl}(x_{a,i}m)\rangle}(x_{a,i}m) \qquad \text{[monotonicity and Fact A.1]} \tag{18c}$$

$$= R(x_{a,i}m) \ , \qquad \text{[definition of } R\text{]} \tag{18d}$$

as desired. With Properties 1-3 established, we conclude that R is a degree-D pseudo-reduction operator for $\mathcal{P}_{\mathbf{A}\to\mathbf{T}}$. □

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*. 5, 10

[ABRW04] Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Pseudorandom generators in propositional proof complexity. *SIAM Journal on Computing*, 34(1):67–88, 2004. Preliminary version in *FOCS '00*. 21

[AD22] Albert Atserias and Víctor Dalmau. Promise constraint satisfaction and width. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '22)*, pages 1129–1153, 2022. 3, 22

[AGH17] Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2+\varepsilon)$-SAT is NP-hard. *SIAM Journal on Computing*, 46(5):1554–1573, 2017. 2

[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*. 4

[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*. 4, 6, 14, 15, 16, 23

[ARS02] R. P. Anstee, Lajos Rónyai, and Attila Sali. Shattering news. *Graphs and Combinatorics*, 18(1):59–73, 2002. 17

[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. 2

[BCIP02] Joshua Buresh-Oppenheim, Matthew Clegg, Russell Impagliazzo, and Toniann Pitassi. Homogenization and the polynomial calculus. *Computational Complexity*, 11(3-4):91–108, 2002. Preliminary version in *ICALP '00*. 5

[BCK15] Boaz Barak, Siu On Chan, and Pravesh K. Kothari. Sum of squares lower bounds from pairwise independence. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC '15)*, pages 97–106, 2015. 4

[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. 4

[BG15] Christoph Berkholz and Martin Grohe. Limitations of algebraic approaches to graph isomorphism testing. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP '15), Part I*, volume 7364 of *Lecture Notes in Computer Science*, pages 155–166. Springer, 2015. 4

[BG17] Christoph Berkholz and Martin Grohe. Linear diophantine equations, group CSPs, and graph isomorphism. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '17)*, pages 327–339, 2017. 2, 4

[BG19] Joshua Brakensiek and Venkatesan Guruswami. An algorithmic blend of LPs and ring equations for promise CSPs. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '19)*, pages 436–455, 2019. 2

[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*. 4, 5

[BGWŽ20] Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Živný. The power of the combined basic linear programming and affine relaxation for promise constraint satisfaction problems. *SIAM Journal on Computing*, 49(6):1232–1248, 2020. Preliminary version in *SODA '20*. 2

[BHK+19] Boaz Barak, Samuel Hopkins, Jonathan Kelner, Pravesh K. Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2):687–735, 2019. Preliminary version in *FOCS '16*. 4

[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*. 4, 5

[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. 9

[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*. 4

[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. Preliminary version in *APPROX-RANDOM '17*. 4, 5

[BKM25] Amey Bhangale, Subhash Khot, and Dor Minzer. On approximability of satisfiable k-CSPs: V. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC '25)*, pages 62–71, 2025. 3

[BKW17] Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and how to use them. In *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl*

Follow-Ups, pages 1–44. 2017. 1

[Bul17] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS '17)*, pages 319–330, 2017. 1

[CdRN+25] Jonas Conneryd, Susanna F. de Rezende, Jakob Nordström, Shuo Pang, and Kilian Risse. Graph colouring is hard on average for polynomial calculus and Nullstellensatz, 2025. arXiv:2503.17022. Preliminary version in *FOCS '23*. 4, 7, 9, 14, 15, 19, 20, 21, 27

[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. 5, 10

[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. 19

[CM13] Siu On Chan and Michael Molloy. A dichotomy theorem for the resolution complexity of random constraint satisfaction problems. *SIAM Journal on Computing*, 42(1):27–60, 2013. Preliminary version in *FOCS '08*. 4, 21, 24

[CN25] Siu On Chan and Hiu Tsun Ng. How random CSPs fool hierarchies: II. Technical Report TR25-026, Electronic Colloquium on Computational Complexity (ECCC), 2025. Preliminary version in *STOC '25*. 3, 4, 5, 7, 21, 22, 23, 24, 25, 26, 27

[CNP24] Siu On Chan, Hiu Tsun Ng, and Sijin Peng. How random CSPs fool hierarchies. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24)*, pages 1944–1955, 2024. 3, 4

[CS88] Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, October 1988. 4

[CŽ23] Lorenzo Ciardo and Stanislav Živný. CLAP: A new algorithm for promise CSPs. *SIAM Journal on Computing*, 52(1):1–37, 2023. Preliminary version in *SODA '22*. 3, 26

[CŽ24a] Lorenzo Ciardo and Stanislav Živný. The periodic structure of local consistency, 2024. arXiv:2406.19685. 3, 27

[CŽ24b] Lorenzo Ciardo and Stanislav Živný. Semidefinite programming and linear equations vs. homomorphism problems. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24)*, pages 1935–1943, june 2024. 27

[CŽ25a] Lorenzo Ciardo and Stanislav Živný. Approximate graph coloring and the crystal with a hollow shadow. *SIAM Journal on Computing*, 54(4):1138–1192, 2025. Combined full version of two papers published in *STOC '23* and *SODA '23*. 2, 3

[CŽ25b] Lorenzo Ciardo and Stanislav Živný. Semidefinite programming and linear equations vs. homomorphism problems. *SIAM Journal on Computing*, 54(3):545–584, 2025. Preliminary version in *STOC '24*. 3

[DO24] Víctor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems. In *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '24)*, 2024. 3, 11

[Fel07] Bálint Felszeghy. *Gröbner Theory of Zero Dimensional Ideals: With a View Towards Combinatorics*. PhD thesis, Budapest University of Technology and Economics, 2007. Available at `https://math.bme.hu/~fbalint/publ/phd_thesis.pdf`. 17

[Fil19] Yuval Filmus. Another look at degree lower bounds for polynomial calculus. *Theoretical Computer Science*, 796:286–293, 2019. 23

[FR09] Bálint Felszeghy and Lajos Rónyai. Some meeting points of Gröbner bases and combinatorics. In *Algorithmic Algebraic Combinatorics and Gröbner Bases*, pages 207–227. Springer, 2009. 17

[FRR06] Bálint Felszeghy, Balázs Ráth, and Lajos Rónyai. The lex game and some applications. *Journal of Symbolic Computation*, 41(6):663–681, 2006. 7, 17, 18

[FV98] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1998. Preliminary version in *STOC '93*. 1

[GJ76] Michael Garey and David S. Johnson. The complexity of near-optimal graph coloring. *Journal of the ACM*, 23(1):43–49, January 1976. 2

[GL10a] Nicola Galesi and Massimo Lauria. On the automatizability of polynomial calculus. *Theory of Computing Systems*, 47(2):491–506, August 2010. 4, 15

[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. 4, 15

[Gri01] Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259(1–2):613–622, May 2001. 4, 5

[HN90] Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48(1):92–110, 1990. 1

[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. 5, 10, 18

[Jea98] Peter Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200(1):185–204, 1998. 1

[JŁR00] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random Graphs*. John Wiley & Sons, Inc., February 2000. 23

[JPR+21] Chris Jones, Aaron Potechin, Goutham Rajendran, Madhur Tulsiani, and Jeff Xu. Sum-of-squares lower bounds for sparse independent set. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS '21)*, pages 406–416, 2021. 4

[KM21] Pravesh K. Kothari and Peter Manohar. A stress-free sum-of-squares lower bound for coloring. In *Proceedings of the 36th Annual Computational Complexity Conference (CCC '21)*, volume 200 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 23:1–23:21, July 2021. 4

[KMOW17] Pravesh K. Kothari, Ryuhei Mori, Ryan O'Donnell, and David Witmer. Sum of squares lower bounds for refuting any CSP, 2017. arXiv:1701.04521. Preliminary version in *STOC '17*. 4, 23

[KO22] Andrei Krokhin and Jakub Opršal. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News*, 9(3):30–59, August 2022. 1

[KOWŽ23] Andrei Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Živný. Topology and adjunction in promise constraint satisfaction. *SIAM Journal on Computing*, 52(1):38–79, 2023. 2, 3, 5

[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. 19

[KPX24] Pravesh K. Kothari, Aaron Potechin, and Jeff Xu. Sum-of-squares lower bounds for independent set on ultra-sparse random graphs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24)*, pages 192–1934, 2024. 4

[KTY24] Ken-ichi Kawarabayashi, Mikkel Thorup, and Hirotaka Yoneda. Better coloring of 3-colorable graphs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24)*, pages 331–339, 2024. 2

[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. 2

[Laz96] Felix Lazebnik. On systems of linear Diophantine equations. *Mathematics Magazine*, 69(4):261–266, 1996. 2

[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. 4

[LP25] Moritz Lichter and Benedikt Pago. Limitations of affine integer relaxations for solving constraint satisfaction problems. In *Proceedings of the 52nd International Colloquium on Automata, Languages, and Programming (ICALP '25)*, volume 334 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 166:1–166:17, 2025. 3, 4

[Mit02a] David G. Mitchell. *The Resolution Complexity of Constraint Satisfaction*. PhD thesis, University of Toronto, 2002. 4

[Mit02b] David G. Mitchell. Resolution complexity of random constraints. In *Principles and Practice of Constraint Programming (CP '02)*, pages 295–310. Springer Berlin Heidelberg, 2002. 4

[MM08] Miklós Maróti and Ralph McKenzie. Existence theorems for weakly symmetric operations. *Algebra Universalis*, 59(3-4):463–489, 2008. 1

[MN24] Mladen Mikša and Jakob Nordström. A generalized method for proving polynomial calculus degree lower bounds. *Journal of the ACM*, 71(6), November 2024. Preliminary version in *CCC '15*. 4, 14, 15, 21, 23

[MS07] Michael Molloy and Mohammad R. Salavatipour. The resolution complexity of random con-

straint satisfaction problems. *SIAM Journal on Computing*, 37(3):895–922, 2007. Preliminary version in *FOCS'03*. 4, 22, 23

[ÓC22] Adam Ó Conghaile. Cohomology in constraint satisfaction and structure isomorphism. In *Proceedings of the 47th International Symposium on Mathematical Foundations of Computer Science (MFCS '22)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 75:1–75:16, 2022. 3, 12

[O'D17] Ryan O'Donnell. SOS is not obviously automatizable, even approximately. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS '17)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 59:1–59:10, 2017. 2

[Pan21] Shuo Pang. SOS lower bound for exact planted clique. In *Proceedings of the 36th Annual Computational Complexity Conference (CCC '21)*, volume 200 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:63, 2021. 4

[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. 2

[PX25] Aaron Potechin and Jeff Xu. Sum-of-squares lower bounds for coloring random graphs. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC '25)*, pages 84–95, 2025. 4

[Raz98] Alexander A. Razborov. Lower bounds for the polynomial calculus. *Computational Complexity*, 7(4):291–324, December 1998. 5, 11, 14, 18

[RT24] Julian Romero and Levent Tunçel. Graphs with large girth and chromatic number are hard for Nullstellensatz. *SIAM Journal on Discrete Mathematics*, 38(3):2108–2131, 2024. 19

[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. 2

[Sch78] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC '78)*, pages 216–226, 1978. 1

[Sch08] Grant Schoenebeck. Linear level Lasserre lower bounds for certain k-CSPs. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '08)*, pages 593–602, October 2008. 4

[Tse68] Grigori Tseitin. The complexity of a deduction in the propositional predicate calculus. *Zapiski Nauchnyh Seminarov Leningradskogo Otdelenija matematicheskogo Instituta im. V. A. Steklova akademii Nauk SSSR (LOMI)*, 8:234–259, 1968. In Russian. 4

[Zhu20] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *Journal of the ACM*, 67(5), August 2020. Preliminary version in *FOCS '17*. 1