Fair Allocation of Items in Multiple Regions

Houyu Zhou[1*], Tianze Wei[1*†], Biaoshuai Tao[2*], Minming Li[1*]

[1]Department of Computer Science, City University of Hong Kong
[2]School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University
{houyuzhou2-c, t.z.wei-8}@my.cityu.edu.hk, bstao@sjtu.edu.cn, minming.li@cityu.edu.hk

Abstract

We initiate the study of fair allocation with the set of divisible or indivisible items distributed in multiple regions. The key requirement is that each agent can only obtain items from one region. In this work, we consider two kinds of fairness concepts: envy-based notions including *envy-freeness* (EF) and *envy-freeness up to one/any item* (EF1/EFX), and share-based notions including *proportionality* (PROP) and *proportionality up to one/any item* (PROP1/PROPX). On the negative side, we show NP-hardness and inapproximability results about the aforementioned fairness notions. On the positive side, we propose several algorithms to compute the partial allocations that satisfy envy-based notions and allocations that approximate the above fairness notions.

Introduction

Fair allocation studies how to *fairly* allocate a set of items (resources) among agents, which has attracted the attention of many scholars in both theoretical and industrial fields (Steinhaus 1949; Moulin 2004). For instance, network technology companies like Google and Meta use schedulers in the cloud to allocate items (e.g., servers, memory, etc.) or tasks (e.g., development, maintenance, etc.) among several self-interested agents who want to maximize the utility of their own allocations. We observe that there are many real-life scenarios, including the above case, where the items are distributed in multiple regions, and each agent can only obtain the items from a single region. Such scenarios raise a natural question.

How to allocate the items to a set of agents where each agent can only obtain items from one region?

Before answering the question, we consider a scenario where a multinational corporation with many branches in different countries and regions wants to recruit employees. The manager not only assigns employees who meet the requirements to different regions but also allocates different items such as wages, housing, insurance, and medical care to them. Different agents have different valuations for those items, e.g., some employees prefer to live in low-floor dormitories because it is convenient, while others prefer to live in high-floor dormitories because it is quiet. A simple way is to allocate those items equally to the employees in the same region. One difficulty is that it may be impossible to divide every item equally, e.g., every dormitory is unique because of the building and the level it belongs to. In addition, the employee may envy those in another region, though we can achieve equity in the same region. Moreover, we cannot allocate those employees to those regions proportionally with respect to the total valuation of each region since they may have different preferences for those regions. For instance, some employees may prefer inland cities because it is not easy to get wet, while others prefer coastal cities because of the cooler weather. Hence, different employees' valuations in the same region may vary.

One may also wonder how about if we divide the whole allocation procedure into two steps: (1) assign agents to different regions based on their preferences; (2) use existing algorithms in fair allocation to allocate items to agents in each region. The main drawback of this approach is that assigning agents to the regions only on their preferences may cause some crowded regions, which means that every agent in this region may get a low value. The situation may worsen, especially when agents assigned to the same region have similar preferences over items. Hence, we need an integrated procedure to account for agents' preferences over possible solutions properly.

In addition, whether the items are divisible or not, the well-established techniques such as Cut and Choose (Steinhaus 1948), Round-Robin and Envy-Cycles Procedure (Lipton et al. 2004) that are often used in the fair allocation of the divisible items or a set of indivisible items cannot be directly applied to most cases in our setting. There are two reasons: (1) in our setting, each agent can only obtain items from one region, and (2) the current techniques cannot decide how to assign agents to regions. Although there have been several works about fair allocation of multiple divisible items (Cloutier, Nyman, and Su 2010; Nyman, Su, and Zerbib 2020; Hosseini, Igarashi, and Searns 2020), their settings are different from that of our paper, and the techniques in these papers cannot be directly used in our model (See details in Related Work). Hence, none of the previous work fully captures the challenge that we face.

[*]The authors are listed in reverse alphabetical order.
[†]Corresponding author.

	Lower Bound	Upper Bound
EF	× (Coro 1$^{\diamondsuit}$)	
	$\frac{1}{2}$ (Thm 3 and Prop 3$^{\clubsuit}$)	$\frac{1}{2}$ (Coro 2$^{\clubsuit}$)
Partial EF	1 (Thm 2$^{\diamondsuit}$)	
PROP	× (Prop 1$^{\diamondsuit}$)	
	$\frac{1}{2}$ (Thm 3 and Prop 3$^{\clubsuit}$)	$\frac{1}{2}$ (Prop 2$^{\clubsuit}$)

Table 1: The results of divisible items. × : inapproximability. ♢ : diverse-region. ♣ : equal-region.

	Lower Bound	Upper Bound
EF1	× (Prop 4$^{\heartsuit}$)	
	$\frac{1}{2}$ (Prop 9$^{\spadesuit}$)	$\frac{1}{2}$ (Coro 3$^{\spadesuit}$)
Partial EF1	1 (Thm 6$^{\heartsuit}$)	
EFX	× (Prop 4$^{\heartsuit}$)	
	× (Prop 6$^{\spadesuit}$)	
Partial EFX	1 (Thm 5$^{\heartsuit}$)	
PROP1	× (Prop 5$^{\heartsuit}$)	
	$\frac{1}{2}$ (Thm 7$^{\spadesuit}$)	$\frac{1}{2}$ (Prop 8$^{\spadesuit}$)
PROPX	× (Prop 5$^{\heartsuit}$)	
	× (Prop 7$^{\spadesuit}$)	

Table 2: The results of indivisible items. × : inapproximability. ♡ : diverse-region. ♠ : equal-region.

Our Contributions

We develop a formal model for fair allocation in multiple regions. Our model is flexible enough to capture both (1) settings where items are divisible and (2) settings where items are indivisible. Further, it can handle both (a) diverse-region case: every agent has different total valuations for different regions, and (b) equal-region case: every agent has the same total valuation for different regions. When the items are *divisible*, we study fairness criteria including *envy-freeness* (EF), EF-with-bounded-charity, and *proportionality* (PROP). When the items are *indivisible*, we are interested in the fairness notions including *envy-freeness up to one/any item* (EF1/EFX), EF1/EFX-with-bounded-charity, and *proportionality up to one/any item* (PROP1/PROPX). The formal definitions of these notions of fairness can be found in the next section.

For the setting of divisible items, the results of our approximation algorithm design are shown in Table 1. More specifically,

- Divisible items with diverse-region setting
  - Different from the classic setting of one divisible item where the EF allocation always exists, we show that EF allocation may not exist in our setting, and determining whether for a given instance there exists an EF allocation is NP-complete.
  - We show that any bounded approximate EF or PROP allocation cannot be guaranteed.
  - We propose an algorithm (About-To-Envy-Satisfied Algorithm) that outputs an EF-with-bounded-charity allocation in polynomial time under piecewise constant functions.
- Divisible items with equal-region setting
  - Under piecewise constant functions, we show that allocating the unallocated parts in any feasible way on top of the output by the About-To-Envy-Satisfied Algorithm can guarantee a $\frac{1}{2}$-EF allocation.
  - We present an approximation algorithm (Adapted Perfect Partition Algorithm) that computes a $\frac{1}{2}$-EF allocation when agents have general valuation functions.
  - We also show an upper bound of $\frac{1}{2}$ for both EF and PROP, implying our algorithms are the best possible.

For the setting of indivisible items, the results of our approximation algorithm design are shown in Table 2. More specifically,

- Indivisible items with diverse-region setting
  - We show that EF1/EFX allocation may not exist in our setting, and determining whether for a given instance there exists an EF1/EFX allocation is NP-complete.
  - We show that any bounded approximate EF1/EFX or PROP1/PROPX allocation cannot be guaranteed.
  - We propose an algorithm (Sequential Envy-Satisfied Algorithm) that outputs an EFX-with-bounded-charity allocation in pseudo-polynomial time.
  - We present an algorithm (Sequential Envy-Graph-Satisfied Algorithm) that can compute an EF1-with-bounded-charity allocation in polynomial time.
- Indivisible items with equal-region setting
  - We show that any bounded approximate EFX or PROPX allocation cannot be guaranteed.
  - We show that allocating the unallocated items in any feasible way on top of the output allocation by both Sequential Envy-Satisfied Algorithm and Sequential Envy-Graph-Satisfied Algorithm can guarantee a $\frac{1}{2}$-EF1 allocation.
  - We propose an algorithm (Adapted Round-Robin Algorithm) that computes a $\frac{1}{2}$-PROP1 allocation in polynomial time.
  - We also show an upper bound of $\frac{1}{2}$ for both EF1 and PROP1, implying our algorithms are the best possible.

Due to the space limit, some proofs in the following sections are omitted.

Related Work

There has been a significant line of work in the research on the fair allocation of divisible items. In terms of *envy-freeness* (EF), Brams and Taylor (1995) gave the first envy-free protocol for any number of agents, but this protocol requires a finite but unbounded number of cuts. Aziz and Mackenzie (2016b) studied the bounded envy-free protocol with four agents, and then they extended the result to any number of agents (Aziz and Mackenzie 2016a). Besides *envy-freeness*, other classic fairness notions such as *proportionality* and *equitablity* have also been widely studied (Dubins and Spanier 1961; Even and Paz 1984; Edmonds and Pruhs 2006; Cechlárová and Pillárová 2012; Procaccia and Wang 2017). A natural generalization of fair allocation of one cake, i.e., there are multiple cakes, and each agent needs to get one piece from each cake, has been explored. Cloutier, Nyman, and Su (2010) studied the envy-free allocation of multiple cakes with two agents. Lebert, Meunier, and Carbonneaux (2013) showed that an EF allocation always exists with two cakes and three agents. Nyman, Su, and Zerbib (2020) proved that for k cakes, an EF allocation always exists when there are $k(n-1)+1$ agents and each cake is cut into k pieces. Besides that, there are also some works related to multi-cake cutting, such as multi-layered cake cutting and fair multi-cake cutting (Hosseini, Igarashi, and Searns 2020; Segal-Halevi 2021).

In the research on the fair allocation of indivisible items, most work focuses on the fair allocation of one set of items instead of several sets distributed in multiple regions. Because the aforementioned fairness notions cannot be guaranteed when items are indivisible, their relaxations are considered. The prominent fairness notions considered include *envy-freeness up to one/any item* (EF1/EFX), *proportionality up to one/any item* (PROP1/PROPX) and *maximin share* (MMS) (Lipton et al. 2004; Caragiannis et al. 2019; Conitzer, Freeman, and Shah 2017; Moulin 2019; Budish 2011). Among these notions, EF1 and PROP1 allocations can be computed in polynomial-time (Lipton et al. 2004). However, PROPX and MMS allocations may not exist (Aziz, Moulin, and Sandomirskiy 2020; Procaccia and Wang 2014). As for the existence of EFX allocation, it is still an open question, but there are some papers about partial EFX, e.g., Caragiannis, Gravin, and Huang (2019) introduced one notion named EFX-with-charity where some items are not allocated to agents, and the partial allocation is EFX. Chaudhury et al. (2021) proposed an algorithm to compute an EFX-with-bounded-charity allocation where the number of unallocated items and their value are both bounded. The recent progress and open problems in this field can be found in the survey of Amanatidis et al. (2023).

Preliminaries

We introduce our model and solution concepts in this section. Let $N=\{1,\dots,n\}$ denote the set of n agents, $D=\{d_1,\dots,d_m\}$ denote the set of m heterogeneous divisible items and $M=\{e_1,\dots,e_m\}$ denote the set of m indivisible items. There are k regions denoted as $R=\{R_1,\dots,R_k\}$, where divisible or indivisible items are distributed. Each agent $i\in N$ can obtain items from a single region.

Fair Allocation with Divisible Items. In this setting, we have $m=k$, i.e., each region has one divisible item, respectively. If $k=n$, our model is approximately equivalent to the fair allocation of indivisible items where the numbers of items and agents are the same. Hence, we mainly consider $k<n$. When $m=k=1$, our model degenerates to the classic cake cutting model. For each divisible item $d_i\in D$, we assume that it is denoted by an interval $[0,1]$. A *piece* of one item $S\subseteq[0,1]$ is a finite union of subintervals of $[0,1]$. Each agent is endowed with a non-negative integrable density function f_i. Given a piece S of one item, the value of S in agent i's perspective is denoted by $v_i(S)=\int_{x\in S} f_i(x)dx$. Without loss of generality, we assume that $v_i(D)=1$ for all agents $i\in N$. A divisible items allocation instance is denoted as $\mathcal{I}=\langle D,N,\boldsymbol{f}\rangle$ where $\boldsymbol{f}=(f_1,\dots,f_n)$.

An allocation $\mathcal{X}=(X_1,\dots,X_n)$, where X_i is the bundle allocated to agent $i\in N$, is an n-partition of the set of m divisible items among n agents, i.e., $\bigcup_{i\in N}X_i=D$, and $X_i\cap X_j=\emptyset$ for any two agents $i\neq j$. Besides that, in each bundle X_i, there is only one piece from one region.

We consider two types of regions in this setting: diverse-region and equal-region. We call the setting equal-region if for every agent $i\in N$, it holds that $v_i(d_1)=\cdots=v_i(d_k)$. If there is no restriction, we call it diverse-region setting.

We study three notions of fairness in this setting.

Definition 1 (**EF**). *For any $\alpha\in[0,1]$, an allocation $\mathcal{X}=(X_1,\dots,X_n)$ is α-approximate envy-free (α-EF), if for any two agents $i,j\in N$, we have $v_i(X_i)\geq\alpha\cdot v_i(X_j)$. If $\alpha=1$, this allocation is EF.*

Definition 2 (**EF-with-bounded-charity**). *An allocation $\mathcal{X}=(X_1,\dots,X_n)$ is envy-free with bounded charity (EF-with-bounded-charity) if (1) $\mathcal{X}$ is EF; (2) for any agent $i\in N$ and any region $R_j\in R$, $v_i(X_i)\geq v_i(d_j\setminus\bigcup_{i=1}^n X_i)$.*

Definition 3 (**PROP**). *For any $\alpha\in[0,1]$, an allocation $\mathcal{X}=(X_1,\dots,X_n)$ is α-approximate proportional (α-PROP), if for any two agents $i,j\in N$, we have $v_i(X_i)\geq\alpha\cdot\frac{1}{n}$. If $\alpha=1$, this allocation is PROP.*

Next, we introduce the definition of *perfect partition*, a classic concept in the cake cutting model.

Definition 4 (**Perfect Partition**). *A partition of $\mathcal{P}=\{P_1,\dots,P_\ell\}$ of a divisible item d is said to be perfect if for all agents $i\in N$ and $j\in[\ell]$, $v_i(P_j)=\frac{v_i(d)}{\ell}$.*

Fair Allocation with Indivisible Items. In this setting, we have $m\geq k$ since the empty region can be omitted naturally. Similar to the divisible item setting, we mainly consider $k<n$. When $k=1$, our model degenerates to the classic fair allocation with a set of indivisible items. Each agent $i\in N$ has a valuation function $v_i:2^N\to\mathbb{R}_{\geq 0}$. The valuation functions are assumed to be additive, i.e., for any $S\subseteq M$, $v_i(S)=\sum_{e\in S}v_i(\{e\})$. For simplicity, we use $v_i(e)$ instead of $v_i(\{e\})$ for any $e\in M$. Without loss of generality, we assume that $v_i(M)=1$ for all agents $i\in N$. In particular, let M_ℓ denote the set of indivisible items in the region R_ℓ where $\ell\in[k]$. An indivisible items allocation instance is denoted as $\mathcal{I}=\langle M,N,R,\boldsymbol{v}\rangle$ where $\boldsymbol{v}=(v_1,\dots,v_n)$.

An allocation $\mathcal{X}=(X_1,\dots,X_n)$, where X_i is the bundle allocated to agent $i\in N$, is an n-partition of the set of m indivisible items among n agents, i.e., $\bigcup_{i=1}^n X_i=M$, and $X_i\cap X_j=\emptyset$ for any two agents $i\neq j$. In addition, in each bundle X_i, there are only items from one region.

We also consider two types of regions in this setting: diverse-region and equal-region. We call the setting equal-region if, for every agent $i\in N$, it holds that

$\sum_{e\in M_1} v_i(e) = \cdots = \sum_{e\in M_k} v_i(e)$. If there is no restriction, we call it diverse-region setting.

We study the following fairness notions in this setting.

Definition 5 (EF1/EFX). *For any $\alpha \in [0,1]$, an allocation $\mathcal{X} = (X_1,\ldots,X_n)$ is α-approximate envy-free up to one item (α-EF1), if for any two agents $i,j \in N$ with $X_j \neq \emptyset$,*

$$v_i(X_i) \geq \alpha \cdot v_i(X_j \setminus \{e\}) \text{ for some item } e \in X_j. \quad (1)$$

If the quantifier "some" in Inequality (1) is changed to "any", the allocation is α-approximate envy-free up to any item (α-EFX). The allocation is EF1 or EFX if $\alpha = 1$.

Definition 6 (EF1/EFX-with-bounded-charity). *An allocation $\mathcal{X} = (X_1,\ldots,X_n)$ is envy-free up to one item with bounded charity (EF1-with-bounded-charity) or envy-free up to any item with bounded charity (EFX-with-bounded-charity), if (1) $\mathcal{X}$ is EF1 or EFX; (2) for any agent $i \in N$ and any region $R_j \in R$, $v_i(X_i) \geq v_i(M_j \setminus \bigcup_{i=1}^n X_i)$.*

Definition 7 (PROP1/PROPX). *For any $\alpha \in [0,1]$, an allocation $\mathcal{X} = (X_1,\ldots,X_n)$ is α-approximate proportional up to one item (α-PROP1), if for any agent $i \in N$,*

$$v_i(X_i \cup \{e\}) \geq \alpha \cdot \frac{1}{n} \text{ for some item } e \in M \setminus X_i. \quad (2)$$

If the quantifier "some" in Inequality (2) is changed to "any", the allocation is α-approximate proportional up to any item (α-PROPX). The allocation is PROP1 or PROPX if $\alpha = 1$.

Divisible Items

In this section, we study the case of allocating divisible items from multiple regions. We first consider the diverse-region setting and then the equal-region setting.

Diverse-Region

We first show that it is NP-complete to decide whether a given instance admits an allocation that satisfies EF. Before our proof, we introduce the definition of piecewise constant function often appearing in cake cutting literature (Bei et al. 2012; Chen et al. 2013; Tao 2022).

Definition 8 (Piecewise Constant Function). *A valuation density function f_i is piecewise constant if the interval $[0,1]$ can be partitioned into a finite number of intervals such that f_i is a constant function on each interval.*

Theorem 1. *Given an instance $\mathcal{I} = \langle D, N, \boldsymbol{f}\rangle$ with piecewise constant valuation functions, it is NP-complete to decide whether there exists an allocation that satisfies EF.*

Proof Sketch. The proof of NP: given the certificate encoding the assignment of agents to multiple regions, we can verify the existence of EF allocation using linear programming. To prove NP-hardness, we reduce from the PARTITION.

Consider an instance from PARTITION with a set of positive integers $S = \{a_1,\ldots,a_n\}$. We create an instance of $\mathcal{I} = \langle D, N, \boldsymbol{f}\rangle$. To make the proof easier to understand, we do not normalize agents' valuation functions since EF only depends on each agent's own valuation function. We have $2n$ agents $\{t_1,\ldots,t_n,u_1,\ldots,u_n\}$ and $(n+2)$ regions $\{R_1,\ldots,R_{n+2}\}$. Let agent t_i have valuation a_i for each region of $R_1,\ldots,R_n$, and have valuation 0 for each region of R_{n+1} and R_{n+2}. Next, we divide each item in R_{n+1} and R_{n+2} into $n+1$ pieces where all f_i are constant functions with respect to each piece. Let K be a sufficiently large number. Let agent u_i have valuation $K + a_i$ for each region of $R_1,\ldots,R_n$. For each piece i in R_{n+1}, let u_i have valuation K for that piece, 0 for the other pieces except the last piece, and $\frac{1}{2}\sum_{i\in N} a_i$ for the last piece. All agents u_i have the same valuation function for the last piece. For R_{n+2}, we have the same setting with R_{n+1}.

Then we can verify that if there is a partition S_1 and S_2, we can allocate t_i to R_i for every $i \in N$, u_i to R_{n+1} with the value of $K + a_i$ if $a_i \in S_1$ or R_{n+2} with the value of $K + a_i$ if $a_i \in S_2$. If there is an EF allocation, then we can put a_i in S_1 if u_i is in R_{n+1} or S_2 if u_i is in R_{n+2}. □

Besides the NP-completeness result, we show that it is impossible to approximate EF and PROP allocations.

Proposition 1. *There exists an instance for which no allocation is α-PROP for any $\alpha \in (0,1]$. The impossible result holds even for piecewise constant functions.*

Since EF implies PROP, we can directly get the following corollary.

Corollary 1. *There exists an instance for which no allocation is α-EF for any $\alpha \in (0,1]$. The impossible result holds even for piecewise constant functions.*

Combining Proposition 1 and Corollary 1 together, any bounded approximate EF or PROP allocation cannot be guaranteed even with piecewise constant functions. However, if we allow partial allocations, we can bypass the impossibility result in Corollary 1. Next, we will present an algorithm named About-To-Envy-Satisfied Algorithm that outputs an EF-with-bounded-charity allocation under piecewise constant functions. Since the algorithm is complicated, we introduce several components first.

Component-1: Item Transformation. In this component, we transform the item in each region to be homogeneous, i.e., each agent i's valuation on d_t is a constant for every $t \in [m]$. To see this, for each d_t, we partition it into multiple intervals on which each agent's valuation is a liner function (this can be done by taking the union of all the points of discontinuity for all agents' valuations). After that, we will restrict ourselves to the allocations where each agent i's share on d_t is parameterized by a parameter $x_{it} \in [0,1]$ such that agent i gets a x_{it} fraction of the length in each of these intervals. For any two agents $i,j \in N$, the valuation of agent j on agent i's allocation on d_t is then given by $x_{it} \cdot v_j(d_t)$. Thus, we can treat each agent i's valuation on each d_t as a constant. To simplify the description, we use v_{it} to denote agent i's valuation on d_t.

Component-2: About-To-Envy-Graph. Given a partial allocation $\mathcal{X} = (X_1,\ldots,X_n)$, we define the corresponding *about-to-envy-graph*, which is a directed edge-weighted graph $G = (V,E,w)$ where vertices represent agents and a directed edge (i,j) represents $v_i(X_i) = v_i(X_j)$. The weight

Algorithm 1: About-To-Envy-Satisfied Algorithm

Input: An instance $\mathcal{I} = \langle D, N, \boldsymbol{f} \rangle$ with piecewise constant functions $\boldsymbol{f}$
Output: An EF-with-bounded-charity allocation $\mathcal{X}$

Let $\mathcal{X} = (\emptyset, \dots, \emptyset)$;
Item Transformation;
Assign each agent i to R_t with the maximum v_{it};
Allocate the same positive length δ to each agent, where δ is the maximum possible such that the allocation is still feasible. In particular, δ is the value such that one region R_t is exhausted, i.e., $P_t = \emptyset$;
while *there is a region R_t such that P_t is envied by at least one agent* **do**
  Let agent i^* be the most envious agent of P_t;
  if *there exists an agent j^* such that there is a path $j^* \to \cdots \to i^*$ on the about-to-envy-graph G where $X_{j^*} \subseteq d_{t'}$ and $P_{t'} = \emptyset$.* **then**
    Let agent j^* return her bundle to the respective region $R_{t'}$;
    For each edge (i, j) on the path $j^* \to \cdots \to i^*$, allocate j's bundle to i;
    Allocate a length of $\frac{v_{i^*}(X_{i^*})}{v_{i^*t}}$ on P_t to i^*;
  else
    Let T be the set of all agents from which there is a path to agent i^*;
    while *no agent in T is in a region that is exhausted* **do**
      For each $i \in T$, allocates $\delta_i = \delta \cdot \frac{X_i}{X_{i^*}}$ to her where δ is the maximum value such that either some agent is added to T or some region becomes exhausted;
  Update the about-to-envy-graph G;
return $\mathcal{X}$

for the edge (i, j) is defined by $w_{ij} = \frac{v_{ir_j}}{v_{ir_i}}$ where $X_i \subseteq d_{r_i}$ and $X_j \subseteq d_{r_j}$.

Intuitively, an edge (i, j) implies that agent i is about to envy agent j if agent j receives any more part of one item. Moreover, if agent j receives an extra length of δ, agent i should obtain at least $w_{ij} \cdot \delta$ to maintain EF.

Component-3: Most Envious Agent[1]**.** Given a partial allocation $\mathcal{X} = (X_1, \dots, X_n)$, let P_t denote the unallocated fraction of the item d_t in R_t for any $t \in [m]$, i.e., $P_t = d_t \setminus \bigcup_{i=1}^n X_i$. If P_t is envied by at least one agent, we say that agent i^* is the *most envious agent* of P_t if $v_{i^*}(X_{i^*}) < v_{i^*}(P_t)$ and $\frac{v_{i^*}(X_{i^*})}{v_{i^*t}} \leq \frac{v_i(X_i)}{v_{it}}$ for any other agent i who also envies P_t (break ties arbitrarily). In other words, among all agents who envy P_t, agent i^* requires a minimum length interval of P_t such that replacing X_{i^*} with this interval keeps the value of agent i's bundle unchanged.

By our definition, if an allocation $\mathcal{X} = (X_1, \dots, X_n)$ is EF and agent i^* is the most envious agent of P_t, replacing X_{i^*} with an interval $I \subseteq P_t$ such that $v_{i^*}(X_{i^*}) = v_{i^*}(I)$ gives another EF allocation. The formal description of our algorithm can be found in Algorithm 1.

[1]This notion differs from that used in the next section.

Theorem 2. *Given an instance $\mathcal{I} = \langle D, N, \boldsymbol{f} \rangle$ with piecewise constant functions $\boldsymbol{f}$, the About-To-Envy-Satisfied Algorithm computes an EF-with-bounded-charity allocation $\mathcal{X} = (X_1, \dots, X_n)$ in polynomial time.*

Proof Sketch. It is easy to check that the **if** and **else** branches do not invalidate EF. Moreover, the algorithm continues as long as someone envies P_t for some region R_t. To conclude our theorem, it remains to show that our algorithm always terminates and runs in polynomial time.

For **if** branch, it is easy to see that it takes polynomial time. For **else** branch, there are at most $O(n)$ times to execute inner while loop and δ in the inner while loop can be done by linear programming. Then we only need to bound the number of the outer **while** loop.

Firstly, we show that **else** can only be applied for a polynomial number of times between two executions of **if**. To see this, each execution of **else** makes some region R_t exhausted (with $P_t = \emptyset$), which takes $O(k)$ times. Hence, we only need to bound the number of the executions of **if**.

Throughout the algorithm, we say that a region R_t is *undesirable* if it is/was exhausted, i.e., $P_t = \emptyset$. We can verify that once a region R_t is undesirable, P_t will no longer be envied by any other agents throughout the algorithm. Therefore, after each execution of **if**, an agent is removed from an undesirable region. In addition, the agents in one undesirable region can only be removed once throughout the algorithm. Since we have k regions, the total number of executions of **if** can be bounded by $O(nk)$. □

Equal-Region

For equal-region setting, We first present the upper bounds of EF and PROP allocations.

Proposition 2. *There exists an instance for which no allocation is $(\frac{1}{2} + \epsilon)$-PROP for any $\epsilon > 0$.*

Corollary 2. *There exists an instance for which no allocation is $(\frac{1}{2} + \epsilon)$-EF for any $\epsilon > 0$.*

Under piecewise constant functions, we observe that if we allocate the unallocated part to any agent in the respective region after the About-To-Envy-Satisfied Algorithm, a $\frac{1}{2}$-EF allocation can be guaranteed since the allocation computed by the About-To-Envy-Satisfied Algorithm is EF-with-bounded-charity.

Proposition 3. *Allocating the unallocated part to any agent in the respective region on top of the output allocation by the About-To-Envy-Satisfied Algorithm can guarantee $\frac{1}{2}$-EF.*

Next, we propose the following algorithm named Adapted Perfect Partition Algorithm to compute an approximate EF allocation with general valuation functions. The intuition of our algorithm relies on *perfect partition* in cake cutting.

Theorem 3. *The Adapted Perfect Partition Algorithm computes a $\frac{1}{2}$-EF allocation.*

We can observe that if $n \bmod k = 0$, the Adapted Perfect Partition Algorithm can compute an exact EF allocation. If $n >> k$, it can compute a nearly exact EF allocation.

Algorithm 2: Adapted Perfect Partition Algorithm

Input: An instance $\mathcal{I} = \langle D, N, \boldsymbol{f} \rangle$
Output: An approximate EF allocation $\mathcal{X}$
Let $\mathcal{X} = (\emptyset, \dots, \emptyset)$;
Assign agents to regions by Round-Robin algorithm;
if $n \bmod k = 0$ **then**
  Divide d_j in R_j into $\frac{n}{k}$ *perfect* pieces for every $j \in [k]$;
else
  Divide d_j in R_j into $\lfloor \frac{n}{k} \rfloor \cdot (\lfloor \frac{n}{k} \rfloor + 1)$ *perfect* pieces for every $j \in [k]$;
Allocate these *perfect* pieces to agent(s) in R_j by Round-Robin algorithm;
return $\mathcal{X}$

Indivisible Items

In this section, we study the case of allocating indivisible items from multiple regions. We first consider the diverse-region setting and then the equal-region setting.

Diverse-Region

We first show that it is NP-complete to decide whether a given instance admits an allocation that satisfies EF1/EFX.

Theorem 4. *Given an instance $\mathcal{I} = \langle M, N, R, \boldsymbol{v} \rangle$, it is NP-complete to decide whether there exists an allocation that satisfies EF1/EFX. The hardness result holds even if all agents have the same valuation function.*

Proof Sketch. The problem is clearly in NP: given an allocation, we can check the property between every pair of agents. To prove NP-hardness, we reduce from the PARTITION.

Consider an instance from PARTITION with a set of positive integers $S = \{a_1, \dots, a_m\}$. We create an instance of $\mathcal{I} = \langle M, N, R, \boldsymbol{v} \rangle$ where $|N| \geq 3$ and $|R| = 2$, and all agents have the same valuation function. For region R_1, we have m items where each has the value of a_i, respectively. For region R_2, we have the items where each has the value of ϵ and the total value is $\frac{|N|-2}{2} \cdot (a_1 + \cdots + a_m)$. Then, we can verify that there is an equitable partition if and only if we can divide items in R_1 into two equitable bundles. □

Next, we present the impossible results of approximate EF1/EFX and PROP1/PROPX allocations.

Proposition 4. *There exists an instance for which no allocation is α-EF1/EFX for any $\alpha \in (0, 1]$.*

Proposition 5. *There exists an instance for which no allocation is α-PROP1/PROPX for any $\alpha \in (0, 1]$.*

Similar to divisible items, we have impossible results for both the decision problem and the approximation. The impossibility results can be bypassed by partial allocations.

Next, we propose the Sequential Envy-Satisfied Algorithm to compute an EFX-with-bounded-charity allocation. Inspired by techniques in Chaudhury et al. (2021), the main idea of our algorithm is that in each iteration, we find the *most envious agent* and reallocate a new bundle to this agent

Algorithm 3: Sequential Envy-Satisfied Algorithm

Input: An instance $\mathcal{I} = \langle M, N, R, \boldsymbol{v} \rangle$
Output: An EFX-with-bounded-charity allocation $\mathcal{X}$
Let $\mathcal{X} = (\emptyset, \dots, \emptyset)$;
while *there is a region R_t such that P_t is envied by at least one agent* **do**
  Let agent i^* be the most envious agent of P_t, and $S \subset P_t$ be an inclusion-wise minimal subset of P_t such that $v_{i^*}(X_{i^*}) < v_{i^*}(S)$;
  Let agent i^* return her bundle to the respective region and allocate S to agent i^*;
return $\mathcal{X}$

while guaranteeing that the other agents only envy this agent up to any item after the reallocation. The formal description of our algorithm can be found in Algorithm 3.

Most Envious Agent (Chaudhury et al. 2021). Given an allocation $\mathcal{X} = (X_1, \dots, X_n)$, let P_t denote the set of unallocated items in R_t for ay $t \in [m]$, i.e., $P_t = M_t \setminus \bigcup_{i=1}^n X_i$. Assume that P_t is envied by at least one agent. If agent i envies P_t, let S_i be any inclusion-wise minimal subset of P_t such that $v_i(X_i) < v_i(S_i)$ (break ties arbitrarily). We say that agent i^*, such that $v_{i^*}(X_{i^*}) < v_{i^*}(S_{i^*})$ for some $S_{i^*} \subset P_t$ and no strict subset of S_{i^*} is envied by any agent, is the *most envious agent* of P_t (break ties arbitrarily).

We can observe that the Sequential Envy-Satisfied Algorithm terminates in finite steps since there is an agent whose bundle's value increases and other agents whose bundles' value is unchanged in each iteration, and the maximum value of each agent's bundle is bounded.

Theorem 5. *Given an instance $\mathcal{I} = \langle M, N, R, \boldsymbol{v} \rangle$, the Sequential Envy-Satisfied Algorithm computes an EFX-with-bounded-charity allocation $\mathcal{X} = (X_1, \dots, X_n)$.*

However, one can see that the Sequential Envy-Satisfied Algorithm is a pseudo-polynomial algorithm. It is still an open problem whether a polynomial algorithm exists. Next, we propose a polynomial algorithm named Sequential Envy-Graph-Satisfied Algorithm that leverages the Envy-Cycles Procedure proposed by Lipton et al. (2004) to compute an EF1-with-bounded-charity allocation.

The intuition of our algorithm leverages both the advantages of the Sequential Envy-Satisfied Algorithm and the Envy-Cycles Procedure. Note that in our setting, the Envy-Cycles Procedure will not proceed if all source agents are in exhausted regions. Our algorithm overcomes this difficulty by moving an agent out without breaking EF1. The formal description of our algorithm can be found in Algorithm 4.

Theorem 6. *Given an instance $\mathcal{I} = \langle M, N, R, \boldsymbol{v} \rangle$, the Sequential Envy-Graph-Satisfied Algorithm computes an EF1-with-bounded-charity allocation $\mathcal{X} = (X_1, \dots, X_n)$ in polynomial time.*

Proof. It is easy to check that the inner **while** loop does not invalidate EF1. For **if** branch, we can verify that no one envies agent i^* up to one item due to the definition of the most envious agent. For the agents who receive the new bundles, the other agents do not envy them up to one item since those bundles are the previous existing bundles. In addition, the agents who receive the new bundles still envy the other agents up to one item due to the increase in their bundles' value and the envy graph. Moreover, the algorithm continues as long as someone envies P_t for some region R_t. To conclude our theorem, it remains to show that our algorithm always terminates and runs in polynomial time.

For the inner **while** loop, it is easy to see that it takes polynomial time. For **if** branch, since there is no envy cycle (otherwise, the algorithm will not jump to the **if** branch). The moving will take at most n steps. Then, we only need to bound the number of the outer **while** loop, which is equivalent to bounding the number of the executions of **if**.

Throughout the algorithm, we say that a region R_t is *undesirable* if it is/was exhausted, i.e., $P_t = \emptyset$. Since the total value of items in P_t will not increase, and the value of all agents' bundles will not decrease. Hence, once a region R_t is undesirable, P_t will no longer be envied by any other agents throughout the algorithm. Therefore, after each execution of **if**, a source agent is removed from an undesirable region. In addition, the agents in one undesirable region can only be removed once throughout the algorithm. There are k regions and the total number of executions of **if** can therefore be bounded by $O(nk)$. □

Equal-Region

We first show that any bounded approximate EFX or PROPX allocation still cannot be guaranteed even for the equal-region. Then, we give constant factor approximation upper bounds for the EF1 and PROP1 allocations.

Proposition 6. *There exists an instance for which no allocation is α-EFX for any $\alpha \in (0, 1]$.*

Proposition 7. *There exists an instance for which no allocation is α-PROPX for any $\alpha \in (0, 1]$.*

Proposition 8. *There exists an instance for which no allocation is $(\frac{1}{2} + \epsilon)$-PROP1 for any $\epsilon > 0$.*

Because EF1 implies PROP1, we can directly have the following corollary.

Corollary 3. *There exists an instance for which no allocation is $(\frac{1}{2} + \epsilon)$-EF1 for any $\epsilon > 0$.*

Next, we show that both the Sequential Envy-Satisfied Algorithm and the Sequential Envy-Graph-Satisfied Algorithm can give a $\frac{1}{2}$-EF1 allocation by allocating the unallocated items to the agents in the respective regions.

Proposition 9. *Allocating the unallocated items to any agent in the respective region on top of the output allocation by the Sequential Envy-Satisfied Algorithm or the Sequential Envy-Graph-Satisfied Algorithm can guarantee $\frac{1}{2}$-EF1.*

Besides the approximate EF1, we are also interested in approximate PROP1. Next, we give an approximate PROP1 algorithm which runs in polynomial time. The main idea of this algorithm is based on the Round-Robin algorithm.

Theorem 7. *The Adapted Round-Robin Algorithm computes a $\frac{1}{2}$-PROP1 allocation in polynomial time.*

Algorithm 4: Sequential Envy-Graph-Satisfied Algorithm

Input: An instance $\mathcal{I} = \langle M, N, R, \boldsymbol{v} \rangle$
Output: An EF1-with-bounded-charity allocation $\mathcal{X}$
Let $\mathcal{X} = (\emptyset, \dots, \emptyset)$;
while *there is a region R_t such that P_t is envied by at least one agent* **do**
  while *there is a directed cycle in the envy graph* ***or*** *there is a source agent in the envy-graph and the region where she picks item before has unallocated items* **do**
    Envy-Cycles Procedure;
  Let agent i^* be the most envious agent of P_t;
  if *there exists a source agent j^* such that there is a path $j^* \rightarrow \cdots \rightarrow i^*$ on the envy-graph G where $X_{j^*} \subseteq M_{t'}$ and $P_{t'} = \emptyset$* **then**
    Let agent j^* return her bundle to the respective region $R_{t'}$;
    For each edge (i, j) on the path $j^* \rightarrow \cdots \rightarrow i^*$, allocate agent j's bundle to agent i;
    Let $S \subset P_t$ be an inclusion-wise minimal subset of P_t such that $v_{i^*}(X_{i^*}) < v_{i^*}(S)$ and allocate S to agent i^*;
  Update envy-graph G;
return $\mathcal{X}$

Algorithm 5: Adapted Round-Robin Algorithm

Input: An instance $\mathcal{I} = \langle M, N, R, \boldsymbol{v} \rangle$
Output: An approximate PROP1 allocation $\mathcal{X}$
Let $\mathcal{X} = (\emptyset, \dots, \emptyset)$;
Assign agents to regions by Round-Robin algorithm;
for $j = 1$ **to** k **do**
  Use Round-Robin algorithm to allocate items in R_j to agents who are assigned to R_j;
return $\mathcal{X}$

Similar to the Adapted Perfect Partition Algorithm, if $n \bmod k = 0$, the Adapted Round-Robin Algorithm computes an exact PROP1 allocation. If $n >> k$, it can compute a nearly exact PROP1 allocation.

Conclusion

In this paper, we initiate the study of fair allocation of items in multiple regions where items can be divisible or indivisible. On the negative side, we show the hardness of the respective decision problem and some inapproximability results. On the positive side, we present the algorithms which can guarantee partial envy-based notions, and tight bounds of approximation ratios in special settings. In future work, we may consider the fair allocation of mixed items in multiple regions, where each region has both one divisible item and some indivisible items.

Acknowledgments

The work described in this paper was partially supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China [Project No. CityU 11213620]. The research of Biaoshuai Tao was supported by the National Natural Science Foundation of China (No. 62102252). We thank the anonymous AAAI-24 reviewers for their very useful feedback.

References

Amanatidis, G.; Aziz, H.; Birmpas, G.; Filos-Ratsikas, A.; Li, B.; Moulin, H.; Voudouris, A. A.; and Wu, X. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence*, 322: 103965.

Aziz, H.; and Mackenzie, S. 2016a. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *IEEE 57th Annual Symposium on Foundations of Computer Science*, 416–427.

Aziz, H.; and Mackenzie, S. 2016b. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of the 48th annual ACM symposium on Theory of Computing*, 454–464.

Aziz, H.; Moulin, H.; and Sandomirskiy, F. 2020. A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation. *Operations Research Letters*, 48(5): 573–578.

Bei, X.; Chen, N.; Hua, X.; Tao, B.; and Yang, E. 2012. Optimal proportional cake cutting with connected pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, 1263–1269.

Brams, S. J.; and Taylor, A. D. 1995. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1): 9–18.

Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6): 1061–1103.

Caragiannis, I.; Gravin, N.; and Huang, X. 2019. Envy-freeness up to any item with high Nash welfare: The virtue of donating items. In *Proceedings of the 20th ACM Conference on Economics and Computation*, 527–545.

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3): 1–32.

Cechlárová, K.; and Pillárová, E. 2012. On the computability of equitable divisions. *Discrete Optimization*, 9(4): 249–257.

Chaudhury, B. R.; Kavitha, T.; Mehlhorn, K.; and Sgouritsa, A. 2021. A little charity guarantees almost envy-freeness. *SIAM Journal on Computing*, 50(4): 1336–1358.

Chen, Y.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1): 284–297.

Cloutier, J.; Nyman, K. L.; and Su, F. E. 2010. Two-player envy-free multi-cake division. *Mathematical Social Sciences*, 59(1): 26–37.

Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation*, 629–646.

Dubins, L. E.; and Spanier, E. H. 1961. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1P1): 1–17.

Edmonds, J.; and Pruhs, K. 2006. Balanced allocations of cake. In *47th Annual IEEE Symposium on Foundations of Computer Science*, 623–634.

Even, S.; and Paz, A. 1984. A note on cake cutting. *Discrete Applied Mathematics*, 7(3): 285–296.

Hosseini, H.; Igarashi, A.; and Searns, A. 2020. Fair Division of Time: Multi-layered Cake Cutting. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, 182–188.

Lebert, N.; Meunier, F.; and Carbonneaux, Q. 2013. Envy-free two-player m-cake and three-player two-cake divisions. *Operations Research Letters*, 41(6): 607–610.

Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, 125–131.

Moulin, H. 2004. *Fair division and collective welfare*. MIT press.

Moulin, H. 2019. Fair division in the internet age. *Annual Review of Economics*, 11: 407–441.

Nyman, K.; Su, F. E.; and Zerbib, S. 2020. Fair division with multiple pieces. *Discrete Applied Mathematics*, 283: 115–122.

Procaccia, A. D.; and Wang, J. 2014. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 15th ACM Conference on Economics and Computation*, 675–692.

Procaccia, A. D.; and Wang, J. 2017. A lower bound for equitable cake cutting. In *Proceedings of the 18th ACM Conference on Economics and Computation*, 479–495.

Segal-Halevi, E. 2021. Fair multi-cake cutting. *Discrete Applied Mathematics*, 291: 15–35.

Steinhaus, H. 1948. The problem of fair division. *Econometrica*, 16: 101–104.

Steinhaus, H. 1949. Sur la division pragmatique. *Econometrica: Journal of the Econometric Society*, 17: 315–319.

Tao, B. 2022. On existence of truthful fair cake cutting mechanisms. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, 404–434.