






Beyond Exact Fairness: Envy-Free Incomplete Connected Fair Division

Ajaykrishnan E S  

University of California, Santa Barbara, CA, USA

Daniel Lokshtanov   

University of California, Santa Barbara, CA, USA

Abstract

We study the problem of ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION, where exactly p vertices of an undirected graph must be allocated to agents such that each agent receives a connected share and does not envy another agent's share. Focusing on agents with additive valuations, we show that the problem remains computationally hard when parameterized by p and the number of agents. This result holds even for star graphs and with the input numbers given in unary representation, thereby resolving an open problem posed by Gahlawat and Zehavi (FSTTCS 2023).

In stark contrast, we show that if one is willing to tolerate even the slightest amount of envy, then the problem becomes efficient with respect to the natural parameters. Specifically, we design an Efficient Parameterized Approximation Scheme parameterized by p and the number of agent types. Our algorithm works on general graphs and remains efficient even when the input numbers are provided in binary representation.

2012 ACM Subject Classification Theory of computation \rightarrow W hierarchy; Theory of computation \rightarrow Fixed parameter tractability

Keywords and phrases Envy-Free Incomplete Connected Fair Division, Efficient Parameterized Approximation Scheme, $W[1]$ -hardness

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2025.29

Funding NSF Grant No. 2505099, Collaborative Research: AF Medium: Structure and Quasi-Polynomial Time Algorithms.

1 Introduction

Fair division is a critical problem in resource allocation, focused on distributing resources among multiple agents in a way that meets certain fairness criteria. The theoretical study of this problem has been a core area of research, with work dating as far back as the 1940s [33]. Fair division, beyond its theoretical appeal, has numerous practical applications, as demonstrated by platforms like Spliddit [24], which use fair allocation algorithms to solve real-world problems such as sharing rent, splitting taxi fares, and assigning goods. However, for many of these practical scenarios, one ends up having to address more complex constraints. An example of this is the case where the goods are *indivisible* and the bundle of items corresponding to each agent must satisfy some *connectivity* requirements. The φ -CONNECTED FAIR DIVISION (φ -CFD) problem, introduced by Bouveret et al. [11] models this using a graph whose vertices represent items, where the bundles allocated to agents are required to form connected subsets of the vertices, where φ specifies some desired notion of fairness.

However, this model requires that every vertex of the graph be allocated to some agent, which fails to cover some natural scenarios where there is a need to be *economic* in the allocation. As an example consider a museum that plans to allocate sections of a large exhibition hall for various temporary art exhibits. Each exhibit needs a connected space to display its artwork effectively, but the museum must also reserve areas for future events or special exhibits. The challenge then becomes to assign some part of the available space fairly among the groups, ensuring each receives a connected section, while leaving enough space



© Ajaykrishnan E S and Daniel Lokshtanov;
licensed under Creative Commons License CC-BY 4.0

45th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2025).

Editors: C. Aiswarya, Ruta Mehta, and Subhajit Roy; Article No. 29; pp. 29:1–29:16



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

unallocated for other purposes. Another scenario where being economic is crucial involves allocating resource-intensive goods, such as when a company rents office spaces or a startup purchases equipment for its employee teams.

Motivated by these considerations, Gahlawat and Zehavi [21] introduced the φ -INCOMPLETE CONNECTED FAIR DIVISION problem (φ -ICFD). In this, the goal is to find a *valid* and φ -allocation, that is, an allocation of exactly p vertices of the graph where each agent receives a distinct connected bundle that satisfies the fairness criterion specified by φ . From the perspective of parameterized complexity, they consider this problem for various well known notions of fairness, one of them being envy-freeness. An allocation of vertices to agents is *envy-free* if, for every agent, the value they assign to their own bundle is at least as large as the value they assign to any other agent's bundle. They show that, if the valuations are provided in binary representation, then the ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION problem becomes W[1] hard parameterized by p and the number of agents. However, their work leaves behind the following open problem: does ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION admit a Fixed-Parameter Tractable (FPT) algorithm parameterized by number of agents and p , if the valuations are given in unary representation?

1.1 Our Contribution

We resolve the open problem posed by Gahlawat and Zehavi [21] by proving that ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION is W[1]-hard when parameterized by $p + |A|$, where A is the set of agents, when the input numbers are in unary representation, even for highly restricted input instances such as star graphs. Our result follows via a reduction from (k, M) -VECTOR-SUM, a multidimensional variant of the classic NP-hard problem SUBSET SUM, shown to be W[1]-hard by Abboud et al. [1], parameterized by $k + d$, where d is the dimension of the input vectors, even when the input numbers are in unary representation.

► **Theorem 1.1.** *ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION is W[1]-hard parameterized by $p + |A|$, with the input numbers provided in unary representation, even on star graphs.*

To address the intractability implied by Theorem 1.1, we approach the problem from the perspective of approximate fairness. The study of finding *approximately fair* allocations, meaning allocations that are at most a small multiplicative (or additive factor) unfair to the agents, was initiated by Lipton et al. [29] in the context of cake-cutting and has since been investigated in other settings. For a comprehensive overview of this topic, we refer the reader to the survey by Amanatidis et al. [4], which discusses approximate fairness in various settings. In our case, an allocation is ε -*envy-free* if, for every agent, the value they assign to their own bundle scaled up by a multiplicative factor of $(1 + \varepsilon)$, is at least as large as the value they assign to any other agent's bundle.

In this setting, it is very natural to apply a geometric rounding scheme to round the involved numbers to the nearest power of $(1 + \varepsilon)$ and make some guesses regarding the agents valuation of allocated items. This is especially tempting given that the hardness of the instance in the proof of Theorem 1.1 arises precisely from the inability to iterate over all possible valuations that agents may assign to the allocated items. This does in fact lead to a straightforward FPT algorithm, parameterized by $p + |A|$, that finds a valid and ε -envy-free allocation or correctly concludes that the input instance has no valid and envy-free allocation, if the input numbers are provided in unary representation. But, in the case where the input numbers are provided in binary representation, the agents could have wildly different valuations for the same items and consequently the number of guesses

will no longer be FPT in terms of the relevant parameters. In fact, due to this reason, for some time we believed that there exists a positive real number c such that, assuming the Exponential Time Hypothesis [27], no FPT algorithm can either find a valid and c -envy-free allocation or correctly conclude that the input instance admits no valid and envy-free allocation. Surprisingly, this is not the case – the problem admits an algorithm with the aforementioned guarantees that runs in FPT time parameterized by p and $|\mathcal{A}|$, the number of *agent types*, which is a smaller parameter than the number of agents.

► **Theorem 1.2.** *There exists an algorithm that, given an instance (G, A, \mathcal{U}, p) of ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION with input numbers in binary representation, and a positive real number ε , runs in time $(\frac{1}{\varepsilon} \log \frac{p}{\varepsilon})^{\mathcal{O}(|\mathcal{A}|p^2)} (mn)^{\mathcal{O}(1)}$ and outputs either a valid and ε -envy-free allocation Π , or correctly states that (G, A, \mathcal{U}, p) does not admit a valid and envy-free allocation.*

We now give a sketch of the algorithm asserted in Theorem 1.2. The algorithm begins by applying a simple reduction rule that allows us to upper bound the number of agents by a function of p and $|\mathcal{A}|$. After this, we proceed to apply the standard technique of color coding [3] (see also [17]) to color the vertices of the graph with p colors, and make some guesses regarding the color of the vertices that end up being assigned to each agent.

At this point, we draw inspiration from the polynomial-time algorithm of Gan et al. [22], which addresses the special case where each agent is allocated exactly one item, sometimes referred to as the *house allocation* problem. To illustrate, consider a more restricted setting in which every agent is assigned a unique color and must be allocated exactly one vertex which has the same color. In this setting, for each agent, we consider the set of items they value the most. If there exists an agent such that none of their most-valued items have the same color as them, then these items can be safely discarded, as allocating them to any agent would make this agent envious. On the other hand, if for every agent there exists at least one most-valued item that has the same color as them, then assigning each agent such an item yields an envy-free allocation. The key idea here is that we can either find an envy-free allocation, or restrict the search-space by a significant amount.

Our setting is more technically challenging since we need to work with approximate envy-freeness, enforce connectivity constraints, and accommodate agents that receive multiple items. For this reason, we encode the restriction of the search space using a *target* vector $\bar{\mu} = (\mu_1, \dots, \mu_{|\mathcal{A}|})$, and an allocation *attains* the target for agent i , if the valuation that agent i has for their own bundle is between $(1 + \varepsilon)^{\mu_i - 1}$ and $(1 + \varepsilon)^{\mu_i}$. The algorithm guesses a set of properties (of which there are only FPT many) of the valid and envy-free allocation we are searching for, sets the target vector so that $(1 + \varepsilon)^{\mu_i} \geq pM_i$, where M_i denotes the valuation of agent i 's favorite item, and runs a subroutine that searches for a valid and envy-free allocation that respects these guesses.

The subroutine has an outer loop which, following the key idea from the simpler setting, either finds a solution or decreases the target in each iteration. To prove correctness, we must show that if there exists a valid and envy-free allocation that respects the guesses and attains a target that is at most the current one, then the algorithm does one of the following: it either returns a valid and ε -envy-free allocation, or identifies an agent i such that no valid and envy-free allocation that respects the guesses and attains a target that is at most the current one attains the current target for agent i . We design a procedure that achieves this by observing that the guesses constrain the instances to the point that the allocations corresponding to each agent can be searched for independently of each other. Finding such a feasible assignment for each individual agent is done by reducing to a weighted version of the well studied GRAPH MOTIF problem, first introduced by Lacroix et al. [28]. This concludes our sketch.

We remark that a subtle but crucial aspect for the correctness of the algorithm is that the feature guessing must be performed outside the subroutine. Additionally, we draw the reader's attention to the fact that we consider two subtly different versions of the problem. In the optional setting, agents may be assigned an empty bundle, whereas in the mandatory setting, each agent must receive a non-empty set of items. Gahlawat and Zehavi [21] study only the mandatory setting. Our results apply to both the optional and mandatory settings. We first establish our results for the optional setting and then show how they extend to the mandatory case.

1.2 Related Work

Fair division of indivisible goods have received a considerable amount of attention in the computational social choice literature, an interested reader can refer some of the recent surveys [4, 34]. The work introducing the fair division of *indivisible* items with *connectivity* requirements on the players' bundles was carried out by Bouveret et al. [11]. Their initial results inspired a series of follow-up studies on the model. This body of work includes studies on conditions guaranteeing the existence of fair allocations [8], achieving various fairness notions such as Pareto-optimal allocations [26] and maximin share allocations [25]. Additionally, Aziz et al. [5] examined scenarios where the items to be allocated could be either goods or bads. There has also been work, examining the *price of connectivity* [6], which evaluates the loss in fairness when allocation are required to be connected, when compared to an optimal disconnected allocation. The study of connected fair division where only a subset of the items are to be allocated, was initiated by Gahlawat and Zehavi [21]. Related notions that have been explored in literature include *donating* a few items to achieve fairness [16] and imposing *cardinality constraints* on the number of items of different types that an agent can receive [9].

The concept of finding *approximately fair* allocations that are at most a small multiplicative (or additive factor) unfair to the agents was introduced by Lipton et al. [29] in the concept of cake-cutting and later on by Budish [15] in the context of indivisible goods. Since then, it has been studied across a variety of settings and fairness notions, including envy-freeness [32], Nash social welfare [2], and max-min fairness [23, 20], among others.

The area of parameterized complexity has found notable use in fair division literature, as many fundamental problems in this domain are (unsurprisingly) NP-hard. The works most relevant to ours include that of Deligkas et al. [18], which provides a comprehensive overview of the parameterized complexity of φ -CFD with respect to the natural parameters as well as several structural parameters of the graph for various fairness notions, and that of Gahlawat and Zehavi [21] which does the same for for the φ -ICFD problem. There is a large body of work separate from this [31, 10], which consider fair division problems from a parameterized perspective with the added requirement of Pareto-optimality. The works of Bredereck et al. [13, 12] apply N -fold Integer Programming to solve problems concerning envy-free and Pareto-efficient allocations of items, with the latter even handling instances where input numbers are given in binary representation. Furthermore, the problem of finding *locally* envy-free allocations, where the input includes a graph in which vertices represent agents, and an agent can only envy their neighbors in the graph have also been explored [7, 14].

2 Preliminaries

For a positive integer n , let $[n]$ be the set $\{1, 2, \dots, n\}$, and $[a, b]$ denote the set of integers $i \in \mathbb{Z}$ that satisfies $i \leq b$ and $i \geq a$. Let $\mathbb{N}_{\geq 0} = \mathbb{N} \cup \{0\}$. For vectors $\bar{v}^{(1)}, \bar{v}^{(2)} \in \mathbb{Z}^n$, we write $\bar{v}^{(1)} \leq \bar{v}^{(2)}$ (respectively, $\bar{v}^{(1)} = \bar{v}^{(2)}$) to mean that $v_i^{(1)} \leq v_i^{(2)}$ (respectively, $v_i^{(1)} = v_i^{(2)}$),

where $v_i^{(1)}$ and $v_i^{(2)}$ denote the i -th coordinates of $\bar{v}^{(1)}$ and $\bar{v}^{(2)}$, for every $i \in [n]$. Furthermore, we use $\bar{v}^{(1)} + \bar{v}^{(2)}$ to denote the vector $\bar{v} \in \mathbb{Z}^n$ that satisfies $v_i = v_i^{(1)} + v_i^{(2)}$ for every $i \in [n]$. We consider finite, simple, undirected graphs. For standard graph terminology we refer to the book by Diestel [19]. We make a slight departure from standard notations by denoting the number of vertices in a graph by m .

An input instance of the ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION problem consists of a tuple (G, A, \mathcal{U}, p) , where,

- G , the *utility graph*, is a graph on m vertices
- $A = [n]$ is the set of *agents*, where $n \in \mathbb{N}$
- \mathcal{U} is an n -tuple of *valuation functions* $\{u_i : V(G) \rightarrow \mathbb{N}_{\geq 0}\}_{i \in [n]}$
- $p \in [0, m]$ is a non-negative integer

The vertices of G are referred to as *items*. As is standard in literature, we assume that agents have *additive valuations* and use $u_i(S)$ to denote $\sum_{v \in S} u_i(v)$ for any $S \subseteq V(G)$. An *allocation* of items is a tuple $\Pi = (\pi_1, \dots, \pi_n)$, where $\pi_i \subseteq V(G)$ is referred to as the *bundle* of agent i . When clear from the context, we sometimes abuse notations and use π to denote $\bigcup_{i \in A} \pi_i$, the set of allocated vertices. An allocation is *valid* if it allocates exactly p vertices in total, with each bundle being connected and no vertex allocated to multiple bundles. An allocation is *envy-free* if no agent desires another agent's bundle more than their own. Formally, an allocation that is both valid and envy-free must satisfy,

- $\sum_{i=1}^n |\pi_i| = p$
- $G[\pi_i]$ is connected
- $\pi_i \cap \pi_j = \emptyset$ for every pair of distinct agents i, j
- $u_i(\pi_i) \geq u_i(\pi_j)$ for every pair of agents i, j

Similarly, an allocation is ε -*envy-free*, for some positive real number ε , if it satisfies $(1 + \varepsilon)u_i(\pi_i) \geq u_i(\pi_j)$ for every agent i and j . Given an instance $\mathcal{I} = (G, A, \mathcal{U}, p)$, the ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION problem asks whether there exists an allocation in \mathcal{I} that is both valid and envy-free. Here we remark that there are two subtly different versions of the problem. In the *mandatory* setting, each agent must be assigned a non-empty bundle; that is, $|\pi_i| > 0$ for every agent $i \in A$. In contrast, the *optional* setting allows agents to receive empty bundles, meaning $|\pi_i| \geq 0$ for all $i \in A$. Throughout this paper, we use the shorthand EF-ICFD to refer to the envy-free incomplete connected fair division problem in the optional setting.

Two agents i, j have the same *type* if $u_i(v) = u_j(v)$ for every $v \in V(G)$. We use $\mathcal{A}_{\mathcal{I}}$ to denote the set of all agent types in an instance \mathcal{I} of EF-ICFD. Thus $\mathcal{A}_{\mathcal{I}}$ is a partition of A such that two agents i, j belong to the same part $\alpha \in \mathcal{A}_{\mathcal{I}}$ if and only if they have the same type. For ease of notation, we shall drop the subscript \mathcal{I} whenever it is clear from the context. We say that an instance (G, A, \mathcal{U}, p) of EF-ICFD is *in unary* (respectively, *in binary*) if the valuation functions in \mathcal{U} are encoded in unary (respectively, binary). This means there exists a polynomial (respectively, exponential) function f such that $\max\{u_i(v) \mid i \in A, v \in V(G)\} \leq f(|V(G)|)$.

For standard terminology in parameterized algorithms, we refer to the book Parameterized Algorithms [17] by Cygan et al. An (n, k) -*perfect hash family* [30] is a family of functions from $[n]$ to $[k]$ such that for every $S \subseteq [n]$ with $|S| = k$, there exists a function in the family that is injective on S . The following is known regarding (n, k) -perfect hash families.

► **Proposition 2.1** ([30], [17]). *For any $n, k \in \mathbb{N}$, one can construct an (n, k) -perfect hash family of size $e^k k^{\mathcal{O}(\log k)} \log n$ in time $e^k k^{\mathcal{O}(\log k)} n \log n$.*

An *Efficient Parameterized Approximation Scheme* (EPAS) for an optimization problem Π with input size n and parameter k is an algorithm that, for every $\varepsilon > 0$, runs in time $f(k, 1/\varepsilon) \cdot n^{\mathcal{O}(1)}$ for some computable function f , and returns a $(1 + \varepsilon)$ -approximate solution (or $(1 - \varepsilon)$ -approximate solution for maximization problems) to Π .

3 W[1] Hardness of EF-ICFD Parameterized by p and $|A|$

To show W[1]-hardness of EF-ICFD parameterized by p and $|A|$, we reduce from (k, M) -VECTOR-SUM, a multidimensional version of the classic NP-hard problem Subset Sum. We begin by formally defining the problem and stating a known result regarding its complexity.

► **Definition 3.1.** For positive integers k, n, M , and d , the (k, M) -VECTOR-SUM problem is defined as follows: given a tuple $(\mathcal{W}, \bar{t}, k)$, where $\mathcal{W} = \{w^{(i)}\}_{i \in [n]} \subseteq [0, kM]^d$, and $\bar{t} \in [0, kM]^d$, determine whether there exists a subset $S \subseteq [n]$ of cardinality k , such that $\sum_{i \in S} w^{(i)} = \bar{t}$.

► **Proposition 3.2** (Lemma 4.1 - [1]). (k, M) -VECTOR-SUM is W[1]-hard parameterized by $k + d$ even when $M = kn^{1+o(1)}$.

With the definition of (k, M) -VECTOR-SUM in place, we are now prepared to describe the reduction.

3.1 The Reduction

Let $\mathcal{I}_{\text{vs}} = (\mathcal{W}, \bar{t}, k)$ be an instance of $(k, n^{1+o(1)})$ -VECTOR-SUM with d dimensions. We define the following instance $\mathcal{I}_{\text{FD}} = (G, A, \mathcal{U}, p)$ of EF-ICFD:

- Let G be a star with center c , and $2d + n$ leaves. We assign a unique label to each leaf using elements of the set $V_\alpha \cup V_\beta \cup \{v_1, v_2, \dots, v_n\}$, where $V_\alpha = \{a_1, a_2, \dots, a_d\}$ and $V_\beta = \{b_1, b_2, \dots, b_d\}$. With a slight abuse of notation, we use the label to refer to the corresponding vertex.
- Let $A = [2d + 1]$. We assign a unique label to each agent using elements of the set $\{A_1, A_2, \dots, A_d\} \cup \{B_1, B_2, \dots, B_d\} \cup \{C\}$. Here as well, we slightly abuse notation, and use the label to refer to the corresponding agent.
- $p = 2d + k + 1$.

- For every $l \in [d]$, we define $u_{A_l}(x) = \begin{cases} n^3 + w_i^{(l)} & x = v_i \\ n^3 & x = c \\ 0 & x \in V_\beta \\ 0 & x \in V_\alpha \setminus \{a_l\} \\ (k+1)n^3 + t_l & x = a_l \end{cases}$
- For every $l \in [d]$, we define $u_{B_l}(x) = \begin{cases} n^3 - w_i^{(l)} & x = v_i \\ n^3 & x = c \\ 0 & x \in V_\alpha \\ 0 & x \in V_\beta \setminus \{b_l\} \\ (k+1)n^3 - t_l & x = b_l \end{cases}$
- $u_C(v) = 1$ for every $v \in V(G)$

3.2 Proof of Correctness

We shall prove that (G, p, A, \mathcal{U}) has a valid and envy-free allocation if and only if $(\mathcal{W}, \bar{t}, k)$ has a subset $S \subseteq [n]$ of cardinality k , such that $\sum_{i \in S} w^{(i)} = \bar{t}$. To this end, we make some observations about the properties that any valid and envy-free allocation in the instance \mathcal{I}_{FD} must satisfy. Note that since the graph is a star and since the bundles must be connected, only the agent who is allocated the center of the star gets more than one vertex. Furthermore, every agent must receive at least one vertex. Otherwise, as each agent has zero valuation for at most $2d - 1$ vertices, and since a total of $2d + k + 1$ vertices must be allocated, any agent with an empty bundle would be envious of some other agent. Hence, for exactly $2d + k + 1$ vertices to be assigned in total, one agent must receive $k + 1$ and the others must receive one vertex each. Now, since C values every vertex equally, they must receive exactly $k + 1$ vertices; otherwise, they would envy the agent who has more than one vertex. With this, we are ready to prove a key property of this instance.

► **Lemma 3.3.** *In any valid and envy-free allocation Π in (G, p, A, \mathcal{U}) , every agent except C must be allocated the vertex they value the most.*

Proof. Let us consider agent C . Observe that, if π_C contains a vertex a_l for some $l \in [d]$, then it cannot be assigned to A_l , since Π is valid. Therefore, irrespective of which vertex is assigned to A_l , we have $u_{A_l}(\pi_C) \geq (k + 1)n^3 > u_{A_l}(\pi_{A_l})$, which contradicts the assumption that Π is envy-free. Hence π_C must not contain any vertex from the set V_α and by a similar argument from V_β . Now observe that, according to any agent X distinct from C , we have $u_X(\pi_C) \geq (k + 1)n^3 - kn^{1+o(1)} \geq kn^3$. Therefore, for X to not be envious of C , we must have $u_X(\pi_X) \geq kn^3$. But this can only be achieved if X is assigned the vertex they value the most. Hence, in the allocation Π , agent A_l must receive a_l and agent B_l must receive b_l for every $l \in [d]$. This concludes the proof of the lemma. ◀

This gives us all the tools we need to prove the main lemma.

► **Lemma 3.4.** *The instance (G, p, A, \mathcal{U}) admits a valid and envy-free allocation Π if and only if there exists a subset $S \subseteq [n]$ of cardinality k such that $\sum_{i \in S} w^{(i)} = \bar{t}$ in the instance $(\mathcal{W}, \bar{t}, k)$.*

Proof. Let Π be a valid and envy-free allocation in (G, p, A, \mathcal{U}) . According to lemma 3.3, in Π , agent A_l and B_l must be allocated a_l and b_l respectively, for every $l \in [d]$. Furthermore, we also have that agent C must be allocated exactly $k + 1$ vertices. Let these vertices be $\{c, v_{i_1}, v_{i_2}, \dots, v_{i_k}\}$. We claim that the set $S = \{i_1, i_2, \dots, i_k\} \subseteq [n]$ satisfies $\sum_{i \in S} w^{(i)} = \bar{t}$ in the instance $(\mathcal{W}, \bar{t}, k)$. Observe that, as Π is envy-free, $u_{A_l}(\pi_C)$ is at most $u_{A_l}(\pi_{A_l})$ and $u_{B_l}(\pi_C)$ is at most $u_{B_l}(\pi_{B_l})$. This implies that,

$$\begin{aligned} n^3 + \sum_{j=1}^k (n^3 + w_l^{(i_j)}) &\leq (k + 1)n^3 + t_l \implies \sum_{j=1}^k w_l^{(i_j)} \leq t_l \\ n^3 + \sum_{j=1}^k (n^3 - w_l^{(i_j)}) &\leq (k + 1)n^3 - t_l \implies \sum_{j=1}^k w_l^{(i_j)} \geq t_l \end{aligned}$$

Hence if (G, p, A, \mathcal{U}) has a valid and envy-free allocation, then there exists a set $S = \{i_1, i_2, \dots, i_k\} \subseteq [n]$ that satisfies $\sum_{i \in S} w^{(i)} = \bar{t}$ in the instance $(\mathcal{W}, \bar{t}, k)$.

Let $S = \{i_1, i_2, \dots, i_k\} \subseteq [n]$ be a set of cardinality k , such that, $\sum_{i \in S} w^{(i)} = \bar{t}$ in the instance $(\mathcal{W}, \bar{t}, k)$. Consider an allocation Π which allocates $\{c, v_{i_1}, v_{i_2}, \dots, v_{i_k}\}$ to C , a_l to A_l , and b_l to B_l for every $l \in [d]$. We claim that Π is a valid and envy-free

allocation in (G, p, A, \mathcal{U}) . Observe that, for every agent $X \neq C$, we have, $u_C(\pi_X) = 1 < (k+1) = u_C(\pi_C)$. Furthermore, for every $l \in [d]$ and agent $X \neq A_l$, we have $u_{A_l}(\pi_X) \leq (k+1)n^3 + \sum_{j=1}^k w_l^{(i_j)} = (k+1)n^3 + t_l = u_{A_l}(\pi_{A_l})$. Finally, for every $l \in [d]$ and agent $X \neq B_l$, we have $u_{B_l}(\pi_X) \leq (k+1)n^3 - \sum_{j=1}^k w_l^{(i_j)} = (k+1)n^3 - t_l = u_{B_l}(\pi_{B_l})$. Consequently, Π is envy-free for the instance (G, p, A, \mathcal{U}) . Since Π allocates a total of p vertices as disjoint, connected subsets to the agents by definition, its validity follows. This completes the proof of the lemma. \blacktriangleleft

By combining Proposition 3.2 and Lemma 3.4 with the properties of the reduction, we obtain Theorem 1.1

► **Theorem 1.1.** ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION is $W[1]$ -hard parameterized by $p + |A|$, with the input numbers provided in unary representation, even on star graphs.

Proof. Given an instance $\mathcal{I}_{\text{VS}} = (\mathcal{W}, \bar{t}, k)$ of $(k, n^{1+o(1)})$ -VECTOR-SUM with d dimensions, let $\mathcal{I}_{\text{FD}} = (G, A, \mathcal{U}, p)$ be an instance of EF-ICFD constructed according to Subsection 3.1. Observe that in our construction G is a star graph, $p = 2d + k + 1$ and $|A| = 2d + 1$. Furthermore, the range of all valuation functions are bounded above by $\mathcal{O}(n^3)$ which implies that the instance is in unary and that $|\mathcal{I}_{\text{FD}}| = n^{\mathcal{O}(1)}$. Finally, given \mathcal{I}_{VS} , we can construct \mathcal{I}_{FD} in $n^{\mathcal{O}(1)}$ time. Therefore, the statement of the theorem follows from Proposition 3.2 and Lemma 3.4. \blacktriangleleft

We conclude by remarking that this hardness proof extends to the mandatory setting without modification, as the properties of the constructed instance ensure that, in any envy-free allocation, every agent is assigned a non-empty set. Consequently, both directions of the proof remain valid even in this context.

4 Approximation Scheme for EF-ICFD Parameterized by p and $|A|$

We begin this section by stating and proving the reduction rule that helps us upper bound the number of agents in an instance of EF-ICFD by a function of p and $|A|$.

► **Lemma 4.1.** *There exists an algorithm that takes as input an instance $\mathcal{I} = (G, A, \mathcal{U}, p)$ of EF-ICFD, runs in $(mn)^{\mathcal{O}(1)}$ time, and outputs an instance $\mathcal{I}' = (G', A', \mathcal{U}', p')$ of EF-ICFD such that \mathcal{I}' admits a valid and envy-free allocation if and only if \mathcal{I} admits one, and furthermore, the number of agents of type \mathbf{a} in \mathcal{I}' is at most $p' + 1$, for every $\mathbf{a} \in \mathcal{A}_{\mathcal{I}'}$.*

Proof. We define $G' := G$, $p' := p$, and $\mathcal{U}' := \mathcal{U}$. To define A' , consider each agent type $\mathbf{a} \in \mathcal{A}_{\mathcal{I}}$. If there are at most $p + 1$ agents of type \mathbf{a} in A , we include all of them in A' . Otherwise, we include an arbitrary selection of $p + 1$ agents of type \mathbf{a} in A' . It is clear from definition of \mathcal{I}' that the number of agents of type \mathbf{a} in \mathcal{I}' is at most $p' + 1$, for every $\mathbf{a} \in \mathcal{A}_{\mathcal{I}'}$. Hence to prove the lemma, it suffice to prove that \mathcal{I}' admits a valid and envy-free allocation if and only if \mathcal{I} admits one.

To this end, let Π' be a valid and envy-free allocation in \mathcal{I}' . We define an allocation Π in \mathcal{I} , where for agent $i \in A$, we have $\pi_i := \pi'_i$ if $i \in A' \cap A$ and $\pi_i := \emptyset$ otherwise. Since the validity of Π follows directly from that of Π' we are left with proving envy-freeness. Note that no agent is envious of agents in $A \setminus A'$, since they are allocated the empty bundle. Furthermore, agents in $A \cap A'$ are not envious of each other since Π' is an envy-free allocation. Finally, agents of $A \setminus A'$ are not envious of agents in $A \cap A'$, since corresponding to an agent $i \in A \setminus A'$, there exists an agent $j \in A \cap A'$ who has the same type and is not envious of anyone in $A \cap A'$. Therefore, we conclude that Π is a valid and envy-free allocation in \mathcal{I} .

Now, let Π be a valid and envy-free allocation in \mathcal{I} . Let $\sigma : A' \rightarrow A$ be an injective function with the property that $\sigma(i)$ and i are of the same type, and every agent in A who is allocated a non-empty bundle under Π lies in the image of σ . Such a function exists because at most p agents receive non-empty bundles under Π , and A' includes all agents from types with at most $p + 1$ agents, as well as at least $p + 1$ agents from types that has more than $p + 1$ agents in A . We define an allocation Π' in \mathcal{I}' by setting $\pi'_i := \pi_{\sigma(i)}$ for each $i \in A'$. Since Π is a valid and envy-free allocation in \mathcal{I} , and Π' is obtained by relabeling agents via an injection that preserves types, it follows that Π' is also valid and envy-free in \mathcal{I}' . This concludes the proof of the lemma. \blacktriangleleft

► **Observation 4.2.** *If $\mathcal{I} = (G, A, \mathcal{U}, p)$ is an instance of EF-ICFD obtained after applying the algorithm of Lemma 4.1, then $|A| \leq (p + 1) \cdot |A|$*

► **Observation 4.3.** *If \mathcal{I}' is the instance of EF-ICFD obtained after applying the algorithm of Lemma 4.1 to \mathcal{I} , then given a valid and envy-free allocation Π' for \mathcal{I}' , we can construct one for \mathcal{I} in $(mn)^{O(1)}$ time.*

Observation 4.2 follows from the fact that each agent is associated with a type in \mathcal{A} , and the lemma guarantees that no type in the output instance has more than $p + 1$ agents. Observation 4.3 follows from the construction described in the proof of Lemma 4.1, which shows how to obtain a valid and envy-free allocation for the original instance from one for the output instance.

Now, we state a lemma that serves as a subroutine in the algorithm. Since it is not central to the main discussion, we defer its proof to the end of the section.

► **Lemma 4.4.** *There exists an algorithm that, given a graph G , a weight function $w : V(G) \rightarrow \mathbb{N}_{\geq 0}$, and a partition $\mathcal{C} = \{C_1, \dots, C_k\}$ of $V(G)$ into k parts, runs in time $3^k \cdot |V(G)|^{O(1)}$ and outputs a connected subgraph H of G that maximizes $\sum_{v \in V(H)} w(v)$ subject to the constraint that $|V(H) \cap C_i| = 1$ for every $i \in [k]$, if such a subgraph exists, and correctly reports that no such subgraph exists otherwise.*

Now, let $\mathcal{I} = (G, A, \mathcal{U}, p)$ be an instance of EF-ICFD and ε be a positive real number. We make the following definitions corresponding to \mathcal{I} and ε . Let $\varepsilon'_\varepsilon > 0$ denote the positive root of the equation $x^2 + 3x = \varepsilon$, and define $q_\varepsilon := 1 + \varepsilon'_\varepsilon$ and $t_\varepsilon := 1 + \left\lceil \log_{q_\varepsilon} \left(\frac{p}{\varepsilon'_\varepsilon} \right) \right\rceil$. Finally, for positive real number ρ and $\alpha \in \mathbb{N}_{\geq 0}$, let $\overline{\log}_\rho \alpha$ denote $\lceil \log \alpha \rceil$, if $\alpha > 0$ and -1 otherwise.

► **Definition 4.5.** *We define a bucket profile to be a tuple $(\mathcal{C}, \mathcal{S}, \Lambda)_{\varepsilon, \mathcal{I}}$, where $\mathcal{C} = \{C_c\}_{c \in [p]}$ is a partition of $V(G)$, $\mathcal{S} = \{S_i\}_{i \in A}$ is a partition of $[p]$, and $\Lambda : A \times [p] \rightarrow [0, t_\varepsilon]$ is a function such that, for every pair of distinct agents $i, j \in A$, we have $\sum_{\substack{c \in S_j \\ \Lambda(i, c) \neq t_\varepsilon}} q_\varepsilon^{-\Lambda(i, c)} \leq q_\varepsilon$.*

Referring back to the high-level sketch of the algorithm in Section 1, we note that the set \mathcal{C} corresponds to the partitioning obtained through color coding, \mathcal{S} captures the guess that maps colors to agents, and Λ encodes the guessed features of the allocation. In particular, $\Lambda(i, c)$ represents the guess on how far apart (in terms of powers of q) the valuation of a vertex in color class c and the valuation of the bundle allocated to agent i is, with respect to u_i , in the valid and envy-free allocation we are searching for. The following definition formalizes what it means for an allocation to be consistent with the guesses.

► **Definition 4.6.** *An allocation Π in \mathcal{I} respects a bucket profile $(\mathcal{C}, \mathcal{S}, \Lambda)_{\varepsilon, \mathcal{I}}$ if it satisfies the following:*

- *For every agent $i \in A$ and $c \in [p]$, we have $|\pi_i \cap C_c| = 1$ if $c \in S_i$, and $|\pi_i \cap C_c| = 0$ otherwise;*

29:10 Beyond Exact Fairness: Envy-Free Incomplete Connected Fair Division

- For every agent $i \in A$, $c \in [p]$, and $v \in \pi \cap C_c$, we have
 - $\overline{\log}_{q_\varepsilon}(u_i(\pi_i)) - \overline{\log}_{q_\varepsilon}(u_i(v)) \geq \Lambda(i, c)$ or $u_i(v) = 0$ if $\Lambda(i, c) = t_\varepsilon$, and,
 - $\overline{\log}_{q_\varepsilon}(u_i(\pi_i)) - \overline{\log}_{q_\varepsilon}(u_i(v)) = \Lambda(i, c)$ and $u_i(v) > 0$ otherwise.

Finally, we formally define the target vector and what it means for an allocation to be consistent with it.

► **Definition 4.7.** Let $M := \max\{u_i(v) \mid v \in V(G), i \in A\}$. We define a target vector $\bar{\mu}_{\varepsilon, \mathcal{I}} \in [-1, \lceil \log_{q_\varepsilon}(pM) \rceil]^A$ and use μ_i to denote the entry of $\bar{\mu}_{\varepsilon, \mathcal{I}}$ associated with agent $i \in A$. An allocation Π attains the target $\bar{\mu}_{\varepsilon, \mathcal{I}}$ for agent $i \in A$, if $\overline{\log}_{q_\varepsilon}(u_i(\pi_i)) = \mu_i$.

For the rest of the discussion, for ease of notation, we drop the subscript \mathcal{I} and ε , whenever it is clear from context.

The following lemma is the main subroutine of the algorithm.

► **Lemma 4.8.** There exists an algorithm that, given an instance (G, A, \mathcal{U}, p) of EF-ICFD, a positive real number ε , a bucket profile $(\mathcal{C}, \mathcal{S}, \Lambda)$, and a target vector $\bar{\mu}$, runs in time $3^p \cdot m^{\mathcal{O}(1)}$ and returns one of the following:

- A valid and ε -envy-free allocation Π , or
- An agent $i \in A$ such that no valid and envy-free allocation that respects the bucket profile $(\mathcal{C}, \mathcal{S}, \Lambda)$ attains a target vector $\bar{\mu}'$, where $\bar{\mu}' \leq \bar{\mu}$ with $\mu'_i = \mu_i$.

Proof. We define vertex $v \in V(G)$ to be *eligible* if, letting $c \in [p]$ be such that $v \in C_c$, for every agent $i \in A$ with $c \notin S_i$, we have $\overline{\log}_q(u_i(v)) \leq \mu_i - \Lambda(i, c)$ or $u_i(v) = 0$. Let G_i be the subgraph of G induced by the eligible vertices in $\bigcup_{c \in S_i} C_c$. We define a subgraph H_i of G_i to be *colorful*, if $|V(H_i) \cap C_c| = 1$ for every $c \in S_i$. Before describing the algorithm, we prove the following claims:

▷ **Claim 4.9.** If for every agent $i \in A$ there exists a connected, colorful subgraph H_i of G_i satisfying the inequality $\overline{\log}_q(u_i(V(H_i))) \geq \mu_i$, then the allocation Π defined by $\pi(i) := V(H_i)$ is a valid and ε -envy-free allocation.

Proof. We begin by observing that $\{\pi_i\}_{i \in A}$ are pairwise disjoint. Otherwise, if there exists a vertex $v \in \pi_i \cap \pi_j$ for some pair of distinct agents $i, j \in A$, then $v \in V(G_i) \cap V(G_j)$ which implies that either v belongs to two distinct sets in \mathcal{C} , which contradicts the fact that \mathcal{C} is a partition of $V(G)$, or the set C_c containing v belongs to both S_i and S_j , which contradicts the fact that \mathcal{S} is a partition of $[p]$. Moreover, we have $\sum_{i \in A} |\pi_i| = \sum_{i \in A} |S_i| = p$, since \mathcal{S} is a partition of $[p]$. Finally, as each $G[\pi_i]$ is connected by definition, it follows that Π is a valid allocation.

Now, to prove ε -envy-freeness, consider $c \in [p]$, agent $i \in A$ such that $c \notin S_i$, and vertex $v \in \pi \cap C_c$. If $u_i(v)$ is positive, then $\overline{\log}_q(u_i(v)) = \lceil \log_q(u_i(v)) \rceil \geq 0$. Moreover, since $\Lambda(i, c)$ is non-negative and v is eligible, it follows that $0 \leq \mu_i \leq \overline{\log}_q(u_i(\pi_i)) = \lceil \log_q(u_i(\pi_i)) \rceil$. Therefore, whenever $u_i(v)$ is positive, we have the following inequality, :

$$\begin{aligned}
 u_i(v) \cdot q^{\Lambda(i, c)} &\leq u_i(v) \cdot q^{\mu_i - \overline{\log}_q(u_i(v))} \\
 &\leq u_i(v) \cdot q^{\log_q(u_i(\pi_i)) + 1 - \log_q(u_i(v))} \\
 &= u_i(\pi_i) \cdot q
 \end{aligned} \tag{1}$$

Thus, if $i, j \in A$ are distinct agents, then,

$$\begin{aligned}
 u_i(\pi_j) &= \sum_{\substack{c \in S_j \\ v \in \pi_j \cap C_c}} u_i(v) \leq \sum_{\substack{c \in S_j \\ \Lambda(i,c)=t}} u_i(\pi_i) \cdot q^{1-\Lambda(i,c)} + \sum_{\substack{c \in S_j \\ \Lambda(i,c) \neq t}} u_i(\pi_i) \cdot q^{1-\Lambda(i,c)} \\
 &\leq p \cdot u_i(\pi_i) \cdot q^{1-t} + q^2 u_i(\pi_i) \\
 &\leq \varepsilon' u_i(\pi_i) + (1 + \varepsilon')^2 u_i(\pi_i) \\
 &\leq (1 + \varepsilon) u_i(\pi_i)
 \end{aligned}$$

where the first transition follows from the fact that π_j contains exactly one vertex from each C_c for $c \in S_j$; the second follows from inequality (1); the third uses the observation that there are at most p elements $c \in S_j$ with $\Lambda(i, c) = t$ and the assumption that $(\mathcal{C}, \mathcal{S}, \Lambda)$ is a bucket profile, which implies $\sum_{\substack{c \in S_j \\ \Lambda(i,c) \neq t}} q^{-\Lambda(i,c)} \leq q$; the fourth and fifth follow from the definitions of q , t , and ε' . Therefore, Π is ε -envy-free, which completes the proof of the claim. \triangleleft

\triangleright **Claim 4.10.** If there exists an agent $i \in A$ such that no connected, colorful subgraph H_i of G_i satisfies $\overline{\log}_q(u_i(V(H_i))) \geq \mu_i$, then there does not exist a valid and envy-free allocation that respects $(\mathcal{C}, \mathcal{S}, \Lambda)$ and attains a target vector $\bar{\mu}'$, where $\bar{\mu}'$ is a target vector that satisfies $\bar{\mu}' \leq \bar{\mu}$ and $\mu'_i = \mu_i$.

Proof. Let $i \in A$ be an agent such that no connected, colorful subgraph H_i of G_i satisfies $\overline{\log}_q(u_i(V(H_i))) \geq \mu_i$. Let $\bar{\mu}'$ be a target vector that satisfies $\bar{\mu}' \leq \bar{\mu}$ and $\mu'_i = \mu_i$. Suppose for contradiction that there exists a valid and envy-free allocation Π that respects $(\mathcal{C}, \mathcal{S}, \Lambda)$, and attains $\bar{\mu}'$.

Since Π is valid, $G[\pi_i]$ is connected. Moreover, as Π attains $\bar{\mu}'$, we have $\overline{\log}_q(u_i(\pi_i)) = \mu'_i = \mu_i$. Furthermore, observe that since Π respects $(\mathcal{C}, \mathcal{S}, \Lambda)$, each vertex in π_i belongs to a distinct set in \mathcal{C} whose indices belong to S_i , which ensures that $G[\pi_i]$ is colorful. Finally, to show that $G[\pi_i]$ is a subgraph of G_i , note that for every $c \in S_i$, $v \in C_c$, and agent $j \in A$ that is distinct from i , if $u_j(v) > 0$ then we have $\overline{\log}_q(u_j(\pi_j)) - \overline{\log}_q(u_j(v)) \geq \Lambda(j, c)$. As Π attains $\bar{\mu}'$, it follows that $\overline{\log}_q(u_j(\pi_j)) = \mu'_j \leq \mu_j$, which implies $\overline{\log}_q(u_j(v)) \leq \mu_j - \Lambda(j, c)$. Hence, every vertex in π_i is eligible, and consequently, $G[\pi_i]$ is a subgraph of G_i . Hence, $G[\pi_i]$ is a connected, colorful subgraph of G_i satisfying $\overline{\log}_q(u_i(V(H_i))) \geq \mu_i$, which contradicts the definition of i . This completes the proof of the claim. \triangleleft

Algorithm. For each agent $i \in A$, the algorithm computes the graph G_i , defined to be the subgraph of G induced by the eligible vertices in $\bigcup_{c \in S_i} C_c$. It calls the algorithm of Lemma 4.4 with input $(G_i, u_i, \{C_c\}_{c \in S_i})$ as a subroutine to compute a maximum-weight connected, colorful subgraph H_i of G_i . If for every agent $i \in A$ a subgraph H_i is found, for which $\overline{\log}_q(u_i(V(H_i))) \geq \mu_i$, then the algorithm returns the allocation Π defined by $\pi(i) := V(H_i)$. Otherwise, it returns an agent $i \in A$ for which no connected, colorful subgraph satisfying $\overline{\log}_q(u_i(V(H_i))) \geq \mu_i$ could be found.

The correctness of the algorithm follows from Claims 4.9 and 4.10, along with the correctness of the subroutine. The running time bound follows from the fact that the subroutine takes $3^p m^{\mathcal{O}(1)}$ time, along with the observation that all other steps of the algorithm run in $m^{\mathcal{O}(1)}$ time. This completes the proof of Lemma 4.8. \blacktriangleleft

With this we are ready to prove Theorem 1.2.

► **Theorem 1.2.** *There exists an algorithm that, given an instance (G, A, \mathcal{U}, p) of ENVY-FREE INCOMPLETE CONNECTED FAIR DIVISION with input numbers in binary representation, and a positive real number ε , runs in time $(\frac{1}{\varepsilon} \log \frac{p}{\varepsilon})^{\mathcal{O}(|A|p^2)} (mn)^{\mathcal{O}(1)}$ and outputs either a valid and ε -envy-free allocation Π , or correctly states that (G, A, \mathcal{U}, p) does not admit a valid and envy-free allocation.*

Proof. If the instance (G, A, \mathcal{U}, p) satisfies $|A| > (p+1)|\mathcal{A}|$, then we apply the algorithm from Lemma 4.1 to obtain a new instance \mathcal{I}' , for which we have $|A_{\mathcal{I}'}| \leq (p+1)|\mathcal{A}|$, by Observation 4.2. We then solve \mathcal{I}' and, using Observation 4.3 and Lemma 4.1, reconstruct a solution for the original instance (G, A, \mathcal{U}, p) , incurring only a negligible overhead of $m^{\mathcal{O}(1)}$ time. Therefore, for the remainder of the proof, we may assume that we are working with an instance that satisfies $|A| \leq (p+1)|\mathcal{A}|$.

Now, let us fix an arbitrary labeling l on $V(G)$ using the set $[m]$. We say that a partition \mathcal{C} of $V(G)$ is induced by a function $f : [m] \rightarrow [p]$, if for every $v \in V(G)$, we have that $v \in C_c$ if and only if $f(l(v)) = c$, where $l(v)$ denotes the label of v . Let \mathcal{B} denote the set of all bucket profiles that belong to the set $\mathcal{F}_{\mathcal{C}} \times \mathcal{F}_{\mathcal{S}} \times \mathcal{F}_{\Lambda}$, where $\mathcal{F}_{\mathcal{C}}$ is the set of partitions of $V(G)$ induced by elements of an (m, p) -perfect hash family of size $e^p p^{\mathcal{O}(\log p)} \log m$ whose existence is guaranteed by Proposition 2.1, $\mathcal{F}_{\mathcal{S}}$ is the set of all partitions of the set $[p]$ into n parts, where some parts may be empty, and \mathcal{F}_{Λ} is the set of all functions from $A \times [p]$ to $[0, t]$. Let $M := \max\{u_i(v) \mid v \in V(G), i \in A\}$, and define a target vector $\bar{\mu}' = (\mu'_1, \dots, \mu'_n)$, where $\mu'_i := \lceil \log_q(p \cdot M) \rceil$ for each $i \in A$.

Algorithm. For each bucket profile in \mathcal{B} , the algorithm initializes the target vector $\bar{\mu}$ to $\bar{\mu}'$. It then invokes the algorithm from Lemma 4.8 as a subroutine. If it returns a valid and ε -envy-free allocation Π , the algorithm outputs Π and terminates. Otherwise, upon receiving an agent $i \in A$, it updates $\bar{\mu}$ by reducing μ_i by one and iterates by invoking the algorithm from Lemma 4.8 again. This loop continues until some entry of $\bar{\mu}$ becomes equal to negative two, at which point the algorithm proceeds to the next bucket profile in \mathcal{B} . If no allocation is found after all elements of \mathcal{B} have been considered, the algorithm concludes that the instance (G, A, \mathcal{U}, p) admits no valid and envy-free allocation.

To prove correctness of the algorithm, assume that (G, A, \mathcal{U}, p) is an instance of EF-ICFD that has a valid and envy-free allocation Π . We only need to show that the algorithm returns an allocation, since by Lemma 4.8, any allocation that is returned by the subroutine and consequently the algorithm will be valid and ε -envy-free. Toward this end, we make the following claim.

► **Claim 4.11.** If (G, A, \mathcal{U}, p) is an instance of EF-ICFD that has a valid and envy-free allocation Π , then there exists a bucket profile $(\mathcal{C}, \mathcal{S}, \Lambda) \in \mathcal{B}$ that is respected by Π .

Proof. Let \mathcal{C} be an element of $\mathcal{F}_{\mathcal{C}}$, such that for every $c \in [p]$, we have that $|C_c \cap \bigcup_{i \in A} \pi_i| = 1$. To see that such an element exists, observe that since Π is a valid allocation, the sets π_i are pairwise disjoint and $\sum_{i \in A} |\pi_i| = p$, which implies that $|\bigcup_{i \in A} \pi_i| = p$. Thus, the set of labels assigned to vertices in $\bigcup_{i \in A} \pi_i$ forms a subset of $[m]$ of size p . By the definition of an (m, p) -perfect hash family, there exists a function that is injective on this set, and this function induces a partition of $V(G)$ satisfying the desired property. Let \mathcal{S} be a partition of $[p]$ such that $c \in [p]$ belongs to S_i if and only if $C_c \cap \pi_i \neq \emptyset$. Finally, for every agent $i \in A$, $c \in [p]$ and $v \in \pi_i \cap C_c$, define, $\Lambda(i, c) = \overline{\log}_q(u_i(\pi_i)) - \overline{\log}_q(u_i(v_c))$ if it is less than t and $u_i(v_c) > 0$, and $\Lambda(i, c) = t$ otherwise. The tuple $(\mathcal{C}, \mathcal{S}, \Lambda)$ is a bucket profile, since for any pair of distinct agents $i, j \in A$, we have:

$$\sum_{\substack{c \in S_j \\ \Lambda(i,c) \neq t}} q^{-\Lambda(i,c)} \leq \sum_{\substack{c \in S_j \\ v \in \pi \cap C_c}} q^{-\lceil \log_q(u_i(\pi_i)) \rceil + \lceil \log_q(u_i(v_c)) \rceil} \leq q \cdot \sum_{v \in \pi_j} \frac{u_i(v)}{u_i(\pi_i)} \leq q$$

Finally, since Π respects the bucket profile $(\mathcal{C}, \mathcal{S}, \Lambda)$ by definition, the lemma follows. \triangleleft

Let $(\mathcal{C}, \mathcal{S}, \Lambda)$ be the bucket profile in \mathcal{B} that is respected by Π , whose existence is guaranteed by Claim 4.11. Let $\bar{\mu}^o$ denote the target vector attained by Π , where $\mu_i^o := \lceil \log_q(u_i(\pi_i)) \rceil$ for each $i \in A$. We claim that, while the algorithm processes the bucket profile $(\mathcal{C}, \mathcal{S}, \Lambda)$, it never considers any target vector $\bar{\mu}$ for which $\mu_i < \mu_i^o$ for some $i \in A$. Suppose, for contradiction, that it does. Let $\bar{\mu}^\alpha$ be the first such target vector considered by the algorithm and $\bar{\mu}^\beta$ be the target vector of the previous iteration. Let $i \in A$ be the agent returned by the subroutine while considering $(\mathcal{C}, \mathcal{S}, \Lambda)$ and $\bar{\mu}^\beta$. By definition of $\bar{\mu}^\alpha$, we have that $\mu_i^\alpha = \mu_i^\beta - 1$ and $\mu_j^\alpha \geq \mu_j^\beta$ for every $j \in A$ such that $j \neq i$, which implies $\bar{\mu}^o \leq \bar{\mu}^\beta$ and $\mu_i^\beta = \mu_i^o$. But, since Π is a valid and envy-free allocation that respects $(\mathcal{C}, \mathcal{S}, \Lambda)$ and attains $\bar{\mu}^o$ for agent i , by the correctness of Lemma 4.8, agent i cannot be returned by the subroutine in the iteration corresponding to $(\mathcal{C}, \mathcal{S}, \Lambda)$ and target vector $\bar{\mu}^\beta$. Thus, we conclude that while the algorithm processes the bucket profile $(\mathcal{C}, \mathcal{S}, \Lambda)$, it never considers any target vector $\bar{\mu}$ for which $\mu_i < \mu_i^o$ for some $i \in A$. Observe that this suffices to prove that the algorithm returns an allocation, since the algorithm begins the iteration corresponding to $(\mathcal{C}, \mathcal{S}, \Lambda)$ with $\bar{\mu} = \bar{\mu}'$, and in every iteration where it does not return an allocation, it decreases μ_i for some $i \in A$. Therefore, if the algorithm never encounters a target vector $\bar{\mu}$ with $\mu_i < \mu_i^o$ for any $i \in A$, it must eventually consider $\bar{\mu} = \bar{\mu}^o$, before reaching a target vector with at least one entry equal to negative two. By the correctness of the subroutine in Lemma 4.8, an allocation must be returned at this point, since Π is a valid and envy-free allocation that respects $(\mathcal{C}, \mathcal{S}, \Lambda)$ and attains $\bar{\mu}^o$ for every agent i . This concludes the proof of correctness of the algorithm.

Finally, the running time bound of $(\frac{1}{\epsilon} \log \frac{p}{\epsilon})^{\mathcal{O}(|A|p^2)} m^{\mathcal{O}(1)}$ follows from the fact that the subroutine runs in $3^p m^{\mathcal{O}(1)}$ time, together with the observation that the algorithm considers at most $e^p p^{\mathcal{O}(\log p)} \log m \cdot (p|A|)^p \cdot t^{|A|p^2}$ bucket profiles in \mathcal{B} , and that the entries of $\bar{\mu}$ can be decreased at most $\mathcal{O}(pm|A|)$ times before at least one of them becomes negative. \blacktriangleleft

We remark that this algorithm extends to the mandatory setting with only minor modification. In particular, it suffices to strengthen the definition of a bucket profile to ensure that the partition \mathcal{S} only contains non-empty sets. The correctness of the algorithm then follows with straightforward changes to account for this additional constraint.

Finally, we conclude the section by restating Lemma 4.4 and describing the algorithm it refers to.

► Lemma 4.4. *There exists an algorithm that, given a graph G , a weight function $w : V(G) \rightarrow \mathbb{N}_{\geq 0}$, and a partition $\mathcal{C} = \{C_1, \dots, C_k\}$ of $V(G)$ into k parts, runs in time $3^k \cdot |V(G)|^{\mathcal{O}(1)}$ and outputs a connected subgraph H of G that maximizes $\sum_{v \in V(H)} w(v)$ subject to the constraint that $|V(H) \cap C_i| = 1$ for every $i \in [k]$, if such a subgraph exists, and correctly reports that no such subgraph exists otherwise.*

Proof. Let H be a subgraph of G . We use the term *weight of H* to mean the quantity $\sum_{v \in V(H)} w(v)$. We describe a dynamic programming algorithm that computes the weight of the maximum-weight connected subgraph H of G that satisfies $|V(H) \cap C_i| = 1$ for every $i \in [k]$, if such a subgraph exists. Using standard backtracking techniques or self-reduction, a subgraph that realizes this value can be recovered with, in the worst case, an additional overhead of $n^{\mathcal{O}(1)}$ in the running time.

First, let us describe a function $f : V(G) \times 2^{[k]} \rightarrow \mathbb{N}_{\geq 0}$, where $2^{[k]}$ denotes the set of all subsets of $[k]$.

$$f(v, S) = \begin{cases} w(v), & \text{if } S = \{i\} \text{ and } v \in C_i, \\ -\infty, & \text{if } v \notin \bigcup_{i \in S} C_i, \\ \max_{\substack{\emptyset \subsetneq S_1 \subsetneq S \\ u \in N[v]}} \{f(u, S_1) + f(v, S \setminus S_1)\}, & \text{otherwise.} \end{cases} \quad (2)$$

It is straightforward to observe that $f(v, S)$ computes the weight of the maximum-weight connected subgraph H of $G[\bigcup_{i \in S} C_i]$ that contains the vertex v . The algorithm constructs a dynamic programming table $F[v, S]$, where $v \in V(G)$ and $S \subseteq [k]$, by filling entries in increasing order of $|S|$ starting with singleton sets. The entries corresponding to larger sets, are computed using the recurrence of function f in Definition (2), after replacing recursive calls to f with lookups to the table F at the corresponding indices. After populating the table states, the algorithm checks the value of $\max_{v \in V(G)} F[v, [k]]$. If it is $-\infty$ the algorithm reports that no such subgraph exists, otherwise it returns the value. The correctness of the algorithm follows from an induction on the size of the subset S in the definition of the table entries. Finally, we observe that the running time is upper bounded by $\sum_{\substack{v \in V(G) \\ S \subseteq [p]}} 2^{|S|} \cdot |V(G)| \leq \sum_{v \in V(G)} \sum_{i \in [k]} \binom{k}{i} \cdot 2^i \cdot 1^{k-i} \cdot |V(G)| = \sum_{v \in V(G)} 3^k \cdot |V(G)| = 3^k |V(G)|^2$, which concludes the proof of the lemma. ◀

5 Concluding Remarks and Discussion

We prove that EF-ICFDis $W[1]$ -hard parameterized by $p + |A|$, even when the input numbers are in unary representation and provide an EPAS parameterized by $p + |A|$ for the problem even when input is in binary representation. We remark that our algorithm outperforms the brute force algorithm running in $\mathcal{O}(m^p \cdot n^p)$ when $p < \mathcal{O}(\log m + \log n)$. This regime remains practically relevant, particularly in scenarios involving the allocation of expensive or limited goods.

Beyond envy-freeness, other fairness notions have been explored in the incomplete and connected setting. Gahlawat and Zehavi [21] investigated three such notions, including EF1 and EFX, which relax envy-freeness by requiring that envy disappears once the most (respectively, least) valuable item is removed from every other agent's bundle. They showed that finding an EF1 (EFX) allocation parameterized by $p + |A|$ is $W[1]$ -hard when the input is given in binary representation. Our results may provide a starting point for answering related questions to these relaxed fairness concepts.

References

- 1 Amir Abboud, Kevin Lewi, and Ryan Williams. Losing weight by gaining edges. In Andreas S. Schulz and Dorothea Wagner, editors, *Algorithms - ESA 2014*, pages 1–12, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg. doi:10.1007/978-3-662-44777-2_1.
- 2 Hannaneh Akrami, Bhaskar Ray Chaudhury, Martin Hoefer, Kurt Mehlhorn, Marco Schmalhofer, Golnoosh Shahkarami, Giovanna Varricchio, Quentin Vermande, and Ernest van Wijland. Maximizing nash social welfare in 2-value instances. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5):4760–4767, June 2022. doi:10.1609/aaai.v36i5.20402.
- 3 Noga Alon, Raphael Yuster, and Uri Zwick. Color-coding. *J. ACM*, 42(4):844–856, July 1995. doi:10.1145/210332.210337.

- 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: Recent progress and open questions. *Artificial Intelligence*, 322:103965, 2023. doi:10.1016/j.artint.2023.103965.
- 5 Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. *Auton. Agents Multi Agent Syst.*, 36(1):3, 2022. doi:10.1007/s10458-021-09532-8.
- 6 Xiaohui Bei, Ayumi Igarashi, Xinhang Lu, and Warut Suksompong. The price of connectivity in fair division. *SIAM journal on Discrete Mathematics*, 36(2):1156–1186, 2022. doi:10.1137/20m1388310.
- 7 Aurélie Beynier, Yann Chevaleyre, Laurent Gourvès, Ararat Harutyunyan, Julien Lesca, Nicolas Maudet, and Anaëlle Wilczynski. Local envy-freeness in house allocation problems. *Autonomous Agents and Multi-Agent Systems*, 33:591–627, 2019. doi:10.1007/s10458-019-09417-x.
- 8 Vittorio Bilò, Ioannis Caragiannis, Michele Flammini, Ayumi Igarashi, Gianpiero Monaco, Dominik Peters, Cosimo Vinci, and William S Zwicker. Almost envy-free allocations with connected bundles. *Games and Economic Behavior*, 131:197–221, 2022. doi:10.1016/j.geb.2021.11.006.
- 9 Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, IJCAI’18, pages 91–97. ijcai.org, 2018. doi:10.24963/ijcai.2018/13.
- 10 Bernhard Bliem, Robert Brederick, and Rolf Niedermeier. Complexity of efficient and envy-free resource allocation: Few agents, resources, or utility levels. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, IJCAI’16, pages 102–108. AAAI Press, 2016. URL: <http://www.ijcai.org/Abstract/16/022>.
- 11 Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. Fair division of a graph. In Carles Sierra, editor, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, IJCAI’17, pages 135–141. ijcai.org, 2017. doi:10.24963/ijcai.2017/20.
- 12 Robert Brederick, Andrzej Kaczmarczyk, Dušan Knop, and Rolf Niedermeier. High-multiplicity fair allocation using parametric integer linear programming. In *ECAI 2023*, pages 303–310. IOS Press, 2023. doi:10.3233/FAIA230284.
- 13 Robert Brederick, Andrzej Kaczmarczyk, Dušan Knop, and Rolf Niedermeier. High-multiplicity fair allocation: Lenstra empowered by n-fold integer programming. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC ’19*, pages 505–523, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3328526.3329649.
- 14 Robert Brederick, Andrzej Kaczmarczyk, and Rolf Niedermeier. Envy-free allocations respecting social networks. *Artificial Intelligence*, 305:103664, 2022. doi:10.1016/j.artint.2022.103664.
- 15 Eric Budish. The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. In *Proceedings of the Behavioral and Quantitative Game Theory: Conference on Future Directions, BQGT ’10*, New York, NY, USA, 2010. Association for Computing Machinery. doi:10.1145/1807406.1807480.
- 16 Ioannis Caragiannis, Evi Micha, and Nisarg Shah. A little charity guarantees fair connected graph partitioning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4908–4916, 2022. doi:10.1609/aaai.v36i5.20420.
- 17 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Daniel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer Publishing Company, Incorporated, 1st edition, 2015. doi:10.1007/978-3-319-21275-3.
- 18 Argyrios Deligkas, Eduard Eiben, Robert Ganian, Thekla Hamm, and Sebastian Ordyniak. The parameterized complexity of connected fair division. In Zhi-Hua Zhou, editor, *Proceedings*

- of the *Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 139–145. International Joint Conferences on Artificial Intelligence Organization, August 2021. Main Track. doi:10.24963/ijcai.2021/20.
- 19 Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate texts in mathematics*. Springer, 5th edition, 2012.
 - 20 Uriel Feige, Ariel Sapir, and Laliv Tauber. A tight negative example for mms fair allocations. In *International Conference on Web and Internet Economics*, pages 355–372. Springer, 2021. doi:10.1007/978-3-030-94676-0_20.
 - 21 Harmender Gahlawat and Meirav Zehavi. Parameterized Complexity of Incomplete Connected Fair Division. In Patricia Bouyer and Srikanth Srinivasan, editors, *43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2023)*, volume 284 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 14:1–14:18, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.FSTTCS.2023.14.
 - 22 Jiarui Gan, Warut Suksompong, and Alexandros A Voudouris. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019. doi:10.1016/j.mathsocsci.2019.07.005.
 - 23 Jugal Garg and Setareh Taki. An improved approximation algorithm for maximin shares. In *Proceedings of the 21st ACM Conference on Economics and Computation, EC '20*, pages 379–380, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3391403.3399526.
 - 24 Jonathan Goldman and Ariel D. Procaccia. Spliddit: unleashing fair division algorithms. *SIGecom Exch.*, 13(2):41–46, January 2014. doi:10.1145/2728732.2728738.
 - 25 Gianluigi Greco and Francesco Scarcello. The complexity of computing maximin share allocations on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2006–2013, 2020. doi:10.1609/aaai.v34i02.5572.
 - 26 Ayumi Igarashi and Dominik Peters. Pareto-optimal allocation of indivisible goods with connectivity constraints. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 2045–2052, 2019. doi:10.1609/aaai.v33i01.33012045.
 - 27 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *Journal of Computer and System Sciences*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.
 - 28 Vincent Lacroix, Cristina G Fernandes, and Marie-France Sagot. Motif search in graphs: application to metabolic networks. *IEEE/ACM transactions on computational biology and bioinformatics*, 3(4):360–368, 2006. doi:10.1109/TCBB.2006.55.
 - 29 Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 125–131, 2004. doi:10.1145/988772.988792.
 - 30 Moni Naor, Leonard J. Schulman, and Aravind Srinivasan. Splitters and near-optimal derandomization. In *36th Annual Symposium on Foundations of Computer Science, FOCS 1995, Milwaukee, Wisconsin, USA, 23-25 October 1995*, FOCS '95, pages 182–191, USA, 1995. IEEE Computer Society. doi:10.1109/SFCS.1995.492475.
 - 31 Trung Thanh Nguyen and Jörg Rothe. Fair and efficient allocation with few agent types, few item types, or small value levels. *Artificial Intelligence*, 314:103820, 2023. doi:10.1016/j.artint.2022.103820.
 - 32 Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2):1039–1068, 2020. doi:10.1137/19M124397X.
 - 33 Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
 - 34 Warut Suksompong. Constraints in fair division. *SIGecom Exch.*, 19(2):46–61, December 2021. doi:10.1145/3505156.3505162.