

Consensus Division in an Arbitrary Ratio

Paul Goldberg   

University of Oxford, UK

Jiawei Li  

The University of Texas at Austin, TX, USA

Abstract

We consider the problem of partitioning a line segment into two subsets, so that n finite measures all have the same ratio of values for the subsets. Letting $\alpha \in [0, 1]$ denote the desired ratio, this generalises the PPA-complete *consensus-halving* problem, in which $\alpha = \frac{1}{2}$. Stromquist and Woodall [30] showed that for any α , there exists a solution using $2n$ cuts of the segment. They also showed that if α is irrational, that upper bound is almost optimal. In this work, we elaborate the bounds for rational values α . For $\alpha = \frac{\ell}{k}$, we show a lower bound of $\frac{k-1}{k} \cdot 2n - O(1)$ cuts; we also obtain almost matching upper bounds for a large subset of rational α .

On the computational side, we explore its dependence on the number of cuts available. More specifically,

1. when using the minimal number of cuts for each instance is required, the problem is NP-hard for any α ;
2. for a large subset of rational $\alpha = \frac{\ell}{k}$, when $\frac{k-1}{k} \cdot 2n$ cuts are available, the problem is in PPA- k under Turing reduction;
3. when $2n$ cuts are allowed, the problem belongs to PPA for any α ; more generally, the problem belong to PPA- p for any prime p if $2(p-1) \cdot \frac{\lfloor p/2 \rfloor}{p} \cdot n$ cuts are available.

2012 ACM Subject Classification Theory of computation \rightarrow Problems, reductions and completeness

Keywords and phrases Consensus Halving, TFNP, PPA- k , Necklace Splitting

Digital Object Identifier 10.4230/LIPIcs.ITCS.2023.57

Funding *Paul Goldberg*: Currently supported by a J.P. Morgan Faculty Research Award.

Jiawei Li: Supported by Scott Aaronson’s Simons award “It from Qubit”.

Acknowledgements We thank Alexandros Hollender for pointing out the paper [30] to us. We thank the anonymous ITCS reviewers for their helpful comments. Jiawei Li wants to thank Xiaotie Deng for introducing the consensus-halving problem to him.

1 Introduction

The complexity class TFNP (standing for total functions, computable in nondeterministic polynomial time), refers to problems of computing a solution that is guaranteed to exist, and once found can be easily checked for correctness. Such problems are of particular interest when they appear to be computationally hard, due to the fact that they cannot be NP-hard unless $\text{NP} = \text{coNP}$ [24]. Due to this point and the fact that TFNP does not seem to have complete problems, hard problems in TFNP have been classified via certain syntactic subclasses corresponding to the combinatorial existence principles that guarantee the existence of solutions, without indicating an efficient algorithm for their construction. They include the well-known classes PPAD, PPA, and PPP, introduced by Papadimitriou [25] in 1994. PPAD represents the complexity of Nash equilibrium computation and related problems, and more recently, PPA has been shown to capture the complexity of certain problems of consensus division, discussed in more detail below. Papadimitriou [25] also pointed out a collection of classes PPA- k (where $k \geq 2$ is a natural number). PPA- k (Definition 4 below) consists of problems where the existence guarantee of solutions is due to a modulo- k



© Paul Goldberg and Jiawei Li;

licensed under Creative Commons License CC-BY 4.0

14th Innovations in Theoretical Computer Science Conference (ITCS 2023).

Editor: Yael Tauman Kalai; Article No. 57; pp. 57:1–57:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

counting argument, and these classes turn out to be relevant for the problem we study here. Note that PPA is the same as PPA-2: PPA stands for “polynomial parity argument”, i.e. a modulo-2 counting argument.

We study a problem arising from a 1985 result of Stromquist and Woodall [30] saying roughly, that if we have n valuation measures of an interval A , and we want to divide A into two shares whose values – with respect to each of the n measures – are in some given ratio $\alpha : 1 - \alpha$ (for $\alpha \in (0, 1)$), then it’s possible to divide up the interval using at most $2n$ cut points, so that the pieces can indeed be allocated to the two shares to get the desired outcome. This is what we mean by consensus division: if each valuation measure comes from a separate agent, then the objective is to split up a resource (the interval) in such a way that all agents agree on the values of the partition. We call this problem *imbalanced-consensus-division*. What is the complexity of computing such a partition of the interval? This problem generalizes *consensus-halving* (discussed in more detail in Section 1.2), the special case when $\alpha = 1/2$. In this case of $\alpha = 1/2$, just n cuts suffice to find a solution, and the problem of computing a suitable set of n cuts was recently shown to be PPA-complete [16, 17, 12, 10] (the latter paper shows PPA-hardness even for additive-constant approximation). That gives the consensus-halving problem a novel complexity-theoretic status and also indicates that it is highly unlikely to be solvable in polynomial time.

A closely related problem is CONSENSUS- k -DIVISION, which also generalizes *consensus-halving* ($k = 2$) and has been studied in recent works [19, 18] due to its connection to PPA- k . CONSENSUS- k -DIVISION is the problem of splitting the interval into $k \geq 2$ shares all of the equal value with respect to all measures. We study the connection between CONSENSUS- k -DIVISION and our model. In one direction, several results in our paper are based on existing results of CONSENSUS- k -DIVISION (Theorems 10, 11); in the other direction, we show a novel complexity result for CONSENSUS- k -DIVISION itself as a by-product of studying *imbalanced-consensus-division* (Corollary 15).

We also define *imbalanced-necklace-splitting* as a generalization of the *necklace-splitting* problem [3, 1], which could be taken as a discrete version of *consensus-halving*. The equivalence between *necklace-splitting* and *consensus-halving* (with inverse-polynomial approximation) [16] can be generalized to our *imbalanced* variant (Theorem 3). We focus on *imbalanced-consensus-division* in this paper, while all our results can be extended to *imbalanced-necklace-splitting* via Theorem 3.


1.1 Our Contribution

On the combinatorial (as opposed to computational) side, we obtain more detailed bounds on the number of cuts that may be needed in the worst case for *rational* ratio α . Recall that Stromquist and Woodall [30] showed a general upper bound of $2n$ cuts for any ratio, and proved the tightness of this bound for all irrational α . A series of instances are constructed in Section 3 which provides the lower bound.

► **Informal Theorem 1.** *For any rational ratio $\alpha = \ell/k \in (0, 1)$, roughly $\frac{2(k-1)}{k} \cdot n$ cuts are needed in the worst case.*

A visualization of the lower bound is in Figure 1. We believe our lower bound for any rational ratio is tight.

¹ Also called the Riemann Function, see https://en.wikipedia.org/wiki/Thomae%27s_function



■ **Figure 1** Lower bounds for all rationals with denominators smaller than 50. Constant terms are ignored. This plot is identical to turning Thomae’s function upside down¹.

In Section 4, we generalize the existence proof of [30] and build a reduction to the CONSENSUS- k -DIVISION problem for some specific choices of α and k . We, therefore, improve the previous general upper bound of $2n$ for any rational α .

► **Informal Theorem 2.** *For a large subset Q^* of rationals $\alpha = \ell/k$ (formally specified in Section 4), $\frac{2(k-1)}{k} \cdot n$ cuts are enough in the worst case; for any other rational ratios, there is an upper bound strictly smaller than $2n$ by a gap linear in n .*

Notice that the bound for any rationals in set Q^* is almost tight.

► **Remark.** Stromquist and Woodall [30] also showed a lower bound of $2n - 2$ cuts for some rational α and n . We clarify that their lower bounds for any rational α only work for constant-size n (depending on α). In our paper, we treat ratio α as a fixed constant and all the bounds are asymptotic in n . Our upper bound shows that for any rational ratio α , it is impossible to require $2n - 2$ cuts in the worst case for arbitrarily large n .

On the computational side, the reduction to CONSENSUS- k -DIVISION also reveals an interesting connection with the complexity classes PPA- k . The most commonly studied setting is that the minimum number of cuts (as a function of n and α) is given to make the solution always exists, i.e., to make the problem a total problem. We study the complexity in this setting for rationals in set Q^* , since we know the tight bound for them.

► **Informal Theorem 3.** *For a large subset Q^* of rationals $\alpha = \ell/k$, finding a solution using $\frac{2(k-1)}{k} \cdot n$ cuts lies in PPA- k under Turing reductions. In particular, if $k = p^r$ for a prime p , the problem lies in PPA- p .*

When more cuts are allowed, the problem should become easier. Here we show that as more cuts are allowed, the problem is contained in more and more of the complexity classes PPA- p .

► **Informal Theorem 4.** *For any $\alpha \in (0, 1)$ and any prime p , finding an inverse-polynomial approximate solution using $2(p-1) \cdot \frac{\lceil p/2 \rceil}{\lfloor p/2 \rfloor} \cdot n$ cuts is in PPA- p .*

For example, solving consensus-halving ($\alpha = 1/2$) with $8n$ cuts is in PPA-3. This fills the blank of previous results in an intriguing way: on the one side, consensus-halving (with inverse-polynomial approximation) remains PPA-complete for $n + n^{1-\delta}$ cuts for any

small constant δ [15, 18]; on the other side, Alon and Graur [2] showed that finding an inverse-polynomial approximate solution with $O(n \log n)$ cuts is in P. Also, these are the first natural² problems in the intersection of multiple PPA- p classes.

Our results also yield further PPA- p containment for CONSENSUS- k -DIVISION as extra cuts are allowed.

► **Informal Theorem 5.** *For any prime p , solving CONSENSUS- k -DIVISION with $2(k-1) \cdot (p-1) \cdot \frac{\lceil p/2 \rceil}{\lfloor p/2 \rfloor} \cdot n$ cuts with inverse-polynomial error lies in PPA- p .*

In Section 5, we study the hardest setting, i.e., a solution with a minimal number of cuts for each instance is required. Notice that this problem may not be a TFNP problem anymore since there is no easy way to verify whether a solution does use a minimal number of cuts.

► **Informal Theorem 6.** *For any $\alpha \in (0, 1)$, finding a solution using the minimum number of cuts is NP-hard.*

1.2 Background, Related Work

We have mentioned consensus-halving as an important special case of the problem of interest in the present paper. Consensus-halving is the computational analog of the Hobby-Rice theorem [22]. It was shown to be PPA-complete in [16]; this result was subsequently strengthened to apply to the (two-thief) necklace-splitting problem [17]. (Necklace-splitting is a discretized version of consensus-halving, and the extension to necklace-splitting required PPA-hardness for inverse-polynomial additive error in the values of the two shares). This line of work on PPA-completeness also highlights the close connection of the problem with the ham-sandwich theorem from topology, as well as the Borsuk-Ulam theorem. Further work has extended this to showing PPA-completeness even for simple measures (unions of uniform distributions over just two sub-intervals). But at present, little is known about how much it helps if we allow ourselves more than n cuts. CONSENSUS- k -DIVISION is another natural generalization of consensus-halving, consisting of consensus division into $k \geq 2$ shares, all of equal value. Alon [1] identified the number of cuts needed to achieve this (namely $(k-1)n$), in the context of the necklace-splitting problem, and its computational complexity is recently studied in [19], in which context the classes PPA- k are also important.

Goldberg et al. [20] and Segal-Halev [28] study versions of consensus-division where the “cake” being partitioned is not a line segment, but an unordered collection of items on which the agents have diverse valuations. Deligkas, Filos-Ratsikas, and Hollender [13] study consensus-halving with a constant number of agents and more general valuation functions. The complexity of computing the exact solution of consensus-halving is considered in Deligkas et al. [11] and Batziou, Hansen, and Høgh [6].

Most of the literature on cake-cutting is about the search for a *fair division* into pieces that get allocated to the agents (such that, for example, no agent values someone else’s pieces more than his own), as opposed to consensus division, as considered here. In the context of fair division, there is an “arbitrary proportions” analog to the problem studied in this paper: Segal-Halevi [27] and Crew et al. [9] have studied an analogous generalization of fair division in which each agent has a (non-negative fractional) claim on the cake, all claims summing to 1. In common with consensus division, it is found that in the more general case of unequal proportions, more cuts may be required than in the special case of equal proportions. Segal-Halevi shows via a simple construction that $2n - 2$ cuts may be required (when proportions are equal, it is known that $n - 1$ are sufficient).

² By natural we mean there is no explicit circuit in the input of the problem.

There is a similar recent interest in considering unequal (or weighted) sharing, in the context of indivisible items. Here, envy-freeness is unachievable in general, and instead one considers envy-freeness subject to being able to remove a small number of items. This work addresses the performance of standard algorithms such as round-robin picking sequences, and whether various desiderata can be satisfied [5, 7, 8, 4].

2 Preliminaries

In this section, we give detailed definitions of the problem and complexity classes of interest here.

The CONSENSUS-HALVING problem involves a set of n agents each of whom has a valuation function on a 1-dimensional line segment A (here we set A to be the unit interval $[0, 1]$). Consider the problem of selecting k “cut points” in A that partition A into $k + 1$ pieces, then label each piece either “positive” or “negative” in such a way that each agent values the positive pieces equally to the negative ones. In 2003, Simmons and Su [29] showed that this can always be done for $k = n$; their proof applies the Borsuk-Ulam theorem and is a proof of existence analogous to Nash’s famous existence proof of equilibrium points of games, proved using Brouwer’s or Kakutani’s fixed point theorem. Significantly, Borsuk-Ulam is the *undirected* version of Brouwer, and already from [25] we know that it relates to PPA. The CONSENSUS-HALVING problem was shown to be PPA-complete in [16]. As detailed in Definition 1, we assume that valuations are presented as step functions using the logarithmic cost model of numbers.

Consensus-splitting in an arbitrary ratio $\alpha : 1 - \alpha$ is one of the open problems raised in [29]. Clearly, a single cut suffices for $n = 1$, and for $n = 2$ it remains the case that 2 cuts suffice (to see this, assume by rescaling so that agent 1’s distribution is uniform, and consider sliding an interval of length α along the unit interval, keeping track of agent 2’s value for the interval). For $n > 2$, things get more complicated. For any real number $\alpha \in [0, 1]$, we define the following variant of the CONSENSUS-HALVING problem.

► **Definition 1** (α -IMBALANCED-CONSENSUS-DIVISION, α -ICD). *Input:* $\varepsilon > 0$ and n continuous probability measures μ_1, \dots, μ_n on $[0, 1]$, representing the valuation function of each agent. We assume the probability measures are presented as piecewise constant functions on $[0, 1]$, i.e., step functions (explicitly given in the input).

Output: A partition of the unit interval into two (not necessarily connected) subsets A_+ and A_- , such that for any $i \in [n]$, we have $|\mu_i(A_+) - \alpha| \leq \varepsilon$.

It is easy to see that α -ICD is equivalent to $(1 - \alpha)$ -ICD. Definition 1 is not quite complete, since we care about the number of cuts needed to make the partition, which in general is dependent on α . The most commonly studied setting is when the number of cuts allowed, which is a function of n and α in our case, is minimal to make the problem a total problem, i.e., the solution always exists. Without further specification, we solve α -ICD in this setting. We also study the cases where more or fewer cuts are available in this paper, in which the number of cuts will be specified explicitly.

As noted in [31], piecewise-constant functions have been used in various previous works on cake-cutting and can approximate natural real-valued functions. Another advantage of piecewise-constant functions is that by the same argument as Theorem 5.2 of [14], an exact solution ($\varepsilon = 0$) of α -ICD with rational α could be efficiently calculated from an approximated solution with inverse-exponential ε . Most of our results could be extended to additive valuation function under inverse-polynomial approximation.

In this paper, we describe each piecewise-constant function by a set of *value blocks*. Each value block represents an interval on which the valuation function takes a constant value. Naturally, the total weights of value blocks add up to one for each agent.

If a subinterval of the partition belongs to A_+ or A_- , we say it has the label “+” or “−” respectively. We assume the label of each subinterval alternates after each cut without loss of generality (two consecutive subintervals having the same label can be merged).

We similarly define α -IMBALANCED-NECKLACE-SPLITTING (α -INS) problem for any rational α .

► **Definition 2** (α -IMBALANCED-NECKLACE-SPLITTING, α -INS). **Input:** An open necklace with t beads, each of which has one of n colors. There are a_i beads of color i , where $a_i, \alpha \cdot a_i \in \mathbb{N}$ for any $i \in [n]$.

Output: A partition of the necklace into two (not necessarily connected) pieces A_+ and A_- , such that for any $i \in [n]$, piece A_+ contains exactly $\alpha \cdot a_i$ beads of color i .

► **Theorem 3** (Essentially from Section 6 of [16]). For any rational α , there is a many-to-one reduction from α -INS to α -ICD, and vice versa. Moreover, the reductions in both directions preserve the number of agents/colors and the number of cuts.

The complexity classes PPA- k are defined as follows [25, 23]. For any integer $k \geq 2$, PPA- k is the set of problems reducible in polynomial time to the problem BIPARTITE-MOD- k :

► **Definition 4** (the problem BIPARTITE-MOD- k). We are given a bipartite graph on the vertices $(0 \times \{0, 1\}^n, 1 \times \{0, 1\}^n)$ represented concisely via a circuit C , that given as input a vertex in $0 \times \{0, 1\}^n$, outputs a set of $\leq k$ potential neighbours in $1 \times \{0, 1\}^n$, and vice versa. An edge (u, v) is present provided that v is one of the potential neighbors of u , and vice versa. Suppose that the number of neighbours of 0^{n+1} lies in $\{1, 2, \dots, k-1\}$. A solution consists of some other vertex having a degree in $\{1, 2, \dots, k-1\}$.

The existence of at least one solution to any instance of BIPARTITE-MOD- k follows from a modulo- k counting argument. In the case of $k = 2$ we have complexity class PPA, in which the corresponding problem is called LEAF, consisting of a concisely-represented undirected graph of degree ≤ 2 , in which the all-zeroes vector is a leaf (a degree-1 vertex), and the problem is the find another leaf of the graph. PPA- p is closed under Turing reductions for any prime p , while PPA- k for general k (except for prime or power of prime) is believed to be not closed under Turing reductions [23, 21].

3 Lower Bound

We assume that all fractions discussed in this paper are written down as rational numbers ℓ/k , where ℓ, k are coprime and $\ell < k$.

► **Theorem 5** (Lower Bound). For any rational number $\alpha = \ell/k \in [0, 1]$, $\frac{2(k-1)}{k} \cdot n - O(1)$ cuts are necessary (in the worst case) for an exact solution of α -ICD with n agents.

► **Remark.** The $O(1)$ term in the lower bound is bounded by k , the denominator of α .

Before introducing the construction of α -ICD instances establishing this lower bound, Lemmas (6,7) identify two simple properties of fractions. Their proofs are left in Appendix A.

Let $\alpha_1 := \ell_1/k_1$, $\alpha_2 := \ell_2/k_2$, $\alpha_1 < \alpha_2$ be two fractions. We say α_1 and α_2 are *adjacent* if $\ell_2 \cdot k_1 - \ell_1 \cdot k_2 = 1$.

► **Lemma 6.** Let $\alpha_1 := \ell_1/k_1$, $\alpha_2 := \ell_2/k_2$, $\alpha_1 < \alpha_2$ be two adjacent fractions. For any fraction $\alpha := \ell/k$, $\alpha \in (\alpha_1, \alpha_2)$, we have $k \geq k_1 + k_2$.

► Remark. It's easy to verify that by taking $\alpha := (\ell_1 + \ell_2)/(k_1 + k_2)$, we have $\alpha \in (\alpha_1, \alpha_2)$, and α is adjacent to both α_1, α_2 .

► Lemma 7. Given a fraction $\alpha := \ell/k, k \geq 2$, let (α_1, α_2) be the smallest interval satisfying that $\alpha \in (\alpha_1, \alpha_2)$ and $k_1, k_2 < k$, where $\alpha_1 := \ell_1/k_1, \alpha_2 := \ell_2/k_2$. Then

1. α_1 and α_2 are adjacent;
2. $\ell = \ell_1 + \ell_2, k = k_1 + k_2$;
3. α is adjacent to both α_1, α_2 .

► Corollary 8. Define set $\mathbb{Q}_k := \{a/b \in [0, 1] : a, b \text{ coprime}, b \leq k\}$, i.e., all the fractions in $[0, 1]$ with denominator smaller than or equal to k .

When elements in \mathbb{Q}_k are listed in increasing order, every two consecutive elements are adjacent.

Now we are ready to introduce the construction of the α -ICD instances matching the bound.

Proof of Theorem 5. We first fix a rational number $\alpha := \ell/k$. Let (α_1, α_2) be the smallest interval satisfying that $\alpha \in (\alpha_1, \alpha_2)$ and $k_1, k_2 < k$, where $\alpha_1 := \ell_1/k_1, \alpha_2 := \ell_2/k_2$. (So, Lemma 7 applies.)


We construct an α -ICD instance with $k_1 + k_2$ agents, partitioned into two types:

Type-1 There are k_1 type-1 agents. A type-1 agent has k_2 value blocks with $1/k_2$ weight each;

Type-2 There are k_2 type-2 agents. A type-2 agent has k_1 value blocks with $1/k_1$ weight each.

There are in total $2k_1 \cdot k_2$ value blocks, half of which belong to type-1 agents and the other half to type-2 agents. The arrangement of these value blocks is specified by the following rules, see Figure 2 for a visualization:

- all the value blocks are disjoint;
- all the value blocks from the i -th type-1 (respectively, type-2) agent are on the left of any value blocks from the $(i + 1)$ -th type-1 (respectively, type-2) agent;
- from left to right, the value blocks from type-1 and type-2 agents occur in turn; the first block comes from the first type-2 agent.



■ **Figure 2** The arrangement of value blocks in a $(2/5)$ -ICD instance. There are three type-1 agents and two type-2 agents. Each color corresponds to a different agent. Note that $\alpha_1 = 1/3 < 2/5 < 1/2 = \alpha_2$.

To lower-bound the number of cuts needed for this instance, we start with the following observations, recalling that A_+ is supposed to have measure α :

► **Observation 9.**

1. For each type-1 agent, at most $\ell_2 - 1$ of its value blocks are entirely included in the set A_+ . Otherwise, the set A_+ will contain at least $\ell_2/k_2 = \alpha_2$ fraction of measure for that agent, which is strictly greater than α .

2. For each type-2 agent, at least $\ell_1 + 1$ of its value blocks have positive measure in the set A_+ . Otherwise, the set A_+ will contain at most $\ell_1/k_1 = \alpha_1$ fraction of measure for that agent, which is strictly smaller than α .

Now assuming in a valid solution, there are t intervals labeled with “+”, denoted by I_1, \dots, I_t . For each $i \in [t]$ define a_i to be the number of value blocks from type-2 agents that are intersected with the i -th interval I_i . Then from our construction, it follows that there are at least $a_i - 1$ of value blocks from type-1 agents that are *entirely* included in the interval I_i .

We have the following constraints regarding the observation above:

$$\sum_{i=1}^t \max\{a_i - 1, 0\} \leq (\ell_2 - 1) \cdot k_1; \quad (1)$$

$$\sum_{i=1}^t a_i \geq (\ell_1 + 1) \cdot k_2, \quad (2)$$

where the LHS of inequality (1) lower-bounds the total number of value blocks from type-1 agents that are *entirely* included in the set A_+ ; the LHS of the inequality (2) is the total number of value blocks from type-2 agents that are intersected with the set A_+ .

Reformulating inequality (1) as

$$\sum_{i=1}^t a_i \leq \sum_{i=1}^t \max\{a_i - 1, 0\} + t \leq (\ell_2 - 1) \cdot k_1 + t,$$

and combining it with inequality (2), we know that both inequalities could be satisfied only when

$$t \geq (\ell_1 + 1) \cdot k_2 - (\ell_2 - 1) \cdot k_1 = k_1 + k_2 - 1 = k - 1,$$

where the last two equalities come from Lemma 7. Notice that $k - 1$ intervals need at least $2(k - 1) - 2$ cuts to separate them. Therefore, we prove that this α -ICD instance, which has $k_1 + k_2 = k$ agents, needs at least $2(k - 1) - 2$ cuts.

We prove the lower bound for arbitrary n by first copying the instance above for $c = \lfloor n/k \rfloor$ times. The first value block of the next copy is on the right of the last value block of the previous copy. By a similar argument, this new instance needs at least $c \cdot t$ intervals labeled with “+”. Thus, $2 \cdot c \cdot t - 2 = 2c \cdot (k - 1) - 2$ cuts are needed for these c copies with $c \cdot k$ agents. Finally, inserting $n - c \cdot k$ dummy agents with non-overlapping value blocks if k is not a factor of n . Each dummy agent needs at least one cut and the total number of cuts needed are

$$2c \cdot (k - 1) - 2 + (n - c \cdot k) \geq \frac{2(k - 1)}{k} \cdot n - k.$$

This concludes the lower bound for any rational number α . ◀

► **Remark.** The proof suggests that the lower bound also applies to the ε -approximate version of α -ICD with rational α for a sufficiently small constant ε . Precisely, let (α_1, α_2) be the smallest interval satisfying that $\alpha \in (\alpha_1, \alpha_2)$ and $k_1, k_2 < k$, where k, k_1, k_2 are the denominators of $\alpha, \alpha_1, \alpha_2$ respectively. The lower bound still holds as long as $\varepsilon < \min\{\alpha - \alpha_1, \alpha_2 - \alpha\}$.

4 Upper Bound and Implications on Complexity

We generalize the technique from [30] to improve the upper bounds for rational ratios α .

We define the set $Q^* \subset [0, 1]$ by the following generating rules:

1. $0, 1 \in Q^*$;
2. If $\frac{\ell}{k} \in Q^*$, then for any prime p that is not a factor of ℓ , we have $\frac{\ell}{k \cdot p}, 1 - \frac{\ell}{k \cdot p} \in Q^*$.

Notice that if k is prime, ℓ/k belongs to Q^* only when $\ell \in \{1, k-1\}$. For other values of k , there are still usually some fractions ℓ/k missing from Q^* , for example, $4/9$.

We show an upper bound for the set of ratios α in Q^* that differs from the corresponding lower bound of Theorem 5 by at most k . For any other rational number α , we give an upper bound that is smaller than $2n$ by a margin linear in n and thus separate the case of rational and irrational numbers α .

► **Theorem 10 (Upper Bound).** *For getting an exact solution of an α -ICD instance with n agents,*

1. *if α is rational and $\alpha = \ell/k \in Q^*$, $\frac{2(k-1)}{k} \cdot n$ cuts are always sufficient;*
2. *if α is rational and $\alpha \notin Q^*$, there is a constant $c_\alpha > 0$ such that $(2 - c_\alpha) \cdot n$ cuts are always sufficient;*
3. *if α is irrational, $2n$ cuts are always sufficient.*

Proof. Since [30] provided a general upper bound of $2n$ cuts for α -ICD with any ratio $\alpha \in [0, 1]$, it remains to consider the case with rational α .

For technical simplicity, we assume the measure function is defined on the circle S^1 , parameterized by $[0, 1]$ where the points 0 and 1 are identified as the same point. Any solution to the α -ICD defined on S^1 is also a valid solution to the corresponding α -ICD defined on the interval $[0, 1]$. Thus, any upper bound for the S^1 case is also valid for the interval case.

Case 1: $\alpha \in Q^*$.

We proceed by induction following the generating rule of Q^* .

Base Case: 0 cuts are needed when $\alpha = 0$ or $\alpha = 1$.

Induction Step: Assume we already have that for $\alpha = \frac{\ell}{k} \in Q^*$, $\frac{2(k-1)}{k} \cdot n$ cuts are enough. In other words, there is a solution with at most $\frac{k-1}{k} \cdot n$ intervals of “+” label.

For any prime p that is not a factor of ℓ , let $\alpha' = \frac{\ell}{k \cdot p}$. To get a solution for the α' -ICD instance, we first get a solution for the α -ICD with the same set of measure functions. We then solve the CONSENSUS- p -DIVISION problem, which uses $(p-1) \cdot n$ cuts, on the intervals with “+” labels in the first step.

Now, $\frac{k-1}{k} \cdot n$ intervals of “+” label are divided into $(\frac{k-1}{k} + p - 1) \cdot n$ intervals with label $+_1, \dots, +_p$. Intervals of each label take $\alpha/p = \alpha'$ fraction of measure and thus correspond to a possible solution for the α' -ICD instance. By averaging principle, there must be a label $+_i$ that labels at most $((\frac{k-1}{k} + p - 1) \cdot n)/p$ intervals. Observing that

$$((\frac{k-1}{k} + p - 1) \cdot n)/p = \frac{(k-1) + (p-1) \cdot k}{k \cdot p} \cdot n = \frac{k \cdot p - 1}{k \cdot p} \cdot n,$$

there is a solution to the α' -ICD instance using at most $\frac{2(k \cdot p - 1)}{k \cdot p} \cdot n$ cuts (two cuts for each interval with label $+_i$). Finally, $(1 - \alpha')$ -ICD is equivalent to α' -ICD. This concludes the proof for any ratio α in the set Q^* .

Case 2: $\alpha \notin Q^*$.

Let $\alpha_1 = \alpha = \ell_1/k_1 \notin Q^*$. We recursively generate a sequence of ratios as follows:

1. take $\alpha_i = (k_{i-1} \bmod \ell_{i-1})/k_{i-1}$; let ℓ_i and k_i be the numerator and the denominator of α_i in the simplest term;
2. if $\alpha_i \in Q^*$, take $t = i$ and stop this process.

Since $\ell_i < \ell_{i-1}$, the process above will stop in finite steps. We now prove by induction in the reverse order of index i .

Base Case: Since $\alpha_t \in Q^*$, there is a constant $c_{\alpha_t} = 2/k_t > 0$ such that $(2 - c_{\alpha_t}) \cdot n$ cuts are always sufficient for α_t -ICD.

Induction Step: Assume we already know that $(2 - c_{\alpha_i}) \cdot n$ cuts are always sufficient for α_i -ICD with a constant $c_{\alpha_i} > 0$. The same constant also holds for $(1 - \alpha_i)$ -ICD.

Notice that $1 - \alpha_i = (k_i - \ell_i)/k_i = (d \cdot \ell_{i-1})/k_{i-1}$ for some positive integer d . Therefore, α_{i-1} -ICD could be solved by first finding a solution of the $(1 - \alpha_i)$ -ICD instance with the same set of measure functions, and then solving the CONSENSUS- d -DIVISION problem on the previous solution set.

We use a similar argument in case 1 to count the number of cuts needed. The intervals labeled with “+” in the solution of $(1 - \alpha_i)$ -ICD is $(1 - c_{\alpha_i}/2) \cdot n$. Solving CONSENSUS- d -DIVISION incurs $(d - 1) \cdot n$ cuts and divides the previous solution set into $(d - c_{\alpha_i}/2) \cdot n$ intervals, with labels ranging from $+_1, \dots, +_d$. Each label corresponds to a valid solution for α_{i-1} -ICD, and there must be a label $+_i$ that labels at most $(1 - c_{\alpha_i}/2d) \cdot n$ intervals. Therefore, there is a constant $c_{\alpha_{i-1}} = c_{\alpha_i}/d > 0$ such that $(2 - c_{\alpha_{i-1}}) \cdot n$ cuts are always sufficient for α_{i-1} -ICD. ◀

The proof of Theorem 10 is essentially a reduction from α -ICD to CONSENSUS- k -DIVISION, and we thus derive several results on the complexity of α -ICD.

► **Theorem 11.** *For any $\alpha = \ell/k \in Q^*$, $k \geq 2$, solving exact α -ICD with $\lfloor \frac{2(k-1)}{k} \cdot n \rfloor$ cuts is in PPA- k under Turing reductions. In particular, if $k = p^r$ for a prime p , the problem lies in PPA- p .*

Proof. Let $k = p_1^{r_1} p_2^{r_2} \dots p_t^{r_t}$, where each p_i is prime and $r_i \geq 1$. In the proof of Theorem 10, α -ICD is solved by calling r_1 times of CONSENSUS- p_1 -DIVISION, r_2 times of CONSENSUS- p_2 -DIVISION, \dots , r_t times of CONSENSUS- p_t -DIVISION in a specific order.

The exact CONSENSUS- p -DIVISION problem is in the class PPA- p for any prime p [19]. PPA- p is a subset of PPA- q if p is a factor of q [23, 21], so we can deduce that solving exact α -ICD with $\lfloor \frac{2(k-1)}{k} \cdot n \rfloor$ cuts lies in PPA- k under Turing reductions.

In particular, when $k = p^r$ for a prime p , PPA- k is equal to PPA- p and PPA- p is closed under Turing reductions [23, 21]. ◀

We also consider the complexity of α -ICD when there are more available cuts than necessary.

► **Theorem 12.** *For any prime p and any ratio $\alpha \in [0, 1]$, solving α -ICD for inverse-polynomial approximation error ε with $2(p-1) \cdot \frac{\lceil p/2 \rceil}{\lfloor p/2 \rfloor} \cdot n$ cuts lies in PPA- p .*

Fixing a prime p , for any ratio α , our plan is to find a ratio $\alpha' = \ell'/p^m$ such that $|\alpha - \alpha'| \leq 1/n$, and showing that solving exact α' -ICD is in PPA- p . By definition 1, the exact solution for an α' -ICD instance is also a $(1/n)$ -approximate solution for the α -ICD instance with the same set of measure functions. The choice of m will be specified later.

To this end, we define a sequence of $m + 1$ sets $Q_p^0, Q_p^1, \dots, Q_p^m \subset [0, 1]$ as follows:

1. $Q_p^0 = \{0, 1\}$;
2. if $\alpha \in Q_p^i$, we have $\alpha, \frac{\lceil p/2 \rceil}{p} \cdot \alpha$ and $1 - \frac{\lceil p/2 \rceil}{p} \cdot \alpha$ in the set Q_p^{i+1} .

We now present two lemmas on the properties of sets $Q_p^0, Q_p^1, \dots, Q_p^m$.

► **Lemma 13.** *For any $\alpha \in Q_p^t, t \in [m]$, an α -ICD instance with n agents can be solved exactly by calling CONSENSUS- p -DIVISION at most t times, while using at most $2(p-1) \cdot \frac{\lceil p/2 \rceil}{\lceil p/2 \rceil} \cdot n$ cuts.*

Proof. We prove our claim by an induction argument on t , which is similar to case 1 of Theorem 10.

Base Case: When $\alpha \in \{0, 1\}$, α -ICD is trivial.

Induction Step: Now assume $\alpha \in Q_p^{t-1}$ and we have an exact solution of α -ICD by calling at most $t-1$ times of CONSENSUS- p -DIVISION and using at most $2(p-1) \cdot \frac{\lceil p/2 \rceil}{\lceil p/2 \rceil} \cdot n$ cuts. By definition, we also have $\alpha \in Q_p^t$.

Let $\alpha' = \frac{\lceil p/2 \rceil}{p} \cdot \alpha$. We apply CONSENSUS- p -DIVISION to the intervals with “+” label in the solution of α -ICD, which results in p labels $\{+1, \dots, +p\}$, and each of them corresponds to a set of intervals that consists of α/p fraction of measure. The total number of intervals that have one of labels $\{+1, \dots, +p\}$ are $(p-1) \cdot \frac{\lceil p/2 \rceil}{\lceil p/2 \rceil} \cdot n + (p-1) \cdot n$. We can find $\lceil p/2 \rceil$ labels $(+i_1, \dots, +i_{\lceil p/2 \rceil})$ and they correspond to at most

$$\left((p-1) \cdot \frac{\lceil p/2 \rceil}{\lceil p/2 \rceil} \cdot n + (p-1) \cdot n \right) \cdot \frac{\lceil p/2 \rceil}{p} = (p-1) \cdot \frac{\lceil p/2 \rceil}{\lceil p/2 \rceil} \cdot n$$

intervals. Therefore, we show that α' -ICD can be solved by calling at most t times of CONSENSUS- p -DIVISION and using at most $2(p-1) \cdot \frac{\lceil p/2 \rceil}{\lceil p/2 \rceil} \cdot n$ cuts. The case of $(1-\alpha')$ -ICD also holds by the equivalence. Our claim is now proven. ◀

► **Lemma 14.** *Let g_i ($0 \leq i \leq m$) to be the size of the largest hole of set Q_p^i , formally,*

$$g_i = \max\{d \in \mathbb{R} : \exists \alpha \in [0, 1] \text{ s.t. } (\alpha - d/2, \alpha + d/2) \cap Q_p^i = \emptyset\}.$$

Then, there is $g_i \leq g_{i-1} \cdot \frac{\lceil p/2 \rceil}{p}$.

Proof. Since for any $\alpha \in Q_p^{i-1}$, there is $\frac{\lceil p/2 \rceil}{p} \cdot \alpha \in Q_p^i$, the largest *hole* of Q_p^i within interval $[0, \frac{\lceil p/2 \rceil}{p}]$ is at most $g_{i-1} \cdot \frac{\lceil p/2 \rceil}{p}$. On the other hand, $\alpha \in Q_p^i$ implies $1 - \alpha \in Q_p^i$. Thus, the largest *hole* of Q_p^i within interval $[1 - \frac{\lceil p/2 \rceil}{p}, 1]$ is also at most $g_{i-1} \cdot \frac{\lceil p/2 \rceil}{p}$. Combining the two parts, we prove that $g_i \leq g_{i-1} \cdot \frac{\lceil p/2 \rceil}{p}$. ◀

Proof of Theorem 12. Notice that $\frac{\lceil p/2 \rceil}{p} \leq 2/3$ for any prime p . By taking $m = 2 \log_2 n$, for any ratio $\alpha \in [0, 1]$, there exists $\alpha' \in Q_p^m$ such that $|\alpha - \alpha'| \leq 1/n$ by Lemma 14. We also have $|Q_p^m| \leq 2 \cdot 3^m \leq O(n^4)$, which means that we can calculate the whole set Q_p^m and remember how each value in Q_p^m is generated. Thus, we can find α' and know how to reduce α' -ICD to m times of calling CONSENSUS- p -DIVISION in polynomial time. Again, using the fact CONSENSUS- p -DIVISION is in the class PPA- p for any prime p [19] and PPA- p is closed under Turing reductions [23, 21], we get solving exact α' -ICD with $2(p-1) \cdot \frac{\lceil p/2 \rceil}{\lceil p/2 \rceil} \cdot n$ cuts lies in PPA- p .

Finally, the exact solution of α' -ICD is also an inverse-polynomial approximated solution for α -ICD, which completes our proof. ◀

57:12 Consensus Division in an Arbitrary Ratio

► **Corollary 15.** *For any prime p , solving CONSENSUS- k -DIVISION for inverse-polynomial ε with $2(k-1) \cdot (p-1) \cdot \frac{\lceil p/2 \rceil}{\lfloor p/2 \rfloor} \cdot n$ cuts lies in PPA- p .*

Proof. CONSENSUS- k -DIVISION can be solved by calling α -ICD $k-1$ times such that we carve out $1/k$ fraction in each time.

Formally, in the i -th ($i \in [k-1]$) step,

1. we first rescale every measure function to weight 1;
2. we then solve $(1/(k-i+1))$ -ICD on these measure functions and mark intervals in the set A_+ (from the solution of $(1/(k-i+1))$ -ICD) with label i ;
3. finally, we set the function value in A_+ to be zero for all measure functions.

In the k -th step, we mark all the remaining intervals with the label k . The approximation error for solving each α -ICD instance will add up linearly. The overall approximation error of CONSENSUS- k -DIVISION is still inverse-polynomial if the error of solving each α -ICD instance is also inverse-polynomial. By Theorem 12, we conclude that solving CONSENSUS- k -DIVISION with $2(k-1) \cdot (p-1) \cdot \frac{\lceil p/2 \rceil}{\lfloor p/2 \rfloor} \cdot n$ cuts lies in PPA- p . ◀

5 NP-Hardness for the Exact Number of Cuts

The lower bound and upper bound results in Theorems 5 and 10 focus on the worst-case scenario. In practice, the number of cuts needed for a specific instance could be much fewer than the bound, and a solution with the minimum number of cuts would be preferred. In this section, we show that it's hard to decide the minimum number of cuts needed for a given α -ICD instance, which also implies finding such a solution is hard.

► **Theorem 16.** *For any $\alpha \in (0, 1)$, deciding the minimum number of cuts needed for a given α -ICD instance is NP-hard.*

Filos-Ratsikas et al. [15] show that deciding whether a CONSENSUS-HALVING instance with n agents has a solution using $n-1$ cuts is NP-complete. Therefore, we only need to consider the case for $\alpha < 1/2$.

We reduce the EXACTLY-1-3SAT problem to α -ICD for any $\alpha < 1/2$. EXACTLY-1-3SAT is a variant of 3SAT where the problem is to determine whether there exists a satisfying assignment such that *exactly one* literal in each clause is true, instead of at least one as in ordinary 3SAT. The NP-completeness of EXACTLY-1-3SAT is first shown in [26] as a special case of Schaefer's Dichotomy Theorem.

Suppose the given EXACTLY-1-3SAT instance ϕ has N variables x_1, \dots, x_N and M clauses c_1, \dots, c_M . Let k be an integer such that $\alpha \in [1/(k+1), 1/k)$. The α -ICD instance constructed is based on the construction in Section 3 for the lower bound, but has two more types of agents representing variables and clauses.

Type-1 agent

There are N type-1 agents and each of them has $2k$ value blocks with $1/2k$ weight each.

All the value blocks of type-1 agents are evenly spaced from left to right. Formally, the j -th block of i -th type-1 agent locates at interval $[T_1(i, j), T_1(i, j) + 1]$, where

$$T_1(i, j) = (3M + 3) \cdot ((i - 1) \cdot 2k + (j - 1)) .$$

Type-2 agent

There are $N \cdot k - 1$ type-2 agents and each of them has a single value block of weight 1.

The value block of the i -th type-2 agent is placed between the $2i$ -th and the $(2i + 1)$ -th blocks among all $N \cdot 2k$ blocks from type-1 agents. More specifically, it locates at $[T_2(i), T_2(i) + 1]$, where $T_2(i) = (3M + 4) + (i - 1) \cdot (6M + 6)$.

Variable agent

There are N variable agents and each of them has five value blocks. The first four blocks weigh $\alpha/2$ each, and the last block weighs $1 - 2\alpha$.

The first two blocks of i -th variable agent are located at

$$[T_1(i, 1) + 1, T_1(i, 1) + 2], [T_1(i, 1) + 3M + 2, T_1(i, 1) + 3M + 3]$$

respectively, both of which are between the first two blocks of the i -th type-1 agent; the next pair of blocks are placed between the third and the fourth blocks of the i -th type-1 agent, that are

$$[T_1(i, 3) + 1, T_1(i, 3) + 2], [T_1(i, 3) + 3M + 2, T_1(i, 3) + 3M + 3]$$

respectively.

The last block is located at $[T_1(N, 2k) + i, T_1(N, 2k) + i + 1]$, which is on the right of any blocks of type-1 agents.

Define intervals

$$I_{i,0}^v := [T_1(i, 1) + 1, T_1(i, 1) + 3M + 3] \text{ and } I_{i,1}^v := [T_1(i, 3) + 1, T_1(i, 3) + 3M + 3],$$

where $I_{i,0}^v$ covers the first two blocks, and $I_{i,1}^v$ covers the next two blocks. By definition, $I_{i,0}^v, I_{i,1}^v$ are between the first two and the next two blocks of the i -th type-1 agents respectively. Roughly speaking, only one of $I_{i,0}^v, I_{i,1}^v$ would be labeled with “+”, which corresponds to the value of variable x_i taking 0 or 1.

Clause agent

There are M clause agents and each of them has four value blocks. If $\alpha \geq 1/3$, the first three blocks have weight $(1 - \alpha)/2$ each and the last block has weight $(3\alpha - 1)/2$; otherwise, the first three blocks have weight α each and the last block has weight $1 - 3\alpha$.


Assuming the j -th ($j \leq 3$) literal of clause c_i is x_k or $\neg x_k$, let $\ell_{i,j} = 1$ if it's x_k and $\ell_{i,j} = 0$ if it's $\neg x_k$. The j -th block of the i -th clause agent is then placed at

$$[T_c(i, j, k, \ell_{i,j}), T_c(i, j, k, \ell_{i,j}) + 1] \in I_{k, \ell_{i,j}}^v,$$

where

$$T_c(i, j, k, \ell) = T_1(k, 1 + 2\ell) + 2 + (j - 1) \cdot M + i - 1, \text{ for } \ell \in \{0, 1\}.$$

If $\alpha \geq 1/3$, the last block of i -th clause agent is located in $[T_2(2) + i, T_2(2) + i + 1]$, which is next to the value block of the second type-2 agent; otherwise, it is located in $[T_1(N, 2k) + N + i, T_1(N, 2k) + N + i + 1]$, which is on the right of any blocks from type-1 agents.



■ **Figure 3** An incomplete view of value blocks in an α -ICD instance with $\alpha < 1/3$. Each color corresponds to a different agent. The clause agent in dark blue has $\neg x_i$ as one of its literals. The two black dashed lines represent possible locations of cuts that satisfy $x_i = 0$. The scale of each interval is not preserved.

Analysis

We first ignore all the variable agents and clause agents.

► **Lemma 17.** *At least $2 \cdot (N \cdot k - 1)$ cuts are needed if only type-1 and type-2 agents are considered. Also, such a solution must satisfy the following conditions,*

1. *there are exactly $(N \cdot k - 1)$ intervals labeled with “+”, denoted by $I_1^+, \dots, I_{N \cdot k - 1}^+$ from left to right. The length of $[T_2(i), T_2(i) + 1] \cap I_i^+$ is exactly α for any i .*
2. *the first and the last interval are labeled with “-”;*
3. *for any $i \in [N]$, at most one of intervals $I_{i,0}^v, I_{i,1}^v$ could possibly have intersection with “+” labeled intervals.*

Proof. Similar to the proof of Theorem 5, notice that each block from type-2 agents is separated by two blocks of type-1 agents, and the sum of their weight is $1/k > \alpha$. Thus, each “+” labeled interval could only intersect with at most one type-2 agent; also, if a “+” labeled interval intersects with any type-2 agent, it can not be the first or the last interval in the solution.

Since each type-2 agent should be intersected with at least one “+” labeled interval, there are at least $(N \cdot k - 1)$ intervals labeled with “+”, and those intervals need at least $2 \cdot (N \cdot k - 1)$ cuts. On the other hand, any solution with $2 \cdot (N \cdot k - 1)$ cuts must satisfy conditions 1 and 2 by our discussion above.

For the last condition, if both $I_{i,0}^v, I_{i,1}^v$ are intersected by “+” labeled intervals, then at least two value blocks of i -th type-1 agent are fully covered with “+” label, which is impossible since $1/k > \alpha$. ◀

Next, we show that the EXACTLY-1-3SAT instance ϕ is satisfiable if and only if $2 \cdot (N \cdot k - 1)$ cuts are enough for the whole α -ICD instance, including variable and clause agents.

► **Lemma 18.** *If the EXACTLY-1-3SAT instance ϕ is satisfiable, then $2 \cdot (N \cdot k - 1)$ cuts are enough for the α -ICD instance.*

Proof. Let (y_1, \dots, y_N) be a set of satisfying assignments of ϕ . We construct a valid solution for the α -ICD instance by specifying each “+” labeled intervals $I_1^+, \dots, I_{N \cdot k - 1}^+$.

- for any $i \in [N]$, if $y_i = 0$, interval $I_{(i-1) \cdot k + 1}^+$ is chosen as

$$[T_1(i, 1) + 2k \cdot ((1/k) - \alpha), T_2((i - 1) \cdot k + 1) + \alpha],$$

otherwise, it's set as

$$[T_2((i-1) \cdot k + 1) + (1 - \alpha), T_1(i, 4) + 1 - 2k \cdot ((1/k) - \alpha)];$$

- for any $j \neq (i-1) \cdot k + 1$ for all $i \in [N]$, interval I_j^+ is set to be $[T_2(j) + (1 - \alpha), T_2(j) + M + 1]$.

Now let's verify that all types of agents are satisfied by the solution above.

Type-1 For any $i \in [N]$, interval $I_{(i-1) \cdot k + 1}^+$ always covers exactly α fraction of the i -th type-1 agent, and any other “+” labeled intervals have no intersection with the i -th type-1 agent.

Type-2 For any $j \in [N \cdot k - 1]$, interval I_j^+ covers exactly α fraction of the j -th type-2 agent, and any other “+” labeled intervals have no intersection with the j -th type-2 agent.

Variable For any $i \in [N]$, interval $I_{(i-1) \cdot k + 1}^+$ always covers one of the $I_{i,0}^v, I_{i,1}^v$, and has no intersection with the other one. Thus, exactly two blocks of weight $\alpha/2$ from the i -th variable agent are covered with the label “+”. The last block is always not covered.

Clause For any $j \in [M]$, exactly one of the first three blocks of the j -th clause agent is covered, since (y_1, \dots, y_N) is a set of satisfying assignment of ϕ . If $\alpha \geq 1/3$, the last block is covered by I_2^+ ; if $\alpha < 1/3$, the last block is not covered. In both cases, exactly α fraction of the j -th clause agents are covered by “+” label. ◀

In the reverse direction, we have the following characterization for any valid solution with $2 \cdot (N \cdot k - 1)$ cuts.

▶ **Lemma 19.** *If there is a valid solution to the α -ICD instance with $2 \cdot (N \cdot k - 1)$ cuts, then*

1. *for any $i \in [N]$, exactly one of the $I_{i,0}^v, I_{i,1}^v$ is fully covered with label “+”, while the other one is fully covered with label “-”;*
2. *for any $i \in [M]$, exactly one of the first three value blocks of the clause agent i is fully covered with the label “+”, while the other two are fully covered with the label “-”.*

Proof. By the second condition in Lemma 17, we know that the last block of any variable agents will not be covered by the “+” label. Then by the third condition in Lemma 17, we know that the i -th variable agent could be satisfied only when one of the $I_{i,0}^v, I_{i,1}^v$ is fully covered with the label “+”. This concludes our first statement.

Any one of the first three value blocks of the clause agent i would be either fully covered with the label “+”, or fully covered with the label “-”. given by the first statement.

- When $\alpha < 1/3$, exactly one of the first three blocks should be fully covered with the label “+” since the last block will never be covered.
- When $\alpha \geq 1/3$, at most one of the first three blocks should be fully covered with the label “+”, since $2 \cdot ((1 - \alpha)/2) > \alpha$; also, at least one of the first three blocks should be fully covered with the label “+”, since the weight of the last block is $(3\alpha - 1)/2 < \alpha$. ◀

By Lemma 19, we could retrieve a satisfying assignment (y_1, \dots, y_N) for the EXACTLY-1-3SAT instance ϕ by letting $y_i = \ell$ if interval $I_{i,\ell}^v$ is covered with label “+”. Combining with Lemma 18, we conclude the correctness of our reduction.

6 Future Work

The number of cuts needed for α -ICD

There is still a small gap between our lower bound and upper bound (Theorem 5, 10) for any rational ratio α that is not in the set Q^* . We believe the lower bound is tight, while the upper bound could be further improved, e.g., by applying an appropriate \mathbb{Z}_p variant of the Borsuk-Ulam theorem.

Computational complexity

Given the minimal number of cuts that make α -ICD a total problem, we show the connection of solving α -ICD with the complexity classes PPA- k in Theorem 11 when α is in the set Q^* . To extend this result to all rational ratios, the same technique which improves the upper bound may be required. When more cuts than necessary are available, we conjecture that with $2n$ cuts, for any ratio $\alpha \in [0, 1]$, solving α -ICD will be in the intersection of PPA- p for all prime p .

Any hardness result for α -ICD would also be of interest. α -ICD seems a promising candidate to be a natural complete problem of classes PPA- k , while very few natural complete problems for PPA- k are known [21]. Moreover, α -ICD is closely related to CONSENSUS- k -DIVISION, as the proof of Theorem 10 builds a reduction from α -ICD to CONSENSUS- k -DIVISION. The only hardness result currently known for CONSENSUS- k -DIVISION is that CONSENSUS-3-DIVISION is PPAD-hard, given by [18].

References

- 1 Noga Alon. Splitting necklaces. *Advances in Mathematics*, 63(3):247–253, 1987.
- 2 Noga Alon and Andrei Graur. Efficient splitting of measures and necklaces. *CoRR*, abs/2006.16613, 2020. [arXiv:2006.16613](#).
- 3 Noga Alon and Douglas B West. The borsuk-ulam theorem and bisection of necklaces. *Proceedings of the American Mathematical Society*, 98(4):623–628, 1986.
- 4 Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair division of indivisible goods: A survey. *CoRR*, abs/2208.08782, 2022. [doi:10.48550/arXiv.2208.08782](#).
- 5 Haris Aziz, Hervé Moulin, and Fedor Sandomirskiy. A polynomial-time algorithm for computing a pareto optimal and almost proportional allocation. *Oper. Res. Lett.*, 48(5):573–578, 2020. [doi:10.1016/j.orl.2020.07.005](#).
- 6 Eleni Batziou, Kristoffer Arnsfelt Hansen, and Kasper Høgh. Strong approximate consensus halving and the borsuk-ulam theorem. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPICs*, pages 24:1–24:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. [doi:10.4230/LIPICs.ICALP.2021.24](#).
- 7 Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. Weighted envy-freeness in indivisible item allocation. *ACM Trans. Economics and Comput.*, 9(3):18:1–18:39, 2021. [doi:10.1145/3457166](#).
- 8 Mithun Chakraborty, Ulrike Schmidt-Kraepelin, and Warut Suksompong. Picking sequences and monotonicity in weighted fair division. *Artif. Intell.*, 301:103578, 2021. [doi:10.1016/j.artint.2021.103578](#).
- 9 Logan Crew, Bhargav Narayanan, and Sophie Spirkl. Disproportionate division. *Bull. London Math. Soc.*, 52:885–890, 2020.
- 10 Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. Constant inapproximability for PPA. *CoRR*, abs/2201.10011, 2022. [arXiv:2201.10011](#).
- 11 Argyrios Deligkas, John Fearnley, Themistoklis Melissourgos, and Paul G. Spirakis. Computing exact solutions of consensus halving and the borsuk-ulam theorem. *J. Comput. Syst. Sci.*, 117:75–98, 2021. [doi:10.1016/j.jcss.2020.10.006](#).
- 12 Argyrios Deligkas, Aris Filos-Ratsikas, and Alexandros Hollender. Two’s company, three’s a crowd: Consensus-halving for a constant number of agents. In *EC ’21: The 22nd ACM Conference on Economics and Computation*, pages 347–368. ACM, 2021.

- 13 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.*, 313:103784, 2022. doi:10.1016/j.artint.2022.103784.
- 14 Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010.
- 15 Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Paul W. Goldberg, and Jie Zhang. Hardness results for consensus-halving. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS*, volume 117 of *LIPIcs*, pages 24:1–24:16, 2018.
- 16 Aris Filos-Ratsikas and Paul W Goldberg. Consensus halving is PPA-complete. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 51–64, 2018.
- 17 Aris Filos-Ratsikas and Paul W Goldberg. The complexity of splitting necklaces and bisecting ham sandwiches. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 638–649, 2019.
- 18 Aris Filos-Ratsikas, Alexandros Hollender, Katerina Sotiraki, and Manolis Zampetakis. Consensus-halving: Does it ever get easier? In *EC '20: The 21st ACM Conference on Economics and Computation*, pages 381–399. ACM, 2020.
- 19 Aris Filos-Ratsikas, Alexandros Hollender, Katerina Sotiraki, and Manolis Zampetakis. A topological characterization of modulo- p arguments and implications for necklace splitting. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms*, pages 2615–2634. SIAM, 2021.
- 20 Paul W. Goldberg, Alexandros Hollender, Ayumi Igarashi, Pasin Manurangsi, and Warut Suksompong. Consensus halving for sets of items. In *Web and Internet Economics - 16th International Conference, WINE*, volume 12495 of *LNCS*, pages 384–397. Springer, 2020.
- 21 Mika Göös, Pritish Kamath, Katerina Sotiraki, and Manolis Zampetakis. On the complexity of modulo- q arguments and the Chevalley–Warning theorem. In *35th Computational Complexity Conference, CCC*, volume 169 of *LIPIcs*, pages 19:1–19:42. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- 22 Charles R. Hobby and John R. Rice. A moment problem in l_1 approximation. *Proceedings of the American Mathematical Society*, 16(4):665–670, August 1965.
- 23 Alexandros Hollender. The classes PPA- k : Existence from arguments modulo k . *Theor. Comput. Sci.*, 885:15–29, 2021. doi:10.1016/j.tcs.2021.06.016.
- 24 Nimrod Megiddo and Christos H Papadimitriou. On total functions, existence theorems and computational complexity. *Theoretical Computer Science*, 81(2):317–324, 1991.
- 25 Christos H Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences*, 48(3):498–532, 1994.
- 26 Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing*, pages 216–226. ACM, 1978. doi:10.1145/800133.804350.
- 27 Erel Segal-Halevi. Cake-cutting with different entitlements: How many cuts are needed? *J. Math. Anal. Appl.*, 480:1–5, 2019.
- 28 Erel Segal-Halevi. Fair multi-cake cutting. *Discret. Appl. Math.*, 291:15–35, 2021.
- 29 Forest W Simmons and Francis Edward Su. Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Mathematical social sciences*, 45(1):15–25, 2003.
- 30 Walter Stromquist and Douglas R Woodall. Sets on which several measures agree. *Journal of mathematical analysis and applications*, 108(1):241–248, 1985.
- 31 Biaoshuai Tao. On existence of truthful fair cake cutting mechanisms. *CoRR*, abs/2104.07387, 2021. arXiv:2104.07387.

A Proof of Lemmas 6,7

► **Lemma 6.** *Let $\alpha_1 := \ell_1/k_1$, $\alpha_2 := \ell_2/k_2$, $\alpha_1 < \alpha_2$ be two adjacent fractions. For any fraction $\alpha := \ell/k$, $\alpha \in (\alpha_1, \alpha_2)$, we have $k \geq k_1 + k_2$.*

57:18 Consensus Division in an Arbitrary Ratio

Proof.

$$\frac{1}{k_1 \cdot k_2} = \frac{\ell_2}{k_2} - \frac{\ell_1}{k_1} = \left(\frac{\ell_2}{k_2} - \frac{\ell}{k} \right) + \left(\frac{\ell}{k} - \frac{\ell_1}{k_1} \right) \geq \frac{1}{k \cdot k_2} + \frac{1}{k \cdot k_1} = \frac{k_1 + k_2}{k \cdot k_1 \cdot k_2}.$$

Thus, we have $k \geq k_1 + k_2$. ◀

► **Lemma 7.** Given a fraction $\alpha := \ell/k, k \geq 2$, let (α_1, α_2) be the smallest interval satisfying that $\alpha \in (\alpha_1, \alpha_2)$ and $k_1, k_2 < k$, where $\alpha_1 := \ell_1/k_1, \alpha_2 := \ell_2/k_2$. Then

1. α_1 and α_2 are adjacent;
2. $\ell = \ell_1 + \ell_2, k = k_1 + k_2$;
3. α is adjacent to both α_1, α_2 .

Proof. We first prove the first statement by induction on k .

Base Case $k = 2$: The numerator ℓ takes value 1 in this case. It's easy to verify the correctness as we denote 0, 1 by 0/1, 1/1 respectively.

Now for any $k > 2$, we prove by contradiction. Assume $\ell_2 \cdot k_1 - \ell_1 \cdot k_2 = t > 1$, and consider the following three cases:

1. $k_1 < k_2$: take (α_3, α_4) as the smallest interval satisfying that $\alpha_2 \in (\alpha_3, \alpha_4)$ and α_3, α_4 have smaller denominator than that of α_2 . We know that $\alpha_3 \geq \alpha_1$ since k_1 , the denominator of α_1 , is also smaller than k_2 . By the induction hypothesis, α_2 and α_3 are *adjacent*, while α_1 and α_2 are not *adjacent*, indicating that α_3 is strictly larger than α_1 . Thus, either (α_1, α_3) or (α_3, α_2) contains α , contradicting to the fact that (α_1, α_2) is the smallest interval;
2. $k_1 > k_2$: this case follows from similar argument as the case of $k_1 < k_2$;
3. $k_1 = k_2$: since $k > 2$, we know that $k_1 > 1$ and $0 < \ell_1 < \ell_2 < k_1$. Take $\alpha_3 := \ell_1/(k_1 - 1)$, and it's easy to verify that $\alpha_3 \in (\alpha_1, \alpha_2)$. Thus, either (α_1, α_3) or (α_3, α_2) contains α , a contradiction.

We conclude the first statement by induction. Next, we claim that $k = k_1 + k_2$; otherwise, by the Lemma 6, we must have $k > k_1 + k_2$ and either $(\alpha_1, (\ell_1 + \ell_2)/(k_1 + k_2))$ or $((\ell_1 + \ell_2)/(k_1 + k_2), \alpha_2)$ will contain α , which again contradicts to the fact that (α_1, α_2) is the smallest interval.

Now, we argue that $\ell = \ell_1 + \ell_2$. If $\ell > \ell_1 + \ell_2$, then

$$\frac{1}{k} \leq \frac{\ell}{k} - \frac{\ell_1 + \ell_2}{k} < \frac{\ell_2}{k_2} - \frac{\ell_1 + \ell_2}{k} = \frac{1}{k_2 \cdot k},$$

which is impossible; the case with $\ell < \ell_1 + \ell_2$ could be ruled out with same argument.

The third statement follows from statements 1 and 2. ◀