Electronic Colloquium on Computational Complexity, Report No. 26 (2025)

HOW RANDOM CSPS FOOL HIERARCHIES: II

SIU ON CHAN* AND HIU TSUN NG†

Abstract. Relaxations for the constraint satisfaction problem (CSP) include bounded width (BW), linear program (LP), semidefinite program (SDP), affine integer program (AIP), and their combinations. Tightening relaxations systematically leads to hierarchies and stronger algorithms. For LP+AIP and SDP+AIP hierarchies, various lower bounds were shown by Ciardo and Živný (STOC 2023, STOC 2024) and by Chan, Ng, and Peng (STOC 2024).

This paper continues on this line of work to show lower bounds for related hierarchies. We prove new lower bounds to BW and AIP hierarchies. We also show the first lower bounds to the cohomological consistency hierarchy of Ó Conghaile (MFCS 2022) and the C(BLP+AIP) hierarchy of Ciardo and Živný (SODA 2022). Our lower bounds are for linear level and optimal. They make partial progress towards an open question of Lichter and Pago (arXiv 2407.09097v1) concerning the power of these hierarchies.

We prove our results using new closure and boundary. We generalize closure and boundary to streamline proofs across hierarchies.

1. Introduction

Relaxations, such as linear program and affine integer program, are powerful tools for solving constraint satisfaction problems (CSP). A relaxation can be tightened systematically into a hierarchy based on variable subsets of size at most d in an instance. The level of the hierarchy is d. A number of hierarchies have been used for CSPs:

- Bounded width hierarchy (BW) in the local consistency algorithm.
- Sherali–Adams linear programming (LP) hierarchy.
- Lasserre semidefinite programming (SDP) hierarchy, which is the Lagrangian dual of the sum-of-squares SDP hierarchy.
- Affine integer programming (AIP) hierarchy.

For each hierarchy except SDP, the hierarchy of level d can be solved in time $n^{O(d)}$, where n is the number of variables in the instance; for SDP, the hierarchy is only known to be approximately solved in the same running time. For $d = o(n)$, this gives subexponential-time algorithms for solving CSPs. If a CSP cannot be solved by a hierarchy even up to level $\Omega(n)$, a class of subexponential-time algorithms is ruled out. It is of interest to study the power and limitations of each hierarchy for not just specific CSPs, but also for classes of CSPs.

Previous works have shown limitations of these hierarchies for general classes of CSPs. Both LP [BGMT12] and SDP hierarchies [BCK15, KMOW17] cannot solve pairwise uniform CSPs even at level $\Omega(n)$. There are also average-case upper bounds for LP and SDP hierarchies for predicate CSPs that are not pairwise uniform [AOW15]. A CSP is pairwise uniform if each constraint has a pairwise uniform distribution of satisfying assignments. Pairwise uniform CSPs were introduced by Austrin and Mossel [AM09] in the context of hardness of approximation. So there is a fairly good understanding of the power of LP and SDP hierarchies.

The situation is far from complete for the AIP hierarchy. Affine integer program was introduced by Barto, Bulín, Krokhin, and Opršal [BBKO21] in the context of promise CSP (PCSP), and was inspired by its precursor proposed by Brakensiek and Guruswami [BG19]. AIP is stronger

*sochan@gmail.com.
†htng0@link.cuhk.edu.hk.

ISSN 1433-8092

than Gaussian elimination and solves linear equations over finite fields. Interestingly, AIP also solves promise CSPs that are not linear equations, such as PCSP(1-in-3-SAT, 3NAE) [BG19]. Constant-level AIP hierarchy cannot solve approximate graph coloring [CŽ23a] or its generalization, approximate graph homomorphism [CŽ24a]. $\Omega(n)$-level AIP hierarchy cannot solve CSPs that are "pairwise neutral" for AIP (a condition inspired by "pairwise uniform"), of which a sufficient condition is that the satisfying assignments of every constraint contain a Hamming ball of radius 2 [CNP24].

Which (P)CSPs can be solved in subexponential time by the $o(n)$-level bounded width hierarchy? Approximate graph coloring cannot [AD22], and neither can approximate graph homomorphism [CŽ24b], nor CSPs pairwise neutral for BW [CNP24]. These are the only results formally proved so far, but in fact a more satisfactory answer is suggested by [CM13] in the context of resolution. Chan and Molloy [CM13] showed exponential resolution complexity lower bounds for CSPs that are "null-constraining", that is, CSPs whose sufficiently long simple constraint paths have no effect on their endpoints. Their exponential resolution lower bounds for such CSPs follow from linear lower bounds on the *resolution width*, as is common for resolution lower bounds [BSW01]. On the other hand, a different notion of width, *relational width*, equals the level of BW hierarchy needed to solve k-SAT [AD08]. This suggests that null-constraining CSPs should have linear-level BW hierarchy lower bounds, too, if one blurs the line between resolution width and relational width, and if the equivalence in [AD08] generalizes to CSPs beyond k-SAT. Inspired by this connection, we directly prove linear-level BW lower bounds for null-constraining CSPs.

Theorem 1.1 (Special case of Theorem 6.24)**.** *Let $k \geqslant 2$. Suppose a k-CSP $(D, \mathcal{R})$ is null-constraining. Then, for any constraint density $\Delta > 0$, with uniformly positive probability, a random instance of the CSP with n variables and Δn constraints has a BW hierarchy solution of level $\Omega(n)$.*

Theorem 6.24 is the counterpart of [CM13, Theorem 1.2], translating their exponential resolution complexity (or linear resolution width) lower bounds into BW hierarchy level lower bounds. Our theorem generalizes [CŽ24b, Theorem 1] that in turn implies [AD22, Corollary 1] and [CŽ24b, Corollary 2]. We prove our theorem using a new notion of "boundary" (BW boundary) inspired by a similar notion in [CM13], and marrying BW boundary with the closure technique in [KMOW17, CNP24]. Our theorem has a shorter and simpler proof than [CŽ24b, Theorem 1], while being more general. Much like [CM13, Theorem 1.2], Theorem 6.24 identifies the most general class of random CSPs having a linear-level BW lower bound with uniformly positive probability at every constant constraint density; see Section 4.1.[1]

BW boundary can be applied not just to the BW hierarchy, but also to AIP. We introduce a new condition for CSP called "lax", and prove linear-level AIP lower bounds for CSPs that are null-constraining and lax, using the notion of BW boundary.

Theorem 1.2 (Theorem 7.18)**.** *Let $\ell, k \geqslant 2$. Suppose a k-CSP is null-constraining and lax. For any constraint density $\Delta > 0$, with uniformly positive probability, a random instance of the CSP with n variables and Δn constraints has an AIP hierarchy solution of level $\Omega(n)$.*

A CSP is lax if every constraint C and every variable v of C has an assignment b to the remaining variables $V(C) \setminus \{v\}$, so that C is satisfied by b regardless of the value for v. A sufficient condition for "lax" is that the satisfying assignments of every constraint contain a Hamming ball of radius 1. Our theorem implies linear-level lower bounds not just for random k-SAT and Max-kCSP considered in [CNP24], but also for random 3NAE, k-hypergraph q-coloring, and its stronger variant k-hypergraph rainbow q-coloring (for $k > q$), whose average-case AIP hierarchy lower bounds do not follow from [CNP24].

[1]"Uniformly positive probability" cannot be strengthened to "asymptotically almost surely" without further assumptions about the CSP, because with uniformly positive probability a random instance has constant-length cycles that may be unsatisfiable.

Combined Hierarchy:

Rather than considering each hierarchy in isolation, recent works have combined them for stronger algorithms. A weakness of BW, LP and SDP hierarchies is that they cannot solve linear equations over finite fields. Brakensiek, Guruswami, Wrochna, and Živný [BGWŽ20] proposed a new relaxation, basic LP+AIP, that fixes the weakness of LP with AIP. In LP+AIP, a solution is a pair $(s^{\mathrm{LP}}, s^{\mathrm{AIP}})$ of solutions to the constituent hierarchies, and s^{AIP} is further restricted to be contained inside the support of s^{LP}. This corresponds to first finding an LP solution s^{LP}, and then strengthening every constraint by $\mathrm{supp}(s^{\mathrm{LP}})$, and finally finding an AIP solution s^{AIP} to the strengthened constraints. [BGWŽ20] showed that basic LP+AIP solves all tractable boolean CSPs, and asked if the LP+AIP hierarchy also solves all tractable (non-boolean) CSPs.

Bulatov [Bul17, Bul20] and Zhuk [Zhu20] independently gave algorithms for tractable CSPs, by skillfully combining local consistency and Gaussian elimination, the two main techniques for CSPs. Encouraged by [BGWŽ20], a number of papers [CŽ22, ÓC22, DO24] proposed various combinations of BW or LP with AIP, and asked whether their hierarchies also solve all tractable CSPs, while being simpler to describe than Bulatov and Zhuk's. See [CŽ22, Page 5], [ÓC22, Question 13], [DO24, Conjecture 4.10] for questions and conjectures about the power of combined hierarchies for solving CSPs. Lower bounds to LP+AIP [CŽ23b] and to SDP+AIP hierarchies [CŽ24a, CNP24] were recently obtained by Ciardo and Živný, and by Chan, Ng, and Peng.

Lichter and Pago [LP24] recently disproved the conjectures in [BGWŽ20, CŽ22, DO24], showing the weaker hierarchies ($\mathbb{Z}$-affine consistency, LP+AIP, and CLAP) cannot solve some tractable coset-CSP. They were unable to show limitations for the stronger hierarchies, cohomological consistency of Ó Conghaile [ÓC22] and C(BLP+AIP) of Ciardo and Živný [CŽ22], and even conjectured these two hierarchies solve the CSP they considered.[2] In fact there are currently no lower bounds for cohomological consistency and C(BLP+AIP). It was unknown whether these two stronger hierarchies at constant level solve 3SAT in polynomial time.

Lower bounds have been evasive so far because cohomological consistency and C(BLP+AIP) have a new form of restriction not present in others. These two hierarchies require, in place of an AIP solution s^{AIP}, a *family* r^{AIP} of AIP hierarchy solutions, one for each local assignment b in the support of $s^{\mathfrak{A}}(K)$ for each small variable subset K (where $\mathfrak{A}$ is BW or LP). Intuitively, $r^{\mathrm{AIP}}(K, b)$ represents an AIP solution after conditioning on an assignment b to a subset K. Further, these hierarchies ask for valid BW+AIP or LP+AIP solutions after conditioning on every b; see Section 3.2 and Appendix B for their formal definitions. Conditioning is common in LP and SDP hierarchy approximation algorithms. BW conditioning appeared in Singleton Arc Consistency [CDG11] that motivated the AIP conditioning step in C(BLP+AIP) and cohomological consistency.

Even though BW, LP, and SDP solutions remain valid after conditioning, the same is not true for AIP. This obstacle explains why similar lower bounds have been evasive for these stronger hierarchies based on AIP conditioning. We overcome this obstacle using the "augmented closure" idea from [CNP24] to construct conditional AIP solutions, and show the first lower bounds to cohomological consistency and C(BLP+AIP).

Theorem 1.3 (Corollary 8.7 and Proposition B.1)**.** *Let $k \geqslant 2$. Suppose a k-CSP is null-contraining and lax. Then, for any constraint density $\Delta > 0$, with uniformly positive probability, a random instance of the CSP with n variables and Δn constraints has a solution in the cohomological consistency hierarchy of level $\Omega(n)$.*

Theorem 1.4 (Theorem 8.8)**.** *Let $k \geqslant 3$. Suppose a k-CSP is pairwise uniform and lax. Then, for any constraint density $\Delta > 0$, with uniformly positive probability, a random instance of the CSP with n variables and Δn constraints has a solution in the C(BLP+AIP) hierarchy of level $\Omega(n)$.*

---

[2]See Remark 1.5 for updates after our paper was accepted to STOC 2025.

Examples of CSPs satisfying Corollary 8.7 are 3NAE, k-hypergraph q-coloring, and k-hypergraph rainbow q-coloring for $k > q$ (Section 4.3). Examples of CSPs satisfying Theorem 8.8 are 3-SAT, and the Max-kCSP considered in [CNP24, Lemma C.4].

Our lower bounds are for level $\Omega(n)$, which is optimal, and rule out algorithms solving 3SAT in subexponential time based on these hierarchies. This lends weight to the exponential time hypothesis that 3SAT instances on n variables require $\exp(\Omega(n))$ time to solve. This answers an open problem of Pago [Pag24], by giving an NP-complete CSP not solvable by any constant-level cohomological consistency. Via reductions [CNP24, Appendix C], these theorems also lead to linear-level lower bounds in the cohomological consistency and C(BLP+AIP) hierarchies for C-vs-$2^{\Omega(C)}$ graph coloring for sufficiently large constant C,[3] among others.

We hope our techniques will be combined with those in [LP24] in the future to resolve the remaining conjectures about cohomological consistency and C(BLP+AIP), and to answer an open question in [LP24], which asks whether these two hierarchies solve all tractable CSPs.

Remark 1.5. After our paper was accepted to STOC 2025, Lichter and Pago updated their arXiv preprint [LP24], proving that cohomological consistency solves the coset-CSP in [LP24], and independently gave an NP-complete CSP not solvable by level-$O(1)$ cohomological consistency.

Remark 1.6. Our results also apply to promise CSP, including many cases whose computational complexity is open. Consider the PCSP($\mathbf{A}, \mathbf{B}$) with relational structures $\mathbf{A}, \mathbf{B}$ of arity k such that $\mathbf{A} \to \mathbf{B}$; see [KO22] for the definitions. Suppose CSP($\mathbf{A}$) is pairwise uniform and lax, and CSP($\mathbf{B}$) is not trivially satisfiable by a single value. Then with uniformly positive probability, random instances of PCSP($\mathbf{A}, \mathbf{B}$) have a linear-level $\mathbf{A}$-solution in the C(BLP+AIP) hierarchy without being $\mathbf{B}$-satisfiable, by Theorem 8.8 and [CNP24, Lemma B.3].

1.1. **Proof Overview.**

BW Closure:

A promising way to construct a hierarchy solution s is to consider the local closure in an instance I of a variable subset S. A BW solution s is a map from every small variable subset S to a subset $s(S) \subseteq A_S$ of satisfying assignments to the subinstance $I[S]$ of I induced by S, so that $\pi_T(s(S)) = s(T)$ (i.e. the projection of $s(S)$ onto T agrees with $s(T)$) for small subsets S, T such that $T \subseteq S$. When constructing the local solution $s(S)$ on S, it is sometimes necessary to look at constraints of I that are not contained in S. The local closure $\mathrm{cl}_S^t(I)$ of S identifies the subinstance of I based on which the solution $s(S)$ is constructed. For τ-wise uniform CSPs, local closure was introduced by Kothari, Mori, O'Donnell, and Witmer [KMOW17] for LP and SDP solutions, and in a slightly different form by Benabbas, Georgiou, Magen and Tulsiani [BGMT12] earlier for LP. Chan, Ng, and Peng [CNP24] exploited this idea to construct solutions for BW and AIP using the same closure, which we will call local τ-wise closure. They showed lower bounds for CSPs that are "τ-wise neutral", a notion inspired from "τ-wise uniform". The "τ-wise neutral" condition means that a constraint "disappears" in the eyes of the hierarchy when looking at subsets of size at most τ. For example, 3-SAT is pairwise neutral to all these hierarchies.

Pairwise closure cannot directly show lower bounds for graph coloring, as a graph coloring constraint is an inequality constraint on two variables. On the other hand, given a simple constraint path P of length at least 2, P have completely no effect on its endpoints, in terms of the projection of satisfying assignments of P onto the endpoints. That means sufficiently long constraint paths "disappear" for the BW hierarchy when looking at the endpoints only. Chan and Molloy [CM13] exploited this "null-constraining" property to show lower bounds for resolution complexity. Motivated by their idea, we define a new closure called BW closure to prove BW hierarchy lower bounds. The

[3]*Linear*-level hierarchy lower bound for C-vs-$2^{\Omega(C)}$ graph coloring does not follow from similar lower bound for 3SAT, despite the former problem being NP-hard [KOWŽ23], because the direct sum step [Hua13] in the reduction blows up the instance size by a super-linear factor.

local BW closure of S identifies the variables and constraints of I that must be considered when defining $s(S)$, assuming the CSP is null-constraining. Rather than defining the BW closure directly, it is easier to instead define what constraints should *not* appear in the closure. Such constraints belong to the "BW boundary" of S (Definition 5.20), a notion generalizing the boundary notion in [CM13]. Crucially, when removing a constraint subset $\mathcal{Q}$ in the BW boundary from a (sub)instance J, any satisfying assignment to $J \setminus \mathcal{Q}$ can be extended to a satisfying assignment of J (Lemma 6.19). This property is key to constructing BW hierarchy solutions.

AIP Solution:

Armed with BW closure, we construct new AIP hierarchy solutions. For AIP, the local solution $s(S)$ is a function assigning an integer to every satisfying assignment $a \in A_S$ of the subinstance $I[S]$. A new approach is needed when working with BW closure in place of τ-wise closure. We construct $s(S)$ iteratively, starting from smaller subsets S. When constructing the AIP solution $s(S)$ on S, we first identify subsets $T \subseteq S$ that have all the information T needs for constructing $s(T)$. Such subsets T equal the vertices in its closure $\mathrm{cl}_T(I[S])$. We call such T "insular". Having constructed AIP solutions $s(T)$ on all insular subsets $T \subseteq S$, we are required to construct $s(S)$ consistent with all the $s(T)$, i.e. $\pi_T(s(S)) = s(T)$. The set of solutions $s(S)$ satisfying this consistency requirement is a coset of a subgroup X_S in $\mathbb{Z}^{A_S}$ (Definition 7.12). Constructing $s(S)$ based on $s(T)$'s amounts to choosing $w_S \in X_S$. Having chosen w_T (for $s(T)$ earlier), can we choose $s(S)$ consistent with all the $s(T)$'s?

Each w_T belongs to $\mathbb{Z}^{A_T}$ and is the projection of some $w_{S,T} \in \mathbb{Z}^{A_S}$, and we want to define $s(S) = w_S + \sum_T w_{S,T}$, where T runs over insular subsets of S. The challenge is that $\pi_{T'}(w_{S,T})$ may be nonzero for some other insular $T' \subseteq S$ even if $T' \not\subseteq T$. We have to ensure the extensions $w_{S,T}$ of w_T across different insular subsets T of S do not interfere with each other, despite the w_T's being chosen independently. Our approach is to take $w_{S,T} \stackrel{\text{def}}{=} w_T \otimes \mathbb{1}_{b^{S,T}}$ for some assignment $b^{S,T} \in D^{S \setminus T}$. When the CSP satisfies an additional property we call lax, $w_{S,T}$ can be made to be supported on satisfying assignments of S. Its projection to T' is $\pi_{T'}(w_{S,T}) = \pi_{T' \cap T}(w_T) \otimes \pi_{T' \setminus T}(\mathbb{1}_{b_{S,T}})$. Since insular subsets are closed under intersection, $T' \cap T$ is also insular in S when T, T' are. Then $\pi_{T' \cap T}(w_T) = 0$ as desired because $w_T \in X_T$; this ensures $w_{S,T}$ has zero projection to other insular $T' \subseteq S$ whenever $T \not\subseteq T'$. The actual proof in Section 7 is a lot more complicated because we need to further ensure the neighborhood $B(T, 1)$ of T is insular. For this we will instead relate the insular property to path-closedness (Definition 7.1).

Combined Hierarchy:

To construct solutions $(s^{\mathfrak{A}}, s^{\mathfrak{B}})$ to a combined $\mathfrak{A} + \mathfrak{B}$ hierarchy subject to the support restriction $\mathrm{supp}(s^{\mathfrak{A}}(S)) \supseteq \mathrm{supp}(s^{\mathfrak{B}}(S))$ for every subset S, we follow the approach in [CNP24]: We ensure $\mathrm{supp}(s^{\mathfrak{B}}(R)) \subseteq A_R = \mathrm{supp}(s^{\mathfrak{A}}(R))$ for subsets R from a well-chosen family $\mathcal{Z} \subseteq 2^V$. In other words, $s^{\mathfrak{A}}$ has full support on R. We also need to ensure every small subset S is contained in some $R \in \mathcal{Z}$, so that $s(S) \stackrel{\text{def}}{=} \pi_S(s(R))$. In [CNP24], the family $\mathcal{Z}$ consists of the vertex sets of the local closures of small subsets $S \subseteq V$. [CNP24] worked with only one closure, the τ-wise closure, so working directly with the local closures was manageable. In this work, to handle multiple notions of closures at the same time (e.g. pairwise closure for $\mathfrak{A} = \mathrm{LP}$ and BW closure for $\mathfrak{B} = \mathrm{AIP}$, or BW closure for both $\mathfrak{A} = \mathrm{BW}$ and $\mathfrak{B} = \mathrm{AIP}$), we need a more modular approach. We identify the abstract properties enjoyed by $\mathcal{Z}$, and call any family with the same properties an *insular family* (Definition 6.3), because $\mathcal{Z}$ contains precisely the insular sets mentioned before. We also identify common properties shared by τ-wise closure and BW closure, and work instead with a generic closure (Definition 5.1). We unify proofs about different closures, as well as common steps in constructing different hierarchy solutions. Thanks to this abstraction, our proofs are relatively short and modular.

To tackle the conditioning restriction in stronger hierarchies such as cohomological consistency, we have to construct a conditional AIP solution $r^{\mathrm{AIP}}(K, b)$ given a local assignment b to a variable subset K. We reuse the augmented closure idea from [CNP24]: To construct the local solution on subset S conditioned on some assignment b on a different subset K, we need to consider the closure of $K \cup S$ (or rather, a set R in $\mathcal{Z}$ containing both K and S; see Eq. (8)). Conditioning presents a new challenge not faced by [CNP24]: the AIP solution to R must now be supported on assignments agreeing with b, so that the AIP solution is an extension of b (Definition 6.6). And for BW and LP, we require an even stronger property that $\mathrm{supp}(s^{\mathfrak{A}}(R))$ *equals* the set of satisfying assignments on R agreeing with b, i.e., the solution is a *full* extension of b (Definition 6.6). If we represent $s^{\mathfrak{A}}_{K,b}$ as the conditional hierarchy $\mathfrak{A}$ solution given b, this property ensures the support restriction $\mathrm{supp}(s^{\mathfrak{A}}_{K,b}(S)) \supseteq \mathrm{supp}(r^{\mathrm{AIP}}(K,b)(S))$ holds after conditioning, so that the conditional BW+AIP or LP+AIP solutions are valid. This way we overcome the conditioning hurdle and construct solutions fooling even the C(BLP+AIP) hierarchy, the strongest combination of BW, LP, and AIP in the literature.

1.2. Paper Organization.

We define all hierarchies in Section 3, and various classes of CSPs in our theorem statements in Section 4. Section 5 introduces generalizations of closure and boundary operators; it also introduces BW boundary and reviews its expansion property. Section 6 introduces insular family and a framework for hierarchy lower bounds, and proves lower bounds for BW and LP. Section 7 proves AIP lower bounds. Section 8 proves lower bounds to strong combined hierarchies. Appendix B shows the connection between our strong combined hierarchies and cohomological consistency and C(BLP+AIP).

2. Preliminaries

2.1. Miscellaneous.

A set is t-small if its cardinality is at most t. $\mathbb{N} \overset{\text{def}}{=} \{0, 1, 2, \dots\}$ denotes the set of natural numbers, and $[k] \overset{\text{def}}{=} \{1, 2, \dots, k\}$.

An event E_n happens with uniformly positive probability if there is $\varepsilon > 0$ such that $\mathbb{P}[E_n] \geqslant \varepsilon$ for all sufficiently large n.

Given a function $a : S \to D$, $a_T : T \to D$ denotes its restriction to $T \subseteq S$.

The notation $f : (x \in X) \to Y_x$ means that f is a dependent function mapping every $x \in X$ to $f(x) \in Y_x$, where the codomain Y_x varies with x. Given $Z \subseteq X$, $f|_Z : (z \in Z) \to Y_z$ is the restriction of f to Z, given by $f|_Z(z) \overset{\text{def}}{=} f(z)$ for $z \in Z$.

A family $\mathcal{T}$ of subsets over a universe is a down set (aka downward closed set), if $A \in \mathcal{T}$ and $B \subseteq A$ imply $B \in \mathcal{T}$.

Given a family $\mathcal{F} \subseteq 2^U$ of subsets over a universe U and a subset $S \subseteq U$, the subfamily of $\mathcal{F}$ contained in S is $\mathcal{F}[S] \overset{\text{def}}{=} \{T \in \mathcal{F} \mid T \subseteq S\}$.

Given functions $a : S \to D$ and $b : T \to D$ defined on subsets $S, T \subseteq V$ such that $a_{S\cap T} = b_{S\cap T}$, $(a \cup b)$ denotes the combined function: $(a \cup b)(v) \overset{\text{def}}{=} a(v)$ if $v \in S$, and $(a \cup b)(v) \overset{\text{def}}{=} b(v)$ if $v \in T$.

Given a subset $S \subseteq V$ and assignment $a \in D^S$, denote by $\mathbb{1}_a : D^V \to \mathbb{Z}$ the indicator function of a, defined as $\mathbb{1}_a(b) = 1$ if $b_S = a_S$, and $\mathbb{1}_a(b) = 0$ otherwise.

$\mathbb{0}$ denotes the empty function, i.e. given any set D, $\mathbb{0}$ is the unique function from D to $\emptyset$.

Given commutative semiring $\mathcal{M}$ and sets S, T and D, the tensor product of $u \in \mathcal{M}^{D^S}$ and $v \in \mathcal{M}^{D^T}$ is $u \otimes v : D^{S\cup T} \to \mathcal{M}$, given by

$$(u \otimes v)(a, b) \overset{\text{def}}{=} u(a)v(b) \quad \text{for } (a, b) \in D^S \times D^T \cong D^{S\cup T}.$$

2.2. Hypergraph.

A hypergraph H consists of a finite vertex set and a family of subsets (each called a hyperedge) over the vertex set. $V(H)$ denotes its vertex set and $\mathcal{E}(H)$ denotes its hyperedge set. Given a hyperedge set $\mathcal{E}$, $V(\mathcal{E}) \overset{\text{def}}{=} \bigcup_{e\in\mathcal{E}} e$ denotes its set of vertices. Write $H = (V, \mathcal{E})$ if a hypergraph H has vertex set V and hyperedge set $\mathcal{E}$.

A hyperedge set $\mathcal{E}$ is k-uniform if every $e \in \mathcal{E}$ has cardinality exactly k. A hypergraph is k-uniform if its hyperedge set is.

The degree $\deg(v)$ of a vertex $v \in V$ in a hyperedge set $\mathcal{E}$ is the number of hyperedges in $\mathcal{E}$ containing v.

A walk in a hypergraph $H = (V, \mathcal{E})$ is a sequence of vertices $v_0, \dots, v_\ell$ in V and a sequence of hyperedges $(e_0, \dots, e_{\ell-1})$ in $\mathcal{E}$ satisfying $\{v_i, v_{i+1}\} \subseteq V(e_i)$ for $0 \leqslant i < \ell$. The length of the walk is ℓ. $v_0, \dots, v_\ell$ are its connecting vertices. v_0 and v_ℓ are its endpoints.

A walk is a Berge path if $v_0, \dots, v_\ell$ are distinct and $e_0, \dots, e_{\ell-1}$ are also distinct. A walk is a Berge cycle if (1) $\ell \geqslant 2$; (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 hypergraph H (or a hyperedge set $\mathcal{E}$), denoted $\mathrm{girth}(H)$ (or $\mathrm{girth}(\mathcal{E})$), is the infimum length of a Berge cycle in H (or $\mathcal{E}$).

Given a Berge path (resp. cycle), its hyperedge e_i is simple if every vertex in $V(e_i) \setminus \{v_i, v_{i+1}\}$ has degree 1 in $\{e_j\}_{0 \leqslant j < \ell}$, and its connecting vertex v_i is simple if it has degree 2 in $\{e_j\}_{0 \leqslant j < \ell}$. A Berge path is simple[4] if all its hyperedges and non-endpoint connecting vertices are. A Berge cycle is simple if all its hyperedges and connecting vertices are.

A simple path is pendant in $\mathcal{E}$ if none of its non-endpoint belongs to any hyperedge in $\mathcal{E} \setminus \{e_0, \dots, e_{\ell-1}\}$.

The distance $\mathrm{dist}(u, v)$ between vertices u and v in a hypergraph $H = (V, \mathcal{E})$ is the infimum length of a Berge path in H with endpoints u and v. Also $\mathrm{dist}(u, S) \stackrel{\text{def}}{=} \inf\{\mathrm{dist}(u, v) \mid v \in S\}$ for $S \subseteq V$. Given $r \in \mathbb{N}$, the r-neighborhood $B(v, r) \stackrel{\text{def}}{=} \{u \in V \mid \mathrm{dist}(u, v) \leqslant r\}$ of $v \in V$ consists of vertices in H of distance at most r from v. And $B(S, r) \stackrel{\text{def}}{=} \bigcup_{v \in S} B(v, r)$ for $S \subseteq V$.

Given a hypergraph $H = (V, \mathcal{E})$ and a vertex subset $S \subseteq V$, the subhypergraph of H induced by S is $H[S] \stackrel{\text{def}}{=} (S, \mathcal{E}[S])$, which has vertex set S and hyperedge set $\mathcal{E}[S] = \{e \in \mathcal{E} \mid e \subseteq S\}$.

Write $\deg_{\mathcal{E}}(v), \mathrm{dist}_{\mathcal{E}}(u, v)$, and $B_{\mathcal{E}}(v, r)$ if $\mathcal{E}$ is not clear from the context.

Given a hyperedge set $\mathcal{E}$, call $e, f \in \mathcal{E}$ "intersecting" if $e \cap f \neq \emptyset$. This reflexive and symmetric binary relation on $\mathcal{E}$ has a transitive closure, called "connected in $\mathcal{E}$," whose equivalence class is a (connected) component of $\mathcal{E}$. $\mathcal{E}$ is connected if it is its connected component.

2.3. **Constraint Satisfaction.**

For simplicity, we consider only constraint satisfaction problems with k *distinct* variables per constraint (k-CSP). A constraint satisfaction problem is a pair $(D, \mathcal{R})$. Its domain D is a nonempty finite set. $\mathcal{R}$ is a nonempty family of k-ary relations over D, i.e. $R \subseteq D^k$ for every $R \in \mathcal{R}$. An instance $I = (V, \mathcal{C})$ of the CSP $(D, \mathcal{R})$ consists of a finite set V of variables and a finite set $\mathcal{C}$ of constraints over V. Every constraint $C \in \mathcal{C}$ is of the form $C = (S, R)$, where the scope $S \in V^k$ of C is a sequence of k distinct variables, and $R \in \mathcal{R}$ is the set of accepting assignments of C. Write $V(C)$ for the set of variables in the scope of C. In other words, an instance I is a hypergraph, where each hyperedge C is associated with (S, R).

Given an instance I, write $V(I)$ for the set of variables in I and $\mathcal{C}(I)$ for the set of constraints in I. Given a constraint set $\mathcal{C}$, write $V(\mathcal{C}) \stackrel{\text{def}}{=} \bigcup_{C \in \mathcal{C}} V(C)$ for the set of variables in $\mathcal{C}$.

A partial assignment $a \in D^{V(C)}$ on a constraint $C = (S, R)$ with scope $S = (v_1, \dots, v_k) \in V^k$ naturally corresponds to an assignment $a' \in D^k$ via $a' \stackrel{\text{def}}{=} a \circ S$, that is, $a'(i) = a(S(i)) = a(v_i)$ for $i \in [k]$. The set of satisfying assignments for C is $A_C \stackrel{\text{def}}{=} \left\{a \in D^{V(C)} \mid a \circ S \in R\right\}$. An assignment $a \in D^V$ satisfies a constraint $C = (S, R)$ if $a_{V(C)} \in A_C$; otherwise a violates C. An assignment $a \in D^V$ satisfies a constraint set $\mathcal{C}$ or an instance $I = (V, \mathcal{C})$ if a satisfies every constraint $C \in \mathcal{C}$; otherwise a violates $\mathcal{C}$ and I. $A_I \subseteq D^V$ denotes the set of satisfying assignments to the instance $I = (V, \mathcal{C})$.

[4]What we call simple path is called path in [MS07, CM13] and fiber in [CŽ24b].

Given a k-CSP $(D, \mathcal{R})$ and a subdomain $D' \subseteq D$, the CSP restricted to D' is $(D, \mathcal{R})|_{D'} \stackrel{\text{def}}{=} (D', \mathcal{R}|_{D'})$, where $\mathcal{R}|_{D'} \stackrel{\text{def}}{=} \{R \in \mathcal{R} \mid R \in D'^k\}$.

A collection $\mathcal{R}$ of k-ary relations is permutation-closed if for any permutation σ of $[k]$, any $R \in \mathcal{R}$, the relation $\sigma(R) \stackrel{\text{def}}{=} \{(\delta_{\sigma(1)}, \ldots, \delta_{\sigma(k)}) \mid (\delta_1, \ldots, \delta_k) \in R\}$ obtained by permuting the coordinates of R according to σ is also in $\mathcal{R}$.

A predicate CSP is a CSP whose domain D is an abelian group, and there is a set $Q \subseteq D^k$ of assignments satisfying a predicate, such that $\mathcal{R} = \{Q + z \mid z \in D^k\}$, i.e. each constraint is formed from any other by a shift. Examples of such CSPs are k-SAT, k-XOR, and k-NAE.

2.4. **Random Instance.**

We consider the model of random instances in [CM13]. In this model, there is a distribution $\mathcal{P}$ over $\mathcal{R}$. When $\mathcal{P}$ is uniform over $\mathcal{R}$, this model coincides with the natural model of random CSPs in many previous works [BGMT12, KMOW17, CNP24].

Definition 2.1 (Random instance)**.** Fix a k-CSP $(D, \mathcal{R})$, a distribution $\mathcal{P}$ over $\mathcal{R}$ with $\operatorname{supp}(\mathcal{P}) = \mathcal{R}$, a finite set V of variables, and $m \in \mathbb{N}$.

(1) A random constraint $C = (S, R)$ is chosen by picking a sequence $S \in V^k$ of k distinct variables uniformly, and R according to $\mathcal{P}$.
(2) A random instance consists of choosing with replacement m uniformly random constraints.

Equivalently, a random instance consists of first picking a k-uniform hypergraph whose hyperedges are chosen uniformly and independently, and then choosing the relation R of each hyperedge (i.e. constraint) independently from $\mathcal{R}$ according to $\mathcal{P}$.

3. Hierarchy

3.1. **Elementary Hierarchy.**

Let us define a common generalization of the following four hierarchies: bounded width (BW), linear program (LP), semidefinite program (SDP), and affine integer program (AIP).

Given a CSP with domain D and an instance $I = (V, \mathcal{C})$ of the CSP, a hierarchy is a collection $(\mathcal{D}_S)_S$ of sets, one for each d-small variable subset $S \subseteq V$. $\mathcal{D}_S$ is the *relaxed domain* or the *set of relaxed assignments* on S, and varies across hierarchies.

$A_S \subseteq D^S$ denotes the set of partial assignments on S that satisfy all constraints $\mathcal{C}[S]$ in I contained in S.

We now define the relaxed domain $\mathcal{D}_S$ of each elementary hierarchy on a subset $S \subseteq V$ of variables.

BW: $\mathcal{D}_S^{\text{BW}} \stackrel{\text{def}}{=} 2^{A_S} \setminus \{\emptyset\}$, the family of nonempty subsets over A_S; see e.g. [AD08, Definition 1].

LP: $\mathcal{D}_S^{\text{LP}} \stackrel{\text{def}}{=} \Delta(A_S)$, the set of distributions over A_S [BGMT12, Lemma 2.3]. The LP hierarchy is known as the Sherali–Adams hierarchy.

SDP: $\mathcal{D}_S^{\text{SDP}} \stackrel{\text{def}}{=} \{\alpha_S \in \mathcal{X}^{A_S} \setminus \{0^{A_S}\} \mid \alpha_S(a) \perp \alpha_S(b) \text{ for } a, b \in A_S, a \neq b\}$, the set of orthogonal vectors (not all zero) indexed by A_S in an arbitrary inner product space $\mathcal{X}$. Our formulation is equivalent to [CŽ23c, Definition 7], as well as the formulation of pseudo-expectation/moment matrix in [KMOW17, Definition 2.7]. The SDP hierarchy is also known as the Lasserre hierarchy and its Lagrangian dual as the sum-of-squares hierarchy; see e.g. [Lau09, Section 6].

AIP: $\mathcal{D}_S^{\text{AIP}} \stackrel{\text{def}}{=} \left\{w : A_S \to \mathbb{Z} \;\middle|\; \sum_{a \in A_S} w(a) = 1\right\}$, containing affine integer weights supported on A_S. Our definition here is identical to [CNP24] but different from [CŽ23a, CŽ23c]; see [CNP24, Remark 3.2].

For each hierarchy, we consider a relaxed assignment $\alpha_S \in \mathcal{D}_S$ to be a function from the set of satisfying assignments A_S on S to a commutative monoid $\mathcal{M}$. In each of the elementary hierarchy (BW/LP/AIP) except SDP, $\mathcal{M}$ is additionally a commutative semiring.

BW: $\mathcal{M}_{\text{BW}} \stackrel{\text{def}}{=} \mathbb{B}$. The Boolean algebra $\mathbb{B}$ has join $\vee$ as the addition and meet $\wedge$ as the multiplication.

LP: $\mathcal{M}_{\text{LP}} \stackrel{\text{def}}{=} \mathbb{R}_+$, the nonnegative reals.

SDP: $\mathcal{M}_{\text{SDP}} \stackrel{\text{def}}{=} \mathcal{X}$.

AIP: $\mathcal{M}_{\text{AIP}} \stackrel{\text{def}}{=} \mathbb{Z}$.

For any subsets $T \subseteq S \subseteq V$, there is a projection $\pi_{S\to T} : \mathcal{M}^{D^S} \to \mathcal{M}^{D^T}$ of functions taking values in $\mathcal{M}$:

$$\pi_{S\to T}(\alpha_S)(b) \stackrel{\text{def}}{=} \sum_{\substack{a\in D^S\\ a_T=b}} \alpha_S(a) \qquad \text{for } \alpha_S \in \mathcal{M}^{D^S}, b \in D^T. \tag{1}$$

Here the sum is over $\mathcal{M}$.

If $\alpha_S : A \to \mathcal{M}$ is a function from $A \subseteq D^S$, we also think of it as a function $\alpha_S : D_S \to \mathcal{M}$ from D_S supported on A:

$$\alpha_S(b) = 0 \qquad \text{for } b \in D^S \setminus A. \tag{2}$$

Here 0 denotes the additive identity of the commutative monoid $\mathcal{M}$. Given any $\alpha : X \to \mathcal{M}$, denote by $\operatorname{supp}(\alpha) \stackrel{\text{def}}{=} \{x \in X \mid \alpha(x) \neq 0\}$ the support of α.

Eq. (1) immediately implies that projection commutes with addition:

$$\pi_{S\to T}(\alpha_S + \beta_S) = \pi_{S\to T}(\alpha_S) + \pi_{S\to T}(\beta_S) \qquad \text{for any } \alpha_S, \beta_S \in \mathcal{M}^{D^S}. \tag{3}$$

It is also easy to verify that composition of compatible projections is a projection:

$$\pi_{S\to R} = \pi_{T\to R} \circ \pi_{S\to T} \qquad \text{for any } R \subseteq T \subseteq S. \tag{4}$$

We sometimes abbreviate $\pi_{S\to T}$ as π_T when S is clear from the context.

Definition 3.1. A level-d hierarchy solution s maps every d-small vertex subset $S \subseteq V$ to $\mathcal{D}_S$ and is consistent with projection:

$$s(T) = \pi_{S\to T}(s(S)) \qquad \text{for every } T \subseteq S \subseteq V, |S| \leqslant d.$$

Level d of any of the BW, LP, and AIP hierarchy can be solved in time $|V|^{O(d)}$, and the SDP hierarchy can be approximately solved in time $|V|^{O(d)}$. See [CNP24, Section 3.2] for a more thorough discussion.

3.2. **Combined Hierarchy.**

Starting with [BGWŽ20], a number of hierarchies were proposed by combining the elementary ones (BW, LP, SDP, and AIP). In the literature, two ways of strengthening the solution $s^{\mathfrak{A}}$ of another hierarchy with AIP were proposed: support restriction and conditioning.

Support restriction: Each combined hierarchy strengthens AIP by requiring its solution to be contained in the support of the solution of some other hierarchy $\mathfrak{A}$.

Formally, a level-d solution to the combined $\mathfrak{A}$ + AIP hierarchy $s \stackrel{\text{def}}{=} (s^{\mathfrak{A}}, s^{\text{AIP}})$ consists of

- a level-d solution $s^{\mathfrak{A}}$ of an elementary hierarchy $\mathfrak{A}$ (e.g. BW, LP, or SDP), and
- a level-d solution s^{AIP} of the AIP hierarchy,

such that for every subset $S \subseteq V$ of size at most d,

- $\operatorname{supp}(s^{\mathfrak{A}}(S)) \supseteq \operatorname{supp}(s^{\text{AIP}}(S))$.

Examples are:

LP+AIP: The LP+AIP relaxation was introduced in [BGWŽ20]; see also [BG19]. The LP+AIP hierarchy was formally defined in [CŽ23b, Section 3.2].

BW+AIP: The BW+AIP hierarchy was introduced by Dalmau and Opršal [DO24] as the $\mathbb{Z}$-affine consistency hierarchy.

SDP+AIP: The SDP+AIP hierarchy was introduced concurrently (in equivalent forms) in [CNP24, CŽ24a].

Conditioning: Starting with [CŽ22], an even stronger restriction to the combined relaxation and hierarchies was considered: The restriction requires a *family* r^{AIP} of AIP solutions contained in the support of the solution $s^{\mathfrak{A}}$ of another hierarchy $\mathfrak{A}$, one for each local assignment b to a small subset S in the support of $s^{\mathfrak{A}}$. These AIP solutions are further required to be an indicator function of the corresponding local assignment, so that $r^{\mathrm{AIP}}(S, b)$ represents an AIP solution conditioned on b. Intuitively, this restriction corresponds to conditioning $s^{\mathfrak{A}}$ on b to get a new solution $s^{\mathfrak{A}}_{S,b}$ of $\mathfrak{A}$ (which can be done if $\mathfrak{A}$ is BW, LP, and SDP), and then requiring $\mathrm{supp}(s^{\mathfrak{A}}_{S,b}(T)) \supseteq \mathrm{supp}(r^{\mathrm{AIP}}(S, b)(T))$ for every S, b and T.

Formally, a level-d solution to the strong $\mathfrak{A}$ + AIP hierarchy $s \stackrel{\mathrm{def}}{=} (s^{\mathfrak{A}}, r^{\mathrm{AIP}})$ consists of

- a level-d solution $s^{\mathfrak{A}}$ of another hierarchy (BW, LP, or SDP), and
- a function r^{AIP} mapping every $S \in \binom{V}{\leqslant d}$ and $b \in \mathrm{supp}(s^{\mathfrak{A}}(S))$ to a level-d AIP hierarchy solution,

such that for every $S, T \in \binom{V}{\leqslant d}, b \in \mathrm{supp}(s^{\mathfrak{A}}(S))$,

- $\mathrm{supp}(s^{\mathfrak{A}}(T)) \supseteq \mathrm{supp}(r^{\mathrm{AIP}}(S, b)(T))$, and
- $r^{\mathrm{AIP}}(S, b)(S) = \mathbb{1}_b$.

Examples are:

Strong LP+AIP: Our Strong LP+AIP hierarchy is closely related to, though weaker than, the C(BLP+AIP) hierarchy [CŽ22]; see Appendix B.2 for the latter. We also manage to translate our hierarchy lower bound for strong LP+AIP to C(BLP+AIP).

Strong BW+AIP: A solution to our strong BW+AIP hierarchy is a solution to cohomological consistency [ÓC22]; see Appendix B.1.

Of course, one can also consider strong SDP+AIP. Strong SDP+AIP is related, though not identical, to the approximation algorithm in [BKM24].

4. CSP Class

In the following subsections, we discuss different classes of CSPs appearing in our theorem statements.

4.1. Null-Constraining.

Null-constraining CSPs were introduced by Chan and Molloy [CM13] in the context of resolution complexity.

Definition 4.1 ([CM13, Definition 2.1])**.** A simple constraint path P *permits* $(d_1, d_2) \in D^2$ if some assignment $a \in A_P$ of P satisfies all constraints of P, and the endpoints $x, y \in V(P)$ of P receives the assignment $a(x) = d_1$, $a(y) = d_2$.

Definition 4.2 (Null-constraining)**.** A collection $\mathcal{R}$ of k-ary relations over D is *ℓ-null-constraining* if every simple constraint path over $\mathcal{R}$ of length at least ℓ permits every $(d_1, d_2) \in D^2$. $\mathcal{R}$ is *null-constraining* if it is ℓ-null-constraining for some $\ell \in \mathbb{N}$. A k-CSP $(D, \mathcal{R})$ is null-constraining if $\mathcal{R}$ is.

[CM13, Theorem 1.2] showed that random instances of null-constraining k-CSPs have exponential resolution complexity with uniformly positive probability, at every constant constraint density. Our Theorem 6.24 translates their result into linear level lower bound for the BW hierarchy.

If a k-CSP $(D, \mathcal{R})$ has a subdomain $D' \subseteq D$ such that the CSP $(D, \mathcal{R})|_{D'}$ restricted to D' is null-constraining, then the CSP also has a linear level BW lower bound. This is because any BW hierarchy solution of an instance of $(D, \mathcal{R})|_{D'}$ is also a BW hierarchy solution of the corresponding instance of $(D, \mathcal{R})$, which means the lower bounds apply to $(D, \mathcal{R})$ as well.

What if the CSP has no subdomain $\mathcal{D}'$ whose restricted CSP $(D, \mathcal{R})|_{D'}$ is null-constraining? [CM13] showed that random instances of such a CSP have at most polylogarithmic resolution complexity, asymptotically almost surely. To describe their result, we need a few more definitions.

Definition 4.3 ([CM13, Definition 2.2])**.** Given a set of constraint relations $\mathcal{R}$ and two values $d_1, d_2 \in D$ in the domain D, write $d_1 \sim_{\mathcal{R}} d_2$ if there is $\ell \in \mathbb{N}$ such that every simple constraint path over $\mathcal{R}$ of length at least ℓ permits (d_1, d_2).

Without loss of generality, we may assume $\mathcal{R}$ is always permutation-closed because our random model randomly permutes variables in $R \in \mathcal{R}$ when imposing R on each constraint.

Proposition 4.4 ([CM13, Proposition 2.3][5])**.** *For every permutation-closed $\mathcal{R}$, $\sim_{\mathcal{R}}$ is transitive and symmetric.*

It is a bit deceptive to only look at the set of constraint relations $\mathcal{R}$ itself, as constraint paths over $\mathcal{R}$ can sometimes strengthen other constraints in $\mathcal{R}$ to effectively impose a stronger constraint that is not in $\mathcal{R}$.

If $d_1 \not\sim_{\mathcal{R}} d_2$, then an arbitrarily long simple constraint path over $\mathcal{R}$ can still impose a restriction on its endpoints: We cannot assign (d_1, d_2) to its endpoints. If these endpoints x, y lie inside some constraint C outside the path, then C is effectively strengthened by this restriction, and cannot be satisfied with assignments that assign (d_1, d_2) to x, y.

Definition 4.5. [CM13, Definition 2.4] Given a set of constraint relations $\mathcal{R}$. For any $d_1, d_2 \in D$ such that $d_1 \not\sim_{\mathcal{R}} d_2$, $R \in \mathcal{R}$ and $(i, j) \in [k]^2$, consider the relation

$$R_{i,j,d_1,d_2} \stackrel{\text{def}}{=} R \cap \{a \in D^k \mid a(i) \neq d_1 \vee a(j) \neq d_2\}$$

formed by additionally banning the partial assignment (d_1, d_2) on the i^{th} and the j^{th} variable.

The set of constraint relations $\mathcal{R}$ is *closed* if for any $d_1 \not\sim_{\mathcal{R}} d_2$, $R \in \mathcal{R}$ and $(i, j) \in [k]^2$, the relation R_{i,j,d_1,d_2} is also in $\mathcal{R}$. The closure $\mathrm{cl}(\mathcal{R})$ of $\mathcal{R}$ is the smallest closed set of constraint relations containing $\mathcal{R}$.

Definition 4.6 ([CM13, Definition 2.5])**.** A closed set of constraint relations $\mathcal{R}$ is *complete* if $d_1 \not\sim_{\mathcal{R}} d_2$ for all $d_1, d_2 \in D$. Equivalently, $\mathcal{R}$ is complete if it contains the empty relation $\emptyset \subseteq D^k$. A constraint set that is not complete is *incomplete*.

Lemma 4.7 ([CM13, Lemma 2.7])**.** *Given a closed and permutation-closed collection $\mathcal{R}$ of k-ary relations over D, $\mathcal{R}$ is incomplete if and only if $\mathcal{R}|_{D'}$ is null-constraining for some nonempty $D' \subseteq D$.*

[CM13] showed that if $\mathcal{R}$ has incomplete closure, then a random instance of the CSP $(D, \mathcal{R})$ (with n variables and $m = \Theta(n)$ constraints) will have exponential *resolution complexity* with uniformly positive probability, at any constant constraint density. In this paper, we will show the analogous result in terms of linear-level lower bounds of the *BW hierarchy*.

Theorem 6.24 says that CSPs with *incomplete* $\mathrm{cl}(\mathcal{R})$ have a linear level BW lower bound with uniformly positive probability, at every constant constraint density. This is the most general class of CSPs with that conclusion, because if a CSP has *complete* $\mathrm{cl}(\mathcal{R})$, [CM13, Theorem 1.1] shows that random instances of CSP have an unsatisfiable "forbidding flower" subinstance of polylogarithmic size with probability $1 - o(1)$, at sufficiently large constant constraint density. This implies a polylogarithmic-level BW upper bound. It is also not difficult to show that such a "forbidding flower" in fact leads to a *constant*-level BW upper bound.

[5]What we call permutation-closed is called symmetric in [CM13].

4.2. **τ-wise Uniform.**

Definition 4.8 (τ-wise uniform)**.** Given natural numbers $k \geqslant \tau$, a k-CSP $(D, \mathcal{R})$ is τ-wise uniform if for every $R \in \mathcal{R}$, there is a τ-wise uniform distribution $\mu_R \in \Delta(R)$ over the relation R. A distribution $\mu \in \Delta(D^k)$ over k-ary tuples in D is τ-wise uniform if its marginal $\mu_T(\mu)$ on every subset $T \subseteq [k]$ of size at most τ is uniform.

Pairwise uniform CSPs were introduced by Austrin and Mossel [AM09], who showed such CSPs are approximation resistant under the Unique Games Conjecture (UGC). There is also partial progress towards removing UGC from their result [Cha16, BK22]. For pairwise uniform CSPs, linear-level LP hierarchy lower bound was shown by Benabbas, Georgiou, Magen, and Tulsiani [BGMT12], and SDP hierarchy lower bound by Barak, Chan, and Kothari [BCK15]. For τ-wise uniform CSPs, Kothari, Mori, O'Donnell, and Witmer [KMOW17] showed an average-case lower bound whose level has a better dependence on τ, which is tight up to constants [AOW15, RRS17, Ahn20, dT23]. See also [CNP24] for a simplification of [KMOW17].

4.3. **Lax.**

Lax is a new class of CSPs introduced in this paper. We will prove AIP lower bounds for CSPs that are null-constraining and lax. A CSP is lax if every constraint C and every variable v of C has an assignment b to the remaining variables $V(C) \setminus \{v\}$, so that C is satisfied by b regardless of the value for v.

Definition 4.9 (Lax)**.** A k-CSP $(D, \mathcal{R})$ is lax if for every $R \in \mathcal{R}$, $v \in [k]$, there exists $b \in D^{[k]\setminus\{v\}}$ such that $D^v \times \{b\} \subseteq R$.

Example 4.10. Given $k \geqslant 3$, k-NAE is lax. Given any constraint C of k-NAE and $v \in V(C)$, one can find an assignment $b \in \{0,1\}^L$ to the remaining variables $L \stackrel{\text{def}}{=} V(C) \setminus \{v\}$ such that not all literals (of variables in L) are assigned equal. Then given any value $a \in \{0,1\}$ assigned to v, the assignment $\alpha \stackrel{\text{def}}{=} (v \mapsto a) \cup b$ satisfies C.

Example 4.11. Given $k \geqslant q \geqslant 2$, a k-uniform hypergraph is q-colorable if there is a vertex coloring of the hypergraph such that no hyperedge is monochromatic. This problem is a weaker version of the rainbow coloring problem described next. Therefore k-hypergraph q-coloring is lax when $k > q$. That is, when the number of vertices per edge is strictly greater than the number of colors.

Example 4.12. Given $k \geqslant q \geqslant 2$, a k-uniform hypergraph is rainbow q-colorable if there is a vertex coloring of the hypergraph such that every hyperedge receives all q colors. When $k > q$, the rainbow q-coloring problem for k-uniform hypergraph is lax, again by choosing an assignment $b \in [q]^L$ containing all q colors to the remaining variables L. The computational complexity of rainbow coloring was studied in [GL18, GS17, ABP20, GS20]. For example, Guruswami and Sandeep [GS20] showed that for $k \geqslant 4$, given a k-uniform $(k-1)$-rainbow colorable hypergraph, it is NP-hard find a $\lceil \frac{k-2}{2} \rceil$-rainbow coloring. By contrast, Theorem 7.18 and Remark 1.6 imply linear-level AIP lower bound for k-uniform hypergraph q-rainbow coloring even when the hypergraph has an arbitrarily large chromatic number, for every $k > q \geqslant 2$ (and the lower bound also holds for cohomological consistency by Corollary 8.7 and Proposition B.1). Our lower bounds hold for parameters that go beyond current NP-hardness results.

In all of the examples above (k-NAE for $k \geqslant 3$, k-uniform hypergraph (rainbow) q-coloring for $k > q \geqslant 2$), the CSP is also 1-null-constraining, i.e. for every constraint C, every pair u, v of distinct variables in C, every pair of values $d_1, d_2 \in D$, there is a satisfying assignment α to C such that $\alpha(u) = d_1$ and $\alpha(v) = d_2$. Therefore these CSPs have AIP lower bounds by Theorem 7.18.

Lemma 4.13. *If the satisfying assignment of every constraint of a k-CSP $(D, \mathcal{R})$ contains a Hamming ball of radius* 1*, then the CSP is lax.*

Proof. For $R \in \mathcal{R}$, suppose R contains some Hamming ball $B = \{a \in D^k \mid \mathrm{dist_H}(a,c) \leqslant 1\}$ centered at $c = c^R \in D^k$ of radius 1. Here $\mathrm{dist_H}(a,c) \overset{\text{def}}{=} |\{i \in [k] \mid a_i \neq c_i\}|$ is the Hamming distance between a and c. Then for every $i \in [k]$ every $a \in D^k$ with $a_{[k]\setminus\{i\}} = c_{[k]\setminus\{i\}}$ belongs to $B \subseteq R$. □

The converse is not true; 3-NAE is lax without containing any Hamming ball of radius 1. If a hamming ball B of radius 1 is contained in the satisfying assignments of a 3-NAE constraint C, the center c of B cannot satisfy all literals or violates all literals of C. If c satisfies exactly 1 literal, then changing that literal would violate C. If c violates exactly 1 literals, then changing that literal would violate C.

[CNP24] showed AIP lower bounds for CSPs that are pairwise neutral for AIP, a sufficient condition of which is that the satisfying assignments of every constraint of the CSP contain a Hamming ball of radius 2. Lemma 4.13 shows that Hamming balls of radius 1 (together with null-constraining) already suffice for AIP lower bounds. However, the class of CSPs having AIP lower bound in this paper does not strictly generalize the class of "pairwise neutral" CSPs in [CNP24]. We have an example of a CSP being pairwise neutral for AIP without being lax that we do not include in this paper.

5. Closure and Boundary

5.1. Closure.

In this subsection, we study a generic notion of closure capturing those in the literature. Throughout this subsection, fix a vertex set V and a collection $\mathcal{H} \subseteq 2^V$ of hyperedges over V.

To construct hierarchy lower bounds, a number of closure operators have been used.[6] The closure of a variable subset S identifies the constraint subset on which S's local solution depends; see [CNP24, Examples 1.8 and 1.9] for concrete motivating examples.

Examples of closure operators include:

- (What we call) τ-wise closure for τ-wise neutral CSPs from [CNP24, Definition 5.1], which is the non-local version of [KMOW17, Definition 5.3]
- Augmented τ-wise closure for combined hierarchies [CNP24, Definition 9.1]
- BW closure introduced in this work for BW and AIP solutions

To unify proofs about them, we now introduce a common generalization.

Definition 5.1. A closure operator cl for $(V,\mathcal{H})$ maps every vertex subset $S \subseteq V$ and hyperedge set $\mathcal{E} \subseteq \mathcal{H}$ to a hyperedge subset $\mathrm{cl}(S,\mathcal{E}) \subseteq \mathcal{E}$. Also write $\mathrm{cl}_S(\mathcal{E})$ for $\mathrm{cl}(S,\mathcal{E})$. For $S \subseteq V$, call a hyperedge set $\mathcal{E} \subseteq \mathcal{H}$ S-closed if $\mathrm{cl}_S(\mathcal{E}) = \mathcal{E}$. Further, cl satisfies all of the following:

- (Monotone in S) $\mathrm{cl}_T(\mathcal{E}) \subseteq \mathrm{cl}_S(\mathcal{E})$ for $T \subseteq S \subseteq V$, $\mathcal{E} \subseteq \mathcal{H}$.
- (Monotone in $\mathcal{E}$) $\mathrm{cl}_S(\mathcal{F}) \subseteq \mathrm{cl}_S(\mathcal{E})$ for $S \subseteq V$, $\mathcal{F} \subseteq \mathcal{E} \subseteq \mathcal{H}$.
- (Idempotent) $\mathrm{cl}_S(\mathrm{cl}_S(\mathcal{E})) = \mathrm{cl}_S(\mathcal{E})$ for $S \subseteq V, \mathcal{E} \subseteq \mathcal{H}$.
- (Containing $\mathcal{E}[S]$) $\mathrm{cl}_S(\{e\}) = \{e\}$ for $e \subseteq S \subseteq V$, $e \in \mathcal{H}$.
- (Chaining) $\mathcal{F}$ is $(S \cup V(\mathcal{E}))$-closed iff $\mathcal{E} \cup \mathcal{F}$ is S-closed, for $S \subseteq V$, $\mathcal{F} \subseteq \mathcal{H}$, S-closed $\mathcal{E} \subseteq \mathcal{H}$.

Remark 5.2. Note that $\mathrm{cl}(S,\mathcal{E})$ is always a *subset* of $\mathcal{E}$. $\mathrm{cl}(S,\mathcal{E})$ is called a closure operator (whose name suggests being a *superset*) because $\mathrm{cl}(S,\mathcal{E})$ is an intermediate step towards defining a subhypergraph $\mathrm{cl}_S(H)$ that contains S in Definition 5.5. $\mathrm{cl}_S(H)$ represents the subhypergraph that contains enough information to determine a relaxed assignment on S. The map $S \mapsto V(\mathrm{cl}_S(H)) = S \cup V(\mathrm{cl}_S(\mathcal{E}))$ justifies the name "closure".

[6] The notion of closure of a hyperedge set in a hypergraph is completely unrelated to the notion of closure of a constraint relation in Section 4.1.

Intuitively, the first two items in Definition 5.1 are natural consequences of that fact that $\mathrm{cl}_S(\mathcal{E})$ is the constraint subset of $\mathcal{E}$ that the local solution $\sigma(S)$ to S depends on. For instance, "monotonicity in S" says that the constraints $\sigma(T)$ depends on is a subset of those that $\sigma(S)$ depends on. The last three items ensures, among others, the vertex-closure map $f \overset{\text{def}}{=} S \mapsto S \cup V(\mathrm{cl}_S(\mathcal{E}))$ is idempotent, i.e., $f(f(S)) = f(S)$, which will be useful for constructing "insular family" in Section 6.1.

Lemma 5.3. *The union of any collection* $\{\mathcal{E}_\alpha\}_{\alpha\in\mathcal{I}}$ *of* S*-closed hyperedge sets is again* S*-closed.*

Proof. $\mathrm{cl}_S(\bigcup_{\alpha\in\mathcal{I}}\mathcal{E}_\alpha) \subseteq \bigcup_{\alpha\in\mathcal{I}}\mathcal{E}_\alpha$ by Definition 5.1. Also, $\bigcup_{\beta\in\mathcal{I}}\mathcal{E}_\beta \supseteq \mathcal{E}_\alpha$ for any $\alpha \in \mathcal{I}$, implying $\mathrm{cl}_S(\bigcup_{\beta\in\mathcal{I}}\mathcal{E}_\beta) \supseteq \mathrm{cl}_S(\mathcal{E}_\alpha) = \mathcal{E}_\alpha$ by monotonicity in $\mathcal{E}$ and S-closedness of $\mathcal{E}_\alpha$. Taking the union over $\alpha \in \mathcal{I}$, we get $\mathrm{cl}_S(\bigcup_{\beta\in\mathcal{I}}\mathcal{E}_\beta) \supseteq \bigcup_{\alpha\in\mathcal{I}}\mathcal{E}_\alpha$. □

Lemma 5.4. $\mathrm{cl}_S(\mathcal{E})$ *equals the union of* S*-closed* $\mathcal{F} \subseteq \mathcal{E}$*.*

Proof. Given any S-closed $\mathcal{F} \subseteq \mathcal{E}$, $\mathrm{cl}_S(\mathcal{E}) \supseteq \mathrm{cl}_S(\mathcal{F}) = \mathcal{F}$ by monotonicity in $\mathcal{E}$ and S-closedness of $\mathcal{F}$. Therefore the union of such $\mathcal{F}$ is contained in $\mathrm{cl}_S(\mathcal{E})$. Conversely, $\mathrm{cl}_S(\mathcal{E})$ is S-closed since cl_S is idempotent, and $\mathrm{cl}_S(\mathcal{E}) \subseteq \mathcal{E}$ by Definition 5.1, so $\mathrm{cl}_S(\mathcal{E})$ is contained in the union of S-closed $\mathcal{F} \subseteq \mathcal{E}$. □

In Definition 5.1, for every fixed S, the closure operator maps hyperedge sets $\mathcal{E}$ to hyperedge sets, indicating which hyperedges/constraints in $\mathcal{E}$ are relevant in defining the local solution on S. We now extend it to a map from hypergraphs H to hypergraphs, indicating which vertices and hyperedges (i.e. subhypergraph) of H are relevant to the local solution on S.

Definition 5.5. Given a hypergraph $H = (V, \mathcal{E})$ and $S \subseteq V$, define $\mathrm{cl}_S(H) \overset{\text{def}}{=} (S \cup V(\mathrm{cl}_S(\mathcal{E})), \mathrm{cl}_S(\mathcal{E}))$.

The next definition generalizes the local τ-wise closure introduced in [KMOW17, Definition 5.3] that the authors of [KMOW17] believed is the "right" definition for τ-wise uniform CSPs.

Definition 5.6 (Local closure)**.** Given a closure operator cl for $(V, \mathcal{H})$, $S \subseteq V$, $\mathcal{E} \subseteq \mathcal{H}$, $t \in \mathbb{N}$, the t-local S-closure $\mathrm{cl}^t_S(\mathcal{E})$ of $\mathcal{E}$ is the union of t-small S-closed $\mathcal{F} \subseteq \mathcal{E}$.

The local closure operator cl^t also becomes an operator mapping hypergraphs to hypergraphs, via Definition 5.5, i.e. $\mathrm{cl}^t_S(H) = (S \cup \mathrm{cl}^t_S(\mathcal{E}), \mathrm{cl}^t_S(\mathcal{E}))$ for a hypergraph $H = (V, \mathcal{E})$.

We now prove a few lemmas relating closure and local closure, generalizing those in [CNP24, Section 8.1].

Lemma 5.7. *For any* $t \in \mathbb{N}$*,* $T \subseteq S \subseteq V$*,* $\mathcal{E} \subseteq \mathcal{H}$*,* $\mathrm{cl}^t_T(\mathcal{E}) \subseteq \mathrm{cl}^t_S(\mathcal{E})$*.*

Proof. Given any t-small T-closed $\mathcal{F} \subseteq \mathcal{E}$, $\mathcal{F} = \mathrm{cl}_T(\mathcal{F}) \subseteq \mathrm{cl}_S(\mathcal{F}) \subseteq \mathcal{F}$, where the first inclusion uses monotonicity in S and the second is Definition 5.1. Therefore $\mathcal{F}$ is also S-closed, and $\mathcal{F} \subseteq \mathrm{cl}^t_S(\mathcal{E})$ by Definition 5.6. Taking the union over all t-small T-closed $\mathcal{F} \subseteq \mathcal{E}$, we have $\mathrm{cl}^t_T(\mathcal{E}) \subseteq \mathrm{cl}^t_S(\mathcal{E})$. □

Lemma 5.8. *For any* $\mathcal{F} \subseteq \mathcal{E} \subseteq \mathcal{H}$*,* $T \subseteq V$*, if* $\mathcal{F}$ *is* t*-small, then*

$$\mathrm{cl}_T(\mathcal{F}) \subseteq \mathrm{cl}^t_T(\mathcal{E}).$$

Proof. $\mathrm{cl}_T(\mathcal{F})$ is T-closed as cl_T is idempotent. Definition 5.1 implies $\mathrm{cl}_T(\mathcal{F}) \subseteq \mathcal{F}$, so $\mathrm{cl}_T(\mathcal{F})$ is t-small. Since $\mathrm{cl}_T(\mathcal{F}) \subseteq \mathcal{F} \subseteq \mathcal{E}$, the result follows by Definition 5.6. □

Lemma 5.9. *For any* $T \subseteq S \subseteq V$*, if* $\mathrm{cl}^t_S(\mathcal{E})$ *is* t*-small, then*

$$\mathrm{cl}_T(\mathrm{cl}^t_S(\mathcal{E})) = \mathrm{cl}^t_T(\mathcal{E}). \tag{5}$$

As a result, $\mathrm{cl}_T(\mathrm{cl}^t_S(H)) = \mathrm{cl}^t_T(H)$ *for* $H = (V, \mathcal{E})$ *if* $\mathrm{cl}^t_S(\mathcal{E})$ *is* t*-small.*

Proof. $\mathrm{cl}^t_T(\mathcal{E}) \subseteq \mathrm{cl}^t_S(\mathcal{E})$ by Lemma 5.7. Further, $\mathrm{cl}^t_T(\mathcal{E})$ is T-closed by Lemma 5.3, so $\mathrm{cl}^t_T(\mathcal{E}) \subseteq \mathrm{cl}_T(\mathrm{cl}^t_S(\mathcal{E}))$ by Lemma 5.4.

The reverse inclusion is Lemma 5.8 (with $\mathcal{F} \overset{\text{def}}{=} \mathrm{cl}^t_S(\mathcal{E})$). □

5.2. **Boundary.**

Throughout this subsection, fix a vertex set V and a collection $\mathcal{H} \subseteq 2^V$ of hyperedges over V.

Often a closure operator is defined in terms of boundary: Instead of directly identifying a constraint subset on which the local solution depends, it's easier to first identify a constraint subset (or a collection of constraint subsets) on which the local solution does NOT depend, using a boundary operator.[7] Examples of boundary operators include:

- The "exterior" constraints for τ-wise neutral CSPs [CNP24, Section 5.1]
- The BW boundary introduced in this work for BW and AIP

To unify proofs about them, we now introduce a common generalization.

Definition 5.10. A boundary operator $\mathcal{B}$ for $(V, \mathcal{H})$ maps every vertex subset $S \subseteq V$ and hyperedge set $\mathcal{E} \subseteq \mathcal{H}$ to a family of nonempty subsets of $\mathcal{E}$, i.e. $\mathcal{B}(S, \mathcal{E}) \subseteq 2^{\mathcal{E}} \setminus \{\emptyset\}$. Also write $\mathcal{B}_S(\mathcal{E})$ for $\mathcal{B}(S, \mathcal{E})$. Further,

- (Anti-monotone in S) $\quad \mathcal{B}_T(\mathcal{E}) \supseteq \mathcal{B}_S(\mathcal{E})$ for $T \subseteq S \subseteq V$, $\mathcal{E} \subseteq \mathcal{H}$.
- (Independent of $\mathcal{E}[S]$) $\quad \mathcal{B}_S(\mathcal{E}) = \mathcal{B}_S(\mathcal{E} \setminus \mathcal{F})$ for $S \subseteq V$, $\mathcal{E} \subseteq \mathcal{H}$, $\mathcal{F} \subseteq \mathcal{E}[S]$.
- (Chaining) $\quad \mathcal{B}_S(\mathcal{E}) = \mathcal{B}_{S \cup V(\mathcal{F})}(\mathcal{E})$ whenever $\mathcal{B}_S(\mathcal{F}) = \emptyset$, for $S \subseteq V$, $\mathcal{F} \subseteq \mathcal{E} \subseteq \mathcal{H}$.
- (Indifferent) $\quad \mathcal{F} \in \mathcal{B}_{S \cup V(\mathcal{E} \setminus \mathcal{F})}(\mathcal{E})$ iff $\mathcal{F} \in \mathcal{B}_S(\mathcal{E})$ for $S \subseteq V$, $\mathcal{E} \subseteq \mathcal{H}$.

Remark 5.11. For $S \subseteq V$, $\mathcal{F} \subseteq \mathcal{E} \subseteq \mathcal{H}$, if $\mathcal{B}_S(\mathcal{F}) = \emptyset$, then $\mathcal{B}_S(\mathcal{E}) = \mathcal{B}_{S \cup V(\mathcal{F})}(\mathcal{E}) = \mathcal{B}_{S \cup V(\mathcal{F})}(\mathcal{E} \setminus \mathcal{F}) \subseteq 2^{\mathcal{E} \setminus \mathcal{F}}$, where the first equality is chaining and the second is independence of $\mathcal{E}[S]$.

Example 5.12. [CNP24, Section 5.1] considered

$$\mathcal{B}_S(\mathcal{E}) \stackrel{\text{def}}{=} \{\{e\} \mid e \in \mathcal{E},\ \text{at most } \tau \text{ vertices } u \in e \text{ satisfy } u \in S \text{ or } \deg_{\mathcal{E}}(u) > 1\}.$$

Lemma 5.13. *$\mathcal{B}$ in the previous example is a boundary operator for* $(V, \binom{V}{>\tau})$.

Proof. Anti-monotone in S: Follows immediately from definition.

Independent of $\mathcal{E}[S]$: Fix any $S \subseteq V$, $\mathcal{F} \subseteq \mathcal{E}[S] \subseteq \binom{V}{>\tau}$.

$\{e\} \in \mathcal{B}_S(\mathcal{E})$ iff $e \in \mathcal{E}$ has at most τ vertices $u \in e$ that satisfy $u \in S$ or $\deg_{\mathcal{E}}(u) > 1$, iff e has at most τ vertices $u \in e$ that satisfy $u \in S$ or $\deg_{\mathcal{E} \setminus \mathcal{F}}(u) > 1$ (since $V(\mathcal{F}) \subseteq S$), iff $\{e\} \in \mathcal{B}_S(\mathcal{E} \setminus \mathcal{F})$ ($e \notin \mathcal{F}$ because $|e| > 1$ and e has at least one vertex outside of $S \supseteq V(\mathcal{E}[S]) \supseteq V(\mathcal{F})$).

Chaining: Consider any $\mathcal{F} \subseteq \mathcal{E} \subseteq \binom{V}{>\tau}$ such that $\mathcal{B}_S(\mathcal{F}) = \emptyset$. Fix any $S \subseteq V$.

If $\{e\} \in \mathcal{B}_S(\mathcal{E})$, e has at most τ vertices u such that $u \in S$ or $\deg_{\mathcal{E}}(u) > 1$. Since $\deg_{\mathcal{G}}(w) \leqslant \deg_{\mathcal{E}}(w)$ for $\mathcal{G} \subseteq \mathcal{E}, w \in e$, we have $e \in \mathcal{B}_S(\mathcal{G})$ whenever $\mathcal{G}$ contains e. Since $\mathcal{B}_S(\mathcal{F}) = \emptyset$, $e \in \mathcal{E} \setminus \mathcal{F}$. This further implies $\{e\} \in \mathcal{B}_{S \cup V(\mathcal{F})}(\mathcal{E})$, as vertices $w \in e$ with $\deg_{\mathcal{E}}(w) = 1$ must be outside $V(\mathcal{F})$.

The reverse inclusion is anti-monotonicity in S.

Indifferent: Consider any $S \subseteq V$, $\mathcal{E} \subseteq \binom{V}{>\tau}$.

$\{e\} \in \mathcal{B}_S(\mathcal{E})$ iff $e \in \mathcal{E}$ has at most τ vertices $u \in e$ that satisfy $u \in S$ or $\deg_{\mathcal{E}}(u) > 1$, iff $e \in \mathcal{E}$ has at most τ vertices $u \in e$ that satisfy $u \in S \cup V(\mathcal{E} \setminus \{e\})$ or $\deg_{\mathcal{E}}(u) > 1$ (because vertices $v \in e$ with $\deg_{\mathcal{E}}(v) = 1$ does not belong to any hyperedge in $\mathcal{E} \setminus \{e\}$), iff $\{e\} \in \mathcal{B}_{S \cup V(\mathcal{E} \setminus \{e\})}(\mathcal{E})$. □

Inspired by [CNP24, Proposition 5.4], define closure to be the remaining hyperedges after repeatedly removing those in the boundary:

Definition 5.14 (Interior)**.** Given a boundary operator $\mathcal{B}$ for $(V, \mathcal{H})$, define the interior operator by

$$\operatorname{int}_S(\mathcal{E}) \stackrel{\text{def}}{=} \mathcal{E} \setminus \bigcup \mathcal{B}_S(\mathcal{E}) \qquad \text{for } S \subseteq V, \mathcal{E} \subseteq \mathcal{H},$$

and the i-th iterated interior operator by $\operatorname{int}_S^0(\mathcal{E}) \stackrel{\text{def}}{=} \mathcal{E}$, and $\operatorname{int}_S^{i+1}(\mathcal{E}) \stackrel{\text{def}}{=} \operatorname{int}_S(\operatorname{int}_S^i(\mathcal{E}))$ for $i \in \mathbb{N}$.

[7]The concept of (*hyperedge*) boundary operator is closely related to, but distinct from, boundary *vertices* in [BGMT12, KMOW17, CNP24].

Definition 5.15 (Closure operator cl of $\mathcal{B}$)**.** For $S \subseteq V$, cl_S is the fixed-point of int_S^i, i.e. $\mathrm{cl}_S(\mathcal{E}) \stackrel{\text{def}}{=} \mathrm{int}_S^i(\mathcal{E})$, where $i \in \mathbb{N}$ is smallest such that $\mathrm{int}_S^{i+1}(\mathcal{E}) = \mathrm{int}_S^i(\mathcal{E})$.

Definition 5.15 is well-defined: $\mathrm{int}_S(\mathcal{E}) \subseteq \mathcal{E}$ for every $\mathcal{E}$, and the decreasing sequence $\mathcal{E} \supseteq \mathrm{int}_S(\mathcal{E}) \supseteq \mathrm{int}_S^2(\mathcal{E}) \supseteq \dots$ must hit a fixed-point.

Remark 5.16. $\mathcal{E}$ is S-closed if and only if $\mathcal{B}_S(\mathcal{E}) = \emptyset$. Indeed, $\mathcal{E}$ is S-closed implies $\mathrm{cl}_S(\mathcal{E}) = \mathcal{E}$, so $\mathrm{int}_S^{i+1}(\mathcal{E}) = \mathrm{int}_S^i(\mathcal{E}) = \mathcal{E}$ for some i, and $\mathrm{int}_S(\mathcal{E}) = \mathrm{int}_S(\mathrm{int}_S^i(\mathcal{E})) = \mathrm{int}_S^{i+1}(\mathcal{E}) = \mathcal{E}$, implying $\mathcal{B}_S(\mathcal{E}) = \emptyset$. Conversely, $\mathcal{B}_S(\mathcal{E}) = \emptyset$ implies $\mathrm{int}_S(\mathcal{E}) = \mathcal{E} = \mathrm{int}_S^0(\mathcal{E})$, so $\mathcal{E}$ is a fixed-point and $\mathrm{cl}_S(\mathcal{E}) = \mathcal{E}$.

Lemma 5.17. *The closure operator* cl *of* $\mathcal{B}$ *satisfies Definition 5.1.*

Proof. Since $\mathrm{int}_S(\mathcal{E}) \subseteq \mathcal{E}$ for every S and $\mathcal{E}$, an induction on $i \in \mathbb{N}$ shows that $\mathrm{int}_S^i(\mathcal{E}) \subseteq \mathcal{E}$ for every i, S, and $\mathcal{E}$. In particular $\mathrm{cl}_S(\mathcal{E}) \subseteq \mathcal{E}$.

Anti-monotonicity in S implies $\mathrm{int}_T(\mathcal{E}) \subseteq \mathrm{int}_S(\mathcal{E})$ for $T \subseteq S \subseteq V, \mathcal{E} \subseteq \mathcal{H}$. An induction on i shows $\mathrm{int}_T^i(\mathcal{E}) \subseteq \mathrm{int}_S^i(\mathcal{E})$ for every i, and in particular $\mathrm{cl}_T(\mathcal{E}) \subseteq \mathrm{cl}_S(\mathcal{E})$.

For $S \subseteq V, \mathcal{F} \subseteq \mathcal{E} \subseteq \mathcal{H}$, it follows by an induction on i that $\mathrm{int}_S^i(\mathcal{F}) \subseteq \mathcal{F}$ for every $i \in \mathbb{N}$, and in particular $\mathrm{cl}_S(\mathcal{F}) = \mathrm{int}_S^i(\mathcal{F}) \subseteq \mathcal{F} \subseteq \mathcal{E}$ for some i. $\mathrm{int}_S^{i+1}(\mathcal{F}) = \mathrm{int}_S^i(\mathcal{F})$ implies $\mathcal{B}_S(\mathrm{cl}_S(\mathcal{F})) = \mathcal{B}_S(\mathrm{int}_S^i(\mathcal{F})) = \emptyset$. By Remark 5.11, it follows by an induction on j that $\mathrm{cl}_S(\mathcal{F}) \subseteq \mathrm{int}_S^j(\mathcal{E})$ for every j, and in particular $\mathrm{cl}_S(\mathcal{F}) \subseteq \mathrm{cl}_S(\mathcal{E})$.

cl_S is idempotent, because $\mathrm{cl}_S(\mathcal{E}) = \mathrm{int}_S^i(\mathcal{E})$ implies $\mathrm{int}_S(\mathrm{int}_S^i(\mathcal{E})) = \mathrm{int}_S^{i+1}(\mathcal{E}) \stackrel{(*)}{=} \mathrm{int}_S^i(\mathcal{E})$, where $(*)$ uses the definition of $\mathrm{cl}_S(\mathcal{E})$. Thus $\mathrm{cl}_S(\mathrm{cl}_S(\mathcal{E})) = \mathrm{cl}_S(\mathrm{int}_S^i(\mathcal{E})) \stackrel{(\star)}{=} \mathrm{int}_S^i(\mathcal{E}) = \mathrm{cl}_S(\mathcal{E})$, where $(\star)$ uses the definition of $\mathrm{cl}_S(\mathrm{int}_S^i(\mathcal{E}))$.

Given $e \subseteq S \subseteq V$, $\mathcal{B}_S(\{e\}) = \mathcal{B}_S(\emptyset) = \emptyset$ by independence of $\mathcal{E}[S]$. Then $\mathrm{int}_S(\{e\}) = \{e\}$, so $\mathrm{cl}_S(\{e\}) = \{e\}$ by Definition 5.15.

If $\mathcal{E}$ is S-closed, then $\mathcal{B}_S(\mathcal{E}) = \emptyset$, so $\mathcal{B}_S(\mathcal{E} \cup \mathcal{F}) = \mathcal{B}_{S \cup V(\mathcal{E})}(\mathcal{E} \cup \mathcal{F}) = \mathcal{B}_{S \cup V(\mathcal{E})}(\mathcal{F})$, where the first equality is chaining and second is independence of $\mathcal{E}[S]$. Therefore $\mathcal{E} \cup \mathcal{F}$ is S-closed if and only if $\mathcal{F}$ is $(S \cup V(\mathcal{E}))$-closed by Remark 5.16, proving chaining. □

The following lemma was used implicitly in [BGMT12, CNP24] for constructing BW and LP solutions for τ-wise uniform CSPs. It says that starting from a hyperedge set $\mathcal{E}$ and a vertex subset S, one can iteratively remove boundaries from the current hyperedge set, *one at a time*, until the resulting hyperedge set becomes S-closed.

Lemma 5.18. *Given boundary* $\mathcal{B}$ *and its closure* cl *for* $(V, \mathcal{H})$*, for every* $S \subseteq V$*,* $\mathcal{E} \subseteq \mathcal{H}$*, there is a strictly decreasing sequence* $\mathcal{E} = \mathcal{E}_0 \supsetneq \cdots \supsetneq \mathcal{E}_r = \mathrm{cl}_S(\mathcal{E})$ *such that for every* $0 \leqslant i < r$*,* $\mathcal{E}_i \setminus \mathcal{E}_{i+1} \in \mathcal{B}_{S \cup V(\mathcal{E}_{i+1})}(\mathcal{E}_i)$*.*

Proof. We prove by induction on $\mathcal{E}$ along the inclusion order. The base case is $\mathcal{E} = \mathrm{cl}_S(\mathcal{E})$ (which also covers the case $\mathcal{E} = \emptyset$). Then the conclusion holds with $r = 0$.

If $\mathcal{E} \supsetneq \mathrm{cl}_S(\mathcal{E})$, then $\mathrm{int}_S(\mathcal{E}) \subsetneq \mathcal{E}$ by Definition 5.14, and $\bigcup \mathcal{B}_S(\mathcal{E}) \neq \emptyset$. Take any $\mathcal{F} \in \mathcal{B}_S(\mathcal{E})$ and let $\mathcal{K} \stackrel{\text{def}}{=} \mathcal{E} \setminus \mathcal{F}$. Then $\mathcal{F} \in \mathcal{B}_{S \cup V(\mathcal{K})}(\mathcal{E})$ by indifference. We claim that $\mathrm{cl}_S(\mathcal{K}) = \mathrm{cl}_S(\mathcal{E})$. Our claim implies the lemma by applying the Induction Hypothesis to $\mathcal{K} \subsetneq \mathcal{E}$ to yield a sequence s and prepending $\mathcal{E}$ to the front of s.

It remains to prove the claim. cl satisfies Definition 5.1 by Lemma 5.17. $\mathrm{cl}_S(\mathcal{K}) \subseteq \mathrm{cl}_S(\mathcal{E})$ by monotonicity in $\mathcal{E}$. Also, $\mathcal{G} \stackrel{\text{def}}{=} \mathrm{cl}_S(\mathcal{E})$ is S-closed by idempotence of cl_S, i.e. $\mathrm{cl}_S(\mathcal{G}) = \mathcal{G}$. So $\mathrm{int}_S(\mathcal{G}) = \mathcal{G}$ by Definition 5.15, and hence $\bigcup \mathcal{B}_S(\mathcal{G}) = \emptyset$. $\mathcal{G} \cap \mathcal{F} \subseteq \mathcal{G} \cap \bigcup \mathcal{B}_S(\mathcal{E}) = \emptyset$ using $\mathcal{F} \in \mathcal{B}_S(\mathcal{E})$ and Remark 5.11. Thus $\mathcal{G} = \mathrm{cl}_S(\mathcal{E}) \subseteq \mathcal{K}$. Therefore $\mathrm{cl}_S(\mathcal{E}) \subseteq \mathrm{cl}_S(\mathcal{K})$ by Lemma 5.4. □

5.3. **BW Boundary and Closure.**

In this subsection, we introduce BW boundary and BW closure operators. [MS07, CM13] considered the concept of boundary when proving resolution complexity lower bound:

Definition 5.19 ([CM13, Definition 5.2])**.** Let $\ell \geqslant 1$ and $\mathcal{E}$ be a hyperedge set.

(1) The first boundary of $\mathcal{E}$, $\mathcal{B}^1(\mathcal{E})$, is the set of hyperedges e in $\mathcal{E}$ which contain at most one vertex of degree greater than 1 in $\mathcal{E}$.
(2) The second boundary of $\mathcal{E}$, $\mathcal{B}^2_\ell(\mathcal{E})$, is the set of pendant paths[8] of length ℓ in $\mathcal{E}$.
(3) The boundary of $\mathcal{E}$ is $\mathcal{B}_\ell(\mathcal{E}) \stackrel{\text{def}}{=} \mathcal{B}^1(\mathcal{E}) \cup \mathcal{B}^2_\ell(\mathcal{E})$.

Their definition motivates the following generalization:

Definition 5.20. Let $\ell \geqslant 1$, $S \subseteq V$ be a vertex subset, and $\mathcal{E}$ a hyperedge set over V.

- The first S-boundary of $\mathcal{E}$, $\mathcal{B}^1_S(\mathcal{E})$, is the collection of $\{e\}$ over hyperedges $e \in \mathcal{E} \setminus \mathcal{E}[S]$ such that e contains at most one vertex u satisfying $u \in S$ or $\deg_{\mathcal{E}}(u) > 1$.
- The second S-boundary of $\mathcal{E}$, $\mathcal{B}^2_{S,\ell}(\mathcal{E})$, is the collection of hyperedge sets of pendant paths P of length ℓ in $\mathcal{E} \setminus \mathcal{E}[S]$, such that non-endpoints of P are outside S.
- The S-boundary of $\mathcal{E}$ is $\mathcal{B}_{S,\ell}(\mathcal{E}) \stackrel{\text{def}}{=} \mathcal{B}^1_S(\mathcal{E}) \cup \mathcal{B}^2_{S,\ell}(\mathcal{E})$.

Of course, when $S = \emptyset$, Definition 5.20 is equivalent to Definition 5.19.

Lemma 5.21. *For any fixed $\ell \geqslant 2$, $(S, \mathcal{E}) \mapsto \mathcal{B}_{S,\ell}(\mathcal{E})$ is a boundary operator for $(V, \binom{V}{\geqslant 2})$ as in Definition 5.10.*

Proof. Anti-monotone in S: For $T \subseteq S \subseteq V$, $\mathcal{B}^1_T(\mathcal{E}) \supseteq \mathcal{B}^1_S(\mathcal{E})$ and $\mathcal{B}^2_{T,\ell}(\mathcal{E}) \supseteq \mathcal{B}^2_{S,\ell}(\mathcal{E})$ because Definition 5.20 is anti-monotone in S, so $\mathcal{B}_{T,\ell}(\mathcal{E}) \supseteq \mathcal{B}_{S,\ell}(\mathcal{E})$.

Independent of $\mathcal{E}[S]$: Fix any $S \subseteq V$, $\mathcal{F} \subseteq \mathcal{E}[S] \subseteq \binom{V}{\geqslant 2}$.

$\{e\} \in \mathcal{B}^1_S(\mathcal{E})$ if and only if $e \in \mathcal{E}$ has at most one vertex u such that $u \in S$ or $\deg_{\mathcal{E}}(u) > 1$, if and only if e has at most one vertex $u \in S$ or $\deg_{\mathcal{E}\setminus\mathcal{F}}(u) > 1$ (since $V(\mathcal{F}) \subseteq S$), if and only if $\{e\} \in \mathcal{B}^1_S(\mathcal{E} \setminus \mathcal{F})$ ($e \notin \mathcal{F}$ because $|e| > 1$ and e contains at least one other vertex outside of $S \supseteq V(\mathcal{F})$).

$\mathcal{E}(P) \in \mathcal{B}^2_{S,\ell}(\mathcal{E})$ if and only P is a pendant path of length ℓ in $\mathcal{E}$ whose non-endpoints are outside S, if and only if P is a pendant path of length ℓ in $\mathcal{E} \setminus \mathcal{F}$ whose non-endpoints are outside S, if and only if $\mathcal{E}(P) \in \mathcal{B}^2_{S,\ell}(\mathcal{E} \setminus \mathcal{F})$ (every $e \in \mathcal{E}(P)$ is outside $\mathcal{F}$ because $\ell > 1$ and e contains at least one non-endpoint of P that is outside $S \supseteq V(\mathcal{F})$).

Chaining: Consider any $\mathcal{F} \subseteq \mathcal{E} \subseteq \binom{V}{\geqslant 2}$ such that $\mathcal{B}_{S,\ell}(\mathcal{F}) = \emptyset$. Fix any $S \subseteq V$.

If $\{e\} \in \mathcal{B}^1_S(\mathcal{E})$, e contains at most one vertex u such that $u \in S$ or $\deg_{\mathcal{E}}(u) > 1$. Since $\deg_{\mathcal{G}}(w) \leqslant \deg_{\mathcal{E}}(w)$ for $\mathcal{G} \subseteq \mathcal{E}, w \in e$, we have $e \in \mathcal{B}^1_S(\mathcal{G})$ whenever $\mathcal{G}$ contains e. Since $\mathcal{B}_S(\mathcal{F}) = \emptyset$, $e \in \mathcal{E} \setminus \mathcal{F}$. This further implies $\{e\} \in \mathcal{B}^1_{S \cup V(\mathcal{F})}(\mathcal{E})$, as vertices $w \in e$ with $\deg_{\mathcal{E}}(w) = 1$ must be outside $V(\mathcal{F})$.

Suppose $\mathcal{E}(P) \in \mathcal{B}^2_{S,\ell}(\mathcal{E})$ for some pendant path $P = (e_0, \ldots, e_{\ell-1})$ in $\mathcal{E}$, whose non-endpoints are outside S. We claim that $\mathcal{E}(P) \cap \mathcal{F} = \emptyset$. If not, some $e_i \in \mathcal{F}$. Let a be minimal such that all the hyperedges $e_a, e_{a+1}, \ldots, e_i$ are in $\mathcal{F}$. Similarly, let b be maximal such that all the hyperedges $e_i, \ldots, e_b$ are in $\mathcal{F}$. Consider the following cases:

(1) $a = 0$ and $b = \ell - 1$. Then $(e_0, \ldots, e_{\ell-1})$ completely lies inside $\mathcal{F}$ and therefore is a pendant path in $\mathcal{F}$, whose non-endpoints are outside S; or
(2) $a > 0$. Then $e_{a-1} \notin \mathcal{F}$ and e_a has at most one vertex u satisfying $u \in S$ or $\deg_{\mathcal{F}}(u) > 1$.
(3) $b < \ell - 1$. Similar to the case above.

In any of the three cases, $\mathcal{B}_{S,\ell}(\mathcal{F}) \neq \emptyset$, a contradiction. So $\mathcal{E}(P) \subseteq \mathcal{E} \setminus \mathcal{F}$. This further implies $\mathcal{E}(P) \in \mathcal{B}^2_{S \cup V(\mathcal{F}),\ell}(\mathcal{E})$.

The reverse inclusion is anti-monotonicity in S.

[8]A pendant path P in $\mathcal{E}$ is a simple path in $\mathcal{E}$ such that none of the non-endpoint of P belongs to any hyperedge in $\mathcal{E} \setminus \mathcal{E}(P)$. See Section 2.2 for the relevant definitions.

Indifferent: Consider any $S \subseteq V$, $\mathcal{E} \subseteq \binom{V}{\geqslant 2}$.

$\{e\} \in \mathcal{B}^1_S(\mathcal{E})$ if and only if $e \in \mathcal{E}$ has at most one vertex u such that $u \in S$ or $\deg_{\mathcal{E}}(u) > 1$, if and only if $e \in \mathcal{E}$ has at most one vertex u such that $u \in S \cup V(\mathcal{E} \setminus \{e\})$ or $\deg_{\mathcal{E}}(u) > 1$, if and only if $\{e\} \in \mathcal{B}^1_{S \cup V(\mathcal{E}\setminus\{e\})}(\mathcal{E})$.

$\mathcal{E}(P) \in \mathcal{B}^2_{S,\ell}(\mathcal{E})$ if and only P is a pendant path of length ℓ in $\mathcal{E}$ whose non-endpoints are outside S, if and only P is a pendant path of length ℓ in $\mathcal{E}$ whose non-endpoints are outside $S \cup V(\mathcal{E} \setminus \mathcal{E}(P))$, if and only if $\mathcal{E}(P) \in \mathcal{B}^2_{S \cup V(\mathcal{E}\setminus\mathcal{E}(P)),\ell}(\mathcal{E})$. □

Call the boundary operator $\mathcal{B}_{S,\ell}(\mathcal{E})$ in Lemma 5.21 "BW boundary", since it will be used to show BW lower bounds (and in fact AIP too). BW closure, denoted $\mathrm{cl}_{S,\ell}(\mathcal{E})$, is the closure operator of $\mathcal{B}_{S,\ell}(\mathcal{E})$. Call a hyperedge set $\mathcal{F} \subseteq 2^V$ (S,ℓ)-closed if $\mathcal{F}$ is S-closed under $\mathrm{cl}_{S,\ell}(\mathcal{E})$.

Lemma 5.22 (Implicit in [CM13, Lemma 5.8])**.** *If $\mathcal{E}$ is (S,ℓ)-closed, then $|S| \geqslant |\mathcal{B}_\ell(\mathcal{E})|/\ell$.*

Proof. Since $\mathcal{E}$ is (S,ℓ)-closed, $\mathcal{B}_{S,\ell}(\mathcal{E}) = \emptyset$. Any hyperedge $e \in \mathcal{B}^1(\mathcal{E})$ contains at least $|e|-1$ vertices u of degree 1 in $\mathcal{E}$. One of these u must be in S, or else $e \in \mathcal{B}^1_S(\mathcal{E})$. Similarly, any pendant path $P \in \mathcal{B}^2_\ell(\mathcal{E})$ has a non-endpoint in S, for otherwise $P \in \mathcal{B}^2_{S,\ell}(\mathcal{E})$.

No two hyperedges in $\mathcal{B}^1(\mathcal{E})$ share a degree 1 vertex, and therefore there are $|\mathcal{B}^1(\mathcal{E})|$ vertices in S to take all hyperedges of $\mathcal{B}^1(\mathcal{E})$ out of $\mathcal{B}^1_S(\mathcal{E})$. Likewise, a vertex can only be a non-endpoint for at most ℓ pendant paths of length ℓ, so there are at least $|\mathcal{B}^2_\ell(\mathcal{E})|/\ell$ vertices in S to take all pendant paths in $\mathcal{B}^2_\ell(\mathcal{E})$ out of $\mathcal{B}^2_{S,\ell}(\mathcal{E})$.

When $\ell > 1$, these two types of vertices are disjoint, and $|S| \geqslant |\mathcal{B}^1(\mathcal{E})| + |\mathcal{B}^2_\ell(\mathcal{E})|/\ell \geqslant |\mathcal{B}_\ell(\mathcal{E})|/\ell$. When $\ell = 1$, $\mathcal{B}^1(\mathcal{E}) \subseteq \mathcal{B}^2_\ell(\mathcal{E}) = \mathcal{B}_\ell(\mathcal{E})$, so $|S| \geqslant |\mathcal{B}^2_\ell(\mathcal{E})|/\ell = |\mathcal{B}_\ell(\mathcal{E})|/\ell$. □

5.4. **Sparse and Expanding.**

Definition 5.23 (Sparse)**.** A k-uniform hyperedge set $\mathcal{E}$ is δ-sparse if $|\mathcal{E}| \leqslant \frac{1+\delta}{k-1}|V(\mathcal{E})|$. $\mathcal{E}$ is (t,δ)-sparse if every t-small subset of $\mathcal{E}$ is δ-sparse. $\mathcal{E}$ is hereditarily δ-sparse if it is (∞,δ)-sparse. A k-uniform hypergraph is δ-sparse (resp. (t,δ)-sparse) if its hyperedge set is.

We now state two lemmas from [MS07].

Lemma 5.24 ([MS07, Lemma 10])**.** *Let H be a random k-uniform hypergraph with n vertices and Δn edges. For any $\delta > 0$, there is $\mu = \mu(\Delta,k,\delta) > 0$ such that a.s. H is $(\mu n,\delta)$-sparse.*

Remark 5.25. Our Lemma 5.24 differs from [MS07, Lemma 10] in that (1) our δ-sparse condition has a non-strict inequality, while [MS07] is strict; and (2) our μn upper bounds the number of hyperedges in the subhypergraph, while [MS07] upper bounds the number of vertices. Despite the differences, the proof of [MS07, Lemma 10] implies Lemma 5.24.

Lemma 5.26 ([MS07, Lemma 11])**.** *Let $\ell \geqslant 2$, $\mathcal{E}$ be a nonempty k-uniform hyperedge set without any simple cycle component. If $|\mathcal{B}_\ell(\mathcal{E})| \leqslant |V(\mathcal{E})|/(72\ell^2k^3)$, then $\mathcal{E}$ is not δ-sparse for some $\delta = \delta(k,\ell) > 0$.*

Definition 5.27. Let $t \geqslant 0$ and $\gamma > 0$. A hyperedge set $\mathcal{E}$ over V is (ℓ,t,γ)-*expanding* if $|\mathcal{B}_\ell(\mathcal{F})| \geqslant \gamma|\mathcal{F}|$ for any t-small $\mathcal{F} \subseteq \mathcal{E}$. A hypergraph is (ℓ,t,γ)-*expanding* if its hyperedge set is.

The goal for the rest of this section is to show that a random hypergraph is $(\ell,\Omega(n),\Omega(1))$-expanding with uniformly positive probability.

Lemma 5.28. *Let $k \geqslant 2, \ell \geqslant 2, \delta = \delta(k,\ell) > 0$ be given by Lemma 5.26. There is $\gamma = \gamma(k,\ell) > 0$ such that any (t,δ)-sparse k-uniform hypergraph H with* $\mathrm{girth}(H) > \ell$ *is (ℓ,t,γ)-expanding.*

Proof. Let $\gamma \overset{\text{def}}{=} \frac{k-1}{(72\ell^2k^3)(1+\delta)} \leqslant 1$. Let $\mathcal{F}$ be any t-small hyperedge subset of $\mathcal{E}(H)$. We need to show that $|\mathcal{B}_\ell(\mathcal{F})| \geqslant \gamma|\mathcal{F}|$. Partition $\mathcal{F}$ into its connected components $\mathcal{F}_i$'s.

Claim 5.29. *If a component $\mathcal{F}_i$ is a simple cycle, then $|\mathcal{B}_\ell(\mathcal{F}_i)| \geqslant |\mathcal{F}_i|$.*

Proof. Let $\mathcal{F}_i$ be a simple cycle $(e_0, \ldots, e_{r-1})$ of length $r > \ell$. As a Berge cycle, $|\mathcal{B}^1(\mathcal{F}_i)| = 0$. Since $r > \ell$, $\mathcal{F}_i$ has r distinct simple subpaths $P_i \stackrel{\text{def}}{=} (e_i, \ldots, e_{i+\ell-1 \bmod r})$ of length ℓ, for $0 \leqslant i < r$. Each P_i is pendant in $\mathcal{F}$ because the $\mathcal{F}_i$ is a simple cycle component. Therefore $|\mathcal{B}_\ell(\mathcal{F}_i)| \geqslant |\mathcal{F}_i|$. □

Claim 5.30. *If a component $\mathcal{F}_i$ is not a simple cycle, then* $|\mathcal{B}_\ell(\mathcal{F}_i)| \geqslant \frac{k-1}{(72\ell^2k^3)(1+\delta)}|\mathcal{F}_i|$.

Proof. Since $\mathcal{F}_i$ is δ-sparse, $|\mathcal{F}_i| \leqslant \frac{1+\delta}{k-1}|V(\mathcal{F}_i)|$. Lemma 5.26 implies $|\mathcal{B}_\ell(\mathcal{F}_i)| > \frac{1}{72\ell^2k^3}|V(\mathcal{F}_i)| \geqslant \frac{k-1}{72\ell^2k^3(1+\delta)}|\mathcal{F}_i|$. □

The two claims imply $|\mathcal{B}_\ell(\mathcal{F}_i)| \geqslant \gamma|\mathcal{F}_i|$ for every connected component $\mathcal{F}_i$ of $\mathcal{F}$. Hence,

$$|\mathcal{B}_\ell(\mathcal{F})| \stackrel{(*)}{=} \sum_i |\mathcal{B}_\ell(\mathcal{F}_i)| \geqslant \sum_i \gamma|\mathcal{F}_i| = \gamma|\mathcal{F}|,$$

where $(*)$ is due to all $\mathcal{F}_i$ being vertex-disjoint. □

It is well-known that the numbers of Berge cycles of length at most ℓ in a random hypergraph converge in distribution to independent Poisson distributions; see [JLR00, Theorems 3.19, 6.10] for the calculations for graphs. The next lemma follows from such calculations.

Lemma 5.31. *Let $\Delta > 0, \ell \in \mathbb{N}$ and $k \geqslant 2$. There is $\varepsilon = \varepsilon(\Delta, \ell, k) > 0$ such that with probability at least $\varepsilon - o(1)$, a random k-uniform hypergraph with n vertices and Δn hyperedges has no Berge cycles of length at most ℓ.*

Corollary 5.32. *Let $\Delta > 0$, $k \geqslant 2$. There are $\mu > 0$ and $\gamma > 0$ such that with uniformly positive probability, a random k-uniform hypergraph H with n variables and Δn constraints is $(\ell, \mu n, \gamma)$-expanding.*

Proof. By union bound, H satisfies both Lemmas 5.24 and 5.31 with probability $\varepsilon - o(1)$. The corollary now follows from Lemma 5.28. □

Lemma 5.33. *Let $S \subseteq V$, $\mathcal{E} \stackrel{\text{def}}{=} \mathcal{E}_1 \cup \mathcal{E}_2$ be the union of two t-small (S, ℓ)-closed hyperedge sets $\mathcal{E}_1$ and $\mathcal{E}_2$. If $\mathcal{E}$ is $(\ell, 2t, \gamma)$-expanding and $|S| \leqslant r\gamma/\ell$ for some $r \geqslant 0$, then $\mathcal{E}$ is r-small.*

Proof. $\mathcal{E}$ is $2t$-small, so $\gamma|\mathcal{E}| \leqslant |\mathcal{B}_\ell(\mathcal{E})|$. $\mathcal{E}$ is (S, ℓ)-closed by Lemma 5.3, which implies $|S| \geqslant |\mathcal{B}_\ell(\mathcal{E})|/\ell$ by Lemma 5.22. Combining both inequalities gives $\gamma|\mathcal{E}| \leqslant |\mathcal{B}_\ell(\mathcal{E})| \leqslant \ell|S| \leqslant r\gamma$. □

6. Lower Bound

6.1. Insular Family.

In this section, we identify the common properties of the collection $\{V(\mathrm{cl}^{2t}_S(H))\}$ of local closure of small vertex subsets S in a hypergraph H. We call any family satisfying the same properties an *insular family* (Definition 6.3).

Let $\mathfrak{A}$ denote either the BW, LP, or AIP hierarchy throughout this subsection. Recall that the hierarchy has a commutative monoid $\mathcal{M} = \mathcal{M}_{\mathfrak{A}}$. Also fix a closure operator cl for $(V, \mathcal{H})$.

Constructing a hierarchy solution amounts to recursively defining local solutions on variable subsets T in an insular family $\mathcal{Z} \subseteq 2^V$, starting from the minimal subset in $\mathcal{Z}$. Then local solutions propagate to every small subset S (including those not in $\mathcal{Z}$) by projection (Definition 6.13). We have to first construct local solutions on subsets T in $\mathcal{Z}$, rather than on every small subset S directly, because some subset S may not contain all the constraints on which S's local solution depend. By contrast, subsets T in $\mathcal{Z}$ are all "insular" (defined next) and contain all the constraints that T's local solution depends on.

Definition 6.1 (Insular)**.** A vertex subset $S \subseteq V$ is insular in a hyperedge set $\mathcal{E} \subseteq \mathcal{H}$ if $\mathrm{cl}_S(\mathcal{E}) \subseteq \mathcal{E}[S]$.

Since $\mathrm{cl}_S(\{e\}) = \{e\}$ for every $e \in \mathcal{E}[S]$, Lemma 5.4 implies $\mathcal{E}[S] \subseteq \mathrm{cl}_S(\mathcal{E})$. So a subset S is insular if and only if $\mathrm{cl}_S(\mathcal{E}) = \mathcal{E}[S]$. This also implies $V(\mathrm{cl}_S(H)) = S$ for an insular family S in a hypergraph $H = (V, \mathcal{E})$.

Lemma 6.2. *If vertex subsets S and T are both insular in $\mathcal{E}$, then so is $S \cap T$.*

Proof. $\mathrm{cl}_{S\cap T}(\mathcal{E}) \subseteq \mathrm{cl}_S(\mathcal{E}) \subseteq \mathcal{E}[S]$, where the first inclusion is monotonicity in S, and the second is because S is insular in $\mathcal{E}$. Likewise $\mathrm{cl}_{S\cap T}(\mathcal{E}) \subseteq \mathrm{cl}_T(\mathcal{E}) \subseteq \mathcal{E}[T]$. Thus $\mathrm{cl}_{S\cap T}(\mathcal{E}) \subseteq \mathcal{E}[S] \cap \mathcal{E}[T] = \mathcal{E}[S \cap T]$. □

Definition 6.3 (Insular family)**.** Given a hyperedge set $\mathcal{E} \subseteq \mathcal{H}$ or a hypergraph $H = (V, \mathcal{E})$ such that $\mathcal{E} \subseteq \mathcal{H}$, a family $\mathcal{Z} \subseteq 2^V$ of vertex subsets is an insular family if:

- T is insular in $\mathcal{E}[S]$ whenever $S, T \in \mathcal{Z}$ and $T \subseteq S$; and
- $S \cap T \in \mathcal{Z}$ for $S, T \in \mathcal{Z}$.

In particular, if $\mathcal{Z}$ is nonempty, then for any $\mathcal{Y} \subseteq \mathcal{Z}$, the common intersection $\bigcap \mathcal{Y}$ also belongs to $\mathcal{Z}$. Denote by $\min \mathcal{Z} \overset{\text{def}}{=} \bigcap \mathcal{Z}$ the common intersection over all of $\mathcal{Z}$.

Definition 6.1 implicitly depends on a closure operator cl. If we want to emphasize the closure operator cl, we sometimes call a subset S insular in $\mathcal{E}$ *under* cl, or S cl-*insular* in $\mathcal{E}$. The same remark also holds for Definition 6.3.

Given $t \geqslant 1$ and a hypergraph $H = (V, \mathcal{E})$ with $\mathcal{E} \subseteq \mathcal{H}$, define

$$\mathcal{Z}_t \overset{\text{def}}{=} \{S \subseteq V \mid \mathrm{cl}^t_S(\mathcal{E}) \text{ is } t\text{-small and } \mathrm{cl}^t_S(\mathcal{E}) = \mathcal{E}[S]\}.$$

Lemma 6.4. *$\mathcal{Z}_t$ is an insular family.*

Proof. Let $S, T \in \mathcal{Z}_t$ and $T \subseteq S$. Then $\mathrm{cl}_T(\mathcal{E}[S]) = \mathrm{cl}_T(\mathrm{cl}^t_S(\mathcal{E})) \overset{(*)}{=} \mathrm{cl}^t_T(\mathcal{E}) = \mathcal{E}[T] = (\mathcal{E}[S])[T]$, where $(*)$ is Lemma 5.9. Thus T is insular in $\mathcal{E}[S]$.

If $S, T \in \mathcal{Z}_t$, Lemma 5.7 implies $\mathrm{cl}^t_{S\cap T}(\mathcal{E}) \subseteq \mathrm{cl}^t_S(\mathcal{E}) \subseteq \mathcal{E}[S]$, and likewise $\mathrm{cl}^t_{S\cap T}(\mathcal{E}) \subseteq \mathcal{E}[T]$. So $\mathrm{cl}^t_{S\cap T}(\mathcal{E}) \subseteq \mathcal{E}[S] \cap \mathcal{E}[T] = \mathcal{E}[S \cap T]$ and is t-small. Since $\mathrm{cl}_{S\cap T}(\{e\}) = \{e\}$ for every $e \in \mathcal{E}[S \cap T]$, $\mathrm{cl}^t_{S\cap T}(\mathcal{E}) \supseteq \mathcal{E}[S \cap T]$. Therefore $S \cap T \in \mathcal{Z}_t$, and $\mathcal{Z}_t$ is closed under intersection. □

Towards constructing a hierarchy solution, first construct its restriction to an insular family:

Definition 6.5 (Scheme)**.** Given a family $\mathcal{Z}$ of vertex subsets in an instance, a scheme for $\mathcal{Z}$ is a function σ mapping every $S \in \mathcal{Z}$ to $\sigma(S) \in \mathcal{M}^{A_S}$, so that

$$\pi_T(\sigma(S)) = \sigma(T) \qquad \text{whenever } T, S \in \mathcal{Z} \text{ and } T \subseteq S. \tag{6}$$

Given an instance I, for $U \subseteq S \subseteq V$, $q \in A_U$, let $A_{S,q} \overset{\text{def}}{=} \{b \in A_S \mid b_U = q\}$ be the set of satisfying assignments b to $I[S]$ that agree with q.

Definition 6.6 (Extension)**.** Let $\mathcal{Z}$ be an insular family of an instance. An extension of $q \in A_{\min \mathcal{Z}}$ over $\mathcal{Z}$ is a scheme σ for $\mathcal{Z}$ such that $\sigma(\min \mathcal{Z}) = \mathbb{1}_q$ and $\mathrm{supp}(\sigma(S)) \subseteq A_{S,q}$ for $S \in \mathcal{Z}$. σ is full if $\mathrm{supp}(\sigma(S)) = A_{S,q}$ for $S \in \mathcal{Z}$.

Definition 6.7 (Extensible)**.** A CSP is $\mathfrak{A}$-extensible if for every instance $I = (V, \mathcal{C})$ with $\mathcal{C} \subseteq \mathcal{H}$ and insular family $\mathcal{Z}$ of I, every $q \in A_{\min \mathcal{Z}}$ has an extension σ over $\mathcal{Z}$. The CSP is $\mathfrak{A}$-fully-extensible if additionally σ is full.

When the hierarchy is clear from the context, simply say extensible to mean $\mathfrak{A}$-extensible.

The next definition is inspired by [CNP24, Lemma 8.10]. It concerns a property guaranteeing the local closure of a small subset to be small.

Definition 6.8 (Confined)**.** A hyperdege set $\mathcal{E} \subseteq \mathcal{H}$ is (t, β)-confined if for every $r \geqslant 0$, $(r\beta)$-small $S \subseteq V$, and two t-small S-closed $\mathcal{E}_1, \mathcal{E}_2 \subseteq \mathcal{E}$, the union $\mathcal{E}_1 \cup \mathcal{E}_2$ is r-small. A hypergraph is (t, β)-confined if its hyperedge set is.

Remark 6.9. If cl is the BW closure operator with parameter ℓ, Lemma 5.33 implies any $(\ell, 2t, \gamma)$-expanding k-uniform hypergraph is $(t, \gamma/\ell)$-confined.

Remark 6.10. If cl is the τ-wise closure operator, [CNP24, Remark 8.3, Lemma 8.10] imply any $(2t, 2\gamma)$-expanding[9] k-uniform hypergraph is (t, γ)-confined.

Lemma 6.11. *If $\mathcal{E}$ is (t, β)-confined and $|S| \leqslant r\beta$ for some $r \leqslant t$, then $\mathrm{cl}^t_S(\mathcal{E})$ is r-small.*

Proof. Since the local closure is a finite union of t-small S-closed $\mathcal{F} \subseteq \mathcal{E}$, the result follows by induction and Definition 6.8. ☐

Lemma 6.12. *If a k-uniform hypergraph $H = (V, \mathcal{E})$ is $(2t, \beta)$-confined, then $V(\mathrm{cl}^{2t}_S(H)) \in \mathcal{Z}_t$ for every (t/α)-small $S \subseteq V$, and $\min \mathcal{Z}_t = \emptyset$, where $\alpha \overset{\text{def}}{=} 1 + k/\beta$.*

Proof. For any (t/α)-small $S \subseteq V$, $|V(\mathrm{cl}^{2t}_S(H))| \leqslant |S| + |V(\mathrm{cl}^{2t}_S(\mathcal{E}))| \leqslant |S| + k|\mathrm{cl}^{2t}_S(\mathcal{E})| \overset{(*)}{\leqslant} \alpha|S| \leqslant t$, where $(*)$ is Lemma 6.11. In particular $\mathrm{cl}^{2t}_S(\mathcal{E})$ is t-small.

Let $T \overset{\text{def}}{=} V(\mathrm{cl}^{2t}_S(H)) = S \cup V(\mathrm{cl}^{2t}_S(\mathcal{E}))$. We now show that $\mathrm{cl}^t_T(\mathcal{E}) \subseteq \mathrm{cl}^{2t}_S(\mathcal{E})$. Since $\mathrm{cl}^{2t}_S(\mathcal{E})$ is S-closed by Lemma 5.3 and t-small, given any t-small T-closed $\mathcal{F} \subseteq \mathcal{E}$, $\mathcal{F} \cup \mathrm{cl}^{2t}_S(\mathcal{E})$ is S-closed by chaining and $2t$-small, so $\mathcal{F} \subseteq \mathcal{F} \cup \mathrm{cl}^{2t}_S(\mathcal{E}) \subseteq \mathrm{cl}^{2t}_S(\mathcal{E})$ by definition of $\mathrm{cl}^{2t}_S(\mathcal{E})$. Taking the union over all such $\mathcal{F}$, $\mathrm{cl}^t_T(\mathcal{E}) \subseteq \mathrm{cl}^{2t}_S(\mathcal{E})$, and $\mathrm{cl}^t_T(\mathcal{E})$ is also t-small. Also, $\mathrm{cl}^t_T(\mathcal{E}) \subseteq \mathcal{E}[T]$ because $T \supseteq V(\mathrm{cl}^{2t}_S(\mathcal{E}))$. $\mathrm{cl}^t_T(\mathcal{E}) \supseteq \mathcal{E}[T]$ because $\{e\}$ is T-closed for $e \in \mathcal{E}[T]$ and $t \geqslant 1$. Therefore $\mathrm{cl}^t_T(\mathcal{E}) = \mathcal{E}[T]$ and $T \in \mathcal{Z}_t$.

In particular, when $S = \emptyset$, $V(\mathrm{cl}^{2t}_\emptyset(H)) = \emptyset \in \mathcal{Z}_t$. ☐

6.2. **Hierarchy Solution.**

In this subsection, we construct hierarchy solutions based on the notion of insular family.

A function σ for an insular family $\mathcal{Z}$ "trickles down" to sets in the downward closure of $\mathcal{Z}$. Denote by $\mathcal{Z}_\downarrow \overset{\text{def}}{=} \{S \subseteq V \mid S \subseteq T \text{ for some } T \in \mathcal{Z}\}$ the downward closure of $\mathcal{Z}$. For $S \in \mathcal{Z}_\downarrow$, denote by $S^\uparrow \overset{\text{def}}{=} \bigcap\{T \in \mathcal{Z} \mid T \supseteq S\} \in \mathcal{Z}$ the common intersection of all subsets in $\mathcal{Z}$ containing S.

Definition 6.13. Given an insular family $\mathcal{Z}$ and a function $\sigma : (S \in \mathcal{Z}) \to \mathcal{M}^{D^S}$, define $\sigma_\downarrow : (S \in \mathcal{Z}_\downarrow) \to \mathcal{M}^{D^S}$ by

$$\sigma_\downarrow(S) \overset{\text{def}}{=} \pi_S \circ \sigma(S^\uparrow) \qquad \text{for } S \in \mathcal{Z}_\downarrow. \tag{7}$$

More generally, we introduce a generalization useful for combined hierarchies later:

Definition 6.14. Given insular family $\mathcal{Z} \subseteq 2^V$, $K \subseteq V$, let $\mathcal{Z}_{\downarrow K} \overset{\text{def}}{=} \{S \subseteq V \mid K \cup S \in \mathcal{Z}_\downarrow\}$. Further given function $\sigma : (S \in \mathcal{Z}) \to \mathcal{M}^{D^S}$, $K \subseteq V$, define $\sigma_{\downarrow K} : (S \in Z_{\downarrow K}) \to \mathcal{M}^{D^S}$ by

$$\sigma_{\downarrow K}(S) \overset{\text{def}}{=} \pi_S \circ \sigma\left((K \cup S)^\uparrow\right) \qquad \text{for } S \in \mathcal{Z}_{\downarrow K}. \tag{8}$$

Note that $\sigma_\downarrow = \sigma_{\downarrow\emptyset}$. Eq. (8) amounts to the "augmented closure" idea in [CNP24]: when defining the local solution to S conditioned on a local assignment b to K, it is necessarily to consider K and S together.

Lemma 6.15. *If σ is a scheme for $\mathcal{Z}$, then $\sigma_{\downarrow K}$ is a scheme for $\mathcal{Z}_{\downarrow K}$.*

Proof. For any $T \subseteq S \in \mathcal{Z}_{\downarrow K}$,

$$\begin{aligned}\pi_T(\sigma_{\downarrow K}(S)) &\overset{(8)}{=} \pi_T \circ \pi_S \circ \sigma\left((K \cup S)^\uparrow\right) \overset{(4)}{=} \pi_T \circ \sigma\left((K \cup S)^\uparrow\right) \overset{(4)}{=} \pi_T \circ \pi_{(K\cup T)^\uparrow} \circ \sigma\left((K \cup S)^\uparrow\right)\\ &\overset{(6)}{=} \pi_T \circ \sigma\left((K \cup T)^\uparrow\right) \overset{(8)}{=} \sigma_{\downarrow K}(T).\end{aligned}$$

☐

[9]See [CNP24, Definition 8.4] for the definition.

Theorem 6.16. *Suppose a k-CSP is extensible and its instance $I = (V, \mathcal{C})$ with $\mathcal{C} \subseteq \mathcal{H}$ is $(2t, \beta)$-confined. There is a hierarchy solution of level $d \stackrel{\text{def}}{=} t/(1 + k/\beta)$.*

Proof. Lemma 6.12 implies $\mathcal{Z}_t$ is nonempty and $\min \mathcal{Z}_t = \emptyset$. Lemma 6.4 implies $\mathcal{Z}_t$ is an insular family. Since $\mathbb{0} \in A_\emptyset$ and the CSP is extensible, there is an extension σ of $\mathbb{0}$ over $\mathcal{Z}_t$. $\sigma_\downarrow(\emptyset) \stackrel{(7)}{=} \pi_\emptyset \circ \sigma(\emptyset) = \pi_\emptyset(\mathbb{1}_\mathbb{0}) = \mathbb{1}_\mathbb{0}$. Since every d-small $S \subseteq V$ belongs to $(\mathcal{Z}_t)_\downarrow$ (because $S \subseteq V(\text{cl}_S^{2t}(\mathcal{H})) \in \mathcal{Z}_t$ by Lemma 6.12), Lemmas 6.15 and A.1 imply $\sigma_\downarrow$ is a level-d solution. □

The next two lemmas are useful when we consider combined hierarchies later.

Lemma 6.17. *For any extension σ of $q \in A_{\min \mathcal{Z}}$ over an insular family $\mathcal{Z} \subseteq 2^V$, $K \subseteq V$,*

$$\text{supp}(\sigma_{\downarrow K}(S)) \subseteq \pi_S \left(A_{(K \cup S)^\uparrow, q} \right). \tag{9}$$

Proof. For $S \in \mathcal{Z}_{\downarrow K}$,

$$\text{supp}(\sigma_{\downarrow K}(S)) \stackrel{(8)}{=} \text{supp} \circ \pi_S \circ \sigma \left((K \cup S)^\uparrow \right) \subseteq \pi_S \circ \text{supp} \circ \sigma \left((K \cup S)^\uparrow \right) \subseteq \pi_S \left(A_{(K \cup S)^\uparrow, q} \right), \tag{10}$$

where the inclusions are [CNP24, Lemma A.7, Lemma A.6]. □

Lemma 6.18. *For the previous lemma, if additionally $\mathfrak{A}$ is BW or LP, and σ is full, then*

$$\text{supp}(\sigma_{\downarrow K}(S)) = \pi_S \left(A_{(K \cup S)^\uparrow, q} \right) \tag{11}$$

Proof. The first inclusion in Eq. (10) becomes an equality by Eq. (17) when $\mathfrak{A}$ is BW or LP. The last inclusion in Eq. (10) also becomes an equality because σ is full. □

6.3. BW Lower Bound.

We prove BW hierarchy lower bounds in this subsection.

Throughout this subsection, fix $\ell \geqslant 2$ and an ℓ-null-constraining CSP. Consider the BW boundary operator $\mathcal{B}_{S,\ell}(\mathcal{E})$ and BW closure operator $\text{cl}_{S,\ell}(\mathcal{E})$ with parameter ℓ.

Lemma 6.19 (Implicit in [CM13, Lemma 5.8])**.** *Given an instance $I = (V, \mathcal{C})$, for any $S \subseteq V$, $\mathcal{C} \subseteq 2^V$, $\mathcal{Q} \in \mathcal{B}_{S,\ell}(\mathcal{C})$, any satisfying assignment $b \in A_{S \cup V(\mathcal{C} \setminus \mathcal{Q})}$ can be extended to a satisfying assignment $f \in A_I$, i.e. $f_{S \cup V(\mathcal{C} \setminus \mathcal{Q})} = b$.*

Proof. If $\mathcal{Q} = \{C\} \in \mathcal{B}_S^1(\mathcal{C})$, then C has at most one variable u in $S \cup V(\mathcal{C} \setminus \mathcal{Q})$. (If no such u exists, any $u \in V(C)$ will do.) Pretending C is the first constraint in a simple constraint path of length ℓ with endpoint u, the ℓ-null-constraining property yields a satisfying assignment $h \in A_e$ for e such that h and b agree at the intersection of their domains.

If $\mathcal{Q} = \mathcal{C}(P) \in \mathcal{B}_{S,\ell}^2(\mathcal{C})$, then P is a pendant path of length ℓ in $\mathcal{C}$ whose non-endpoints are outside $S \cup V(\mathcal{C} \setminus \mathcal{Q})$. The ℓ-null-constraining property yields a satisfying assignment h for P, such that h and b agree at the intersection of their domains.

In both cases, $f \stackrel{\text{def}}{=} b \cup h$ satisfies $\mathcal{C}$. □

Definition 6.20. Given any instance $J = (V, \mathcal{C})$, $U \subseteq V$, $q \in A_U$, consider the set of satisfying assignments of J whose projection to U equals q,

$$A_{J,q} \stackrel{\text{def}}{=} \{b \in A_J \mid b_U = q\}.$$

Lemma 6.21. *For any instance $J = (V, \mathcal{C})$ with $\mathcal{C} \subseteq \binom{V}{\geqslant 2}$, any $S \subseteq V$, any $\mathcal{Q} \in \mathcal{B}_{S,\ell}(\mathcal{C})$,*

$$\pi_S(A_{J,q}) = \pi_S(A_{J \setminus \mathcal{Q}, q}),$$

where $J \setminus \mathcal{Q} \stackrel{\text{def}}{=} (V, \mathcal{C} \setminus \mathcal{Q})$ is the subinstance of J with $\mathcal{Q}$ removed.

Proof. "$\subseteq$": Because $A_{J,q} \subseteq A_{J\setminus\mathcal{Q},q}$.

"$\supseteq$": Suppose $c = b_S$ for some $b \in A_{J\setminus\mathcal{Q},q}$. Let $T \stackrel{\text{def}}{=} S \cup V(J \setminus \mathcal{Q})$. Then $\mathcal{Q} \in \mathcal{B}_{T,\ell}(\mathcal{C})$ by indifference. Lemma 6.19 yields an extension $f \in A_J$ of b_T such that $f_T = b_T$. In particular $f_S = c$ and $f_U = b_U = q$. ☐

Theorem 6.22. *For $k, \ell \geqslant 2$, any ℓ-null-constraining k-CSP is BW-fully-extensible.*

Proof. Let $I = (V, \mathcal{C})$ be an instance and $\mathcal{Z}$ an insular family. Given $q \in A_U$, where $U \stackrel{\text{def}}{=} \min \mathcal{Z}$, define $\sigma(S) \stackrel{\text{def}}{=} A_{S,q} = \{b \in A_S \mid b_U = q\}$ for $S \in \mathcal{Z}$. Then $\operatorname{supp}(\sigma(S)) = A_{S,q}$ for $S \in \mathcal{Z}$ and $\sigma(U) = \mathbb{1}_q$. The result now follows by applying the following claim to $T, S \in \mathcal{Z}$ with $T \subseteq S$. ☐

Claim 6.23. *If $T \subseteq S$ and T is insular in $\mathcal{C}[S]$, then $\sigma(T) = \pi_T(\sigma(S))$.*

Proof. $\mathrm{cl}_T(\mathcal{C}[S]) = (\mathcal{C}[S])[T] = \mathcal{C}[T]$. Lemma 5.18 yields a strictly decreasing sequence $\mathcal{C}[S] = \mathcal{C}_0 \supsetneq \cdots \supsetneq \mathcal{C}_r = \mathcal{C}[T]$ such that $\mathcal{C}_i \setminus \mathcal{C}_{i+1} \in \mathcal{B}_{T\cup V(\mathcal{C}_{i+1}),\ell}(\mathcal{C}_i) \subseteq \mathcal{B}_{S,\ell}(\mathcal{C}_i)$. Let $J_i \stackrel{\text{def}}{=} (V, \mathcal{C}_i)$ for $0 \leqslant i \leqslant r$. Lemma 6.21 implies $\sigma(T) = \pi_T(A_{J_r,q}) = \cdots = \pi_T(A_{J_0,q}) = \pi_T(\sigma(S))$. ☐

Our next theorem is the BW counterpart of [CM13, Theorem 1.2].

Theorem 6.24. *Let $k \geqslant 2$. Suppose a k-CSP $(D, \mathcal{R})$ has incomplete $\mathrm{cl}(\mathcal{R})$. Then, for any constraint density $\Delta > 0$, with uniformly positive probability, a random instance of the CSP with n variables and Δn constraints has a BW hierarchy solution of level $\Omega(n)$.*

Proof. By Lemma 4.7,[10] we may assume the CSP is ℓ-null-constraining for some ℓ. We may assume $\ell \geqslant 2$, since 1-null-constraining CSPs are also 2-null-constraining.

By Corollary 5.32, there is $\gamma > 0$ and $t = \Omega(n)$ such that with uniformly positive probability, a random instance I of the CSP with n variables and Δn constraints is $(\ell, 4t, \gamma)$-expanding, and hence $(2t, \gamma/\ell)$-confined by Remark 6.9. Theorems 6.16 and 6.22 construct a BW hierarchy solution for I of level $\Omega(n)$. ☐

The linear-level BW lower bound for approximate graph homomorphism [CŽ24b, Corollary 2] follows easily from Theorem 6.24. Consider any non-bipartite graph $G = (V, E)$, which must have a non-bipartite connected component $U \subseteq V$. The induced subgraph $G[U]$ is connected and non-bipartite. As is well known, this is equivalent to the markov chain of the random walk on $G[U]$ being irreducible and aperiodic.[11] Therefore $E|_U$ is null-constraining [Oll02, Corollary 4.1]. Now the linear-level BW lower bound in [CŽ24b, Corollary 2] follows from Theorem 6.24 and Lemma 4.7 and [CNP24, Lemma B.3].

6.4. **LP Lower Bound.**

We prove LP hierarchy lower bounds in this subsection.

Consider a τ-wise uniform CSP. For each constraint C, there is a τ-wise uniform distribution μ_C of satisfying assignments to C. Consider the τ-wse boundary operator $\mathcal{B}_S(\mathcal{E})$ and τ-wise closure operator $\mathrm{cl}_S(\mathcal{E})$.

Previous works used the following distribution to construct LP solutions [BGMT12, KMOW17, CNP24]:

Definition 6.25. Suppose $J = (V, \mathcal{C})$ is an instance. Define the μ_J to be the following distribution over satisfying assignments $b \in A_J$:

(1) Draw b_v uniformly from $D^{\{v\}}$ independently for isolated variable $v \in V \setminus V(\mathcal{C})$; and

[10] As mentioned in Section 4.1, we may assume $\mathcal{R}$ is permutation-closed because the random model permutes variables in $R \in \mathcal{R}$ when imposing R on each constraint.

[11] Aperiodic in the sense of markov chain [Oll02, Definition 4.2]. Not to be confused with a different notion of the same name in [CŽ24b, Definitions 3 and 4].

(2) Draw $b_{V(C)}$ from μ_C independently for $C \in \mathcal{C}$, conditioned on agreeing at their common variables.

We now introduce a variant:

Definition 6.26. Suppose $J = (V, \mathcal{C})$ is an instance, $U \subseteq V$, $q \in A_U$. Define $\mu_{J,q}$ to be the distribution of a random satisfying assignment b from μ_J conditioned on $b_U = q$.

$\mu_{J,q}$ is well-defined if and only if there is $b \in D^V$ such that $b_U = q$ and $b_{V(C)} \in \operatorname{supp}(\mu_C)$ for every $C \in \mathcal{C}$. If $\mu_{J,q}$ is well-defined, then $\operatorname{supp}(\mu_{J,q}) \subseteq A_{J,q}$.

The next lemma generalizes [CNP24, Proposition 10.5].

Lemma 6.27. *For any instance $J = (V, \mathcal{C})$ with $\mathcal{C} \subseteq \binom{V}{>\tau}$, $U \subseteq S \subseteq V$, $q \in A_U$, $\mathcal{Q} \in \mathcal{B}_S(\mathcal{C})$, then $\mu_{J\setminus\mathcal{Q},q}$ is well-defined if and only if $\mu_{J,q}$ is, and*

$$\pi_S(\mu_{J,q}) = \pi_S(\mu_{J\setminus\mathcal{Q},q}). \tag{12}$$

Proof. Denote by $\mathbb{1}_q(b_U)$ the $\{0,1\}$-indicator function of whether $b_U = q$. For every $b \in D^V$,

$$\mu_{J,q}(b) = \frac{p_{J,q}(b)}{\sum_{a\in D^V} p_{J,q}(a)}, \quad \text{where } p_{J,q}(b) \stackrel{\text{def}}{=} \mathbb{1}_q(b_U) \prod_{B\in\mathcal{C}} \mu_B\left(b_{V(B)}\right) \quad \text{for } b \in D^V.$$

The $p_{J,q}(b)$'s are well-defined regardless of whether $\mu_{J,q}$ is. Let $\{C\} \stackrel{\text{def}}{=} \mathcal{Q}$. For every $a \in D^T$,

$$\pi_T(p_{J,q})(a) = \sum_{\substack{b\in D^V\\ b_T=a}} p_{J,q}(b) = \mathbb{1}_q(a_U) \prod_{B\in\mathcal{C}\setminus\mathcal{Q}} \mu_B\left(a_{V(B)}\right) \left[\sum_{\substack{b\in D^V\\ b_T=a}} \mu_C\left(b_{V(C)}\right)\right]. \tag{13}$$

Let $U \stackrel{\text{def}}{=} V(C) \cap T$. Since $\{C\} \in \mathcal{B}_S(\mathcal{C})$, $|U| \leqslant \tau$. The term in square brackets in Eq. (13) equals $\pi_U(\mu_C)(a_U) = 1/|D^U|$ since μ_C is τ-wise uniform. So that term is independent of a. Therefore

$$\pi_T(p_{J,q})(a) = \frac{\mathbb{1}_q(a_U)}{|D^U|} \prod_{B\in\mathcal{C}\setminus\mathcal{Q}} \mu_B\left(a_{V(B)}\right).$$

On the other hand, for $a \in D^T$,

$$\pi_T(p_{J\setminus\mathcal{Q},q})(a) = \sum_{\substack{b\in D^V\\ b_T=a}} p_{J\setminus\mathcal{Q},q}(b) = \left|D^{V\setminus T}\right| \cdot \mathbb{1}_q(a_U) \prod_{B\in\mathcal{C}\setminus\mathcal{Q}} \mu_B\left(a_{V(B)}\right) = \left|D^{V(C)}\right| \pi_T(p_{J,q}(a)). \tag{14}$$

$\mu_{J\setminus\mathcal{Q},q}$ is well-defined iff $\sum_{b\in D^V} p_{J\setminus\mathcal{Q},q}(b) > 0$, iff $\sum_{a\in D^T} \pi_T(p_{J\setminus\mathcal{Q},q})(a) > 0$, iff $\sum_{a\in D^T} \pi_T(p_{J,q})(a) > 0$, iff $\sum_{b\in D^V} p_{J,q}(b) > 0$, iff $\mu_{J,q}$ is well-defined.

When $\mu_{J\setminus\mathcal{Q},q}$ or $\mu_{J,q}$ is well-defined, Eq. (14) implies $\pi_T(p_{J\setminus\mathcal{Q},q})$ is proportional to $\pi_T(p_{J,q})$. Since both functions sum to one, they are identical. Eq. (12) follows by further projecting to S. □

Theorem 6.28. *For $k > \tau$, any τ-wise uniform k-CSP is LP-extensible.*

Proof. Let $I = (V, \mathcal{C})$ be an instance and $\mathcal{Z}$ an insular family. Given $q \in A_U$, where $U \stackrel{\text{def}}{=} \min \mathcal{Z}$, define $\sigma(S) \stackrel{\text{def}}{=} \mu_{I[S]\setminus\mathcal{C}[U],q}$ for $S \in \mathcal{Z}$. Note that $A_{I[S]\setminus\mathcal{C}[U],q} = A_{I[S],q} = A_{S,q}$ because every $b \in D^V$ with $b_U = q$ satisfies $\mathcal{C}[S]$ if and only if it satisfies $\mathcal{C}[S] \setminus \mathcal{C}[U]$, knowing that q already satisfies $\mathcal{C}[U]$. Therefore $\operatorname{supp}(\sigma(S)) = \operatorname{supp}(\mu_{I[S]\setminus\mathcal{C}[U],q}) \subseteq A_{I[S]\setminus\mathcal{C}[U],q} = A_{S,q}$ for $S \in \mathcal{Z}$. Also $\sigma(U) = \mathbb{1}_q$ as $\mathcal{C}[U] \setminus \mathcal{C}[U] = \emptyset$. Since $U \subseteq S$ is insular in $\mathcal{C}[S]$ for every $S \in \mathcal{Z}$, Claim 6.29 implies $\sigma(S)$ is well-defined. Claim 6.29 also implies σ is consistent with projection for $\mathcal{Z}$. □

Claim 6.29. *If $T \subseteq S$, T is insular in $\mathcal{C}[S]$ and $\sigma(T)$ is well-defined, then $\sigma(S)$ is also well-defined and $\sigma(T) = \pi_T(\sigma(S))$.*

Proof. $\mathrm{cl}_T(\mathcal{C}[S]) = (\mathcal{C}[S])[T] = \mathcal{C}[T]$. Lemma 5.18 yields a strictly decreasing sequence $\mathcal{C}[S] = \mathcal{C}_0 \supsetneq \cdots \supsetneq \mathcal{C}_r = \mathcal{C}[T]$ such that $\mathcal{C}_i \setminus \mathcal{C}_{i+1} \in \mathcal{B}_{T\cup V(\mathcal{C}_{i+1}),\ell}(\mathcal{C}_i) \subseteq \mathcal{B}_{S,\ell}(\mathcal{C}_i)$. Let $J_i \stackrel{\text{def}}{=} (V, \mathcal{C}_i)$ for $0 \leqslant i \leqslant r$. Lemma 6.27 implies $\sigma(T) = \pi_T(\mu_{J_r,q}) = \cdots = \pi_T(\mu_{J_0,q}) = \pi_T(\sigma(S))$ and every $\mu_{J_i,q}$ is well-defined. □

For our C(BLP+AIP) lower bound, the LP extension has to be full. We are unable to upgrade Theorem 6.28 to "fully-extensible". Fortunately a weaker version suffices: the extension is full on small subsets (Theorem 6.34).

Given a family $\mathcal{Z} \subseteq 2^V$ and $d \geqslant 0$, let $\mathcal{Z}_{\leqslant d} \stackrel{\text{def}}{=} \mathcal{Z} \cap \binom{V}{\leqslant d}$.

Lemma 6.30. *$T^\uparrow \subseteq S^\uparrow$ for any $T \subseteq S \in \mathcal{Z}_\downarrow$.*

Proof. $T \subseteq S \subseteq S^\uparrow$ by definition of $S^\uparrow$. Since $S^\uparrow \in \mathcal{Z}$, $T^\uparrow \subseteq S^\uparrow$ by definition of $T^\uparrow$. □

Lemma 6.31. *$(S^\uparrow)^\uparrow = S^\uparrow$ for any $S \in \mathcal{Z}_\downarrow$.*

Proof. $S^\uparrow \subseteq S^\uparrow \in \mathcal{Z}$, so $(S^\uparrow)^\uparrow \subseteq S^\uparrow$ by definition of $(S^\uparrow)^\uparrow$. $S \subseteq S^\uparrow \in \mathcal{Z}$ by definition of $S^\uparrow$, and hence $S^\uparrow \subseteq (S^\uparrow)^\uparrow$ by Lemma 6.30. □

Lemma 6.32. *$(S \cup T)^\uparrow = (S^\uparrow \cup T^\uparrow)^\uparrow$ for any $S, T \subseteq V$ such that $S \cup T \in \mathcal{Z}_\downarrow$.*

Proof. $S \subseteq S \cup T$ implies $S^\uparrow \subseteq (S \cup T)^\uparrow$ by Lemma 6.30. Likewise $T^\uparrow \subseteq (S \cup T)^\uparrow$. Thus $S^\uparrow \cup T^\uparrow \subseteq (S \cup T)^\uparrow$. Therefore $(S^\uparrow \cup T^\uparrow)^\uparrow \subseteq ((S \cup T)^\uparrow)^\uparrow = (S \cup T)^\uparrow$ using Lemmas 6.30 and 6.31. This also implies $S^\uparrow \cup T^\uparrow \in \mathcal{Z}_\downarrow$.

Definition of $S^\uparrow$ implies $S \subseteq S^\uparrow \subseteq S^\uparrow \cup T^\uparrow$. Likewise $T^\uparrow \subseteq S^\uparrow \cup T^\uparrow$. Therefore $S \cup T \subseteq S^\uparrow \cup T^\uparrow$. Then $(S \cup T)^\uparrow \subseteq (S^\uparrow \cup T^\uparrow)^\uparrow$ by Lemma 6.30. □

Given $d \geqslant 0$, let $V_d \stackrel{\text{def}}{=} \binom{V}{\leqslant d}$. Given a family $\mathcal{Z} \subseteq 2^V$, let $\mathcal{Z}_{\leqslant d} \stackrel{\text{def}}{=} \left\{R^\uparrow \mid R \in \mathcal{Z}_\downarrow \cap V_d\right\}$.

Lemma 6.33. *If $V_{2d} \subseteq \mathcal{Z}_\downarrow$, then $\mathcal{Z}_{\leqslant 2d} = \left\{(S \cup T)^\uparrow \mid S, T \in \mathcal{Z}_{\leqslant d}\right\}$.*

Proof. "$\supseteq$": For $S, T \in \mathcal{Z}_{\leqslant d}$, $S = P^\uparrow$ and $T = Q^\uparrow$ for some $P, Q \in \mathcal{Z}_\downarrow \cap V_d$. Then $P \cup Q \in V_{2d} \subseteq \mathcal{Z}_\downarrow$, and $\mathcal{Z}_{\leqslant 2d} \ni (P \cup Q)^\uparrow = (P^\uparrow \cup Q^\uparrow)^\uparrow = (S \cup T)^\uparrow$, where the first equality is Lemma 6.32.

"$\subseteq$": Given $S \in \mathcal{Z}_{\leqslant 2d}$, $S = T^\uparrow$ for some $T \in \mathcal{Z}_\downarrow \cap V_{2d}$. Take any subset $R \subseteq T$ of size $\max(d, |T|)$. Then R and $T \setminus R$ both belong to $\mathcal{Z}_\downarrow \cap V_d$. Therefore $S = T^\uparrow = (R \cup (T \setminus R))^\uparrow = (R^\uparrow \cup (T \setminus R)^\uparrow)^\uparrow \in \left\{(X \cup Y)^\uparrow \mid X, Y \in \mathcal{Z}_{\leqslant d}\right\}$ by Lemma 6.32. □

Theorem 6.34. *Given $k > \tau \geqslant 0$, $d \geqslant 0$, any τ-wise uniform k-CSP, its instance $I = (V, \mathcal{C})$, insular family $\mathcal{Z}$ such that $V_{2d} \subseteq \mathcal{Z}_\downarrow$, any $q \in A_{\min \mathcal{Z}}$, there is a full LP-extension of q over $\mathcal{Z}_{\leqslant d}$.*

Proof. For $K \in \mathcal{Z}_{\leqslant d}$, $b \in A_{K,q}$, Theorem 6.28 yields an extension σ_b of b over the insular subfamily $\mathcal{Z}_{\supseteq K} \stackrel{\text{def}}{=} \{R \in \mathcal{Z} \mid R \supseteq K\}$. Then $\mathrm{supp}(\sigma_b(R)) \subseteq A_{R,b}$ for every $R \in \mathcal{Z}_{\supseteq K}$. Also $\sigma_b(K) = \mathbb{1}_b$.

Lemma 6.33 implies $(R \cup K)^\uparrow \in \mathcal{Z}_{\leqslant 2d}$ for $R \in \mathcal{Z}_{\leqslant d}$, and hence $(R \cup K)^\uparrow \in \mathcal{Z}_{\leqslant 2d} \cap Z_{\supseteq K}$. Define σ'_b to be the restriction of $(\sigma_b)_{\downarrow K}$ to $\mathcal{Z}_{\leqslant d}$; see Eq. (8). σ'_b is a scheme by Lemma 6.15. And $\sigma'_b(U) = \pi_U \circ \sigma_b(K) = \pi_U(\mathbb{1}_b) = \mathbb{1}_q$. Also for $S \in \mathcal{Z}_{\leqslant d}$, Eq. (9) implies $\mathrm{supp}(\sigma'_b(S)) \subseteq \pi_S\left(A_{(K\cup S)^\uparrow,q}\right) \subseteq A_{S,q}$.

Consider the average $\sigma \stackrel{\text{def}}{=} \mathbb{E}_{K,b}[\sigma'_b]$ over $K \in \mathcal{Z}_{\leqslant d}$ and $b \in A_{K,q}$. σ is a scheme for $\mathcal{Z}_{\leqslant d}$, as every σ'_b is. And $\sigma(U) = \mathbb{1}_q$, because $\sigma'_b(U) = \mathbb{1}_q$ for every K and b. Also for $S \in \mathcal{Z}_{\leqslant d}$, $\mathrm{supp}(\sigma(S)) = \bigcup_{K,b} \mathrm{supp}(\sigma'_b(S)) \subseteq A_{S,q}$. Finally, for $S \in \mathcal{Z}_{\leqslant d}$, $b \in A_{S,q}$, since $\sigma_b(S) = \mathbb{1}_b$, we have $b \in \mathrm{supp}(\sigma'_b(S)) \subseteq \mathrm{supp}(\sigma(S))$. Therefore $\mathrm{supp}(\sigma(S)) = A_{S,q}$ for $S \in \mathcal{Z}_{\leqslant d}$. □

7. AIP Lower Bound

We prove AIP hierarchy lower bound in this section.

7.1. **Path-Closed.**

As outlined in Section 1.1, a key step in constructing AIP solutions is to argue that the neighborhood $B(T,1)$ of an insular subset T is also insular. For this purpose, we introduce a related concept of path-closedness, because path-closedness behaves better with neighborhood and distance.

Recall $\mathcal{B}_{S,\ell}$ and $\mathrm{cl}_{S,\ell}(\mathcal{E})$ denote the BW boundary and closure operators with parameter ℓ. Given $\ell \in \mathbb{N}$ and a hyperedge set $\mathcal{E} \subseteq 2^V$, a vertex subset $S \subseteq V$ is ℓ-insular in $\mathcal{E}$ if it is insular in $\mathcal{E}$ under $\mathrm{cl}_{S,\ell}(\mathcal{E})$, i.e. $\mathrm{cl}_{S,\ell}(\mathcal{E}) \subseteq \mathcal{E}[S]$.

Definition 7.1 (Path-closed)**.** Given $\ell \in \mathbb{N}$ and a hyperedge set $\mathcal{E} \subseteq 2^V$, a vertex subset $S \subseteq V$ is ℓ-path-closed in $\mathcal{E}$ if for every $u, v \in S$ and every simple path P in $\mathcal{E}$ of length at most ℓ with endpoints u and v, we have $V(P) \subseteq S$.

We will show that a $(3\ell+3)$-insular subset T is $(3\ell+2)$-path-closed, and its neighborhood $B(T,1)$ is 3ℓ-path-closed, which in turns implies $B(T,1)$ is ℓ-insular under additional mild assumptions ("ℓ-nice" hyperedge set).

Lemma 7.2. *Let $\ell \geqslant 2$. If $S \subseteq V$ is ℓ-insular in $\mathcal{E}$, then S is $(\ell-1)$-path-closed in $\mathcal{E}$.*

Proof. Let P be a simple path in $\mathcal{E}$ of length at most $(\ell-1)$ from $u \in S$ to $v \in S$. We prove by induction on the length of P that $V(P) \subseteq S$.

Base Case: P has length 0. Then $V(P) = \{u\} = \{v\} \subseteq S$.

Induction Case 1: No vertex in P other than the endpoints u and v belong to S. Then $\mathcal{E}(P)$ is (S,ℓ)-closed. Thus $\mathcal{E}(P) \subseteq \mathrm{cl}_{S,\ell}(\mathcal{E}) \subseteq \mathcal{E}[S]$, where the first inclusion is Lemma 5.4 and the second the definition of ℓ-insular. Therefore $V(P) = V(\mathcal{E}(P)) \subseteq S$.

Induction Case 2: Some vertex $w \notin \{u, v\}$ in P belongs to S.

If w shares a hyperedge e in P with u (i.e. $e \supseteq \{w, u\}$), then $\{e\}$ is (S,ℓ)-closed (using $\ell \geqslant 2$), so $\{e\} \in \mathrm{cl}_{S,\ell}(\mathcal{E}) \subseteq \mathcal{E}[S]$, i.e. $e \subseteq S$. In particular, the connecting variable of P in e other than u also belongs to S. Then $Q \stackrel{\text{def}}{=} P \setminus \{e\}$ is a shorter simple path of length at most $(\ell-1)$ with both endpoints in S. Induction Hypothesis implies $V(Q) \subseteq S$ as well. Therefore $V(P) = e \cup V(Q) \subseteq S$.

If w shares a constraint in P with v, the same conclusion $V(P) \subseteq S$ holds, because u and v have symmetric roles.

If w does not share any constraint in P with u or v, apply the Induction Hypothesis to the shorter simple subpath Q of P from u to w, as well as to the shorter simple subpath R of P from w to v, we get $V(P) = V(Q) \cup V(R) \subseteq S$. ☐

Lemma 7.3. *For any $\ell, m \in \mathbb{N}$, if S is $(\ell+2m)$-path-closed in $\mathcal{E}$ with $\mathrm{girth}(\mathcal{E}) > \ell + 2m$, then $B(S,m)$ is ℓ-path-closed in $\mathcal{E}$.*

Proof. Let P be a simple path in $\mathcal{E}$ from $u \in B(S,m)$ to $v \in B(S,m)$ of length at most ℓ. There is a Berge path P_u of length at most m from $u' \in S$ to u, and a Berge path P_v of length at most m from v to $v' \in S$.

Consider the walk $Q \stackrel{\text{def}}{=} P_u \cup P \cup P_v$ of length at most $\ell + 2m$. For every vertex w in P_u other than u and u', $\deg_{P_u}(w) = \deg_Q(w)$ because $\mathrm{girth}(\mathcal{E}) > \ell + 2m$. Likewise for every vertex w in P_v other than v and v', $\deg_{P_v}(w) = \deg_Q(w)$. If $u \neq u'$, $\deg_Q(u) = 2$ again because $\mathrm{girth}(\mathcal{E}) > \ell + 2m$. If $v \neq v'$, $\deg_Q(v) = 2$. Therefore Q is a simple path.

Since S is $(\ell+2m)$-path-closed, $V(P) \subseteq V(Q) \subseteq S \subseteq B(S,m)$. ☐

Definition 7.4 (Nice)**.** Let $\ell \geqslant 2$. A hyperedge set $\mathcal{E} \subseteq 2^V$ is ℓ-nice if $\mathrm{girth}(\mathcal{E}) > 3\ell$ and every 3ℓ-path-closed $S \subseteq V$ in $\mathcal{E}$ is ℓ-insular in $\mathcal{E}$.

Lemma 7.5. *Given a boundary operator $\mathcal{B}$ for $(V, \mathcal{H})$, $\mathcal{E}$ is S-closed if and only if $\mathcal{E} \setminus \mathcal{E}[S]$ is, for $S \subseteq V$, $\mathcal{E} \subseteq \mathcal{H}$.*

Proof. $\mathcal{B}_S(\mathcal{E}) = \mathcal{B}_S(\mathcal{E} \setminus \mathcal{E}[S])$ by independence of $\mathcal{E}[S]$. The result follows by Remark 5.16. ☐

Theorem 7.6. *If a k-uniform hyperedge set $\mathcal{E} \subseteq 2^V$ is hereditarily δ-sparse with $\delta \stackrel{\text{def}}{=} \delta(k, 2\ell - 1)$ given by Lemma 5.26, and* $\mathrm{girth}(\mathcal{E}) > 3\ell$*, then it is ℓ-nice.*

Proof. Suppose $S \subseteq V$ is 3ℓ-path-closed. We will show that S is ℓ-insular in $\mathcal{E}$. If $\mathcal{G} \subseteq \mathcal{E}$ is (S, ℓ)-closed, then so is $\mathcal{F} \stackrel{\text{def}}{=} \mathcal{G} \setminus \mathcal{G}[S]$ by Lemma 7.5.

Call a pendant path S-external if its non-endpoints are outside S. Let $\mathcal{P}$ be the collection of maximal S-external pendant paths P in $\mathcal{F}$ of length at least 1. Here maximal means that P is not contained in a strictly longer S-external pendant path. Then every endpoint v of P must satisfy (1) $v \in S$; or (2) $\deg_{\mathcal{F}}(v) \geqslant 3$; or (3) v belongs to a hyperedge $e \in \mathcal{F} \setminus \mathcal{E}(P)$ such that e contains at least three vertices u with $\deg_{\mathcal{F}}(u) \geqslant 2$. The possibility $\deg_{\mathcal{F}}(v) = 1$ is ruled out, since $\mathcal{F}$ is (S, ℓ)-closed and cannot contain $\{e\} \in \mathcal{B}^1_S(\mathcal{E})$. Also P has length less than ℓ because $\mathcal{F}$ is (S, ℓ)-closed.

Let $\mathcal{K}$ be obtained from $\mathcal{F}$ by removing all hyperedges of every S-external pendant path P in $\mathcal{F}$ having an endpoint in S. Call a vertex u special if u belongs to S or $\deg_{\mathcal{K}}(u) < \deg_{\mathcal{F}}(u)$.

Claim 7.7. *Let u and w be distinct special vertices. Then* $\mathrm{dist}_{\mathcal{F}}(u, w) \geqslant \ell$*.*

Proof. u is connected to some vertex in S via an S-external pendant path P_u in $\mathcal{F}$ of length less than ℓ. (P_u may have length zero.) Likewise w is connected to some vertex in S via an S-external pendant path P_v in $\mathcal{F}$. If there is a Berge path P_{uw} in $\mathcal{F}$ between u and w of length less than ℓ, then $Q \stackrel{\text{def}}{=} P_u \cup P_{uw} \cup P_w$ is a Berge path of length less than 3ℓ. Since $\mathrm{girth}(\mathcal{E}) > 3\ell$, Q is a simple path in $\mathcal{F}$ from S to S. Since S is 3ℓ-path-closed, $V(Q) \subseteq S$. Since P_u, P_{uw}, and P_w involve only hyperedges in $\mathcal{F} = \mathcal{G} \setminus \mathcal{G}[S]$, none of these paths have any hyperedge, so $u = w$. Since u and w are distinct by assumption, there cannot be any Berge path of length less than ℓ between them. □

Claim 7.8. *Any pendant path P in $\mathcal{K}$ has length at most $2\ell - 2$.*

Proof. If every non-endpoint in P is not special, then P is an S-external pendant path in $\mathcal{F}$, and P has length less than ℓ because $\mathcal{F}$ is (S, ℓ)-closed and a subpath of length ℓ in P would be in $\mathcal{B}^2_{S,\ell}(\mathcal{F})$.

Otherwise P contains a special non-endpoint u. We will show that P has at most one such vertex. Indeed, if w is another special non-endpoint in P, then $\mathrm{dist}_{\mathcal{F}}(u, w) \geqslant \ell$ by Claim 7.7. That means the shortest path P'_{uw} between u and w in P is an S-external pendant path in $\mathcal{F}$ of length at least ℓ, as vertices in P'_{uw} apart from u and w have the same degree in $\mathcal{K}$ and in $\mathcal{F}$ and are outside S. But P'_{uw} cannot exist because $\mathcal{F}$ is (S, ℓ)-closed. Therefore at most one such u exists.

The shortest path P_{uv} in P from u to an endpoint v of P is an S-external pendant path in $\mathcal{F}$, again because all vertices in P_{uv} other than u and v have the same degree in $\mathcal{K}$ and in $\mathcal{F}$ and are outside S. Therefore P_{uv} has length less than ℓ. Since u has distance (in P) at most $\ell - 1$ from any endpoint of P, P has length at most $2\ell - 2$. □

Claim 7.9. *No hyperedge $e \in \mathcal{K}$ has at most one vertex u with* $\deg_{\mathcal{K}}(u) > 1$.

Proof. e cannot contain two special vertices by Claim 7.7.

If e has no special vertex, then having at most one vertex u with $\deg_{\mathcal{K}}(u) = \deg_{\mathcal{F}}(u) > 1$ implies $e \in \mathcal{B}^1_S(\mathcal{F})$, and $\mathcal{F}$ cannot be (S, ℓ)-closed.

The remaining case is e contains exactly one special vertex u. There is an S-external pendant path P_u in $\mathcal{F}$ of connecting u and S (P_u has length zero if $u \in S$). Since $e \in \mathcal{K}$, $e \setminus \{u\}$ has at least two vertices of degree at least 2 in $\mathcal{F}$, for otherwise $P_u \cup \{e\}$ is an S-external pendant path containing an endpoint in S and is thus removed from $\mathcal{K}$. Since vertices in $e \setminus \{u\}$ are not special, they have the same degree in $\mathcal{K}$ and in $\mathcal{F}$. Therefore there are at least two vertices in $e \setminus \{u\}$ whose degree in $\mathcal{K}$ is larger than 1. □

The previous two claims imply $\mathcal{B}_{2\ell-1}(\mathcal{K})$ is empty. Since $\mathcal{E}$ is hereditarily δ-sparse, $\mathcal{K}$ is δ-sparse, and hence $\mathcal{K}$ is empty by the next claim and Lemma 5.26.

Claim 7.10. *$\mathcal{K}$ does not have any simple cycle component.*

Proof. Suppose $\mathcal{K}'$ is a simple cycle component in $\mathcal{K}$. $\mathcal{K}'$ has length more than 3ℓ because $\mathrm{girth}(\mathcal{E}) > 3\ell$. Take any simple subpath P of $\mathcal{K}'$ of length exactly 3ℓ. P is pendant in $\mathcal{K}$ because $\mathcal{K}'$ is a simple cycle component. This contradicts Claim 7.8. ☐

Finally, $\mathcal{F}$ is also empty by the next claim. Thus $\mathcal{G} \subseteq \mathcal{E}[S]$.

Claim 7.11. *If $\mathcal{F}$ is nonempty, then so is $\mathcal{K}$.*

Proof. Since $\mathcal{E}$ is (∞, δ)-sparse and $\mathrm{girth}(\mathcal{E}) > 3\ell$, it is (ℓ, ∞, γ)-expanding for some $\gamma > 0$ by Lemma 5.28, and $(\infty, \gamma/\ell)$-confined by Remark 6.9. Lemmas 5.4 and 6.11 imply $\mathrm{cl}_{\emptyset,\ell}(\mathcal{E}) = \emptyset$. Then $\mathrm{cl}_{\emptyset,\ell}(\mathcal{F}) \subseteq \mathrm{cl}_{\emptyset,\ell}(\mathcal{E}) = \emptyset$ by monotonicity in $\mathcal{E}$. Since $\mathcal{F}$ is nonempty, some $\mathcal{Q} \in \mathcal{B}_{\emptyset,\ell}(\mathcal{F})$ exists. Since $\mathcal{F}$ is (S, ℓ)-closed, $\mathcal{Q} \notin \mathcal{B}_{S,\ell}(\mathcal{F})$, and $v \in V(\mathcal{Q}) \cap S \subseteq V(\mathcal{F}) \cap S$ exists.

Take any $e \in \mathcal{F}$ containing $v \in S$. Since $\mathcal{F}$ is (S, ℓ)-closed, e must have at least one vertex $u \neq v$ with $\deg_{\mathcal{F}}(u) > 1$ ($u \notin S$ by Claim 7.7). If e has another vertex $w \notin \{u, v\}$ with $\deg_{\mathcal{F}}(w) > 1$, then $e \in \mathcal{K}$. Otherwise e has no such w, and e is part of an S-external pendant path in $\mathcal{F}$ of length at least 1 with endpoint v. Take P to be the maximal such path. As argued earlier, the other endpoint x of P must satisfy (1) $x \in S$; or (2) $\deg_{\mathcal{F}}(x) \geqslant 3$; or (3) x belongs to a hyperedge $f \in \mathcal{F} \setminus \mathcal{E}(P)$ such that f contains at least three vertices y with $\deg_{\mathcal{F}}(y) \geqslant 2$. And P must have length less than ℓ. Possibility (1) is ruled out by Claim 7.7. Since $\deg_{\mathcal{F}}(x) \geqslant 2$ in possibilites (2) and (3), x is contained in some hyperedge $g \in \mathcal{F} \setminus \mathcal{E}(P)$. We have $g \in \mathcal{K}$ as desired, for otherwise $g \in \mathcal{E}(Q)$ for some $Q \in \mathcal{P}$ and $P \cup Q$ is a Berge path (and hence a simple path since $\mathrm{girth}(\mathcal{E}) > 3\ell$) in $\mathcal{F} = \mathcal{G} \setminus \mathcal{G}[S]$ of length less than 2ℓ from S to S, contradicting S being 3ℓ-path-closed. ☐

☐

7.2. **AIP Scheme.**

Given an ℓ-null-constraining and lax CSP, let $\ell' \overset{\text{def}}{=} 3\ell + 3$. Recall that for instance $J = (V, \mathcal{C})$, $U \subseteq V$, $q \in A_U$, $A_{J,q} \overset{\text{def}}{=} \{b \in A_J \mid b_U = q\}$ is the set of satisfying assignments to J agreeing with q.

Definition 7.12. Given an instance $J = (V, \mathcal{C})$, ℓ'-insular family $\mathcal{Z}$, $\min \mathcal{Z} \subseteq S \subseteq V$, $q \in A_{\min \mathcal{Z}}$. Define

$$X_S \overset{\text{def}}{=} \{w \in \mathbb{Z}^{A_{S,q}} \mid \pi_T(w) = 0 \text{ for } T \in \mathcal{Z}[S] \setminus \{S\}\}$$

to be the subgroup of $\mathbb{Z}^{A_{S,q}}$ having zero projection to every proper subset T of S in $\mathcal{Z}$.

Given a hypergraph $H = (V, \mathcal{E})$, an insular family $\mathcal{Z}$ is ℓ-nice if $\mathcal{E}[S]$ is ℓ-nice for every $S \in \mathcal{Z}$.

Theorem 7.13. *Let $k, \ell \geqslant 2$. Suppose a k-CSP $(D, \mathcal{R})$ is ℓ-null-constraining and lax, and $J = (V, \mathcal{C})$ is its instance. Given any ℓ-nice ℓ'-insular family $\mathcal{Z}$ in J with $V \in \mathcal{Z}$, any $S \in \mathcal{Z}$, $q \in A_{\min \mathcal{Z}}$, $w \in X_S$, there is $x \in \mathbb{Z}^{A_{V,q}}$ such that $\pi_S(x) = w$ and $\pi_T(x) = 0$ whenever $T \in \mathcal{Z}$ and $T \not\supseteq S$.*

Proof. Consider $B(S, 1) = \{u \in V \mid \mathrm{dist}(u, S) \leqslant 1\}$ the 1-neighborhood of S in J. We first extend w to some weight $y \in \mathbb{Z}^{D^{B(S,1)}}$ on $B(S, 1)$.

Let $Z \overset{\text{def}}{=} B(S, 1) \setminus S$ be vertices in the neighborhood and outside S. Every $u \in Z$ belongs to a unique constraint $C_u \in \mathcal{C}$ such that $C_u \cap S \neq \emptyset$ (and in fact $|C_u \cap S| = 1$). Indeed, $\mathrm{dist}(u, S) = 1$, some C_u exists such that $V(C_u) \cap S \neq \emptyset$. Since S is ℓ'-insular in $\mathcal{E}[V] = \mathcal{E}$, S is $(3\ell + 2)$-path-closed in $\mathcal{E}$ by Lemma 7.2, so such a C_u is unique, and also $|C_u \cap S| = 1$.

Consider the collection $\mathcal{C}' \overset{\text{def}}{=} \{C_u \mid u \in Z\}$ of such constraints. Since the CSP is lax, every $C \in \mathcal{C}'$ has an assignment $b^C \in D^{V(C) \setminus S}$ to its variables outside S such that $D^{v_C} \times \{b^C\} \subseteq A_C$, where $\{v_C\} \overset{\text{def}}{=} V(C) \cap S$.

Let

$$y \overset{\text{def}}{=} w \otimes \bigotimes_{C \in \mathcal{C}'} \mathbb{1}_{b^C} \in \mathbb{Z}^{D^{B(S,1)}}.$$

In fact, y is supported on *satisfying* assignments $\beta \in A_{J[B(S,1)],q}$ of the subinstance $J[B(S, 1)]$ induced by the 1-neighborhood, such that $\beta_U = q$. Indeed, every $\alpha \in \mathrm{supp}(y)$ has $\alpha_S \in \mathrm{supp}(w) \subseteq A_S$, so α

satisfies every constraint $C \in \mathcal{C}[S]$ induced on S. Every constraint $C \in \mathcal{C}[B(S,1)] \setminus \mathcal{C}[S]$ induced on $B(S,1)$ but not induced on S has all its vertices in $B(S,1)$. C cannot contain only vertices in Z, for otherwise (since $k \geqslant 2$) C contains distinct $u, v \in V(C) \cap Z$, and (C_u, C, C_v) forms a walk (and hence a simple path because $\mathrm{girth}(J) > 3$) of length 3 from S to S, contradicting S being $(3\ell+2)$-path-closed. Therefore $C \in \mathcal{C}[B(S,1)] \setminus \mathcal{C}[S]$ contains some vertex in S and some vertex in Z, so $C \in \mathcal{C}'$. This implies $\alpha_{V(C)} \in D^{v_C} \times \{b^C\} \subseteq A_C$.

Further, y is supported on assignments α such that $\alpha_{C\cap Z} = b^C$ for $C \in \mathcal{C}'$. In other words, assignments α whose restriction to Z is completely determined by the lax property and does not depend on w.

S is $(3\ell+2)$-path-closed in $\mathcal{E}$, so $B(S,1)$ is 3ℓ-path-closed in $\mathcal{E}$ by Lemma 7.3 and hence ℓ-insular in $\mathcal{E}$ as $\mathcal{E}$ is ℓ-nice. Any $\alpha \in \mathrm{supp}(y)$ can be extended to $\alpha' \in A_{J,q}$ by Claim 6.23. Thanks to the property of the previous paragraph, given any $\beta \in \mathrm{supp}(y)$, if we replace α' with β in S, the new assignment $\beta \cup \alpha'_{V(J)\setminus S} \in A_{J,q}$ also satisfies J. This is because a constraint C of J either contains some variable in S (and C belongs to $J[B(S,1)]$, and is thus satisfied by β) or does not (and C is satisfied by $\alpha'_{V(J)\setminus S}$).

Define

$$x \stackrel{\text{def}}{=} y \otimes \mathbb{1}_{\alpha'_{V(J)\setminus B(S,1)}}.$$

The previous paragraph implies x is supported on $A_{J,q}$. x extends w, because $\pi_S(x) = w$ by Eq. (16). Also, if $T \in \mathcal{Z}$ and $T \not\supseteq S$, then $\pi_T(x) = \pi_{T\cap S}(w) \otimes \mathbb{1}_{\alpha'_{T\setminus S}}$ by Eq. (16). $T \cap S$ is ℓ'-insular in $\mathcal{E}$ by Lemma 6.2, so $T \cap S \in \mathcal{Z}[S] \setminus \{S\}$. Since $w \in X_S$, $\pi_{T\cap S}(w) = 0$ and hence $\pi_T(x) = 0$. ☐

Given a family $\mathcal{Z} \subseteq 2^V$, an element $S \in \mathcal{Z}$ is maximal if $S \subseteq T$ and $T \in \mathcal{Z}$ imply $S = T$, and minimal if $T \subseteq S$ and $T \in \mathcal{Z}$ imply $S = T$.

Given a family $\mathcal{Z} \subseteq 2^V$, a subfamily $\mathcal{T} \subseteq \mathcal{Z}$ is a down set of $\mathcal{Z}$ (also called downward-closed family) if $A \in \mathcal{T}, B \subseteq A$ and $B \in \mathcal{Z}$ imply $B \in \mathcal{T}$.

Lemma 7.14. *If $\mathcal{Z}$ is an insular family and $\mathcal{T}$ is a down set of $\mathcal{Z}$, then $\mathcal{T}$ is also an insular family.*

Proof. It suffices to verify $\mathcal{T}$ is closed under intersection. Given $S, T \in \mathcal{T} \subseteq \mathcal{Z}$, $S \cap T \in \mathcal{Z}$ as $\mathcal{Z}$ is closed under intersection. Since $S \cap T \subseteq S$, $S \cap T \in \mathcal{T}$ as $\mathcal{T}$ is a down set of $\mathcal{Z}$. ☐

Theorem 7.15. *Let $k, \ell \geqslant 2$. Suppose a k-CSP $(D, \mathcal{R})$ is ℓ-null-constraining and lax, $J = (V, \mathcal{C})$ its instance, and $\mathcal{Z}$ an ℓ-nice ℓ'-insular family in J with $V \in \mathcal{Z}$. Given any extension σ of $q \in A_{\min \mathcal{Z}}$ for $\mathcal{Z}' \stackrel{\text{def}}{=} \mathcal{Z} \setminus \{V\}$, there is an extension σ' of q for $\mathcal{Z}$ such that $\sigma'(T) = \sigma(T)$ for $T \in \mathcal{Z}'$.*

Proof. Given any $w \in \mathbb{Z}^{A_{V,q}}$, let $\mathcal{T}_w \stackrel{\text{def}}{=} \{R \in \mathcal{Z} \mid \pi_T(w) = \sigma(T) \text{ for every } T \in \mathcal{Z}[R] \setminus \{R\}\}$. The previous lemma implies $\mathcal{T}_w$ is a down set of $\mathcal{Z}$, and hence an ℓ'-insular subfamily of $\mathcal{Z}$.

Claim 7.16. *If $\mathcal{T}_w \subsetneq \mathcal{Z}$ for some $w \in \mathbb{Z}^{A_{V,q}}$, then there is $w' \in \mathbb{Z}^{A_{V,q}}$ such that $\mathcal{T}_{w'} \supsetneq \mathcal{T}_w$.*

Proof. Since $\mathcal{T}_w \subsetneq \mathcal{Z}$, pick any minimal $S \in \mathcal{Z} \setminus \mathcal{T}_w$. Then $\pi_S(w) - \sigma(S) \in X_S$ because $\pi_T(\pi_S(w) - \sigma(S)) = \pi_T(w) - \pi_T(\sigma(S)) = \pi_T(w) - \sigma(T) = 0$ for $T \subsetneq S, T \in \mathcal{Z}$. Theorem 7.13 yields $x \in A^{V,q}$ such that $\pi_S(x) = \pi_S(w) - \sigma(S)$ and $\pi_T(x) = 0$ for $T \in \mathcal{Z}, T \not\supseteq S$. Define $w' \stackrel{\text{def}}{=} w - x \in \mathbb{Z}^{A_{V,q}}$. Since $S \in \mathcal{Z} \setminus \mathcal{T}_w$, every $T \in \mathcal{T}_w$ satisfies $T \not\supseteq S$ by definition of $\mathcal{T}_w$. Therefore $\pi_T(w') = \pi_T(w) - \pi_T(x) = \pi_T(w) = \sigma(T)$ whenever $T \in \mathcal{T}_w$. This implies $\mathcal{T}_{w'} \supseteq \mathcal{T}_w$. Further $\pi_S(w') = \pi_S(w) - \pi_S(x) = \sigma(S)$, so $\mathcal{T}_{w'} \supseteq \mathcal{T}_w \cup \{S\} \supsetneq \mathcal{T}_w$. ☐

The theorem now follows from the claim: Since $\min \mathcal{Z}$ is ℓ'-insular in $\mathcal{C}[V] = \mathcal{C}$, $A_{S,q} \neq \emptyset$ by Claim 6.23 and some $w \in \mathbb{Z}^{A_{V,q}}$ exists. Apply the claim iteratively until we get the desired $\hat{w} \in \mathbb{Z}^{A_{V,q}}$ with $\mathcal{T}_{\hat{w}} = \mathcal{Z}$. We can then set $\sigma'(V) \stackrel{\text{def}}{=} \hat{w}$ and $\sigma'(T) \stackrel{\text{def}}{=} \sigma(T)$ for $T \in \mathcal{Z}'$. This implies $\sigma'(\min \mathcal{Z}) = \sigma(\min \mathcal{Z}) = \mathbb{1}_q$.

We have $\sigma'(T) = \pi_T(\sigma'(S))$ for $T, S \in \mathcal{Z}$, $T \subseteq S$. Indeed, if $S \subsetneq V$, then the equality holds since σ is a scheme for $\mathcal{Z}'$. If $S = T = V$ then the equality holds, too. If $S = V \supsetneq T$, then the equality holds since $\mathcal{T}_{\hat{w}} = \mathcal{Z}$. □

Theorem 7.17. *Let $k, \ell \geqslant 2$. Suppose a k-CSP $(D, \mathcal{R})$ is ℓ-null-constraining and lax, $I = (V, \mathcal{C})$ its instance, and $\mathcal{Z}$ an ℓ-nice ℓ'-insular family in I. Given $q \in A_{\min \mathcal{Z}}$, there is an AIP-extension σ of q for $\mathcal{Z}$.*

Proof. We prove by induction on $\mathcal{Z}$ along the inclusion order. The base case is $|\mathcal{Z}| = 1$, then $\mathcal{Z} = \{\min \mathcal{Z}\}$, and $\sigma(\min \mathcal{Z}) = \mathbb{1}_q$ works.

If $|\mathcal{Z}| > 1$, pick any maximal $S \in \mathcal{Z}$. The subfamily $\mathcal{Z}' \overset{\text{def}}{=} \mathcal{Z} \setminus \{S\}$ is an ℓ-nice ℓ'-insular family. Induction Hypothesis yields an extension σ' of q for $\mathcal{Z}'$. Theorem 7.15 (applied to $J \overset{\text{def}}{=} I[S]$, $\mathcal{Z} \overset{\text{def}}{=} \mathcal{Z}[S]$, and $\sigma \overset{\text{def}}{=} \sigma'|_{\mathcal{Z}[S]\setminus\{S\}}$) yields an extension σ'' of q for $\mathcal{Z}[S]$ such that $\sigma''(T) = \sigma'(T)$ for $T \in \mathcal{Z}[S] \setminus \{S\}$.

Then $\sigma = \sigma' \cup \sigma''$ is the desired extension. Let's verify $\sigma(T) = \pi_T(\sigma(R))$ for $T, R \in \mathcal{Z}$ and $T \subseteq R$. If $R = S$, then $T \in \mathcal{Z}[S]$ and the equality follows because σ'' is a scheme for $\mathcal{Z}[S]$. If $R \neq S$, then $T, R \in \mathcal{Z}'$ by maximality of S, and the equality follows because σ' is a scheme for $\mathcal{Z}'$. □

Theorem 7.18. *Let $\ell, k \geqslant 2$. Suppose a k-CSP is ℓ-null-constraining and lax. For any constraint density $\Delta > 0$, with uniformly positive probability, a random instance of the CSP with n variables and Δn constraints has an AIP hierarchy solution of level $\Omega(n)$.*

Proof. By Lemmas 5.24, 5.28 and 5.31, there is $\delta \overset{\text{def}}{=} \delta(k, \ell')$ given by Lemma 5.26, $\gamma > 0$ and $t = \Omega(n)$ such that with uniformly positive probability, a random instance I of the CSP with n variables and Δn constraints is $(4t, \delta)$-sparse and $(\ell', 4t, \gamma)$-expanding, has girth more than 3ℓ, and hence $(2t, \gamma/\ell')$-confined by Remark 6.9. Therefore $\mathcal{Z}_t$ defined with the $\mathrm{cl}_{S,\ell'}(\mathcal{E})$ operator is ℓ-nice by Theorem 7.6. Theorems 6.16 and 7.17 construct an AIP hierarchy solution for I of level $\Omega(n)$. □

8. Combined Lower Bound

In this section, we prove lower bounds to combined hierarchies such as cohomological consistency and C(BLP+AIP).

8.1. Multiple Closures.

Consider boundary operators $\mathcal{B}^{\mathfrak{A}}, \mathcal{B}^{\mathfrak{B}}$ and their closure operators $\mathrm{cl}^{\mathfrak{A}}, \mathrm{cl}^{\mathfrak{B}}$ for $(V, \mathcal{H})$.

Write $\mathcal{B}^{\mathfrak{B}} \subseteq \mathcal{B}^{\mathfrak{A}}$ if $\mathrm{cl}^{\mathfrak{A}}_S(\mathcal{F}) = \emptyset$ for $\mathcal{F} \in \mathcal{B}^{\mathfrak{B}}_S$, $S \subseteq V$. In other words, every hyperedge set $\mathcal{F}$ in $\mathcal{B}^{\mathfrak{B}}$ morally belongs to $\mathcal{B}^{\mathfrak{A}}$, in the sense that $\mathcal{F}$ becomes nothing after iteratively removing boundaries in $\mathcal{B}^{\mathfrak{A}}$. Write $\mathrm{cl}^{\mathfrak{A}} \subseteq \mathrm{cl}^{\mathfrak{B}}$ if $\mathrm{cl}^{\mathfrak{A}}_S(\mathcal{E}) \subseteq \mathrm{cl}^{\mathfrak{B}}_S(\mathcal{E})$ for $S \subseteq V, \mathcal{E} \subseteq \mathcal{H}$.

Lemma 8.1. *Fix any $\ell > 1$ and $\mathcal{H} = \binom{V}{\geqslant 3}$. Let $\mathcal{B}$ be the pairwise boundary ($\tau = 2$) and $\mathcal{B}'$ the BW boundary with parameter ℓ. Then $\mathcal{B}' \subseteq \mathcal{B}$.*

Proof. For $S \subseteq V$, $e \in \mathcal{H}$, if $\{e\} \in \mathcal{B}'_S(\mathcal{E})$, then e has at most one variable u with $u \in S$ or $\deg_{\mathcal{E}}(u) > 1$, so $\{e\} \in \mathcal{B}$.

If $\mathcal{E}(P) \in \mathcal{B}'_S(\mathcal{E})$, then $P = (e_0, \ldots, e_{\ell-1})$ is a pendant path of length ℓ in $\mathcal{E}$, whose non-endpoints are outside S. One can iteratively remove e_i from P starting from $i = 0$, so that e_i always has at most two variables u with $u \in S$ or $\deg_{\mathcal{E}\setminus\{e_0,\ldots,e_i\}}(u) > 1$. Therefore $\mathrm{cl}_S(\mathcal{E}(P)) = \emptyset$, where cl is the closure operator of $\mathcal{B}$. □

Lemma 8.2. *If $\mathcal{B}^{\mathfrak{B}} \subseteq \mathcal{B}^{\mathfrak{A}}$, then $\mathrm{cl}^{\mathfrak{A}} \subseteq \mathrm{cl}^{\mathfrak{B}}$.*

Proof. For any $S \subseteq V$, $\mathcal{E} \subseteq \mathcal{H}$, $\mathcal{B}^{\mathfrak{B}} \subseteq \mathcal{B}^{\mathfrak{A}}$ implies $\mathrm{int}^{\mathfrak{A}}_S(\mathcal{E}) \subseteq \mathrm{int}^{\mathfrak{B}}_S(\mathcal{E})$. Induction on i implies $\mathrm{int}^{\mathfrak{A},i}_S(\mathcal{E}) \subseteq \mathrm{int}^{\mathfrak{B},i}_S(\mathcal{E})$ for every $i \in \mathbb{N}$. Therefore $\mathrm{cl}^{\mathfrak{A}}_S(\mathcal{E}) \subseteq \mathrm{cl}^{\mathfrak{B}}_S(\mathcal{E})$. □

Given a closure operator cl, an cl-insular family is an insular family under cl.

Lemma 8.3. *Given any hypergraph H and $\mathrm{cl}^{\mathfrak{A}} \subseteq \mathrm{cl}^{\mathfrak{B}}$, any $\mathrm{cl}^{\mathfrak{B}}$-insular family $\mathcal{Z}$ is also a $\mathrm{cl}^{\mathfrak{A}}$-insular family.*

Proof. For any $T, S \in \mathcal{Z}$ such that $T \subseteq S$, $\mathrm{cl}^{\mathfrak{A}}_T(\mathcal{E}[S]) \subseteq \mathrm{cl}^{\mathfrak{B}}_T(\mathcal{E}[S]) \subseteq \mathcal{E}[T] = (\mathcal{E}[S])[T]$, where the first inclusion is $\mathrm{cl}^{\mathfrak{A}} \subseteq \mathrm{cl}^{\mathfrak{B}}$, and the second is because $\mathcal{Z}$ is a $\mathrm{cl}^{\mathfrak{B}}$-insular family. Therefore T is $\mathrm{cl}^{\mathfrak{A}}$-insular in $\mathcal{E}[S]$ as well. □

For combined hierarchies $\mathfrak{A} + \mathfrak{B}$, $\mathcal{B}^{\mathfrak{B}} \subseteq \mathcal{B}^{\mathfrak{A}}$ often holds, so $\mathrm{cl}^{\mathfrak{A}} \subseteq \mathrm{cl}^{\mathfrak{B}}$. For the purpose of constructing an insular family, we only need to consider a single closure operator $\mathrm{cl} = \mathrm{cl}^{\mathfrak{B}}$.

8.2. **Combined Lower Bound.**

Throughout this subsection, consider a pair of hierarchies $\mathfrak{A}$ and $\mathfrak{B}$. $\mathcal{A}$ is BW or LP, while $\mathfrak{B}$ is BW, LP, or AIP (in interesting cases $\mathfrak{B}$ is AIP). Fix also closure operators $\mathrm{cl}^{\mathfrak{A}}, \mathrm{cl}^{\mathfrak{B}}$ for $(V, \mathcal{H})$. Recall $V_d \overset{\text{def}}{=} \binom{V}{\leqslant d}$ and $\mathcal{Z}_{\leqslant d} \overset{\text{def}}{=} \left\{ R^{\uparrow} \mid R \in \mathcal{Z}_{\downarrow} \cap V_d \right\}$ defined in Section 6.4.

Theorem 8.4. *Let $\mathfrak{A}$ be BW or LP. Suppose a k-CSP is $\mathfrak{A}$-fully-extensible over $\mathrm{cl}^{\mathfrak{A}}$-insular families and $\mathfrak{B}$-extensible over $\mathrm{cl}^{\mathfrak{B}}$-insular families, where $\mathrm{cl}^{\mathfrak{A}} \subseteq \mathrm{cl}^{\mathfrak{B}}$. Given any instance $I = (V, \mathcal{C})$ with $\mathcal{C} \subseteq \mathcal{H}$ and is $(8t, \beta)$-confined under $\mathrm{cl}^{\mathfrak{B}}$, there is a solution of level $d \overset{\text{def}}{=} t/(1 + k/\beta)$ in the strong $\mathfrak{A} + \mathfrak{B}$ hierarchy.*

Proof. Lemma 6.12 implies $\mathcal{Z}_{4t}$ is nonempty and $\min \mathcal{Z}_{4t} = \emptyset$. Lemmas 6.4 and 8.3 imply $\mathcal{Z}_{4t}$ is an insular family under both $\mathrm{cl}^{\mathfrak{A}}$ and $\mathrm{cl}^{\mathfrak{B}}$. Since $\mathbb{0} \in A_{\emptyset}$ and the CSP is $\mathfrak{A}$-fully-extensible, there is a full $\mathfrak{A}$-extension $\sigma^{\mathfrak{A}}$ of $\mathbb{0}$ for $\mathcal{Z}_t$. Let $s \overset{\text{def}}{=} \sigma^{\mathfrak{A}}_{\downarrow}$. $s(\emptyset) = \sigma^{\mathfrak{A}}_{\downarrow}(\emptyset) \overset{(7)}{=} \pi_{\emptyset} \circ \sigma^{\mathfrak{A}}(\emptyset) = \pi_{\emptyset}(\mathbb{1}_{\mathbb{0}}) = \mathbb{1}_{\mathbb{0}}$. Since $V_{4d} \subseteq (\mathcal{Z}_{4t})_{\downarrow}$ by Lemma 6.12, Lemmas 6.15 and A.1 imply s is a level-$2d$ $\mathfrak{A}$-solution. For $S \in V_{2d}$,

$$\mathrm{supp}(s(S)) = \mathrm{supp}\left(\sigma^{\mathfrak{A}}_{\downarrow}(S)\right) \overset{(11)}{=} \pi_S\left(A_{S^{\uparrow}}\right). \tag{15}$$

For every $T \in \mathcal{Z}_{4t}$, consider the subfamily $\mathcal{Z}_{\supseteq T} \overset{\text{def}}{=} \{S \in \mathcal{Z}_{4t} \mid S \supseteq T\}$. Then $\mathcal{Z}_{\supseteq T}$ is an insular family with $\min \mathcal{Z}_{\supseteq T} = T$. Since the CSP is $\mathfrak{B}$-extensible, for every $q \in A_T$, there is a $\mathfrak{B}$-extension $\sigma^{\mathfrak{B},q}$ of q over $\mathcal{Z}_{\supseteq T}$. For $K \in V_d$, $b \in \pi_K(A_{K^{\uparrow}})$, take any $q \in A_{K^{\uparrow}, b}$, and define $r(K, b) \overset{\text{def}}{=} \sigma^{\mathfrak{B},q}_{\downarrow K}$ by Eq. (8). $r(K, b)$ is a scheme by Lemma 6.15. Then $r(K, b)(K) = \pi_K(\sigma^{\mathfrak{B},q}(K^{\uparrow})) = \pi_K(\mathbb{1}_q) = \mathbb{1}_b$. Lemma A.1 implies $r(K, b)$ is a level-d $\mathfrak{B}$-solution.

Further, for $K, S \in V_d$, $b \in \mathrm{supp}(s(K)) = \pi_K\left(A_{K^{\uparrow}}\right)$, Eq. (9) implies

$$\mathrm{supp}(r(K, b)(S)) \subseteq \pi_S\left(A_{(K \cup S)^{\uparrow}, q}\right) \subseteq \pi_S(A_{S^{\uparrow}}) = \mathrm{supp}(s(S)).$$ □

Remark 8.5. Even though we did not prove τ-wise uniform CSPs to be LP-fully-extensible, they satisfy a weaker property so that Theorem 8.4 applies to them. Theorem 6.34 shows that if an insular family $\mathcal{Z}_{4t}$ satisfies $V_{4d} \subseteq (Z_{4t})_{\downarrow}$ (as in the proof of Theorem 8.4), there is a full LP-extension σ^{LP} of q over $(\mathcal{Z}_{4t})_{\leqslant 2d}$. Therefore Eq. (15) holds for $S \in V_{2d}$, and Theorem 8.4 also applies to pairwise uniform CSPs when $\mathfrak{A}$ is LP.

Remark 8.6. Even though we did not prove ℓ-null-constraining and lax CSPs to be AIP-extensible, they satisfy a weaker property so that Theorem 8.4 applies to them, provided I is additionally $(16t, \delta)$-sparse with $\delta = \delta(k, \ell') > 0$ from Lemma 5.26, $\ell' \overset{\text{def}}{=} 3\ell + 3$ and $\mathrm{girth}(I) > 3\ell$. As in the proof of Theorem 7.18, we take $\mathrm{cl}^{\mathrm{AIP}}$ to be the BW closure with parameter ℓ'. Then $\mathcal{Z}_{4t}$ is additionally ℓ-nice, so that Theorem 7.17 applies. Therefore Theorem 8.4 also applies to ℓ-null-constraining and lax CSPs when $\mathfrak{B}$ is AIP, as long as I is also $(8t, \delta)$-sparse and $\mathrm{girth}(I) > 3\ell$.

Corollary 8.7. *Let $k \geqslant 2$. Suppose a k-CSP is null-contraining and lax. Then, for any constraint density $\Delta > 0$, with uniformly positive probability, a random instance of the CSP with n variables and Δn constraints has a solution in the strong BW+AIP hierarchy of level $\Omega(n)$.*

Proof. Suppose the CSP is ℓ-null-constraining and lax. We may assume $\ell \geqslant 2$ since 1-null-constraining implies 2-null-constraining. Then the CSP is BW-fully-extensible by Theorem 6.22. It also satisfies a weaker notion of AIP-extensibility by Theorem 7.17.

By Lemmas 5.24, 5.28 and 5.31, there is $\delta = \delta(k, \ell') > 0$ from Lemma 5.26, $\gamma > 0$ and $t = \Omega(n)$ such that with uniformly positive probability, a random instance I of the CSP with n variables and Δn constraints is $(16t, \delta)$-sparse and $(\ell', 16t, \gamma)$-expanding, has girth more than 3ℓ, and hence $(8t, \gamma/\ell')$-confined by Remark 6.9. The result now follows by Theorem 8.4 and Remark 8.6. □

Together with Proposition B.1, Corollary 8.7 also shows lower bound for the cohomological consistency hierarchy.

Propositions B.4 and B.6 show that C(BLP+AIP) hierarchy lower bound amounts to the following:

A function $\mathcal{A}^* : (S \in \binom{V}{\leqslant d}) \to 2^{A_S}$, family $u(\cdot,\cdot)$ of level-d LP hierarchy solution, family $r(\cdot,\cdot)$ of level-d AIP hierarchy solution, such that for $T, S \in \binom{V}{\leqslant d}$, $b \in \mathcal{A}^*(S)$,

- $\mathcal{A}^*(S)$ is nonempty
- $(u(S,b), r(S,b))$ is a level-d LP+AIP hierarchy solution
- $\operatorname{supp} \circ u(S,b)(T) \subseteq \mathcal{A}^*(T)$
- $\operatorname{supp} \circ u(S,b)(S) = \mathbb{1}_b$

Theorem 8.8. *Let $k \geqslant 3$. Suppose a k-CSP is pairwise uniform and lax. Then, for any constraint density $\Delta > 0$, with uniformly positive probability, a random instance of the CSP with n variables and Δn constraints has a solution in the C(BLP+AIP) hierarchy of level $\Omega(n)$.*

Proof. A pairwise uniform CSP $(D, \mathcal{R})$ is 1-null-constraining and hence 2-null-constraining. It satisfies a weaker notion of LP-full-extensibility by Theorem 6.34. It also satisfies a weaker notion of AIP-extensibility by Theorem 7.17.

By Lemmas 5.24, 5.28 and 5.31, there is $\delta = \delta(k, \ell') > 0$ from Lemma 5.26, $\gamma > 0$ and $t = \Omega(n)$ such that with uniformly positive probability, a random instance I of the CSP with n variables and Δn constraints is $(16t, \delta)$-sparse and $(\ell', 16t, \gamma)$-expanding, has girth more than 3ℓ, and hence $(8t, \gamma/\ell')$-confined under the BW closure operator $\mathrm{cl}^{\mathfrak{B}}$ by Remark 6.9. We can therefore apply Theorem 8.4 by Remarks 8.5 and 8.6. Let $\mathrm{cl}^{\mathfrak{A}}$ be the pairwise closure. $\mathrm{cl}^{\mathfrak{A}} \subseteq \mathrm{cl}^{\mathfrak{B}}$ follows from Lemmas 8.1 and 8.2.

Let $d, \mathcal{Z}_{4t}$ be defined as in the proof of Theorem 8.4. Theorem 8.4 (with $\mathfrak{A} = \mathrm{LP}$ and $\mathfrak{B} = \mathrm{AIP}$) yields a level-$d$ LP solution $s = \sigma_{\downarrow}^{\mathrm{LP}}$, and a family r of level-d AIP solutions. More precisely, for $S \in V_d, b \in \pi_S(A_{S^\uparrow})$, take any $q \in A_{S^\uparrow,b}$, and define $r(S,b) \overset{\text{def}}{=} \sigma_{\downarrow S}^{\mathrm{AIP},q}$ by Eq. (8). Eq. (9) implies $\operatorname{supp}(r(S,b)(T)) \subseteq \pi_T\left(A_{(S\cup T)^\uparrow,q}\right)$ for $T \in V_d$.

Define $\mathcal{A}^*(S) \overset{\text{def}}{=} \operatorname{supp}(s(S)) = \operatorname{supp}(\sigma_{\downarrow}^{\mathrm{LP}}(S)) \overset{(11)}{=} \pi_S(A_{S^\uparrow})$ for $S \in V_d$. In particular

$$\pi_\emptyset(\mathcal{A}^*(S)) = \pi_\emptyset(\pi_S(A_{S^\uparrow})) \overset{(4)}{=} \pi_\emptyset(A_{S^\uparrow}) \overset{(*)}{=} A_\emptyset = \{\mathbb{0}\} \neq \emptyset,$$

where $(*)$ is Claim 6.29. Therefore $\mathcal{A}^*(S) \neq \emptyset$.

Theorem 8.4 (with $\mathfrak{A} = \mathfrak{B} = \mathrm{LP}$) yields a family u of LP solutions. More precisely, for $S \in V_d, b \in \operatorname{supp}(s(T)) = \pi_S(A_{S^\uparrow})$, take any $q \in A_{S^\uparrow,b}$, Theorem 6.34 yields a full extension $\sigma^{\mathrm{LP},q}$ of q over $\mathcal{Z}_{\leqslant 2d}$, and let $u(S,b) \overset{\text{def}}{=} \sigma_{\downarrow S}^{\mathrm{LP},q}$. Then $u(S,b)(S) = \pi_S \circ \sigma^{\mathrm{LP},q}(S^\uparrow) = \pi_S(\mathbb{1}_q) = \mathbb{1}_b$. Eq. (11) implies $\operatorname{supp}(u(S,b)(T)) = \pi_T\left(A_{(S\cup T)^\uparrow,q}\right) \supseteq \operatorname{supp}(r(S,b)(T))$. Also $\operatorname{supp}(u(S,b)(T)) = \pi_T\left(A_{(S\cup T)^\uparrow,q}\right) \subseteq \pi_T(A_{T^\uparrow}) = \mathcal{A}^*(T)$. This yields a solution to C(BLP+AIP) by Proposition B.4. □

Acknowledgements. Siu On thanks Chris Jones and Standa Živný for helpful discussions, and University of Bocconi where this work was partially done. We thank anonymous STOC reviewers for their valuable feedback and suggestions.

Appendix A. Useful Results

A.1. Hierarchy.

Lemma A.1 ([CNP24, Lemma 3.3]). *For BW, LP, and AIP, a dependent function $s : (S \in \binom{V}{\leqslant d}) \to \mathcal{M}^{A_S}$ with $s(\emptyset) = \mathbb{1}_{\emptyset}$ and is consistent with projection is a hierarchy solution (i.e. takes values in $\mathcal{D}_S$).*

Lemma A.2 ([CNP24, Lemma A.3]). *Suppose $\mathcal{M}$ is a commutative semiring. For any disjoint S and T, $\alpha_S \in \mathcal{M}^{D^S}$, $\alpha_T \in \mathcal{M}^{D^T}$, $R \subseteq S \cup T$,*

$$\pi_{S\cup T \to R}(\alpha_S \otimes \alpha_T) = \pi_{S\to S\cap R}(\alpha_S) \otimes \pi_{T\to T\cap R}(\alpha_T). \tag{16}$$

Lemma A.3. *For BW and LP, for any $S \subseteq T$, $\alpha \in \mathcal{M}^{D^T}$,*

$$\pi_S \circ \operatorname{supp}(\alpha) = \operatorname{supp} \circ \pi_S(\alpha). \tag{17}$$

Proof. See [CNP24, Lemma A.8] for LP. The proof for BW is analogous. □

Appendix B. Algorithm Lower Bound

This section concerns the equivalences of solutions to the combined hierarchies:

$$\text{strong BW+AIP hierarchy} \iff \text{cohomological consistency hierarchy}$$
$$\text{a variant of strong LP+AIP hierarchy} \iff \text{C(BLP+AIP) hierarchy}$$

These equivalences are akin to the equivalence between solutions to the BW hierarchy and to the local consistency checking algorithm (that finds the maximal strategy by repeatedly removing violations).

A level-d hierarchy solution s *is supported on* $\mathcal{A} : (S \in \binom{V}{\leqslant d}) \to 2^{A_S}$ if $s(S) \subseteq \mathcal{A}(S)$ for every $S \in \binom{V}{\leqslant d}$.

B.1. Cohomological Consistency Hierarchy.

The following is a transcription of the cohomological consistency hierarchy algorithm [ÓC22, Section 4.3.1] using our terminology.

- Fix any input instance $I = (V, \mathcal{C})$ of a CSP $(D, \mathcal{R})$. Fix the level-d of the hierarchy.
- For every d-small variable subset $S \subseteq V$, let $\mathcal{A}_0(S) \stackrel{\text{def}}{=} A_S$. Let $i \stackrel{\text{def}}{=} 0$.
- Repeat the following elimination rules until they are not applicable. Let $\mathcal{A}$ denote the resulting $\mathcal{A}_i$.
  - ("**Forth**") Take any d-small $S, T \subseteq V$ such that $T = S \cup \{v\}$ for some $v \in V$. Take any $a \in \mathcal{A}_i(S)$. Check if $a \notin \pi_S \circ \mathcal{A}_i(T)$.
  - ("**$\mathbb{Z}$ext**") Take any d-small $S \subseteq V$, $a \in \mathcal{A}_i(S)$ such that there is no level-d AIP hierarchy solution r supported on $\mathcal{A}_i$ where $r(S) = \mathbb{1}_a$.

    In both cases, let $\mathcal{A}_{i+1}(U) \stackrel{\text{def}}{=} \mathcal{A}_i(U)$ for every d-small $U \subseteq V$ except S, where $\mathcal{A}_{i+1}(S) \stackrel{\text{def}}{=} \mathcal{A}_i(S) \setminus \{a\}$. Increment i.
- If $\mathcal{A}(S)$ is empty for some d-small $S \subseteq V$, **reject**; else **accept**.[12]

Proposition B.1. *If $I = (V, \mathcal{C})$ has a level-d strong* BW+AIP *solution (s, r), then the cohomological d-consistency algorithm accepts I.*

Proof. We claim that assignments in $s(S)$ (for any d-small $S \subseteq V$) cannot be eliminated by cohomological d-consistency.

Claim B.2. *For every $i \geqslant 0$, $\mathcal{A}_i(S) \supseteq s(S)$ for every d-small $S \subseteq V$.*

[12] We can simply check if $\mathcal{A}(\emptyset)$ is empty due to the forth elimination rule.

Proof. We prove by induction in i.

Base Case: $i = 0$. For any d-small $S \subseteq V$, $s(S) \subseteq A_S = \mathcal{A}_0(S)$.

Induction Step: $i \geqslant 1$. Take any d-small $S \subseteq V$, $a \in s(S)$. By induction hypothesis, $a \in \mathcal{A}_{i-1}(S)$. We now claim that a cannot be eliminated in the next iteration.

We first show that a cannot be eliminated by "Forth". Indeed, take any $v \in V$ and $T \stackrel{\text{def}}{=} S \cup \{v\}$. We know that $a \in s(S) = \pi_T \circ s(T)$ and therefore there is some $b \in s(T)$ such that $b_S = a$. Since $b \in s(T) \stackrel{\text{I.H.}}{\subseteq} \mathcal{A}_{i-1}(T)$, we know that $a = b_S \in \pi_S \circ \mathcal{A}_{i-1}(T)$, and "Forth" is not applicable for this S, T and a. The same argument holds for all T, and therefore a cannot be eliminated with "Forth".

We now show that a cannot be eliminated by "$\mathbb{Z}$ext". Indeed, consider the AIP hierarchy solution $r(S,a)$. For every d-small T, we know that $r(S,a)(T)$ is supported on $s(T)$ and therefore supported on $\mathcal{A}_{i-1}(T)$ by induction hypothesis. We also know that $r(S,a)(S) = \mathbb{1}_a$. This means that $r(S,a)$ satisfy the conditions of r in "$\mathbb{Z}$ext", and a cannot be eliminated by "$\mathbb{Z}$ext". □

We are done because $s(S)$ is not empty for every d-small $S \subseteq V$, and so is $\mathcal{A}(S)$ using the claim above. □

Proposition B.3. *If cohomological d-consistency accepts $I = (V, \mathcal{C})$, then I has a level-d strong* BW+AIP *solution (s,r).*

Proof. Assignments in $s \stackrel{\text{def}}{=} \mathcal{A}$ are not eliminated by cohomological d-consistency. Therefore, for each d-small $S \subseteq V$ and $a \in s(S)$, there is a level-d AIP hierarchy solution r' supported on $\mathcal{A}$ such that $r'(S) = \mathbb{1}_a$. For each such (S,a), let $r(S,a) \stackrel{\text{def}}{=} r'$ be this level-d AIP hierarchy solution (that is dependent on (S,a)).

We now claim that (s,r) is a level-d strong BW+AIP solution. Indeed, for every d-small $S,T \subseteq V$, $a \in s(S) = \operatorname{supp} \circ s(S)$,

- $\operatorname{supp} \circ r(S,a)(T) \subseteq \mathcal{A}(T) = s(T) = \operatorname{supp} \circ s(T)$.
- $r(S,a)(S) = \mathbb{1}_a$ by construction.

□

B.2. **C(BLP+AIP) Hierarchy Algorithm.**

The following is a transcription of the C(BLP+AIP) algorithm [CŽ22, Remark 19] using our terminology.

- Fix any input instance $I = (V, \mathcal{C})$ of a CSP $(D, \mathcal{R})$.
- For every constraint $C \in \mathcal{C}$, let $\mathcal{A}_0(C) \stackrel{\text{def}}{=} A_C$. Let $i \stackrel{\text{def}}{=} 0$.
- Repeat the following elimination rule until it is not applicable. Let $\mathcal{A}$ denote the resulting $\mathcal{A}_i$.
  - (Elimination Rule) Take any $C' \in \mathcal{C}$ and $a' \in \mathcal{A}_i(C')$ such that there are no probability distributions and integer weights μ_C's, μ_v's, w_C's, w_v's over all $C \in \mathcal{C}$, $v \in V$ where
    - $\operatorname{supp} \circ w_C \subseteq \operatorname{supp} \circ \mu_C \subseteq \mathcal{A}(C)$ for all $C \in \mathcal{C}$.
    - $\pi_v \circ \mu_C = \mu_v$ and $\pi_v \circ w_C = w_v$ for all $C \in \mathcal{C}$ and $v \in V(C)$.
    - $\mu_{C'} = \mathbb{1}_{a'}$.

    Let $\mathcal{A}_{i+1}(C) \stackrel{\text{def}}{=} \mathcal{A}_i(C)$ for all $C \in \mathcal{C}$ except C', where $\mathcal{A}_{i+1}(C') \stackrel{\text{def}}{=} \mathcal{A}_i(C') \setminus \{a'\}$. Increment i.
- If $\mathcal{A}(C)$ is empty for some $C \in \mathcal{C}$, **reject**; else **accept**.

Similar to how CLAP can be strengthened into a hierarchy [CŽ22, Remark 33], we strengthen C(BLP+AIP) algorithm into a hierarchy algorithm as follows:

- Fix any input instance $I = (V, \mathcal{C})$ of a CSP $(D, \mathcal{R})$. Fix the level-d of the hierarchy.
- For every d-small variable subsets $S \subseteq V$, let $\mathcal{A}_0(S) \stackrel{\text{def}}{=} A_S$. Let $i \stackrel{\text{def}}{=} 0$.

- Repeat the following elimination rule until it is not applicable. Let $\mathcal{A}$ denote the resulting $\mathcal{A}_i$.
  - (Elimination Rule) Take any d-small $S \subseteq V$ and $a \in \mathcal{A}_i(S)$ such that there is no level-d LP+AIP hierarchy solution (s, r) supported on $\mathcal{A}_i$ where $s(S) = \mathbb{1}_a$.
    Let $\mathcal{A}_{i+1}(U) \stackrel{\text{def}}{=} \mathcal{A}_i(U)$ for all d-small $U \subseteq V$ except S, where $\mathcal{A}_{i+1}(S) \stackrel{\text{def}}{=} \mathcal{A}_i(S) \setminus \{a\}$. Increment i.
- If $\mathcal{A}(S)$ is empty for some d-small $S \subseteq V$, **reject**; else **accept**.

Proposition B.4. *Fix $I = (V, \mathcal{C})$. If there is $\mathcal{A}^* : (S \in \binom{V}{\leqslant d}) \to 2^{A_S}$, family of level-$d$ LP hierarchy solution $s(\cdot,\cdot)$, family of level-d AIP hierarchy solution $r(\cdot,\cdot)$ such that for every $T, S \in \binom{V}{\leqslant d}$, $a \in \mathcal{A}^*(S)$,*

- *$\mathcal{A}^*(S)$ is nonempty.*
- *$(s(S,a), r(S,a))$ is a level-d LP+AIP hierarchy solution.*
- *$\operatorname{supp} \circ s(S,a)(T) \subseteq \mathcal{A}^*(T)$.*
- *$\operatorname{supp} \circ s(S,a)(S) = \mathbb{1}_a$.*

Then the level-d C(BLP+AIP) hierarchy accepts I.

Proof. We claim that assignments in $\mathcal{A}^*$ (for any d-small $S \subseteq V$) cannot be eliminated by C(BLP+AIP) hierarchy.

Claim B.5. *For all $i \geqslant 0$, $\mathcal{A}_i(S) \supseteq \mathcal{A}^*(S)$ for all d-small $S \subseteq V$.*

Proof. We prove by induction in i.

Base Case: $i = 0$. For any d-small $S \subseteq V$, $\mathcal{A}^*(S) \subseteq A_S = \mathcal{A}_0(S)$.

Induction Step: $i \geqslant 1$. Take any d-small $S \subseteq V$, $a \in \mathcal{A}^*(S)$. By induction hypothesis, $a \in \mathcal{A}_{i-1}(S)$. We now claim that a cannot be eliminated in the next iteration.

Indeed, $(s(S,a), r(S,a))$ is a level-d LP+AIP hierarchy solution that satisfy the restrictions of (s, r) in the elimination rule:

- For all d-small $T \subseteq V$, $\operatorname{supp} \circ r(S,a)(T) \stackrel{(*)}{\subseteq} \operatorname{supp} \circ s(S,a)(T) \subseteq \mathcal{A}^*(T) \stackrel{\text{I.H.}}{\subseteq} \mathcal{A}_{i-1}(T)$, where $(*)$ is due to $(s(S,a), r(S,a))$ being an LP+AIP hierarchy solution. This means that $s(S,a)$ and $r(S,a)$ are both supported on $\mathcal{A}_{i-1}$.
- $\operatorname{supp} \circ s(S,a)(S) = \mathbb{1}_a$ □

We are done because $\mathcal{A}^*(S)$ is not empty for all d-small $S \subseteq V$, and so is $\mathcal{A}(S)$ using the claim above. □

C(BLP+AIP) is in fact equivalent to the level-d hierarchy in Proposition B.4. That is, the C(BLP+AIP) hierarchy accepts I if and only if there is a level-d hierarchy solution.

Proposition B.6. *If the level-d C(BLP+AIP) hierarchy accepts $I = (V, \mathcal{C})$, then there is $\mathcal{A}^*$, $s(\cdot,\cdot)$ and $r(\cdot,\cdot)$ that satisfy the restrictions listed in Proposition B.4.*

Proof. Assignments in $\mathcal{A}^* \stackrel{\text{def}}{=} \mathcal{A}$ are not eliminated by C(BLP+AIP) hierarchy. Therefore, for each d-small $S \subseteq V$ and $a \in \mathcal{A}^*(S)$, there is a level-d LP+AIP hierarchy solution (s', r') supported on $\mathcal{A}^*$ where $s'(S) = \mathbb{1}_a$. For each such (S,a), let $s(S,a) \stackrel{\text{def}}{=} s'$ and $r(S,a) \stackrel{\text{def}}{=} r'$, where (s', r') is the level-d LP+AIP hierarchy solution (that is dependent on (S,a)) given above.

We now claim that $\mathcal{A}^*$, $s(\cdot,\cdot)$ and $r(\cdot,\cdot)$ satisfy the restrictions. Indeed, for every d-small $T, S \subseteq V$, $a \in \mathcal{A}^*(S)$,

- $\mathcal{A}^*(S)$ is nonempty because C(BLP+AIP) accepts I.
- $(s(S,a), r(S,a))$ is a level-d LP+AIP hierarchy solution by construction.
- $\operatorname{supp} \circ s(S,a)(T) \subseteq \mathcal{A}^*(T)$ by construction.
- $\operatorname{supp} \circ s(S,a)(S) = \mathbb{1}_a$ by construction. □

References

[ABP20] Per Austrin, Amey Bhangale, and Aditya Potukuchi. *Improved Inapproximability of Rainbow Coloring*, pages 1479–1495. 2020.

[AD08] Albert Atserias and Víctor Dalmau. A Combinatorial Characterization of Resolution Width. *Journal of Computer and System Sciences*, 74(3):323–334, 2008. Computational Complexity 2003.

[AD22] Albert Atserias and Víctor Dalmau. Promise Constraint Satisfaction and Width. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1129–1153, 2022.

[Ahn20] Kwangjun Ahn. A Simpler Strong Refutation of Random k-XOR. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 2:1–2:15, 2020.

[AM09] Per Austrin and Elchanan Mossel. Approximation Resistant Predicates from Pairwise Independence. *Computational Complexity*, 18:249–272, 2009.

[AOW15] Sarah R. Allen, Ryan O'Donnell, and David Witmer. How to Refute a Random CSP. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 689–708, 2015.

[BBKO21] Libor Barto, Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic Approach to Promise Constraint Satisfaction. *Journal of the ACM*, 68(4), 2021.

[BCK15] Boaz Barak, Siu On Chan, and Pravesh K. Kothari. Sum of Squares Lower Bounds from Pairwise Independence. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, page 97–106. ACM, 2015.

[BG19] Joshua Brakensiek and Venkatesan Guruswami. An Algorithmic Blend of LPs and Ring Equations for Promise CSPs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, page 436–455. SIAM, 2019.

[BGMT12] Siavosh Benabbas, Konstantinos Georgiou, Avner Magen, and Madhur Tulsiani. SDP Gaps from Pairwise Independence. *Theory of Computing*, 8(12):269–289, 2012.

[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.

[BK22] Amey Bhangale and Subhash Khot. UG-hardness to NP-hardness by Losing Half. *Theory of Computing*, 18(5):1–28, 2022.

[BKM24] Amey Bhangale, Subhash Khot, and Dor Minzer. On Approximability of Satisfiable k-CSPs: V. *CoRR*, abs/2408.15377, 2024.

[BSW01] Eli Ben-Sasson and Avi Wigderson. Short Proofs Are Narrow—Resolution Made Simple. *Journal of the ACM*, 48(2):149–169, 2001.

[Bul17] Andrei A. Bulatov. A Dichotomy Theorem for Nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science*, pages 319–330, 2017.

[Bul20] Andrei A. Bulatov. A Dichotomy Theorem for Nonuniform CSPs Simplified. *CoRR*, abs/2007.09099, 2020.

[CDG11] Hubie Chen, Victor Dalmau, and Berit Grußien. Arc Consistency and Friends. *Journal of Logic and Computation*, 23(1):87–108, 11 2011.

[Cha16] Siu On Chan. Approximation Resistance from Pairwise-Independent Subgroups. *Journal of the ACM*, 63(3), 2016.

[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.

[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 2024, page 1944–1955, New York, NY, USA, 2024. Association for Computing Machinery.

[CŽ22] Lorenzo Ciardo and Stanislav Živný. CLAP: A New Algorithm for Promise CSPs. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1057–1068, 2022.

[CŽ23a] Lorenzo Ciardo and Stanislav Živný. Approximate Graph Colouring and Crystals. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2256–2267, 2023.

[CŽ23b] Lorenzo Ciardo and Stanislav Živný. Approximate Graph Colouring and the Hollow Shadow. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, page 623–631. ACM, 2023.

[CŽ23c] Lorenzo Ciardo and Stanislav Živný. Hierarchies of Minion Tests for PCSPs through Tensors. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 568–580, 2023.

[CŽ24a] 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*, 2024. To appear.

[CŽ24b] Lorenzo Ciardo and Stanislav Živný. The Periodic Structure of Local Consistency. *CoRR*, abs/2406.19685, 2024.

[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, New York, NY, USA, 2024. Association for Computing Machinery.

[dT23] Tommaso d'Orsi and Luca Trevisan. A Ihara-Bass Formula for Non-Boolean Matrices and Strong Refutations of Random CSPs. In *Proceedings of the Conference on Proceedings of the 38th Computational Complexity Conference*, CCC '23, Dagstuhl, DEU, 2023. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[GL18] Venkatesan Guruswami and Euiwoong Lee. Strong Inapproximability Results on Balanced Rainbow-Colorable Hypergraphs. *Combinatorica*, 38(3):547–599, June 2018.

[GS17] Venkatesan Guruswami and Rishi Saket. Hardness of Rainbow Coloring Hypergraphs. In Satya V. Lokam and R. Ramanujam, editors, *37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2017, December 11-15, 2017, Kanpur, India*, volume 93 of *LIPIcs*, pages 33:33–33:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[GS20] Venkatesan Guruswami and Sai Sandeep. Rainbow Coloring Hardness via Low Sensitivity Polymorphisms. *SIAM Journal on Discrete Mathematics*, 34(1):520–537, 2020.

[Hua13] Sangxia Huang. Improved Hardness of Approximating Chromatic Number. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 233–243. Springer Berlin Heidelberg, 2013.

[JLR00] Svante Janson, Tomasz Łuczak, and Andrzej Rucinski. *Random Graphs*. John Wiley & Sons, 2000.

[KMOW17] Pravesh K. Kothari, Ryuhei Mori, Ryan O'Donnell, and David Witmer. Sum of Squares Lower Bounds for Refuting Any CSP. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, page 132–145. ACM, 2017.

[KO22] Andrei Krokhin and Jakub Opršal. An Invitation to the Promise Constraint Satisfaction Problem. *ACM SIGLOG News*, 9(3):30–59, Aug 2022.

[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.

[Lau09] Monique Laurent. *Sums of Squares, Moment Matrices and Optimization Over Polynomials*, pages 157–270. Springer, 2009.

[LP24] Moritz Lichter and Benedikt Pago. Limitations of Affine Integer Relaxations for Solving Constraint Satisfaction Problems. *CoRR*, abs/2407.09097v1, 2024.

[MS07] Michael Molloy and Mohammad R. Salavatipour. The Resolution Complexity of Random Constraint Satisfaction Problems. *SIAM Journal on Computing*, 37(3):895–922, 2007.

[ÓC22] Adam Ó Conghaile. Cohomology in Constraint Satisfaction and Structure Isomorphism. In *47th International Symposium on Mathematical Foundations of Computer Science, 2022*, volume 241, pages 75:1–75:16, 2022.

[Oll02] Häggström Olle. Finite Markov Chains and Algorithmic Applications, 2002.

[Pag24] Benedikt Pago. Limitations of Affine CSP Algorithms. CSP World Congress 2024, 2024. https://www.cl.cam.ac.uk/~btp26/slidesCWC24handout.pdf.

[RRS17] Prasad Raghavendra, Satish Rao, and Tselil Schramm. Strongly Refuting Random CSPs below the Spectral Threshold. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, page 121–131. ACM, 2017.

[Zhu20] Dmitriy Zhuk. A Proof of the CSP Dichotomy Conjecture. *Journal of the ACM*, 67(5), 2020.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il