



Near-optimal Lower Bounds on Quantifier Depth and Weisfeiler–Leman Refinement Steps

CHRISTOPH BERKHOLZ, Humboldt-Universität zu Berlin, Germany

JAKOB NORDSTRÖM, University of Copenhagen, Denmark and Lund University, Sweden

We prove near-optimal tradeoffs for quantifier depth (also called quantifier rank) versus number of variables in first-order logic by exhibiting pairs of n -element structures that can be distinguished by a k -variable first-order sentence but where every such sentence requires quantifier depth at least $n^{\Omega(k/\log k)}$. Our tradeoffs also apply to first-order counting logic and, by the known connection to the k -dimensional Weisfeiler–Leman algorithm, imply near-optimal lower bounds on the number of refinement iterations.

A key component in our proof is the hardness condensation technique introduced by Razborov in the context of proof complexity. We apply this method to reduce the domain size of relational structures while maintaining the minimal quantifier depth needed to distinguish them in finite variable logics.

CCS Concepts: • **Theory of computation** → **Finite Model Theory**; *Proof complexity*;

Additional Key Words and Phrases: First-order logic, first-order counting logic, bounded variable fragment, quantifier depth, Weisfeiler–Leman, refinement iterations, lower bounds, tradeoffs, hardness condensation, XORification

ACM Reference format:

Christoph Berkholz and Jakob Nordström. 2023. Near-optimal Lower Bounds on Quantifier Depth and Weisfeiler–Leman Refinement Steps. *J. ACM* 70, 5, Article 32 (October 2023), 32 pages.

<https://doi.org/10.1145/3195257>

1 INTRODUCTION

The k -variable fragment of first-order logic L^k consists of those first-order sentences that use at most k different variables. A simple example is the L^2 sentence

$$\exists x \exists y (Exy \wedge \exists x (Eyx \wedge \exists y (Exy \wedge \exists x Eyx))) \quad (1.1)$$

A preliminary version of this work appeared in *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'16)*.

Part of the work of the first author was performed while at KTH Royal Institute of Technology supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD). The research of the second author was also partially done at KTH Royal Institute of Technology, and was supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013) / ERC grant agreement no. 279611, by the Swedish Research Council grants 621-2010-4797, 621-2012-5645, and 2016-00782, and by the Independent Research Fund Denmark grant 9040-00389B.

Authors' addresses: C. Berkholz, Technische Universität Ilmenau Ilmenau, Germany; email: christoph.berkholz@tu-ilmenau.de; J. Nordström, University of Copenhagen, Department of Computer Science, Copenhagen, Denmark and Lund University, Department of Computer Science, Lund, Sweden; email: jn@di.ku.dk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

0004-5411/2023/10-ART32 \$15.00

<https://doi.org/10.1145/3195257>

stating that there exists a directed path of length 4 in a digraph. Extending \mathbb{L}^k with counting quantifiers $\exists^{\geq 1}x$ yields \mathbb{C}^k , which can be more economical in terms of variables. As an illustration, the \mathbb{L}^8 sentence

$$\exists x \exists y_1 \cdots \exists y_7 (\bigwedge_{i \neq j} y_i \neq y_j \wedge \bigwedge_i E x y_i) \quad (1.2)$$

stating the existence of a vertex of degree at least 7 in a graph can be written more succinctly as the \mathbb{C}^2 sentence

$$\exists x \exists^{\geq 7} y E x y. \quad (1.3)$$

Bounded variable fragments of first-order logic have found numerous applications in finite model theory and related areas (see Reference [23] for a survey). Their importance stems from the fact that the model checking problem (given a finite relational structure \mathcal{A} and a sentence φ , does \mathcal{A} satisfy φ ?) can be decided in polynomial time [31, 46]. Moreover, the equivalence problem (given two finite relational structures \mathcal{A} and \mathcal{B} , do they satisfy the same sentences?) for \mathbb{L}^k and \mathbb{C}^k can be decided in time $n^{O(k)}$ [32], i.e., polynomial for constant k .

1.1 Quantifier Depth

If \mathcal{A} and \mathcal{B} are not equivalent in \mathbb{L}^k or \mathbb{C}^k , then there exists a sentence φ that defines a distinguishing property, i.e., such that $\mathcal{A} \models \varphi$ and $\mathcal{B} \not\models \varphi$, which certifies that the structures are non-isomorphic. But how complex can such a sentence be? In particular, what is the minimal quantifier depth of an \mathbb{L}^k or \mathbb{C}^k sentence that distinguishes two n -element relational structures \mathcal{A} and \mathcal{B} ? The best upper bound for the quantifier depth of \mathbb{L}^k and \mathbb{C}^k is n^{k-1} [32], while, to the best of our knowledge, the strongest lower bounds have been only linear in n [17, 20, 24]. In this article, we present a near-optimal lower bound of $n^{\Omega(k/\log k)}$.

THEOREM 1.1. *There exist $\varepsilon > 0$ and $K_0 \in \mathbb{N}$ such that for all integers k and n with $K_0 \leq k \leq n^{1/12}$ there is a pair of n -element $(k-1)$ -ary relational structures $\mathcal{A}_n, \mathcal{B}_n$ that can be distinguished in k -variable first-order logic but satisfy the same \mathbb{L}^k and \mathbb{C}^k sentences up to quantifier depth $n^{\varepsilon k / \log k}$.*

Note that any two non-isomorphic n -element σ -structures \mathcal{A} and \mathcal{B} can always be distinguished by a simple n -variable first-order sentence of quantifier depth n , namely,

$$\exists x_1 \cdots \exists x_n \left(\bigwedge_{i \neq j} x_i \neq x_j \wedge \bigwedge_{\substack{R \in \sigma, \\ (v_{i_1}, \dots, v_{i_r}) \in R^{\mathcal{A}}}} R x_{i_1}, \dots, x_{i_r} \wedge \bigwedge_{\substack{R \in \sigma, \\ (v_{i_1}, \dots, v_{i_r}) \notin R^{\mathcal{A}}}} \neg R x_{i_1}, \dots, x_{i_r} \right). \quad (1.4)$$

Since our $n^{\Omega(k/\log k)}$ lower bound for k -variable logics grows significantly faster than this trivial upper bound n on the quantifier depth as the number of variables increases, Theorem 1.1 also describes a tradeoff in the supercritical regime above worst-case investigated by Razborov [43]: If one reduces one complexity measure (the number of variables), then the other complexity parameter (the quantifier depth) increases sharply even beyond its worst-case upper bound.

The equivalence problem for \mathbb{C}^{k+1} is known to be closely related to the ***k-dimensional Weisfeiler–Leman algorithm (k-WL)*** for testing non-isomorphism of graphs and, more generally, relational structures. It was shown by Cai, Fürer, and Immerman [17] that two structures are distinguished by k -WL if and only if there exists a \mathbb{C}^{k+1} sentence that differentiates between them. Moreover, the quantifier depth of such a sentence also relates to the complexity of the WL algorithm in that the number of iterations k -WL needs to tell \mathcal{A} and \mathcal{B} apart coincides with the minimal quantifier depth of a distinguishing \mathbb{C}^{k+1} sentence. Therefore, Theorem 1.1 also implies a near-optimal lower bound on the number of refinement steps required in the Weisfeiler–Leman algorithm. We discuss this next.

1.2 The Weisfeiler–Leman Algorithm

The Weisfeiler–Leman algorithm, independently introduced by Babai in 1979 and by Immerman and Lander in Reference [32] (cf. References [17] and [3] for historic notes), is a hierarchy of methods for isomorphism testing that iteratively refine a partition (or coloring) of the vertex set, ending with a *stable coloring* that classifies *similar vertices*. Since no isomorphism can map non-similar vertices to each other, this reduces the search space. Moreover, if two structures end up with different stable colorings, then we can immediately deduce that the structures are non-isomorphic. The 1-dimensional Weisfeiler–Leman algorithm, better known as *color refinement*, initially colors the vertices according to their degree (clearly, no isomorphism identifies vertices of different degree). The vertex coloring is then refined based on the color classes of the neighbors. For example, two degree-5 vertices get different colors in the next step if they have a different number of degree-7 neighbors. This refinement step is repeated until the coloring stays stable (i.e., every pair of equally colored vertices have the same number of neighbors in every other color class). This algorithm is already quite strong and is extensively used in practical graph isomorphism algorithms.

In k -dimensional WL this idea is generalized to colorings of k -tuples of vertices. Initially the k -tuples are colored by their isomorphism type, i.e., two tuples $\vec{v} = (v_1, \dots, v_k)$ and $\vec{w} = (w_1, \dots, w_k)$ get different colors if the mapping $v_i \mapsto w_i$ is not an isomorphism on the substructures induced on $\{v_1, \dots, v_k\}$ and $\{w_1, \dots, w_k\}$. In the refinement step, we consider for each k -tuple $\vec{v} = (v_1, \dots, v_k)$ and every vertex v the isomorphism type $i(\vec{v}, v)$ of the $(k+1)$ -tuple (v_1, \dots, v_k, v) and the current colors $c(\vec{v}_j)$ of the k -tuples $\vec{v}_j := (v_1, \dots, v_{j-1}, v, v_{j+1}, \dots, v_k)$, where v is substituted at the j th position in the tuple \vec{v} . We refer to the tuple $(i(\vec{v}, v), c(\vec{v}_1), \dots, c(\vec{v}_k))$ as the *color type* $t(\vec{v}, v)$ and let v be a t -neighbor of \vec{v} if $t = t(\vec{v}, v)$. Now, two k -tuples \vec{v} and \vec{w} get different colors if they are already colored differently or if there exists a color type t such that \vec{v} and \vec{w} have a different number of t -neighbors. The refinement step is repeated until the coloring remains stable. Since in every round the number of color classes grows, the process stops after at most n^k steps. The color names can be chosen in such a way that the stable coloring is canonical, which means that two isomorphic structures end up with the same coloring, and such a canonical stable coloring can be computed in time $n^{O(k)}$.

This simple combinatorial algorithm is surprisingly powerful. Grohe [25] showed that for every nontrivial graph class that excludes some minor (such as planar graphs or graphs of bounded treewidth) there exists some k such that k -WL computes a different coloring for all non-isomorphic graphs and hence solves graph isomorphism in polynomial time on that graph class. Weisfeiler–Leman has also been used as a subroutine in algorithms that solve graph isomorphism on all graphs. As one part of his recent graph isomorphism algorithm, Babai [3] applies k -WL for polylogarithmic k to relational (k -ary) structures and makes use of the quasi-polynomial running time of this algorithm.

Given the importance of the Weisfeiler–Leman procedure, it is a natural question whether the trivial n^k upper bound on the number of refinement steps is tight. By the correspondence between the number of refinement steps of k -WL and the quantifier depth of C^{k+1} [17], our main result implies a near-optimal lower bound even up to polynomial, but still sublinear, values of k (i.e., $k = n^\delta$ for small enough constant δ).

THEOREM 1.2. *There exist $\varepsilon > 0, K_0 \in \mathbb{N}$ such that for all integers k, n with $K_0 \leq k \leq n^{1/12}$ there is an n -element k -ary relational structure \mathcal{A}_n for which the k -dimensional Weisfeiler–Leman algorithm needs $n^{\varepsilon k / \log k}$ refinement steps to compute the stable coloring.*

In addition to the near-optimal lower bounds for a specific dimension (or number of variables) k , we also obtain the following tradeoff between the dimension and the number of refinement steps: If we fix two parameters ℓ_1 and ℓ_2 (possibly depending on n) satisfying $\ell_1 \leq \ell_2 \leq n^{1/6} / \ell_1$, then there

are n -element structures such that k -WL requires $n^{\Omega(\ell_1/\log \ell_2)}$ refinement steps for all $\ell_1 \leq k \leq \ell_2$. A particularly interesting choice of parameters is $\ell_1 = \log^c n$ for some constant $c > 1$ and $\ell_2 = n^{1/7}$. This implies the following quasi-polynomial lower bound on the number of refinement steps for Weisfeiler–Leman from polylogarithmic dimension all the way up to dimension $n^{1/7}$:

THEOREM 1.3. *For every $c > 1$ there is a sequence of n -element relational structures \mathcal{A}_n for which the k -dimensional Weisfeiler–Leman algorithm needs $n^{\Omega(\log^{c-1} n)}$ refinement steps to compute the stable coloring for all k with $\log^c n \leq k \leq n^{1/7}$.*

1.3 Previous Lower Bounds and Subsequent Developments

In their seminal work [17], Cai, Fürer, and Immerman established the existence of non-isomorphic n -vertex graphs that cannot be distinguished by any first-order counting sentence with $o(n)$ variables. Since every pair of non-isomorphic n -element structures can be distinguished by a C^n (or even L^n) sentence (as shown in (1.4) above), this result also implies a linear lower bound on the quantifier depth of C^k if $k = \Omega(n)$. For all constant $k \geq 2$, a linear $\Omega(n)$ lower bound on the quantifier depth of C^k follows implicitly from an intricate construction of Grohe [24], which was used to show that the equivalence problems for L^k and C^k are complete for polynomial time. An explicit linear lower bound based on a simplified construction was subsequently presented by Fürer [20].

For the special case of $k = 2$, Kiefer and McKay [33] showed that for all large enough n there are examples of n -vertex graphs that need quantifier depth either $n - 2$ or $n - 1$ to be distinguished in C^2 . This matches the general upper bound of $n - 1$ and improves the previous $(1 - o(1))n$ lower bound of Krebs and Verbitsky [38]. For $k = 3$ the situation is quite different: The trivial upper bound of $n^2 - 1$ on the quantifier depth of a C^3 -sentence distinguishing two n -vertex graphs is not tight. This has first been shown by Kiefer and Schweitzer [34, 35], who obtained an $O(n^2/\log n)$ upper bound. The upper bound on quantifier depth for distinguishing n -vertex graphs in C^3 has later been further improved to $O(n \log n)$ by Lichter et al. [39]. Thus, while the trivial linear upper bound is tight for $k = 2$, substantial savings over the quadratic upper bound for $k = 3$ are possible. In a way, our lower bounds rule out the possibility of a significant asymptotical improvement for larger k and are the first lower bounds that are super-linear in the domain size n .

As the final version of this article was being prepared, a result [26] was announced that improves our lower bound from $n^{\Omega(k/\log k)}$ to $n^{\Omega(k)}$, thus making it asymptotically optimal.

1.4 Discussion of Techniques

The hard instances we construct are based on propositional XOR (exclusive or) formulas, which can alternatively be viewed as systems of linear equations over $\text{GF}(2)$. There is a long history of using XOR formulas for proving lower bounds in different areas of theoretical computer science such as, e.g., finite model theory, proof complexity, and combinatorial optimization/hardness of approximation. Our main technical insight is to combine two methods that, to the best of our knowledge, have not been used together before, namely, Ehrenfeucht–Fraïssé games on structures based on XOR formulas and hardness amplification by variable substitution.

More than three decades ago, Immerman [30] presented a way to encode an XOR formula into two graphs that are isomorphic if and only if the formula is satisfiable. This can then be used to show that the two graphs cannot be distinguished by a sentence with few variables or low quantifier depth using Ehrenfeucht–Fraïssé games. Arguably the most important application of this method is the result in Reference [17] establishing that a linear number of variables is needed to distinguish two graphs in first-order counting logic. Graph constructions based on XOR formulas have also been used to prove lower bounds on the quantifier depth of C^k [20, 30]. We remark

that for our result we have to use a slightly different encoding of XOR formulas into relational structures rather than graphs.

In proof complexity, various flavors of XOR formulas (usually called *Tseitin formulas* when used to encode the *handshaking lemma* saying that the sum of all vertex degrees in an undirected graph has to be an even number) have been employed to obtain lower bounds for proof systems such as resolution [45], polynomial calculus [15], and bounded-depth Frege [6, 28, 42]. Such formulas have also played an important role in many lower bounds for the Positivstellensatz/sums-of-squares proof system [22, 36, 44] corresponding to the Lasserre semidefinite programming hierarchy, which has been the focus of much recent interest in the context of combinatorial optimization.¹ Another use of XOR in proof complexity has been for hardness amplification, where one takes a (typically non-XOR) formula that is moderately hard with respect to some complexity measure, substitutes all variables by exclusive ors over pairwise distinct sets of variables, and then shows that the new *XORified* formula must be very hard with respect to some other (more important) complexity measure. This technique was perhaps first made explicit in Reference [7] (attributed there to personal communication with Alekhovich and Razborov, with a note that it is also very similar in spirit to an approach used in Reference [11]) and has later appeared in, e.g., References [5, 9, 10, 14, 19]. An even more crucial role in proof complexity is played by well-connected so-called *expander graphs*. For instance, given a formula in conjunctive normal form (CNF) one can look at its bipartite clause-variable incidence graph, or some variant of the CVIG derived from the combinatorial structure of the formula, and prove that if this graph is an expander, then this implies that the formula must be hard for proof systems such as resolution [11] and polynomial calculus [1, 40].

In a striking paper [43], the author combines XORification and expansion in a simple (with hindsight) but amazingly powerful way. Namely, instead of replacing every variable by an XOR over new, fresh variables, Razborov recycles variables from a much smaller pool, thus decreasing the total number of variables. This means that the hardness amplification proofs no longer work, at least not in their current form, since they crucially use that all new substitution variables are distinct. But here expansion comes into play. If the pattern of variable substitutions is described by a strong enough bipartite expander, then it turns out that locally there is enough “freshness” even among the recycled variables to make the hardness amplification go through over a fairly wide range of the parameter space. And, since the formula has not only become harder but has also had the number of variables decreased, this can be viewed as a kind of *hardness compression* or *hardness condensation*.

What we do in this article is to first revisit Immerman’s old quantifier depth lower bound for first-order counting logic [30] and observe that the construction can be used to obtain an improved scalable lower bound for the k -variable fragment. We then translate Razborov’s hardness condensation technique [43] into the language of finite variable logics and use it—perhaps somewhat amusingly applied to XORification of XOR formulas, which is usually not the case in proof complexity—to reduce the domain size of relational structures while maintaining the minimal quantifier depth required to distinguish them.

1.5 Outline of This Article

The rest of this article is organized as follows: In Section 2, we describe how to translate XOR formulas to relational structures and play combinatorial games on these structures. This then allows

¹No proof complexity is needed in this article, and so readers unfamiliar with these proof systems need not worry—this is just an informal overview. However, readers interested in a more in-depth treatment of this topic can consult the survey chapter cited in Reference [16] or the book cited in Reference [37].

us to state our main technical lemmas in Section 3 and show how these lemmas yield our results. Turning to the proofs of these technical lemmas, in Section 4, we present a version of Immerman's quantifier depth lower bound for XOR formulas, and in Section 5, we apply Razborov's hardness condensation technique to these formulas. Finally, in Section 6, we make some concluding remarks and discuss possible directions for future research. Some proofs of technical results needed in the article are deferred to Appendix A.

2 FROM XOR FORMULAS TO RELATIONAL STRUCTURES

In this article all structures are finite and defined over a relational signature σ . We use the letters X , E , and R for unary, binary, and r -ary relation symbols, respectively, and let $X^{\mathcal{A}}$, $E^{\mathcal{A}}$, and $R^{\mathcal{A}}$ be their interpretation in a structure \mathcal{A} . We write $V(\mathcal{A})$ to denote the domain of the structure \mathcal{A} . The k -variable fragment of first-order logic L^k consists of all first-order formulas that use at most k different variables (possibly re-quantifying them as in Equation (1.1)). We also consider k -variable first-order counting logic C^k , which is the extension of L^k by counting quantifiers $\exists^{\geq i} x \varphi(x)$, stating that there exist at least i elements $u \in V(\mathcal{A})$ such that $(\mathcal{A}, u) \models \varphi(x)$. For a survey of finite variable logics and their applications, we refer the reader to, e.g., Reference [23].

An ℓ -XOR constraint is a tuple (x_1, \dots, x_ℓ, a) consisting of ℓ distinct Boolean variables and a Boolean value $a \in \{0, 1\}$. We refer to ℓ as the *width* of the constraint. An assignment α satisfies (x_1, \dots, x_ℓ, a) if $\alpha(x_1) + \dots + \alpha(x_\ell) \equiv a \pmod{2}$. An ℓ -XOR formula F is a conjunction of XOR constraints of width at most ℓ and is satisfied by an assignment α if α satisfies all constraints in F .

For every ℓ -XOR formula F on n variables, we can define a pair of $2n$ -element structures $\mathcal{A} = \mathcal{A}(F)$ and $\mathcal{B} = \mathcal{B}(F)$ that are isomorphic if and only if F is satisfiable. The domain of the structures contains two elements x_i^0 and x_i^1 for each Boolean variable x_i . There is one unary predicate X_i for every variable x_i satisfied by the corresponding two elements x_i^0 and x_i^1 . Hence, these unary relations partition the domain of the structures into two-element sets, i.e., $X_i^{\mathcal{A}} = X_i^{\mathcal{B}} = \{x_i^0, x_i^1\}$. To encode the XOR constraints, we introduce one m -ary relation R_m for every $1 \leq m \leq \ell$ and set

$$R_m^{\mathcal{A}} = \left\{ (x_{i_1}^{a_1}, \dots, x_{i_m}^{a_m}) \mid (x_{i_1}, \dots, x_{i_m}, a) \in F, \sum_i a_i \equiv 0 \pmod{2} \right\} \quad (2.1a)$$

and

$$R_m^{\mathcal{B}} = \left\{ (x_{i_1}^{a_1}, \dots, x_{i_m}^{a_m}) \mid (x_{i_1}, \dots, x_{i_m}, a) \in F, \sum_i a_i \equiv a \pmod{2} \right\}. \quad (2.1b)$$

Every bijection β between the domains of $\mathcal{A}(F)$ and $\mathcal{B}(F)$ that preserves the unary relations X_i can be translated to an assignment α for the XOR formula via the correspondence

$$\alpha(x_i) = 0 \Leftrightarrow \beta(x_i^0) = x_i^0 \Leftrightarrow \beta(x_i^1) = x_i^1 \quad (2.2a)$$

and

$$\alpha(x_i) = 1 \Leftrightarrow \beta(x_i^0) = x_i^1 \Leftrightarrow \beta(x_i^1) = x_i^0. \quad (2.2b)$$

Moreover, it is not hard to show that such a bijection defines an isomorphism between $\mathcal{A}(F)$ and $\mathcal{B}(F)$ if and only if the corresponding assignment satisfies F . See Figure 1 for a small example illustrating the construction.

This kind of encodings of XOR formulas into relational structures has been very useful for proving lower bounds for finite variable logics in the past. Our transformation of XOR constraints of width ℓ into ℓ -ary relational structures resembles the way Gurevich and Shelah [27] encode XOR formulas as hypergraphs. It is also closely related to the way Cai, Fürer, and Immerman [17] obtain two non-isomorphic graphs \mathcal{G} and \mathcal{H} from an unsatisfiable 3-XOR formula F in the sense that \mathcal{G} and \mathcal{H} can be seen to be the incidence graphs of our structures $\mathcal{A}(F)$ and $\mathcal{B}(F)$.




Fig. 1. Structure encoding of $F = \{(x_7, x_8, 1)\}$.

To prove our main result, we make use of the combinatorial characterization of quantifier depth of finite-variable logics in terms of pebble games for L^k and C^k , which are played on two given relational structures. Since in our case the structures are based on XOR formulas, for convenience, we will consider a simplified combinatorial game that is played directly on the XOR formulas rather than on their structure encodings. We first describe this game and then show in Lemma 2.1 that this yields an equivalent characterization.

The r -round k -pebble game is played on an XOR formula F by two players, whom we will refer to as Player 1 and Player 2. A position in the game is a partial assignment α of at most k variables of F and the game starts with the empty assignment. In each round, Player 1 can delete some variable assignments from the current position (he chooses some $\alpha' \subseteq \alpha$). If the current position assigns values to exactly k variables, then Player 1 has to delete at least one variable assignment. Afterwards, he chooses some currently unassigned variable x and asks Player 2 for its value. She answers by either 0 or 1 (independently of any previous answers to the same question) and adds this variable assignment to the current position.

A winning position for Player 1 is an assignment falsifying some constraint in F . Player 1 wins the r -round k -pebble game if he has a strategy to win every play of the k -pebble game within at most r rounds. Otherwise, we say that Player 2 wins (or survives) the r -round k -pebble game. Player 1 *wins the k -pebble game* if he wins the r -round k -pebble game within a finite number of rounds r . Note that if Player 1 wins the k -pebble game, then he can always win the k -pebble within $2^k n^{k+1}$ rounds, because there are at most $\sum_{i=0}^k 2^i \binom{n}{i} \leq 2^k n^{k+1}$ different positions with at most k pebbles on n -variable XOR formulas. We say that Player 1 *can reach a position β* from a position α within r rounds if he has a strategy such that in every play of the r -round k -pebble game starting from position α he either wins or ends up with position β .

As a side remark, we note that if we expand the XOR formula to CNF, then our pebble game is the same as the so-called *Boolean existential pebble game* played on this CNF encoding and therefore also characterizes the resolution width required for the corresponding CNF formula as shown in Reference [2]. Intuitively, it is this correspondence that enables us to apply the proof complexity techniques from Reference [43] in our setting. We will not need to use any concepts from proof complexity in this article, however, but will present a self-contained proof, and so we do not elaborate further on this connection.

Let us now show that the game described above is equivalent to the pebble game for L^k and to the bijective pebble game for C^k played on the structures $\mathcal{A}(F)$ and $\mathcal{B}(F)$.

LEMMA 2.1. *Let k, p, r be integers such that $r > 0$ and $k \geq p$ and let F be a p -XOR formula giving rise to structures $\mathcal{A} = \mathcal{A}(F)$ and $\mathcal{B} = \mathcal{B}(F)$ as described in the paragraph preceding Equations (2.1a) and (2.1b). Then the following statements are equivalent:*

- (a) *Player 1 wins the r -round k -pebble game on F .*
- (b) *There is a k -variable first-order sentence $\varphi \in L^k$ of quantifier depth r such that $\mathcal{A}(F) \models \varphi$ and $\mathcal{B}(F) \not\models \varphi$.*

- (c) There is a k -variable sentence in first-order counting logic $\varphi \in C^k$ of quantifier depth r such that $\mathcal{A}(F) \models \varphi$ and $\mathcal{B}(F) \not\models \varphi$.
- (d) The $(k-1)$ -dimensional Weisfeiler–Leman procedure can distinguish between $\mathcal{A}(F)$ and $\mathcal{B}(F)$ within r refinement steps.

PROOF. Let us start by briefly recalling known characterizations in terms of Ehrenfeucht–Fraïssé games of L^k [4, 31] and C^k [17, 29]. In both cases the game is played by two players, called Spoiler and Duplicator, on the two structures \mathcal{A} and \mathcal{B} . Positions in the games are partial mappings $p = \{(u_1, v_1), \dots, (u_i, v_i)\}$ from $V(\mathcal{A})$ to $V(\mathcal{B})$ of size at most k . The games start from the empty position and proceed in rounds. At the beginning of each round in both games, Spoiler chooses $p' \subseteq p$ with $|p'| < k$.

- In the L^k -game, Spoiler then selects either some $u \in V(\mathcal{A})$ or some $v \in V(\mathcal{B})$ and Duplicator responds by choosing an element $v \in V(\mathcal{B})$ or $u \in V(\mathcal{A})$ in the other structure.
- In the C^k -game, Duplicator first selects a global bijection $f : V(\mathcal{A}) \rightarrow V(\mathcal{B})$ and Spoiler chooses some pair $(u, v) \in f$. (If $|V(\mathcal{A})| \neq |V(\mathcal{B})|$, then Spoiler wins the C^k -game immediately.)

The new position is $p' \cup \{(u, v)\}$. Spoiler wins the r -round L^k / C^k game if he has a strategy to reach within r rounds a position p that does not define an isomorphism on the induced substructures.² Both games characterize equivalence in the corresponding logics: Spoiler wins the r -round L^k / C^k game if and only if there is a sentence $\varphi \in L^k / C^k$ of quantifier depth r such that $\mathcal{A} \models \varphi$ and $\mathcal{B} \not\models \varphi$.

When these games are played on the two structures $\mathcal{A}(F)$ and $\mathcal{B}(F)$ obtained from an XOR formula F , it is not hard to verify that both games are equivalent to the k -pebble game on F . To see this, we identify Spoiler with Player 1, Duplicator with Player 2, and partial mappings $p = \{(x_i^{a_i}, x_i^{b_i}) \mid i \leq \ell\}$ with partial assignments $\alpha = \{x_i \mapsto a_i \oplus b_i \mid i \leq \ell\}$. Because of the X_i -relations, we can assume that partial assignments of any other form will not occur, as they are losing positions for Duplicator. To formalize this proof strategy, we observe that the direction (b) \Rightarrow (c) is trivial and handle the two remaining cases (a) \Rightarrow (b) and (c) \Rightarrow (a) by proving the contrapositive. Since the equivalence between (c) and (d) was proven in Reference [17], this completes the proof.

To obtain $\neg(\text{b}) \Rightarrow \neg(\text{a})$, assume that Duplicator has a strategy to survive the L^k -game for r rounds on $\mathcal{A}(F)$ and $\mathcal{B}(F)$. To devise a strategy for Player 2 on F , we identify every position $\alpha = \{x_i \mapsto b_i \mid i \leq \ell\}$ in the k -pebble game on F with the position $p(\alpha) = \{(x_i^0, x_i^{b_i}) \mid i \leq \ell\}$ in the L^k -game. The strategy of Player 2 is now as follows: If Player 1 asks for x_i , then Player 2 answers with the value $b \in \{0, 1\}$ such that $x_i^b \in V(\mathcal{B}(F))$ is the answer of Duplicator when Spoiler pebbles $x_i^0 \in V(\mathcal{A}(F))$. Recall that the X_i -relations force Duplicator to answer with one of the two elements x_i^0, x_i^1 in her strategy. By the definition of the structures, a position α on F is non-losing for Player 2 if and only if $p(\alpha)$ is non-losing for Duplicator.

For the direction $\neg(\text{a}) \Rightarrow \neg(\text{c})$ in the lemma, let us suppose that Player 2 has a strategy to survive the k -pebble game for r rounds. For every position $p = \{(x_i^{a_i}, x_i^{b_i}) \mid i \leq \ell\}$ in the C^k -game, we let $\alpha(p) = \{x_i \mapsto a_i \oplus b_i \mid i \leq \ell\}$ be the corresponding position in the k -pebble game. We develop Duplicator's strategy in the C^k -game and let p with $|p| < k$ be the position in the current round. We have to show that there is a bijection $f : V(\mathcal{A}) \rightarrow V(\mathcal{B})$ such that for every choice of $x_i^a \in V(\mathcal{A}(F))$

²This implies that in the C^k -game, we can assume that Duplicator always chooses a bijection f that is consistent with p' , as otherwise Spoiler can win immediately.

the position $\alpha(p \cup \{(x_i^a, f(x_i^a))\})$ is the successor of $\alpha(p)$ in the k -pebble game on F when Player 1 asks for x_i . To construct the bijection f , we let c_i be the answer according to the strategy of Player 2 when asked for x_i in position $\alpha(p)$ and set $f(x_i^a) = x_i^{a \oplus c_i}$ for all $i \in [n]$ and $a \in \{0, 1\}$. Note that with this definition, we get $\alpha(p \cup \{(x_i^a, f(x_i^a))\}) = \alpha(p) \cup \{x_i \mapsto c_i\}$ for all $i \in [n]$ and $a \in \{0, 1\}$. Since a position p is non-losing for Duplicator if and only if $\alpha(p)$ is non-losing for Player 2, the lemma is proven. \square

3 TECHNICAL LEMMAS AND PROOFS OF MAIN THEOREMS

To prove our lower bounds of the quantifier depth of finite variable logics in Theorem 1.1 and the number of refinement steps of the Weisfeiler–Leman algorithm in Theorems 1.2 and 1.3, we utilize the characterization in Lemma 2.1 and show that there are n -variable XOR formulas on which Player 1 is able to win the k -pebble game but cannot do so in significantly less than $n^{k/\log k}$ rounds. The next lemma states this formally and also provides a tradeoff as the number of pebbles increases.

LEMMA 3.1 (MAIN TECHNICAL LEMMA). *There is an absolute constant $K_0 \in \mathbb{N}^+$ such that for integers $k_{\text{lo}}, k_{\text{hi}}$, and n satisfying $K_0 \leq k_{\text{lo}} \leq k_{\text{hi}} \leq n^{1/6}/k_{\text{lo}}$ there is an XOR formula F with n variables such that Player 1 wins the k_{lo} -pebble game on F , but does not win the k_{hi} -pebble game within $n^{k_{\text{lo}}/(10 \log k_{\text{hi}})-1/5}$ rounds.*

Note that there is a limit to how far k_{lo} and k_{hi} can be from each other for the lemma to make sense—the statement becomes vacuous if $k_{\text{lo}} \leq 2 \log k_{\text{hi}}$. Let us see how this lemma yields the theorems in Section 1.

PROOF OF THEOREM 1.1. This theorem follows immediately from Lemmas 2.1 and 3.1, but let us write out the details for clarity. By setting $k_{\text{lo}} = k_{\text{hi}} = k$ in Lemma 3.1, we can find XOR formulas with n variables such that Player 1 wins the k -pebble game on F_n but needs more than $n^{\varepsilon k/\log k}$ rounds to do so (provided we choose $\varepsilon < 1/10$ and K_0 large enough). We can then plug these XOR formulas into Lemma 2.1 to obtain n -element structures $\mathcal{A}_n = \mathcal{A}(F_n)$ and $\mathcal{B}_n = \mathcal{B}(F_n)$ that can be distinguished in the k -variable fragments of first-order logic L^k and first-order counting logic C^k , but where this requires sentences of quantifier depth at least $n^{\varepsilon k/\log k}$. \square

PROOF OF THEOREM 1.2. If we let F_n be the XOR formula from Lemma 3.1 for $k_{\text{lo}} = k_{\text{hi}} = k + 1$, then by Lemma 2.1 it holds that the structures $\mathcal{A}(F_n)$ and $\mathcal{B}(F_n)$ will be distinguished by the k -dimensional Weisfeiler–Leman algorithm, but only after $n^{\varepsilon(k+1)/\log(k+1)} \geq n^{\varepsilon k/\log k}$ refinement steps. Hence, computing the stable coloring of either of these structures requires at least $n^{\varepsilon k/\log k}$ refinement steps (since they would be distinguished earlier if at least one of the computations terminated earlier). \square

PROOF OF THEOREM 1.3. This is similar to the proof of Theorem 1.2, but setting $k_{\text{lo}} = \lfloor \log n^\varepsilon \rfloor + 1$ and $k_{\text{hi}} = \lceil n^{1/7} \rceil + 1$ in Lemma 3.1. \square

The proof of the tradeoff between the number of pebbles versus number of rounds in Lemma 3.1 splits into two steps. We first establish a rather weak lower bound on the number of rounds in the pebble game played on suitably chosen m -variable XOR formulas for $m \gg n$. We then transform this into a much stronger lower bound for formulas over n variables using hardness condensation. To help the reader keep track of which results are proven in which setting, in what follows, we will write ℓ_{lo} and ℓ_{hi} to denote parameters depending on m and k_{lo} and k_{hi} to denote parameters depending on n .

To implement the first step in our proof plan, we use tools developed by Immerman [30] to establish a lower bound as stated in the next lemma.

LEMMA 3.2. *For all $\ell_{\text{hi}}, m \geq 3$ there is an m -variable 3-XOR formula $F_m^{\ell_{\text{hi}}}$ on which Player 1*

- (a) *wins the 3-pebble game, but*
 (b) *does not win the ℓ_{hi} -pebble game within $\max(3, \frac{1}{\lceil \log \ell_{\text{hi}} \rceil}) m^{1/(1+\lceil \log \ell_{\text{hi}} \rceil)} - 2$ rounds.*

We defer the proof of Lemma 3.2 to Section 4, but at this point an expert reader might wonder why we would need to prove this lower bound at all, since a much stronger $\Omega(m)$ bound on the number of rounds in the pebble game on 4-XOR formulas was already obtained by Fürer [20]. The reason is that in Fürer's construction Player 1 cannot win the game with few pebbles. However, it is crucial for the second step of our proof, where we boost the lower bound but also significantly increase the number of pebbles that are needed to win the game, that Player 1 is able to win the original game with very few pebbles.

The second step in the proof of our main technical lemma is carried out by using the techniques developed by Razborov [43] and applying them to the XOR formulas in Lemma 3.2. Roughly speaking, if we set $k_{\text{lo}} = k_{\text{hi}} = k$ for simplicity, then the number of variables decreases from m to $n \approx m^{1/k}$, whereas the $m^{1/\log k}$ round lower bound for the k -pebble game stays essentially the same and hence becomes $n^{k/\log k}$ in terms of the new number of variables n . The properties of hardness condensation are summarized in the next lemma, which we prove in Section 5. To demonstrate the flexibility of this tool, we state the lemma in its most general form—readers who want to see an example of how to apply it to the XOR formulas in Lemma 3.2 can mentally fix $p = 3$, $\ell_{\text{lo}} = 3$, $r \approx m^{1/\log \ell_{\text{hi}}}$, and $\Delta \approx \ell_{\text{hi}}/3$ when reading the statement of the lemma below.

LEMMA 3.3 (HARDNESS CONDENSATION LEMMA). *There exists an absolute constant $\Delta_0 \in \mathbb{N}^+$ such that the following holds: Let F be an m -variable p -XOR formula and suppose that we can choose parameters $\ell_{\text{lo}} > 0$, $\ell_{\text{hi}} \geq \Delta_0 \ell_{\text{lo}}$ and r such that Player 1*

- (a) *has a winning strategy for the ℓ_{lo} -pebble game on F , but*
 (b) *does not win the ℓ_{hi} -pebble game on F within r rounds.*

Then for any Δ satisfying $\Delta_0 \leq \Delta \leq \ell_{\text{hi}}/\ell_{\text{lo}}$ and $(2\ell_{\text{hi}}\Delta)^{2\Delta} \leq m$ there is a (Δp) -XOR formula H with $\lceil m^{3/\Delta} \rceil$ variables such that Player 1

- (a) *has a winning strategy for the $(\Delta\ell_{\text{lo}})$ -pebble game on H , but*
 (b) *does not win the ℓ_{hi} -pebble game on H within $r/(2\ell_{\text{hi}})$ rounds.*

Taking Lemmas 3.2 and 3.3 on faith for now, we are ready to prove our main technical lemma yielding an $n^{\Omega(k/\log k)}$ lower bound on the number of rounds in the k -pebble game.

PROOF OF LEMMA 3.1 Let Δ_0 be the constant in Lemma 3.3. We let

$$K_0 \geq 3\Delta_0 + 9 \tag{3.1}$$

be an absolute constant, the precise value of which will be determined by calculations later in the proof. We are given k_{hi} , k_{lo} , and n satisfying the conditions

$$K_0 \leq k_{\text{lo}} \leq k_{\text{hi}} \leq n^{1/6}/k_{\text{lo}} \tag{3.2}$$

in Lemma 3.1. Let us set

$$\ell_{\text{hi}} := k_{\text{hi}} \tag{3.3a}$$

and

$$m := n^{\lfloor k_{\text{lo}}/9 \rfloor} \tag{3.3b}$$

and apply Lemma 3.2 (which is in order, since $\ell_{\text{lo}} \geq 3$ and $m \geq 3$ by inequalities (3.1) and (3.2)). This yields an m -variable 3-XOR formula on which Player 1 wins the 3-pebble game but cannot

win the ℓ_{hi} -pebble game within

$$r := \frac{1}{\lceil \log \ell_{\text{hi}} \rceil} m^{1/(1+\lceil \log \ell_{\text{hi}} \rceil)} - 2 \quad (3.3c)$$

rounds. As a side remark, we note that this lower bound term might vanish if k_{lo} and k_{hi} were to far apart from each other ($k_{\text{lo}} \leq 2 \log k_{\text{hi}}$), but recall that in this case the statement of Lemma 3.1 becomes vacuous anyway. Therefore, in what follows, we assume without loss of generality that $r \geq 1$. Now, we can apply hardness condensation as in Lemma 3.3 to the formula provided by Lemma 3.2, where we fix parameters

$$p := 3, \quad (3.4a)$$

$$\ell_{\text{lo}} := 3, \quad (3.4b)$$

and

$$\Delta := 3 \lfloor k_{\text{lo}}/9 \rfloor. \quad (3.4c)$$

To verify that our choice of parameters is legal, note that in addition to $r \geq 1$ we also have $\ell_{\text{lo}} > 0$ and

$$\ell_{\text{hi}} = k_{\text{hi}} \geq K_0 > 3\Delta_0 = \Delta_0 \ell_{\text{lo}}. \quad (3.5)$$

Thus, the assumptions needed for (a) and (b) are satisfied by the XOR formula obtained from Lemma 3.2. To confirm that Δ chosen as in Equation (3.4c) satisfies the conditions in Lemma 3.3, observe that

$$\Delta_0 \leq 3 \lfloor K_0/9 \rfloor \leq 3 \lfloor k_{\text{lo}}/9 \rfloor = \Delta \leq k_{\text{lo}}/3 \leq k_{\text{hi}}/3 = \ell_{\text{hi}}/\ell_{\text{lo}}. \quad (3.6)$$

Furthermore, since $\Delta \leq k_{\text{lo}}/3$ and $\ell_{\text{hi}} = k_{\text{hi}} \leq n^{1/6}/k_{\text{lo}}$, we get

$$(2\ell_{\text{hi}}\Delta)^{2\Delta} \leq \left(\frac{2}{3}n^{1/6}\right)^{2\Delta} \leq n^{\Delta/3} = m. \quad (3.7)$$

Since $m^{3/\Delta} = n$, Lemma 3.3 provides us with an n -variable XOR formula on which according to (a') Player 1 has a winning strategy for the (3Δ) -pebble game and hence also for the game with $k_{\text{lo}} \geq 9 \lfloor k_{\text{lo}}/9 \rfloor = 3\Delta$ pebbles. Moreover, by (b') it holds that Player 1 needs more than $r/(2\ell_{\text{hi}}) = r/(2k_{\text{hi}})$ rounds to win the k_{hi} -pebble game. To complete the proof, we observe that if we choose K_0 large enough, then for $n > k_{\text{hi}} \geq k_{\text{lo}} \geq K_0$ it holds that

$$\begin{aligned} \frac{r}{2k_{\text{hi}}} &= \frac{1}{2k_{\text{hi}} \lceil \log k_{\text{hi}} \rceil} n^{\lfloor k_{\text{lo}}/9 \rfloor / (1 + \lceil \log k_{\text{hi}} \rceil)} - \frac{1}{k_{\text{hi}}} && \left[\text{by (3.3b) and (3.3c)} \right] \\ &\geq \frac{6n^{1/5}}{n^{1/6} \log n} n^{\lfloor k_{\text{lo}}/9 \rfloor / (1 + \lceil \log k_{\text{hi}} \rceil) - 1/5} - \frac{1}{k_{\text{hi}}} && \left[\text{since } k_{\text{hi}} \leq n^{1/6} \right] \\ &\geq n^{k_{\text{lo}} / (10 \log k_{\text{hi}}) - 1/5} && \left[\text{for large enough } n, k_{\text{hi}}, \text{ and } k_{\text{lo}}. \right] \end{aligned} \quad (3.8)$$

We now choose the constant K_0 large enough so all conditions encountered in the calculations above are valid. This establishes the lemma. \square

4 XOR FORMULAS OVER HIGH-DIMENSIONAL PYRAMIDS

We now proceed to establish the k -pebble game lower bound stated in Lemma 3.2. Our XOR formulas will be constructed over directed acyclic graphs (DAGs), as described in the following definition:

Definition 4.1. Let \mathcal{G} be a DAG with sources S and a unique sink z . The XOR formula $\text{xor}(\mathcal{G})$ contains one variable v for every vertex $v \in V(\mathcal{G})$ and consists of the following constraints:

- (a) $(s, 0)$ for every source $s \in S$,

- (b) $(v, w_1, \dots, w_\ell, 0)$ for all non-sources $v \in V(\mathcal{G}) \setminus S$ with in-neighbors $N^-(v) = \{w_1, \dots, w_\ell\}$,
- (c) $(z, 1)$ for the unique sink z .

Note that the formula $xor(\mathcal{G})$ is always unsatisfiable, since all source vertices are forced to 0 by (a), which forces all other vertices to 0 in topological order by (b), contradicting (c) for the sink. Incidentally, these formulas are somewhat similar to the *pebbling formulas* defined in Reference [11], which have been very useful in proof complexity (see the surveys cited in References [16, 41] for more details). The difference is that pebbling formulas state that a vertex v is true if and only if all of its in-neighbors are true, whereas $xor(\mathcal{G})$ states that v is true if and only if the parity of the number of true in-neighbors is odd.

It is clear that one winning strategy for Player 1 is to ask first about the sink z , for which Player 2 has to answer 1 (or lose immediately) and then about all the in-neighbors of the sink until the answer for one vertex v is 1 (if there is no such vertex, then Player 2 again loses immediately). At this point, Player 1 can forget all other vertices and then ask about the in-neighbors of v until a 1-labelled vertex w is found and then continue in this way to trace a path of 1-labelled vertices backwards through the DAG until some source s is reached, which contradicts the requirement that s should be labelled 0. Formalizing this as an induction proof on the depth of \mathcal{G} shows that if the in-degree is bounded, then Player 1 can win the pebble game on $xor(\mathcal{G})$ with few pebbles, as stated in the next lemma.

LEMMA 4.2. *Let \mathcal{G} be a DAG with a unique sink and maximal in-degree d . Then Player 1 wins the $(d + 1)$ -pebble game on $xor(\mathcal{G})$.*

As a warm-up for the proof of Lemma 3.2, let us describe a weak lower bound from Reference [30] for the complete binary tree of height h (with edges directed from the leaves to the root), which we will denote \mathcal{T}_h . By the lemma above, Player 1 wins the 3-pebble game on $xor(\mathcal{T}_h)$ in $O(h)$ steps by propagating 1 from the root down to some leaf. However, Player 2 has the freedom to decide on which path she answers 1. Hence, she can safely respond 0 for a vertex v as long as it holds for the lowest 1-labelled vertex w that we can find a non-pebbled leaf with a pebble-free path leading to w without passing through v . In particular, if Player 2 is asked about vertices at least ℓ layers below the lowest pebbled vertex for which the answer 1 was given, then she can answer 0 for $2^\ell - 1$ queries. It follows that the height h provides a lower bound on the number of rounds Player 1 needs to win the game, even if he has an infinite amount of pebbles. We remark that this proof in terms of pebble-free paths is somewhat reminiscent of an argument by Cook [18] for the so-called black pebble game corresponding to the pebbling formulas in Reference [11] briefly discussed above.

The downside of this lower bound is that the height is only logarithmic in the number of vertices and thus too weak for us, as we are shooting for a lower bound of the order of $n^{1/\log k}$. To get a better bound for the black pebble game, Cook instead considered so-called pyramid graphs as in Figure 2(a). These will not be sufficient to obtain strong enough lower bounds for our pebble game, however.³ Instead, following Immerman, we consider a kind of high-dimensional

³For readers knowledgeable in pebbling, we comment that the problem is that the open-path argument in Reference [18] does not work in a DAG-like setting for the XOR pebble game. To see this, consider a pyramid with a vertex row u, v, w and a second row p, q, r, s immediately below such that the edges are $(p, u), (q, u), (q, v), (r, v), (r, w), (s, w)$. Then, if the values of u, w on the upper row and p, s on the lower row are known, there is still an open path via (q, v) or (r, v) , which is enough for the black pebbling lower bound for pyramids in Reference [18]. But in the XOR pebble game this means that r and q are already fixed because of the XOR constraints, and so there is no “open path” with unconstrained vertices.

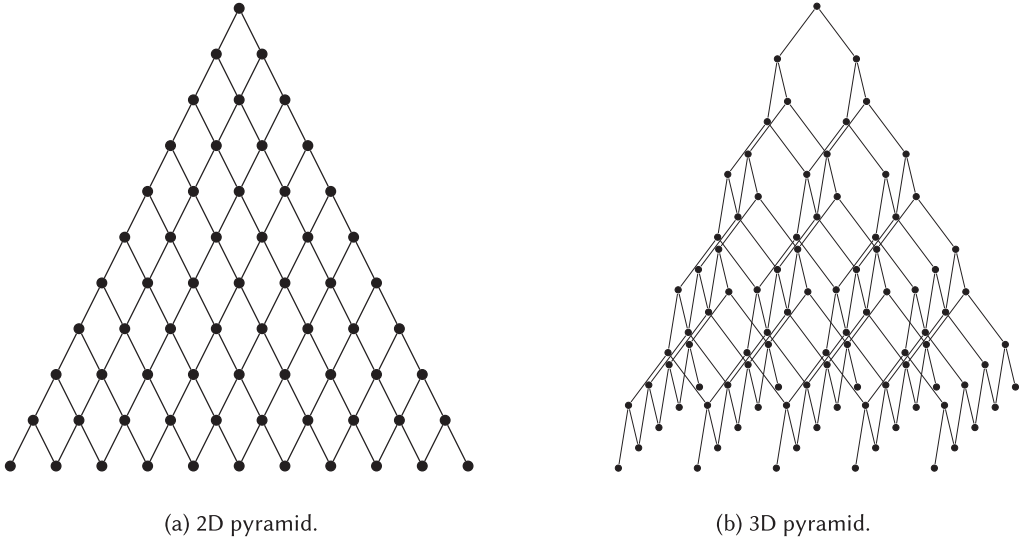


Fig. 2. Examples of high-dimensional pyramids (where all edges are directed upwards).

generalization of these graphs, for which the lower bound on the number of rounds in the k -pebble game is still linear in the height h while the number of vertices is roughly $h^{\log k}$.

Definition 4.3 ([30]). For $d \geq 1$, we define the $(d + 1)$ -dimensional pyramid of height h , denoted by \mathcal{P}_h^d , to be the following layered DAG. We let L , $0 \leq L \leq h$ be the layer number and set $q_d(L) := \lfloor L/d \rfloor$ and $r_d(L) := L \pmod{d}$. Hence, for any L , we have $L = q_d(L) \cdot d + r_d(L)$. For integers $x_i \geq 0$ the vertex set is

$$V(\mathcal{P}_h^d) = \{(x_0, \dots, x_{d-1}, L) \mid L \leq h; x_i \leq q_d(L) + 1 \text{ if } i < r_d(L); x_i \leq q_d(L) \text{ if } i \geq r_d(L)\}, \quad (4.1a)$$

where we say that L is the layer of the vertex (x_0, \dots, x_{d-1}, L) . The edge set $E(\mathcal{P}_h^d)$ consists of the set of vertex pairs

$$\begin{aligned} & \left((x_0, \dots, x_{r_d(L)}, \dots, x_{d-1}, L + 1), (x_0, \dots, x_{r_d(L)}, \dots, x_{d-1}, L) \right), \\ & \left((x_0, \dots, x_{r_d(L) + 1}, \dots, x_{d-1}, L + 1), (x_0, \dots, x_{r_d(L)}, \dots, x_{d-1}, L) \right) \end{aligned} \quad (4.1b)$$

for all vertices $(x_0, \dots, x_{d-1}, L) \in V(\mathcal{P}_h^d)$ and layers $L < h$, so every vertex in layer L has exactly two in-neighbors from layer $L + 1$.

It might be easier to parse Definition 4.3 by noting that the (kd) th layer of \mathcal{P}_h^d is a d -dimensional cube of side length k . Intuitively, we then want to have incoming edges to each vertex u at the (kd) th layer from all vertices v in the d -dimensional cube of side length $k + 1$ such that all coordinates of v are at distance 0 or +1 from the coordinates of u . This would give a fan-in larger than 2, however, and to avoid this, we expand in one dimension at a time to obtain a sequence of multidimensional cuboids where in each consecutive cuboid the side length increases by one in one dimension, until d layers later, we have a complete cube with side length $k + 1$. We refer the reader to Figure 2(a) for an illustration of a 2-dimensional pyramid (i.e., a standard pyramid graph) generated by stacking 1-dimensional cubes on top of one another and to Figure 2(b) for a 3-dimensional pyramid generated from 2-dimensional cuboids (where all the edges in the figures are assumed to be directed upwards). The vertex $(0, \dots, 0)$ at the top of the pyramid is the unique

sink and all vertices at the bottom layer h are sources. Observe that it follows from the definition that $|V(\mathcal{P}_h^d)| \leq (h+1)^{d+1}$.

As high-dimensional pyramids have in-degree 2, Lemma 4.2 implies that Player 1 wins the 3-pebble game on \mathcal{P}_h^d . Recall that, as discussed in the proof sketch of the lemma, Player 1 starts his winning strategy in the 3-pebble game by pebbling the sink of the pyramid and its two in-neighbors. One of them has to be labelled 1. Then he picks up the two other pebbles and pebbles the two in-neighbors of the vertex marked with 1 and so on. Continuing this strategy, he is able to “move” the 1 all the way to the bottom, reaching a contradiction, in a number of rounds that is linear in the height of the pyramid. This strategy turns out to be nearly optimal in the sense that to move a 1 from the top to the bottom in \mathcal{P}_h^d , as long as the total number of available pebbles is at most 2^d , it makes no sense for Player 1 at any point in the game to pebble a vertex that is d or more levels away from the lowest level containing a pebble.

The next lemma states a key property of pyramids in this regard. To state it, we first need to make a definition.

Definition 4.4. We refer to a partial assignment \mathcal{M} of Boolean values to the vertices of a DAG \mathcal{G} as a *labelling* or *marking* of \mathcal{G} . We say that \mathcal{M} is *consistent* if no constraint of type (b) or (c) (i.e., a constraint on a non-source vertex) in the XOR formula $xor(\mathcal{G})$ in Definition 4.1 is falsified by \mathcal{M} . We also say that \mathcal{M}' is *consistent with \mathcal{M}* if $\mathcal{M} \cup \mathcal{M}'$ is a partial assignment yielding a consistent labelling of \mathcal{G} .

That is, a consistent labelling does not violate any constraint on any non-source vertex, but source vertex constraints (a) may be falsified. Such labellings are easy to find for high-dimensional pyramids.

LEMMA 4.5 ([30]). *Let \mathcal{M} be any consistent labelling of all vertices in a pyramid \mathcal{P}_h^d from layer 0 to layer L . Then for every set S of $2^d - 1$ vertices on or below layer $L + d$ there is a consistent labelling of the entire pyramid that extends \mathcal{M} and labels all vertices in S with 0.*

To get some intuition why Lemma 4.5 holds, note that the d -dimensional pyramids are constructed in such a way that they locally look like binary trees. In particular, every vertex $v \in V(\mathcal{P}_h^d)$ together with all its predecessors at distance at most d form a complete binary tree. By the same argument as for the binary trees above, it follows that if v is labelled with 1, then Player 2 can safely answer 0 up to $2^d - 1$ times when asked about vertices d layers below v . However, the full proof of Lemma 4.5 is more challenging and requires some quite subtle reasoning. For the convenience of the reader, we now present a slightly modified version of the proof in Reference [30] with notation and terminology adapted to this article.

To formalize the intuitive argument above, we need some additional notation and technical definitions. Let us use the shorthand $\vec{x} = (x_0, \dots, x_{d-1})$. For a pyramid \mathcal{P}_h^d , a coordinate $j \in \{0, \dots, d-1\}$, and a layer L , we let

$$slen(j, L) := \max\{x_j \mid (\vec{x}, L) \in V(\mathcal{P}_h^d)\} \quad (4.2)$$

be the side length in the j th dimension of the cuboid in layer L , i.e., the maximal value that can be achieved in the j th coordinate in layer L , and for $L' \geq L$, we write

$$\Delta_{slen}(j, L, L') := slen(j, L') - slen(j, L) \quad (4.3)$$

to denote how much the cuboids in \mathcal{P}_h^d grow in the j th dimension in between layers L and L' .

We define the *frustum* $\mathcal{P}_{L,h}^d$ to be the subgraph of \mathcal{P}_h^d induced on the set $\{(\vec{x}, L') \mid L' \geq L\}$ of all vertices on layer L and below. We say that the *wedge* $\mathcal{W}(j, a, L)$ is the subgraph of \mathcal{P}_h^d induced




Fig. 3. Wedge $\mathcal{W}(0, 3, 5)$ \circ ; restricted frustums $\mathcal{P}_{5,10}^1[\{0\}, \emptyset, \{0 \mapsto 3\}]$ \times and $\mathcal{P}_{5,10}^1[\emptyset, \{0\}, \{0 \mapsto 3\}]$ \square .

on the vertices $(x_0, \dots, x_{j-1}, a, x_{j+1}, \dots, x_{d-1}, L)$ with fixed j th coordinate $x_j = a$ together with all predecessors of these vertices. That is, the vertex set of $\mathcal{W}(j, a, L)$ is

$$V(\mathcal{W}(j, a, L)) = \{(\vec{x}, L') \in V(\mathcal{P}_h^d) \mid L' \geq L, a \leq x_j \leq a + \Delta_{slen}(j, L, L')\}. \quad (4.4)$$

An important part in our proof will be played by subgraphs obtained by deleting wedges from frustums. We define these subgraphs next.

Fix a frustum $\mathcal{P}_{L,h}^d$, two disjoint subsets of coordinates $I_{l_0}, I_{h_i} \subseteq \{0, \dots, d-1\}$, $I_{l_0} \cap I_{h_i} = \emptyset$, and a mapping $\alpha : I_{l_0} \cup I_{h_i} \rightarrow \mathbb{N}_0$ such that $\alpha(j) \leq slen(j, L)$. We let the *restricted frustum* $\mathcal{P}_{L,h}^d[I_{l_0}, I_{h_i}, \alpha]$ be the subgraph of the frustum $\mathcal{P}_{L,h}^d$ induced on the vertex set

$$\{(\vec{x}, L') \in V(\mathcal{P}_{L,h}^d) \mid \forall j \in I_{l_0} : x_j > \alpha(j) + \Delta_{slen}(j, L, L'); \forall j \in I_{h_i} : x_j < \alpha(j)\}. \quad (4.5)$$

Note that no coordinates in $I_{l_0} \cup I_{h_i}$ are expanded, i.e., the cuboids will not grow in size in dimensions $I_{l_0} \cup I_{h_i}$ as we move down the layers. For dimensions in I_{h_i} the coordinate set stays the same, and for dimensions in I_{l_0} the coordinate set shifts by an additive $+1$ every time the pyramid graph grows in this direction. We say that a layered directed graph is a (d, q) -*frustum* if it is a restricted frustum $\mathcal{P}_{L,h}^d[I_{l_0}, I_{h_i}, \alpha]$ where q coordinates are restricted, i.e., $|I_{l_0} \cup I_{h_i}| = q$.

To see how restricted frustums are obtained by deleting wedges from frustums, note that after removing the wedge $\mathcal{W}(j, a, L)$ from the frustum $\mathcal{P}_{L,h}^d$, the remaining graph is the disjoint union of the restricted frustums $\mathcal{P}_{L,h}^d[\{j\}, \emptyset, \{j \mapsto a\}]$ and $\mathcal{P}_{L,h}^d[\emptyset, \{j\}, \{j \mapsto a\}]$. Figures 3 and 4 show 2D and 3D pyramids with a wedge \circ and a restricted frustums \times and \square .

We prove Lemma 4.5 by inductively cutting the pyramid into a wedge and restricted frustums to the left and right of this wedge. It will be convenient to focus on (d, q) -frustums which grow in the dimensions corresponding to the topmost $d - q$ layers (as the one in Figure 4). More formally, we say that an (d, q) -frustum $\mathcal{P}_{L,h}^d[I_{l_0}, I_{h_i}, \alpha]$ is *top-expanding* if

$$\{(L + j) \bmod d \mid 0 \leq j \leq d - 1 - q\} \cap (I_{l_0} \cup I_{h_i}) = \emptyset. \quad (4.6)$$

As we have done for digraphs with unique sinks in Definition 4.1, we identify with each (restricted) frustum \mathcal{P} the corresponding XOR formula $xor(\mathcal{P})$ containing all constraints given in




Fig. 4. Pyramid with wedge $\mathcal{W}(0, 2, L + 1) \circ$ and restricted frustum $\mathcal{P}_{L+1,10}^2[\emptyset, \{0\}, \{0 \mapsto 2\}]\times$.

Definitions 4.1(a) and (b). We do not include constraints for the vertices in the top layer as in (c), however, but instead provide a labelling assigning values to all vertices in that layer.⁴

We now state our inductive claim. Lemma 4.5 follows immediately once this claim has been established, as the subgraph of a pyramid \mathcal{P}_h^d on or below layer L is an (unrestricted) top-expanding $(d, 0)$ -frustum $\mathcal{P}_{L,h}^d$ and, in particular, $xor(\mathcal{P}_h^d)$ with all vertices on or above layer L consistently labelled is equivalent to $xor(\mathcal{P}_{L,h}^d)$ with the same labelling of layer L .

CLAIM 4.6. *Let \mathcal{P} be a top-expanding (d, q) -frustum and \mathcal{M}_L be any labelling of its top layer L . Then for every set S of $2^{d-q} - 1$ vertices on or below layer $L + d - q$ there is a consistent labelling of all vertices in \mathcal{P} that extends \mathcal{M}_L and labels every vertex in S with 0.*

⁴For readers more familiar with proof complexity language, our subgraphs correspond to formulas obtained by applying restrictions to $xor(\mathcal{P}_h^d)$.




Fig. 5. Shapes of connected component between layers L and $L + 1$ expanding in dimension j .

The following proposition summarizes the core properties of frustums that we will use when establishing Claim 4.6.

PROPOSITION 4.7. *Let $\mathcal{P} = \mathcal{P}_{L,h}^d[I_{lo}, I_{hi}, \alpha]$ be a restricted frustum with a labelling \mathcal{M}_L of all vertices in the top layer L . Then the following holds:*

- (a) *There is a labelling \mathcal{M}_{L+1} of all vertices in layer $L + 1$ of \mathcal{P} that is consistent with \mathcal{M}_L .*
- (b) *Let \mathcal{M}_{L+1} be any labelling of layer $L + 1$ in \mathcal{P} that is consistent with \mathcal{M}_L and suppose that \mathcal{P} expands coordinate j from layer L to layer $L + 1$, i.e., $j = r_d(L)$ and $j \notin I_{lo} \cup I_{hi}$. Then for any vertex $(\vec{y}, L + 1)$ in \mathcal{P} it holds that the labelling $\mathcal{M}_{L+1}^{\vec{y}}$ defined by*

$$\mathcal{M}_{L+1}^{\vec{y}}(\vec{x}, L + 1) := \begin{cases} 1 - \mathcal{M}_{L+1}(\vec{x}, L + 1) & \text{if } x_i = y_i \text{ for all } i \neq j, \\ \mathcal{M}_{L+1}(\vec{x}, L + 1) & \text{otherwise} \end{cases}$$

is also consistent with \mathcal{M}_L .

PROOF. We first note that the set of XOR constraints between two layers L and $L + 1$ can be partitioned into several connected components. Each of the components forms a “one-dimensional line” in the direction of the expanding coordinate $j = r_d(L)$ (disregarding the directions of the edges; see the illustration in Figure 5). More formally, two vertices from layers L and $L + 1$ are on the same line if and only if they agree on the coordinates x_i for all $i \neq j$. These lines form the connected components of the graph induces on layers L and $L + 1$. All such lines between layers L and $L + 1$ isomorphic and their shape depends on whether $j \in I_{lo}$, $j \in I_{hi}$, or $j \notin I_{lo} \cup I_{hi}$. We remark that in Figures 5(b) and 5(c) the vertex with in-degree 1 and its predecessor form a binary XOR constraint in $xor(\mathcal{P})$.

If the layer is not expanding (as depicted in Figures 5(b) and 5(c)) and the upper layer L is completely labelled, then it is not hard to see that there is a unique consistent labelling of the lower-level vertices of each line (determined by propagating values from right to left in Figure 5(b) and from left to right in Figure 5(c)). As all lines are disjoint, this gives a unique labelling of the entire layer $L + 1$ that is consistent with the labelling of layer L . If the layer expands (i.e., if $j \notin I_{lo} \cup I_{hi}$ as illustrated in Figure 5(a)), then we have more freedom. Indeed, if we label either the rightmost or the leftmost vertex at the bottom layer with 0, then we have the same situation as in Figures 5(b) and 5(c), respectively. This concludes the proof of item 1 in the proposition.

For item (b), first observe that the condition $j \notin I_{lo} \cup I_{hi}$ means that we are in the case depicted in Figure 5(a). This means that if we have a consistent labelling of the upper and lower part of a line, then flipping all values at the lower level yields another consistent labelling. This is so, since every XOR constraint contains exactly two vertices from the lower part. Hence, flipping both of these vertices does not change the parity of the variables in the XOR constraint but leaves the constraint satisfied. As all lines between layers L and $L + 1$ are disconnected from each other, flipping all values in one line gives another consistent labelling for the whole layer $L + 1$, which is precisely what is claimed in item (b). The proposition follows. \square

PROOF OF CLAIM 4.6. The proof is by induction over decreasing values of q , the base case being $q = d$. As $|S| = 2^{d-q} - 1 = 0$ if $q = d$, in this case, we only have to ensure that there is a consistent

labelling of the entire frustum that is consistent with the labelling of the top layer. This follows from inductively applying Proposition 4.7(a) layer-by-layer.

For the inductive step, assume that the claim holds for all top-expanding $(d, q + 1)$ -frustums. We want to prove it for a top-expanding (d, q) -frustum $\mathcal{P} = \mathcal{P}_{L,h}^d[I_{lo}, I_{hi}, \alpha]$. Let $j = r_d(L) = L \bmod d$ be the dimension that expands from layer L to $L + 1$. As \mathcal{P} is top-expanding and $q < d$, we have $j \notin I_{lo} \cup I_{hi}$. For some well-chosen $a \in [0, \text{slen}(j, L + 1)]$ to be specified shortly, we partition \mathcal{P} on and below layer $L + 1$ into the wedge

$$\mathcal{W}(j, a, L + 1) \tag{4.7a}$$

and two disjoint $(d, q + 1)$ -frustums: the “right” frustum

$$\mathcal{P}_{L+1,h}^d[I_{lo} \cup \{j\}, I_{hi}, \alpha \cup \{j \mapsto a\}] \tag{4.7b}$$

and the “left” frustum

$$\mathcal{P}_{L+1,h}^d[I_{lo}, I_{hi} \cup \{j\}, \alpha \cup \{j \mapsto a\}] \tag{4.7c}$$

(the latter is depicted by \times in Figure 4). We choose the position a of the wedge such that both $(d, q + 1)$ -frustums in Equations (4.7b) and (4.7c) contain at most $(|S| - 1)/2 \leq 2^{d-(q+1)} - 1$ vertices from S (which implies that the wedge (4.7a) contains at least one vertex from S). To be more specific, we choose the largest $a \geq 0$ such that the left frustum (4.7c) contains at most $(|S| - 1)/2$ vertices $S_a \subseteq S$. Such an a exists as $S_0 = \emptyset$.

If a reached the maximum $\text{slen}(j, L + 1)$, then empty right frustum (4.7b) clearly contains no vertices from S . Otherwise, let S_{a+1} be the set of vertices from S left of the wedge at position $a + 1$. By the choice of a , we have $|S_{a+1}| > (|S| - 1)/2$. Furthermore, because all vertices in S are below layer $L + 1$, it follows that $S_{a+1} \setminus S_a$ is contained in the wedge (4.7a) at position a . Hence, the right frustum (4.7b) contains at most $|S \setminus S_{a+1}| \leq (|S| - 1)/2$ vertices.

Now, we proceed as follows: First, we use Proposition 4.71 to obtain any consistent labelling of all vertices in layer $L + 1$. Consider the set of all vertices $v = (\vec{x}, L + 1)$ in layer $L + 1$ with $x_j = a$, i.e., the topmost vertices in the wedge (4.7a). Note that this set of vertices form a hyperplane through, and perpendicular to, the disconnected parallel lines discussed in Proposition 4.7. We go over these vertices v one-by-one and flip every 1-labelled v to 0. As j is the expanding coordinate from layer L to $L + 1$, after every such flip, we can apply Proposition 4.7(b) to relabel the rest of the line through v as needed. In this way, all vertices in the top layer of the wedge (4.7a), get labelled by 0, and we label all other vertices in the wedge with 0 also. It follows from repeated application of Proposition 4.7(b) that the end result is a labelling of layer $L + 1$ that is consistent with \mathcal{M}_L . Note that this labels every vertex from S within the wedge with 0, and moreover layer $L + 1$ contains no vertices from S outside of the wedge. This is because if some vertex from S is on layer $L + 1$, then we have by the assumption in Claim 4.6 that $|S| = 2^{d-q} - 1 = 1$ for $d = q + 1$, and hence this one labelled vertex is guaranteed to be in the wedge by the choice of a . In this way, we obtain a labelling for the top layer of both $(d, q + 1)$ -frustums, and we then apply induction to consistently label all vertices in both frustums in such a way that every vertex in S is set to 0. Now, we argue that the disjoint union of the all-zero labelling of the wedge and the consistent labellings of both frustums is a consistent labelling of \mathcal{P} . Clearly, every (non-source) constraint in $\text{xor}(\mathcal{P})$ that is entirely contained in the wedge is satisfied by the all-zero labelling of the wedge. In the same way, every constraint that is entirely contained in one of the two frustums is satisfied by their consistent labellings. It remains to consider constraints that contain variables from the wedge as well as from one of the frustums. By construction, those constraints have the form $(v, w_1, w_2, 0)$, for a vertex v with in-neighbors w_1, w_2 , where the vertex v and one of its neighbors (say, w_1) is within the frustum and the other neighbor is inside the wedge. As the edge (w_1, v) inside the frustum forms the binary constraint $(v, w_1, 0)$, it follows that the consistent labelling of the frustum

guarantees that the parity of v and w_1 is even. Because w_2 is labelled 0 in the wedge, the merged labelling satisfies $(v, w_1, w_2, 0)$. The claim follows. \square

As noted above, our proof of Claim 4.6 also establishes Lemma 4.5.

In Reference [30] log n -dimensional pyramids (where n is the number of vertices) are used to prove a $\Omega(2^{\sqrt{\log n}})$ lower bound on the quantifier depth of full first-order counting logic. The next lemma shows that if we instead choose the dimension to be logarithmic in the number of variables (i.e., pebbles) in the game, then we get an improved quantifier depth lower bound for the k -variable fragment.

LEMMA 4.8. *For every $d \geq 2$ and height h , Player 1 does not win the 2^d -pebble game on $\text{xor}(\mathcal{P}_h^d)$ within $h/(d-1) - 1$ rounds.*

PROOF. We show that Player 2 has a counter-strategy to answer consistently for at least $\lfloor h/(d-1) \rfloor - 1$ rounds, and therefore Player 1 needs at least $\lfloor h/(d-1) \rfloor + 1$ rounds to win. Starting at the top layer $L_1 = 0$ in round $r = 1$, Player 2 maintains the invariant that at the start of round r she has a consistent labelling of all vertices from layer 0 to layer L_r with the property that there is no pebble on layers $L_r + 1$ to $L_r + d - 1$.

Whenever Player 1 places a pebble on or above layer L_r Player 2 responds according to the consistent labelling and whenever Player 1 puts a pebble on or below layer $L_r + d$ she answers 0, and in both cases sets $L_{r+1} = L_r$. Note that as long as Player 1 places pebbles in this way, the game can go on forever. Since there are never more than $2^d - 1$ pebbles left on vertices on or below layer $L_r + d$ (when Player 1 runs out of pebbles the next move must be a removal), the conditions needed for Lemma 4.5 to apply are never violated.

Thus, the interesting case is when Player 1 places a pebble between layer $L_r + 1$ and $L_r + d - 1$. Then Player 2 uses Lemma 4.5 to extend her labelling to the first layer $L_{r+1} > L_r$ such that there is no pebble on layers $L_{r+1} + 1$ to $L_{r+1} + d - 1$, after which she answers the query according to the new labelling. It is worth noting that when Player 2 skips downward from layer L_r to layer L_{r+1} she might jump over a lot of layers in one go, but if so, then there is at least one pebble for every $(d-1)$ th layer forcing such a big jump. We see that following this strategy Player 2 survives for at least $\lfloor h/(d-1) \rfloor - 1$ rounds, and this establishes the lemma. \square

Putting the pieces together, we can now present the lower bound for the k -pebble game in Lemma 3.2.

PROOF OF LEMMA 3.2. Recall that we want to prove that for all $\ell_{\text{hi}} \geq 3$ and $m \geq 3$ there is an m -variable 3-XOR formula F on which Player 1 wins the 3-pebble game but cannot win the ℓ_{hi} -pebble game within $\frac{1}{\lceil \log \ell_{\text{hi}} \rceil} m^{1/(1+\lceil \log \ell_{\text{hi}} \rceil)} - 2$ rounds. If $m < (5 \lceil \log \ell_{\text{hi}} \rceil)^{(\lceil \log \ell_{\text{hi}} \rceil + 1)}$, then the round lower bound is trivial and we let F be for instance, the 3-variable formula $\text{xor}(\mathcal{P}_1^1)$ plus $m - 3$ auxiliary variables on which Player 1 needs 3 rounds to win. Otherwise, we choose the formula to be $F = \text{xor}(\mathcal{P}_h^d)$ for parameters $d = \lceil \log \ell_{\text{hi}} \rceil$ and $h = \lfloor m^{1/(d+1)} \rfloor - 1$. Note that \mathcal{P}_h^d contains less than $(h+1)^{d+1} \leq m$ vertices and we can add dummy variables to reach exactly m . Since the graph \mathcal{P}_h^d has in-degree 2, Lemma 4.2 says that Player 1 wins the 3-pebble game as claimed in Lemma 3.2(a). The lower bound for the ℓ_{hi} -pebble game in Lemma 3.2(b) follows from Lemma 4.8 and the observation that because $h \geq 5d - 1$, we have $h/(d-1) \geq (h+1)/d$ and hence

$$h/(d-1) - 1 \geq (h+1)/d - 1 = \frac{1}{\lceil \log \ell_{\text{hi}} \rceil} \left\lfloor m^{1/(1+\lceil \log \ell_{\text{hi}} \rceil)} \right\rfloor - 1 \quad (4.8a)$$

$$\geq \frac{1}{\lceil \log \ell_{\text{hi}} \rceil} m^{1/(1+\lceil \log \ell_{\text{hi}} \rceil)} - 2. \quad (4.8b)$$

The lemma follows. \square

5 HARDNESS CONDENSATION

In this section, we establish Lemma 3.3, which shows how to convert an XOR formula into an equally hard formula over fewer variables. As discussed in the introduction, this part of our construction relies heavily on Razborov's paper [43]. We follow his line of reasoning closely below, but translate it from proof complexity to a pebble game argument for bounded variable logics.

A key technical concept in the proof is graph expansion. Let us define the particular type of expander graphs that we need and then discuss some crucial properties of these graphs. We use standard notation, letting $\mathcal{G} = (U \dot{\cup} V, E)$ denote a bipartite graph with left vertex set U and right vertex set V . We let $N^{\mathcal{G}}(U') = \{v \mid \{u, v\} \in E(\mathcal{G}), u \in U'\}$ denote the set of neighbor vertices on the right of a left vertex subset $U' \subseteq U$ (and vice versa for right vertex subsets).

Definition 5.1 (Boundary Expander). A bipartite graph $\mathcal{G} = (U \dot{\cup} V, E)$ is an $m \times n$ (s, c) -boundary expander graph if $|U| = m$, $|V| = n$, and for every set $U' \subseteq U$, $|U'| \leq s$, it holds that $|\partial^{\mathcal{G}}(U')| \geq c|U'|$, where the *boundary* $\partial^{\mathcal{G}}(U')$ is the set of all $v \in N^{\mathcal{G}}(U')$ having a unique neighbor in U' , meaning that $|N^{\mathcal{G}}(v) \cap U'| = 1$. An (s, Δ, c) -boundary expander is an (s, c) -boundary expander where additionally $|N^{\mathcal{G}}(u)| \leq \Delta$ for all $u \in U$, i.e., the graph has left degree bounded by Δ .

In what follows, we will omit \mathcal{G} from the notation when the graph is clear from context.

In any (s, c) -boundary expander with expansion $c > 0$, it holds that any left vertex subset $U' \subseteq U$ of size $|U'| \leq s$ has a partial matching into V where in addition the vertices in U' can be ordered in such a way that every vertex $u_i \in U'$ is matched to a vertex outside of the neighborhood of the preceding vertices u_1, \dots, u_{i-1} . The proof of this fact is sometimes referred to as a *peeling argument*.

LEMMA 5.2 (PEELING LEMMA). *Let $\mathcal{G} = (U \dot{\cup} V, E)$ be an (s, c) -boundary expander with $s \geq 1$ and $c > 0$. Then for every set $U' \subseteq U$, $|U'| = t \leq s$ there is an ordering u_1, \dots, u_t of its vertices and a sequence of vertices $v_1, \dots, v_t \in V$ such that $v_i \in N(u_i) \setminus N(\{u_1, \dots, u_{i-1}\})$.*

PROOF. The proof is by induction on t . The base case $t = 1$ is immediate, since $s \geq 1$ and $c > 0$ implies that no left vertex can be isolated. For the inductive step, suppose the lemma holds for $t - 1$. To construct the sequence v_1, \dots, v_t , we first fix v_t to be any vertex in $\partial(U')$, which has to exist, since $|\partial(U')| \geq c|U'| > 0$. The fact that v_t is in the boundary of U' means that there is a unique $u_t \in U'$ such that $|N(v_t) \cap U'| = \{u_t\}$. Thus, for this pair (u_t, v_t) , it holds that $v_t \in N(u_t) \setminus N(U' \setminus \{u_t\})$. By the induction hypothesis, we can now find sequences u_1, \dots, u_{t-1} and v_1, \dots, v_{t-1} for $U' \setminus \{u_t\}$ such that $v_i \in N(u_i) \setminus N(\{u_1, \dots, u_{i-1}\})$, to which we can append u_t and v_t at the end. The lemma follows. \square

For a right vertex subset $V' \subseteq V$ in $\mathcal{G} = (U \dot{\cup} V, E)$, we define the *kernel* $\text{Ker}(V') \subseteq U$ to be the set of all left vertices whose entire neighborhood is contained in V' , i.e.,

$$\text{Ker}(V') = \{u \in U \mid N(u) \subseteq V'\}. \quad (5.1)$$

We let $\mathcal{G} \setminus V'$ denote the subgraph of \mathcal{G} induced on $(U \setminus \text{Ker}(V')) \dot{\cup} (V \setminus V')$. In other words, we obtain $\mathcal{G} \setminus V'$ from \mathcal{G} by first deleting V' and afterwards all isolated vertices from U .

The next lemma states that if \mathcal{G} is an expander graph, then for any small enough right vertex set V' , we can always find a *closure* $\gamma(V') \supseteq V'$ with a small kernel such that the subgraph $\mathcal{G} \setminus \gamma(V')$ has good boundary expansion. The proof of this lemma (albeit with slightly different parameters) can be found in Reference [43], but we also include it in Appendix A for completeness.

LEMMA 5.3 ([43]). *Let \mathcal{G} be an $(s, 2)$ -boundary expander. Then for every $V' \subseteq V$ with $|V'| \leq s/2$ there exists a subset $\gamma(V') \subseteq V$ with $\gamma(V') \supseteq V'$ such that $|\text{Ker}(\gamma(V'))| \leq |V'|$ and the induced subgraph $\mathcal{G} \setminus \gamma(V')$ is an $(s/2, 1)$ -boundary expander.*

Note that Lemma 5.3 does not provide us with information about how the closures of different sets are related. In particular, if we want to choose closures of minimal size, then $V_1 \subseteq V_2$ does not necessarily imply $\gamma(V_1) \subseteq \gamma(V_2)$. For Lemmas 5.2 and 5.3 to be useful, we need to know that there exist good enough boundary expanders. To prove this, one can just fix a left vertex set U of size m and a right vertex set V of size n and then for every $u \in U$ choose Δ neighbors from V uniformly and independently at random. A standard probabilistic argument shows that with high probability this random graph is an $m \times n$ $(s, \Delta, 2)$ -boundary expander for appropriately chosen parameters. We state this formally as a lemma below. A similar lemma is proven in Reference [43], but we also provide a proof in Appendix A for the convenience of the reader.

LEMMA 5.4. *There is an absolute constant $\Delta_0 \in \mathbb{N}^+$ such that for all integers Δ, s , and m satisfying $\Delta \geq \Delta_0$ and $(s\Delta)^{2\Delta} \leq m$ there exist $m \times \lceil m^{3/\Delta} \rceil$ $(s, \Delta, 2)$ -boundary expanders.*

For readers familiar with expander graphs from other contexts, it might be worth pointing out that the parameters above are different from what tends to be the standard expander graph settings of $s = \Omega(m)$ and $\Delta = O(1)$. Instead, in Lemma 5.4, we have s growing sublinearly in m and Δ need not be constant (although we still need $\Delta \gtrsim \log m / \log \log m$ to satisfy the conditions of the lemma).

In what follows, unless otherwise stated $\mathcal{G} = (U \dot{\cup} V, E)$ will be an $(s, 2)$ -boundary expander for $s = 2k$. We will use such expanders when we do XOR substitution in our formulas, as described formally in the next definition. In words, variables in the XOR formula are identified with left vertices U in \mathcal{G} , the pool of new variables is the right vertex set V , and every variable $u \in U$ in an XOR constraint is replaced by an exclusive or $\bigoplus_{v \in N(u)} v$ over its neighbors $v \in N(u)$. We emphasize that in “standard” XORification as found in the proof complexity literature all new substituted variables would be distinct, i.e., $N(u_1) \cap N(u_2) = \emptyset$ for $u_1 \neq u_2$. While this often makes formulas harder, it also increases the number of variables. Here, we use the approach in Reference [43] to instead recycle variables from a much smaller set V in the substitutions, thus decreasing the total number of variables.

Definition 5.5 (XOR Substitution with Recycling). Let F be an XOR formula with $\text{Vars}(F) = U$ and let $\mathcal{G} = (U \dot{\cup} V, E)$ be a bipartite graph. For every constraint $C = (u_1, \dots, u_t, a)$ in F , we let $C[\mathcal{G}]$ be the constraint $(v_1^1, \dots, v_1^{z_1}, \dots, v_t^1, \dots, v_t^{z_t}, a)$, where $N(u_i) = \{v_i^1, \dots, v_i^{z_i}\}$ for all $1 \leq i \leq t$. Taking unions, we let $F[\mathcal{G}]$ be the XOR formula $F[\mathcal{G}] = \{C[\mathcal{G}] \mid C \in F\}$.

When using an $m \times m^{3/\Delta}$ $(s, \Delta, 2)$ -boundary expander as in Lemma 5.4 for substitution in an m -variable XOR formula F as described in Definition 5.5, we obtain a new XOR formula $F[\mathcal{G}]$ where the number of variables has decreased significantly to $m^{3/\Delta}$. The next lemma, which is at the heart of our logic-flavored version of hardness condensation, states that a round lower bound for the k -pebble game on F implies a round lower bound for the k -pebble game on $F[\mathcal{G}]$.

LEMMA 5.6. *Let k be a positive integer and let \mathcal{G} be an $m \times n$ $(2k, 2)$ -boundary expander. Then if F is an XOR formula over m variables such that Player 2 wins the r -round k -pebble game on F , she also wins the $r/(2k)$ -round k -pebble game on $F[\mathcal{G}]$.*

By way of comparison with Reference [43], we remark that a straightforward translation of Razborov’s technique would start with formulas on which Player 1 can win with few pebbles, but needs an almost linear number of rounds to win the game, even if he has an infinite amount of pebbles.⁵ Applying this without modification to Immerman’s construction, we would obtain very weak bounds (and, in particular, nothing interesting for constant k). Instead, as input to our

⁵In terms of resolution, this corresponds to formulas that are refutable in small width, but where every resolution refutation has almost linear depth.

hardness condensation lemma, we use a construction that has a round lower bound of $n^{1/\log k}$ and show that for hardness condensation it is not necessary that the original formula is hard if the number of pebbles get too large.

Before embarking on a formal proof of Lemma 5.6, which is rather technical and will take the rest of this section, let us discuss the intuition behind it. The main idea to obtain a good strategy for Player 2 on the substituted formula $F[\mathcal{G}]$ is to think of the game as being played on F and simulate the survival strategy there for as long as possible (which is where boundary expansion comes into play).

Let $\mathcal{G} = (U \cup V, E)$ be an $(2k, 2)$ -boundary expander as stated in the lemma. We have $\text{Vars}(F) = U$ and $\text{Vars}(F[\mathcal{G}]) = V$. Given a strategy for Player 2 in the r -round k -pebble game on F , we want to convert this into a winning strategy for Player 2 for the $r/(2k)$ -round k -pebble game on $F[\mathcal{G}]$. A first approach (which will not quite work) is the following:

While playing on the substituted formula $F[\mathcal{G}]$, Player 2 simulates the game on F . For every position β in the game on $F[\mathcal{G}]$, she maintains a corresponding position α on F , which is defined on all variables whose entire neighborhood in the expander is contained in the domain of β , i.e., $\text{Vars}(\alpha) = \text{Ker}(\text{Vars}(\beta))$. The assignments of α should be defined in such a way that they are *consistent with* β , i.e., such that $\alpha(u) = \bigoplus_{v \in N(u)} \beta(v)$. It then follows from the description of XORification in Definition 5.5 that α falsifies an XOR constraint of F if and only if β falsifies an XOR constraint of $F[\mathcal{G}]$.

Now, Player 2 wants to play in such a way that if β changes to β' in one round of the game on $F[\mathcal{G}]$, then the corresponding position α also changes to α' in one round of the game on F . Intuitively, this should be done as follows: Suppose that starting from a position β , Player 1 asks for a variable $v \in V$. If v is not the last unassigned vertex in a neighborhood of some $u \in U$, i.e., $\text{Ker}(\text{Vars}(\beta)) = \text{Ker}(\text{Vars}(\beta) \cup \{v\})$, then Player 2 can make an arbitrary choice, as $\alpha = \alpha'$ is consistent with both choices. If v is the last free vertex in the neighborhood of exactly one vertex u , i.e., $\{u\} = \text{Ker}(\text{Vars}(\beta) \cup \{v\}) \setminus \text{Ker}(\text{Vars}(\beta))$, then Player 2 assumes that she was asked for u in the simulated game on F . If in her strategy for the r -round k -pebble game on F she would answer with an assignment $a \in \{0, 1\}$ that would yield the new position $\alpha' = \alpha \cup \{u \mapsto a\}$, then in the game on $F[\mathcal{G}]$ she now sets v to the right value $b \in \{0, 1\}$ such that the new position $\beta' = \beta \cup \{v \mapsto b\}$ satisfies the consistency property $\alpha'(u) = \bigoplus_{v \in N(u)} \beta'(v)$. If Player 2 could follow this strategy, then the number of rounds she would survive the game on $F[\mathcal{G}]$ would be lower-bounded by the number of rounds she survives in the game on F .

There is a gap in this intuitive argument, however, namely, how to handle the case when the queried variable v completes the neighborhood of two (or more) vertices u_1, u_2 at the same time. If it holds that $\{u_1, u_2\} \subseteq \text{Ker}(\text{Vars}(\beta) \cup \{v\}) \setminus \text{Ker}(\text{Vars}(\beta))$, then we have serious problems. Following the strategy above for u_1 and u_2 separately can yield two different and conflicting ways of assigning v , meaning that for the new position β' there will be no consistent assignment α' of $\text{Ker}(\text{Vars}(\beta'))$.

To circumvent this problem and implement the proof idea above, we will use the boundary expansion of \mathcal{G} to ensure that this problematic case does not occur. For instance, suppose that the graph $\mathcal{G}' = \mathcal{G} \setminus \text{Vars}(\beta)$, which is the induced subgraph of \mathcal{G} on $U \setminus \text{Vars}(\alpha)$ and $V \setminus \text{Vars}(\beta)$, has boundary expansion at least 1. Then the bad situation described above with two variables u_1, u_2 having neighborhood $N^{\mathcal{G}'}(u_1) = N^{\mathcal{G}'}(u_2) = \{v\}$ in \mathcal{G}' cannot arise, since this would imply $\partial^{\mathcal{G}'}(\{u_1, u_2\}) = \emptyset$, contradicting the expansion properties of \mathcal{G}' . Unfortunately, we cannot ensure boundary expansion of $\mathcal{G} \setminus \text{Vars}(\beta)$ for every position β , but we can apply Lemma 5.3 and extend the current position to a larger one that is defined on the closure $\gamma(\text{Vars}(\beta))$ and has the desired expansion property. Since Lemma 5.3 ensures that the domain $\text{Ker}(\gamma(\text{Vars}(\beta)))$ of our assignment α under construction is bounded by $|\alpha| \leq |\beta| \leq k$, such an extension will still be good enough.

We now proceed to present a formal proof. When doing so, it turns out to be convenient for us to prove the contrapositive of the statement discussed above. That is, instead of transforming a strategy for Player 2 in the r -round k -pebble game on F to a strategy for the $r/(2k)$ -round k -pebble game on $F[\mathcal{G}]$ for an $(2k, 2)$ -boundary expander \mathcal{G} , we will show that a winning strategy for Player 1 in the game on the substituted formula $F[\mathcal{G}]$ can be used to obtain a winning strategy for Player 1 in the game on the original formula F .

Suppose that β is any position in the k -pebble game on $F[\mathcal{G}]$, i.e., a partial assignment of variables in V . Since \mathcal{G} is a $(2k, 2)$ -boundary expander and $|\beta| \leq k$, we can apply Lemma 5.3 to obtain a superset $\gamma(\text{Vars}(\beta)) \supseteq \text{Vars}(\beta)$ having the properties that $|\text{Ker}(\gamma(\text{Vars}(\beta)))| \leq |\text{Vars}(\beta)|$ and the induced subgraph $\mathcal{G} \setminus \gamma(\text{Vars}(\beta))$ is a $(k, 1)$ -boundary expander. For the rest of this section, fix a minimal such set $\gamma(V')$ for every $V' = \text{Vars}(\beta)$ corresponding to a position β in the k -pebble game (where we note for later use that we have $\gamma(\emptyset) = \emptyset$ by minimality). This will allow us to define formally what we mean by *consistent* positions in the two games on F and $F[\mathcal{G}]$, as described next.

Definition 5.7 (Consistent Positions). Let α be a position in the pebble game on F , i.e., a partial assignment of variables in U , and let β be a partial assignment of variables in V corresponding to a position in the pebble game on $F[\mathcal{G}]$. We say that α is *consistent with* β if there exists an extension $\beta_{\text{ext}} \supseteq \beta$ with $\text{Vars}(\beta_{\text{ext}}) = N(\text{Vars}(\alpha)) \cup \text{Vars}(\beta)$ such that for all $u \in \text{Vars}(\alpha)$ it holds that $\alpha(u) = \bigoplus_{v \in N(u)} \beta_{\text{ext}}(v)$.

Let β be a position in the k -pebble game on the XOR-substituted formula $F[\mathcal{G}]$ and let $\gamma(V')$ be the fixed, minimal closure of β chosen above. Then, we let $\text{Cons}(\beta)$ denote the set of all positions α with $\text{Vars}(\alpha) = \text{Ker}(\gamma(\text{Vars}(\beta)))$ in the pebble game on F that are consistent with β .

Observe that for $\alpha_1 \subseteq \alpha_2$ and $\beta_1 \subseteq \beta_2$, it holds that if α_2 is consistent with β_1 , then so is α_1 , and if α_1 is consistent with β_2 , then α_1 is consistent also with β_1 . Furthermore, by Lemma 5.3, we have $|\alpha| \leq |\beta|$ for all $\alpha \in \text{Cons}(\beta)$. The next claim states the core inductive argument.

CLAIM 5.8. *Let β be a position on $F[\mathcal{G}]$ for an $(2k, 2)$ -boundary expander \mathcal{G} and suppose that Player 1 wins the k -pebble game on $F[\mathcal{G}]$ from position β in i rounds. Then Player 1 has a strategy to win the k -pebble game on F within $2ki$ rounds from every position $\alpha \in \text{Cons}(\beta)$.*

We note that this claim is just a stronger version of (the contrapositive of) Lemma 5.6.

PROOF OF LEMMA 5.6 ASSUMING CLAIM 5.8. Note that if $r/(2k) < 1$, then the lemma is trivially true, as Player 1 always needs at least one round to win the pebble game from the empty position. Otherwise, we apply Claim 5.8 with parameters $\beta = \emptyset$ and $i = r/(2k)$. Since $\text{Cons}(\emptyset) = \{\emptyset\}$, we directly get the contrapositive statement of Lemma 5.6 that if Player 1 wins the $r/(2k)$ -round k -pebble game on $F[\mathcal{G}]$, then he wins the r -round k -pebble game on F . \square

All that remains for us to do now is to establish Claim 5.8, after which the hardness condensation lemma will follow easily.

PROOF OF CLAIM 5.8. The proof is by induction on the number of rounds i . For the base case $i = 0$, we have to show that if β falsifies an XOR constraint in $F[\mathcal{G}]$, then every assignment $\alpha \in \text{Cons}(\beta)$ falsifies an XOR constraint in F . But if β falsifies a constraint of $F[\mathcal{G}]$, which by construction has the form $C[\mathcal{G}]$ for some constraint C from F , then by Definitions 5.5 and 5.7 it holds that every $\alpha \in \text{Cons}(\beta)$ falsifies C .

For the induction step, suppose that the statement holds for $i - 1$ and assume that Player 1 wins the k -pebble game on $F[\mathcal{G}]$ from position β in i rounds. The i th round consists of two steps:

- (1) Player 1 first chooses a subassignment $\beta' \subseteq \beta$.
- (2) He then asks for the value of one variable $v \in V \setminus \text{Vars}(\beta')$, to which Player 2 chooses an assignment $b \in \{0, 1\}$ yielding the new position $\beta' \cup \{v \mapsto b\}$.

As Player 1 has a strategy to win from β within i rounds, it follows that he can win from both $\beta' \cup \{v \mapsto 0\}$ and $\beta' \cup \{v \mapsto 1\}$ within $i - 1$ rounds. By the inductive assumption, we then deduce for the set of assignments

$$\text{Cons}(\beta' * v) := \text{Cons}(\beta' \cup \{v \mapsto 0\}) \cup \text{Cons}(\beta' \cup \{v \mapsto 1\}) \quad (5.2)$$

consistent with either $\beta' \cup \{v \mapsto 0\}$ or $\beta' \cup \{v \mapsto 1\}$ that the following statement holds:

SUBCLAIM 5.9. *Player 1 can win the k -pebble game on F within $2k(i - 1)$ rounds from all positions in $\text{Cons}(\beta' * v)$. \square*

Note that a position is in $\text{Cons}(\beta' * v)$ if it is consistent with either $\beta' \cup \{v \mapsto 0\}$ or $\beta' \cup \{v \mapsto 1\}$. Therefore, $\text{Cons}(\beta' * v)$ is the set of all positions over $\text{Ker}(\gamma(\text{Vars}(\beta') \cup \{v\}))$ that are consistent with β' . What remains to show is that from every position $\alpha \in \text{Cons}(\beta)$ Player 1 can reach⁶ some position in $\text{Cons}(\beta' * v)$ within $2k$ rounds. We split the proof into two steps, corresponding to the two steps in the move of Player 1 from position β .

SUBCLAIM 5.10. *From every position $\alpha \in \text{Cons}(\beta)$, Player 1 can reach some position in $\text{Cons}(\beta')$ for $\beta' \subseteq \beta$ within k rounds.*

SUBCLAIM 5.11. *From every position $\alpha \in \text{Cons}(\beta')$, Player 1 can reach some position in $\text{Cons}(\beta' * v)$ within k rounds.*

Let us establish Subclaims 5.10 and 5.11 in reverse order.

PROOF OF SUBCLAIM 5.11. Player 1 starts with an assignment $\alpha_{\text{start}} \in \text{Cons}(\beta')$, which is defined over the variables $U_{\text{start}} = \text{Ker}(\gamma(\text{Vars}(\beta')))$, and wants to reach some assignment $\alpha_{\text{end}} \in \text{Cons}(\beta' * v)$ defined over the variables $U_{\text{end}} = \text{Ker}(\gamma(\text{Vars}(\beta') \cup \{v\}))$.

If $\text{Ker}(\gamma(\text{Vars}(\beta'))) = \text{Ker}(\gamma(\text{Vars}(\beta') \cup \{v\}))$, then Player 1 can choose $\alpha_{\text{end}} = \alpha_{\text{start}}$. To see this, note that if α_{start} assigns a value to some $u \in N(v)$, then, since $\alpha_{\text{start}} \in \text{Cons}(\beta')$, it holds by Definition 5.7 that $N(u) \subseteq \gamma(\text{Vars}(\beta'))$, and thus α_{start} is already consistent with $\beta' \cup \{v \mapsto b\}$ for some $b \in \{0, 1\}$. Hence, Player 1 need not ask any question in this case, but the induction hypothesis immediately yields the desired conclusion. And if α_{start} does not assign values to any $u \in N(v)$, then it is consistent with $\beta' \cup \{v \mapsto b\}$ for any $b \in \{0, 1\}$, since no additional parity constraints are added in Definition 5.7.

The more interesting case is when $\text{Ker}(\gamma(\text{Vars}(\beta'))) \neq \text{Ker}(\gamma(\text{Vars}(\beta') \cup \{v\}))$. Now, Player 1 first deletes all assignments of variables in $U_{\text{start}} \setminus U_{\text{end}}$ from α_{start} to get α_0 . Since $\alpha_0 \subseteq \alpha_{\text{start}}$ and α_{start} is consistent with β' by assumption, α_0 is also consistent with β' . Afterwards, he asks for all variables in $U' = U_{\text{end}} \setminus U_{\text{start}}$. We need to argue that regardless of how Player 2 answers, it holds that Player 1 reaches a position that is consistent with β' . This is where the peeling argument in Lemma 5.2 is needed.

As discussed above, by our choice of the closure $\gamma(\text{Vars}(\beta'))$ (obtained using Lemma 5.3), we know that the bipartite graph $\mathcal{G}' = \mathcal{G} \setminus \gamma(\text{Vars}(\beta'))$ is a $(k, 1)$ -boundary expander and furthermore that for $U' = U_{\text{end}} \setminus U_{\text{start}}$ it holds that $|U'| \leq |U_{\text{end}}| \leq |\text{Vars}(\beta') \cup \{v\}| \leq k$, as observed right after Definition 5.7. Hence, we can apply Lemma 5.2 to \mathcal{G}' and U' to get an ordered sequence u_1, \dots, u_t

⁶Recall that *Player 1 can reach a position α' from a position α* means that he has a strategy such that in every play of game starting from position α he either wins or ends up with position α' .

satisfying $N^{\mathcal{G}'}(u_i) \setminus N^{\mathcal{G}'}(\{u_1, \dots, u_{i-1}\}) \neq \emptyset$. We will think of Player 1 as querying the (at most k) vertices in U' in this order, after which he ends up with a position α_{end} defined on the variables U_{end} .

To argue that the position α_{end} obtained in this way is consistent with β' independently of how Player 2 answers, and is hence contained in $\text{Cons}(\beta' * v)$, we show inductively that all positions encountered during the transition from α_{start} to α_{end} are consistent with β' . As already noted, this holds for the position α_0 obtained from α_{start} by deleting all assignments of variables in $U_{\text{start}} \setminus U_{\text{end}}$. For the induction step, let $i \geq 0$ and assume inductively that the current position α_i over

$$U_i := (U_{\text{start}} \cap U_{\text{end}}) \cup \{u_j \mid 1 \leq j \leq i\} \quad (5.3)$$

is consistent with β' . Now, Player 1 asks about the variable u_{i+1} and Player 2 answers with a value a_{i+1} . Since α_i is consistent with β' , there is an assignment $\beta_{\text{ext}} \supseteq \beta'$ that sets the variables $v \in N(\text{Vars}(\alpha_i))$ to the right values such that $\alpha_i(u) = \bigoplus_{v \in N(u)} \beta_{\text{ext}}(v)$ for all $u \in \text{Vars}(\alpha_i)$. By our ordering of $U' = \{u_1, \dots, u_t\}$ chosen above, we know that u_{i+1} has at least one neighbor on the right-hand side V that is neither contained in $N^{\mathcal{G}}(U_i) = N^{\mathcal{G}}(\text{Vars}(\alpha_i))$ nor in the domain of β' . Hence, regardless of which value a_{i+1} Player 2 chooses for her answer, we can extend the assignment β_{ext} to the variables $N^{\mathcal{G}}(u_{i+1}) \setminus (N^{\mathcal{G}}(\text{Vars}(\alpha_i)) \cup \text{Vars}(\beta'))$ in such a way that $\bigoplus_{v \in N(u_{i+1})} \beta_{\text{ext}}(v) = a_{i+1}$. This shows that α_{i+1} defined over $U_{i+1} = (U_{\text{start}} \cap U_{\text{end}}) \cup \{u_j \mid 1 \leq j \leq i+1\}$ is consistent with β' . Subclaim 5.11 now follows by the induction principle. \dashv

Before proving Subclaim 5.10, we should perhaps point out why this claim is not vacuous. Recalling the discussion just below Lemma 5.3, this is because the condition $V_1 \subseteq V_2$ does not allow us to conclude that $\gamma(V_1) \subseteq \gamma(V_2)$.

PROOF OF SUBCLAIM 5.10. The proof is similar to that of Subclaim 5.11 above. Player 1 starts with an assignment $\alpha_{\text{start}} \in \text{Cons}(\beta)$ and wants to reach some assignment in $\text{Cons}(\beta')$ for $\beta' \subseteq \beta$ within k rounds. By assumption, α_{start} is consistent with β and therefore (since $\beta' \subseteq \beta$) is also consistent with β' . Player 1 deletes all assignments from the domain $U_{\text{start}} = \text{Ker}(\gamma(\text{Vars}(\beta)))$ of α_{start} that do not occur in the domain $U_{\text{end}} = \text{Ker}(\gamma(\text{Vars}(\beta')))$ of positions in $\text{Cons}(\beta')$, resulting in the position $\alpha_0 \subseteq \alpha_{\text{start}}$ that is consistent with β' . Next, he applies Lemma 5.2 to $\mathcal{G}' = \mathcal{G} \setminus \gamma(\text{Vars}(\beta))$ to obtain an ordering of the remaining variables $U_{\text{end}} \setminus U_{\text{start}}$. In the same way as above, he can query the variables in this order while maintaining the invariant that the current position is consistent with β' . \dashv

Combining Subclaims 5.9, 5.10, and 5.11, we conclude that Player 1 wins from every position $\alpha \in \text{Cons}(\beta)$ within $2ki$ rounds. This concludes the proof of Claim 5.8. \square

We are finally in a position to give a formal proof of Lemma 3.3.

PROOF OF LEMMA 3.3. Let $\Delta_0 \in \mathbb{N}^+$ be the constant in Lemma 5.4. Suppose we are given an m -variable p -XOR formula F and parameters $\ell_{\text{lo}}, \ell_{\text{hi}}, r, \Delta$ satisfying the conditions in Lemma 3.3 that $\ell_{\text{hi}}/\ell_{\text{lo}} \geq \Delta \geq \Delta_0$ and $(2\ell_{\text{hi}}\Delta)^{2\Delta} \leq m$.

Fix $k := \ell_{\text{hi}}$ and $s := 2\ell_{\text{hi}}$. Since $(s\Delta)^{2\Delta} \leq m$ and $\Delta \geq \Delta_0$, we appeal to Lemma 5.4 to obtain an $m \times \lceil m^{3/\Delta} \rceil (s, \Delta, 2)$ -boundary expander $\mathcal{G} = (U \dot{\cup} V, E)$, and applying XORification with respect to \mathcal{G} , we construct the formula $H := F[\mathcal{G}]$. Clearly, H is an (Δp) -XOR formula with $\lceil m^{3/\Delta} \rceil$ variables. We want to prove that Player 1 has a winning strategy for the $(\Delta\ell_{\text{lo}})$ -pebble game on H as guaranteed by Lemma 3.3(a'), but that he does not win the ℓ_{hi} -pebble game on H within $r/(2\ell_{\text{hi}})$ rounds, as stated in Lemma 3.3(b').

For the upper bound in Lemma 3.3(a'), we recall that Player 1 has a winning strategy in the ℓ_{lo} -pebble game on F by assumption (a) in the lemma. He can use this strategy to win the $(\Delta\ell_{\text{lo}})$ -pebble game on H as follows: Whenever his strategy tells him to ask for a variable $u \in U = \text{Vars}(F)$, he instead asks for the at most Δ variables in $N(u) \subseteq V = \text{Vars}(H)$ and assigns to u the value that

corresponds to the parity of the answers Player 2 gives for $N(u)$. In this way, he can simulate his strategy on F until he reaches an assignment that contradicts an XOR constraint C from F . As the corresponding assignment of the variables $\{v \mid v \in N(u), u \in \text{Vars}(C)\}$ falsifies the constraint $C[\mathcal{G}] \in H$, at this point Player 1 wins the $(\Delta\ell_{\text{lo}})$ -pebble game on H .

The lower bound in Lemma 3.3(b') follows immediately from Lemma 5.6. By assumption (b) in Lemma 3.3, Player 1 does not win the ℓ_{hi} -pebble game on F within r rounds. Since \mathcal{G} is an $m \times n$ $(2k, 2)$ -boundary expander, Lemma 5.6 says that he does not win the ℓ_{hi} -pebble game on $H = F[\mathcal{G}]$ within $r/(2\ell_{\text{hi}})$ rounds either. This concludes the proof of Lemma 3.3. \square

6 CONCLUDING REMARKS

In this article, we prove an $n^{\Omega(k/\log k)}$ lower bound on the minimal quantifier depth of L^k and C^k sentences that distinguish two finite n -element relational structures, nearly matching the trivial n^{k-1} upper bound. By the known connection to the k -dimensional Weisfeiler–Leman algorithm, this result implies near-optimal $n^{\Omega(k/\log k)}$ lower bounds also on the number of refinement steps of this algorithm. The key technical ingredient in our proof is the hardness condensation technique recently introduced by Razborov [43] in the context of proof complexity, which we translate into the language of finite variable logics and use to reduce the domain size of relational structures while maintaining the minimal quantifier depth required to distinguish them.

An obvious open problem is to improve our lower bound. One way to achieve this would be to strengthen the lower bound on the number of rounds in the k -pebble game on 3-XOR formulas in Lemma 3.2 from $n^{1/\log k}$ to n^δ for some $\delta \gg 1/\log k$. By the hardness condensation lemma, this would directly improve our lower bound from $n^{\Omega(k/\log k)}$ to $n^{\Omega(\delta k)}$.

The structures on which our lower bounds hold are n -element relational structures of arity $\Theta(k)$ and size $n^{\Theta(k)}$. We would have liked to have this results also for structures of bounded arity, such as graphs. However, the increase of the arity is inherent in the method of amplifying hardness by making XOR substitutions. An optimal lower bound of $n^{\Omega(k)}$ on the quantifier depth required to distinguish two n -vertex graphs has been obtained by the first author in an earlier work [12] for the *existential-positive fragment* of L^k . Determining the quantifier rank of full L^k and C^k on n -vertex graphs remains an open problem.

Another open question related to our results concerns the complexity of finite variable equivalence for non-constant k . What is the complexity of deciding, given two structures and a parameter k , whether the structures are equivalent in L^k or C^k ? As this problem can be solved in time $(\|\mathcal{A}\| + \|\mathcal{B}\|)^{O(k)}$, it is in EXPTIME if k is part of the input. It has been conjectured that this problem is EXPTIME-complete [21], but it is not even known whether it is NP-hard. Note that the quantifier depth is connected to the computational complexity of the equivalence problem by the fact that an upper bound of the form $n^{O(1)}$ on n -element structures would have implied that testing equivalence is in PSPACE. Hence, our lower bounds on the quantifier depth can be seen as a necessary requirement for establishing EXPTIME-hardness of the equivalence problem.

APPENDIX

A EXISTENCE AND PROPERTIES OF EXPANDER GRAPHS

In this appendix, we present proofs of Lemmas 5.3 and 5.4, starting with the latter lemma. We again remark that most of this material can already be found in a similar form in Reference [43], although the exact parameters are somewhat different. It also seems appropriate to point out that there is a significant overlap with essentially identical technical lemmas in Reference [13].

Just to avoid ambiguity, let us state explicitly that even though we have the Euler number e appearing below, all logarithms are being taken to base 2 (though this should not really matter too much).

LEMMA 5.4 (RESTATED). *There is an absolute constant $\Delta_0 \in \mathbb{N}^+$ such that for all integers Δ, s, m satisfying $\Delta \geq \Delta_0$ and $(s\Delta)^{2\Delta} \leq m$ there exist $m \times \lceil m^{3/\Delta} \rceil$ $(s, \Delta, 2)$ -boundary expanders.*

PROOF. Let U and V be two disjoint sets of vertices of size $|U| = m$ and $|V| = n = \lceil m^{3/\Delta} \rceil$. For every $u \in U$, we choose Δ times a neighbor $v \in V$ uniformly at random with repetitions. This yields a bipartite graph $\mathcal{G} = (U \dot{\cup} V, E)$ of left-degree at most Δ . In the sequel, we show that \mathcal{G} is likely to be an $(s, \Delta, 2)$ -boundary expander.

First note that for every set $U' \subseteq U$ all neighbors $v \in N(U') \setminus \partial(U')$ that are not in the boundary of U' have at least two neighbors in U' . Since there are at most $\Delta|U'| - |\partial(U')|$ edges between U' and $N(U') \setminus \partial(U')$, it follows that $|N(U') \setminus \partial(U')| \leq (\Delta|U'| - |\partial(U')|)/2$ and hence

$$|N(U')| \leq \frac{|\partial(U')| + \Delta|U'|}{2}. \quad (\text{A.1})$$

If \mathcal{G} is not an $(s, \Delta, 2)$ -boundary expander, then there is a set U' of size $|U'| = \ell \leq s$ that has a boundary $\partial(U')$ of size $|\partial(U')| < 2\ell$ and from Inequality (A.1) it then follows that $|N(U')| < (1 + \Delta/2)\ell$. By a union bound argument (and relaxing to non-strict inequalities), we obtain

$$\Pr[\mathcal{G} \text{ is not an } (s, \Delta, 2)\text{-boundary expander}] \quad (\text{A.2a})$$

$$\leq \sum_{\ell=1}^s \sum_{U' \subseteq U; |U'|=\ell} \Pr[|\partial(U')| \leq 2\ell] \quad (\text{A.2b})$$

$$\leq \sum_{\ell=1}^s \sum_{U' \subseteq U; |U'|=\ell} \Pr[|N(U')| \leq (1 + \Delta/2)\ell] \quad (\text{A.2c})$$

$$\leq \sum_{\ell=1}^s \binom{m}{\ell} \binom{n}{(1 + \Delta/2)\ell} \left(\frac{(1 + \Delta/2)\ell}{n} \right)^{\Delta\ell} \quad (\text{A.2d})$$

$$\leq \sum_{\ell=1}^s m^\ell \left(\frac{en}{(1 + \Delta/2)\ell} \right)^{(1+\Delta/2)\ell} \left(\frac{(1 + \Delta/2)\ell}{n} \right)^{\Delta\ell} \quad (\text{A.2e})$$

$$= \sum_{\ell=1}^s m^\ell (en)^{(1+\Delta/2)\ell} ((1 + \Delta/2)\ell)^{(\Delta/2-1)\ell} n^{-\Delta\ell} \quad (\text{A.2f})$$

$$\leq \sum_{\ell=1}^s n^{(\Delta/3)\ell} (en)^{(1+\Delta/2)\ell} ((1 + \Delta/2)\ell)^{(\Delta/2-1)\ell} n^{-\Delta\ell} \quad (\text{A.2g})$$

$$= \sum_{\ell=1}^s n^{(\Delta/3)\ell} n^{\frac{\log e}{\log n}(1+\Delta/2)\ell} n^{\frac{1}{\log n} \log((\Delta/2+1)\ell)^{(\Delta/2-1)\ell}} n^{(-\Delta/2+1)\ell} \quad (\text{A.2h})$$

$$\leq \sum_{\ell=1}^s n^{\left(\frac{\log e}{\log n} \Delta + \frac{1}{\log n} \log(\Delta s)^{(\Delta/2-1)-\Delta/6+1} \right) \ell}, \quad (\text{A.2i})$$

where in going from Term (A.2d) to Term (A.2e), we use the inequality $\binom{n}{k} \leq \left(\frac{en}{k} \right)^k$ for $e \approx 2.718$ denoting the Euler number, and in going from Term (A.2h) to Term (A.2i), we assume that $\Delta \geq 2$ and also use that $\ell \leq s$.

To show that the expression (A.2i) is bounded away from 1—which implies that \mathcal{G} is an $(s, \Delta, 2)$ -boundary expander with constant probability—it suffices to study the exponent and prove that there is a constant $\varepsilon > 0$ such that

$$\frac{\log e}{\log n} \Delta + \frac{1}{\log n} \log(\Delta s) \left(\frac{\Delta}{2} - 1 \right) - \frac{\Delta}{6} + 1 \leq -\varepsilon < 0, \quad (\text{A.3})$$

which holds if there is an $\varepsilon' = \varepsilon/\Delta$ such that

$$\frac{\log e}{\log n} + \frac{1}{\log n} \log(\Delta s) \left(\frac{1}{2} - \frac{1}{\Delta} \right) - \frac{1}{6} + \frac{1}{\Delta} \leq -\varepsilon' < 0. \quad (\text{A.4})$$

Since $(s\Delta)^{2\Delta} \leq m \leq n^{\Delta/3}$, we have $s\Delta \leq n^{1/6}$, and it follows that

$$\frac{\log e}{\log n} + \frac{\log(\Delta s)(1/2 - 1/\Delta)}{\log n} - \frac{1}{6} + \frac{1}{\Delta} \quad (\text{A.5a})$$

$$\leq \frac{\log e}{\log n} + \frac{\log(n^{1/6})}{2 \log n} - \frac{1}{6} + \frac{1}{\Delta} \quad (\text{A.5b})$$

$$= \frac{\log e}{\log n} - \frac{1}{12} + \frac{1}{\Delta} \quad (\text{A.5c})$$

$$\leq -\varepsilon' < 0, \quad (\text{A.5d})$$

where we can make the last inequality hold for ε' small enough and n and Δ large enough. Formally, assuming (*) $\Delta \geq 13$, we fix constants $0 < \varepsilon' < 1/12$ and n_0 such that the inequality between Term (A.5c) and Term (A.5d) holds for any n satisfying (**) $n \geq n_0$. Then, we obtain that Term (A.2i) is bounded by $\sum_{\ell=1}^s n^{-\varepsilon'\Delta\ell}$. Insisting in addition that (***) $n \geq 3^{1/\varepsilon'}$, we can upper-bound Term (A.2i) by

$$\sum_{\ell=1}^s n^{-\varepsilon'\Delta\ell} \leq \sum_{\ell=1}^{\infty} \left(\frac{1}{3}\right)^\ell \leq \frac{1}{2}. \quad (\text{A.6})$$

It remains to calculate how to set Δ_0 to make sure that conditions (*), (**), and (***) hold. Note that, by assumption, we have $(s\Delta)^{2\Delta} \leq m$, which implies that $\Delta^\Delta \leq m$. It follows that we will always have $n = \lceil m^{3/\Delta} \rceil \geq (\Delta^\Delta)^{3/\Delta} = \Delta^3 \geq (\Delta_0)^3$. Hence, to guarantee $\Delta \geq 13$, $n \geq n_0$, and $n \geq 3^{1/\varepsilon'}$, it is sufficient to choose $\Delta_0 \geq \max(13, n_0^{1/3}, 3^{1/(3\varepsilon')})$. This concludes the proof of the lemma. \square

We next prove that in a good enough boundary expander, it holds that for any small enough right vertex set V' there is a superset $\gamma(V') \supseteq V'$ with a small kernel such that the induced subgraph $\mathcal{G} \setminus \gamma(V')$ (obtained from \mathcal{G} by deleting V' and then all isolated vertices from U) is also a good boundary expander. Recall that we refer to such a set $\gamma(V')$ as a *closure* of V' .

LEMMA 5.3 (RESTATED). *Let \mathcal{G} be an $(s, 2)$ -boundary expander. Then for every $V' \subseteq V$ with $|V'| \leq s/2$ there exists a subset $\gamma(V') \subseteq V$ with $\gamma(V') \supseteq V'$ such that $|\text{Ker}(\gamma(V'))| \leq |V'|$ and the induced subgraph $\mathcal{G} \setminus \gamma(V')$ is an $(s/2, 1)$ -boundary expander.*

PROOF. Let $\mathcal{G} = (U \dot{\cup} V, E)$ be an $(s, 2)$ -boundary expander and let $V' \subseteq V$ have size $|V'| \leq s/2$. We construct an increasing sequence $V' = V_0 \subset V_1 \subset \dots \subset V_r = \gamma(V')$ such that $\mathcal{G} \setminus V_r$ is an $(s/2, 1)$ -boundary expander as follows:

If $\mathcal{G} \setminus V_0$ is not an $(s/2, 1)$ -boundary expander, then there exists a set U_1 of size $|U_1| \leq s/2$ such that $|\partial^{\mathcal{G} \setminus V_0}(U_1)| \leq |U_1|$. Delete $N(U_1)$ and $\text{Ker}(N(U_1))$ from $\mathcal{G} \setminus V_0$. If the resulting graph is not an $(s/2, 1)$ -boundary expander, then we repeat this process and iteratively delete vertex sets that do not satisfy the expansion condition. Formally, for $i \geq 1$ fix U_i to be any set of size $|U_i| \leq s/2$ such that

$$|\partial^{\mathcal{G} \setminus V_{i-1}}(U_i)| \leq |U_i|, \quad (\text{A.7})$$

where we set

$$V_i := V_0 \cup \bigcup_{j=1}^i N^{\mathcal{G}}(U_j) \quad (\text{A.8})$$

(and where we note that at the i th step, we delete what remains of $N^{\mathcal{G}}(U_i)$ and the kernel $\text{Ker}(N^{\mathcal{G}}(U_i))$ of this right vertex set). Since all sets U_i constructed above are non-empty, this process must terminate for some $i = \tau$ and the resulting graph $\mathcal{G} \setminus V_\tau$ is then an $(s/2, 1)$ -boundary expander (note that an empty graph without vertices vacuously satisfies the expansion condition). It remains to verify that the size condition $|\text{Ker}(V_\tau)| \leq |V_0|$ for the kernel of the closure of V' holds. This is immediately implied by the following inductive claim:

CLAIM A.1. *Let $V_{-1} = U_0 = \emptyset$ and suppose that $i \geq 0$. Then for U_i satisfying Inequality (A.7) and V_i defined by Inequality (A.8) the following properties hold:*

- (1) *For all U' such that $\text{Ker}(V_{i-1}) \cup U_i \subseteq U' \subseteq \text{Ker}(V_i)$, we have $|\partial^{\mathcal{G}}(U') \setminus V_0| \leq |\text{Ker}(V_i)|$.*
- (2) *The kernel of V_i has size $|\text{Ker}(V_i)| \leq |V_0|$.*

For $i = 0$, Property 1 in Claim A.1 follows, because $U' \subseteq \text{Ker}(V_0)$ implies that $\partial^{\mathcal{G}}(U') \subseteq V_0$. For Property 2, suppose that $|\text{Ker}(V_0)| \leq s$. Then expansion implies $2|\text{Ker}(V_0)| \leq |\partial^{\mathcal{G}}(\text{Ker}(V_0))|$, and combining this with $\partial^{\mathcal{G}}(\text{Ker}(V_0)) \subseteq V_0$, we obtain $|\text{Ker}(V_0)| \leq \frac{1}{2}|V_0|$. If instead $|\text{Ker}(V_0)| > s$, then we can find a subset $U' \subseteq \text{Ker}(V_0)$ of size $|U'| = s$. By expansion, we have $|\partial^{\mathcal{G}}(U')| \geq 2s$, which is a contradiction, because, as argued above, we should have $|\partial^{\mathcal{G}}(U')| \leq |V_0| \leq s/2$.

For the induction step, suppose that both properties hold for $i - 1$. Let $U^* = \text{Ker}(V_{i-1}) \cup U_i$ and consider any U' satisfying $U^* \subseteq U' \subseteq \text{Ker}(V_i)$. We claim that every boundary element in $\partial^{\mathcal{G}}(U')$ is either a boundary element from $\partial^{\mathcal{G}}(U^*)$ or is contained in V_0 . To see this, note that, since $U' \subseteq \text{Ker}(V_i)$, we have $\partial^{\mathcal{G}}(U') \subseteq V_i = V_0 \cup \bigcup_{j=1}^i N^{\mathcal{G}}(U_j)$. Furthermore, it can be observed that $\bigcup_{j=1}^i U_j \subseteq U^* \subseteq U'$ (this is basically due to the fact that $N(\text{Ker}(V')) \subseteq V'$ for any V'). Hence, if $v \in \partial^{\mathcal{G}}(U') \setminus V_0$, then it must hold that $v \in \bigcup_{j=1}^i N^{\mathcal{G}}(U_j)$, and so the unique neighbor of v on the left is contained in $\bigcup_{j=1}^i U_j$ and therefore also in U^* , implying that $v \in \partial(U^*)$. This yields that

$$\partial^{\mathcal{G}}(U') \setminus V_0 \subseteq \partial^{\mathcal{G}}(U^*) \setminus V_0, \quad (\text{A.9})$$

as claimed, and in what follows, we will show

$$\left| \partial^{\mathcal{G}}(U^*) \setminus V_0 \right| = \left| \partial^{\mathcal{G}}(\text{Ker}(V_{i-1}) \cup U_i) \setminus V_0 \right| \leq |\text{Ker}(V_i)| \quad (\text{A.10})$$

to prove Property 1.

Note that by construction every vertex in $V_{i-1} \setminus V_0$ has at least one neighbor in $\text{Ker}(V_{i-1})$. It follows that all new boundary vertices in $\partial^{\mathcal{G}}(\text{Ker}(V_{i-1}) \cup U_i) \setminus \partial^{\mathcal{G}}(\text{Ker}(V_{i-1}))$ are either from V_0 or from the boundary $\partial^{\mathcal{G} \setminus V_{i-1}}(U_i)$ of U_i outside of V_{i-1} . Therefore, we have

$$\partial^{\mathcal{G}}(U^*) \setminus V_0 = \partial^{\mathcal{G}}(\text{Ker}(V_{i-1}) \cup U_i) \setminus V_0 \subseteq \left(\partial^{\mathcal{G}}(\text{Ker}(V_{i-1})) \setminus V_0 \right) \dot{\cup} \partial^{\mathcal{G} \setminus V_{i-1}}(U_i). \quad (\text{A.11})$$

Since, by assumption, U_i does not satisfy the expansion condition, we know that

$$\left| \partial^{\mathcal{G} \setminus V_{i-1}}(U_i) \right| \leq |U_i| \quad (\text{A.12})$$

and by the inductive hypothesis concerning Property 1 with $U' = \text{Ker}(V_{i-1})$, we have

$$\left| \partial^{\mathcal{G}}(\text{Ker}(V_{i-1})) \setminus V_0 \right| \leq |\text{Ker}(V_{i-1})|. \quad (\text{A.13})$$

Combining Expression (A.9) with Expression (A.11) and (A.13), we deduce that

$$\left| \partial^{\mathcal{G}}(U') \setminus V_0 \right| \leq \left| \partial^{\mathcal{G}}(\text{Ker}(V_{i-1}) \cup U_i) \setminus V_0 \right| \leq |\text{Ker}(V_{i-1})| + |U_i| \leq |\text{Ker}(V_i)|, \quad (\text{A.14})$$

where the final inequality holds, since $\text{Ker}(V_{i-1})$ and U_i are disjoint subsets of $\text{Ker}(V_i)$. This concludes the inductive step for Property 1.

To establish Property 2, assume first that $|\text{Ker}(V_i)| \leq s$. Then by expansion and Property 1 applied to $U' = \text{Ker}(V_i)$, we have

$$2|\text{Ker}(V_i)| \leq \left| \partial^{\mathcal{G}}(\text{Ker}(V_i)) \right| \leq |V_0| + |\text{Ker}(V_i)| \quad (\text{A.15})$$

and hence

$$|\text{Ker}(V_i)| \leq |V_0|, \quad (\text{A.16})$$

as desired. If instead $|\text{Ker}(V_i)| > s$, then by the inductive hypothesis, we know that $|\text{Ker}(V_{i-1})| \leq s/2$ and by construction, we have $|U_i| \leq s/2$. Therefore, we can find a set U' of size $|U'| = s$ satisfying the condition $\text{Ker}(V_{i-1}) \cup U_i \subseteq U' \subseteq \text{Ker}(V_i)$ in Property 1. From the expansion properties of \mathcal{G} , we conclude that $|\partial(U')| \geq 2s$. But this is a contradiction, because for sets U' satisfying the conditions in Property 1, we derived Inequality (A.14), which implies that $|\partial(U')| \leq |V_0| + |\text{Ker}(V_{i-1})| + |U_i| \leq 2 \cdot |V_0| + |U_i| \leq 3s/2$. \square

ACKNOWLEDGMENTS

We are very grateful to Alexander Razborov for patiently explaining the hardness condensation technique in Reference [43] during numerous and detailed discussions. We also want to thank Neil Immerman for helping us find relevant references to previous works and explaining some of the finer details in Reference [30]. The first author wishes to acknowledge useful feedback from the participants of Dagstuhl Seminar 15511 *The Graph Isomorphism Problem*. Finally, we are most indebted to the anonymous *LICS* and *JACM* reviewers for very detailed comments that helped improve the manuscript considerably.

REFERENCES

- [1] Michael Alekhnovich and Alexander A. Razborov. 2003. Lower bounds for polynomial calculus: Non-binomial case. *Proc. Steklov Inst. Math.* 242 (2003), 18–35. Retrieved from <http://people.cs.uchicago.edu/razborov/files/misha.pdf>
- [2] Albert Atserias and Víctor Dalmau. 2008. A combinatorial characterization of resolution width. *J. Comput. Syst. Sci.* 74, 3 (May 2008), 323–334.
- [3] László Babai. 2016. Graph isomorphism in quasipolynomial time. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC'16)*. 684–697.
- [4] Jon Barwise. 1977. On Moschovakis closure ordinals. *J. Symbol. Logic* 42, 2 (June 1977), 292–296.
- [5] Chris Beck, Jakob Nordström, and Bangsheng Tang. 2013. Some trade-off results for polynomial calculus. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC'13)*. 813–822.
- [6] Eli Ben-Sasson. 2002. Hard examples for the bounded depth Frege proof system. *Comput. Complex.* 11, 3-4 (2002), 109–136.
- [7] Eli Ben-Sasson. 2002. Size space tradeoffs for resolution. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC'02)*. 457–464.
- [8] Eli Ben-Sasson. 2009. Size-space tradeoffs for resolution. *SIAM J. Comput.* 38, 6 (May 2009), 2511–2525.
- [9] Eli Ben-Sasson and Jakob Nordström. 2008. Short proofs may be spacious: An optimal separation of space and length in resolution. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS'08)*. 709–718.
- [10] Eli Ben-Sasson and Jakob Nordström. 2011. Understanding space in proof complexity: Separations and trade-offs via substitutions. In *Proceedings of the 2nd Symposium on Innovations in Computer Science (ICS'11)*. 401–416.
- [11] Eli Ben-Sasson and Avi Wigderson. 2001. Short proofs are narrow—resolution made simple. *J. ACM* 48, 2 (Mar. 2001), 149–169.
- [12] Christoph Berkholz. 2014. The propagation depth of local consistency. In *Proceedings of the 20th International Conference on Principles and Practice of Constraint Programming (CP'14)* (Lecture Notes in Computer Science, Vol. 8656). Springer, 158–173.
- [13] Christoph Berkholz and Jakob Nordström. 2020. Supercritical space-width trade-offs for resolution. *SIAM J. Comput.* 49, 1 (Feb. 2020), 98–118.
- [14] Joshua Buresh-Oppenheim and Toniann Pitassi. 2007. The complexity of resolution refinements. *J. Symbol. Logic* 72, 4 (Dec. 2007), 1336–1352.
- [15] Samuel R. Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. 2001. Linear gaps between degrees for the polynomial calculus modulo distinct primes. *J. Comput. Syst. Sci.* 62, 2 (Mar. 2001), 267–289.
- [16] Samuel R. Buss and Jakob Nordström. 2021. Proof complexity and SAT solving. In *Handbook of Satisfiability (2nd ed.)* (Frontiers in Artificial Intelligence and Applications, Vol. 336), Armin Biere, Marijn J. H. Heule, Hans van Maaren, and Toby Walsh (Eds.). IOS Press 233–350.

- [17] Jin-yi Cai, Martin Fürer, and Neil Immerman. 1992. An optimal lower bound on the number of variables for graph identifications. *Combinatorica* 12, 4 (1992), 389–410.
- [18] Stephen A. Cook. 1974. An observation on time-storage trade off. *J. Comput. Syst. Sci.* 9, 3 (1974), 308–316.
- [19] Yuval Filmus, Massimo Lauria, Mladen Mikša, Jakob Nordström, and Marc Vinyals. 2013. Towards an understanding of polynomial calculus: New separations and lower bounds (extended abstract). In *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP'13)* (Lecture Notes in Computer Science, Vol. 7965). Springer, 437–448.
- [20] Martin Fürer. 2001. Weisfeiler-Leman refinement requires at least a linear number of iterations. In *Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP'01)* (Lecture Notes in Computer Science, Vol. 2076). Springer, 322–333.
- [21] Erich Grädel, Phokion G. Kolaitis, Leonid Libkin, Maarten Marx, Joel Spencer, Moshe Y. Vardi, Yde Venema, and Scott Weinstein. 2007. *Finite Model Theory and Its Applications*. Springer.
- [22] Dima Grigoriev. 2001. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theor. Comput. Sci.* 259, 1–2 (May 2001), 613–622.
- [23] Martin Grohe. 1998. Finite variable logics in descriptive complexity theory. *Bull. Symbol. Logic* 4, 4 (1998), 345–398.
- [24] Martin Grohe. 1999. Equivalence in finite-variable logics is complete for polynomial time. *Combinatorica* 19, 4 (Oct. 1999), 507–532.
- [25] Martin Grohe. 2012. Fixed-point definability and polynomial time on graphs with excluded minors. *J. ACM* 59, 5, Article 27 (Oct. 2012), 27:1–27:64 pages.
- [26] Martin Grohe, Moritz Lichter, and Daniel Neuen. 2023. *The Iteration Number of the Weisfeiler-Leman Algorithm*. Technical Report 2301.13317. <https://doi.org/10.48550/arXiv.2301.13317>
- [27] Yuri Gurevich and Saharon Shelah. 1996. On finite rigid structures. *J. Symbol. Logic* 61, 2 (June 1996), 549–562.
- [28] Johan Hästad. 2021. On small-depth Frege proofs for Tseitin for grids. *J. ACM* 68, 1 (Feb. 2021), 1:1–1:31.
- [29] Lauri Hella. 1996. Logical hierarchies in PTIME. *Inf. Comput.* 129, 1 (Aug. 1996), 1–19.
- [30] Neil Immerman. 1981. Number of quantifiers is better than number of tape cells. *J. Comput. Syst. Sci.* 22, 3 (June 1981), 384–406.
- [31] Neil Immerman. 1982. Upper and lower bounds for first order expressibility. *J. Comput. Syst. Sci.* 25, 1 (Aug. 1982), 76–98.
- [32] Neil Immerman and Eric Lander. 1990. Describing graphs: A first-order approach to graph canonization. In *Complexity Theory Retrospective: In Honor of Juris Hartmanis on the Occasion of His Sixtieth Birthday*, Alan L. Selman (Ed.). Springer, 59–81.
- [33] Sandra Kiefer and Brendan D. McKay. 2020. The iteration number of colour refinement. In *Proceedings of the 47th International Colloquium on Automata, Languages and Programming (ICALP'20)* (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 168), 73:1–73:19.
- [34] Sandra Kiefer and Pascal Schweitzer. 2016. Upper bounds on the quantifier depth for graph differentiation in first order logic. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'16)*, 287–296.
- [35] Sandra Kiefer and Pascal Schweitzer. 2019. Upper bounds on the quantifier depth for graph differentiation in first-order logic. *Logic. Meth. Comput. Sci.* 15, 2 (2019).
- [36] Arist Kojevnikov and Dmitry Itsykson. 2006. Lower bounds of static Lovász–Schrijver calculus proofs for Tseitin tautologies. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP'06)* (Lecture Notes in Computer Science, Vol. 4051). Springer, 323–334.
- [37] Jan Krajíček. 2019. *Proof Complexity* (Encyclopedia of Mathematics and Its Applications, Vol. 170). Cambridge University Press.
- [38] Andreas Krebs and Oleg Verbitsky. 2015. Universal covers, color refinement, and two-variable counting logic: Lower bounds for the depth. In *Proceedings of the 30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'15)*, 689–700.
- [39] Moritz Lichter, Ilia Ponomarenko, and Pascal Schweitzer. 2019. Walk refinement, walk logic, and the iteration number of the Weisfeiler-Leman algorithm. In *Proceedings of the 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'19)*, 1–13.
- [40] Mladen Mikša and Jakob Nordström. 2015. A generalized method for proving polynomial calculus degree lower bounds. In *Proceedings of the 30th Annual Computational Complexity Conference (CCC'15)* (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 33), 467–487.
- [41] Jakob Nordström. 2013. Pebble games, proof complexity and time-space trade-offs. *Logic. Meth. Comput. Sci.* 9, 3, Article 15 (Sept. 2013), 15:1–15:63 pages.
- [42] Toniann Pitassi, Benjamin Rossman, Rocco Servedio, and Li-Yang Tan. 2016. Poly-logarithmic Frege depth lower bounds. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC'16)*, 644–657.

- [43] Alexander A. Razborov. 2016. A new kind of tradeoffs in propositional proof complexity. *J. ACM* 63, 2, Article 16 (April 2016), 16:1–16:14 pages.
- [44] Grant Schoenebeck. 2008. Linear level Lasserre lower bounds for certain k -CSPs. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS'08)*. 593–602.
- [45] Alasdair Urquhart. 1987. Hard examples for resolution. *J. ACM* 34, 1 (Jan. 1987), 209–219.
- [46] Moshe Y. Vardi. 1995. On the complexity of bounded-variable queries (extended abstract). In *Proceedings of the 14th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'95)*. 266–276.

Received 4 October 2016; revised 22 May 2023; accepted 28 May 2023