

The Parameterized Complexity of Connected Fair Division

Argyrios Deligkas¹, Eduard Eiben¹, Robert Ganian², Thekla Hamm², Sebastian Ordyniak³

¹Royal Holloway, University of London, UK

²TU Wien, Austria

³University of Leeds, UK

{argyrios.deligkas,eduard.eiben}@rhul.ac.uk, {rganian, thekla.hamm,sordyniak}@gmail.com

Abstract

We study the Connected Fair Division problem (CFD), which generalizes the fundamental problem of fairly allocating resources to agents by requiring that the items allocated to each agent form a connected subgraph in a provided item graph G . We expand on previous results by providing a comprehensive complexity-theoretic understanding of CFD based on new algorithms and lower bounds while taking into account several well-established notions of fairness: proportionality, envy-freeness, EF1 and EFX. In particular, we show that to achieve tractability, one needs to restrict both the agents and the item graph in a meaningful way. We design XP-algorithms for the problem parameterized by (1) clique-width of G plus the number of agents and (2) treewidth of G plus the number of agent types, along with corresponding lower bounds. Finally, with respect to the restrictions considered here, we show that to achieve fixed-parameter tractability one needs to not only use a more restrictive parameterization of G , but also include the maximum item valuation as an additional parameter.

1 Introduction

The task of allocating a set of indivisible resources among participating agents under a suitably defined notion of “fairness” represents a core research topic in the area of computational social choice [Bouveret and Lang, 2008; Bouveret *et al.*, 2016; Brams and Taylor, 1996]. The classical fair resource allocation problem that arises from this task, FAIR DIVISION, has been extensively studied in the literature, and we now have a fairly good understanding of the problem’s complexity under a variety of notions of fairness [Bouveret and Lang, 2008; Budish, 2011; Caragiannis *et al.*, 2019; Plaut and Roughgarden, 2020]. However, in many settings it is desirable to ensure that the allocation respects a certain *connectivity* restriction on the items allocated to each individual agent—this is easy to substantiate when the items represent immovable assets (e.g., when allocating rooms in a new building to university departments, or in case of corporate split-ups), but can also arise when the items are intangible (e.g., when allocating duties organized in a mind map).

Motivated by these considerations, in 2017 Bouveret *et al.* (2017) developed a natural framework designed to capture this additional aspect of fair resource allocation: given (1) a set of items represented as vertices of an input graph G , (2) a set A of agents, and (3) a set U of valuation functions specifying the utility of each item for each agent, determine if there exists an allocation of items to agents which is (I) *fair* (under a suitable notion of fairness), and (II) maintains the *connectivity* of items assigned to each individual agent. They called this the φ -CONNECTED FAIR DIVISION PROBLEM (φ -CFD), where φ specifies the desired notion of fairness. Their initial results inspired a body of follow-up work on the model involving, e.g., the incorporation of chores [Aziz *et al.*, 2019], computing so-called maximin share allocations [Greco and Scarcello, 2020], Pareto-optimal allocations [Igarashi and Peters, 2019] and conditions guaranteeing the existence of fair allocations [Bilò *et al.*, 2019].

In spite of these developments, our understanding of the precise boundaries of computational tractability of φ -CFD remains somewhat limited to date. On one hand, the problem is a generalization of the classical φ -FAIR DIVISION problem [Bouveret and Lang, 2008; Bouveret *et al.*, 2016], since the problems coincide if G is restricted to the class of complete graphs. But G will not be complete in the typical use case of the model—on the contrary, one may often expect it to be rather sparse, and that is why already the original work introducing the model [Bouveret *et al.*, 2017] initiated the investigation of the problem under various restrictions of the structure of G . Unfortunately, the algorithmic and lower-bound results provided in the original paper as well as in follow-up works [Aziz *et al.*, 2019; Bilò *et al.*, 2019; Igarashi and Peters, 2019] leave a wide gap between the known boundaries of tractability and the easiest intractable cases. This provides a rather unsatisfactory contrast between φ -CFD and other variants of the resource allocation problem, for which we often already have complete complexity landscapes [Bliem *et al.*, 2016a; Bredereck *et al.*, 2018; Eiben *et al.*, 2020].

1.1 Contribution

We perform a comprehensive and in-depth study of the complexity of φ -CFD with the aim of identifying the precise cut-off between tractable and intractable classes of instances. In line with previous work on complexity-theoretic

aspects of computational social choice [Chen *et al.*, 2017; Chen *et al.*, 2018; Eiben *et al.*, 2018b] and especially fair allocation [Bliem *et al.*, 2016a; Eiben *et al.*, 2020], we employ the *parameterized complexity* paradigm to obtain a more fine-grained understanding of the (in-)tractability of φ -CFD under various restrictions. We focus our analysis on combinations of restrictions on two core components of φ -CFD: the structure of the item graph (G) and the set of agents (A). Moreover, all of our results are designed to take into account several perspectives on fairness: *proportionality*, *envy-freeness*, and also two prominent relaxations of envy-freeness called *EF1* and *EFX* [Bild *et al.*, 2019; Budish, 2011; Caragiannis *et al.*, 2019; Plaut and Roughgarden, 2020].

We begin by observing that the problem remains hopelessly intractable even under extreme restrictions to all other aspects of the input if the valuation function U is allowed to use binary encodings (Proposition 1). That is why we proceed—in line with several previous works in the field [Eiben *et al.*, 2020; Greco and Scarcello, 2020]—under the reasonable assumption that U is encoded in unary.

But even under this assumption, we show that the problem remains surprisingly difficult when one does not apply at least some restriction to the set A of agents. Indeed, while it was already known that proportional CFD is polynomial-time tractable on graphs which are stars [Bouveret *et al.*, 2017], we prove that for all other considered notions of fairness CFD remains NP-hard even on stars (Theorem 2). Moreover, proportional CFD remains NP-hard on paths [Bouveret *et al.*, 2017]. On the other hand, restricting the set of agents alone does not lead to tractability either—even if we only consider instances with just two agents (Proposition 3). So in order to achieve tractability, one needs to assume restrictions on G as well as on A ; a single-dimensional approach is not viable.

With the above considerations in mind, we can now proceed to our main algorithmic contributions. First, if we parameterize by the number of agents, then we can solve φ -CFD for all considered fairness notions on an extremely broad range of graphs—not only on trees and more generally graphs of bounded treewidth, but also on dense graph classes such as complete graphs, complete bipartite graphs, cographs, distance-hereditary graphs, and many others. We formalize this by showing that φ -CFD is XP-tractable¹ when parameterized by $|A|$ plus the *clique-width* of G (Theorem 4). Clique-width is among the most general structural parameters of graphs and has found widespread algorithmic applications [Bliem *et al.*, 2016b; Eiben *et al.*, 2018a].

Second, we consider a less severe restriction on A : instead of parameterizing by the number of agents, we parameterize only by the number of *agent types*, which are maximal sets of agents with identical preferences. Agent types have been considered as a middle-ground between unrestricted agents and a bounded number of agents in various contexts [Eiben *et al.*, 2020; Nguyen and Rothe, 2020], and also in the original work that introduced φ -CFD [Bouveret *et al.*, 2017]. As our second algorithmic contribution, we show that φ -CFD is XP-tractable when parameterized by the number of agent

¹A problem is XP-tractable when parameterized by k if it can be solved in polynomial time for every fixed, constant value of k .

types plus the *treewidth* of G (Theorem 6). We justify the use of treewidth instead of clique-width when parameterizing by the number of agent types by providing a corresponding lower bound (Theorem 9). Furthermore, Theorem 6 directly generalizes the results of Bouveret *et al.* (2017), who gave XP-algorithms for proportional and envy-free CFD on paths parameterized by the number of agent types.

Both Theorem 4 and Theorem 6 are obtained by dynamic programming along the respective decompositions, an approach that is often used for these structural parameters. While handling proportional and envy-free allocations in this way is not too difficult, dealing with connectivity in the definitions of EF1 and EFX requires a more careful treatment. In particular, our algorithms for these two variants require ideas that are not needed for, e.g., solving maximin CFD parameterized by treewidth [Greco and Scarcello, 2020].

While both of these algorithmic results are “tight” in the sense that it is not reasonable to hope for a relaxation of the requirements on A and or G , the notion of tractability used here—XP-tractability—is weaker than the one which is typically hoped for in the parameterized setting, notably fixed-parameter tractability. As our next result (Theorem 7), we not only rule out the fixed-parameter tractability of the problem in both considered settings, but even for much more restricted cases: in particular, even when parameterizing by the number of agents plus the *vertex cover number* of G . This essentially rules out the fixed-parameter tractability of the problem under any even remotely reasonable restriction on A and G , for unary valuation functions U ; surprisingly, the problem exhibits the same complexity-theoretic behavior under extreme restrictions ($|A|$ plus vertex cover number) as when using parameterizations that are very relaxed ($|A|$ plus clique-width, or the number of agent types plus treewidth).

But is it really impossible to obtain non-degenerate fixed-parameter algorithms for φ -CFD? The final part of our paper focuses on this question and provides a more positive outlook on the problem’s complexity. First of all, one may observe that Theorem 4 provides fixed-parameter algorithms for envy-free and proportional CFD in the case where $|A|$ is bounded by a fixed constant. But what happens if—instead of imposing extreme restrictions on the number of agents—we also parameterize by the maximum valuation in U ? In Theorem 8 we establish the fixed-parameter tractability of φ -CFD parameterized by the maximum valuation in U , the number of agent types, and the vertex cover number of G . As our last contribution, Theorem 9, we show that this final algorithmic result cannot be strengthened to more general graph parameters such as treewidth (or even deletion distance to stars).

Overall, our results provide a comprehensive understanding of the precise boundaries of tractability of φ -CFD, introduce three novel algorithms for the problem, and show that the problem exhibits a rather diverse complexity-theoretic behavior. An overview is provided in Figure 1.

2 Preliminaries

For an integer i , we let $[i] = \{1, 2, \dots, i\}$. We refer to the handbook by Diestel [2012] for standard graph terminology. We also refer to the standard books for a basic overview of

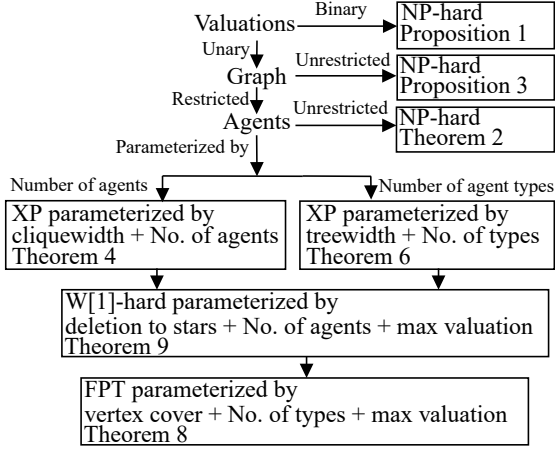


Figure 1: A mind-map of our results.

parameterized complexity theory [Cygan *et al.*, 2015], and assume that readers are aware of the complexity classes FPT, XP and $W[1]$. Readers interested in the full details of Theorem 6 are also expected to have a basic understanding of *treewidth* and *nice tree-decompositions* [Cygan *et al.*, 2015].

Clique-width. Let k be a positive integer. A k -graph is a graph whose vertices are labeled by $[k]$. We denote by $\gamma(v)$ the label of the vertex v . We call the k -graph consisting of exactly one vertex v (say, labeled by α) an initial k -graph and denote it by $\alpha(v)$. The *clique-width* of a graph G is the smallest integer k such that G can be constructed from initial k -graphs by means of repeated application of the following three operations.

1. Disjoint union (denoted \oplus).
2. Relabeling: changing all labels α to β (denoted $p_{\alpha \rightarrow \beta}$).
3. Edge insertion: adding an edge between each vertex labeled α and each vertex labeled β ($\alpha \neq \beta$; denoted $\eta_{\alpha, \beta}$).

A k -expression tree [Courcelle *et al.*, 2000] is a rooted tree representation of how the three operations are used to construct a given graph; specifically, the k -expression tree represents each $\alpha(v)$ as a leaf, each \oplus operator as an \oplus node with two children, and each $p_{\alpha \rightarrow \beta}$ or $\eta_{\beta, \alpha}$ operator by a corresponding node with a single child. It is known that an approximately optimal k -expression tree can be computed in fixed-parameter time [Oum and Seymour, 2006], and that every graph class of bounded treewidth has bounded clique-width.

3 The Connected Fair Division Problem

An instance of the CONNECTED FAIR DIVISION problem (CFD) [Bild *et al.*, 2019; Bouveret *et al.*, 2017] consists of an undirected graph $G = (V, E)$, a set A of n agents, and a valuation function $u_i : V \rightarrow \mathbb{N}$ for every agent $i \in A$. In this setting, every vertex $v \in V$ corresponds to an *item*. We follow the standard assumption in the literature, and we assume that the agents have *additive valuations*; hence for every $X \subseteq V$ and every $i \in A$ we have $u_i(X) = \sum_{v \in X} u_i(v)$. Agents i and j have the *same type* if for every $v \in V$ it holds

$u_i(v) = u_j(v)$. We denote the set of all agent types by \mathcal{A} . The type of an agent i is then the agent type $\mathbf{a} \in \mathcal{A}$ such that $i \in \mathbf{a}$. Hence for an agent type $\mathbf{a} \in \mathcal{A}$ we denote by $u_{\mathbf{a}}$ the valuation function such that $u_{\mathbf{a}} = u_i$ for all $i \in \mathbf{a}$.

An *allocation* of items is a tuple $\pi = (\pi_i)_{i \in A}$ such that (1) for all $i \in A$ the set $\pi_i \subseteq V$ and the subgraph $G[\pi_i]$ of G induced by π_i is connected, (2) $\bigcup_{i \in A} \pi_i = V$, and (3) for all $i, j \in A$ such that $i \neq j$ we have $\pi_i \cap \pi_j = \emptyset$. We also define a *partial allocation* for a subset A' of A as a tuple $\pi = (\pi_i)_{i \in A'}$ satisfying (3). We say that the *bundle* π_i is assigned to agent i . We will focus on *fair* allocations, under several different fairness notions [Bouveret *et al.*, 2017; Bild *et al.*, 2019]. Given a bundle π_i , let $\tau_i = \{v \in \pi_i \mid G[\pi_i \setminus \{v\}] \text{ is connected}\}$. An allocation $\pi_1, \pi_2, \dots, \pi_n$ is:

- *proportional* (PROP) if $u_i(\pi_i) \geq \frac{u_i(V)}{n}$, for every $i \in A$;
- *envy-free* (EF) if $u_i(\pi_i) \geq u_i(\pi_j)$, for every $i, j \in A$;
- *envy-free up to one item* (EF1) if $u_i(\pi_i) \geq u_i(\pi_j) - \max_{v \in \tau_j} u_i(v)$, for every $i, j \in A$;
- *envy-free up to any item* (EFX) if $u_i(\pi_i) \geq u_i(\pi_j) - \min_{v \in \tau_j} u_i(v)$, for every $i, j \in A$;

For a fairness criterion $\varphi \in \mathcal{F} = \{\text{PROP}, \text{EF}, \text{EF1}, \text{EFX}\}$, φ -CFD asks whether there exists an allocation satisfying φ . We remark that here we consider the decision version of the problem for formal reasons only; each of our algorithms can also output a suitable allocation if one exists.

As our first course of action, we show that even extremely restricted instances of φ -CFD remain intractable if the valuation functions u_i are encoded in binary.

Proposition 1. *For all $\varphi \in \mathcal{F}$, φ -CFD is NP-hard when valuation functions are encoded in binary even when restricted to instances having treewidth at most 2 and where $|A| = 2$, and $u_1 = u_2$.*

In view of Proposition 1, from now on we assume every valuation to be encoded in unary (which also implies that all numbers are upper-bounded by the size of the instance).

4 The Futility of Singular Restrictions

Since our aim is to determine when exactly φ -CFD can be solved efficiently, the most natural thing to do is consider whether the problem becomes tractable if either the graph or the set of agents is restricted. In this section, we show that such a single-dimensional approach does not lead to tractability even on highly restricted classes of instances. Note that NP-hardness of EF-CFD on stars was already shown in [Bouveret *et al.*, 2017] using a similar idea, albeit without the restriction on the valuations.

Theorem 2. *For all $\varphi \in \mathcal{F} \setminus \{\text{PROP}\}$, φ -CFD is NP-hard even when restricted to instances in which G is a star, with binary valuations.*

Proof Sketch. For each of the three fairness notions, we give a reduction from INDEPENDENT SET on 4-regular graphs. Let (H, k) be a 4-regular instance of INDEPENDENT SET, where H has n vertices and accordingly $m = 2n$ edges, the vertices of H are indexed as $\{v_1, \dots, v_n\}$, and the edges of H are indexed as $\{e_1, \dots, e_m\}$. We construct an instance of CFD

as follows. Define G to be a star with $m + k$ leaves. We denote the center of the star as c , and call the first n leaves of the star *vertex leaves*, $X = \{x_1, \dots, x_n\}$, and the remaining leaves of the star *dummy leaves*, $D = \{d_1, \dots, d_{m-n+k}\}$. Let $A = [m + 1]$. To complete the reduction, it suffices to define the valuation functions; in this exposition we present only the function for $\varphi = \text{EF}$. Let $u_{m+1}(y) = 1$ if $y = c$ and 0 otherwise, and for each other $i \in [m]$ let $u_i(y) = 1$ if $y \in D \cup \{v_j \mid v_j \in e_i\}$ and 0 otherwise.

Finally, we provide some intuition to help verify the reduction. Consider an independent set W of size k in H . For an arbitrary rooted spanning tree of H , denote for each $v \in V(H)$ the unique edge $e \in E(H)$ which corresponds to an arc toward v in the rooted spanning tree by e_v . Let $\{f_1, \dots, f_{m-n+1+k}\}$ be an enumeration of all edges in $E(H) \setminus \{e_v \mid v \in W\}$. This gives rise to a connected resource allocation π_1, \dots, π_{m+1} for the constructed instance by letting $\pi_{m+1} = \{c\} \cup \{x_i \mid v_i \in W\}$, and for $i \in [m]$

$$\pi_i = \begin{cases} \{x_j\} & \text{if } e_i = e_{v_j} \text{ for } v_j \in V(H) \setminus W \\ \{d_j\} & \text{if } e_i = f_j. \quad \square \end{cases}$$

Our second hardness result is considerably easier, and is obtained by a reduction from a problem called **EQUITABLE CONNECTED PARTITION (ECP)** [Enciso *et al.*, 2009].

Proposition 3. *For every $\varphi \in \mathcal{F}$, φ -CFD is NP-hard when restricted to instances where $|A| = 2$, and $u_1 = u_2 = 1$.*

5 Solving Instances with Few Agents

While Proposition 3 rules out the existence of efficient algorithms for all instances with bounded n , here we show that the situation changes as long as we parameterize by n plus the clique-width of G , denoted $\text{cw}(G)$.

Theorem 4. *For every $\varphi \in \mathcal{F}$, φ -CFD is in XP parameterized by the clique-width and the number of agents.*

Proof Sketch. It is known that an approximate k -expression tree with at most $k^2|V(G)|$ nodes can be computed in fixed-parameter time [Kanté and Rao, 2013; Oum and Seymour, 2006], so it suffices to decide φ -CFD when such a k -expression tree T with $k = f(\text{cw}(G))$ for G is provided.

Let $\mathcal{I} = (G, A, (u_i)_{i \in A})$ be an instance of φ -CFD for $\varphi \in \mathcal{F}$. Let t be a node of T , and recall that t is one of the following four types of nodes: $\alpha(v)$, \oplus , $\eta_{\alpha, \beta}$ or $p_{\alpha \rightarrow \beta}$. Let T_t be the subtree of T rooted at t , and let G_t be the k -graph, which is a subgraph of G , defined by the k -expression tree T_t . For instance, if r is the root of T then $G_r = G$, and for each leaf t in T it holds that G_t is a graph with a single labeled vertex.

The high-level idea of the algorithm is to compute a set of records for every node $t \in V(T)$ and to compute these records in a leaf-to-root fashion. Each record will represent one way an allocation of G can intersect with the vertices of G_t , and the information contained in the records at the root of T is sufficient to determine which of the allocations are connected and satisfy the considered fairness notion.

Formally for each node t of T we define a *record* $((\rho^i)_{i \in A}, (\mu^i)_{i \in A}, (\nu^i)_{i \in A})$ consisting for each $i \in A$ of the following components.

(1) $\rho^i \in 2^{2^{[k]}} \times \{0, 1\}$, such that for each partial allocation $\tilde{\pi}$ corresponding to $((\rho^i)_{i \in A}, (\mu^i)_{i \in A}, (\nu^i)_{i \in A})$ the set $\rho^i[0]$ contains $\mathcal{L} \subseteq [k]$ if and only if $G_t[\tilde{\pi}_i]$ contains a connected component labeled precisely by labels in \mathcal{L} . Formally, $\rho^i[0] = \{\bigcup_{w \in C} \{\lambda(w)\} \mid C \text{ is a connected component of } \tilde{\pi}_i\}$. The number $\rho^i[1] = 1$ if and only if $\tilde{\pi}_i$ is connected.

(2) $\mu^i \in \mathbb{Z}^n$ is a vector of n integers, in which μ_j^i describes the additive valuation under u_j of all items assigned to i by any partial allocation corresponding to $((\rho^i)_{i \in A}, (\mu^i)_{i \in A}, (\nu^i)_{i \in A})$.

(3) ν^i is only included when $\varphi \in \{\text{EF1}, \text{EFX}\}$. It then for each agent $j \neq i$ describes the valuation under u_j of an item v which would be omitted from the bundle of agent i in order to determine whether j envies i in any partial allocation $\tilde{\pi}$ corresponding to $((\rho^i)_{i \in A}, (\mu^i)_{i \in A}, (\nu^i)_{i \in A})$. This is necessary to ensure that connectedness is maintained when removing v from the bundle of i in the root of T . For this purpose $\nu^i : A \times 2^{2^{[k]}} \times \{0, 1\} \rightarrow \mathbb{Z}$ is a partial function for which $\nu^i(j, \mathcal{K}, x)$ describes the valuation under u_j of an item $v \in \tilde{\pi}_i$ that j would remove from $\tilde{\pi}_i$ such that $\mathcal{K} = \{\bigcup_{w \in C} \{\lambda(w)\} \mid C \text{ is a connected component of } \tilde{\pi}_i \setminus \{v\}\}$ and $\tilde{\pi}_i \setminus \{v\}$ is connected if and only if $x = 1$, where $\tilde{\pi}$ is any partial allocation $\tilde{\pi}$ corresponding to $((\rho^i)_{i \in A}, (\mu^i)_{i \in A}, (\nu^i)_{i \in A})$. For a technical reason, if $\tilde{\pi}_i = \{v\}$, then $\nu^i(j, \{\emptyset\}, 1) = u_j(v)$.

Note that the number of records at each node is easily bounded by $2^{(2^k+1)n} \cdot |\mathcal{I}|^{n(1+n2^{2^k+1})}$ for $\varphi \in \{\text{EF1}, \text{EFX}\}$ and even $2^{(2^k+1)n} \cdot |\mathcal{I}|^n$ for $\varphi \in \{\text{PROP}, \text{EF}\}$, which is XP parameterized by $\text{cw}(G)$ and n . To unify the dealing with EF1 and EFX, we let $\text{obj} = \min$ if $\varphi = \text{EFX}$ and $\text{obj} = \max$ if $\varphi = \text{EF1}$. We say a record $((\rho^i)_{i \in A}, (\mu^i)_{i \in A}, (\nu^i)_{i \in A})$ is *valid* whenever there is an allocation π of $V(G_t)$ to A such that for every $i \in A$, all of the following hold.

- (R1) $\rho^i[0] = \{\bigcup_{w \in C} \{\lambda(w)\} \mid C \text{ is a connected component of } \pi_i\}$ and $\rho^i[1] = 1$ if and only if π_i is connected.
- (R2) For all $j \in [n]$, $\mu_j^i = \sum_{v \in \pi_i} u_j(v)$.
- (R3) For all $j \in A$ and all $\mathcal{K} \subseteq 2^{[k]}$, it holds that $\nu^i(j, \mathcal{K}, 0) = \text{obj}(\{u_j(v) \mid v \in \pi_i \wedge \mathcal{K} = \{\bigcup_{w \in C} \{\lambda(w)\} \mid C \neq \pi_i \setminus \{v\} \text{ is a connected component of } \pi_i \setminus \{v\}\}\})$. Furthermore, for all $j \in A$ and all $\Lambda \subseteq 2^{[k]}$, it holds that $\nu^i(j, \{\Lambda\}, 1) = \text{obj}(\{u_j(v) \mid v \in \pi_i \wedge (\pi_i \setminus \{v\} \text{ is a connected}) \wedge \Lambda = \bigcup_{w \in \pi_i \setminus \{v\}} \{\lambda(w)\}\})$ and $\nu^i(j, \mathcal{K}, 1) = \text{undefined}$ otherwise.

For a node $t \in V(T)$ we denote by $\mathcal{R}(t)$ the set of all valid records for t . With this setup there is a φ -CFD for $(G = (V, G), A, (u_i)_{i \in A})$ if and only if at the root r of T , $\mathcal{R}(r)$ contains a valid record $((\rho^i)_{i \in A}, (\mu^i)_{i \in A}, (\nu^i)_{i \in A})$ such that all of the following hold.

- (R*1) For all $i \in A$ and for all $j \in A$: (a) $\mu_i^i \geq \frac{\sum_{v \in V} u_i(v)}{n}$ if we are considering the fairness notion $\varphi = \text{PROP}$; (b) $\mu_i^i \geq \mu_j^j$ if we are considering the fairness notion $\varphi = \text{EF}$; (c) $\mu_i^i \geq \mu_j^j - \text{obj}(\{\nu^j(i, \{\Lambda\}, 1) \mid \Lambda \subseteq 2^{[k]}\})$ if we are considering a fairness notion $\varphi \in \{\text{EF1}, \text{EFX}\}$. This condition ensures that the allocation satisfies the desired fairness notion.

(R*2) For all $i \in A$, $\rho^i = (\{\Lambda\}, 1)$ for some $\Lambda \in 2^{[k]}$. This condition ensures that the allocation is connected.

Both of these conditions can be checked in $\mathcal{O}^*(2^k)$ -time for each record. Thus, to complete the proof it suffices to show how to compute the set of valid records $\mathcal{R}(t)$ at each node t of T , given the sets of valid records of the children of t in T .

This can be done by giving a construction for each case of t .

Using this we can compute the set of all valid records $\mathcal{R}(t)$ for every node t of T using a leaves-to-root dynamic program, achieving a final runtime of $|\mathcal{I}|^{\mathcal{O}(n2^{2^k})}$ for EF1 and EFX and of $2^{\mathcal{O}(2^k n)} \cdot |\mathcal{I}|^{\mathcal{O}(n)}$ for envy-free and proportional CFD. \square

Corollary 5. *If $|A|$ is bounded by a constant, envy-free and proportional CFD is FPT parameterized by clique-width.*

6 Instances with Few Agent Types

Given the state of the art, the obvious question that arises from Theorem 4 is whether we can relax the parameterization by n to a less restrictive parameterization, notably to the number of agent types. We show later in Theorem 9 that this is not possible, i.e., for each $\varphi \in \mathcal{F}$, φ -CFD is NP-hard even for one agent type and clique-width three.

On the other hand, as the main result for this section, we show that offsetting the “less restrictive” agent type parameter by a “more restrictive” graph parameter—notably treewidth—gives rise to an XP-algorithm for the problem.

Theorem 6. *For each $\varphi \in \mathcal{F}$, φ -CFD is in XP when parameterized by the treewidth and the number of agent types.*

Proof Sketch. We show the theorem using a bottom-up dynamic programming algorithm along a tree decomposition (T, χ) of the graph G . As usual the main challenge lies in the definition of the records that we keep for each node t of T , where each record models an equivalence class of partial solutions for the sub-instance induced by the items in G_t , i.e., the graph G induced by all items contained in bags in the subtree rooted at t . To illustrate the main ideas behind the definition of our records, consider a solution, i.e., an allocation $\pi = (\pi_i)_{i \in A}$ for the whole instance $\mathcal{I} = (G, A, (u_i)_{i \in A})$ of φ -CFD for $\varphi \in \mathcal{F}$. Then, our records need to keep information about all bundles that contain at least some items from G_t . There are two main types of those bundles: (1) bundles that do not contain any items from the current bag $\chi(t)$, and (2) bundles that do contain some item from $\chi(t)$. Note that the former bundles must be completely contained in G_t (because the tree decomposition does not allow any edges from $G_t - \chi(t)$ to $G - G_t$), while the latter bundles can still contain items from $G - G_t$ and therefore do not even need to be connected inside G_t . Now, let A_t be the set of all agents being assigned a bundle of type (2), let A_t^F be the set of all agents being assigned a bundle of type (1), and let A_t^I be the set of all other agents, i.e., all agents whose bundles do not intersect with G_t . For the agents in A_t^F , remember the following.

- The number of agents in A_t^F of type α for every $\alpha \in \mathcal{A}$. This is important to ensure that we do not falsely assume more agents of a certain type than are present in the instance.

- We need to ensure that no agent in A_t^F envies an agent in $A_t \cup A_t^I$, i.e., any agent that will be assigned items at a later stage of the algorithm. For this it suffices to remember the minimum value w.r.t. u_α of any bundle assigned to an agent in A_t^F of type α for every $\alpha \in \mathcal{A}$. Note that this is only necessary if $\varphi \in \{\text{EF}, \text{EF1}, \text{EFX}\}$ and can be skipped if $\varphi = \text{PROP}$.
- We need to ensure that no agent in $A_t \cup A_t^I$ envies an agent in A_t^F . Depending on $\varphi \in \{\text{EF}, \text{EF1}, \text{EFX}\}$, it will be sufficient to remember the following for every agent type $\alpha \in \mathcal{A}$: (1) the maximum value w.r.t. u_α of any bundle assigned to an agent in A_t^F (if $\varphi = \text{EF}$), and (2) the maximum value w.r.t. u_α of any bundle assigned to an agent in $i \in A_t^F$ after removing an item in π_i^I of maximum/minimum value w.r.t. u_α (if φ is EF1/EFX).

While the above information for the agents in A_t^F is still fairly straightforward, the real challenge arises for the agents in A_t , since we additionally have to take into account the connectivity of their bundles and in the case of EF1 and EFX even the connectivity of their bundles after removing an item. Consider an agent $i \in A_t$. Then, $G_t[\pi_i]$ is not necessarily connected (it might only become connected because of the items in $\pi_i \setminus G_t$), which means that we have to remember not only which items of π_i are in $\chi(t)$ but also the partition of those items into components of $G_t[\pi_i]$. In the case of $\varphi \notin \{\text{EF1}, \text{EFX}\}$, this together with the value of the partial bundle $\pi_i \cap G_t$ w.r.t. u_α for every agent type $\alpha \in \mathcal{A}$ would actually be sufficient information. However, if $\varphi \in \{\text{EF1}, \text{EFX}\}$ this does not suffice since in order to ensure that no agent in $A_t^F \cup A_t^I$ will envy i , we also need to know the minimum/maximum value of any item that could potentially be removed from the bundle π_i . Since whether or not an item in $\pi_i \cap G_t$ can be removed can really only be determined once we know the whole bundle π_i , we need to store the minimum/maximum value of any item whose removal results in a certain refinement of the current partition of $\pi_i \cap \chi(t)$ into components of $G_t[\pi_i]$. In other words, for every “refinement” of the current partition of $\pi_i \cap \chi(t)$ into components of $G_t[\pi_i]$, we store the minimum/maximum value of any item r whose removal results in this refinement (defined by $G_t[\pi_i \setminus \{r\}]$). \square

7 Towards Fixed-Parameter Tractability

From the perspective of parameterized complexity, the most obvious question that arises from Theorem 4 and 6 is whether one can obtain fixed-parameter algorithms under these parameterizations. We firmly answer this question by providing an involved reduction that establishes the W[1]-hardness of φ -CFD even under significantly stronger restrictions than those required by either of the two algorithms.

Theorem 7. *For each $\varphi \in \mathcal{F}$, φ -CFD is W[1]-hard parameterized by the vertex cover number and the number of agents, even when all agents have identical valuations.*

Proof Sketch. We provide reductions from the W[1]-hard UNARY BIN PACKING problem [Jansen *et al.*, 2013]: given a set of items with unary values and a set of k bins, determine if it is possible to allocate items to bins so that each bin receives

the same total value, which we call B . On a high-level, the reduction constructs an instance of φ -CFD consisting of k (or, for technical reasons that arise in the case of EF1, $2k$) agents and a complete bipartite graph with k bin-vertices on one side each representing one bin, and item-vertices on the other side each representing one item. We then attach one (or, in the case of EF1, two) leaves to the bin-vertices and set the valuations in a way which ensures that every fair allocation in the constructed CFD instance will force k agents to each receive one bin-vertex along with items whose values sum up precisely to B . This yields a correspondence between solutions to the φ -CFD instance and solutions to the original UNARY BIN PACKING instance, completing the reduction. \square

But given Theorem 7, is it really impossible to obtain fixed-parameter algorithms for φ -CFD? Corollary 5 noted that if the number of agents is bounded by a constant, proportional as well as envy-free CFD becomes fixed-parameter tractable parameterized by clique-width alone. However, if we do not wish to restrict the agents or their types in this manner, the problem can be shown to be fixed-parameter tractable if we additionally restrict the number of valuations of a vertex, i.e., $\text{val} = |\{u_i(v) \mid i \in [n], v \in V\}|$.

Theorem 8. *For each $\varphi \in \mathcal{F}$, φ -CFD is fixed-parameter tractable when parameterized by the vertex cover number, the number of agent types, and val .*

Proof Sketch. Let $(G = (V, E), A, (u_i)_{i \in A})$ be an instance of φ -CFD. We compute a minimum size vertex cover S of G in time $\mathcal{O}(2^{\text{vcn}(G)} \cdot |V|)$ [Cygan *et al.*, 2015].

Since $N_G(v) \subseteq S$ for each $v \in V(G) \setminus S$, there are at most $2^{|\text{val}|}$ equivalent items where for $v, w \in V(G) \setminus S$, $v \sim w$ if and only if $N_G(v) = N_G(w)$ and for all $a \in A$ every agent $a \in a$ values v and w equally. An equivalence class under this relation is completely determined by its neighborhood in G and its valuations for each agent type. Now, the algorithm proceeds in two steps. In the first step we branch on a possible structure of the sought solution, as will be described below.

For each leaf of the branching tree, we then construct an instance of ILP with a number of variables bounded in our parameters. Such an instance of ILP can then be solved in FPT time parameterized by the number of variables by Lenstra's well-known algorithm [Lenstra Jr., 1983].

In the branching step we first branch on which vertices of S are allocated to which agent (up to identical agent types), and denote the set of agents that receive at least one item of S as A_S . Moreover, for each $S' \subseteq S$ and each possible valuation q for agent types in $\text{val}^{|A|}$, we branch on whether an agent in $A \setminus A_S$ is allocated an item in the equivalence class determined by S' and q . For each $a \in A_S$ we branch on whether a receives zero, one, or more items in the equivalence class determined by S' and q . For each agent type we also branch on whether any agent of that type receives no items.

At this point, it is not difficult to ensure that each pursued branch remains connected. To complete the allocation, we use the aforementioned ILP formulation to ensure (1) the exact number of items of each equivalence class assigned to each agent $a \in A_S$, when a is branched to receive at least two such items and (2) the exact number of agents in $A \setminus A_S$

of each agent type that receive an item in a specific equivalence class, if it is branched that there is at least one agent of the agent type receiving an item in that equivalence class. \square

As our last result, we show that even using val as an additional parameter does not allow us to lift fixed-parameter tractability to structural parameters such as treewidth.

Theorem 9. *For each $\varphi \in \mathcal{F}$, φ -CFD is W[1]-hard when parameterized by the deletion distance to stars and the number of agents even when $\text{val} = 1$. Moreover, φ -CFD is NP-hard even when $\text{val} = 1$ and the clique-width is at most 3.*

Proof Sketch. We prove our result via a reduction from UNARY BIN PACKING (UBP). Let $(I, (s_i)_{i \in I}, k, B)$ be an instance of UNARY BIN PACKING such that $\sum_{i \in I} s_i = k \cdot B$. We construct an instance $(G = (V, E), A, (u_i)_{i \in A})$ of CFD such that for all $i \in A$ and all $v \in V$ it holds $u_i(v) = 1$ and let $u = u_i$ for all $i \in A$. We let $|A| = k$. For each bin $j \in [k]$, we create a bin vertex b_j . For each item $i \in I$, we create s_i many vertices: one "center" vertex c_i and $s_i - 1$ many "pendant" vertices $p_i^1, p_i^2, \dots, p_i^{s_i-1}$ and we denote $S_i = \{c_i, p_i^1, \dots, p_i^{s_i-1}\}$. The center vertex c_i is adjacent to all pendant vertices $p_i^1, \dots, p_i^{s_i-1}$ and all bin vertices b_1, \dots, b_k . Note that each connected component of $G - \{b_j \mid j \in [k]\}$ is a star of size s_i for some $i \in I$.

Given a solution (I_1, I_2, \dots, I_k) for UBP, we let $\pi_i = \{b_i\} \cup \bigcup_{j \in I_i} S_j$ it is easy to see that $u(\pi_i) = 1 + B$ and the allocation is fair. On the other hand, given a fair allocation π , we can observe that, since $|V| = k \cdot (B + 1)$, π has to be also PROP. To finish the proof, we show that every agent has to get exactly one bin vertex and if an agent gets the center vertex c_i , then she gets all vertices in the star S_i . \square

8 Concluding Remarks

Interestingly, it turns out that the complexity-theoretic behavior of CFD seems to exhibit very little variation between the different considered notions of fairness. It is also surprising that while fairly unrestrictive parameterizations suffice to guarantee XP-tractability (as showcased by Theorems 4 and 6), disproportionately stronger restrictions are required to achieve fixed-parameter tractability (see Theorems 7 and 8). We remark that all of our algorithms can also be straightforwardly extended to handle chores (i.e., negative evaluations).

While our results provide a detailed understanding of the complexity of CFD for several of the most prominent notions of fairness studied in the literature, there remain open questions that deserve further attention. One such question is that it is currently not known whether EF1 allocations always exist on path graphs [Bilò *et al.*, 2019].

Acknowledgments

Ganian and Hamm gracefully acknowledge support from the Austrian Science Fund (FWF, Projects P31336 and Y1329). Hamm also acknowledges support from FWF (Project W1255-N23). Ordyniak acknowledges the support by the EPSRC (EP/V00252X/1).

References

- [Aziz *et al.*, 2019] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. In *IJCAI 2019*, pages 53–59. ijcai.org, 2019.
- [Bilò *et al.*, 2019] 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. In *ITCS 2019*, pages 14:1–14:21. Dagstuhl, 2019.
- [Bliem *et al.*, 2016a] Bernhard Bliem, Robert Brederbeck, and Rolf Niedermeier. Complexity of efficient and envy-free resource allocation: Few agents, resources, or utility levels. In *IJCAI 2016*, pages 102–108. IJCAI/AAAI Press, 2016.
- [Bliem *et al.*, 2016b] Bernhard Bliem, Sebastian Ordyniak, and Stefan Woltran. Clique-width and directed width measures for answer-set programming. In *ECAI 2016*, pages 1105–1113. IOS Press, 2016.
- [Bouveret and Lang, 2008] Sylvain Bouveret and Jérôme Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *J. Artif. Intell. Res.*, 32:525–564, 2008.
- [Bouveret *et al.*, 2016] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In *Handbook of Computational Social Choice*, pages 284–310. Cambridge University Press, 2016.
- [Bouveret *et al.*, 2017] Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. Fair division of a graph. In *IJCAI 2017*, pages 135–141. ijcai.org, 2017.
- [Brams and Taylor, 1996] Steven J Brams and Alan D Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, 1996.
- [Brederbeck *et al.*, 2018] Robert Brederbeck, Andrzej Kaczmarczyk, and Rolf Niedermeier. Envy-free allocations respecting social networks. In *AAMAS 2018*, pages 283–291, 2018.
- [Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2019] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.
- [Chen *et al.*, 2017] Jiehua Chen, Piotr Faliszewski, Rolf Niedermeier, and Nimrod Talmon. Elections with few voters: Candidate control can be easy. *J. Artif. Intell. Res.*, 60:937–1002, 2017.
- [Chen *et al.*, 2018] Jiehua Chen, Rolf Niedermeier, and Piotr Skowron. Stable marriage with multi-modal preferences. In *EC 2018*, pages 269–286. ACM, 2018.
- [Courcelle *et al.*, 2000] Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000.
- [Cygan *et al.*, 2015] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [Diestel, 2012] Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.
- [Eiben *et al.*, 2018a] Eduard Eiben, Robert Ganian, Dusan Knop, and Sebastian Ordyniak. Unary integer linear programming with structural restrictions. In *IJCAI 2018*, pages 1284–1290. ijcai.org, 2018.
- [Eiben *et al.*, 2018b] Eduard Eiben, Robert Ganian, and Sebastian Ordyniak. A structural approach to activity selection. In *IJCAI 2018*, pages 203–209. ijcai.org, 2018.
- [Eiben *et al.*, 2020] Eduard Eiben, Robert Ganian, Thekla Hamm, and Sebastian Ordyniak. Parameterized complexity of envy-free resource allocation in social networks. In *AAAI 2020*, pages 7135–7142. AAAI Press, 2020.
- [Enciso *et al.*, 2009] Rosa Enciso, Michael R. Fellows, Jiong Guo, Iyad A. Kanj, Frances A. Rosamond, and Ondrej Suchý. What makes equitable connected partition easy. In *IWPEC 2009*, pages 122–133. Springer, 2009.
- [Greco and Scarcello, 2020] Gianluigi Greco and Francesco Scarcello. The complexity of computing maximin share allocations on graphs. In *AAAI 2020*, pages 2006–2013. AAAI Press, 2020.
- [Igarashi and Peters, 2019] Ayumi Igarashi and Dominik Peters. Pareto-optimal allocation of indivisible goods with connectivity constraints. In *AAAI 2019*, pages 2045–2052. AAAI Press, 2019.
- [Jansen *et al.*, 2013] Klaus Jansen, Stefan Kratsch, Dániel Marx, and Ildikó Schlotter. Bin packing with fixed number of bins revisited. *J. Comput. Syst. Sci.*, 79(1):39–49, 2013.
- [Kanté and Rao, 2013] Mamadou Moustapha Kanté and Michaël Rao. The rank-width of edge-coloured graphs. *Theory Comput. Syst.*, 52(4):599–644, 2013.
- [Lenstra Jr., 1983] Hendrik W. Lenstra Jr. Integer programming with a fixed number of variables. *Math. Oper. Res.*, 8(4):538–548, 1983.
- [Nguyen and Rothe, 2020] Trung Thanh Nguyen and Jörg Rothe. Approximate Pareto set for fair and efficient allocation: Few agent types or few resource types. In *IJCAI 2020*, pages 290–296. ijcai.org, 2020.
- [Oum and Seymour, 2006] Sang-il Oum and Paul D. Seymour. Approximating clique-width and branch-width. *J. Combin. Theory Ser. B*, 96(4):514–528, 2006.
- [Plaut and Roughgarden, 2020] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. *SIAM J. Discret. Math.*, 34(2):1039–1068, 2020.