

On the PTAS for Maximin Shares in an Indivisible Mixed Manna

Rucha Kulkarni¹, Ruta Mehta¹, Setareh Taki¹

¹ University of Illinois at Urbana-Champaign
ruchark2@illinois.edu, rutameht@illinois.edu, staki2@illinois.edu

Abstract

We study fair allocation of indivisible items, both goods and chores, under the popular fairness notion of maximin share (MMS). The problem is well-studied when there are only goods (or chores), where a PTAS to compute the MMS values of agents is well-known (Woeginger 1997; Jansen, Klein, and Verschae 2016).

In contrast, for the mixed manna, a recent result of (Kulkarni, Mehta, and Taki 2020) showed that finding even an approximate MMS value of an agent up to any approximation factor in $(0, 1]$ is NP-hard for general instances. In this paper, we complement the hardness result by obtaining a PTAS to compute the MMS value, when its absolute value is at least $1/\rho$ times either the total value of all the goods or total cost of all the chores, for some constant $\rho \geq 1$.

1 Introduction

Finding fair and efficient allocations is a fundamental problem in algorithmic game theory. The problem has been extensively studied for divisible resources, phrased as the cake cutting problem; see (Robertson and Webb 1998) for a summary. Here, a division of a *cake* that gave one piece to each of n agents was termed fair if it ensured properties like (a) envy-freeness, meaning every agent values her own piece more than those allocated to other agents, and (b) proportionality, meaning every agent values her piece at least $1/n$ fraction of her total value for the cake.

When there are two agents, the simple cut and choose protocol is known to work since the biblical era, where one agent cuts the cake into two pieces and the other agent gets to choose first. Recent years have seen a surge of works on the fair division of indivisible items, like school/course seats, assets and liabilities, and computing resources on networks, due to their wide applications (Steinhaus 1948; Brams and Taylor 1996; Vossen 2002; Moulin 2004; Etkin, Parekh, and Tse 2007; Budish 2011; Ghodsi et al. 2018). A simple example of allocating a single indivisible item among two agents shows

that both envy-free and proportional allocations may not exist. Therefore, Budish (2011) defined the notion of maximin share (MMS) based on the following extension of the cut and choose protocol. If there are n agents, an agent partitions the items into n bundles assuming she will get to choose last. As she may end up with the least valued bundle, naturally, she will partition the items in such a way that the value of the least valued bundle is maximized. This is called her MMS value. An allocation where every agent receives a bundle of at least her MMS value is called an MMS allocation.

This problem is well studied for the *good manna* where all items are valued non-negatively by every agent, and for the *bad (chore) manna* where items are valued non-positively by everyone (See Section 1.1 for related work). We consider a *mixed manna* setting, where every item can be positively valued by some agents, and negatively valued by some. A natural starting question in the quest to find MMS allocations is,

Q: Given a mixed manna, what is the MMS value of every agent?

This question is NP-hard, even for the good manna. Note that finding the MMS value of an agent is equivalent to finding an MMS allocation when all agents have valuations identical to this agent. Hence, the problem of finding MMS allocations is also NP-hard, even with identical agents and a good manna. However, (Woeginger 1997) gave a PTAS for this setting. This PTAS was later used in several works to find approximate MMS allocations with non-identical agents (Procaccia and Wang 2014; Kurokawa, Procaccia, and Wang 2016; Amanatidis et al. 2017; Ghodsi et al. 2018; Garg and Taki 2020). The best known approximation result for the MMS problem with nonidentical agents in a chore manna (Huang and Lu 2019) also uses a PTAS for finding MMS values (Jansen, Klein, and Verschae 2016) as a subroutine. The next question then is,

Q: Is there a PTAS to find the MMS values of agents in a mixed manna setting?

Surprisingly, (Kulkarni, Mehta, and Taki 2020) showed that even in a highly restricted setting of two identical agents where the mixed manna has only two

chores, it is NP-hard to find *approximate* MMS values within *any* constant factor. Their reduction indicates that perhaps the bottleneck issue that makes the problem hard, is that the absolute MMS value can be arbitrarily small. Intuitively speaking, as the MMS value approaches arbitrarily close to zero, the problem of finding approximate MMS values approaches that of finding exact MMS values. In the limit where $\text{MMS} = 0$, every approximate MMS value is the exact MMS value. We then ask,

Q: If we had a guarantee that the absolute MMS value is greater than some threshold value, say Δ , then does the problem become tractable?

In this paper, we resolve this question positively for a specific value of Δ . Note that since the MMS problem is scale-free, setting Δ to a fixed constant will not make the problem any easier. Therefore, Δ will have to be instance dependent. To be specific, let v^+ denote the sum of values of an agent for all the items she values positively, and v^- the sum of absolute values of all her negatively valued items (chores).

Theorem 1.1 (Informal). *In the case of identical agents, there is an algorithm that: (a) when $|\text{MMS}| \geq \min\{v^+, v^-\}/\rho$ for some constant $\rho \geq 1$, finds an allocation that gives every agent a bundle of value at least $(1 - \epsilon)$ -MMS for any constant $\epsilon > 0$, and (b) when $|\text{MMS}| < \min\{v^+, v^-\}/\rho$, reports this by returning the trivial allocation where all items are given to one agent. The algorithm runs in time $O(mnL)$, where m, n are the number of items and agents, and L is the bit-length of the input.*

We note that our assumption is weaker than having $\min\{v^+, v^-\}$ being a constant. Also, the extensively studied good manna and chore manna are special cases of this setting, hence any algorithmic results here translate to these settings as well.

One of the key tools used by our algorithm is a carefully designed Integer Program (IP) that can be solved in polynomial-time. IPs have been used to solve related problems in several works. (Woeginger 1997) gave a PTAS using this idea for the machine covering problem, which is equivalent to the MMS problem in a good manna. The MMS problem in a chore manna is equivalent to machine scheduling which has a PTAS using IP for bin packing (the dual problem). Several algorithms for the bin packing problem solve a relaxation of an IP as their main idea (De La Vega and Lueker 1981; Johnson 1982; Karmarkar and Karp 1982). Our approach builds on these, but requires several new ideas to handle both goods and chores simultaneously. Next we briefly describe some of these.

Non-constant variables. The variables of the IP will correspond to subsets of items. We show that we only need to consider subsets with total value at most a particular bound. With a good (chore) manna, this restricts to subsets where the number of items with value at least some fraction of the bound is a constant. While

for a mixed manna, subsets of even $O(m)$ size may have a small value due to positive and negative values may cancel each other. Hence, the number of variables of the IP is not a constant for a mixed manna.

We circumvent this issue by reducing the problem to a problem with only goods, where a restricted set of allocations are allowed, called *valid allocations*.

Allow only valid allocations. The next task is to define constraints in the IP that ensure a valid allocation. Towards this, we define a cost function that characterizes valid allocations with a single constraint.

Sign of MMS. Our approach works for both cases $\text{MMS} \geq 0$ (Section 3), and $\text{MMS} < 0$ (Appendix B). These are inherently different problems. The $\text{MMS} \geq 0$ problem maximizes the smallest bundle's value, while the negative MMS case minimizes the absolute value of the largest bundle. We show that our IP for the former case can be modified to work for the later case of $\text{MMS} < 0$.

1.1 Related Work

The MMS problem has been extensively studied for the good manna (Kurokawa, Procaccia, and Wang 2016; Ghodsi et al. 2018; Garg, McGlaughlin, and Taki 2018; Kurokawa, Procaccia, and Wang 2018; Barman and Krishna Murthy 2017; Farhadi et al. 2019; Amanatidis et al. 2017; Garg and Taki 2020) and chore manna (Barman and Krishna Murthy 2017; Huang and Lu 2019) settings. With a good manna, there are several algorithms to find allocations that give every agent a bundle worth a constant fraction of their MMS value; the best factor known so far is $(3/4 + 1/(12n))$ -MMS, by (Garg and Taki 2020). For the chore manna, (Huang and Lu 2019) give a PTAS to find the MMS values of agents, and an $11/9$ approximate MMS allocation. The study of the mixed manna setting started recently. (Kulkarni, Mehta, and Taki 2020) gave a PTAS for the special case of the problem with a constant number of agents, when the total value of goods is some factor away from the total absolute value of chores.

2 Preliminaries and Notation

In this section, we formally define mixed instances and other relevant notions of maximin share. We use $[k]$ to denote the set $\{1, 2, \dots, k\}$, and $(\mathcal{S}_j)_{j \in [k]}$ to denote the (multi-)set $\{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_k\}$.

Definition 2.1. *An MMS instance is a tuple $\langle \mathcal{N}, \mathcal{M}, v \rangle$, where \mathcal{N} is a set of n agents, \mathcal{M} is a set of m indivisible items, and $v : 2^{\mathcal{M}} \rightarrow \mathbb{R}$ is the identical additive valuation function of all agents, represented by $v(S) = \sum_{j \in S} v_j$ for $S \subseteq \mathcal{M}$.*

A partition of all items among all agents is termed an *allocation*, denoted by $\mathcal{A} = \{A_1, A_2, \dots, A_n\}$. Thus, $A_i \cap A_{i'} = \emptyset$ for all distinct i, i' in \mathcal{N} , and $\cup_i A_i = \mathcal{M}$.

Definition 2.2 (MMS value). *Given an MMS instance,*

let $\Pi_n(\mathcal{M})$ be the set of all possible allocations of \mathcal{M} into n sets. The maximin share (MMS) value of an agent, denoted by $\text{MMS}^n(\mathcal{M})$, is defined as,

$$\text{MMS}^n(\mathcal{M}) = \max_{\mathcal{A} \in \Pi_n(\mathcal{M})} \min_{A_k \in \mathcal{A}} v(A_k) .$$

We refer to $\text{MMS}^n(\mathcal{M})$ by MMS when the qualifiers n and \mathcal{M} are clear. Note that MMS can be negative too.

An allocation which gives every agent a set of items worth at least MMS is called an MMS allocation. Note that all agents have the same valuation function v for \mathcal{M} , hence the MMS values are same for all agents. Also, the allocation determining the MMS value for any agent is an MMS allocation. Hence when the agents are identical MMS allocations always exist. However, finding an MMS allocation is known to be NP-Hard (Bouveret and Lemaître 2016). Thus, we search for a PTAS to find almost optimal allocations, termed $(1 - \epsilon)$ -MMS allocations, defined as follows.

Definition 2.3 ($(1 - \epsilon)$ -MMS allocation). *A is called a $(1 - \epsilon)$ -MMS allocation, if for a given $\epsilon > 0$, for each agent $i \in \mathcal{N}$ we have $v(A_i) \geq (1 - \epsilon)\text{MMS}$ if $\text{MMS} \geq 0$, and $v(A_i) \geq (1/(1 - \epsilon))\text{MMS}$, if $\text{MMS} < 0$. Equivalently,*

$$v(A_i) \geq \min\{(1 - \epsilon)\text{MMS}, (1/(1 - \epsilon))\text{MMS}\}.$$

Definition 2.4 (MMS problem). *Given an MMS instance $\langle \mathcal{N}, \mathcal{M}, v \rangle$, the MMS problem is to find a $(1 - \epsilon)$ -MMS allocation of \mathcal{M} among \mathcal{N} .*

Items of a mixed manna can be divided into two sets. *Goods* are the items valued positively according to v . The set of goods is denoted by $\mathcal{M}^+ = \{j \in \mathcal{M} \mid v_j \geq 0\}$. *Chores* are items valued negatively, and the set of chores is termed $\mathcal{M}^- = \{j \in \mathcal{M} \mid v_j \leq 0\}$. We denote by v^+ the sum of values of all goods in the manna. That is, $v^+ = \sum_{j \in \mathcal{M}^+} v_j$. Similarly, we denote by v^- the sum of absolute values of all chores, i.e., $v^- = \sum_{j \in \mathcal{M}^-} |v_j|$.

To circumvent the hardness result of (Kulkarni, Mehta, and Taki 2020) for the MMS problem for any $\epsilon \in [0, 1)$, we make the assumption that $|\text{MMS}| \geq \min\{v^+, v^-\}/\rho$, for some constant $\rho \geq 1$. As we cannot decide before computing the MMS value if a given instance satisfies this property, we pose the following problem, termed the Bounded MMS problem, denoted by B-MMS.

Definition 2.5 (B-MMS problem). *Given an MMS instance $\langle \mathcal{N}, \mathcal{M}, v \rangle$ and an $\epsilon > 0$, return a $(1 - \epsilon)$ -MMS allocation if $\text{MMS} \geq \min\{v^+, v^-\}/\rho$ for some constant $\rho \geq 1$, else report $\text{MMS} < \min\{v^+, v^-\}/\rho$ by returning the trivial allocation where one agent gets all items \mathcal{M} .*

In this paper, we give a polynomial time algorithm that solves the B-MMS problem. In other words, we provide a PTAS to find the MMS values of agents in a mixed manna, when the absolute MMS values are higher than $\min\{v^+, v^-\}/\rho$, for some constant $\rho \geq 1$.

The following lemma by (Kulkarni, Mehta, and Taki 2020) shows an easy way to decide the sign of MMS, al-

lowing us to design separate approaches for the negative and non-negative MMS cases.

Lemma 2.1. $v(\mathcal{M}) \geq 0$ iff $\text{MMS} \geq 0$.

For solving the B-MMS problem, we first find the sign of MMS using Lemma 2.1, then apply the appropriate algorithm for that case.

3 Algorithm for B-MMS when $\text{MMS} \geq 0$

In this section, we describe a PTAS for the B-MMS problem for the $\text{MMS} \geq 0$ case. All missing proofs of this section are provided in Appendix A. The main parts of the PTAS are explained and solved in separate subsections.

3.1 Reducing B-MMS to GC-MMS

We first reduce the given B-MMS problem to a new problem with only goods called the Goods manna Constrained MMS problem, denoted by GC-MMS. At a high level, this is similar to the MMS problem, but it computes optimal allocations over a restricted set of partitions, called *valid allocations* (described shortly).

The intuition behind defining GC-MMS problem is as follows: Suppose we replace every chore by $n - 1$ goods, each of value equal to the absolute value of the chore. Lets call these good-copies of the chores. Every time we want to assign a chore $j \in \mathcal{M}$ to some agent, we instead assign one of the $n - 1$ good-copies of j to the remaining $n - 1$ agents (one copy for each $n - 1$ agents). This adds exactly $|v_j|$ value to every bundle and therefore keeps their relative order the same. Once we do this for every chore, the value added to each bundle is exactly v^- , and is the same for every partition in $\Pi^n(\mathcal{M})$. Therefore, if we restrict the allocations in the new setting to allow an agent to get at most one good-copy of any chore, then MMS allocations in the two settings are equivalent.

Following this intuition, we define a GC-MMS instance and valid allocations as follows.

Definition 3.1 (GC-MMS instance). *A tuple $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle$, where \mathcal{N} is a set of agents, \mathcal{G} is a set of goods, $(\mathcal{S}_j)_{j \in [m^-]}$ are m^- sets of goods, each containing $(n - 1)$ identical copies of a good, and $u : \mathcal{M} \cup (\mathcal{S}_j)_{j \in [m^-]} \rightarrow \mathbb{R}_+$ is the identical valuation function of the agents in \mathcal{N} for all items $\mathcal{G} \cup (\mathcal{S}_j)_{j \in [m^-]}$.*

Definition 3.2 (Valid allocation). *Given a GC-MMS instance $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle$, an allocation \mathcal{A} is valid if no agent receives more than 1 item from any set $\mathcal{S}_j \in (\mathcal{S}_j)_{j \in [m^-]}$, i.e., for all $i \in \mathcal{N}, j \in [m^-]$, $|\mathcal{A}_i \cap \mathcal{S}_j| \leq 1$.*

The GC-MMS problem asks to find a *valid allocation* that maximizes the value of the smallest bundle, i.e., an MMS allocation over the valid allocations. We abuse notation to denote both the problem and the value by GC-MMS, and formally define them as follows.

Definition 3.3 (GC-MMS value). *Given a GC-MMS instance $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle$, let \mathcal{F} be the set of all*

valid allocations. The GC-MMS value of the instance, denoted by GC-MMS, is defined as follows.

$$\text{GC-MMS} = \arg\max_{A \in \mathcal{F}} \min_{A \in \mathcal{A}} v(A)$$

Since it is NP-hard to compute GC-MMS (even when $(\mathcal{S}_j)_{j \in [m^-]} = \emptyset$), define the following approximate version of the problem.

Definition 3.4 (GC-MMS problem). *Given a GC-MMS instance $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle$ and $\epsilon > 0$, return a valid allocation \mathcal{A} such that $\min_{A \in \mathcal{A}} v(A) \geq (1 - \epsilon') \text{GC-MMS}$.*

Next we show that the B-MMS problem can be reduced to GC-MMS problem such that a PTAS for the latter gives a PTAS for the former.

Given an instance $\langle \mathcal{N}, \mathcal{M}, v \rangle$ we define the corresponding GC-MMS instance $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle$ as follows: The set of agents is unchanged, $\mathcal{G} = \mathcal{M}^+$, $m^- = |\mathcal{M}^-|$, and for all $j \in \mathcal{M}^-$, define \mathcal{S}_j to be a set of $(n - 1)$ goods represented as $\mathcal{S}_j := \{(j, k) | k \in [n - 1]\} - \mathcal{S}_j$ consists of good-copies of chore j . Finally, define $u(j) = v(j)$ for all $j \in \mathcal{G}$ and $u((j, k)) = -v(j)$ for all $j \in \mathcal{M}^-$ and $k \in [n - 1]$.

Lemma 3.1. *Allocations of B-MMS are in one-to-one correspondence with valid allocations of GC-MMS, such that if allocation B^π of the former maps to allocation C^π of the later then $u(C_i) = v(B_i) + v^-$, $\forall i \in \mathcal{N}$.*

Proof. Given a B-MMS allocation B^π , add good-copies of each chore to agents who did not receive the chore in B^π , and discard all chores. This gives a GC-MMS allocation C^π . The reverse allocation is obtained by similarly discarding all good-copies and assigning the corresponding chore to the agent who did not receive any good-copy.

Every agent $i \in \mathcal{N}$ receives in C_i all the goods assigned to her in B_i . Every chore that was assigned to her in B_i is discarded in C_i . Due to this, her value increases by the absolute value of chores allotted to her in B_i . Further, for every chore not assigned to her, she receives a good-copy of it in C_i . As for all $j \in C$, $u(j) = |v(j^*)|$ for the corresponding j^* in \mathcal{M} , each good-copy increases her value by the absolute value of the corresponding chore. Her total valuation increases by the absolute value of all chores not assigned to her as well. Hence, the difference $u(C_i) - v(B_i)$ is exactly, $\sum_{j \in B_i} v(j) + \sum_{j \notin B_i} v(j) = v^-$. \square

Corollary 3.1. *GC-MMS, relates to the MMS value of the B-MMS problem as,*

$$\text{GC-MMS} = \text{MMS} + v^-. \quad (1)$$

Equation (1) allows to relate the approximation parameters of B-MMS and GC-MMS allocations as follows.

Theorem 3.1. *If $\text{MMS} \geq v^-/\rho$, then a $(1 - \frac{\epsilon}{(1+\rho)})$ GC-MMS allocation gives a $(1 - \epsilon)$ -MMS allocation, and therefore a PTAS for GC-MMS gives a PTAS for the B-MMS problem.*

Proof. Let $\epsilon' = \frac{\epsilon}{(1+\rho)}$. Take the $(1 - \epsilon')$ GC-MMS allocation, say C^π , and consider the corresponding allocation B^π of the B-MMS instance as described in the proof of Lemma 3.1. From Lemma 3.1, the smallest bundle in B^π has value $(1 - \epsilon')\text{GC-MMS} - v^-$.

If $\text{MMS} \geq v^-/\rho$, we have, $(1 - \epsilon')\text{GC-MMS} - v^- \geq (1 - \epsilon')(\text{MMS} + v^-) - v^- \geq (1 - \epsilon')(\text{MMS} + \rho\text{MMS}) - \rho\text{MMS} = (1 - (1 + \rho)\epsilon')\text{MMS} = (1 - \epsilon)\text{MMS}$. Therefore, B^π is a $(1 - \epsilon)$ -MMS allocation

Since ρ and ϵ are constants in the B-MMS problem, ϵ' is also a constant. Therefore, a PTAS for GC-MMS is indeed a PTAS for the B-MMS problem as well. \square

Due to the above theorem, it suffices to obtain a PTAS for the GC-MMS problem.

3.2 Algorithm for GC-MMS

Algorithm 1 for GC-MMS will perform a search for the highest value μ for which we get an allocation that gives every agent at least a μ -valued bundle. For this we perform a search on a multiplicative grid over all possible values of GC-MMS, obtained as follows. First, we have $v^-/\rho \leq \text{MMS} \leq v(\mathcal{M})/n = (v^+ - v^-)/n$. Combined with Equation (1), we get $v^- + v^-/\rho \leq \text{GC-MMS} \leq (v^+ - v^-)/n + v^-$.

In each iteration of the search, it first checks if there is an item with value more than μ . First, there will be no such chore. Because if there was one, say c , we have $c > \mu \geq \text{GC-MMS} = \text{MMS} + v^-$ which implies $\text{MMS} < c - v^- \leq 0$.

If there is a good j with $v(j) \geq \mu$, we have $\mu - v^- \geq \text{GC-MMS} - v^- = \text{MMS}$ (B-MMS instance). Using this we find B^π , a solution of the B-MMS instance as follows: assign good j and all the chores to an agent, and remove the agent and her bundle. The following allocation for the resulting instance is feasible, and has equal or higher MMS value. From any MMS allocation of the B-MMS instance, (a) remove all chores and add them to the part containing the good j , and (b) remove all goods except j from this part and arbitrarily distribute among the remaining parts. The MMS value of the resulting instance is not lower, and therefore it suffices to find it's $(1 - \epsilon)$ -MMS allocation using the PTAS of (Woeginger 1997). Algorithm 1 returns allocation C^π corresponding to this B-MMS allocation.

If every item has value at most μ , the algorithm applies a subroutine **Exists-GC-MMS**, for which we prove in Section 3.3,

Theorem 3.2. **Exists-GC-MMS** $(\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle, \epsilon, \mu)$ returns a tuple $(\mathcal{A}, \text{flag})$ with $\text{flag} = \text{true}$ and $u(A) \geq (1 - \epsilon)\mu$, $\forall A \in \mathcal{A}$, whenever $\mu \leq \text{GC-MMS}$. And it runs in $O(mn)$ time.

If **Exists-GC-MMS** returns a false flag, the Algorithm resets $\mu \leftarrow (1 - \epsilon)\mu$ and starts the next iteration, else returns the allocation obtained and stops. Theorem 3.2 implies the following. When Algorithm 1 stops, say for a

Algorithm 1: Algorithm for GC-MMS

Input : $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle, \epsilon' > 0$ **Output:** $(1 - \epsilon')$ GC-MMS allocation if
GC-MMS $\geq (1 + 1/\rho)v^-$

```
1  $\bar{\epsilon} \leftarrow \epsilon'/2$ ;  $\mu \leftarrow v^+/n + (1 - 1/n)v^-$ 
2 while  $\mu \geq (1 + 1/\rho)v^-$  do
3   if  $\exists j \in \mathcal{G} : u(j) \geq \mu$  then
4      $\mathcal{A} = (A_1, \dots, A_n), A_n \leftarrow \{j\}$ 
5      $(A_1, \dots, A_{n-1}) \leftarrow (1 - \bar{\epsilon})$ -MMS partition
      of  $\langle \mathcal{N} \setminus \{n\}, \mathcal{G}, u \rangle$  // use PTAS of
      (Woeginger 1997)
6      $A_i \leftarrow A_i \cup \{(j, i)\}$  for all  $i \in [n - 1]$  and
       $j \in [m^-]$ 
7     return  $\mathcal{A}$ 
8    $(\mathcal{A}, \text{flag}) \leftarrow$ 
      Exists-GC-MMS( $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle, \bar{\epsilon}, \mu$ )
9   If  $\text{flag}$  then return  $\mathcal{A}$ 
10  else  $\mu \leftarrow (1 - \bar{\epsilon})\mu$ 
11  $\mathcal{A} = (A_1, \dots, A_n)$  where  $A_i = \{(j, i) : \forall j \in [m^-]\}$ 
    for  $i \in [n - 1], A_n = \mathcal{G}$  // agents 1 to  $n - 1$ 
    each get one good-copy of all chores.
12 return  $\mathcal{A}$ 
```

value μ^* , we know GC-MMS $\leq \mu^*/(1 - \bar{\epsilon})$, from the false flag returned in the previous iteration. From this iteration's output, we have a $(1 - \bar{\epsilon})^2$ GC-MMS allocation. Fixing $\bar{\epsilon}$ as $\epsilon'/2$ gives,

Lemma 3.2. *If GC-MMS $\geq (1 + 1/\rho)v^-$, Algorithm 1 returns a $(1 - \epsilon')$ GC-MMS allocation.*

Now we are ready to show Theorem 1.1 for the case when MMS ≥ 0 .

Theorem 3.3. *There is an algorithm to solve the B-MMS problem for the case MMS ≥ 0 , that runs in time $O(mnL)$, where L is the number of bits needed to represent function v .*

Proof. From Theorem 3.1, it suffices to get a PTAS for the corresponding GC-MMS problem. By Lemma 3.2 Algorithm 1 does solve a GC-MMS problem. The while loop of the algorithm runs for $\frac{1}{\bar{\epsilon}} \log \left(\frac{v^+ + (n-1)v^-}{n} - \frac{(1+\rho)v^-}{\rho} \right) \leq \frac{2(1+\rho)}{\bar{\epsilon}} L$ many times. By Theorem 3.2 and (Woeginger 1997), every iteration of the while loop takes at most $O(mn)$ time, and therefore the overall running time is $O(mnL)$. \square

Remark 3.1. *The proof of Theorem 1.1 for the case when MMS < 0 is similar, and in some sense simpler, and discussed in Appendix B.*

The next section shows Theorem 3.2.

3.3 Algorithm for Exists-GC-MMS

At a high level, we first map the set of items in the Exists-GC-MMS instance to multi-sets of numbers cor-

responding to their values (scaled to have $\mu = 9\lceil \frac{1}{\bar{\epsilon}} \rceil^2$ for technical reasons). *Valid partitions* of these numbers are defined analogously like valid allocations of the GC-MMS items. We then classify the values as BIG or SMALL. The key component of the algorithm is an IP to find a valid partition of the BIG values such that (a) every part has value at least $9(\lceil \frac{1}{\bar{\epsilon}} \rceil^2 - \lceil \frac{1}{\bar{\epsilon}} \rceil)$, and (b) there are enough SMALL values to greedily allocate over this partition and have every part valued at least $9\lceil \frac{1}{\bar{\epsilon}} \rceil^2$. We now discuss the details of the algorithm formally.

Exists-GC-MMS has two steps 1) Pre-processing and 2) Main Algorithm.

Pre-processing. (Algorithm 2, line 1) Let $E := \lceil \frac{1}{\bar{\epsilon}} \rceil$. Note that E is a constant integer that only depends on $\bar{\epsilon}$ and not on parameters in the GC-MMS instance. Scale the valuations v by $9E^2/\mu$. Let $\mathcal{V}^g = (g_j)_{(j \in [m^+])}$ and $\mathcal{V}^c = \cup_{j \in [m^-]} C_j$, where $C_j = (c_j^k)_{k \in [n-1]}$ be multi-sets of numbers corresponding to scaled valuations, respectively of \mathcal{M}^+ and $(\mathcal{S}_j)_{j \in [m^-]}$. Let $\mathcal{T} = \mathcal{V}^g \cup \mathcal{V}^c$.

This completes the pre-processing step. The following lemmas characterize partitions of \mathcal{T} that correspond to approximately optimal GC-MMS allocations.

Definition 3.5 (Valid Partition of \mathcal{T}). *We call a partition $P = (P_1, \dots, P_n)$ of values in \mathcal{T} valid if each P_k contains at most one element from each C_j , i.e., $|P_k \cap C_j| \leq 1$ for all $k \in [n]$ and $j \in [m^-]$.*

It is easy to see that each valid partition of \mathcal{T} is equivalent to a valid allocation in its corresponding GC-MMS instance. With the scaling step, this directly implies,

Lemma 3.3. *Given a GC-MMS instance $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle$, if $\mu \leq$ GC-MMS then there is a valid partition of \mathcal{T} where the sum of values in each part is at least $9E^2$.*

As $E = \lceil 1/\bar{\epsilon} \rceil$, we can show that a part of value at least $9E^2 - 9E$ will correspond to a bundle of value at least $(1 - \bar{\epsilon})\mu$. We use this and Lemma 3.3 to show the next lemma.

Lemma 3.4. *A valid partition of \mathcal{T} where the sum of values in each part is at least $9(E^2 - E)$ is equivalent to a valid allocation for its corresponding GC-MMS instance where each bundle has value at least $(1 - \bar{\epsilon})\mu$.*

Main Algorithm. Call a valid partition of \mathcal{T} optimal if the sum of values in each part is at least $9(E^2 - E)$. This step returns an optimal partition if $\mu \leq$ GC-MMS, else correctly reports $\mu >$ GC-MMS by returning $\text{flag} = \text{false}$. Note that Algorithm 1 runs Exists-GC-MMS only if every item has value at most μ . Hence, after scaling by $9E^2$, we can assume $t \leq 9E^2, \forall t \in T$. The key of the algorithm is an IP. We first explain the IP.

Notation. We define SMALL and BIG values in \mathcal{T} . Call a value $t \in \mathcal{T}$ SMALL if $t < 3E$ and BIG if $t \geq 3E$. For each $T \subseteq \mathcal{T}$ let SMALL(T) be the set of all small values in T and BIG(T) be the set of all big values in T . We call a set $C_j \in \mathcal{V}^c$ *small*

if it contains **SMALL** values and *big* otherwise¹. Let $\sigma, \sigma^+, (n-1) \cdot \sigma^-$ respectively be the sum of all values in $\text{SMALL}(\mathcal{T})$, $\text{SMALL}(\mathcal{V}^g)$ and $\text{SMALL}(\mathcal{V}^c)$, i.e., $\sigma := \sum_{t \in \text{SMALL}(\mathcal{T})} t$, $\sigma^+ := \sum_{t \in \text{SMALL}(\mathcal{V}^g)} t$ and $\sigma^- := (\sum_{t \in \text{SMALL}(\mathcal{V}^c)} t)/(n-1)$. Note that σ^- is equal to the sum of values obtained by picking one value from each small C_j , and $\sigma = \sigma^+ + (n-1)\sigma^-$.

Next, we know that every **BIG** value will be in the range $[3E, 9E^2]$. For all integers r in $[3E, 9E^2]$, let n_r^+, n^- respectively be the number of values in $\text{BIG}(\mathcal{V}^g)$ and the number of sets C_j with integral part of values r . Thus, $(n-1)n_r^- + n_r^+$ items j in $\mathcal{V}^g \cup \mathcal{V}^c$ have $\lfloor j \rfloor = r$.

We now define notation to represent a subset of **BIG** values and their sum. Let X denote a part in a partition of \mathcal{T} . We define the *type* of X by $\tau(X) = \langle \underline{\tau}(X), \bar{\tau}(X) \rangle = (\underline{\tau}_{3E}, \dots, \underline{\tau}_{9E^2}, \bar{\tau}_{3E}, \dots, \bar{\tau}_{9E^2})$; here $\underline{\tau}_r, \bar{\tau}_r$ are resp. the number of values in $\text{BIG}(X \cap \mathcal{V}^g)$ and $\text{BIG}(X \cap \mathcal{V}^c)$ with integer part r . Let $\text{SIZE}(\tau(X)) := \sum_r r(\underline{\tau}_r + \bar{\tau}_r)$ be the total sum of these rounded values in $\text{BIG}(\mathcal{T} \cap X)$.

Using this notation, we design an IP to find an assignment of **BIG** values in an optimal partition. First, observe that every **BIG** value is at most $9E^2$. Thus, if an optimal partition has some part valued more than $18E^2$, we can remove values until the size of this set is in the range $[9E^2, 18E^2]$. Finding a *partial* allocation of **BIG** values that assigns at least $9E^2$ value to all parts suffices to solve **Exists-GC-MMS**, as we can arbitrarily add the unallocated values. Thus, we will only consider types whose size $\text{SIZE}(\cdot)$ is at most $18E^2$.

The variables of the IP correspond to all types τ that satisfy (i) $\text{SIZE}(\tau) \leq 18E^2$, (ii) $\underline{\tau}_r \leq n_r^+$, (iii) $\bar{\tau}_r \leq n_r^-$. Let $\tau^{(1)}, \tau^{(2)}, \dots, \tau^{(\Gamma)}$ be an enumeration of all variables. Intuitively, we consider types that represent valid allocations of items corresponding to the **BIG** values in a **GC-MMS** instance. Every IP variable takes an integer value equal to the number of times the corresponding type is selected. This in turn represents the number of parts in the output allocation that have a subset of **BIG** items as represented by this type.

Lemma 3.5. *The number of IP variables Γ is $O(1)$.*

Proof. Every type with size at most $18E^2$ can have at most $6E$ **BIG** values, as every **BIG** value is at least $3E$. Each value is one of $[3E, 9E^2]$, a constant sized set. Hence, the number of types $\bar{\tau}$ and $\underline{\tau}$ are each at most $(9E^2 - 3E + 1)^{6E}$. The total number of types at most twice this value, hence a constant as E is a constant. The number of variables of the IP is at most the number of types with size at most $18E^2$, hence is constant. \square

Before defining the IP, we define two *cost* functions for every type. These are used to define constraints to allocate **SMALL** items.

¹Note that each $C_j, j \in [m^-]$ contains $n-1$ equal values. i.e., for each C_j , either $\text{SMALL}(C_j) = \emptyset$ or $\text{SMALL}(C_j) = C_j$.

First, define $c(\tau(X)) := \max\{0, 9E^2 - 6E - \text{SIZE}(\tau(X))\}$. The intuition for this function is as follows. Our aim is to create an optimal partition. If the sum of **BIG** values $\text{SIZE}(\tau(X)) < 9(E^2 - E)$, we must add values from **SMALL**(\mathcal{T}) to X . The required sum from **SMALL**, is at least $9E^2 - 9E - \text{SIZE}(\tau(X))$. However, **SMALL**(\mathcal{T}) does not have arbitrarily precise values. As every **SMALL** value is at most $3E$, we may have to add **SMALL** items until the net value of the part becomes $3E$ more than required, i.e., $9E^2 - 6E$. Hence the cost function $c(\tau(X))$ is defined as specified.

The second cost function captures the value that must be added to a part from $\text{SMALL}(\mathcal{V}^g)$. If a part has $c(\tau) > 0$, we can add at most value σ^- to the part from $\text{SMALL}(\mathcal{V}^c)$. Hence, the minimum value from $\text{SMALL}(\mathcal{V}^g)$ is $\sigma^+(\tau(X)) := \max\{0, c(\tau(X)) - \sigma^-\}$.

Using these notions, we define the following IP for finding an allocation of **BIG** values.

$$\sum_{j=1}^{\Gamma} x_j = n; \quad x_j \in \{0\} \cup \mathbb{N}, \forall j \in [\Gamma] \quad (2)$$

$$\sum_{j=1}^{\Gamma} \underline{\tau}_r^{(j)} x_j \leq n_r^+, \forall r \in [3E, 9E^2] \quad (3)$$

$$\sum_{j=1}^{\Gamma} \bar{\tau}_r^{(j)} x_j \leq (n-1)n_r^-, \forall r \in [3E, 9E^2] \quad (4)$$

$$(a) \sum_{j=1}^{\Gamma} c(\tau^{(j)}) x_j \leq \sigma; \quad (b) \sum_{j=1}^{\Gamma} \sigma^+(\tau^{(j)}) x_j \leq \sigma^+ \quad (5)$$

The **Exists-GC-MMS** algorithm is as follows. After applying the pre-processing step, it defines and solves the above IP. If the IP has a solution, then first it considers the items from the **GC-MMS** instance that correspond to the **BIG** values in \mathcal{T} . The algorithm partitions all these items in n bundles by creating n subsets, with x_i subsets corresponding to type $\tau^{(i)}$. After this, it considers the subsets of **BIG** items that do not have total sum of values at least $9E^2 - 9E$. To each of these, it first adds the **SMALL** items corresponding to the small C_j subsets, by adding at most one item from each subset C_j , in any order. If upon adding these, the value of the set is still not $9E^2 - 9E$, it adds items corresponding to the $\text{SMALL}(\mathcal{V}^g)$ set, until the total sum of values is at least $9E^2 - 9E$. The algorithm returns the tuple $(\mathcal{A}, \text{true})$, where \mathcal{A} is the allocation formed by this process. If the IP does not have a solution, it returns the tuple $(\emptyset, \text{flag} = \text{false})$.

Algorithm 2 formally describes **Exists-GC-MMS**. We now analyze the correctness of **Exists-GC-MMS**.

Lemma 3.6. *If $\mu \leq \text{GC-MMS}$, then IP has a solution.*

Proof. As $\mu \leq \text{GC-MMS}$, from Lemma 3.3, there is a valid partition of \mathcal{T} with sum of values of each part at least $9E^2$. Let this partition be T^{IP} . Let $\tau^i = \tau(T_i^{\text{IP}})$

Algorithm 2: Exists-GC-MMS

Input : $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m^-]}, u \rangle, \bar{\epsilon}, \mu$ **Output**: $(\mathcal{A}, \text{True})$ if there exists a $(1 - \bar{\epsilon})$ -GC-MMS allocation \mathcal{A} and $(\emptyset, \text{False})$ otherwise

```
1  $\mathcal{V}^g \leftarrow \{g_1, \dots, g_{m^+}\}, g_j = u(j) \cdot \left(\frac{9E^2}{\mu}\right), j \in \mathcal{G}$   
    $\mathcal{V}^c \leftarrow \bigcup_{j \in [m^-]} C_j, C_j := \{c_j^1, \dots, c_j^{n-1}\},$   
    $c_j^k = u(j, k), \forall (j, k) \in \mathcal{S}_j, \forall \mathcal{S}_j \in (\mathcal{S}_j)_{j \in [m^-]};$   
    $\mathcal{T} \leftarrow \mathcal{V}^g \cup \mathcal{V}^c$   
2 if IP has a solution  $X$  for  $\mathcal{T}$  then  
3    $j \leftarrow 1$   
4   for all  $i : x_i \neq 0$  : do  
5     Create  $x_i$  parts  $P_j$  to  $P_{j+x_i}$   
6     Add BIG values to each  $P_k, k \in [j, j+x_i]$   
       as per  $\tau^{(i)}$  ;  $j \leftarrow j + x_i + 1$   
7   while  $\exists k : \sum_{j \in P_k} j < (9E^2 - 9E)$  do  
8     while  $\sum_{j \in P_k} j < (9E^2 - 9E)$  do  
9       if  $P_k \cap C_j = \emptyset$  for any  $j \in [m^-]$  then  
10        add one value from  $C_j$  to  $P_k$   
11       else add any value from  $\text{SMALL}(\mathcal{V}^g)$   
12        to  $P_k$   
13   while there is an unallocated value  $k$  from  
14      $C_j$  for any  $j \in [m^-]$  do  
15     Add  $k$  to any  $P_i : P_i \cap C_j = \emptyset$   
16   Add remaining unallocated values arbitrarily  
17    $\mathcal{A} \leftarrow$  allocation corresponding to  
18      $P = (P_1, \dots, P_n)$  // use Lemma 3.3  
19   return  $(\mathcal{A}, \text{True})$   
20 return  $(\emptyset, \text{False})$ 
```

be the type of each part, and $\tau^{\text{IP}} = [\tau^1 \dots, \tau^n]$, be the multi-set of types of all parts.

Constraints (2), (3) and (4) hold for τ^{IP} by definition of a valid partition. For any $\tau^i \in \tau^{\text{IP}}$ with $c(\tau^i) = 0$, we have $\sum_{t \in \text{SMALL}(T_i^{\text{IP}})} t \geq c(\tau^i) = 0$, and for any $\tau^i \in \tau^{\text{IP}}$ with $c(\tau^i) > 0$, we have $\sum_{t \in \text{SMALL}(T_i^{\text{IP}})} t \geq 9E^2 - \sum_{t \in \text{BIG}(T_i^{\text{IP}})} t \geq 9E^2 - 6E - \text{SIZE}(\tau^i) \geq c(\tau^i)$. The second inequality holds because the SIZE function rounds down all values, and there are at most $6E$ BIG values in each T_i^{IP} . By adding the above inequality for all $\tau^i \in \tau^{\text{IP}}$, we obtain (5) of the IP.

Since each T_i^{IP} is a subset of a valid part, its corresponding type τ^i has at most one value from each C_j . Therefore, for any $\tau^i \in \tau^{\text{IP}}$ with $\sigma^+(\tau^i) > 0$ we have, for $\sum_{t \in \text{SMALL}(\tau^i \cap \mathcal{V}^g)} t \geq c(\tau^i) - \sum_{t \in \text{SMALL}(\tau^i \cap \mathcal{V}^c)} t \geq c(\tau^i) - \sigma^- \geq \sigma^+(\tau^i)$. Moreover, for any $\tau^i \in \tau^{\text{IP}}$ with $\sigma^+(\tau^i) = 0$ we have, for $\sum_{t \in \text{SMALL}(\tau^i \cap \mathcal{V}^g)} t \geq \sigma^+(\tau^i) = 0$. By adding the above inequality for all $T_i^{\text{IP}} \in T^{\text{IP}}$ we get constraint (5). Thus, T^{IP} is a solution of the IP. \square

Lemma 3.7. *If the IP has a solution, then the allo-*

cation returned by Exists-GC-MMS is an allocation that gives every agent a bundle of value at least $(1 - \bar{\epsilon})\mu$.

Proof. Let τ^{sol} be the solution of the IP and P^{sol} be the partition of the values formed by Exists-GC-MMS after finding τ^{sol} . We show that each part of P^{sol} has value at least $(9E^2 - 9E)$. From Lemma 3.4, we get that in \mathcal{A} , every agent gets a bundle of value at least $(1 - \bar{\epsilon})\mu$.

After assigning BIG values to P_i as per the type τ^i , suppose there are parts with value less than $9E^2 - 9E$.

Consider any such part P . The algorithm first adds SMALL values from \mathcal{V}^g . As τ^{sol} satisfies constraint (5) (a) of the IP, then $c(\tau(P)) \leq \sigma^+$. That is, the value to add to P so that the sum of values in P is at least $9E^2 - 9E$ is at most the sum of all SMALL values. We first add values from $\text{SMALL}(\mathcal{V}^c)$. Suppose after receiving one value from each set in \mathcal{V}^c , P still has value less than $(9E^2 - 9E)$. As τ^{sol} satisfies constraint (b) of (5) of the IP, the total cost from $\text{SMALL}(\mathcal{V}^c)$ for all parts together is at most σ^+ . As the cost function is monotonic with number of parts, the total cost from $\text{SMALL}(\mathcal{V}^c)$ for P also is at most $\text{SMALL}(\mathcal{V}^c)$. Hence, there are enough values in $\text{SMALL}(\mathcal{V}^g)$ to add to P_i to increase its value to at least $(9E^2 - 9E)$.

After adding values to P , its total value is at most $9E^2 - 6E$, as every item has value at most $3E$. Thus, the value added to it from SMALL values is at most $c(\tau(P))$. The total cost of the remaining parts is $\sum_{P' \neq P} c(\tau(P')) = \sum_{P \in P^{\text{sol}}} c(\tau(P)) - c(\tau(P)) \leq \sigma^+ - (\text{the sum of SMALL values assigned to } P)$, which is exactly the total value of unassigned SMALL values. Hence, constraint 5 (a) is satisfied for the smaller set $\tau^{\text{sol}} \setminus \tau(P)$. Similarly, we can show constraint 5 (b) also is satisfied. The initial constraints 2, 4 and 3 are satisfied for $\tau^{\text{sol}} \setminus \tau(P)$ by the validity of τ^{sol} . Hence $\tau^{\text{sol}} \setminus \tau(P)$ is a solution to the IP for the smaller case after removing P and its assigned values. By induction, we can assign values to every part until all parts are satisfied. Adding any unallocated values arbitrarily in Line 13 only increases the value of each bundle.

Hence, the partition P^{sol} obtained has every bundle of value at least $9E^2 - 9E$. From Lemma 3.4, the corresponding allocation \mathcal{A} gives every agent a bundle of value at least $(1 - \bar{\epsilon})\mu$. \square

Lemma 3.8. *Exists-GC-MMS runs in time $O(mn)$.*

Proof. The time to run Exists-GC-MMS is asymptotically equal to the time for constructing and solving the IP. Lenstra's algorithm (Lenstra Jr 1983) takes time exponential in the number of variables, $O(2^{1/\bar{\epsilon}^2}) = O(2^{4/\epsilon'^2})$ here, and polynomial in the largest coefficient of any variable in all inequalities, $m^+ + (n - 1)m^- = O(mn)$ here. Note that σ and σ^+ are at most $n \cdot 9E^2$. Hence, the IP requires $O(2^{1/\epsilon'^2} mn) = O(mn)$ time. \square

Lemmas 3.6, 3.7 and 3.8 together prove Theorem 3.2.

Acknowledgements

Rucha Kulkarni and Ruta Mehta thank the support of NSF Grant CCF-1750436 (CAREER). Setareh Taki is partially supported by NSF Grant CCF-1942321 (CAREER).

References

- Amanatidis, G.; Markakis, E.; Nikzad, A.; and Saberi, A. 2017. Approximation Algorithms for Computing Maximin Share Allocations. *ACM Trans. Algorithms* 13(4): 52:1–52:28.
- Barman, S.; and Krishna Murthy, S. K. 2017. Approximation algorithms for maximin fair division. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 647–664. ACM.
- Bouveret, S.; and Lemaître, M. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems* 30(2): 259–290.
- Brams, S. J.; and Taylor, A. D. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6): 1061–1103.
- De La Vega, W. F.; and Lueker, G. S. 1981. Bin packing can be solved within $1 + \varepsilon$ in linear time. *Combinatorica* 1(4): 349–355.
- Etkin, R.; Parekh, A.; and Tse, D. 2007. Spectrum sharing for unlicensed bands. *IEEE Journal on selected areas in communications* 25(3): 517–528.
- Farhadi, A.; Ghodsi, M.; Hajiaghayi, M. T.; Lahaie, S.; Pennock, D. M.; Seddighin, M.; Seddighin, S.; and Yami, H. 2019. Fair Allocation of Indivisible Goods to Asymmetric Agents. *J. Artif. Intell. Res.* 64: 1–20.
- Garg, J.; McGlaughlin, P.; and Taki, S. 2018. Approximating Maximin Share Allocations. In *2nd Symposium on Simplicity in Algorithms (SOSA 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Garg, J.; and Taki, S. 2020. An Improved Approximation Algorithm for Maximin Shares. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 379–380.
- Ghodsi, M.; Hajiaghayi, M.; Seddighin, M.; Seddighin, S.; and Yami, H. 2018. Fair Allocation of Indivisible Goods: Improvements and Generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*.
- Huang, X.; and Lu, P. 2019. An algorithmic framework for approximating maximin share allocation of chores. *CoRR* abs/1907.04505. URL <http://arxiv.org/abs/1907.04505>.
- Jansen, K.; Klein, K.; and Verschae, J. 2016. Closing the Gap for Makespan Scheduling via Sparsification Techniques. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP*, volume 55, 72:1–72:13.
- Johnson, D. S. 1982. The NP-completeness column: An ongoing guide. *Journal of Algorithms* 3(4): 381–395.
- Karmarkar, N.; and Karp, R. M. 1982. An efficient approximation scheme for the one-dimensional bin-packing problem. In *23rd Annual Symposium on Foundations of Computer Science (sfcs 1982)*, 312–320. IEEE.
- Kulkarni, R.; Mehta, R.; and Taki, S. 2020. Approximating Maximin Shares with Mixed Manna. *CoRR* abs/2007.09133. URL <https://arxiv.org/abs/2007.09133>.
- Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2016. When Can the Maximin Share Guarantee Be Guaranteed? In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI’16*, 523–529. AAAI Press.
- Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2018. Fair Enough: Guaranteeing Approximate Maximin Shares. *J. ACM* 65(2): 8:1–8:27.
- Lenstra Jr, H. W. 1983. Integer programming with a fixed number of variables. *Mathematics of operations research* 8(4): 538–548.
- Moulin, H. 2004. *Fair division and collective welfare*. MIT press.
- Procaccia, A. D.; and Wang, J. 2014. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 675–692. ACM.
- Robertson, J.; and Webb, W. 1998. *Cake-cutting algorithms: Be fair if you can*. CRC Press.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica* 16: 101–104.
- Vossen, T. 2002. *Fair allocation concepts in air traffic management*. Ph.D. thesis, Supervisor: MO Ball, University of Maryland, College Park, Md.
- Woeginger, G. J. 1997. A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Operations Research Letters* 20(4): 149–154.

A Missing Proofs

Lemma 3.2. *If $\text{GC-MMS} \geq (1 + 1/\rho)v^-$, Algorithm 1 returns a $(1 - \epsilon')$ GC-MMS allocation.*

Proof. First, suppose there is a good $j \in \mathcal{G} : u(j) \geq \mu$, the B-MMS allocation corresponding to the GC-MMS allocation returned gives the single good with all chores to one agent, say i . The value of i ’s B-MMS bundle is

$\mu - v^-$. Note that in every iteration of the algorithm, $\mu \geq \text{GC-MMS}$, as the algorithm stops when it finds the first μ which gives every agent a bundle of value at least μ . Hence, from Equation (1), the value of i 's B-MMS bundle is $\mu - v^- \geq \text{GC-MMS} - v^- \geq \text{MMS}$.

By allocating all chores to i , and a single good, the MMS values of the remaining agents over the remaining goods only increases. Hence, a $(1 - \epsilon')$ -MMS allocation of these, combined with i 's bundle, is a $(1 - \epsilon')$ -GC-MMS allocation.

Now consider the second case when for all items $j \in \mathcal{G} : u(j) < \mu$. Denote the value of μ considered by Algorithm 1 for which `Exists-GC-MMS` returns `flag = true` by μ^* . From Theorem 3.2, the corresponding allocation \mathcal{A} returned by `Exists-GC-MMS` has all bundles of value at least $(1 - \bar{\epsilon})\mu^*$.

If `Exists-GC-MMS` returns `flag = true` for the first value of μ considered, then as this value is the upper bound on μ , we have $\text{GC-MMS} \leq \mu^*$. Hence $(1 - \bar{\epsilon})\mu^* \geq (1 - \bar{\epsilon})\text{GC-MMS} \geq (1 - \epsilon')\text{GC-MMS}$, where the last inequality follows as $\bar{\epsilon} = \epsilon'/2$.

Otherwise, for the value of μ considered in the previous iteration before μ^* , that is, for $\mu^*/(1 - \bar{\epsilon})$, `Exists-GC-MMS` returned `flag = false`. Hence $\text{GC-MMS} < \mu^*/(1 - \bar{\epsilon})$. Hence, the smallest valued bundle in \mathcal{A} has value at least $(1 - \bar{\epsilon})\mu^* \geq (1 - \bar{\epsilon})^2 \text{GC-MMS} \geq (1 - \epsilon')\text{GC-MMS}$.

Thus, in both cases, \mathcal{A} is a $(1 - \epsilon')$ -GC-MMS allocation. \square

Lemma 3.3. *Given a GC-MMS instance $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m]}, u \rangle$, if $\mu \leq \text{GC-MMS}$ then there is a valid partition of \mathcal{T} where the sum of values in each part is at least $9E^2$.*

Proof. Since $\mu \leq \text{GC-MMS}$, there is a valid allocation for $\langle \mathcal{N}, \mathcal{G}, (\mathcal{S}_j)_{j \in [m]}, u \rangle$ where all bundles have value at least μ . If we scale the value of each item by $9E^2/\mu$, the value of all bundles will also be scaled by $9E^2/\mu$ because v is additive. Hence we get a valid partition where all parts have a total sum of at least $9E^2$. \square

Lemma 3.4. *A valid partition of \mathcal{T} where the sum of values in each part is at least $9(E^2 - E)$ is equivalent to a valid allocation for its corresponding GC-MMS instance where each bundle has value at least $(1 - \bar{\epsilon})\mu$.*

Proof. Let $P = (P_1, \dots, P_n)$ be the valid partition and let $A = (A_1, \dots, A_n)$ be its corresponding valid allocation.

We have $\sum_{p \in P_i} p = (9E^2/\mu)u(A_i)$. Hence,

$$\begin{aligned} \sum_{p \in P_i} p &\geq 9(E^2 - E) & \forall i \in [n] \\ \left(\frac{9E^2}{\mu}\right)u(A_i) &\geq 9(E^2 - E) & \forall i \in [n] \\ \implies u(A_i) &\geq \mu \left(\frac{9E^2 - 9E}{9E^2}\right) & \forall i \in [n] \\ \implies u(A_i) &\geq \mu(1 - \frac{1}{E}) & \forall i \in [n] \\ \implies u(A_i) &\geq (1 - \bar{\epsilon})\mu & \forall i \in [n]. \end{aligned}$$

The last inequality holds because $E = \lceil \frac{1}{\bar{\epsilon}} \rceil \geq \frac{1}{\bar{\epsilon}}$. \square

B Algorithm for B-MMS when $\text{MMS} < 0$

In this section, we will give an overview of the algorithm for the B-MMS problem when $\text{MMS} < 0$. That is, v^+ is less than v^- , and $|\text{MMS}| \geq v^+/\rho$, for some constant $\rho \geq 1$. To obtain this result, we reduce the B-MMS problem to the CC-MMS problem, defined shortly. We prove that a PTAS for the CC-MMS problem implies a PTAS for the B-MMS problem. To complete the result, in the major part of this Section, we show a PTAS for the CC-MMS problem.

B.1 Technical Overview

In the entire discussion, we will consider the value of each chore as its absolute value, intuitively representing the *cost* of doing the chore.

The outline of the algorithm is the same as the counter case when $\text{MMS} \geq 0$. To avoid repetition, we will focus on highlighting the main ideas that differ from the approach used in Section 3. We first reduce B-MMS to a problem with only chores, denoted by CC-MMS. We substitute every good of B-MMS by a set of $(n - 1)$ chores in CC-MMS, and call them chore-copies of the good. An optimal allocation here has the *lowest* value for the *largest* bundle (call this value CC-MMS). We find via a search algorithm, that runs a subroutine `Exists-CC-MMS` in each iteration, the lowest value μ for which there exists an allocation that gives every agent at most a μ valued bundle.

`Exists-CC-MMS` (Algorithm 4), scales the values of items by $9E^2/\mu$ and calls the set of all scaled valuations (of chores and chore-copies) \mathcal{T} . We prove in Lemma B.5 that a *valid partition* of \mathcal{T} (defined in Definition B.6), defined as one where the sum of values in each part is at most $9E^2 + 9E$, implies a $\mu/(1 - \bar{\epsilon})$ valued valid allocation for the corresponding CC-MMS instance. Then, `Exists-CC-MMS` classifies the set of values in \mathcal{T} as `BIG`(\mathcal{T}) or `SMALL`(\mathcal{T}), solves an IP to find a suitable allocation of values in `BIG`(\mathcal{T}), and allocates the `SMALL`(\mathcal{T}) values greedily upon the allocation of `BIG`(\mathcal{T}).

We then define *types* of subsets of $\text{BIG}(\mathcal{T})$ with a vector $\tau = [\bar{\tau}, \underline{\tau}]$. However, when $\text{MMS} < 0$, we consider the closest integer to the values when they are rounded *up*. The variables of the IP designed for **Exists-CC-MMS** correspond to all type vectors τ that represent all bundles of $\text{BIG}(\mathcal{T})$ with total value *at most* $9E^2 + 9E$, and those that contain at most one chore-copy of every BIG good.

The IP has constraints to select n types that ensure the selection *at least* covers all given chores and chore-copies. We now define two *surplus* functions. The first, denoted by $c(\tau(X))$, represents the cost that can be added to every bundle X from $\text{SMALL}(\mathcal{T})$ while keeping the bundle's value at most $9E^2 + 6E$. This highest allowed value is $3E$ lower than our desired bound of $9E^2 + 9E$. This is because we cannot form bundles of $\text{SMALL}(\mathcal{T})$ values of arbitrary precision. At a high level, if the IP finds a solution where all the $\text{SMALL}(\mathcal{T})$ values can be filled while keeping every bundle's total cost at most $9E^2 + 6E$, then as every $\text{SMALL}(\mathcal{T})$ value is at most $3E$, we can greedily allocate all these without any bundle's cost exceeding $9E^2 + 9E$. Formally, if $\text{Size}(\tau(X))$ is the cost of a bundle from the $\text{BIG}(\mathcal{T})$ values, we define $c(\tau(X))$ for **Exists-CC-MMS** as $\max\{0, 9E^2 + 6E - \text{Size}(\tau(X))\}$. Next, $\sigma^-(\tau(X)) := \min\{c(\tau(X)), \sigma^-\}$. σ^- is the sum of the values of one chore-copy of every good. Thus, $\sigma^-(\tau(X))$ is the maximum surplus in X for adding SMALL chores, whose values are in the set $\text{SMALL}(\mathcal{V}^g)$.

With two constraints using these functions, the IP ensures that we can allocate all the SMALL items. Using the notation from the $\text{MMS} \geq 0$ case, the IP for **Exists-CC-MMS** is as follows.

$$\sum_{j=1}^{\Gamma} x_j = n; \quad x_j \in \{0\} \cup \mathbb{N}, \forall j \in [\Gamma] \quad (6)$$

$$\sum_{j=1}^{\Gamma} \underline{\tau}_r^{(j)} x_j \geq (n-1)n_r^+, \quad \forall r \in [3E, 9E^2] \quad (7)$$

$$\sum_{j=1}^{\Gamma} \bar{\tau}_r^{(j)} x_j \geq n_r^-, \quad \forall r \in [3E, 9E^2] \quad (8)$$

$$\sum_{j=1}^{\Gamma} c(\tau^{(j)}) x_j \geq \sigma; \quad \sum_{j=1}^{\Gamma} \sigma^-(\tau^{(j)}) x_j \geq \sigma^- \quad (9)$$

We prove that **Exists-CC-MMS** solves this IP and returns an allocation which gives every agent a bundle of value at most $9E^2 + 9E$, if $\mu \geq \text{CC-MMS}$.

The detailed algorithm for CC-MMS and its analysis are described in the next section.

B.2 Detailed discussion

Since we consider the negated value of chores, i.e., their cost, all chores have a non-negative cost. All associated values MMS , CC-MMS are also non-negative. Two small

but key implications of this are (a) because of this negation, the MMS value is *at least* the average sum of all values (b) the sum of values of all items of the B-MMS instance is $v^- - v^+$.

We now introduce all notation for the CC-MMS problem and its formal definition.

Definition B.1 (CC-MMS instance). A tuple $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle$, where \mathcal{N} is a set of agents, \mathcal{M} is a set of chores, $(\mathcal{S}_j)_{j \in [m^+]}$ are m^+ sets of chores, each containing $(n-1)$ identical copies of a good, and $u : \mathcal{C} \cup (\mathcal{S}_j)_{j \in [m^+]} \rightarrow \mathbb{R}_+$ is the identical additive valuation function of the agents in \mathcal{N} for all items $\mathcal{C} \cup (\mathcal{S}_j)_{j \in [m^+]}$.

Definition B.2 (Valid allocation). Given a CC-MMS instance $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle$, an allocation \mathcal{A} is valid if no agent receives more than 1 item from any set $\mathcal{S}_j \in (\mathcal{S}_j)_{j \in [m^+]}$, i.e., for all $i \in \mathcal{N}, j \in [m^+]$, $|\mathcal{A}_i \cap \mathcal{S}_j| \leq 1$.

Definition B.3 (CC-MMS value). Given a CC-MMS instance $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle$, let \mathcal{F} be the set of all valid allocations. The CC-MMS value of the instance, denoted by CC-MMS , is defined as follows.

$$\text{CC-MMS} = \argmin_{\mathcal{A} \in \mathcal{F}} \max_{A \in \mathcal{A}} u(A).$$

Finding allocations that give every agent a bundle worth at most her CC-MMS value is NP-Hard, as the MMS problem in a chore manna is a special case of this problem where $(\mathcal{S}_j)_{j \in [m^+]} = \emptyset$. Hence, we define approximate CC-MMS allocations as follows.

Definition B.4 $((1 - \epsilon')\text{CC-MMS allocation})$. Given a CC-MMS instance $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle$, a valid allocation \mathcal{A} is called a $(1 - \epsilon')\text{CC-MMS allocation}$ if

$$u(A) \leq \text{CC-MMS} / (1 - \epsilon') \quad \forall A \in \mathcal{A}.$$

Definition B.5 (CC-MMS problem). Given a CC-MMS instance $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle$ and $\epsilon' > 0$, return a valid allocation \mathcal{A} such that $\max_{A \in \mathcal{A}} u(A) \leq \text{CC-MMS} / (1 - \epsilon')$.

In Section B.3, we show how to reduce the B-MMS problem to a CC-MMS problem, and why a PTAS for CC-MMS implies a PTAS for B-MMS.

B.3 Reducing B-MMS to CC-MMS

Given an instance $\langle \mathcal{N}, \mathcal{M}, v \rangle$ we define the corresponding CC-MMS instance $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle$ as follows. First, $\mathcal{C} = \mathcal{M}^-$. Moreover, for all $j \in \mathcal{M}^+$, define \mathcal{S}_j to be a set of $(n-1)$ chores, each represented by tuples as $\mathcal{S}_j := \{(j, k) \mid k \in [n-1]\}$. Let $(\mathcal{S}_j)_{j \in [m^+]} := \bigcup_{j \in [m^+]} \mathcal{S}_j$ where $m^+ = |\mathcal{M}^+|$. Finally, define $u(j) = |v(j)|$ for all $j \in \mathcal{M}^-$ and $u((j, k)) = v(j)$ for all $j \in \mathcal{M}^+$ and $k \in [n-1]$. Let the chores in \mathcal{S}_j be called chore-copies of good j .

Lemma B.1. *Allocations of B-MMS are in one-to-one correspondence with valid allocations of CC-MMS, such that if allocation B^π of B-MMS corresponds to allocation C^π of CC-MMS, then $u(C_i) = v(B_i) + v^+, \forall i \in \mathcal{N}$.*

Proof. Given a B-MMS allocation B^π , add chore-copies of each good to agents who did not receive the good in B^π , and discard all goods. This gives a valid CC-MMS allocation. The reverse allocation is obtained by similarly discarding all chore-copies and assigning the corresponding good to the agent who did not receive any chore-copy.

Every agent $i \in \mathcal{N}$ receives in C^π all the chores assigned to her in B^π . Every good that was assigned to her in B^π is discarded in C^π . Due to this, the cost of her bundle (negated value) increases by the value of goods allotted to her in B^π . Further, for every good not assigned to her, she receives a chore-copy of it in C^π . Each chore-copy increases her cost by the value of the corresponding good. Hence, her cost increases by the value of all goods not assigned to her as well. Her bundle's total negated value in C^π is exactly, $\sum_{j \in \mathcal{M} \cap B_i} v(j) + \sum_{j \in \mathcal{M} \setminus B_i} v(j) = v^+$ more than in B^π . \square

Corollary B.1. CC-MMS, relates to the MMS value of the B-MMS problem as,

$$\text{CC-MMS} = \text{MMS} + v^+. \quad (10)$$

Lemma B.1 implies the following lemma.

Lemma B.2. If $\text{MMS} \geq v^+/\rho$, a $(1 - \epsilon')$ CC-MMS allocation implies a $(1 - \epsilon)$ -MMS allocation, for $\epsilon = \epsilon'(1 + \rho)/(1 + \rho\epsilon')$. Thus, an algorithm for the CC-MMS problem implies one for the B-MMS problem.

Proof. We take the $(1 - \epsilon')$ CC-MMS allocation, say C^π , and consider the B-MMS allocation B^π corresponding to it, according to the one-to-one correspondence described in the proof of Lemma B.1. From Equation (10), the largest bundle in B^π has value at most $\text{CC-MMS}/(1 - \epsilon') - v^+$.

If $\text{MMS} \geq v^+/\rho$, we have, $\text{CC-MMS}/(1 - \epsilon') - v^+ = (\text{MMS} + v^+)/(1 - \epsilon') - v^+ = \text{MMS}/(1 - \epsilon') + v^+(\epsilon'/(1 - \epsilon')) \leq \text{MMS}/(1 - \epsilon') + (\epsilon'/(1 - \epsilon'))\rho\text{MMS} = (1 + \rho\epsilon')(1 - \epsilon')\text{MMS}$. To have this value at most $\text{MMS}/(1 - \epsilon)$, we solve $(1 + \rho\epsilon')(1 - \epsilon') = 1/(1 - \epsilon)$, and get $\epsilon = \epsilon'(1 + \rho)/(1 + \rho\epsilon')$. As ρ is a constant in the B-MMS problem, ϵ also is.

Given an algorithm for the CC-MMS problem, we find a $(1 - \epsilon')$ CC-MMS allocation for $\epsilon' = \epsilon/(1 + (1 - \epsilon)\rho)$. From the above relation, this is a $(1 - \epsilon)$ -MMS allocation solving B-MMS. \square

Lemma B.2 shows that solving CC-MMS is sufficient to solve B-MMS when $v^- > v^+$ (the other case is explained in Section 3). Algorithm 3 outlines the PTAS for CC-MMS. As described in Section B.1, the algorithm for CC-MMS will perform a search over a range of possible values of CC-MMS. We now define this range. First, we know $\text{MMS} \geq v^+/\rho$. We also have the following two trivial lower bounds on MMS. First, $\text{MMS} \geq (v^- - v^+)/n$. Also, the MMS value is at least the value of the chore with the largest absolute value. This

is because in any allocation, the largest chore, of value say v , has to be added to some bundle. This bundle has value at least v . Hence the value of the highest valued bundle, that is the MMS value, is at least v . Thus, $\text{MMS} \geq \max_{j \in \mathcal{M}} v(j)$. Hence, the lower bound on the MMS value is $\max\{v^+/\rho, (v^- - v^+)/n, \max_{j \in \mathcal{M}} v(j)\}$. A trivial upper bound on the MMS value is the sum of all items, $(v^- - v^+)$. From Lemma B.1, the corresponding bounds for CC-MMS are $\text{CC-MMS} \geq v^+ + \max\{v^+/\rho, (v^- - v^+)/n, \max_{j \in \mathcal{M}} v(j)\}$, and $\text{CC-MMS} \leq (v^- - v^+) + v^+$.

Note that we want an allocation with the smallest cost of the largest bundle. Hence, the best value for MMS is the lower bound.

Each iteration of the search in Algorithm 3 takes a candidate value μ in above range, and works as follows. Note that as we consider values greater than $\max_j v_j$ in the search, every item has value at most μ in every iteration. The algorithm calls a subroutine **Exists-GC-MMS**, which checks if there is a valid allocation where all bundles have value at most $\mu/(1 - \bar{\epsilon})$ for $\bar{\epsilon} = \epsilon'/2$. If such an allocation exists the algorithm returns it and if it does not exist it correctly reports $\mu < \text{CC-MMS}$. We will discuss the **Exists-CC-MMS** algorithm in Section B.4 and prove the following theorem.

Theorem B.1. **Exists-CC-MMS**($\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle, \epsilon, \mu$) returns a tuple $(\mathcal{A}, \text{flag})$, where \mathcal{A} is a $(1 - \epsilon')$ CC-MMS allocation and $\text{flag} = \text{true}$ when $\mu \geq \text{CC-MMS}$, and runs in $O(mn)$ time.

Algorithm 3: Algorithm for CC-MMS

Input : $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle, \epsilon' > 0$
Output: $(1 - \epsilon')$ CC-MMS allocation if
 $\text{CC-MMS} \geq (1 + 1/\rho)v^+$

- 1 $\bar{\epsilon} \leftarrow \epsilon'/2$,
 $\mu \leftarrow \max\{v^+/\rho, \max_j v(j), (v^- - v^+)/n\}$
- 2 **while** $\mu \leq v^-$ **do**
- 3 $(\mathcal{A}, \text{flag}) \leftarrow$
 Exists-CC-MMS($\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle, \bar{\epsilon}, \mu$)
- 4 **if** flag **then**
- 5 **return** \mathcal{A}
- 6 **else**
- 7 $\mu \leftarrow \mu/(1 - \bar{\epsilon})$
- 8 $\mathcal{A} = (A_1, \dots, A_n)$ where $A_i = \{(j, i) : \forall j \in [m^+]\}$
 for $i \in [n - 1]$, $A_n = \mathcal{C}$ // agents 1 to $n - 1$
 each get one chore-copy of all goods.
- 9 **return** \mathcal{A}

Exists-CC-MMS, together with Algorithm 3, is a PTAS for CC-MMS. We search in the range defined above and return a trivial allocation if the value is not found.

Finally, Theorem B.1 and Lemma B.2 are used to show that Algorithm 3 gives an efficient algorithm for the B-MMS problem.

Lemma B.3. *For any CC-MMS instance, Algorithm 3 returns $(1 - \epsilon')$ CC-MMS allocation.*

Proof. While running Algorithm 3, let μ^* be the μ for which Exists-CC-MMS returns $flag = true$. From Theorem B.1, the corresponding allocation \mathcal{A} returned by Exists-CC-MMS has all bundles of value at most $\mu^*/(1 - \bar{\epsilon})$. If Exists-CC-MMS returns $flag = true$ for the first value of μ considered, then as the first value is the lower bound on μ , we have $CC-MMS \geq \mu^*$, hence $\mu^*/(1 - \bar{\epsilon}) \leq CC-MMS/(1 - \bar{\epsilon}) \leq CC-MMS/(1 - \epsilon')$, where the last inequality follows as $\bar{\epsilon} = \epsilon'/2$.

Otherwise, for the value of μ considered in the previous iteration before μ^* , that is, for $\mu^*(1 - \bar{\epsilon})$, Exists-CC-MMS returned $flag = false$. Hence $CC-MMS > \mu^*(1 - \bar{\epsilon})$. Hence, the largest valued bundle in \mathcal{A} has value at least $\mu^*/(1 - \bar{\epsilon}) \leq CC-MMS/(1 - \bar{\epsilon})^2 \leq CC-MMS/(1 - \epsilon')$.

Therefore, in both cases, \mathcal{A} is a $(1 - \epsilon')$ GC-MMS allocation. \square

Lemmas B.2 and B.3, together with Theorem B.1, prove the second half of Theorem 1.1, specified as follows.

Theorem B.2. *There is an algorithm for the B-MMS problem for the case $MMS < 0$ (before negating valuations) that runs in time $O(mnL)$, where L is the bit-length of the input.*

It now remains to discuss the Algorithm Exists-CC-MMS and prove Theorem B.1. The next section discusses this.

B.4 Algorithm for Exists-CC-MMS

In this section, we describe the algorithm Exists-CC-MMS, which given a CC-MMS instance, a value μ , and a constant $\bar{\epsilon} > 0$, either outputs a valid allocation where all bundles have value at least $\mu/(1 - \bar{\epsilon})$ or correctly reports $\mu < CC-MMS$.

The Algorithm has two steps 1) Pre-processing and 2) Main Algorithm.

Pre-processing. (Algorithm 4, line 1) Let $E := \lfloor \frac{1}{\bar{\epsilon}} \rfloor$. Note that E is a constant integer that only depends on $\bar{\epsilon}$ and not on parameters in the CC-MMS instance. Scale the valuations v by $9E^2/\mu$.

Let \mathcal{V}^c and \mathcal{V}^g be multi-sets of numbers corresponding to scaled valuations, respectively of all the chores in \mathcal{C} , and all the chore-copies in $(\mathcal{S}_j)_{j \in [m^+]}$. Formally, we define $\mathcal{V}^c, \mathcal{V}^g$ and sets G_j for $j \in [m^+]$ as follows.

$$c_j := u(j) \cdot \left(\frac{9E^2}{\mu} \right), j \in \mathcal{C},$$

$$g_j^k := u(j, k) \cdot \left(\frac{9E^2}{\mu} \right), \forall (j, k) \in \mathcal{S}_j, \forall \mathcal{S}_j \in (\mathcal{S}_j)_{j \in [m^+]}$$

$$G_j := \{g_j^1, \dots, g_j^{n-1}\}, j \in \mathcal{M}^+$$

$$\mathcal{V}^c := \{c_1, \dots, c_{m-}\}, \quad \mathcal{V}^g := \bigcup_{j \in [m^+]} G_j.$$

Finally, let $\mathcal{T} = \mathcal{V}^g \cup \mathcal{V}^c$.

This completes the pre-processing step. The following lemmas characterize partitions of \mathcal{T} that correspond to approximately optimal CC-MMS allocations.

Definition B.6 (Valid Partition of \mathcal{T}). *We call a partition $P = (P_1, \dots, P_n)$ of values in \mathcal{T} valid if each P_k contains at most one element from each G_j , i.e., $|P_k \cap G_j| \leq 1$ for all $k \in [n]$ and $j \in [m^+]$.*

It is easy to see that each valid partition of \mathcal{T} is equivalent to a valid allocation in its corresponding CC-MMS instance. The next lemma is a direct implication of this observation.

Lemma B.4. *Given a CC-MMS instance $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle$, if $\mu \geq CC-MMS$ then there is a valid partition of \mathcal{T} where the sum of values in each part is at most $9E^2$.*

Proof. Since $\mu \geq CC-MMS$, there is a valid allocation for $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle$ where all bundles have value at most μ . If we scale the value of each item by $9E^2/\mu$, the value of all bundles will also be scaled by $9E^2/\mu$ because v is additive. Hence we get a valid partition where all parts have a total sum of at most $9E^2$. \square

Lemma B.5. *Finding a valid partition of \mathcal{T} where the sum of values in each part is at most $9(E^2 + E)$ is equivalent to finding a valid allocation for its corresponding CC-MMS instance where each bundle has value at most $\mu/(1 - \bar{\epsilon})$.*

Proof. Let $P = (P_1, \dots, P_n)$ be the valid partition and let $A = (A_1, \dots, A_n)$ be its corresponding valid allocation. We have $\sum_{p \in P_i} p = (9E^2/\mu)u(A_i)$. We have

$$\begin{aligned} \sum_{p \in P_i} p &\leq 9(E^2 + E) & \forall i \in [n] \\ \left(\frac{9E^2}{\mu} \right) u(A_i) &\leq 9(E^2 + E) & \forall i \in [n] \\ \implies u(A_i) &\leq \mu \left(\frac{9E^2 + 9E}{9E^2} \right) & \forall i \in [n] \\ \implies u(A_i) &\leq \mu \left(1 + \frac{1}{E} \right) & \forall i \in [n] \\ \implies u(A_i) &\leq (1 + \bar{\epsilon})\mu & \forall i \in [n]. \\ \implies u(A_i) &\leq \mu/(1 - \bar{\epsilon}) & \forall i \in [n]. \end{aligned}$$

The second-last inequality holds because $E = \lfloor \frac{1}{\bar{\epsilon}} \rfloor \leq \frac{1}{\bar{\epsilon}}$. \square

Main Algorithm. This step obtains a valid partition of \mathcal{T} where the sum of values in each part is at most $9(E^2 + E)$, if $\mu \geq CC-MMS$, else correctly reports this by returning $flag = false$. Note that after scaling by $9E^2$, we can assume $t \leq 9E^2 \forall t \in \mathcal{T}$.

The key of the algorithm is an IP to allocate the BIG items. We first explain the IP.

Notation. We define SMALL and BIG values in \mathcal{T} . Call a value $t \in \mathcal{T}$ SMALL if $t < 3E$ and BIG if $t \geq 3E$. For each $T \subseteq \mathcal{T}$ let $\text{SMALL}(T)$ be the set of all small values in T and $\text{BIG}(T)$ be the set of all big values in T . We abuse notation and call a set $G_j \in \mathcal{V}^g$ *small* if it contains SMALL values and *big* otherwise². Let σ be the sum of all values t in $\text{SMALL}(\mathcal{T})$, i.e., $\sigma := \sum_{t \in \text{SMALL}(\mathcal{T})} t$. Let $\sigma^- := \sum_{t \in \text{SMALL}(\mathcal{V}^c)} t$ be the total sum of all values t in $\text{SMALL}(\mathcal{V}^c)$, and $\sigma^+ := (\sum_{t \in \text{SMALL}(\mathcal{V}^g)} t)/(n-1)$ be $1/(n-1)$ of the sum of SMALL values in \mathcal{V}^g . Note that σ^+ is equal to the sum of values obtained by picking one value from each small G_j , and $\sigma = \sigma^- + (n-1)\sigma^+$.

Next, we know that every BIG value will be in the range $[3E, 9E^2]$. Consider all integers r in $[3E, 9E^2]$. Let n_r^- be the number of values in $\text{BIG}(\mathcal{V}^c)$ with $\lceil g_j \rceil = r$ and let n_r^+ be the number of *big* G_j sets which have $\lceil g_j^k \rceil = r$ for all $k \in [n-1]$. Thus, there are $(n-1)n_r^+ + n_r^-$ items in $\mathcal{V}^g \cup \mathcal{V}^c$ with values whose integral part is r .

We now define the notation to represent a subset of BIG values and their sum. Let X denote a part in a partition of \mathcal{T} . We define the *type* of X by $\tau(X) = \langle \underline{\tau}(X), \bar{\tau}(X) \rangle = (\underline{\tau}_{3E}, \underline{\tau}_{3E+1}, \dots, \underline{\tau}_{9E^2}, \bar{\tau}_{3E}, \bar{\tau}_{3E+1}, \dots, \bar{\tau}_{9E^2})$; here $\underline{\tau}_r$ is the number of values in $\text{BIG}(X \cap \mathcal{V}^c)$ with integer part r and $\bar{\tau}_r$ is the number of values in $\text{BIG}(X \cap \mathcal{V}^g)$ with integer part r . Let $\text{SIZE}(\tau(X)) := \sum_{r=3E}^{9E^2} r(\underline{\tau}_r + \bar{\tau}_r)$ be the total sum of these rounded values in $\text{BIG}(\mathcal{T} \cap X)$.

Using this notation, we design an IP to find an assignment of BIG values in a $(1 - \bar{\epsilon})\text{CC-MMS}$ partition. That is, a selection of n types that allocate all BIG values such that, there is a valid way to add SMALL values and obtain an allocation where each part has value at most $9(E^2 + E)$. Thus, we will only consider types whose size $\text{SIZE}(\cdot)$ is at most $9E^2 + 9E$.

The variables of the IP correspond to all types that satisfy (i) $\text{SIZE}(\tau^{(j)}) \leq 9E^2 + 9E$, (ii) $\underline{\tau}_r^{(j)} \leq n_r^-$, (iii) $\bar{\tau}_r^{(j)} \leq n_r^+$. That is, every type represents a subset of BIG values that have at most n_r^- values from \mathcal{V}^g and n_r^+ values from \mathcal{V}^c with integral part of value r . Let us call the items of CC-MMS corresponding to the BIG (resp. SMALL) values as BIG (resp. SMALL) items. Intuitively, we allow types that represent valid allocations of BIG items of the CC-MMS instance. In the solution, every IP variable takes an integer value equal to the number of times the corresponding type is selected. This in turn represents the number of parts in the output allocation that have a subset of BIG items as represented by this type.

²Note that each G_j , $j \in [m^+]$ contains $(n-1)$ equal values. i.e., for each G_j , either $\text{SMALL}(G_j) = \emptyset$ or $\text{SMALL}(G_j) = G_j$.

Lemma B.6. *The IP has $O(1)$ number of variables.*

Proof. Every type with size at most $9E^2 + 9E$ can have at most $3E + 3$ BIG values, as every BIG value is at least $3E$. Each value is one of $[3E, 9E^2]$, a constant sized set. Hence, the number of types are at most $(9E^2 - 3E + 1)^{(3E+3)}$, which is a constant as E is a constant. The number of variables of the IP is at most the number of types with size at most $9E^2 + 9E$, hence is constant. \square

Before defining the IP, we define two *cost* functions for every type. The intuition for this function is described in the technical overview (Section B.1). These are used to define constraints to allocate SMALL items.

Define $c(\tau(X)) := \{0, 9E^2 + 6E - \text{SIZE}(\tau(X))\}$, and $\sigma^-(\tau(X)) := \min\{c(\tau(X)), \sigma^-\}$.

Using the above notation, the IP is as defined in Section B.1.

Algorithm 4: Exists-CC-MMS

Input : $\langle \mathcal{N}, \mathcal{C}, (\mathcal{S}_j)_{j \in [m^+]}, u \rangle, \bar{\epsilon}, \mu$

Output: $(\mathcal{A}, \text{True})$ if there exists a $(1 - \bar{\epsilon})\text{-CC-MMS}$ allocation \mathcal{A} and $(\emptyset, \text{False})$ otherwise

```

1  $\mathcal{V}^g \leftarrow \bigcup_{j \in [m^-]} G_j, G_j := \{g_j^1, \dots, g_j^{n-1}\};$ 
    $g_j^k = u(j, k), \forall (j, k) \in \mathcal{S}_j, \forall \mathcal{S}_j \in (\mathcal{S}_j)_{j \in [m^-]};$ 
    $\mathcal{V}^c \leftarrow \{c_1, \dots, c_{m^+}\}, c_j = u(j) \cdot \left(\frac{9E^2}{\mu}\right), j \in \mathcal{G};$ 
    $\mathcal{T} \leftarrow \mathcal{V}^g \cup \mathcal{V}^c$ 
2  $P = \{P_1, \dots, P_n\} \leftarrow (\emptyset, \dots, \emptyset)$ 
3 if IP has a solution  $X$  for  $\mathcal{T}$  then
4    $j \leftarrow 1$ 
5   for all  $i : x_i \neq 0$  : do
6     Add BIG values to each  $P_k$ ,  $k \in [j, j + x_i]$ 
     as per  $\tau^{(i)}$ ;  $j \leftarrow j + x_i + 1$ 
7   while there is  $i : \sigma^-(\tau^{(i)}) > 0$  and some
     unallocated value from  $\mathcal{V}^g$  do
8     while  $\exists j : G_j \cap P_i \neq \emptyset$  and  $\sigma^-(\tau^{(i)}) > 0$  do
9       assign one value from  $G_j$  to  $P_i$ 
10       $\sigma^-(\tau^{(i)}) \leftarrow \sigma^-(\tau^{(i)}) - u(j)$ 
11      while  $\exists j : G_j \cap P_i = \emptyset$  do
12        assign one value from  $G_j$  to each part
        in  $\mathcal{N} \setminus \{i\}$ 
13   while there is an unallocated  $j \in \mathcal{V}^c$  do
14     for any  $i : c(\tau^{(i)}) > 0 : P_i \leftarrow P_i \cup \{j\}$ 
15      $c(\tau^{(i)}) \leftarrow \max\{c(\tau^{(i)}) - u(j), 0\}$ 
16    $\mathcal{A} \leftarrow$  Allocation of items corresponding to  $P$ 
   return  $(\mathcal{A}, \text{True})$ 
17 return  $(\emptyset, \text{False})$ 
```

The Exists-CC-MMS algorithm is as follows. After applying the pre-processing step, it defines and solves the IP from Section B.1. If the IP has a solution, then first it considers the items from the CC-MMS instance that

correspond to the BIG values in \mathcal{T} . The algorithm partitions these items in n bundles by creating x_i bundles corresponding to type $\tau^{(i)}$. After this, it considers the bundles of BIG items that do not have total sum of values at most $9E^2 + 9E$. To each of these, it first adds the SMALL items corresponding to the small G_j subsets, by adding at most one item from each subset G_j , in any order. After all items from all sets G_j are allocated, it considers the parts whose sum is still less than $9E^2 + 9E$. It adds items corresponding to $\text{SMALL}(\mathcal{V}^c)$ to any of these, until all items are exhausted. The algorithm returns the tuple \mathcal{A}, flag , where \mathcal{A} is the allocation formed by this process, and $\text{flag} = \text{true}$.

If the IP does not have a solution, it returns the tuple $(\emptyset, \text{flag} = \text{false})$.

Algorithm 4 formally describes **Exists-CC-MMS**. We now analyze the correctness of **Exists-CC-MMS**, and the run time of Algorithm 3.

Lemma B.7. *If $\mu \geq \text{CC-MMS}$, then the IP has a solution.*

Proof. Lemma B.4 shows that there is a partition of \mathcal{T} where all bundles have value at most $9E^2$, as $\mu \geq \text{CC-MMS}$. Let $T^* = (T_1^*, \dots, T_n^*)$ be this partition. Let $\tau^* = [\tau^1, \tau^2, \dots, \tau^n] = [\tau(T_i^*)]_{(i \in [n])}$ be the multi-set of types of each part in T^* .

For τ^* , constraints (6), (7), and (8) hold by definition of a valid partition. For any $\tau^i \in \tau^*$ we have,

$$\begin{aligned} \sum_{t \in \text{SMALL}(T_i^*)} t &\leq 9E^2 - \sum_{t \in \text{BIG}(T_i^*)} t \\ &\leq 9E^2 + 3E - \text{SIZE}(\tau^i) \\ &\leq c(\tau^i). \end{aligned}$$

The second inequality holds because the SIZE function rounds up all values, and there are at most $3E$ BIG values in each T_i^* . By adding the above inequality for all $\tau^i \in \tau^*$, we obtain (9) of the IP.

Since each T_i^* is a subset of a valid part, its corresponding type τ^i has at most one value from each C_j . Therefore, for any $\tau^i \in \tau^*$ we have,

$$\begin{aligned} \sum_{t \in \text{SMALL}(\tau^i \cap \mathcal{V}^g)} t &\leq c(\tau^i) - \sum_{t \in \text{SMALL}(\tau^i \cap \mathcal{V}^c)} t \\ &\leq c(\tau^i) - \sigma^- \leq \sigma^+(\tau^i). \end{aligned}$$

By adding the above inequality for all $T_i^* \in T^*$ we obtain the constraint (9) for IP. Thus, T^* is a solution of the IP. \square

Lemma B.8. *If the IP has a solution, then the allocation returned by **Exists-CC-MMS** after finding this solution is a valid allocation of **CC-MMS** where all bundles have value at most $9E^2 + 9E$.*

Proof. Let τ^{sol} be the solution of the IP, P^{sol} be the partition of the BIG values corresponding to these

types, and \mathcal{A} be the allocation returned by Algorithm **Exists-GC-MMS** after finding τ^{sol} .

The algorithm adds **SMALL** items by first adding chore-copies. Initially, when all chore-copies are unallocated, as τ^{sol} satisfies constraint (b) of (9) of the IP, the total cost $\sigma^-(\tau(.))$ for $\text{SMALL}(\mathcal{V}^g)$ for all bundles is at least σ^- . Hence, there is at least one agent with cost σ^- more than 0. The algorithm assigns as items from different sets G_j to this agent, until her cost $\sigma^-(\tau^i)$ becomes zero. At this point, either i has exhausted one chore-copy of every good, or her cost $c(\tau^i) = 0$. If the former is true, we have a smaller problem with $(n - 2)$ copies of every good remaining, and n agents, hence by inductive reasoning, the algorithm assigns all chore-copies. If $c(\tau^i) = 0$, then first we assign one chore-copy of each good j that i did not get, to every remaining agent, to obtain a valid partition. Note that the sum of costs $\sigma^-(\tau^k)$ for the remaining agents $\mathcal{N} \setminus \{i\}$, is equal to the space in each bundle for adding at most one copy of *every* chore, hence at least j , is more than the total cost of the unallocated items, including j , in \mathcal{V}^g , hence we can do this. The cost $\sigma^-\tau$ for all bundles now is $\sum_{k \in \mathcal{N}, k \neq i} \sigma^-\tau^k = \sum_{k \in \mathcal{N}} \sigma^-\tau^k - \sigma^-\tau^i - (n - 1) \sum_{j: G_j \cap A_i = \emptyset} u(j) \geq \sigma^- - \sigma^-\tau^i - (n - 1) \sum_{j: G_j \cap A_i = \emptyset} u(j)$, which is exactly the cost of unallocated chore-copies. Hence, again by induction, we can repeat the process.

Finally, as constraint (9) is satisfied, until there is some unallocated chore, there will be some bundle A_i with $c(\tau(A_i)) > 0$. Line 8 and 14 can add one **SMALL** chore to A_i . Hence, all **SMALL** chores can be allocated. \square

Lemma B.9. ***Exists-CC-MMS** runs in time $O(mn)$.*

Proof. The time to run **Exists-CC-MMS** is asymptotically equal to the time for constructing and solving the IP. Lenstra's algorithm (Lenstra Jr 1983) takes time exponential in the number of variables, $O(2^{1/\epsilon^2}) = O(2^{1/\epsilon})$ here, and polynomial in the largest coefficient of any variable in all inequalities, $m^- + (n - 1)m^+ = O(mn)$ here. Hence, it requires $O(2^{1/\epsilon}mn) = O(mn)$ time to solve the IP. \square