



# On Fair Division with Binary Valuations Respecting Social Networks

Neeldhara Misra<sup>(✉)</sup> and Debanuj Nayak

Indian Institute of Technology, Gandhinagar, India  
neeldhara.m@iitgn.ac.in, debanuj.nayak@alumni.iitgn.ac.in

**Abstract.** We study the computational complexity of finding fair allocations of indivisible goods in the setting where a social network on the agents is given. Notions of fairness in this context are “localized”, that is, agents are only concerned about the bundles allocated to their neighbors, rather than every other agent in the system. We comprehensively address the computational complexity of finding locally envy-free and Pareto efficient allocations in the setting where the agents have binary valuations for the goods and the underlying social network is modeled by an undirected graph. We study the problem in the framework of parameterized complexity.

We show that the problem is computationally intractable even in fairly restricted scenarios, for instance, even when the underlying graph is a path. We show NP-hardness for settings where the graph has only two distinct valuations among the agents. We demonstrate W-hardness with respect to the number of goods or the size of the vertex cover of the underlying graph. We also consider notions of proportionality that respect the structure of the underlying graph.

**Keywords:** Fair division · Social networks · Envy-freeness · Parameterized complexity

## 1 Introduction

The problem of fairly allocating resources among a set of agents with (possibly distinct) interests in said resources is a fundamental problem with important and varied practical applications. We focus on the problem of allocating indivisible items: in this setting, we have  $n$  agents and  $m$  resources, and every agent expresses their utilities for the resources, either as a ranking over the resources or by specifying a valuation function. The goal is to determine an *allocation* of the items to the agents that respects some notion of “fairness” and “efficiency”. We use the term *bundle* to refer to the set of items that an agent receives in an allocation.

Envy-freeness is one of the most widely used notions of fairness. Given an allocation, an agent envies another if it perceives the bundle of the other agent to be more valuable than her own. An allocation is *envy-free* if no agent envies

another. Note that the trivial allocation that leaves every agent empty-handed is always envy-free. Therefore, one is typically interested in fair allocations that also satisfy some criteria of economic *efficiency*, such as completeness (every good should be allocated to some agent), non-wastefulness (no agent receives a piece of cake that is worth nothing to her and worth something to another agent), or Pareto-efficiency (there is no other feasible agreement that would make at least one agent strictly better off while not making any of the others worse off). We remark here that just as there are trivial allocations that are fair, it is also possible to trivially achieve efficiency if we had no fairness considerations involved: for instance, the allocation that gives all goods to a single agent is Pareto-efficient assuming that the agent has a strictly monotonic utility function over the items.

The question of finding allocations that respect fairness and efficiency demands simultaneously is non-trivial: in particular, such allocations may not exist (if there are two agents and one good, and both agents have positive utility for this single resource), and can be computationally hard to find (for instance, the problem of finding a complete envy-free allocation between even two agents who hold identical valuations over  $m$  goods is equivalent to the PARTITION problem).

The focus of this work is the notion of local envy-freeness. In this setting, the agents are related by a graph, which might be thought of as modeling a social network over the agents, and we explore notions of fairness that account for the structure of this network. For instance, the notion of envy is now restricted: it only manifests between agents who are friends in the network. This is a compelling model of fairness, since agents are likely to not envy agents about whom they have little or no information. We note that the problem of fair division respecting a social network generalizes the classical notion, which can be captured by considering a complete graph on the agents. Thus, the problem of finding allocations that are “locally fair” is a generalization of the classical allocation problem.

## 1.1 Related Work

The model of local envy-freeness has been proposed and considered in several recent lines of work. Some of the earliest considerations for incorporating a graph structure on the agents were made in the context of the *cake-cutting* problem, which is the closely related setting of allocating a divisible resource among agents [1, 4]. Abebe et al. [1] consider both directed and undirected graphs and focus on characterizing the structure of graphs that admit algorithms with certain bounds. They also consider the issue of the *price of envy-freeness* in this setting, which compares the total utility of an optimal allocation to the best utility of an allocation that is envy-free. Bei et al. [4], on the other hand, propose a moving-knife algorithm that outputs an envy-free allocation on trees and an algorithm for computing a proportional allocation on descendant graphs.

We now turn to the literature in the context of indivisible items. Beynier et al. [5] study the fair division problem in the setting of “house allocation”:

here agents have (strict) preferences over items, and each agent must receive exactly one item. An agent envies another in this setting if she prefers the item received by the other agent over her own. In the case of a complete network, for an allocation to be envy-free, each agent must get her top object, and this assignment is automatically Pareto-efficient as well. This motivates the setting of local envy-freeness with respect to a graph on the agents. The authors consider the case when the underlying graph is undirected, and they also consider a variant of the problem where agents themselves can be located on the network by the central authority. These problems turn out to be computationally intractable even on very simple graph structures.

Bredereck et al. [10] consider the problem of graph-based envy-freeness in the context of *directed* graphs and for various classes of valuations: including binary, identical, additive, and even valuations that are both identical and binary. They also consider the complexity of the allocation problem in the framework of parameterized complexity.<sup>1</sup> Somewhat surprisingly, it turns out that finding complete envy-free allocations in the setting of a graph is NP-hard even when the valuations are binary *and* identical. Note that in this setting, every agent in every strongly connected component must get the same number of items: thus, the allocation problem is trivial for directed graphs that are strongly connected, but NP-hard for general directed graphs. Also, it turns out that for general binary preferences, the problem of finding a complete envy-free allocation is NP-hard even when the graph is strongly connected. The problem is also tractable for DAGs: indeed, allocating all resources to a single source agent (corresponding to a vertex with no incoming arcs) is both complete and locally envy-free since nobody can envy a source agent, and empty-handed agents have no envy for each other.

More recently, Eiben et al. [13] consider the problem of finding locally envy-free allocations and envy-free allocations that are additionally proportional in the setting of directed graphs in the framework of parameterized complexity, and specifically considering parameters such as treewidth, cliquewidth, and vertex cover — all of these reflect the structure of the underlying network. It turns out that the problem of finding fair and efficient allocations is tractable for networks that have bounded values for these parameters with some additional assumptions that bound the number of item types or the size of the largest bundle received by an agent. The authors also show hardness results in both the parameterized and classical settings. For instance, the authors show that finding a locally envy-free allocation is NP-hard even when the underlying network is a star, but we note that this is in the setting of general utilities.

The work of Bredereck et al. [8,9] demonstrates that the problem of finding fair and efficient allocations in various settings (including graph-based constraints) is fixed-parameter tractable in the combined parameter “number of agents” and “number of item types” for general utilities. In contrast, our work here focuses on smaller parameters for the special case of binary utilities.

In [11], Chevaleyre, Endriss, and Maudet consider *distributed* mechanisms for allocating indivisible goods, in which agents can locally agree on deals to

---

<sup>1</sup> The terminology relevant to this framework is introduced in the next section.

exchange some of the goods in their possession. This study focuses on convergence properties for such distribution mechanisms both in the context of the classical setting and the setting involving social constraints coming from an underlying undirected graph. Here, the notions of fairness localized according to the graph, and the network also constraints the exchanges that can take place — agents can engage in an exchange only if they are friends in the network. There are also some lines of work that suggest eliminating envy by some mechanism for *hiding* information [14].

## 1.2 Our Contributions

Our focus in this paper is on the setting when agents have binary valuations over the goods and the underlying social network is modeled by an undirected graph. Our focus is on exploring the computational complexity of finding locally envy-free allocations that are also Pareto efficient (EEF) in the framework of parameterized complexity, building most closely on the works of [6, 10, 13].

*Bounded Agent Types.* We begin by noting that the setting of undirected graphs can be significantly different from their directed counterparts: indeed, recall that finding a complete and locally envy-free allocation was NP-hard for even identical binary valuations for directed graphs, but the analogous question is easily seen to be tractable for undirected graphs (indeed, observe that the notions of strong connectivity and connectivity coincide). This motivates the question of whether the problem of finding locally EEF allocations is easier for undirected graphs with a bounded number of agent types. We answer this question in the negative by showing that the problem of determining locally envy-free allocations is NP-hard even when there are only two distinct binary valuations among the agents by a reduction from a graph separation problem called CUTTING  $\ell$  VERTICES (Theorem 1).

*Sparse and Dense Graphs.* In contrast with the result for DAGs, we show that finding locally envy-free allocations that are Pareto efficient (EEF) is NP-hard even when the underlying graph is a path (Theorem 4 and Corollary 1). Although Beynier et al. [5] also show hardness results for very sparse graphs, we note that our methods are significantly different since the models for the valuations are different and additionally, the allocations we seek need not give every agent exactly one item. Moving away from sparsity, we recall that finding complete envy-free allocations for binary valuations is known to be NP-hard even for complete graphs [3, 14], which justifies the need for using additional parameters<sup>2</sup> in the XP<sup>3</sup> algorithm for finding locally envy-free allocations shown by [13, Theorem 10].

<sup>2</sup> The algorithm referred to is XP in the cliquewidth of the underlying graph, the number of agent types and item types.

<sup>3</sup> XP is the class of parameterized problems that can be solved in time  $n^{f(k)}$  for some computable function  $f$ .

*Structural Parameters I: Treewidth and Cliquewidth.* Informally speaking, the parameters treewidth and cliquewidth of graphs quantitatively capture the sparsity and density of the graph by measuring their “likeness” to trees and complete graphs. The results we have already for sparse and dense graphs demonstrate that these parameters being bounded alone is not enough to obtain tractable algorithms. On the other hand, the results of [13] imply that the problem of finding complete and locally envy-free allocations admits XP algorithms when parameterized by either the treewidth or cliquewidth of the underlying graph jointly with the number of item types and agent types. Since their model allows for bidirectional edges, these results apply to the setting of undirected graphs as well. We note that the algorithms described in [13] focus on complete allocations, but can be adapted to account for Pareto efficiency as well.

*Structural Parameters II: Vertex Cover and Twin Cover.* In the setting of directed graphs and general utilities, we note that the problem of finding a complete and locally envy-free allocation is NP-hard even when the underlying graph is a star. In particular, this demonstrates hardness on graphs with a constant-sized vertex cover.<sup>4</sup> It is not clear if this is the case for undirected graphs and binary utilities. We show that the problem of finding locally EEF allocations is  $W[1]$ -hard when parameterized by the vertex cover number (Theorem 3). We remark that a stronger hardness result can be observed for the closely related parameter of twin cover<sup>5</sup> —indeed, the known NP-hardness of finding envy-free allocations for binary valuations on complete graphs [3, 14] implies hardness for graphs that have a twin cover of size zero.

*Few Resources or Agents.* We also consider the cases where the number of goods or the number of agents are relatively small. When considering these parameters, the work of Bliem et al. [6] shows that the computation of EEF allocations is FPT when parameterized by the number of goods or the number of agents for additive 0/1 valuations. In contrast, we show that finding EEF allocations respecting the structure of an underlying undirected graph is  $W[1]$ -hard when parameterized by the number of goods (Theorem 2). On the other hand, the FPT algorithm when parameterized by the number of agents can be extended to account for the graph constraints (noted in Observation 1 in the full version [15]).

*Other Notions of Fairness.* Finally, we also consider notions of proportionality in the context of graphs — we refer to these as local and quasi-global proportionality concepts, representing the extent to which the definitions account for the underlying graph. We demonstrate that computing a locally proportional allocation is NP-hard (Theorem 5), while computing a proportional allocation that is quasi-global is tractable (Theorem 6). Notions of local proportionality

<sup>4</sup> A vertex cover of a graph is a subset of vertices that contains at least one endpoint of every edge. A graph with a bounded vertex cover also has bounded treewidth.

<sup>5</sup> A twin cover of a graph is a subset of vertices  $S$  such that  $G \setminus S$  is a disjoint union of cliques, and further, every pair of vertices  $u, v$  in any clique of  $G \setminus S$  are “twins”, that is,  $N[v] = N[u]$ .

have been proposed and studied in several of the papers that were summarized in the previous section.

## 2 Preliminaries

We use standard terminology from graph theory and fair division. Unless mentioned otherwise, the graphs we consider are simple and undirected. For a graph  $G = (V, E)$ , consisting of a set  $V$  of vertices and a set  $E$  of edges, by  $N(v)$  we denote the *neighborhood* of vertex  $v \in V$ , i.e., the set  $W \subset V$  of vertices such that for each vertex  $w \in W$  there exists an edge  $e = (v, w) \in E$ . The *closed neighborhood* of a vertex  $v$  is  $N(v) \cup \{v\}$  and is denoted  $N[v]$ . The *degree* of a vertex  $v$ , denoted  $d(v)$ , is  $|N(v)|$ . A *clique* is a subset of vertices which are pairwise adjacent. An *independent set* is a subset of vertices, no two of which are adjacent. For  $X \subseteq V$ , the *induced subgraph*  $G[X]$  denotes the subgraph whose vertex set is  $X$  and the edge set consists of all edges whose both end points are in  $X$ .

An instance of fair division for indivisible goods consists of  $n$  agents  $A = \{1, \dots, n\}$  and  $m$  goods (also called *items* or *resources*),  $R = \{o_1, \dots, o_m\}$ . Further, we are also given *valuations* (also called *preference functions* or *utilities*)  $v_\ell : 2^R \rightarrow \mathbb{Z}$  for every agent  $\ell \in A$ . We will assume throughout that the valuation functions are *additive*, i.e., for each agent  $\ell \in A$  and any set of goods  $S \subseteq R$ ,  $v_\ell(S) := \sum_{o \in S} v_\ell(\{o\})$ . A *0/1 valuation* is a function that takes values in  $\{0, 1\}$ , while valuations are said to be *identical* if every agent has the same preference function. In the context of 0/1 valuations, we say that an agent *values* or *approves* a good if her utility for the good is 1. We will use  $\mathcal{V}$  to denote the valuations of the agents  $A$  over  $R$ . When considering fair division in the context of social networks, we are also given an undirected graph  $G$  over the agents  $A$ .

Every subset  $S \subseteq R$  is called a *bundle*. An allocation is a function  $\pi : A \rightarrow 2^R$  mapping each agent to the bundle she receives, such that  $\pi(i) \cap \pi(j) = \emptyset$  when  $i \neq j$  because the items cannot be shared. When  $\bigcup_{a \in A} \pi(a) = R$ , the allocation  $\pi$  is said to be *complete*, otherwise it is *partial*. An allocation is *non-wasteful* if every good is allocated to an agent that assigns positive utility to it.

An allocation  $\pi'$  *dominates*  $\pi$  if for all  $\ell \in A$  it holds that  $v_\ell(\pi(\ell)) \leq v_\ell(\pi'(\ell))$  and for some  $a_j \in A$  it holds that  $v_{a_j}(\pi(a_j)) < v_{a_j}(\pi'(a_j))$ . An allocation  $\pi$  is *Pareto-efficient* if there exists no allocation  $\pi'$  that dominates  $\pi$ . In the case of 0/1 preferences, we note that an allocation is Pareto-efficient if and only if it is complete and non-wasteful, assuming that each resource provides a value of 1 to at least one agent.

Given an instance of fair division  $(A, R, G = (A, E), \mathcal{V})$  as described above, we now introduce the following fairness notions:

- ▷ **Graph Envy-Freeness (GEF)**. We call allocation  $\pi$  graph-envy-free if for each pair of (distinct) agents  $i, j \in A$  such that  $j \in N(i)$ , it holds that  $v_i(\pi(i)) \geq v_i(\pi(j))$ .
- ▷ **Quasi-Global Proportionality (QP)**. We say that an allocation  $\pi$  achieves quasi-global proportionality if for each agent  $\ell \in A$ ,  $v_\ell(\pi(i)) \geq \frac{1}{d(\ell)+1} v_\ell(R)$ .

▷ **Local Proportionality (LP).** We say that an allocation  $\pi$  achieves local proportionality if for each agent  $\ell \in \mathcal{A}$ ,  $v_\ell(\pi(\ell)) \geq \frac{1}{d(\ell)+1} \sum_{j \in N(\ell)} v_\ell(\pi(j))$ .

Note that the graph versions of variants of envy-freeness (such as EF1 or EFX) can be defined analogously in a straightforward manner. It is easy to see that any graph envy-free allocation is also locally proportional and that if the underlying graph is complete, then local proportionality coincides with the standard notion of proportionality. For the problems we consider, we are typically given an instance of fair division on a graph, and the goal is to determine if there exists an allocation that satisfies some notion of fairness and efficiency. For instance, consider the following problems:

GRAPH ENVY-FREE ALLOCATION ( $\mathcal{E}$ -GEFA)  
**Input:** An instance of fair division on a graph  $(\mathcal{A}, \mathcal{R}, \mathcal{G} = (\mathcal{A}, \mathcal{E}), \mathcal{V})$ .  
**Question:** Does there exist an envy-free, Pareto-efficient allocation?

LOCALLY PROPORTIONAL ALLOCATION ( $\mathcal{E}$ -LPA)  
**Input:** An instance of fair division on a graph  $(\mathcal{A}, \mathcal{R}, \mathcal{G} = (\mathcal{A}, \mathcal{E}), \mathcal{V})$ .  
**Question:** Does there exist a Pareto-efficient allocation that achieves local proportionality?

For any efficiency concept (X) and fairness notion (Y), the X-YA problem is defined in a similar fashion. Although our questions are posed as decision versions, we note that most of our algorithms can be easily adapted to handle the natural “search” version of these problems. We refer the reader to the books [7, 16] and the article [11] for additional background on fair division.

A problem *parameterized* by  $k$  is fixed-parameter tractable if it is solvable in  $f(k)|I|^{O(1)}$  time for some computable function  $f$  and the input size  $|I|$  according to the problem’s encoding. Informally,  $W$ -hard problems are presumably not fixed-parameter tractable. The problem of finding a clique on at least  $k$  vertices is  $W[1]$ -hard when parameterized by  $k$ . We call a problem para-NP-hard if it is NP-hard even for a constant value of the parameter. For a comprehensive introduction to the paradigm of parametrized complexity and algorithms, we refer the reader to the book [12].

### 3 Envy-Freeness

#### 3.1 NP-Hardness for Two Agent Types

In this section, we show that finding  $\mathcal{E}$ -GEFA allocations is NP-hard even in the setting of near-identical binary valuations: in particular, when all agents have one of two possible utilities over the items. Note that in the setting of identical

binary valuations when the graph  $G$  is connected, it is easy to see that all agents must value all goods without loss of generality, and that desirable allocations are the ones that allocates the same number of goods to each agent, where the goods themselves may be arbitrarily chosen. Indeed, it is clear that an allocation with equal bundle sizes is  $\mathcal{E}$ -GEFA. On the other hand, consider a  $\mathcal{E}$ -GEFA allocation that does not allocate bundles of equal size to all agents. Let  $\mathbf{a}_i$  and  $\mathbf{a}_j$  be two agents that receive bundles of different size. We can always find two adjacent agents on a path from  $\mathbf{a}_i$  to  $\mathbf{a}_j$  who have received bundles of different sizes, contradicting envy-freeness.

We now show that even a slightly more general situation is computationally intractable — in particular, if all agents have one of two valuations over the goods, the problem of identifying  $\mathcal{E}$ -GEFA allocations is NP-hard. Due to lack of space, the proof is deferred to the full version [15].

**Theorem 1.** *The  $\mathcal{E}$ -GEFA problem is NP-complete even when there are two agent types, and further, agents have 0/1 valuations over the goods.*

### 3.2 W-Hardness Parameterized by Goods

In this section, we demonstrate the hardness of finding  $\mathcal{E}$ -GEFA allocations even when the number of goods is bounded by showing that the problem is  $W[1]$ -hard when parameterized by the number of goods.

**Theorem 2.** *The  $\mathcal{E}$ -GEFA problem is  $W[1]$ -hard when parameterized by the number of goods, even when agents have 0/1 valuations over the goods.*

We describe a reduction from the  $W[1]$ -hard problem CLIQUE, given a graph  $G$  and an integer  $k$ , does there exist a clique of  $k$  vertices in  $G$ . Let  $\mathcal{J} = (G, k)$  be an instance of clique, where  $G = (V, E)$  and further,  $V = \{v_1, \dots, v_n\}$  and  $E = \{e_1, \dots, e_m\}$ . We assume, without loss of generality, that  $m \geq \binom{k}{2}$ , since we can always return a trivial NO-instance when this is not the case. We begin by describing the construction of the reduced instance  $\mathcal{J}_{\mathcal{J}} := (\mathbf{A}, \mathbf{R}, \mathbf{H} = (\mathbf{A}, \mathbf{F}), \mathcal{V})$ . We define the set of goods  $\mathbf{R}$  as follows:

$$\mathbf{R} = \{q_1, \dots, q_\ell, p_1, \dots, p_k, d_1, \dots, d_{\ell+1}\},$$

where  $\ell = \binom{k}{2}$ . For ease of discussion, we call the first  $\ell$  goods *popular* and the next  $k$  goods *specialized*. The remaining are *dummy* items. We now define the set of agents as  $\mathbf{A} = \mathbf{V} \cup \mathbf{E} \cup \mathbf{S} \cup \mathbf{W}$ , where:

$$\mathbf{S} := \{s_1, \dots, s_{\ell+1}\} \text{ and } \mathbf{W} := \{w_{ij} \mid i \in [n], j \in [\ell + 1]\}.$$

We indulge in a mild abuse of notation and use  $v_i$  to refer to both an element of  $\mathbf{V}$  from the clique instance and an agent of  $\mathbf{A}$  in the reduced instance (similarly for edges). The edges of  $\mathbf{H}$  are as follows:

- ▷  $e = (u, v) \in E$  is adjacent to all vertices of  $\mathbf{S}$  and  $u, v$ .
- ▷  $v_i \in \mathbf{V}$  is adjacent to all vertices  $w_{ij}$  for  $j \in [\ell + 1]$ .
- ▷ For each  $1 \leq i \leq n$ ,  $\mathbf{H}[\cup_{j=1}^{\ell+1} w_{ij}]$  induces a clique.



The preferences of the agents are as follows:

- ▷ All agents have an utility of 1 for the popular goods.
- ▷ All agents in  $V$  have an utility of 1 for the specialized goods.
- ▷ The agent  $s_i \in S$  has an utility of 1 for  $d_i$ , for all  $i \in [\ell + 1]$ .

This completes the construction of the instance (Fig. 1). Note that the number of goods is a function of  $k$  alone. We now turn to the argument for equivalence, a clique  $X \subseteq V$ , of size  $k$  exists in  $G$  *iff*, there is an GEF allocation for the instance constructed  $\mathcal{J}_j := (A, R, H = (A, F), \mathcal{V})$ .

*Proof.* In the forward direction, let  $X \subseteq V$  be a clique in  $G$  and let  $Y := G[X] \subseteq G$ . Consider now the following allocation  $\pi$ . We let each agent corresponding to  $Y$  receive one popular item, each agent corresponding to  $X$  receive one specialized item, and finally allocate the item  $d_i$  to  $s_i$  for all  $i \in [\ell + 1]$ . It is straightforward to verify that the allocation  $\pi$  is Pareto-efficient and envy-free with respect to  $H$ .


This concludes our description of a fair and complete allocation strategy given a clique in  $G$ . We now turn to the reverse direction, where we are given an allocation  $\pi$  that is Pareto-efficient and envy-free with respect to  $H$ . It is useful to make the following observation about  $\pi$  to begin with.

*Claim.* Let  $H$  be defined as above, and let  $\pi$  be an allocation that is Pareto-efficient and envy-free with respect to  $H$ . Then, any popular good is assigned by  $\pi$  to an agent from  $E$ . Further, no agent in  $E$  can receive more than one popular good in the allocation  $\pi$ .

Since  $\pi$  is non-wasteful, the specialized goods must be distributed among agents corresponding to  $V$ . The following is easy to see.

*Claim.* Let  $H$  be defined as above, and let  $\pi$  be an allocation that is Pareto-efficient and envy-free with respect to  $H$ . No agent in  $V$  can receive more than one specialized good in the allocation  $\pi$ .

Let  $X \subseteq V$  be the subset of  $k$  agents that receive at least one specialized item and let  $Y \subseteq E$  be the subset of  $\ell$  agents that receive at least one popular item



**Fig. 1.** A sketch of the reduced instance based on an instance  $G = (V, E), k$  of CLIQUE. Recall that  $\ell$  denotes  $\binom{k}{2}$ , and only some vertices of  $W$  are shown for clarity. The shaded vertices induce a complete subgraph. The edge  $e_t = (v_i, v_j)$  is adjacent only to  $v_i$  and  $v_j$  among vertices in  $V$ .

with respect to  $\pi$ . We claim that  $G[X]$  is a clique. In particular, we claim that every edge of  $Y$  has both its endpoints in  $X$ . Indeed, suppose not, and let  $e \in Y$  be an edge with at least one endpoint (say  $v$ ) outside  $X$ . Then,  $v$  envies  $e$ , which contradicts our assumption about  $\pi$  being envy-free with respect to  $H$ .  $\square$

### 3.3 W-Hardness Parameterized by Vertex Cover

Recall that a vertex cover of a graph  $G = (V, E)$  is a subset  $S \subseteq V$  such that  $G \setminus S$  is an independent set (i.e., for any pair of vertices  $u, v \in G \setminus S$ ,  $(u, v) \notin E$ ). In the setting of directed graphs with arbitrary utilities, finding  $\mathcal{E}$ -GEFA allocations is NP-hard even for graphs that have a constant-sized vertex cover. Here, we show that in the setting of binary utilities, finding a  $\mathcal{E}$ -GEFA allocation is  $W[1]$ -hard when parameterized by the vertex cover of the underlying graph. Due to lack of space, the proof is deferred to the full version [15].

**Theorem 3** (\*). *The  $\mathcal{E}$ -GEFA problem is  $W[1]$ -hard when parameterized by the vertex cover of the underlying graph, even when agents have 0/1 valuations over the goods.*

### 3.4 NP-Hardness on Paths

To show the hardness of  $\mathcal{E}$ -GEFA even when the underlying graph is a path, we reduce from a variant of SAT called LINEAR SAT (abbreviated LSAT). In an LSAT instance, each clause has at most three literals, and further the literals of the formula can be sorted such that every clause corresponds to at most three consecutive literals in the sorted list, and each clause shares at most one of its literals with another clause, in which case this literal is extreme in both clauses. The hardness of LSAT was shown in [2]. In fact, by studying the reduced instance, one may assume that a “hard” instance of LSAT has the following structure: the first  $2q$  clauses have two literals each and are of the following form:

$$A_i = \{s_i, \ell_i\}, B_i = \{\ell_i, t_i\}; 1 \leq i \leq q,$$

where  $s_i, \ell_i$ , and  $t_i$  denote literals, while the remaining  $p$  clauses have three literals each and are mutually disjoint from each other as well as the first  $2q$  clauses. For ease of description, we will assume that the LSAT formula that we reduce from has this particular structure. We are now ready to describe our reduction — in the interest of simplicity, our proof is designed to address the case when the graph is a disjoint union of paths, although it is easy to “stitch” these components into a single, longer path, as we will explain later.

**Theorem 4.** *The  $\mathcal{E}$ -GEFA problem is NP-complete even when the graph induced by the agents is a disjoint union of paths, and further, agents have 0/1 valuations over the goods.*

Membership in NP is straightforward to check. We focus here on the reduction demonstrating hardness. Let  $\phi$  be an instance of LSAT over variables  $\hat{X} := \{x_1, \dots, x_n\}$  and clauses:

$$C := \{A_1, B_1, \dots, A_q, B_q, C_1, \dots, C_p\},$$

as described above. We refer to the first  $2q$  clauses as the *coupled* clauses and the remaining as *isolated* clauses. We now turn to the construction of the reduced instance  $\mathcal{J}_\phi := (V, R, H, \mathcal{V})$ . We define the set of goods  $R$  as  $R_{\hat{X}} \cup R_C$ , where:

$$R_{\hat{X}} = \{y_1, \dots, y_n, x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\},$$

and:

$$R_C = \{g_1, \dots, g_q, d_1, \dots, d_p\}.$$


The set of agents  $V$  is given by  $X \cup C \cup Y \cup G \cup D$ , where  $C$  is denoted in the same way as in the LSAT instance, and further:

$$X = \{X_1, \dots, X_n\}, Y = \{Y_1, \dots, Y_n\},$$


$$G = \{G_1, \dots, G_q\} \text{ and } D = \{D_1, \dots, D_p\}.$$

We now simultaneously describe the structure of the graph  $H$  and the preferences of the agents.


**Assignment Gadgets.** For each  $1 \leq i \leq n$ , add an edge between  $X_i$  and  $Y_i$ . The agent  $X_i$  values  $\{x_i, \bar{x}_i, y_i\}$ , while  $Y_i$  values the good  $y_i$  (and nothing else).



**Isolated Clause Gadgets.** For each  $1 \leq i \leq p$ , we add an edge between agents  $C_i$  and  $D_i$ . The agent  $C_i$  values the literal  $\ell$  if and only if  $\ell \in C_i$  along with  $d_i$ , while  $D_i$  values the good  $d_i$  (and nothing else).



**Coupled Clause Gadgets.** For each  $1 \leq i \leq q$ , we add an edge between agents  $A_i$  and  $B_i$ , and also an edge between  $G_i$  and  $A_i$ . The agent  $G_i$  values the good  $g_i$  (and nothing else). Agents  $A_i$  and  $B_i$  value, respectively, the goods  $\{g_i, s_i, t_i, \ell_i\}$  and  $\{s_i, t_i\}$ .



**Fig. 2.** A schematic of the reduced instance in the proof of Theorem 4, which is a disjoint union of paths.  $l_1, l_2, l_3$  are the literals of clause  $C_i$

This completes the description of the construction (Fig.2). We defer the proof of equivalence to the full version [15] due to lack of space.

We remark that it is possible to combine the connected components in the reduced instance above by simply introducing “dummy connector agents” that each value a corresponding dummy item and nothing else, leading to the following consequence.

**Corollary 1.** *The  $\mathcal{E}$ -GEFA problem is NP-complete even when the graph induced by the agents is a path, and further, agents have 0/1 valuations over the goods.*

## 4 Proportionality for Graphs

### 4.1 Local Proportionality: NP-hardness

**Theorem 5.** *The  $\mathcal{E}$ -LPA problem is NP-complete on undirected graphs, even when all agents have 0/1 valuations over the resources.*

*Proof.* We defer the proof to the full version [15] due to lack of space. □

### 4.2 Quasi-Global Proportionality: Efficient Algorithms

To obtain efficient algorithms for finding Pareto-efficient allocations that respect quasi-global proportionality, We model the problem of finding Pareto-efficient allocations respecting quasi-global proportionality using an integer linear program (ILP) with a structured constraint matrix. In particular, it is well-known that if the constraint matrix of an ILP is totally unimodular,<sup>6</sup> then the corresponding instance can be solved in polynomial time. We turn to an explanation of our encoding.

**Theorem 6.** *The problem of finding a Pareto-efficient allocation that is quasi-globally proportional with respect to an underlying undirected graph on the agents can be solved in polynomial time if all agents have 0/1 valuations.*

*Proof.* Let us assume there are  $n$  agents and  $m$  goods. We will introduce a variable  $x_{ij}$  which indicates whether agent  $i$  gets good  $j$ , and  $a_{ij}$  indicates whether agent  $i$  likes good  $j$ . These constraints are as follows.

- ▷ We encode the fact that the allocation defined by  $x$  is well-defined by introducing the following constraint for each good  $j$ :
- ▷ For each agent  $i$ , let  $s_i$  be the number of items that have utility 1 for agent  $i$ . For each agent  $i$ , introduce the following proportionality constraint:

$$\forall j \sum_{i=1}^n x_{ij} \leq 1 \text{ and } \forall i \sum_{j=1}^m a_{ij} x_{ij} \geq \frac{s_i}{(d_i + 1)}$$

We let the objective function be  $\sum_{i=1}^n a_{ij} x_{ij}$ . Note that any assignment for which this function achieves a value of  $m$  is complete and non-wasteful, and also respects quasi-global proportionality. It is straightforward to verify that the constraint matrix for the ILP described above is totally unimodular for any underlying graph  $H$ . □

<sup>6</sup> A unimodular Matrix is a square integer matrix having determinant  $+1$  or  $-1$ . A totally unimodular matrix is a matrix for which every square non-singular submatrix is unimodular.

We remark that the problem of assigning goods in a proportional fashion (for any of the notions of proportionality that we have introduced) beyond 0/1 valuations is NP-hard even when there are only two agents with identical valuations, by a standard reduction from PARTITION, with the graph being a single edge on two agents.

**Acknowledgments.** The first author is grateful to SERB for their support of this project through the ECR grant ECR/2018/002967.

## References

1. Abebe, R., Kleinberg, J.M., Parkes, D.C.: Fair division via social comparison. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS, pp. 281–289 (2017)
2. Arkin, E.M., et al.: Choice is hard. In: Elbassioni, K., Makino, K. (eds.) ISAAC 2015. LNCS, vol. 9472, pp. 318–328. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-48971-0\\_28](https://doi.org/10.1007/978-3-662-48971-0_28)
3. Aziz, H., Gaspers, S., Mackenzie, S., Walsh, T.: Fair assignment of indivisible objects under ordinal preferences. *Artif. Intell.* **227**, 71–92 (2015)
4. Bei, X., Qiao, Y., Zhang, S.: Networked fairness in cake cutting. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI, pp. 3632–3638 (2017)
5. Beynier, A., Chevaleyre, Y., Gourvès, L., Lesca, J., Maudet, N., Wilczynski, A.: Local envy-freeness in house allocation problems. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS, pp. 292–300 (2018)
6. Bliem, B., Bredereck, R., Niedermeier, R.: Complexity of efficient and envy-free resource allocation: few agents, resources, or utility levels. In: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI, pp. 102–108 (2016)
7. Brams, S.J., Taylor, A.D.: Fair Division - from Cake-Cutting to Dispute Resolution. Cambridge University Press, Cambridge (1996)
8. Bredereck, R., Kaczmarczyk, A., Knop, D., Niedermeier, R.: High-multiplicity fair allocation using parametric integer linear programming. *CoRR* [arXiv:2005.04907](https://arxiv.org/abs/2005.04907) (2020)
9. Bredereck, R., Kaczmarczyk, A., Knop, D., Niedermeier, R.: High-multiplicity fair allocation: Lenstra empowered by n-fold integer programming. In: Proceedings of the 2019 ACM Conference on Economics and Computation. Association for Computing Machinery (2019)
10. Bredereck, R., Kaczmarczyk, A., Niedermeier, R.: Envy-free allocations respecting social networks. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS, 2018, pp. 283–291 (2018)
11. Chevaleyre, Y., Endriss, U., Maudet, N.: Distributed fair allocation of indivisible goods. *Artif. Intell.* **242**, 1–22 (2017)
12. Cygan, M., et al.: Parameterized Algorithms. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-21275-3>
13. Eiben, E., Ganian, R., Hamm, T., Ordyniak, S.: Parameterized complexity of envy-free resource allocation in social networks. In: Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI, pp. 7135–7142. AAAI Press (2020)

14. Hosseini, H., Sikdar, S., Vaish, R., Wang, H., Xia, L.: Fair division through information withholding. In: Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI, pp. 2014–2021 (2020)
15. Misra, N., Nayak, D.: On fair division with binary valuations respecting social networks - full version. CoRR [arXiv:2111.11528](https://arxiv.org/abs/2111.11528) (2021)
16. Robertson, J.M., Webb, W.A.: Cake-Cutting Algorithms - Be Fair If You Can. A K Peters, Natick (1998)