

Weighted Maxmin Fair Share Allocation of Indivisible Chores

Haris Aziz¹, Hau Chan² and Bo Li³

¹UNSW Sydney and Data61 CSIRO, Australia

²Department of Computer Science and Engineering, University of Nebraska-Lincoln, USA

³Department of Computer Science, Stony Brook University, USA

haziz@cse.unsw.edu.au, hchan3@unl.edu, boli2@cs.stonybrook.edu

Abstract

We initiate the study of indivisible chore allocation for agents with asymmetric shares. The fairness concept we focus on is the weighted natural generalization of maxmin share: WMMS fairness and OWMMS fairness. We first highlight the fact that commonly-used algorithms that work well for the allocation of goods to asymmetric agents, and even for chores to symmetric agents do not provide good approximations for allocation of chores to asymmetric agents under WMMS. As a consequence, we present a novel polynomial-time constant-approximation algorithm, via linear program, for OWMMS. For two special cases: the binary valuation case and the 2-agent case, we provide exact or better constant-approximation algorithms.

1 Introduction

We consider fair allocation of indivisible chores when agents have asymmetric shares. In contrast to the case of goods for which agents have positive value, chores are disliked by agents and they have negative values for them. The fairness concept we focus on is the maxmin share (MMS) fairness which was designed for allocation of indivisible items. MMS is based on the thought experiment that if the items are partitioned into bundles and an agent would always get the least preferred bundle of items, what is the best way she can partition the items. The value of such a bundle is the maxmin share of the agent. An allocation is deemed MMS fair if each agent gets her required share.

Maxmin share fairness was proposed by [Budish, 2011] as a fairness concept for allocation of indivisible items. It is a relaxation of *proportionality* fairness that requires each of the n agents should get a value that is at least $1/n$ of the total value she has for the set of all items. When items are divisible, maxmin share fairness coincides with proportionality. Maxmin share fairness is a weaker concept when items are indivisible. It was conjectured that a maxmin fair allocation always exists but [Procaccia and Wang, 2014] identified a counter-example. Since the work of [Procaccia and Wang, 2014], there are several papers on algorithms that find an approximate MMS allocation [Amanatidis *et al.*, 2015;

Barman and Murthy, 2017; Ghodsi *et al.*, 2018; Aziz *et al.*, 2017]. All these works make a typical assumption that agents are symmetric and should be treated in a similar manner.

[Farhadi *et al.*, 2017] were the first to consider MMS fairness for the case where indivisible goods are allocated and the agents are *not* symmetric because they may have different entitlement share of the goods. Ideally, an agent would expect to get a share of the total value that is proportional to her entitlement. However, when items are indivisible, MMS fairness needs to be suitably generalized to cater for asymmetric entitlement shares. Farhadi *et al.* generalized MMS fairness to that of the more general MMS concept as *weighted MMS (WMMS)* that caters for entitlements. They devised a simple ordinal (that only used the qualitative ranking information of items) algorithm that ensures an n -approximation guarantee for WMMS where each agent's allocation is at least $1/n$ of her value in a WMMS allocation. Beyond the results for goods [Farhadi *et al.*, 2017; Farhadi *et al.*, 2019], not much is known about chore allocation when the agents are asymmetric despite the recent active research in fair allocation of goods and chores. Furthermore, it is not clear whether the results for goods from one setting could carry over the other [Aziz, 2016].

In this paper, we focus on the fair allocation of *chores* rather than goods for asymmetric agents. In the case of chores, agents do not have entitlements but relative shares. If an agent has a higher share, she is expected to take a higher load of the chores. Treating agents asymmetrically may be a requirement for several reasons. For example, countries with a larger population and CO2 emission may be liable to undertake more responsibility to clean up the environment. In this paper, the central research question we examine is the following one. *When indivisible chores are to be allocated among agents with asymmetric shares, for what approximation factor do approximately WMMS fair allocations exist and how efficiently can they be computed?*

Our contributions are as follows. We consider a model of allocation of chores in which agents have relative shares as compared to entitlements. Different to the case of symmetric agents, we first prove that even with only two agents, no algorithm can simultaneously guarantee each agent's value to be higher than $\frac{4}{3}$ of her weighted maxmin share. Moreover, we show that many greedy algorithms widely used in the literature, including [Farhadi *et al.*, 2017] and [Aziz *et al.*, 2017],

may have arbitrarily bad performance.

Then we design a polynomial-time algorithm which provides a 4-approximation to the minimal relaxation of WMMS value (OWMMS) under which a WMMS allocation exists. To present this algorithm, we first study a special case when all agents have an identical valuation. The algorithm combines (1) the use of a greedy algorithm for the case of identical valuations and (2) linear programming and rounding techniques.

Finally, we study two restricted cases: a two-agent setting and a binary valuation setting. For the two-agent case, we present a variant of divide-and-choose protocol which ensures each agent's value is at least $\frac{3}{2}$ of her weighted maxmin share; For binary valuations, we show that a WMMS allocation exists and can be efficiently computed. For asymmetric agents and indivisible items, this is the first algorithmic result for binary valuations.

2 Related Work

The fair allocation problem has been extensively studied in the cake cutting literature [Dubins and Spanier, 1961; Stromquist, 1980; Alon, 1987; Brams and Taylor, 1995; Brams and Taylor, 1996; Robertson and Webb, 1998; Aziz and Mackenzie, 2016]. In this line of work, researchers study how to fairly allocate a divisible item (e.g., cake) among a number of agents. Solution concepts such as envy-freeness and proportionality are prominent criteria for fairness. In the context of divisible goods, researchers have extended results for the case of equal entitlements to those of unequal entitlements (see e.g., [Cseh and Fleiner, 2018]).

As for MMS fairness, it is already known that even for additive valuations, there exists an instance such that no allocation can simultaneously guarantee each agent receives at least her MMS [Kurokawa *et al.*, 2018]. But approximate MMS can be efficiently computed; see [Barman and Murthy, 2017; Ghodsi *et al.*, 2018; Kurokawa *et al.*, 2018]. Computing WMMS shares is an NP-hard problem for both goods and for chores even for the case of 2 agents and for equal shares. The statement can be derived via a reduction from the integer partition problem [Garey and Johnson, 1979].

Most of the work on fair allocation of items is for the case of goods although recently, fair allocation of chores [Aziz *et al.*, 2017] or combinations of goods and chores [Aziz *et al.*, 2018] has received attention as well. It is shown by [Aziz *et al.*, 2017] that MMS allocations for chores do not always exist but can be 2-approximated by a simple round-robin algorithm. [Aziz *et al.*, 2017] also presented a PTAS for relaxation of MMS called optimal MMS. [Barman and Murthy, 2017] presented an improved approximation algorithm for MMS allocation of chores. Fair allocation of indivisible goods and asymmetric agents has also been studied [Farhadi *et al.*, 2017; Farhadi *et al.*, 2019]. We take a similar approach and study the chore allocation problem when the agents are not symmetric. [Babaioff *et al.*, 2017] considered the allocation of indivisible goods where agents have different entitlements. One of the concepts that they propose is called ℓ -out-of-d MMS that can also apply to agents having ordinal preferences over bundles of chores. However, the paper focusses on results for goods.

3 Preliminaries

We begin by presenting our setting formally and discussing fairness concepts as well as some basic notations in the paper.

3.1 Setting

Let $N = \{1, 2, \dots, n\}$ be a set of n agents, and $M = \{1, 2, \dots, m\}$ be a set of m indivisible items. In this work, we always use $i \in N$ and $j \in M$ to indicate an agent and an item, separately. Each agent has a valuation function $V_i : 2^M \rightarrow \mathbb{R}$. Denote by $V_{ij} = V_i(\{j\})$. We assume that items are chores to every agent, i.e., $V_{ij} \leq 0$ for all $j \in M$ and the valuations are additive, i.e., for any $S \subseteq M$, $V_i(S) = \sum_{j \in S} V_{ij}$. Without loss of generality and just for ease of presentation, throughout this paper except Section 6.2, it is assumed that all of the valuations are normalized, i.e. $V_i(\emptyset) = 0$ and $V_i(M) = -1$.

In this work, we consider the case when agents are asymmetric. Particularly, every agent has a share for the chores, namely $s_i \in (0, 1]$. The shares add up to 1, i.e., $\sum_{i \in N} s_i = 1$.

Letting $\mathbf{V} = (V_1, \dots, V_n)$ and $\mathbf{s} = (s_1, \dots, s_n)$, we use $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ to denote a chore allocation instance and $\mathcal{I} = (N, M, \mathbf{s}, V)$ when all agents have the identical valuation V . Note that when all agents have identical valuation V , $V(\{j\})$ is simplified as V^j for any $j \in M$. Let $\Pi(M)$ be the set of all n -partitions of the items. A generic allocation will be denoted by $X = \langle X_1, \dots, X_n \rangle$ where X_i is the bundle of agent i .

3.2 WMMS Fairness

Before presenting the WMMS fairness concept that takes into account the shares of the agents, we first present the standard MMS fairness concept that assumes the shares of the agents are equal. For symmetric agents, the classical *maxmin share* (MMS) of an agent i with valuation V_i is defined as

$$\text{MMS}_i = \max_{\langle X_i \rangle_{i \in N} \in \Pi(M)} \min_{j \in N} V_i(X_j).$$

Intuitively, when allocating items to n agents, each agent should get an allocation with a value that is $1/n$ of the total value they have for all the items. Since the items are not divisible, this proportionality requirement may be not achievable for the agents. In view of this, MMS_i can be viewed as a relaxed lower bound on the value that agent i hopes for if she has the chance to partition the items into n bundles and every other agent adversarially chooses a bundle before i . Next, we generalize the classical MMS notion to the setting with asymmetric agents.

Definition 1 (Weighted MMS) *Given any chore allocation instance $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$, for every agent $i \in N$, the weighted maxmin share (WMMS) value of i is defined as:*

$$\text{WMMS}_i(\mathcal{I}) = \max_{\langle X_i \rangle_{i \in N} \in \Pi(M)} \min_{j \in N} V_i(X_j) \frac{s_i}{s_j}.$$

Any partition achieves $\text{WMMS}_i(\mathcal{I})$ is called a P - i partition.

When the instance \mathcal{I} is clear from the context, we may use WMMS_i for short. The definition above for WMMS fairness is exactly the same as that of WMMS as formalized by

[Farhadi *et al.*, 2017] for the case of goods except that the entitlement e_i of an agent i is replaced by her share s_i . As mentioned in the introduction, whereas a higher entitlement for goods is desirable for an agent, a higher share for chores is undesirable for the agent.

We call an allocation WMMS if the value of the allocation to each agent i is worth at least WMMS_i to her. Similarly, an allocation is called α -WMMS, if the total value of items allocated to each agent i is at least αWMMS_i for $\alpha \geq 1$.

Note that when all shares are equal, WMMS coincides with MMS fairness so it is a proper generalization of MMS. Secondly, we spell out an insight that also provides justification for the WMMS concept that was defined by [Farhadi *et al.*, 2017]. We note that when the items are *divisible*, then $\text{WMMS}_i = s_i V_i(M)$. Hence, for divisible chores, WMMS fairness also implies a natural generalization of proportionality that takes into account the shares of agents. We call the latter requirement as *weighted proportionality*.

In the following, we define some more notation that will be used in the paper. Given a chore allocation instance $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$, for any agent i and any partition $X = \langle X_i \rangle_{i \in N}$, let $W_i^{\mathcal{I}}(X) = \min_{k \in N} \frac{V_i(X_k)}{s_k}$. That is, $W_i^{\mathcal{I}}(X)$ is the unfairness degree of allocation X to i . Let $W_i(\mathcal{I}) = \max_{\langle X_1, \dots, X_n \rangle \in \Pi(\mathcal{M})} W_i^{\mathcal{I}}(X)$. Thus $W_i(\mathcal{I})$ is the smallest degree of unfairness and $\text{WMMS}_i(\mathcal{I}) = s_i W_i(\mathcal{I})$. Moreover, we have the following simple properties.

Lemma 1 *Given any instance $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$, for any $i \in N$, $W_i(\mathcal{I}) \leq -1$, and $\text{WMMS}_i(\mathcal{I}) \leq -s_i$.*

Proof: Note that for any agent i and any allocation $\langle X_i \rangle_{i \in N}$, $\sum_{k \in N} \frac{V_i(X_k)}{s_k} \cdot s_k = \sum_{k \in N} V_i(X_k) = -1$, which is the weighted arithmetic mean of the terms $\frac{V_i(X_k)}{s_k}$, with weights s_k (whose sum is 1). As the mean equals -1 , the smallest must be at most -1 and $W_i(\mathcal{I}) = \min_{k \in N} \frac{V_i(X_k)}{s_k} \leq -1$. \square

Next we show a simple algorithm, Naive, which returns an n -WMMS allocation. Algorithm Naive produces an allocation that allocates all of the items to a single agent who has the highest share (ties are broken arbitrarily).

Lemma 2 *Let $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ be any chore allocation instance and $\langle X_i \rangle_{i \in N}$ be the output of Algorithm Naive. Then $V_i(X_i) \geq n \text{WMMS}_i(\mathcal{I})$ for any $i \in N$.*

Proof: Let i^* be the agent who has the largest share, thus $s_{i^*} \geq \frac{1}{n}$. It is easy to see that for any agent $i \neq i^*$, $V_i(X_i) = 0$, which is trivially at least as large as $n \text{WMMS}_i(\mathcal{I})$. By Lemma 1, $\text{WMMS}_{i^*}(\mathcal{I}) \leq -s_{i^*} \leq -\frac{1}{n}$. Accordingly, $V_{i^*}(M) \geq n \text{WMMS}_{i^*}(\mathcal{I})$. \square

4 Optimal WMMS Fairness

It is well known that for symmetric agents, no matter the items are goods or chores, an MMS allocation always exists for the 2-agent case. But for asymmetric agents, we note that an exact WMMS allocation may not exist even when there are only two agents. Indeed, by the following lemma, we see that the lower bound of the problem is at least $\frac{4}{3}$, which means that

there is no allocation that can guarantee each agent's value to be greater than $\frac{4}{3} \text{WMMS}_i(\mathcal{I})$ for every $i \in N$.

Lemma 3 *In the chore allocation problem, any algorithm has an approximation ratio of at least $\frac{4}{3}$ for WMMS fairness.*

Proof: In the following we construct an instance $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ with $N = \{1, 2\}$, $M = \{1, 2\}$ and the shares and valuations are shown in Table 1.

Agent	Share	Items	
		1	2
1	$\frac{3}{4}$	$-\frac{3}{4}$	$-\frac{1}{4}$
2	$\frac{1}{4}$	$-\frac{1}{2}$	$-\frac{1}{4}$

Table 1: Instance to establish the $\frac{4}{3}$ lower bound for 2 agents.

We first note that for agent 1, the unique P-1 partition is $X = \langle X_1, X_2 \rangle$ with $X_1 = \{1\}$ and $X_2 = \{2\}$, since $W_1^{\mathcal{I}}(X) = \min \left\{ \frac{-\frac{3}{4}}{\frac{3}{4}}, \frac{-\frac{1}{4}}{\frac{1}{4}} \right\} = -1$, which is the largest among all possible allocations. Accordingly, $\text{WMMS}_1(\mathcal{I}) = \frac{3}{4} \times (-1) = -\frac{3}{4}$. For agent 2, to maximize $W_2^{\mathcal{I}}$, the only way is to set $X' = \langle X'_1, X'_2 \rangle$ with $X'_1 = \{1, 2\}$ and $X'_2 = \emptyset$, since $W_2^{\mathcal{I}}(X') = \min \left\{ \frac{-1}{\frac{3}{4}}, \frac{0}{\frac{1}{4}} \right\} = -\frac{4}{3}$, which is the largest among all possible allocations. Accordingly, $\text{WMMS}_2(\mathcal{I}) = \frac{1}{4} \times (-\frac{4}{3}) = -\frac{1}{3}$.

However, X is bad to agent 2, since $V_2(X_2) = -\frac{1}{2} < \text{WMMS}_2(\mathcal{I})$ and X' is bad to agent 1, since $V_1(X'_1) = -1 < \text{WMMS}_1(\mathcal{I})$. Therefore the best tradeoff to satisfy the two agents simultaneously would be allocation X' , since $V_2(X_2) = \frac{3}{2} \text{WMMS}_2(\mathcal{I})$ and $V_1(X'_1) = \frac{4}{3} \text{WMMS}_1(\mathcal{I})$. Indeed, we need to take all possible allocations into consideration, but it is easy to see that all other allocations can only be worse.

Thus, no algorithm could provide an allocation with each agent i 's value being strictly larger than $\frac{4}{3} \text{WMMS}_i(\mathcal{I})$, which finishes the proof of Lemma 3. \square

Accordingly, it is natural to consider a relaxed version of WMMS, *optimal WMMS (OWMMS) fairness*, which is similar to the one introduced in [Aziz *et al.*, 2017].

Definition 2 (Optimal WMMS) *Let $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ be a chore allocation instance. The optimal WMMS (OWMMS) ratio α^* is defined as the minimal $\alpha \in [1, \infty)$ for which an α -WMMS allocation always exists. Let $\text{OWMMS}_i(\mathcal{I}) = \alpha^* \text{WMMS}_i$ for any $i \in N$. A partition $X = \langle X_1, \dots, X_n \rangle$ is called an OWMMS allocation, if $V_i(X_i) \geq \text{OWMMS}_i(\mathcal{I})$ for all $i \in N$.*

It is easy to see that $\text{WMMS}_i(\mathcal{I}) \geq \text{OWMMS}_i(\mathcal{I})$ for any instance \mathcal{I} and any agent i . For any partition $X = \langle X_1, \dots, X_n \rangle$, if $V_i(X_i) \geq c \cdot \text{OWMMS}_i(\mathcal{I})$ for all $i \in N$, then X is called c -approximation to the OWMMS allocation.

5 Approximation Algorithms

For the case of goods allocation, the greedy *round robin* algorithm considered by [Farhadi *et al.*, 2017] gives the best guarantee (of n -approximation for goods). Interestingly, the same algorithm was proved to provide a 2-approximation for MMS allocation of chores when agents are symmetric [Aziz *et al.*, 2017]. However, when agents have different shares, such an algorithm can be arbitrarily poor. Due to space limit, we provide the bad example in the full version, where we also show that natural attempts to ‘fix’ the bad performance of the greedy algorithm do not help.

In the following, we give our polynomial-time $(4 + \epsilon)$ -approximation algorithm. That is, for any $\epsilon > 0$, it returns an allocation $\langle X_i \rangle_{i \in N}$ such that for any agent i , $V_i(X_i) \geq (4 + \epsilon)\text{OWMMS}_i$. In order to present the main algorithm, we first present a polynomial-time algorithm which guarantees each agent i ’s value to be at least 2WMMS_i , when all of the agents have an identical valuation.

5.1 Identical Valuation

When all agents have an identical valuation, we show the algorithm, *EgalGreedy* defined in Algorithm 1, is a 2-approximation to an exact WMMS allocation.

Algorithm 1 *EgalGreedy* - An Algorithm for Identical Valuations

Require: Chore allocation instance (N, M, s, V)

Ensure: Allocation $X = \langle X_1, \dots, X_n \rangle$.

- 1: Initially, $X_i = \emptyset$ for all $i \in N$.
 - 2: Order all chores from the lowest value to the highest value such that $V^1 \leq V^2 \leq \dots \leq V^m$.
 - 3: **for** $j = 1$ to m **do**
 - 4: $i^* \in \arg \max_{i \in N} \frac{V(X_i \cup \{j\})}{s_i}$;
 - 5: $X_{i^*} = X_{i^*} \cup \{j\}$.
 - 6: **return** Allocation X .
-

The next lemma relies on a connection to the parallel processors scheduling problem. In this problem, there is a set of jobs and a set of processors. Each job has to be processed exactly once on exactly one processor. Processors may have different speeds [Gonzalez *et al.*, 1977; Friesen, 1987]. The problem specifies the time required to process a given job on a given machine. Typically, the goal of scheduling problems is to find an assignment of the jobs such that the longest finishing time (i.e., makespan) is minimized. A detailed survey of this line of work can be found in [Pinedo, 2016]. Due to space limit, we prove Lemma 4 in the full version.

Lemma 4 *For any chore allocation instance $\mathcal{I} = (N, M, s, V)$, where all agents have the identical valuation V , let $\langle X_i \rangle_{i \in N}$ be the allocation outputted by *EgalGreedy*. We have $V(X_i) \geq 2\text{WMMS}_i(\mathcal{I})$ for any $i \in N$.*

One may suspect that a natural generalization of *EgalGreedy* to the case that agents have different valuations may work well. Unfortunately, in the full version, we provide an example that such an algorithm cannot have any constant approximation ratio.

5.2 General Valuations

Now we are ready to study the general case when agents may have different valuations. For any chore allocation instance $\mathcal{I} = (N, M, s, V)$, let variable α represent the WMMS ratio, and variable $x_{ij} \in \{0, 1\}$ represent whether agent i gets item j . Let $\mathbf{x} = (x_{ij})_{i \in N, j \in M}$. Then the problem of computing its OWMMS ratio and an OWMMS allocation can be formalized as the following integer program.

$$\mathcal{IP} : \begin{array}{ll} \min & \alpha \\ \text{s.t.} & \begin{cases} \sum_{j \in M} V_{ij} x_{ij} \geq \alpha \text{WMMS}_i(\mathcal{I}), & \forall i \in N \\ \sum_{i \in N} x_{ij} = 1, & \forall j \in M \\ x_{ij} \in \{0, 1\}, & \forall i \in N, j \in M \\ \alpha \geq 1. \end{cases} \end{array}$$

To solve \mathcal{IP} , in what follows, we first prove a key technical lemma by using the rounding technique introduced by [Lenstra *et al.*, 1990], which gives us the tool to round a fractional assignment to an integer assignment.

Lemma 5 *Let (N, M, s, V) be any chore allocation instance, $\mathbf{w} = (w_1, w_2, \dots, w_n) \in (\mathbb{R}_-)^N$, and $\mathbf{t} = (t_1, t_2, \dots, t_n) \in (\mathbb{R}_-)^N$. Denote by $M_i = \{j \in M \mid V_{ij} \geq t_i\}$ and $N_j = \{i \in N \mid j \in M_i\}$. If the following linear program*

$$\mathcal{P} : \begin{cases} \sum_{j \in M_i} V_{ij} x_{ij} \geq w_i, & \forall i \in N \\ \sum_{i \in N_j} x_{ij} = 1, & \forall j \in M \\ x_{ij} \geq 0, & \forall i \in N, j \in M_i \end{cases}$$

has a feasible solution, then any extreme point $\tilde{\mathbf{x}}$ of this polytope (defining the solution space) can be rounded to a feasible solution $\bar{\mathbf{x}}$ of the integer program

$$\mathcal{P}' : \begin{cases} \sum_{j \in M_i} V_{ij} x_{ij} \geq w_i + t_i, & \forall i \in N \\ \sum_{i \in N_j} x_{ij} = 1, & \forall j \in M \\ x_{ij} \in \{0, 1\}, & \forall i \in N, j \in M_i. \end{cases}$$

Proof: Let $\tilde{\mathbf{x}}$ be an extreme point of the polytope defined by \mathcal{P} . Then $\tilde{\mathbf{x}}$ contains at most $m + n$ nonzero variables due to the number of constraints in \mathcal{P} . We construct a bipartite graph via $\tilde{\mathbf{x}}$, $G = (N, M, E)$, where N is the set of agents, M is the set of chores and $E = \{(i, j) \mid \tilde{x}_{ij} > 0, i \in N, j \in M\}$. Using the same argument with the proof of Theorem 1 in [Lenstra *et al.*, 1990], we know that G is a pseudoforest, i.e., each connected component of G is a tree or a tree plus one additional edge.

Next, we round $\tilde{\mathbf{x}}$ to $\bar{\mathbf{x}}$. In any connected component C of G , whenever there is a chore j whose degree is 1, then it must be that for some i , $\tilde{x}_{ij} = 1$. Then set $\bar{x}_{ij} = 1$ and delete this chore from C . Denote by C' the remained graph. Note that in C' , any remaining chore has a degree at least 2. Therefore C' must contain a matching which covers all chores, due to the fact that C' is a tree or a tree plus one additional edge. According to this matching, if (i, j) is matched, set $\bar{x}_{ij} = 1$; otherwise, set \bar{x}_{ij} to be 0.

Next, it suffices to verify that $\bar{\mathbf{x}}$ is a feasible solution of \mathcal{P}' . For each chore j , $\bar{\mathbf{x}}$ assigns it to exactly one agent. Thus, $\sum_{i \in N_j} \bar{x}_{ij} = 1$ for any $j \in M$. For each agent $i \in N$, there

is at most 1 chore j such that \tilde{x}_{ij} is increased to 1. Since $0 \geq V_{ij} \geq t_i$,

$$\sum_{j \in M_i} V_{ij} \tilde{x}_{ij} \geq \sum_{j \in M_i} V_{ij} \tilde{x}_{ij} + t_i \geq w_i + t_i,$$

which completes the proof of Lemma 5. \square

Note that, solving the optimal α for integer program \mathcal{IP} is equivalent to finding the minimum value of α such that \mathcal{IP} has a feasible integer solution \mathbf{x} . However, \mathcal{P} is not the relaxation of \mathcal{IP} since in \mathcal{P} , there is not a variable x_{ij} for which $V_{ij} < t_i$. Equivalently, we can add these variables to \mathcal{P} and set them to zero. In the following, we discuss the relationship between the solutions of \mathcal{IP} , \mathcal{P} and \mathcal{P}' .

Recall that α^* is the OWMMS ratio, which is also the optimal value of \mathcal{IP} . Let

$$c^* = \min\{c \in \mathbb{R}_0^+ \mid \mathcal{P} \text{ has a feasible solution with } t_i = w_i = c\text{WMMS}_i(\mathcal{I}) \text{ for all } i \in N\}.$$

Note that c^* always exists as $c = n$ is always feasible by Algorithm Naive. Moreover, although \mathcal{P} is not the relaxation of \mathcal{IP} , c^* is still a lower bound of α^* .

Lemma 6 $\alpha^* \geq c^*$.

Lemma 6 shows that to approximate α^* , it suffices to find a feasible solution of \mathcal{IP} whose value is a good approximation to c^* . Next, we show that a feasible solution of \mathcal{P}' is naturally a feasible solution of \mathcal{IP} .

Lemma 7 Let $c \in \mathbb{R}_0^+$. If \mathbf{x} is a feasible solution of \mathcal{P}' with $t_i = w_i = c\text{WMMS}_i(\mathcal{I})$ for all $i \in N$, then $(\mathbf{x}, 2c)$ is a feasible solution of \mathcal{IP} .

Both Lemmas 6 and 7 are proved in the full version. Before we show our main algorithm, let us discuss the following intuitive procedure. First, compute c^* and its corresponding fractional allocation $\tilde{\mathbf{x}}$. Then use Lemma 5 to round $\tilde{\mathbf{x}}$ to an integer solution $\bar{\mathbf{x}}$. By Lemma 7, $\bar{\mathbf{x}}$ is also a feasible solution of \mathcal{IP} . Let $X = \langle X_i \rangle_{i \in N}$ be the final allocation, where $X_i = \{j \in M \mid \bar{x}_{ij} = 1\}$ for every $i \in N$. Thus,

$$V_i(X_i) \geq 2c^*\text{WMMS}_i(\mathcal{I}) \geq 2\alpha^*\text{WMMS}_i(\mathcal{I}),$$

where the first inequality is by Lemma 5 and the second inequality is by Lemma 6. That is, X is a 2-approximation to the optimal WMMS allocation.

However, there are two computational issues with respect to the procedure above: (1) The computation of $\text{WMMS}_i(\mathcal{I})$ may need exponential time¹; (2) Even if we know all the $\text{WMMS}_i(\mathcal{I})$'s, there is a problem of computing c^* .

To resolve (1), we use Algorithm EgalGreedy to compute an approximate value WMMS'_i for each $\text{WMMS}_i(\mathcal{I})$, where $\text{WMMS}'_i \geq 2\text{WMMS}_i(\mathcal{I})$. Then we replace all $\text{WMMS}_i(\mathcal{I})$ by WMMS'_i in above procedure.

To resolve (2), we use binary search to find a near optimal value of c^* . Initially, we first use Algorithm Naive to get an

¹ The computation of WMMS_i is NP-hard, even when $n = 2$ and $s_1 = s_2 = \frac{1}{2}$, via a reduction from the Integer Partition Problem.

upper bound n of α and 1 is a trivial lower bound. Let $\delta > 0$ be the desired precision. Denote by u and l the current upper and lower bounds, respectively. Set $c = \frac{u+l}{2}$ and $w_i = t_i = c\text{WMMS}'_i$, and check if \mathcal{P} has a feasible solution. If \mathcal{P} has a feasible solution, reset u to be $\frac{u+l}{2}$; Otherwise reset $l = \frac{u+l}{2}$. Repeat this process until $u - l \leq \delta$.

We formally describe the algorithm described above as Algorithm 2, denoted by LinPro.

Algorithm 2 LinPro - An Algorithm for General Valuations

Require: Chore allocation instance $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ and $\epsilon > 0$.
Ensure: Allocation $X = \langle X_1, \dots, X_n \rangle$

- 1: Initially, $X_i = \emptyset$ for all $i \in N$.
- 2: For each $i \in N$, run Algorithm EgalGreedy on instance \mathcal{I} and obtain allocation $X^i = \langle X^i_1, X^i_2, \dots, X^i_n \rangle$.
- 3: Set $\text{WMMS}'_i = V_i(X^i)$ for all $i \in N$.
- 4: Let $u = n$ and $l = 1$. % n is the upper bound by Naive.
- 5: **while** $u - l > \frac{\epsilon}{4}$ **do**
- 6: Set $c = \frac{u+l}{2}$.
- 7: Check if \mathcal{P} has a feasible solution by setting $w_i = t_i = c\text{WMMS}'_i$ for all $i \in N$.
- 8: **if** \mathcal{P} has a feasible solution **then**
- 9: Reset $u = c$.
- 10: **else**
- 11: Reset $l = c$.
- 12: Set $w_i = t_i = u\text{WMMS}'_i$ for all $i \in N$ and compute an extreme point $\tilde{\mathbf{x}}$ of \mathcal{P} .
- 13: Use Lemma 5 to round $\tilde{\mathbf{x}}$ to $\bar{\mathbf{x}}$.
- 14: Set $X_i = \{j \in M \mid \bar{x}_{ij} = 1\}$ for all $i \in N$.
- 15: **return** Allocation X .

Theorem 1 Given any chore allocation instance $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ with α^* being its OWMMS ratio. For any $\epsilon > 0$, Algorithm LinPro runs in polynomial time (for any number of agents) and returns an allocation $\langle X_i \rangle_{i \in N}$ such that for any agent i , $V_i(X_i) \geq (4 + \epsilon)\text{OWMMS}_i(\mathcal{I})$.

Proof: By Lemma 7, X is a feasible solution of \mathcal{IP} . At Step 12, as $l \leq c^* \leq u$ and $u - l \leq \frac{\epsilon}{4}$, we have $u \leq c^* + \frac{\epsilon}{4}$. Thus,

$$\begin{aligned} V_i(X_i) &\geq 2u\text{WMMS}'_i \geq 2(c^* + \frac{\epsilon}{4})\text{WMMS}'_i \\ &\geq 4(c^* + \frac{\epsilon}{4})\text{WMMS}_i(\mathcal{I}) \geq 4(\alpha^* + \frac{\epsilon}{4})\text{WMMS}_i(\mathcal{I}) \\ &\geq (4 + \epsilon)\text{OWMMS}_i(\mathcal{I}), \end{aligned}$$

where the first inequality is by Lemma 5, the second inequality is by Lemma 4, and the last inequality is by Lemma 6.

As LinPro requires us to run EgalGreedy and solve at most $O(\log(\frac{n}{\epsilon}))$ numbers of (polynomial-sized) linear program, LinPro runs in polynomial time. \square

Note that the role of EgalGreedy in LinPro can be replaced by other (polynomial-time) approximation algorithms (such the PTAS in [Hochbaum and Shmoys, 1988]) and the approximation ratio (Lemma 4) is improved accordingly.

6 Restricted Cases

In this section, we consider two important restricted cases: (1) two agents and (2) all agents have binary valuations (in which case agents have value 0 or -1 for each item).

6.1 WMMS for Two Agents

Given any instance $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ with $N = \{1, 2\}$, we prove that it is always possible to guarantee each agent i 's value to be at least $\frac{3}{2}\text{WMMS}_i(\mathcal{I})$. Thus, by Lemma 3, the OWMMS ratio α^* for the 2-agent case is within $[\frac{4}{3}, \frac{3}{2}]$.

Divide-and-choose algorithms are widely studied in the literature, especially for the case of two agents. Roughly speaking, the algorithm starts by letting one of the agents divide the whole items (either goods or chores) into two bundles, and the other agent chooses one from the two bundles. Such an algorithm gives an exact MMS allocation for symmetric agents and a 2-WMMS allocation for asymmetric agents (the agent with smaller entitlement divides and the other agent chooses) when the items are goods. However, it is not hard to see that generic divide-and-choose algorithms could be arbitrarily bad when the items are chores.

In the following, we show that with some modification, a divide-and-choose style algorithm, DivCho (defined in Algorithm 3), works well and guarantees each agent i 's value to be at least $\frac{3}{2}\text{WMMS}_i$. Without loss of generality, assume $s_1 \leq s_2$.

Algorithm 3 DivCho -An Algorithm for the 2-Agent Case

Require: Chore allocation instance $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ with $N = \{1, 2\}$.

Ensure: Allocation $X = \langle X_1, X_2 \rangle$.

- 1: Initially, set $X_i = \emptyset$ for both $i \in N$.
 - 2: If $s_1 \leq \frac{1}{3}$ and $s_2 \geq \frac{2}{3}$, set $X_1 = \emptyset$ and $X_2 = M$. Go to Step 5.
 - 3: Let agent 2 partition M into A_1 and A_2 according to a P-2 partition with respect to $\text{WMMS}_2(\mathcal{I})$.
 - 4: Let agent 1 select his favorite bundle from A_1 and A_2 . Denote by X_1 the one chosen by agent 1 and by X_2 the one left for Agent 2.
 - 5: $X = \langle X_1, X_2 \rangle$.
 - 6: **return** Allocation X .
-

Theorem 2 Let $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ with $N = \{1, 2\}$, and $X = \langle X_1, X_2 \rangle$ be the output of Algorithm DivCho on \mathcal{I} . Then, for any agent $i \in N$, $V_i(X_i) \geq \frac{3}{2}\text{WMMS}_i(\mathcal{I}) \geq \frac{3}{2}\text{OWMMS}_i(\mathcal{I})$.

Proof: If $s_1 \leq \frac{1}{3}$ and $s_2 \geq \frac{2}{3}$, Algorithm DivCho allocates all chores to agent 2. Thus $V_1(X_1) = 0$ and $V_2(X_2) = -1$, where agent 1 is trivially satisfied as $V_1(X_1) \geq \frac{3}{2}\text{WMMS}_1(\mathcal{I})$. By Lemma 1, we have $\text{WMMS}_2(\mathcal{I}) \leq -s_2 \leq -\frac{2}{3}$. Accordingly, $V_2(X_2) \geq \frac{3}{2}\text{WMMS}_2(\mathcal{I})$.

As we assume that $s_2 \geq s_1$, our last case is to consider $\frac{1}{2} \leq s_2 \leq \frac{2}{3}$. By the definition of $\text{WMMS}_2(\mathcal{I}) = s_2 \cdot \min\{\frac{V_2(A_1)}{s_1}, \frac{V_2(A_2)}{s_2}\}$, $V_2(A_1) \geq V_2(A_2) \geq \text{WMMS}_2(\mathcal{I})$. As a result, no matter which allocation agent 2 eventually receives after the divide-and-choose procedure, the value of the allocation will always be at least as much as WMMS_2 . For agent 1, since $V(X_1) + V(X_2) = -1$ by assumption and he selects his favorite allocation X_1 , $V_1(X_1) \geq -\frac{1}{2}$. By Lemma 1 and the fact that $s_1 > \frac{1}{3}$, $\text{WMMS}_1(\mathcal{I}) < -\frac{1}{3}$. Therefore, $V_1(X_1) > \frac{3}{2}\text{WMMS}_1(\mathcal{I})$. \square

6.2 Binary Valuations

In this section, we study the case with any number of agents, but every agent's valuation is binary: $V_{ij} \in \{0, -1\}$ for all $i \in N$ and $j \in M$. Note that, throughout this section, we do not impose normalization for ease of exposition. As will be clear later, for this case, we show that it is always possible to guarantee each agent i 's value to be at least WMMS_i , i.e., the optimal WMMS ratio for binary valuation case is exactly 1.

We first prove the following lemma.

Lemma 8 Let $\mathcal{I} = (N, M, \mathbf{s}, \mathbf{V})$ be a chore allocation instance where all agents have an identical valuation V . If V is uniform, (i.e. $V(S) = -|S|$ for any $S \subseteq M$), an exact WMMS allocation can be computed in polynomial time.

Proof: It suffices to show if V is uniform, Algorithm EgalGreedy returns an exact WMMS allocation. Suppose $X = \langle X_i \rangle_{i \in N}$ is the output of EgalGreedy with respect to V . Recall $W(X) = \min_{k \in N} \frac{V(X_k)}{s_k} = \min_{k \in N} \frac{-|X_k|}{s_k}$.

Let $S = \{k \in N \mid \frac{-|X_k|}{s_k} = W(X)\}$ be the set of indices where the minimum is obtained. In the following we show $W = W(X)$. Note that $S \neq \emptyset$. If X is not an optimal partition, then there is a partition X^* such that $W(X^*) > W(X)$.

Thus, every $k \in S$, X_k has to contain a smaller number of chores compared with X_k^* . Accordingly, for some $t \in N \setminus S$, X_t has to contain more chores than X_t^* , i.e., $V(X_t^*) \leq -|X_t| - 1$. If $\frac{-|X_t| - 1}{s_t} \leq W(X)$, $W(X^*)$ cannot be larger than $W(X)$. Thus, $\frac{-|X_t| - 1}{s_t} > W(X)$. But this is a contradiction with the fact that EgalGreedy always allocates greedily, (i.e., Step 4 of EgalGreedy), since the last chore cannot be allocated to X_k for $k \in S$ instead of X_t . That is $W = W(X)$.

Therefore, $\frac{V(X_i)}{s_i} \geq W(X) = W$ and $V(X_i) \geq s_i W = \text{WMMS}_i$ for any $i \in N$, which competes the proof. \square

Thus, by allocating all chores for which some agent has zero value to one such agent, we are left with only the chores for which all agents have value -1. As the modified instance is uniform, by Lemma 8, we have the following theorem.

Theorem 3 For any binary valuation case, a WMMS allocation exists and can be found efficiently.

7 Conclusions

We initiated the study on chore allocation with asymmetric agents. We show that many widely studied greedy algorithms in the literature performs badly and even for the 2-agent case an exact WMMS allocation may not exist. We then presented a constant approximation polynomial time algorithm for OWMMS allocations, and several algorithmic results for the case of identical utilities, binary utilities, and for 2 agents. Finding a stronger lower bound for WMMS allocations for any number of agents remains an open problem.

Acknowledgements

This work is partially supported by NSF CAREER Award No. 1553385. Haris Aziz is supported by a Scientia Fellowship.

References

- [Alon, 1987] N. Alon. Splitting necklaces. *Advances in Mathematics*, 63(3):247–253, 1987.
- [Amanatidis *et al.*, 2015] G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. Approximation algorithms for computing maximin share allocations. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, pages 39–51, 2015.
- [Aziz and Mackenzie, 2016] H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th IEEE Symposium on Foundations of Computer Science*, pages 416–427, 2016.
- [Aziz *et al.*, 2017] H. Aziz, G. Raucher, G. Schryen, and T. Walsh. Algorithms for max-min share fair allocation of indivisible chores. In *Proceedings of 31st AAAI Conference on Artificial Intelligence*, pages 335–341, 2017.
- [Aziz *et al.*, 2018] H. Aziz, I. Caragiannis, and A. Igarashi. Fair allocation of combinations of indivisible goods and chores. *CoRR*, abs/1807.10684, 2018.
- [Aziz, 2016] H. Aziz. Computational social choice: Some current and new directions. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 4054–4057, 2016.
- [Babaioff *et al.*, 2017] M. Babaioff, N. Nisan, and I. Talgam-Cohen. Competitive equilibria with indivisible goods and generic budgets. *CoRR*, abs/1703.08150, 2017.
- [Barman and Murthy, 2017] S. Barman and S. Kumar Krishna Murthy. Approximation algorithms for maximin fair division. In *Proceedings of the 18th ACM conference on Economics and Computation*, pages 647–664, 2017.
- [Brams and Taylor, 1995] S. Brams and A. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995.
- [Brams and Taylor, 1996] S. Brams and A. Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, 1996.
- [Budish, 2011] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Cseh and Fleiner, 2018] A. Cseh and T. Fleiner. The complexity of cake cutting with unequal shares. In *Proceedings of the 11th International Symposium on Algorithmic Game Theory*, pages 19–30, 2018.
- [Dubins and Spanier, 1961] L. Dubins and E. Spanier. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1):1–17, 1961.
- [Farhadi *et al.*, 2017] A. Farhadi, M. Hajiaghayi, M. Ghodsi, S. Lahaie, D. M. Pennock, M. Seddighin, S. Seddighin, and H. Yami. Fair allocation of indivisible goods to asymmetric agents. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems*, pages 1535–1537, 2017.
- [Farhadi *et al.*, 2019] A. Farhadi, M. Hajiaghayi, M. Ghodsi, S. Lahaie, D. M. Pennock, M. Seddighin, S. Seddighin, and H. Yami. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research (JAIR)*, 2019.
- [Friesen, 1987] D. Friesen. Tighter bounds for lpt scheduling on uniform processors. *SIAM Journal on Computing*, 16(3):554–560, 1987.
- [Garey and Johnson, 1979] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [Ghodsi *et al.*, 2018] M. Ghodsi, M. Hajiaghayi, M. Seddighin, S. Seddighin, and H. Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 19th ACM conference on Economics and Computation*. ACM Press, 2018.
- [Gonzalez *et al.*, 1977] T. Gonzalez, O. Ibarra, and S. Sahni. Bounds for lpt schedules on uniform processors. *SIAM Journal on Computing*, 6(1):155–166, 1977.
- [Hochbaum and Shmoys, 1988] D. Hochbaum and D. Shmoys. A polynomial approximation scheme for scheduling on uniform processors: Using the dual approximation approach. *SIAM journal on computing*, 17(3):539–551, 1988.
- [Kurokawa *et al.*, 2018] D. Kurokawa, A. Procaccia, and J. Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM (JACM)*, 65(2):8, 2018.
- [Lenstra *et al.*, 1990] J. Lenstra, D. Shmoys, and E. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical programming*, 46(1-3):259–271, 1990.
- [Pinedo, 2016] M. Pinedo. *Scheduling: theory, algorithms, and systems*. Springer, 2016.
- [Procaccia and Wang, 2014] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 15th ACM conference on Economics and Computation*, pages 675–692. ACM Press, 2014.
- [Robertson and Webb, 1998] J. Robertson and W. Webb. Cake-cutting algorithms: Be fair if you can. 1998.
- [Stromquist, 1980] W. Stromquist. How to cut a cake fairly. *The American Mathematical Monthly*, 87(8):640–644, 1980.