The Chromatic Number of Kneser Hypergraphs via Consensus Division

Ishay Haviv
School of Computer Science, The Academic College of Tel Aviv-Yaffo, Israel

Abstract

We show that the Consensus Division theorem implies lower bounds on the chromatic number of Kneser hypergraphs, offering a novel proof for a result of Alon, Frankl, and Lovász (Trans. Amer. Math. Soc., 1986) and for its generalization by Kriz (Trans. Amer. Math. Soc., 1992). Our approach is applied to study the computational complexity of the total search problem KNESER^p, which given a succinct representation of a coloring of a p-uniform Kneser hypergraph with fewer colors than its chromatic number, asks to find a monochromatic hyperedge. We prove that for every prime p, the KNESER^p problem with an extended access to the input coloring is efficiently reducible to a quite weak approximation of the Consensus Division problem with p shares. In particular, for $p = 2$, the problem is efficiently reducible to *any non-trivial* approximation of the Consensus Halving problem on normalized monotone functions. We further show that for every prime p, the KNESER^p problem lies in the complexity class PPA-p. As an application, we establish limitations on the complexity of the KNESER^p problem, restricted to colorings with a bounded number of colors.

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness; Mathematics of computing → Graph theory

Keywords and phrases Kneser hypergraphs, consensus division, the complexity classes PPA-p

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.60

Related Version *Full Version*: `http://arxiv.org/abs/2311.09016`

Funding Research supported in part by the Israel Science Foundation (grant No. 1218/20).

Acknowledgements We are grateful to Aris Filos-Ratsikas for a fruitful discussion and for a clarification on [9] and to the anonymous reviewers for their useful comments.

1 Introduction

This paper is concerned with two classic problems: the graph-theoretic problem of determining the *chromatic number of Kneser hypergraphs* and the *consensus division* problem from the area of fair division, which lies at the intersection of economics, mathematics, and computer science. We present a novel direct connection between the two problems, offering a new proof for a result of Alon, Frankl, and Lovász [3] on the chromatic number of Kneser hypergraphs as well as for its generalization by Kriz [28]. We use this connection to study the computational complexity of the total search problems associated with Kneser hypergraphs and with approximate consensus division. In what follows, we provide some background on the two problems and on their computational aspects, and then describe our contribution.

Kneser hypergraphs

For an integer $r \geq 2$ and a set family $\mathcal{F}$, the r-uniform Kneser hypergraph $K^r(\mathcal{F})$ is the hypergraph on the vertex set $\mathcal{F}$, whose hyperedges are all the r-subsets of $\mathcal{F}$ whose members are pairwise disjoint. For integers n and k with $n \geq r \cdot k$, let $K^r(n,k)$ denote the hypergraph $K^r(\mathcal{F})$, where $\mathcal{F} = \binom{[n]}{k}$ is the family of all k-subsets of $[n] = \{1, 2, \ldots, n\}$. When $r = 2$, the superscript r may be omitted. The chromatic number of a hypergraph H, denoted by $\chi(H)$, is the minimum number of colors that allow a proper coloring of its vertices, that is, a coloring with no monochromatic hyperedge.

The study of the chromatic number of Kneser graphs was initiated in 1955 by Kneser [27], who observed that the graph $K(n,k)$ admits a proper coloring with $n-2k+2$ colors and conjectured that fewer colors do not suffice, that is, $\chi(K(n,k)) = n-2k+2$. Lovász [29] confirmed Kneser's conjecture in 1978, and his result was extended in multiple ways over the years. One extension, due to Schrijver [34], showed that the subgraph $S(n,k)$ of $K(n,k)$, induced by the k-subsets of $[n]$ that are stable (i.e., include no two consecutive elements modulo n) has the same chromatic number. Another extension, due to Alon, Frankl, and Lovász [3], confirmed a conjecture of Erdös [13] by showing that $\chi(K^r(n,k)) = \lceil \frac{n-r(k-1)}{r-1} \rceil$ for all integers $r \geq 2$. Their lower bound on $\chi(K^r(n,k))$ was further generalized by Kriz [28], as stated below, using a quantity of set families, called the r-colorability defect and denoted by cd_r (see Definition 7; see also [12]).

▶ **Theorem 1** ([28]). *For every integer $r \geq 2$ and for every family $\mathcal{F}$ of non-empty sets,*

$$\chi(K^r(\mathcal{F})) \geq \left\lceil \frac{\mathrm{cd}_r(\mathcal{F})}{r-1} \right\rceil.$$

While Theorem 1 implies the tight lower bound of Alon et al. [3] on the chromatic number of $K^r(n,k)$ (see Lemma 8), it does not cover the aforementioned result of Schrijver [34]. In an attempt to simultaneously generalize both the results, additional extensions were established by several authors, e.g., Meunier [32], Alishahi and Hajiabolhassan [1], Aslam et al. [5], and Frick [18].

It is interesting to mention that although the statements of the above results are all purely combinatorial, their proofs rely on topological tools. Lovász's lower bound [29] on the chromatic number of $K(n,k)$ was based on the Borsuk–Ulam theorem [7], a fundamental result in algebraic topology, and his approach pioneered the area of topological combinatorics. The extension to hypergraphs by Alon et al. [3] was based on a theorem of Bárány, Shlosman, and Szücs [6] that generalizes the Borsuk–Ulam theorem. It was shown by Matoušek [30] that the chromatic number of $K(n,k)$ can also be determined as an application of Tucker's lemma, a combinatorial analogue of the Borsuk–Ulam theorem. His machinery was further developed by Ziegler [36] and later by Meunier [32] to provide alternative proofs for the results of [3, 28, 34].

Consensus division

Another area that extensively applies topological tools is fair division, where the goal is to find fair allocations of resources among several parties. In the consensus division scenario, given m continuous valuation functions $v_1, \ldots, v_m$ defined on subsets of the unit interval $[0,1]$, we aim to divide the interval into r (not necessarily connected) pieces $A_1, \ldots, A_r$ using as few cuts as possible, such that each function assigns the same value to the r pieces, namely, $v_i(A_t) = v_i(A_{t'})$ for all $i \in [m]$ and $t, t' \in [r]$. For the case $r = 2$, referred to as consensus halving, the Hobby–Rice theorem [25] asserts that for additive valuation functions, there always exists such a division with at most m cuts. For a general $r \geq 2$, Alon [2] used the generalization of [6] of the Borsuk–Ulam theorem to show that if the valuation functions are probability measures on $[0,1]$, then there exists a consensus division with at most $(r-1)\cdot m$ cuts. Note that this result played a central role in his proof of the Splitting Necklaces theorem [2]. In fact, it turns out that if the number of desired pieces is a prime p, then the existence of a consensus division with at most $(p-1)\cdot m$ cuts is guaranteed for any given continuous valuation functions, which unlike probability measures, are not required to be additive nor non-negative. This extension was provided for $p = 2$ by Simmons and Su [35] and for any prime p by Filos-Ratsikas, Hollender, Sotiraki, and Zampetakis [16] (see Theorem 9).

Computational aspects

The complexity class TFNP, introduced in [31], consists of the total search problems in NP, i.e., the search problems for which every instance is guaranteed to have a solution, where the solutions can be verified in polynomial time. In 1994, Papadimitriou [33] introduced several sub-classes of TFNP that express the mathematical arguments which yield the existence of the solutions for their problems. In particular, he defined for every prime p, the complexity class PPA-p, associated with the principle that every bipartite graph that has a vertex whose degree is not a multiple of p, has another such vertex. Consequently, PPA-p is the class of problems in TFNP that are efficiently reducible to the corresponding canonical problem: Given a Boolean circuit that represents a bipartite graph and given a vertex whose degree is not a multiple of p, find another such vertex. This family of classes was further studied and developed by Hollender [26] and by Göös, Kamath, Sotiraki, and Zampetakis [21], who independently introduced the classes PPA-r for composites r. Note that the class PPA-2 is denoted by PPA, and that its analogue for directed graphs, called PPAD, is known to be contained in PPA-r for all integers $r \geq 2$ [26, 21].

For any integer $r \geq 2$, let Con-r-Division denote the computational problem, which given an access to m continuous valuation functions $v_1, \ldots, v_m$ over $[0, 1]$, asks to find a consensus division into r pieces $A_1, \ldots, A_r$ with at most $(r-1) \cdot m$ cuts. In its approximate version with precision parameter ε, the pieces should satisfy $|v_i(A_t) - v_i(A_{t'})| \leq \varepsilon$ for all $i \in [m]$ and $t, t' \in [r]$. The problem may be considered for various families of input functions, ranging from piecewise-constant functions and probability measures to monotone functions and general functions. While piecewise-constant functions can be explicitly represented by the endpoints and values of the intervals on which they are nonzero, in the more general settings the functions are given by some succinct representation, e.g., arithmetic circuits or efficient Turing machines that compute them. It follows from [16] that for all primes p, every instance of the Con-p-Division problem has a solution even for $\varepsilon = 0$, hence the problem is total.

A considerable attention has been devoted in recent years to the Con-r-Division problem with $r = 2$, referred to as Con-Halving. Filos-Ratsikas and Goldberg [14, 15] proved that the problem is PPA-complete, even for piecewise-constant functions and inverse-polynomial precision parameter ε, and derived the PPA-completeness of the Splitting Necklaces with two thieves and Discrete Sandwich problems. Their result was strengthened and extended in several ways. Deligkas, Filos-Ratsikas, and Hollender [9] proved that the problem remains PPA-hard when the number of valuation functions is a fixed constant: for two or more general functions, and for three or more monotone functions. Filos-Ratsikas, Hollender, Sotiraki, and Zampetakis [17] proved that the problem remains PPA-hard when the valuation functions are piecewise-uniform with only two blocks. More recently, Deligkas, Fearnley, Hollender, and Melissourgos [8] proved that the problem is PPA-hard for any precision parameter $\varepsilon < 1/5$, even when the valuation functions are piecewise-uniform with three blocks. An additional version of the Con-Halving problem, where the goal is not to partition an interval but an unordered collection of items, was studied by Goldberg, Hollender, Igarashi, Manurangsi, and Suksompong [20].

The complexity of the Con-r-Division problem for a general r is much less understood. It was proved by Filos-Ratsikas, Hollender, Sotiraki, and Zampetakis [16] that for every prime p, Con-p-Division lies in the complexity class PPA-p, and the question of whether it is PPA-p-hard was left open. Yet, for $p = 3$, it was shown in [17] that Con-3-Division is PPAD-hard for an inverse-exponential precision parameter ε. On the algorithmic side, Alon and Graur [4] proved that given an access to m probability measures over $[0, 1]$, it is possible to find in polynomial time a partition of the unit interval into r pieces with at most $(r-1)\cdot m$ cuts, such that every piece receives at least $\frac{1}{m\cdot r}$ of each measure. They also considered the case where a larger number of cuts is allowed, and showed that for a given precision parameter ε, it is possible to efficiently find a solution with $O(\varepsilon^{-2}\cdot m\log m)$ cuts. Goldberg and Li [19] proved that the Con-r-Division problem on m probability measures, where the number of allowed cuts is $2(r-1)(p-1)\cdot\frac{\lceil p/2\rceil}{\lfloor p/2\rfloor}\cdot m$ for a prime p, lies in the complexity class PPA-p. For example, the Con-Halving problem on m probability measures with $8m$ allowed cuts lies in PPA-3.

Before turning to our results, let us introduce, for any integer $r\geq 2$, the computational search problem Kneser^r associated with the family of r-uniform Kneser hypergraphs $K^r(n,k)$. The input consists of integers n and k with $n\geq r\cdot k$ along with a Boolean circuit that represents a coloring of the vertices of the hypergraph $K^r(n,k)$ with $\lfloor\frac{n-r(k-1)-1}{r-1}\rfloor$ colors, which is smaller by one than its chromatic number [3]. The goal is to find a monochromatic hyperedge, that is, r pairwise disjoint vertices that are assigned the same color by the input coloring. Omitting the superscript r when $r=2$, the Kneser problem is known to lie in PPA, and it is an open question whether it is PPA-hard, as suggested by Deng, Feng, and Kulkarni [10]. More generally, it was asked in [16] whether for every prime p, the Kneser^p problem lies in PPA-p and if it is PPA-p-hard. While no hardness result is known for these problems, it was shown in [22] that the Schrijver problem, which asks to find a monochromatic edge in a graph $S(n,k)$ given a coloring of its vertices with $n-2k+1$ colors, is PPA-complete. It was recently shown in [24] that the problem of finding a monochromatic edge in a graph $S(n,k)$, given a coloring of its vertices with only $\lfloor n/2\rfloor-2k+1$ colors, is efficiently reducible to the Kneser problem.

1.1 Our Contribution

This paper presents a novel direct connection between the chromatic number of Kneser hypergraphs and the consensus division problem. As our first contribution, we offer a new proof of Kriz's lower bound on the chromatic number of Kneser hypergraphs [28], stated earlier as Theorem 1. The proof relies on the Consensus Division theorem of Filos-Ratsikas et al. [16]. Our technique borrows and extends ideas that were applied in [20] and in [23].

We then adopt a computational perspective and explore our approach to Theorem 1 as a reduction from the Kneser^p problem to the Con-p-Division problem for any prime p. Our main contribution is an efficient reduction from the Kneser^p problem with an extended access to the input coloring to a quite weak approximation of the Con-p-Division problem. Before the precise statements, let us introduce the following variants of the studied computational problems. Their formal definitions are given in Sections 4.1 and 4.2.

- The Kneser^p problem with *subset queries*: As in the standard Kneser^p problem, the input is a coloring of the vertices of a hypergraph $K^p(n,k)$ with fewer colors than its chromatic number, and the goal is to find a monochromatic hyperedge. Here, however, the access to the coloring allows, in addition to queries for the colors of the vertices, another type of queries called subset queries. Such a query involves a subset D of $[n]$ and a color i, and the answer on the pair (D,i) determines whether D contains a vertex colored i. The notion of subset queries was proposed in [23].
- The Con-p-Division$[<\varepsilon]$ problem: As in the standard Con-p-Division problem, the input consists of m continuous valuation functions $v_1,\dots,v_m$ over $[0,1]$, and a solution is a partition of the unit interval into p pieces $A_1,\dots,A_p$ with at most $(p-1)\cdot m$ cuts.

Here, however, the solution is required to satisfy

$$|v_i(A_t) - v_i(A_{t'})| < \varepsilon \quad (1)$$

for all $i \in [m]$ and $t, t' \in [p]$. Namely, the difference from the standard CON-p-DIVISION problem with precision parameter ε is that the inequality in (1) is strict. In particular, for $\varepsilon = 1$, assuming that the valuation functions are normalized (i.e., return values in $[0, 1]$), the solution is just required to be *non-trivial*. This means that the solution is just not allowed to include two pieces A_t and $A_{t'}$ such that $v_i(A_t) = 1$ and $v_i(A_{t'}) = 0$ for some $i \in [m]$, but the value of $|v_i(A_t) - v_i(A_{t'})|$ may approach 1 as m grows. When $p = 2$, the problem is denoted by CON-HALVING[$< \varepsilon$].

We prove the following theorem, which concerns the case $p = 2$.

▶ **Theorem 2.** *There exists a polynomial-time reduction from the KNESER problem with subset queries to the CON-HALVING[< 1] problem on normalized monotone functions.*

As alluded to before, the complexity of the KNESER problem (with or without subset queries) is not well understood. Theorem 2 implies that any hardness result for the KNESER problem with subset queries would imply a very strong hardness result for the CON-HALVING problem on normalized monotone functions, ruling out the possibility to obtain an efficient algorithm for any non-trivial approximation of the problem. On the other hand, an efficient algorithm for CON-HALVING that guarantees some non-trivial approximation on normalized monotone functions would imply an efficient algorithm for the KNESER problem with subset queries. We find these consequences of the reduction quite surprising and unusual, especially because of the discrete nature of the KNESER problem. For comparison, the efficient reduction from the (discrete) Splitting Necklaces problem with two thieves to the CON-HALVING problem with precision parameter ε, which is given in [14] and builds on an argument of [2], requires ε to be inverse-polynomial in the number of valuation functions (note that those functions are additive, though). Let us stress that Theorem 2 addresses the CON-HALVING[< 1] problem when restricted to normalized monotone valuation functions but not to probability measures. Recall that for this stronger restriction, an algorithm of [4] does provide a non-trivial solution in polynomial time. We finally note that the proof of Theorem 2 essentially supplies a reduction to the version of CON-HALVING, studied in [20], of finding a consensus halving of an unordered collection of items rather than of the unit interval. This makes the result stronger, as this version is efficiently reducible to the standard one.

Our next result relates the KNESER^p problem with subset queries to the CON-p-DIVISION problem for every prime $p \geq 3$. Here, the precision parameter of the latter is $\frac{1}{2}$.

▶ **Theorem 3.** *For every prime $p \geq 3$, there exists a polynomial-time reduction from the KNESER^p problem with subset queries to the CON-p-DIVISION[$< \frac{1}{2}$] problem on normalized monotone functions.*

It is noteworthy that Theorems 2 and 3 are proved in a more general form. Namely, we reduce from a generalized variant of the KNESER^p problem, where the input is a coloring of a hypergraph $K^p(\mathcal{F})$ for some set family $\mathcal{F}$ taken from a prescribed sequence of set families, which is assumed to be efficiently computable (see Definitions 12 and 13). The number of colors used by the input coloring may be any number smaller than Kriz's lower bound on the chromatic number of $K^p(\mathcal{F})$, as given by Theorem 1. For the precise statements, see Theorems 16 and 17.

We proceed with additional results on the computational complexity of the $\textsc{Kneser}^p$ problem (in its standard version, without subset queries). The following theorem settles a question of [16].

▶ **Theorem 4.** *For every prime p, the* $\textsc{Kneser}^p$ *problem lies in* PPA-p.

In fact, we provide two results that strengthen Theorem 4 in two incomparable forms. Firstly, as before, we present a generalized result handling a variant of the $\textsc{Kneser}^p$ problem of finding a monochromatic hyperedge in a hypergraph $K^p(\mathcal{F})$, given a coloring that uses fewer colors than Kriz's lower bound on its chromatic number. Again, $\mathcal{F}$ is taken from a prescribed sequence of set families satisfying certain computational assumptions (see Definition 20 and Corollary 23). Secondly, we show that the membership in PPA-p holds for the $\textsc{Kneser}^p_{\widetilde{\text{stab}}}$ problem. This problem asks to find a monochromatic hyperedge in the sub-hypergraph of $K^p(n,k)$ induced by the k-subsets of $[n]$ that are *almost stable* (i.e., include no two consecutive elements, but can include both 1 and n). The input coloring may use here any number of colors smaller than the chromatic number of this hypergraph, which, as proved in [32], is equal to that of $K^p(n,k)$.

▶ **Theorem 5.** *For every prime p, the* $\textsc{Kneser}^p_{\widetilde{\text{stab}}}$ *problem lies in* PPA-p.

The proofs of Theorems 4 and 5 crucially rely on the mathematical arguments due to Ziegler [36] and Meunier [32], which imply the totality of the studied problems. We verify that their arguments can be transformed into efficient reductions to a computational problem associated with a $\mathbb{Z}_p$-variant of Tucker's lemma, which was shown to lie in the complexity class PPA-p by Filos-Ratsikas et al. [16].

Finally, we apply Theorems 4 and 5 to derive limitations on the complexity of variants of the $\textsc{Kneser}^r$ problem, restricted to colorings with a bounded number of colors (see Theorem 26 and Corollary 28). In particular, we show that the problem of finding a monochromatic edge in a graph $S(n,k)$ given a coloring of its vertices with $\lfloor n/2 \rfloor - 2k + 1$ colors, lies in the complexity class PPA-3 (see Corollary 29). As mentioned earlier, it was shown in [24] that the latter is efficiently reducible to the $\textsc{Kneser}$ problem. It thus follows that, unless PPA $\subseteq$ PPA-3, this reduction cannot yield a PPA-hardness result for the $\textsc{Kneser}$ problem. We note that it is common to believe that the classes PPA-p for primes p do not contain each other, and that an unconditional separation between their black-box versions was provided in [26, 21].

1.2 Outline

The rest of the paper is organized as follows. In Section 2, we gather some definitions that will be used throughout the paper. In Section 3, we present our novel proof of Theorem 1. In Section 4, we consider this proof from a computational perspective and state generalized forms of Theorems 2 and 3. Finally, in Section 5, we state a generalized form of Theorem 4 and obtain some limitations on the complexity of the $\textsc{Kneser}^r$ problem. All missing proofs can be found in the full version of the paper.

2 Preliminaries

For integers n and k, let $\binom{[n]}{k}$ denote the family of k-subsets of $[n] = \{1, 2, \dots, n\}$. A subset of $[n]$ is called *stable* if it does not include two consecutive elements modulo n, equivalently, it forms an independent set in the cycle on the vertex set $[n]$ with the natural order along the cycle. Let $\binom{[n]}{k}_{\text{stab}}$ denote the family of stable k-subsets of $[n]$. A subset of $[n]$ is called *almost stable* if it does not include two consecutive elements, equivalently, it forms an independent set in the path on the vertex set $[n]$ with the natural order along the path. Let $\binom{[n]}{k}_{\widetilde{\text{stab}}}$ denote the family of almost stable k-subsets of $[n]$. Note that $\binom{[n]}{k}_{\text{stab}} \subseteq \binom{[n]}{k}_{\widetilde{\text{stab}}} \subseteq \binom{[n]}{k}$.

The family of Kneser hypergraphs is defined as follows.

▶ **Definition 6** (Kneser Hypergraphs). *For an integer $r \geq 2$ and a set family $\mathcal{F}$, the r-uniform Kneser hypergraph $K^r(\mathcal{F})$ is the hypergraph on the vertex set $\mathcal{F}$, whose hyperedges are all the r-subsets of $\mathcal{F}$ whose members are pairwise disjoint. For integers n and k with $n \geq r \cdot k$, let $K^r(n,k)$, $K^r(n,k)_{\widetilde{\text{stab}}}$, and $K^r(n,k)_{\text{stab}}$, respectively, denote the hypergraphs $K^r(\binom{[n]}{k})$, $K^r(\binom{[n]}{k}_{\widetilde{\text{stab}}})$, and $K^r(\binom{[n]}{k}_{\text{stab}})$. When $r = 2$, the superscript r may be omitted.*

For an integer $t \geq 1$, a hypergraph $H = (V, E)$ is said to be *t-colorable* if it admits a proper t-coloring, that is, a coloring of the vertices of V with t colors such that no hyperedge of E is monochromatic. The *chromatic number* of H, denoted by $\chi(H)$, is the smallest integer t for which H is t-colorable. It is known (see [3, 32, 18]) that for all integers $r \geq 2$, k, and $n \geq r \cdot k$, it holds that

$$\chi(K^r(n,k)) = \chi(K^r(n,k)_{\widetilde{\text{stab}}}) = \chi(K^r(n,k)_{\text{stab}}) = \left\lceil \frac{n - r(k-1)}{r-1} \right\rceil. \tag{2}$$

Theorem 1, proved by Kriz [28], relates the chromatic number of Kneser hypergraphs to the notion of *colorability defect*, defined as follows.

▶ **Definition 7** (Colorability Defect). *For an integer $r \geq 2$ and a family $\mathcal{F}$ of non-empty subsets of a set X, the r-*colorability defect *of $\mathcal{F}$, denoted by $\mathrm{cd}_r(\mathcal{F})$, is the smallest size of a set $Y \subseteq X$, such that the sub-hypergraph of $(X, \mathcal{F})$ induced by $X \setminus Y$ is r-colorable, that is,*

$$\mathrm{cd}_r(\mathcal{F}) = \min \left\{ |Y| \;\middle|\; (X \setminus Y, \{e \in \mathcal{F} \mid e \cap Y = \emptyset\}) \text{ is } r\text{-colorable} \right\}.$$

The following lemma, given in [36, Lemma 3.2], determines the r-colorability defect of the family $\binom{[n]}{k}$.

▶ **Lemma 8** ([36]). *For integers n, k, r with $n \geq r \cdot k$, it holds that $\mathrm{cd}_r(\binom{[n]}{k}) = n - r(k-1)$.*

We next state the Consensus Division theorem due to Filos-Ratsikas et al. [16, Theorem 6.5]. In what follows, let $\mathcal{B}([0,1])$ denote the set of Lebesgue-measurable subsets of the interval $[0,1]$, and let $\mu : \mathcal{B}([0,1]) \to [0,1]$ denote the Lebesgue measure on $[0,1]$. In addition, let $\triangle$ stand for the symmetric difference of sets, defined by $E_1 \triangle E_2 = (E_1 \setminus E_2) \cup (E_2 \setminus E_1)$.

▶ **Theorem 9** (Consensus Division Theorem [16]). *Let p be a prime, and let $m \geq 1$ be an integer. Let $v_1, \ldots, v_m : \mathcal{B}([0,1]) \to \mathbb{R}$ be functions such that for each $i \in [m]$, v_i satisfies the following continuity condition: for any $\varepsilon > 0$ there exists $\delta > 0$ such that for all $E_1, E_2 \in \mathcal{B}([0,1])$ with $\mu(E_1 \triangle E_2) \leq \delta$, it holds that $|v_i(E_1) - v_i(E_2)| \leq \varepsilon$. Then, there exists a consensus-p-division, that is, it is possible to partition the unit interval into p (not necessarily connected) pieces $A_1, \ldots, A_p$ using at most $(p-1) \cdot m$ cuts, such that $v_i(A_t) = v_i(A_{t'})$ for all $i \in [m]$ and $t, t' \in [p]$.*

As for the continuity property, we will sometimes consider the stronger notion of Lipschitz-continuity. For $L \geq 0$, a function $v : \mathcal{B}([0,1]) \to \mathbb{R}$ is said to be *L-Lipschitz-continuous*, if for all $E_1, E_2 \in \mathcal{B}([0,1])$, it holds that $|v(E_1) - v(E_2)| \leq L \cdot \mu(E_1 \triangle E_2)$.

3 The Chromatic Number of Kneser Hypergraphs

In this section, we present our new proof of Kriz's lower bound on the chromatic number of Kneser hypergraphs [28], stated as Theorem 1. The proof applies the Consensus Division theorem, stated as Theorem 9. Note that we focus here on the mathematical proof rather than on its computational aspects, which will be explored in the next section. As is usual for results of this type, we first consider the case where the uniformity r of the hypergraphs is a prime number.

▶ **Theorem 10.** *For every prime p and for every family $\mathcal{F}$ of non-empty sets,*

$$\chi(K^p(\mathcal{F})) \geq \left\lceil \frac{\mathrm{cd}_p(\mathcal{F})}{p-1} \right\rceil.$$

Proof. Let p be a prime, and let $\mathcal{F}$ be a family of non-empty sets. Assume without loss of generality that all members of $\mathcal{F}$ are subsets of $[n]$ for some integer n. Let m be an integer satisfying $m < \frac{\mathrm{cd}_p(\mathcal{F})}{p-1}$, and suppose for the sake of contradiction that there exists a proper coloring $c : \mathcal{F} \to [m]$ of the p-uniform Kneser hypergraph $K^p(\mathcal{F})$.

For each $i \in [m]$, let $\widetilde{u}_i : \{0,1\}^n \to \{0,1\}$ denote the indicator function that determines for every subset of $[n]$ whether it contains a set of $\mathcal{F}$ whose color according to c is i, that is, for every $D \subseteq [n]$, define $\widetilde{u}_i(D) = 1$ if there exists a set $B \in \mathcal{F}$ satisfying $B \subseteq D$ and $c(B) = i$, and define $\widetilde{u}_i(D) = 0$ otherwise. Here and throughout, we identify the subsets of $[n]$ with their characteristic vectors in $\{0,1\}^n$. Note that the function $\widetilde{u}_i$ is monotone with respect to inclusion.

For each $i \in [m]$, consider the extension $u_i : [0,1]^n \to [0,1]$ of $\widetilde{u}_i$ that maps any vector $x \in [0,1]^n$ to the largest value $a \in [0,1]$ such that the set $\{j \in [n] \mid x_j \geq a\}$ contains a set of $\mathcal{F}$ colored i by c if such a value exists, and to 0 otherwise. Equivalently, for any $x \in [0,1]^n$, let π be a permutation of $[n]$ with $x_{\pi(1)} \leq x_{\pi(2)} \leq \cdots \leq x_{\pi(n)}$, and define $u_i(x) = x_{\pi(j)}$ for the largest integer j satisfying $\widetilde{u}_i(\{\pi(j), \pi(j+1), \ldots, \pi(n)\}) = 1$ if such a j exists, and $u_i(x) = 0$ otherwise. Observe that, under the convention $x_{\pi(0)} = 0$, the value of $u_i(x)$ can be written as

$$u_i(x) = \sum_{j=1}^{n} (x_{\pi(j)} - x_{\pi(j-1)}) \cdot \widetilde{u}_i(\{\pi(j), \pi(j+1), \ldots, \pi(n)\}).$$

Notice that the function u_i is an extension of the function $\widetilde{u}_i$. Notice further that u_i is monotone, in the sense that the value of $u_i(x)$ does not decrease when the value of some entry of x increases. Finally, notice that changing some entry of x by ε changes $u_i(x)$ by at most ε, hence the function u_i is continuous.

Now, for each $j \in [n]$, consider the open sub-interval $I_j = (\frac{j-1}{n}, \frac{j}{n})$, and associate it with the element j. For each $i \in [m]$, let $v_i : \mathcal{B}([0,1]) \to [0,1]$ be the function defined as follows. For every $E \in \mathcal{B}([0,1])$, let $x^E \in [0,1]^n$ denote the vector that consists of the normalized Lebesgue measures of E on the sub-intervals $I_1, \ldots, I_n$, that is, $x^E_j = n \cdot \mu(E \cap I_j)$ for all $j \in [n]$, and define $v_i(E) = u_i(x^E)$. Note that the function v_i is the composition of the function u_i with the function that maps any set $E \in \mathcal{B}([0,1])$ to the vector x^E. Since the function u_i is monotone and continuous, it follows that so is v_i

Applying Theorem 9, we obtain that there exists a consensus-p-division of $v_1, \ldots, v_m$, that is, a partition of the unit interval into p pieces $A_1, \ldots, A_p$ using at most $(p-1) \cdot m$ cuts, such that $v_i(A_t) = v_i(A_{t'})$ for all $i \in [m]$ and $t, t' \in [p]$. Let Y denote the set of indices $j \in [n]$ for which the open sub-interval I_j includes a cut, and notice that $|Y| \leq (p-1) \cdot m < \mathrm{cd}_p(\mathcal{F})$. Thus, every sub-interval I_j with $j \in [n] \setminus Y$ is fully contained in one of the pieces $A_1, \ldots, A_p$. Consider the coloring that assigns to every $j \in [n] \setminus Y$ the index $t \in [p]$ of the piece A_t that contains I_j. By $|Y| < \mathrm{cd}_p(\mathcal{F})$, the hypergraph $([n] \setminus Y, \{e \in \mathcal{F} \mid e \cap Y = \emptyset\})$ is not p-colorable. This implies that for some $t \in [p]$, there exists a set $B \in \mathcal{F}$ all of whose elements share the color t, and thus, for all $j \in B$, it holds that $I_j \subseteq A_t$. Denoting by $\ell = c(B)$ the color assigned to B by the given coloring c, it follows from the definition of v_ℓ that $v_\ell(\cup_{j \in B} I_j) = 1$. By monotonicity, it further follows that $v_\ell(A_t) = 1$, which yields, by Theorem 9, that $v_\ell(A_{t'}) = 1$ for all $t' \in [p]$.

Finally, for every $t' \in [p]$, the fact that $v_\ell(A_{t'}) = 1$ implies that there exists a set $B_{t'} \in \mathcal{F}$ with $c(B_{t'}) = \ell$ such that $\mu(A_{t'} \cap I_j) = \frac{1}{n}$ for all $j \in B_{t'}$. Since the pieces $A_1, \ldots, A_p$ are pairwise disjoint, so are the sets $B_1, \ldots, B_p$. It thus follows that these sets form a monochromatic hyperedge in $K^p(\mathcal{F})$, in contradiction to the assumption that the coloring c is proper. This completes the proof. ◀

It is well known that Theorem 10 implies Theorem 1 (see, e.g., [36]). We provide the quick proof for completeness.

Proof of Theorem 1. Theorem 10 shows that Theorem 1 holds whenever r is prime. It therefore suffices to prove that for every pair of integers $r_1, r_2 \geq 2$, if the theorem holds for $r \in \{r_1, r_2\}$ then it holds for $r = r_1 r_2$. So suppose that it holds for r_1 and r_2. Let $\mathcal{F}$ be a family of non-empty subsets of $[n]$ for an integer n, and let $c : \mathcal{F} \to [m]$ be a proper coloring of the hypergraph $K^{r_1 r_2}(\mathcal{F})$ for an integer m. Our goal is to prove that $m \geq \frac{\mathrm{cd}_{r_1 r_2}(\mathcal{F})}{r_1 r_2 - 1}$.

Consider the family

$$\mathcal{G} = \Big\{ G \subseteq [n] \;\Big|\; \mathrm{cd}_{r_1}(\mathcal{F}|_G) > m(r_1 - 1) \Big\},$$

where $\mathcal{F}|_G = \{F \in \mathcal{F} \mid F \subseteq G\}$. Define a coloring $c' : \mathcal{G} \to [m]$ as follows. For every $G \in \mathcal{G}$, let $c'(G)$ be a color of some monochromatic hyperedge in $K^{r_1}(\mathcal{F}|_G)$ with respect to the coloring c. The existence of such a hyperedge for every $G \in \mathcal{G}$ follows from our assumption that the theorem holds for r_1, which yields that $\chi(K^{r_1}(\mathcal{F}|_G)) \geq \frac{\mathrm{cd}_{r_1}(\mathcal{F}|_G)}{r_1 - 1} > m$, where the second inequality is due to the definition of $\mathcal{G}$. We claim that c' is a proper coloring of $K^{r_2}(\mathcal{G})$, and thus $\chi(K^{r_2}(\mathcal{G})) \leq m$. Indeed, otherwise there would exist r_2 pairwise disjoint sets $G_1, \ldots, G_{r_2} \in \mathcal{G}$ that are assigned the same color by c'. Since each set G_i with $i \in [r_2]$ contains r_1 pairwise disjoint sets of $\mathcal{F}|_{G_i}$ that are assigned by c the color $c'(G_i)$, this gives us $r_1 r_2$ pairwise disjoint sets of $\mathcal{F}$ with the same color by c, contradicting the assumption that c is a proper coloring of $K^{r_1 r_2}(\mathcal{F})$.

Finally, by the assumption that the theorem holds for r_2, we have $\chi(K^{r_2}(\mathcal{G})) \geq \frac{\mathrm{cd}_{r_2}(\mathcal{G})}{r_2 - 1}$, which implies that $\mathrm{cd}_{r_2}(\mathcal{G}) \leq m(r_2 - 1)$. This means that it is possible to remove at most $m(r_2 - 1)$ of the elements of $[n]$ and to partition the remaining elements into r_2 sets $F_1, \ldots, F_{r_2}$, such that no set of $\mathcal{G}$ is contained in any of them. In particular, for each $i \in [r_2]$, it holds that $F_i \notin \mathcal{G}$, which implies that $\mathrm{cd}_{r_1}(\mathcal{F}|_{F_i}) \leq m(r_1 - 1)$. This means that it is possible to remove at most $m(r_1 - 1)$ elements from each F_i and to partition the remaining elements into r_1 sets, such that no set of $\mathcal{F}|_{F_i}$ is contained in any of them. It thus follows that one can remove at most $m(r_2 - 1) + r_2 \cdot m(r_1 - 1) = m(r_1 r_2 - 1)$ elements from $[n]$ and partition the remaining elements into $r_1 r_2$ sets, such that no set of $\mathcal{F}$ is contained in any of them. This implies that $\mathrm{cd}_{r_1 r_2}(\mathcal{F}) \leq m(r_1 r_2 - 1)$, providing the desired lower bound on m. ◀

We conclude this section by verifying that $\chi(K^r(n,k)) = \lceil \frac{n - r(k-1)}{r-1} \rceil$ for all $r \geq 2$. The lower bound on $\chi(K^r(n,k))$, originally proved in [3], follows by combining Theorem 1 with Lemma 8. For the upper bound, which is given in [13], put $t = \lceil \frac{n - r(k-1)}{r-1} \rceil$, let $X_1, \ldots, X_{t-1}$ be $t - 1$ pairwise disjoint $(r-1)$-subsets of $[n]$, and observe that the number of elements of $[n[$ that do not belong to any of these sets is $n - (t-1)(r-1) \leq rk - 1$. Consider the coloring that assigns to every k-subset B of $[n]$ a color $i \in [t-1]$ such that $B \cap X_i \neq \emptyset$ if such an i exists, and the color t otherwise. It can be easily checked that this coloring of $K^r(n,k)$ is proper.

4 Reduction from Kneser^p to Approximate Con-p-Division

In this section, we use the proof technique presented in Section 3 to obtain efficient reductions from the computational problems associated with Kneser hypergraphs to those associated with approximate consensus division. We start by formally introducing the involved computational problems.

4.1 The Kneser^r Problem

In order to make our results as strong as possible, we introduce a problem of finding monochromatic hyperedges in general Kneser hypergraphs, defined as follows.

▶ **Definition 11** (The $\textsc{Kneser}^r(\mathcal{F}, m)$ Problem)**.** *For a set $\mathcal{A}$, let $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha \in \mathcal{A}}$ be a sequence of set families, where for each $\alpha \in \mathcal{A}$, $\mathcal{F}_\alpha$ is a family of non-empty subsets of $[n_\alpha]$ for some integer n_α, and let $m : \mathcal{A} \to \mathbb{N}$ be a function. In the $\textsc{Kneser}^r(\mathcal{F}, m)$ problem, the input consists of*

- *an element $\alpha \in \mathcal{A}$ and*
- *a Boolean circuit $C : \{0,1\}^{n_\alpha} \to [m(\alpha)]$ that represents a coloring $c : \mathcal{F}_\alpha \to [m(\alpha)]$ of the sets of $\mathcal{F}_\alpha$ with $m(\alpha)$ colors, in the sense that for every $B \in \mathcal{F}_\alpha$, it holds that $C(B) = c(B)$.*

The goal is to find a monochromatic hyperedge in $K^r(\mathcal{F}_\alpha)$, that is, r pairwise disjoint sets $B_1, B_2, \ldots, B_r \in \mathcal{F}_\alpha$ such that $C(B_1) = C(B_2) = \cdots = C(B_r)$.

Note that for every sequence $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha \in \mathcal{A}}$ and for every function $m : \mathcal{A} \to \mathbb{N}$ for which it holds that $m(\alpha) < \chi(K^r(\mathcal{F}_\alpha))$ for all $\alpha \in \mathcal{A}$, the $\textsc{Kneser}^r(\mathcal{F}, m)$ problem is total. In particular, by Theorem 1, the $\textsc{Kneser}^r(\mathcal{F}, \lfloor \frac{\mathrm{cd}_r(\mathcal{F}_\alpha)-1}{r-1} \rfloor)$ problem is total.

We next define a variant of the $\textsc{Kneser}^r(\mathcal{F}, m)$ problem with an extended access to the input coloring. This variant is referred to as the $\textsc{Kneser}^r(\mathcal{F}, m)$ problem with *subset queries*.

▶ **Definition 12** (The $\textsc{Kneser}^r(\mathcal{F}, m)$ Problem with Subset Queries)**.** *For a set $\mathcal{A}$, let $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha \in \mathcal{A}}$ be a sequence of set families, where for each $\alpha \in \mathcal{A}$, $\mathcal{F}_\alpha$ is a family of non-empty subsets of $[n_\alpha]$ for some integer n_α, and let $m : \mathcal{A} \to \mathbb{N}$ be a function. In the $\textsc{Kneser}^r(\mathcal{F}, m)$ problem* with subset queries*, the input consists of*

- *an element $\alpha \in \mathcal{A}$,*
- *a Boolean circuit $C : \{0,1\}^{n_\alpha} \to [m(\alpha)]$ that represents a coloring $c : \mathcal{F}_\alpha \to [m(\alpha)]$ of the sets of $\mathcal{F}_\alpha$ with $m(\alpha)$ colors, in the sense that for every $B \in \mathcal{F}_\alpha$, it holds that $C(B) = c(B)$, and*
- *a Boolean circuit $S : \{0,1\}^{n_\alpha} \times [m(\alpha)] \to \{0,1\}$ that is supposed to allow subset queries to the coloring c, namely, for every set $D \subseteq [n_\alpha]$ and a color $i \in [m(\alpha)]$, it is supposed to satisfy $S(D,i) = 1$ if there exists a set $B \in \mathcal{F}_\alpha$ such that $B \subseteq D$ and $c(B) = i$, and $S(D,i) = 0$ otherwise.*

The goal is to find either a monochromatic hyperedge in $K^r(\mathcal{F}_\alpha)$ or a violation of the circuit S, namely,

- *(false negative) two sets $B, D \subseteq [n_\alpha]$ and a color $i \in [m(\alpha)]$ such that $B \in \mathcal{F}_\alpha$, $B \subseteq D$, $C(B) = i$, and yet $S(D,i) = 0$, or*

- *(false positive) a set $D \subseteq [n_\alpha]$ and a color $i \in [m(\alpha)]$ such that $S(D,i) = 1$ whereas for every set D' obtained from D by removing a single element it holds that $S(D',i) = 0$, and in addition, either $D \notin \mathcal{F}_\alpha$ or $D \in \mathcal{F}_\alpha$ and $C(D) \neq i$.*

Note that by allowing the solutions of the $\textsc{Kneser}^r(\mathcal{F}, m)$ problem with subset queries to form violations of the circuit S, we obtain a non-promise problem.

We will be particularly interested in the following special cases of Definitions 11 and 12. For any integer $r \geq 2$, let $\mathcal{A}^{(r)}$ denote the set of all pairs of integers (n,k) with $n \geq r \cdot k$, and let $\mathcal{F}^{(r)}$, $\mathcal{F}^{(r,\text{stab})}$, and $\mathcal{F}^{(r,\widetilde{\text{stab}})}$ denote, respectively, the sequences of set families defined by

$$\mathcal{F}^{(r)}_{(n,k)} = \binom{[n]}{k}, \qquad \mathcal{F}^{(r,\text{stab})}_{(n,k)} = \binom{[n]}{k}_{\text{stab}}, \qquad \mathcal{F}^{(r,\widetilde{\text{stab}})}_{(n,k)} = \binom{[n]}{k}_{\widetilde{\text{stab}}}$$

for all $(n,k) \in \mathcal{A}^{(r)}$. Consider the function $m^{(r)} : \mathcal{A}^{(r)} \to \mathbb{N}$ defined by $m^{(r)}(n,k) = \lfloor \frac{n-r(k-1)-1}{r-1} \rfloor$, which by (2), satisfies $m^{(r)}(n,k) = \chi(K^r(n,k)) - 1$. We define the three problems $\textsc{Kneser}^r$, $\textsc{Kneser}^r_{\text{stab}}$, and $\textsc{Kneser}^r_{\widetilde{\text{stab}}}$, respectively, to be $\textsc{Kneser}^r(\mathcal{F}^{(r)}, m^{(r)})$, $\textsc{Kneser}^r(\mathcal{F}^{(r,\text{stab})}, m^{(r)})$, and $\textsc{Kneser}^r(\mathcal{F}^{(r,\widetilde{\text{stab}})}, m^{(r)})$. For any function $m : \mathcal{A}^{(r)} \to \mathbb{N}$, we let $\textsc{Kneser}^r_{\text{stab}}(n,k,m)$ denote the $\textsc{Kneser}^r(\mathcal{F}^{(r,\text{stab})}, m)$ problem. When $r = 2$, the superscript r may be omitted. For any function $m : \mathcal{A}^{(2)} \to \mathbb{N}$, we let $\textsc{Kneser}(n,k,m)$ and $\textsc{Schrijver}(n,k,m)$ denote the problems $\textsc{Kneser}(\mathcal{F}^{(2)}, m)$ and $\textsc{Kneser}(\mathcal{F}^{(2,\text{stab})}, m)$ respectively.

We will use the notion of polynomially computable sequences of set families, defined below.

▶ **Definition 13.** *For a set $\mathcal{A}$, let $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha \in \mathcal{A}}$ be a sequence of set families, where for each $\alpha \in \mathcal{A}$, $\mathcal{F}_\alpha$ is a family of non-empty subsets of $[n_\alpha]$ for some integer n_α. The sequence $\mathcal{F}$ is* polynomially computable *if there exist polynomials q_1, q_2 such that*

1. *there exists an algorithm that given an element $\alpha \in \mathcal{A}$ and a set $B \subseteq [n_\alpha]$ runs in time $q_1(n_\alpha)$ and determines whether $B \in \mathcal{F}_\alpha$, and*
2. *there exists an algorithm that given an element $\alpha \in \mathcal{A}$ and a set $D \subseteq [n_\alpha]$ runs in time $q_2(n_\alpha)$, returns a subset of D that belongs to the family $\mathcal{F}_\alpha$ if such a subset exists, and declares that no such subset exists otherwise.*

The following lemma gives simple examples of polynomially computable sequences.

▶ **Lemma 14.** *For every integer $r \geq 2$, the sequence $\mathcal{F}^{(r)}$ is polynomially computable.*

Proof. Fix an integer $r \geq 2$. Recall that $\mathcal{F}^{(r)}_{(n,k)} = \binom{[n]}{k}$ for all $(n,k) \in \mathcal{A}^{(r)}$. We show that $\mathcal{F}^{(r)}$ satisfies the conditions of Definition 13. For Item 1, consider the algorithm that given a pair $(n,k) \in \mathcal{A}^{(r)}$ and a set $B \subseteq [n]$, checks whether $|B| = k$. For Item 2, consider the algorithm that given a pair $(n,k) \in \mathcal{A}^{(r)}$ and a set $D \subseteq [n]$, checks whether $|D| \geq k$, and if so, returns an arbitrary subset of D of size k. Clearly, these algorithms can be implemented in time polynomial in n, hence $\mathcal{F}^{(r)}$ is polynomially computable. ◀

4.2 The Con-p-Division Problem

We present now the formal definition of the approximate $\textsc{Con}$-p-$\textsc{Division}$ problem for any prime p. The definition essentially extends the one given for $p = 2$ in [9, Appendix B, Definition 8]. One difference between the definitions is that we require a strict inequality in the condition (3) below. For this reason, we denote the problem by $\textsc{Con}$-p-$\textsc{Division}[< \varepsilon]$ and deviate from the notation ε-$\textsc{Con}$-p-$\textsc{Division}$ used in the literature. Note that we do not make any assumptions on the input valuation functions, and therefore allow solutions that demonstrate violations of their expected properties.

▶ **Definition 15** (The Con-p-Division Problem). *For an $\varepsilon \in (0,1]$, a prime p, and a fixed polynomial q, the* Con-p-Division$[< \varepsilon]$ *problem on normalized monotone functions is defined as follows. The input consists of*

- *a Lipschitz parameter $L \geq 0$ and*
- *m Turing machines $v_1, \dots, v_m$ that are supposed to compute L-Lipschitz-continuous normalized monotone valuation functions in $\mathcal{B}([0,1]) \to [0,1]$.*

The goal is to find either a partition of $[0,1]$ into p pieces $A_1, \dots, A_p$ using at most $(p-1) \cdot m$ cuts, such that

$$|v_i(A_t) - v_i(A_{t'})| < \varepsilon \qquad \text{for all } i \in [m] \text{ and } t, t' \in [p], \tag{3}$$

or a violation of some valuation function v_i, namely,

- *a violation of the normalization of v_i, or*
- *a violation of the running time of v_i, i.e., an input E on which v_i does not terminate within $q(|E| + |v_i|)$ steps, or*
- *a violation of the monotonicity of v_i, or*
- *a violation of the L-Lipschitz-continuity of v_i.*

When $p = 2$, we refer to the Con-2-Division$[< \varepsilon]$ *problem as* Con-Halving$[< \varepsilon]$.

Theorem 9 implies that for any $\varepsilon \in (0,1]$ and for every prime p, the Con-p-Division$[< \varepsilon]$ problem is total and thus lies in TFNP.

4.3 From Kneser to Con-Halving$[< 1]$

We consider now the case $p = 2$ and state the following result that provides a reduction from the general Kneser$(\mathcal{F}, m)$ problem with subset queries, where m is smaller by one than the bound given by Theorem 1, to the Con-Halving$[< \varepsilon]$ problem with $\varepsilon = 1$. As a consequence, we will obtain Theorem 2. The proof can be found in the full version of the paper.

▶ **Theorem 16.** *Let $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha \in \mathcal{A}}$ be a polynomially computable sequence of set families. Then, there exists a polynomial-time reduction from the* Kneser$(\mathcal{F}, \mathrm{cd}_2(\mathcal{F}_\alpha) - 1)$ *problem with subset queries to the* Con-Halving$[< 1]$ *problem on normalized monotone functions.*

We are ready to derive Theorem 2.

Proof of Theorem 2. By Lemma 14, the sequence $\mathcal{F}^{(2)}$ is polynomially computable. This allows us to apply Theorem 16, which yields that there exists a polynomial-time reduction from Kneser$(\mathcal{F}^{(2)}, \mathrm{cd}_2(\mathcal{F}^{(2)}_{(n,k)}) - 1)$ with subset queries to Con-Halving$[< 1]$ on normalized monotone functions. By Lemma 8, it holds that $\mathrm{cd}_2(\mathcal{F}^{(2)}_{(n,k)}) = n - 2k + 2$. It thus follows that the Kneser$(\mathcal{F}^{(2)}, \mathrm{cd}_2(\mathcal{F}^{(2)}_{(n,k)}) - 1)$ problem coincides with the Kneser problem, and we are done. ◀

4.4 From Kneser^p to Con-p-Division$[< \frac{1}{2}]$

We next state the following result that provides a reduction, for any prime $p \geq 3$, from the Kneser$^p(\mathcal{F}, m)$ problem with subset queries, where m is smaller than the bound given by Theorem 1, to the Con-p-Division$[< \varepsilon]$ problem with $\varepsilon = \frac{1}{2}$. As a consequence, we will obtain Theorem 3. The proof can be found in the full version of the paper.

▶ **Theorem 17.** *Let $p \geq 3$ be a prime, and let $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha \in \mathcal{A}}$ be a polynomially computable sequence of set families. Then, there exists a polynomial-time reduction from the* $\textsc{Kneser}^p(\mathcal{F}, \lfloor \frac{\mathrm{cd}_p(\mathcal{F}_\alpha)-1}{p-1} \rfloor)$ *problem with subset queries to the* $\textsc{Con-}p\textsc{-Division}[< \frac{1}{2}]$ *problem on normalized monotone functions.*

We are ready to derive Theorem 3.

Proof of Theorem 3. Fix a prime $p \geq 3$. By Lemma 14, the sequence $\mathcal{F}^{(p)}$ is polynomially computable, allowing us to apply Theorem 17, which yields that there exists a polynomial-time reduction from $\textsc{Kneser}^p(\mathcal{F}^{(p)}, \lfloor \frac{\mathrm{cd}_p(\mathcal{F}^{(p)}_{(n,k)})-1}{p-1} \rfloor)$ with subset queries to $\textsc{Con-}p\textsc{-Division}[< \frac{1}{2}]$ on normalized monotone functions. By Lemma 8, it holds that $\mathrm{cd}_p(\mathcal{F}^{(p)}_{(n,k)}) = n - p(k-1)$. It thus follows that the $\textsc{Kneser}^p(\mathcal{F}^{(p)}, \lfloor \frac{\mathrm{cd}_p(\mathcal{F}^{(p)}_{(n,k)})-1}{p-1} \rfloor)$ problem coincides with the $\textsc{Kneser}^p$ problem, and we are done. ◀

5 Kneser^p lies in PPA-p

This section is concerned with the membership of the $\textsc{Kneser}^p$ problem in the complexity class PPA-p for every prime p (see Theorem 4). The result is proved in two stronger forms through known connections between the chromatic number of Kneser hypergraphs and a $\mathbb{Z}_p$-variant of Tucker's lemma. We then establish limitations on the complexity of variants of the $\textsc{Kneser}^r$ problem, restricted to colorings with a bounded number of colors. We start by presenting the computational search problem associated with the $\mathbb{Z}_p$-Tucker lemma.

5.1 The $\mathbb{Z}_p$-Tucker Problem

The definition of the $\mathbb{Z}_p$-$\textsc{Tucker}$ problem requires a few notations. For a prime p, we denote the elements of the cyclic group $\mathbb{Z}_p$ of order p by ω^t for $t \in [p]$. A *signed set* over $\mathbb{Z}_p$ is a set whose elements are associated with signs from $\mathbb{Z}_p$. A signed subset of $[n]$ over $\mathbb{Z}_p$ can be represented by a vector $X \in (\mathbb{Z}_p \cup \{0\})^n$, where the subset consists of the elements $j \in [n]$ with $X_j \neq 0$, and the sign of every such j is X_j. For two signed sets $X, Y \in (\mathbb{Z}_p \cup \{0\})^n$, we denote by $X \preceq Y$ the fact that for every $j \in [n]$, if $X_j \neq 0$ then $X_j = Y_j$.

▶ **Definition 18.** *(The $\mathbb{Z}_p$-$\textsc{Tucker}$ Problem) For a prime p, the $\mathbb{Z}_p$-$\textsc{Tucker}$ problem is defined as follows. Its input consists of two integers n and s satisfying $s \leq \lfloor \frac{n-1}{p-1} \rfloor$ along with a Boolean circuit that represents a $\mathbb{Z}_p$-equivariant map $\lambda : (\mathbb{Z}_p \cup \{0\})^n \setminus \{0\}^n \to \mathbb{Z}_p \times [s]$, that is, a function that maps every nonzero $X \in (\mathbb{Z}_p \cup \{0\})^n$ to a pair $\lambda(X) = (\lambda_1(X), \lambda_2(X))$ in $\mathbb{Z}_p \times [s]$, where for each $t \in [p]$ it holds that $\lambda(\omega^t X) = (\omega^t \lambda_1(X), \lambda_2(X))$. The goal is to find a chain of p signed sets $X_1 \preceq X_2 \preceq \cdots \preceq X_p$ in $(\mathbb{Z}_p \cup \{0\})^n \setminus \{0\}^n$ that are assigned by λ the same absolute value with pairwise distinct signs, that is, for some permutation π of $[p]$ and some $\ell \in [s]$, it holds that $\lambda(X_t) = (\omega^{\pi(t)}, \ell)$ for all $t \in [p]$.*

Note that the assumption that the map λ is $\mathbb{Z}_p$-equivariant can be enforced syntactically. The existence of a solution for every instance of the $\mathbb{Z}_p$-$\textsc{Tucker}$ problem was proved by Ziegler [36]. Its membership in PPA-p, stated below, follows from a much more general result due to Filos-Ratsikas et al. [16, Theorem 5.2].

▶ **Theorem 19** ([16]). *For every prime p, the $\mathbb{Z}_p$-$\textsc{Tucker}$ problem lies in* PPA-p.

In order to obtain the membership of the $\textsc{Kneser}^p$ problem in PPA-p for general sequences of set families, we need the following definition.

▶ **Definition 20.** *For a set $\mathcal{A}$, let $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha \in \mathcal{A}}$ be a sequence of set families, where for each $\alpha \in \mathcal{A}$, $\mathcal{F}_\alpha$ is a family of non-empty subsets of $[n_\alpha]$ for some integer n_α. The sequence $\mathcal{F}$ is* strongly polynomially computable *if it is possible to associate with each family $\mathcal{F}_\alpha$ a linear order on its members, denoted $\leq$, such that there exist polynomials q_1, q_2, q_3 satisfying that*

1. *there exists an algorithm that given an element $\alpha \in \mathcal{A}$ runs in time $q_1(n_\alpha)$ and returns a Boolean circuit $C_1 : \{0,1\}^{2n_\alpha} \to \{0,1\}$ such that for every pair of sets $B_1, B_2 \in \mathcal{F}_\alpha$, it holds that $C_1(B_1, B_2) = 1$ if and only if $B_1 \leq B_2$,*
2. *there exists an algorithm that given an element $\alpha \in \mathcal{A}$ runs in time $q_2(n_\alpha)$ and returns a Boolean circuit $C_2 : \{0,1\}^{n_\alpha} \to \{0,1\}^{n_\alpha}$ such that for every set $D \subseteq [n_\alpha]$, $C_2(D)$ is the minimal subset of D, with respect to the order $\leq$, that belongs to the family $\mathcal{F}_\alpha$ if such a subset exists, and the empty set otherwise, and*
3. *for every prime p, there exists an algorithm that given an element $\alpha \in \mathcal{A}$ runs in time $q_3(n_\alpha)$ and returns the value of $\mathrm{cd}_p(\mathcal{F}_\alpha)$.*

The following lemma gives simple examples of strongly polynomially computable sequences.

▶ **Lemma 21.** *For every integer $r \geq 2$, the sequence $\mathcal{F}^{(r)}$ is strongly polynomially computable.*

Proof. Fix an integer $r \geq 2$. Recall that $\mathcal{F}^{(r)}_{(n,k)} = \binom{[n]}{k}$ for all $(n,k) \in \mathcal{A}^{(r)}$. We associate with the sets of $\mathcal{F}^{(r)}_{(n,k)}$ the linear order $\leq$, defined by $B_1 \leq B_2$ if $B_1 = B_2$ or the smallest element of $B_1 \triangle B_2$ belongs to B_1. Similarly to the proof of Lemma 14, it can be verified that $\mathcal{F}^{(r)}$ satisfies the conditions of Definition 20. Note that Item 3 follows from Lemma 8. ◀

5.2 From Kneser^p to $\mathbb{Z}_p$-Tucker

The following theorem asserts that the KNESER$^p(\mathcal{F}, m)$ problem is efficiently reducible to the $\mathbb{Z}_p$-TUCKER problem for every strongly polynomially computable sequence $\mathcal{F}$, whenever the number of colors m is smaller than the bound given by Theorem 1. Its proof verifies that a mathematical argument of Ziegler [36] can be transformed into an efficient reduction. We present it with the details for completeness in the full version of the paper.

▶ **Theorem 22.** *Let p be a prime, and let $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha \in \mathcal{A}}$ be a strongly polynomially computable sequence of set families. Then, KNESER$^p(\mathcal{F}, \lfloor \frac{\mathrm{cd}_p(\mathcal{F}_\alpha)-1}{p-1} \rfloor)$ is polynomial-time reducible to $\mathbb{Z}_p$-TUCKER.*

By combining Theorem 19 with Theorem 22, we derive the following corollary.

▶ **Corollary 23.** *Let p be a prime, and let $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha \in \mathcal{A}}$ be a strongly polynomially computable sequence of set families. Then, the KNESER$^p(\mathcal{F}, \lfloor \frac{\mathrm{cd}_p(\mathcal{F}_\alpha)-1}{p-1} \rfloor)$ problem lies in* PPA-p.

As a special case of Corollary 23, we derive Theorem 4.

Proof of Theorem 4. Fix a prime p. By Lemma 21, the sequence $\mathcal{F}^{(p)}$ is strongly polynomially computable, hence by Corollary 23, the KNESER$^p(\mathcal{F}^{(p)}, \lfloor \frac{\mathrm{cd}_p(\mathcal{F}^{(p)}_{(n,k)})-1}{p-1} \rfloor)$ problem lies in PPA-p. By Lemma 8, it holds that $\mathrm{cd}_p(\mathcal{F}^{(p)}_{(n,k)}) = n - p(k-1)$, hence the latter problem coincides with the KNESER^p problem, and we are done. ◀

We next consider the KNESER$^p_{\widetilde{\mathrm{stab}}}$ problem associated with the hypergraph $K^p(n,k)_{\widetilde{\mathrm{stab}}}$, that is, the sub-hypergraph of $K^p(n,k)$ induced by the almost stable k-subsets of $[n]$. Corollary 23 can be applied to obtain a membership result in PPA-p for this setting, however, it does not give the largest possible number of colors. To obtain the result with an optimal number of colors, which is smaller than the chromatic number only by one, we apply a modified argument of Meunier [32], verifying that it provides an efficient reduction. Let us mention, though, that the proof of [32] applies a slightly different variant of the $\mathbb{Z}_p$-Tucker lemma, whose proof relies on Dold's theorem [11]. Our proof, which can be found in the full version of the paper, uses the version of the lemma that corresponds to our definition of the $\mathbb{Z}_p$-Tucker problem, which lies in PPA-p. See [16, Remark 1] for a discussion on the computational aspects of Dold's theorem.

▶ **Theorem 24.** *For every prime p, $\textsc{Kneser}^p_{\widetilde{\text{stab}}}$ is polynomial-time reducible to $\mathbb{Z}_p$-Tucker.*

By combining Theorem 19 with Theorem 24, the proof of Theorem 5 is completed.

5.3 Limitations on the Complexity of Kneser^r Problems

We next show that the results from the previous section imply limitations on the complexity of variants of the Kneser^r problem, restricted to colorings with a bounded number of colors. We start with the following simple lemma, which says that the Kneser$^r(\mathcal{F}, m)$ problem does not become easier when r increases.

▶ **Lemma 25.** *Let $r_1 \leq r_2$ be integers, let $\mathcal{F} = (\mathcal{F}_\alpha)_{\alpha\in\mathcal{A}}$ be a sequence of set families, and let $m : \mathcal{A} \to \mathbb{N}$ be a function such that $m(\alpha) < \chi(K^{r_2}(\mathcal{F}_\alpha))$ for all $\alpha \in \mathcal{A}$. Then, Kneser$^{r_1}(\mathcal{F}, m)$ is polynomial-time reducible to Kneser$^{r_2}(\mathcal{F}, m)$.*

Proof. Fix two integers $r_1 \leq r_2$. Consider an instance of the Kneser$^{r_1}(\mathcal{F}, m)$ problem, that is, an element $\alpha \in \mathcal{A}$ and a Boolean circuit that represents a coloring of $\mathcal{F}_\alpha$ with $m(\alpha)$ colors. We simply apply the identity reduction to the Kneser$^{r_2}(\mathcal{F}, m)$ problem. A solution for the obtained instance of Kneser$^{r_2}(\mathcal{F}, m)$, whose existence follows by $m(\alpha) < \chi(K^{r_2}(\mathcal{F}_\alpha))$, is a collection of r_2 pairwise disjoint sets of $\mathcal{F}_\alpha$ with the same color. By $r_1 \leq r_2$, any r_1 sets from this collection form a solution for the same input as an instance of Kneser$^{r_1}(\mathcal{F}, m)$. The correctness of the reduction follows. ◀

By combining Lemma 25 with Theorem 4, we obtain the following result.

▶ **Theorem 26.** *For every integer r and for every prime p such that $r \leq p$,*

$$\textsc{Kneser}^r(\mathcal{F}^{(p)}, \lfloor \tfrac{n-p(k-1)-1}{p-1} \rfloor)$$

lies in PPA-p.

Proof. Fix an integer r and a prime p such that $r \leq p$. Put $m(n,k) = \lfloor \frac{n-p(k-1)-1}{p-1} \rfloor$, and apply Lemma 25 to obtain that Kneser$^r(\mathcal{F}^{(p)}, m)$ is polynomial-time reducible to Kneser$^p(\mathcal{F}^{(p)}, m)$. The latter coincides with the Kneser^p problem, which, by Theorem 4, lies in PPA-p. It thus follows that Kneser$^r(\mathcal{F}^{(p)}, m)$ lies in PPA-p, as required. ◀

Theorem 26 yields, for any integer r, a limitation on the complexity of the Kneser^r problem, restricted to colorings with a bounded number of colors. For example, consider the Kneser problem, which asks to find a monochromatic edge in a graph $K(n, k)$ colored with $n - 2k + 1$ colors, and recall that it lies in PPA. By Theorem 26, applied with $r = 2$ and $p = 3$, the Kneser$(n, k, \lfloor \frac{n-3k+2}{2} \rfloor)$ problem, which asks to find a monochromatic edge in a graph $K(n, k)$ colored with only $\lfloor \frac{n-3k+2}{2} \rfloor$ colors, lies in PPA-3. This implies that the latter problem is not PPA-hard, unless PPA $\subseteq$ PPA-3. We next present analogue consequences for the Schrijver problem.

We need the following simple lemma, whose proof is given in the full version of the paper.

▶ **Lemma 27.** *For every integer r, $\textsc{Kneser}^r_{\mathrm{stab}}(n,k,\lfloor\frac{n-rk}{r-1}\rfloor)$ is polynomial-time reducible to $\textsc{Kneser}^r_{\widetilde{\mathrm{stab}}}$.*

By combining Lemma 27 with Theorem 5, we derive the following.

▶ **Corollary 28.** *For every prime p, the $\textsc{Kneser}^p_{\mathrm{stab}}(n,k,\lfloor\frac{n-pk}{p-1}\rfloor)$ problem lies in* PPA-*p.*

We further derive the following corollary.

▶ **Corollary 29.** *The $\textsc{Schrijver}(n,k,\lfloor\frac{n-3k}{2}\rfloor)$ problem lies in* PPA-3*.*

Proof. Put $m(n,k)=\lfloor\frac{n-3k}{2}\rfloor$. By Lemma 25, the $\textsc{Schrijver}(n,k,m)$ problem, which can be written as $\textsc{Kneser}_{\mathrm{stab}}(n,k,m)$, is polynomial-time reducible to $\textsc{Kneser}^3_{\mathrm{stab}}(n,k,m)$. By Corollary 28, the latter lies in PPA-3. It thus follows that $\textsc{Schrijver}(n,k,m)$ lies in PPA-3 as well. ◀

We finally state the following consequence of Corollary 29 regarding the $\textsc{Schrijver}(n,k,m)$ problem with the function $m(n,k)=\lfloor n/2\rfloor-2k+1$ considered in [24]. We use here the fact that for all integers n and $k\geq 2$, it holds that $\lfloor n/2\rfloor-2k+1\leq\lfloor\frac{n-3k}{2}\rfloor$.

▶ **Corollary 30.** *The $\textsc{Schrijver}(n,k,\lfloor n/2\rfloor-2k+1)$ problem lies in* PPA-3*.*

References

1 Meysam Alishahi and Hossein Hajiabolhassan. On the chromatic number of general Kneser hypergraphs. *J. Comb. Theory, Ser. B*, 115:186–209, 2015.
2 Noga Alon. Splitting necklaces. *Adv. Math.*, 63(3):247–253, 1987.
3 Noga Alon, Peter Frankl, and László Lovász. The chromatic number of Kneser hypergraphs. *Trans. Amer. Math. Soc.*, 298(1):359–370, 1986.
4 Noga Alon and Andrei Graur. Efficient splitting of necklaces. In *Proc. of the 48th International Colloquium on Automata, Languages, and Programming, (ICALP'21)*, pages 14:1–14:17, 2021.
5 Jai Aslam, Shuli Chen, Ethan Coldren, Florian Frick, and Linus Setiabrata. On the generalized Erdös-Kneser conjecture: Proofs and reductions. *J. Comb. Theory, Ser. B*, 135:227–237, 2019.
6 Imre Bárány, Senya B. Shlosman, and András Szücs. On a topological generalization of a theorem of Tverberg. *J. London Math. Soc.*, s2-23(1):158–164, 1981.
7 Karol Borsuk. Drei Sätze über die n-dimensionale euklidische Sphäre. *Fundamenta Mathematicae*, 20(1):177–190, 1933.
8 Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. Constant inapproximability for PPA. In *Proc. of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC'22)*, pages 1010–1023, 2022.
9 Argyrios Deligkas, Aris Filos-Ratsikas, and Alexandros Hollender. Two's company, three's a crowd: Consensus-halving for a constant number of agents. *Artif. Intell.*, 313(C):103784, 2022. Preliminary version in EC'21.
10 Xiaotie Deng, Zhe Feng, and Rucha Kulkarni. Octahedral Tucker is PPA-complete. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:118, 2017.
11 Albrecht Dold. Simple proofs of some Borsuk–Ulam results. *Contemp. Math.*, 19:65–69, 1983.
12 Vladimir Dol'nikov. Transversals of families of sets. In *Studies in the theory of functions of several real variables (Russian)*, volume 29, pages 30–36. Yaroslavl', 1981.
13 Paul Erdös. Problems and results in combinatorial analysis. *Colloq. int. Teorie comb., Roma 1973*, pages 3–17, 1976.
14 Aris Filos-Ratsikas and Paul W. Goldberg. Consensus halving is PPA-complete. In *Proc. of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC'18)*, pages 51–64, 2018.

15 Aris Filos-Ratsikas and Paul W. Goldberg. The complexity of splitting necklaces and bisecting ham sandwiches. In *Proc. of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC'19)*, pages 638–649, 2019.

16 Aris Filos-Ratsikas, Alexandros Hollender, Katerina Sotiraki, and Manolis Zampetakis. A topological characterization of modulo-p arguments and implications for necklace splitting. In *Proc. of the 32nd ACM-SIAM Symposium on Discrete Algorithms (SODA'21)*, pages 2615–2634, 2021.

17 Aris Filos-Ratsikas, Alexandros Hollender, Katerina Sotiraki, and Manolis Zampetakis. Consensus-halving: Does it ever get easier? *SIAM J. Comput.*, 52(2):412–451, 2023. Preliminary version in EC'20.

18 Florian Frick. Chromatic numbers of stable Kneser hypergraphs via topological Tverberg-type theorems. *Int. Math. Res. Not.*, 2020(13):4037–4061, 2020.

19 Paul Goldberg and Jiawei Li. Consensus division in an arbitrary ratio. In *Proc. of the 14th Innovations in Theoretical Computer Science Conference (ITCS'23)*, pages 57:1–57:18, 2023.

20 Paul W. Goldberg, Alexandros Hollender, Ayumi Igarashi, Pasin Manurangsi, and Warut Suksompong. Consensus halving for sets of items. *Math. Oper. Res.*, 47(4):2547–3399, 2022. Preliminary version in WINE'20.

21 Mika Göös, Pritish Kamath, Katerina Sotiraki, and Manolis Zampetakis. On the complexity of modulo-q arguments and the Chevalley–Warning theorem. In *Proc. of the 35th Computational Complexity Conference (CCC'20)*, pages 19:1–19:42, 2020.

22 Ishay Haviv. The complexity of finding fair independent sets in cycles. *Comput. Complex.*, 31(2):14, 2022. Preliminary version in ITCS'21.

23 Ishay Haviv. A fixed-parameter algorithm for the Kneser problem. In *Proc. of the 49th International Colloquium on Automata, Languages, and Programming (ICALP'22)*, pages 72:1–72:18, 2022.

24 Ishay Haviv. On finding constrained independent sets in cycles. In *Proc. of the 50th International Colloquium on Automata, Languages, and Programming (ICALP'23)*, pages 73:1–73:16, 2023.

25 Charles R. Hobby and John R. Rice. A moment problem in L_1 approximation. *Proc. Amer. Math. Soc.*, 16(4):665–670, 1965.

26 Alexandros Hollender. The classes PPA-k: Existence from arguments modulo k. *Theor. Comput. Sci.*, 885:15–29, 2021. Preliminary version in WINE'19.

27 Martin Kneser. Aufgabe 360. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 58(2):27, 1955.

28 Igor Kriz. Equivariant cohomology and lower bounds for chromatic numbers. *Trans. Amer. Math. Soc.*, 333(2):567–577, 1992.

29 László Lovász. Kneser's conjecture, chromatic number, and homotopy. *J. Comb. Theory, Ser. A*, 25(3):319–324, 1978.

30 Jiří Matoušek. A combinatorial proof of Kneser's conjecture. *Combinatorica*, 24(1):163–170, 2004.

31 Nimrod Megiddo and Christos H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theor. Comput. Sci.*, 81(2):317–324, 1991.

32 Frédéric Meunier. The chromatic number of almost stable Kneser hypergraphs. *J. Comb. Theory, Ser. A*, 118(6):1820–1828, 2011.

33 Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.

34 Alexander Schrijver. Vertex-critical subgraphs of Kneser graphs. *Nieuw Arch. Wiskd.*, 26(3):454–461, 1978.

35 Forest W. Simmons and Francis Edward Su. Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Math. Soc. Sci.*, 45(1):15–25, 2003.

36 Günter M. Ziegler. Generalized Kneser coloring theorems with combinatorial proofs. *Invent. Math.*, 147(3):671–691, 2002.