

Dynamic Fair Division with Minimal Disruptions

Eric Friedman*

Christos-Alexandros Psomas[†]

Shai Vardi[‡]

Abstract

In this paper we present an analysis of dynamic fair division of a divisible resource, with arrivals and departures of agents. Our key requirement is that we wish to disrupt the allocation of a small number of existing agents whenever a new agent arrives. We construct optimal recursive mechanisms to compute the allocations and provide tight analytic bounds. Our analysis relies on a linear programming formulation and a reduction of the feasible region of the LP into a class of “harmonic allocations”, which play a key role in the trade-off between the fairness of current allocations and the fairness of potential future allocations. We show that there exist mechanisms that are optimal with respect to fairness and are also Pareto efficient, which is of fundamental importance in computing applications, as system designers loathe to waste resources. In addition, our mechanisms satisfy a number of other desirable game theoretic properties.

1 Introduction

The study of fair allocations has a long history in economics and mathematics (e.g., [1, 3, 4, 6, 15, 16, 17]), and it has recently seen a surge of interest in computer science due to its applications to resource sharing in data centers and the cloud (e.g., [2, 5, 8, 9, 10, 14, 19]). For example, one commonly used protocol for allocating tasks is the “dominant resource fairness” protocol [9], which (approximately) implements the famous Kalai-Smorodinsky solution from classical fair division theory [11]. There are several notions of what it means for an allocation to be “fair”. The two most widely accepted notions are (1) *proportionality* (or simply *fairness*): if there are n agents in the system, each agent is allocated at least a $1/n$ fraction of what she would receive if she were allocated all of the resource, and (2) *envy freeness*: no agent would like to “switch” allocations with any other agent. In this paper, we consider both notions, but focus on the first, for several reasons. First, in the allocation of a single homogeneous resource, fairness immediately implies envy freeness. Second, in the dynamic setting, the notion of envy-freeness is more fleeting, as an agent may wish to switch with another agent over some time period but not over another. We discuss this in more detail later on.

The classical theory typically focuses on static allocations and does not consider dynamics, which are an important consideration in modern computing systems. Walsh [18] was the first to study the problem of online fair cake cutting when agents arrive, get a piece allocated to them and depart. He showed how several well-known fair division solutions (cut-and-choose, Dubins-Spanier, etc) can be adapted to satisfy desirable properties in an online setting with a single (heterogeneous) divisible cake. More recently, and closer to our problem, Kash et al. [12] introduced

*International Computer Science Institute and UC Berkeley. Email: ejf@icsi.berkeley.edu. Supported by grants NSF-1216073, NSF-1161813.

[†]International Computer Science Institute and UC Berkeley. Email: alexpsomi@cs.berkeley.edu. Supported by the Leventis Foundation.

[‡]Tel Aviv University. Email: shaivar1@post.tau.ac.il. Supported by the Google Europe Fellowship in Game Theory.

a model of dynamic allocations. However, their model only considers arrivals and their main mechanism reserves resources for future arrivals; it does not allow the reallocation of resources, or agents to depart. This leads to allocations that satisfy neither fairness nor Pareto efficiency, as resources are left idle.

In this paper, we consider the problem of maintaining fairness and Pareto efficiency in a setting with uncertain arrivals and departures. At first glance, this appears quite simple: at every arrival or departure, simply compute a Pareto optimal fair division of resources among the current agents. However, this is problematic in practice due to the “expense” of reallocating resources. Imagine that whenever an agent arrives, all jobs would be stopped, the resources reallocated, and the jobs restarted from where they left off. If there are many successive arrivals, most of the time would be spent on reallocation, instead of on actually performing the jobs! For this reason, most resource allocation protocols in data centers rarely move existing jobs from one machine to another [13]. We therefore focus on dynamic mechanisms that minimize the disruption to existing agents due to arrivals and departures. As reductions of resources are typically more disruptive to existing computations than additions, we consider the construction of dynamic mechanisms that reduce the allocations of only a small number of existing agents.

To see some of the issues involved in this problem, consider the case where there is 1 unit of a single resource to be shared, and we only allow a single agent to have her resources reduced upon an arrival. When the first agent arrives, she is given the entire resource to use, in order to satisfy fairness. When the second agent arrives, we give her half of the first agent’s resource; this is, of course, a fair division of the resource. When a third agent arrives, however, one cannot maintain fairness while only disturbing one of the existing agents. In this case, the fairest reallocation leads to one agent with half of the resources and 2 agents with one quarter of the resources each, an unfair allocation of $(1/2, 1/4, 1/4)$. The only fair division in this case would be $(1/3, 1/3, 1/3)$ (or, more generally, $(1/t, \dots, 1/t)$ when there are t agents in the system). We can define the unfairness of this allocation using the *fairness ratio* - the ratio of the smallest share to $1/t$. The example above has a fairness ratio of $3/4$.

It is possible to reduce the unfairness of this final allocation by allowing some unfairness earlier in the process. For example, when the second agent arrives, we could divide the resources to create the allocation $(4/7, 3/7)$, and then when the third agent arrives, divide the first agent’s share by 2, giving the allocation $(2/7, 3/7, 2/7)$. This gives a fairness ratio of $6/7$, a significant improvement. In fact, we will see in Section 4 that this is the fairest possible series of allocations for this particular scenario.

Note the above example is for the case where we allow a single disruption upon arrival. However, if we allowed 2 disruptions per arrival, it would be trivial to compute fair allocations for 3 agents. Nonetheless, the problem would return as soon as a 4th agent arrived. When dealing with departures, we would like to be even stricter than arrivals: when an agent leaves, we do not want any of the remaining agents to have their share decreased: it would be unreasonable to allow one agent’s departure to adversely affect another’s share.

In the following, we analyze this allocation problem as a function of d , the maximum number of disruptions allowed per arrival, and N , the maximal number of users that can be concurrently present in the system. We prove upper and lower bounds for the fairness ratios and present mechanisms for attaining optimal fairness. In particular, we show the importance of harmonic sequences in allocations that balance current fairness ratios with potential future ratios. Our analysis relies on a linear programming formulation for computing allocations with optimal fairness in the case that we only allow arrivals, and a reduction of the feasible region of the LP into the class of “sound” allocations.¹ We provide a recursive technique for preserving soundness in the presence of departures, and use this to construct a mechanism for computing these allocations in

¹See Section 4.1 for a formal definition.

a “bounded” system (one with an upper bound on the number of agents present in it). We show how to extend this mechanism to the unbounded case, by combining it with a simple “harmonic” mechanism.

Although our main focus is the *fairness* of our mechanisms, we also bound the *envy*: the ratio of the largest share and the smallest share.

Overview of Results

In Section 2 we present our definitions and some algebraic results that we use extensively later. In Section 3 we formally define our model and prove our first (negative) result: requiring that only a constant number of users be disrupted at every round is incompatible with fairness. We then provide a notion of approximate fairness, and focus on maximizing it.

In Section 4 we address a special case: assume the agents are non-departing, and the number of agents the system can service at any time is bounded by N (which is known in advance); i.e., the system has bounded capacity. We describe a linear program that solves this problem optimally. We then show that we can guarantee the same fairness even if we demand Pareto optimality at every round. Unfortunately, the allocations output by the LP are not necessarily implementable if we allow departures. Nevertheless, we notice that the LP allocations satisfy a certain property we call *soundness*. Informally, for some fairness ratio σ , an allocation for t agents is (t, σ, d) -sound if: the $d + 1$ poorest agents are allocated at least $\frac{\sigma}{t}$, the next $d + 1$ poorest are allocated at least $\frac{\sigma}{t-1}$, and so on – a staggered harmonic progression.

In Section 5 we describe a mechanism that handles departures, the SOUND mechanism. This mechanism takes as input a (N, σ, d) -sound allocation, for N agents, and builds allocations $\mathcal{A} = \{\mathcal{A}^1, \mathcal{A}^2, \dots, \mathcal{A}^N\}$ backwards in a way that: (1) it is possible to “legally”² convert \mathcal{A}^t to \mathcal{A}^{t+1} and vice versa, (2) the fairness ratio does not decrease.

In Section 6 we show how to apply SOUND to the solutions given by LP 4.1, and hence obtain an optimal mechanism (SOUND – LP) for bounded systems. We provide upper and lower numerical bounds on the optimal fairness ratio, and bound the envy. Our first positive result is the following:

Theorem 1.1 (Informal). *The SOUND – LP mechanism has a worst case fairness ratio of $\frac{1}{(d+1) \ln(\frac{d+1}{d})}$ and a worst case envy ratio of $\frac{3(d+1)^2}{2d^2}$.*

This mechanism requires a preprocessing time that is polynomial in N and then is linear in the number of users in the system at arrivals and departures to compute the new allocations. Note that we achieve the same fairness as the arrivals-only case. Thus, even if we allow our mechanisms to disrupt users on departures, the fairness cannot improve.

One can compare this mechanism to a simple myopic mechanism (for simplicity, we define the mechanism only with respect to arrivals):

Mechanism 1.2 (Myopic). *As long as there are d or fewer agents in the system, split the resource evenly among them. Otherwise, upon arrival take the resource allocated to the d wealthiest agents, and split the total into $d + 1$ equal shares.*

It is easy to show, using simple calculus that Mechanism 1.2 has a fairness ratio of $\frac{d}{d+1}$, which is significantly worse than that of SOUND – LP, and an envy ratio of $\frac{d}{d+1}$ as well. Due to space considerations we do not formally prove this, but give the following example as intuition. Consider the case $d = 1$. The shares will always be integer powers of $\frac{1}{2}$. When there are 2^k agents

²By legally, we mean that when an agent arrives, we may disrupt a specified number of agents, and when an agent departs, the other agents’ shares can only increase.

in the system, all will be allocated $\frac{1}{2^k}$, giving a fairness ratio of 1. When another agent arrives, she will be allocated $\frac{1}{2^{k+1}}$, as opposed to the optimal $\frac{1}{2^{k+1}}$. The ratio $\frac{2^k+1}{2^{k+1}}$ tends to $\frac{1}{2}$ as $k \rightarrow \infty$.

In Section 7 we remove the assumption that the system’s capacity is bounded and present an elegant procedure, the HARMONIC mechanism (Subsection 7.1): At all times, the mechanism maintains an allocation similar to the tail of the harmonic series. The fairness ratio of HARMONIC improves as agents arrive and approaches $((d+1) \ln(\frac{d+1}{d}))^{-1}$ as the number of agents approaches infinity. This is the same as the fairness ratio of the bounded capacity mechanism, thus HARMONIC is asymptotically optimal in the presence of many agents. The downside of HARMONIC is that if only a small number of agents are present in the system, it has a low fairness ratio. We show how to get around this issue in Subsection 7.2, combining it with the SOUND mechanism to prove the following:

Theorem 1.3 (Informal). *The SOUND-HARMONIC mechanism can attain a fairness ratio arbitrarily close to $\frac{1}{(d+1) \ln(\frac{d+1}{d})}$, for any number of agents in the system.*

We also bound the envy ratio of this mechanism by $3(d+1)^2/2d^2$, the same as for the bounded case (we also show that when there are many agents in the system, the envy is actually much better).

2 Notation and Preliminaries

Denote by $[a, b]$ the set of integers between (and including) a and b .

There is a single unit of a divisible resource to be allocated among arriving and departing agents. We call each arrival or departure of an agent a “round”. Throughout the paper, we denote the number of agents currently present in the system by t . An allocation $\mathcal{A}^t = (\mathcal{A}_1^t, \mathcal{A}_2^t, \dots, \mathcal{A}_t^t)$ gives a fraction \mathcal{A}_i^t of the resource to agent i ; each agent’s allocation (or *share*) is a number between 0 and 1. We say that \mathcal{A}^t is *feasible* if $\sum_{i \in [1, t]} \mathcal{A}_i^t \leq 1$. In each round t , we renumber the agents so that 1 has the largest allocation and t the smallest. The more resource an agent is allocated, the *wealthier* she is said to be. An allocation \mathcal{A}^t is *Pareto optimal (PO)* if there is no other allocation \mathcal{B}^t in which everyone is at least as well off and someone is strictly better off, i.e., $\sum_{i \in [1, t]} \mathcal{A}_i^t = 1$.

An allocation is *fair* if $\mathcal{A}_i^t \geq \frac{1}{t}$ for all agents i . Note that a fair allocation is always Pareto optimal. An allocation has *fairness ratio* σ_t in the presence of t agents, if it is σ_t close to the “equal split”, i.e.,

$$\sigma_t = \frac{\mathcal{A}_t^t}{1/t}.$$

An allocation \mathcal{A}^t is *envy-free (EF)* if each agent prefers her allocation to that of any other agent. This is equivalent to the condition $\mathcal{A}_1^t = \mathcal{A}_t^t$. We define the *envy ratio* of an allocation to be

$$\xi_t = \frac{\mathcal{A}_1^t}{\mathcal{A}_t^t}.$$

The allocation is envy-free if the envy ratio is 1. Notice that a fair allocation is always envy-free, but not vice versa (a non-PO allocation can be envy-free and unfair, e.g., $(1/3, 1/3)$).

An allocation mechanism is *d-disruptive* if whenever an agent arrives, the mechanism reduces (or *disrupts*) the shares of at most d agents. We refer to the agents whose allocation is reduced as *donors*. An allocation mechanism satisfies *population monotonicity (PM)* if whenever an agent departs, the allocation of the remaining agents can only increase.

We say that a set of allocations $\{\mathcal{A}^t, t \in [1, N]\}$ for some (possibly unbounded) N is *valid* for a d -disruptive system if

1. For every $t \in [1, N - 1]$, we can obtain \mathcal{A}^{t+1} from \mathcal{A}^t by reducing the shares of at most d agents in \mathcal{A}^t and allocating the difference among the $t + 1 - d$ other agents (including the arriving agent).
2. For every $t \in [2, N]$, we can obtain \mathcal{A}^{t-1} from \mathcal{A}^t by removing any agent and distributing its share among the rest of the agents, without reducing any other allocation.

In other words, a set of allocations is valid if we can move “back and forth” between them when an agent arrives or departs, while satisfying d -disruptiveness and population monotonicity.

A system is said to have a *bounded capacity* if there is some upper bound N on the number of agents it can serve.

Some Simple Results on Harmonic Numbers

We require the following well-known inequalities for the n^{th} harmonic number H_n , (where $\gamma' = 0.57721 \dots$ is the Euler-Mascheroni constant):

Fact 2.1.

$$H_n \geq \ln n + \gamma', \quad (1)$$

$$H_n \leq \ln n + \gamma' + \frac{1}{2n}, \quad (2)$$

$$H_n \geq \ln n + \gamma' + \frac{1}{2n} - \frac{1}{12n^2}. \quad (3)$$

From Fact 2.1, we can derive the following proposition, which we use later on.

Proposition 2.2. For natural numbers $b > a > 1$,

$$\ln(b) - \ln(a - 1) - \frac{1}{2a - 2} \leq \sum_{x=a}^b \frac{1}{x} \leq \ln(b) - \ln(a - 1).$$

Proof. We only prove the upper bound; the proof for the lower bound is similar. From Inequalities (2) and (3) (for the lower bound proof, use (1) and (2)),

$$\begin{aligned} \sum_{x=a}^b \frac{1}{x} &= \sum_{x=1}^b \frac{1}{x} - \sum_{x=1}^{a-1} \frac{1}{x} \\ &\leq \ln(b) + \gamma' + \frac{1}{2b} - \left(\ln(a - 1) + \gamma' + \frac{1}{2(a - 1)} - \frac{1}{12(a - 1)^2} \right) \\ &\leq \ln(b) - \ln(a - 1), \end{aligned}$$

for $b > a > 1$. □

3 The Model

We consider a system with a single divisible resource, in which agents arrive and depart. The agents have no limit on the amount of the resource they can use: they are always interested in using as much of the resource as the system can give them. We would like to design dynamic mechanisms that satisfy a number of properties:

- **d -disruptiveness**, for some constant d (given as an input to the mechanism).

- **Fairness:** The output allocation has fairness ratio of 1 in *every round* (this property is also known as *proportionality, sharing incentives* or *individual rationality*).
- **Strongly dynamic Pareto optimality (SDPO):** The output allocation is Pareto Optimal in every round.
- **Envy freeness:** The output allocation is envy-free in every round.
- **Population monotonicity.**

Unfortunately, these properties are incompatible even in very simple scenarios:

Theorem 3.1. *There is no mechanism that satisfies d -disruptiveness and fairness, even in the presence of $d + 2$ agents, for any number of donors d .*

Proof. Assume that $d = 1$. The first agent to arrive must be allocated the whole resource, in order for fairness to be satisfied. Once the second agent arrives, the only action that results in fairness ratio of 1 is splitting the resource equally between the agents: $\frac{1}{2}$ to the first and $\frac{1}{2}$ to the second. A problem arises upon the arrival of a third agent: We can only take resources from one agent, thus there will still be an agent with (at least) $\frac{1}{2}$ of the resource. The best fairness ratio we can get is $3/4$, by splitting the donor's share evenly. It is easy to extend this proof to any d . \square

This simple example shows that once we fix the number of donors we allow, we cannot satisfy fairness. The same example can be used to show that d -disruptiveness, SDPO and EF are incompatible. We therefore relax fairness and envy freeness to σ -fairness and ξ -envy freeness:

Definition 3.2 (σ -fairness, σ -fair). *A mechanism satisfies σ -fairness if throughout its execution the fairness ratio of all of its output allocations is at least σ . Such a mechanism is called σ -fair.*

Definition 3.3 ($\xi - EF$). *A mechanism satisfies $\xi - EF$ if throughout the execution of the mechanism the envy ratio of all of its output allocations is at most ξ . Such a mechanism is called ξ -envy-free.*

Before designing allocation mechanisms, we need to establish a benchmark, against which we can compare the performance of our mechanisms. What is the optimal fairness attainable in a d -disruptive system, even if we don't allow departures? We first make the following observation.

Observation 3.4. *The optimal fairness in a bounded-capacity d -disruptive system, where N is the maximum number of agents, is a function of both d and N .*

Proof. In the proof of Theorem 3.1 we showed that achieving fairness ratio of 1 is possible with $d = 1$ donors and 2 agents, but impossible with 1 donor and 3 agents. Notice that for $d = 2$ donors this is not the case anymore: when there are two agents in the system, we can have the allocation vector $(1/2, 1/2)$, and when there are three agents, we can have $(1/3, 1/3, 1/3)$. \square

As we only consider d -disruptive mechanisms in this paper, we omit “ d -disruptive” from all future results.

In Section 4, we show how to use a linear program to compute the benchmark fairness ratios for any given d and N . To help the reader get a better intuition of how the fairness behaves, we give some of these values in Table 1.

N	d					
	1	2	3	5	10	50
2	1	1	1	1	1	1
3	0.857	1	1	1	1	1
5	0.811	0.909	0.952	1	1	1
10	0.774	0.868	0.913	0.957	1	1
100	0.727	0.827	0.874	0.919	0.958	0.995
1,000	0.722	0.823	0.870	0.915	0.954	0.991
100,000	0.721	0.822	0.869	0.914	0.954	0.990

Table 1: Optimal fairness for some values of d and N

4 A Linear Programming Formulation

In this section we study a special case of our problem. Assume we have a bounded-capacity system and that we are limited to arrivals. Then, we can compute an optimal solution using LP 4.1.

Linear Program 4.1.

$$\begin{aligned}
& \text{maximize} && \sigma \\
& \text{subject to} && \sum_{j=1}^t \mathcal{A}_j^t \leq 1, && t \in [1, N] \\
& && \mathcal{A}_1^t \geq \mathcal{A}_2^t \geq \dots \geq \mathcal{A}_t^t \geq \frac{\sigma}{t}, && t \in [1, N] \\
& && \mathcal{A}_j^{t+1} \geq \mathcal{A}_{j+d}^t, && t \in [d+1, N-1], j \in [1, t-d]
\end{aligned}$$

For clarity, we use the same notation in LP 4.1 as in the rest of the paper: \mathcal{A}_j^t is the j^{th} largest allocation when there are t agents in the system, and σ denotes the fairness ratio. A feasible solution to the LP is a set of N feasible allocations \mathcal{A}^i , such that the allocations are consistent with d -disruptiveness. Consider the constraints of LP 4.1: The first line of constraints guarantee feasibility: when there are $t \in [1, N]$ agents in the system, we require the sum of their shares to be at most 1. The second line of constraints guarantees that the fairness ratio is σ : when there are t agents in the system, their shares are all at least $\frac{\sigma}{t}$. (We also use the second set of constraints to keep the allocations sorted.) Finally, we need to make sure that at most d donors are used in each round. For this, we require that for each t , there are at least $t-d$ agents whose allocation did not decrease from \mathcal{A}^{t-1} . This is true if and only if the $t-d$ largest allocations in round $t+1$ are at least as large as the smallest $t-d$ allocations in round t . We can immediately deduce the following lemma.

Lemma 4.2. *The set of feasible solutions of LP 4.1 is exactly the set of all feasible allocations of a bounded-capacity d -disruptive system. Furthermore, an optimal solution to LP 4.1 is a set of feasible allocations of a bounded-capacity d -disruptive system that optimizes fairness.*

The solution to Linear Program 4.1 consists of a set of N allocations and the optimal fairness ratio σ . For example, the solution to the LP for $d=1, N=3$ can be written as follows:

$$\mathcal{A}_1^1 = 1, \mathcal{A}_1^2 = 4/7, \mathcal{A}_2^2 = \mathcal{A}_1^3 = 3/7, \mathcal{A}_2^3 = \mathcal{A}_3^3 = 2/7, \sigma = 6/7.$$

We can also present these allocations in table form (Table 2).

Interestingly, we do not sacrifice fairness by requiring strongly dynamic Pareto optimality.

No. of Agents	Allocation
1	1
2	4/7, 3/7
3	3/7, 2/7, 2/7

Table 2: Optimal allocation for $d = 1$, $N = 3$. The fairness ratios are 1 , $\frac{3/7}{1/2} = \frac{6}{7}$ and $\frac{2/7}{1/3} = \frac{6}{7}$.

Theorem 4.3. *If there is a solution to LP 4.1 that guarantees a fairness ratio of σ , there is also a solution that guarantees a fairness ratio of σ and is SDPO.*

Proof. Assume that there is a solution to LP 4.1 that is not SDPO, and there is no SDPO allocation that achieves the same fairness ratio σ . For any allocation \mathcal{A} , let $\text{last}_{\mathcal{A}}$ be the last round on which \mathcal{A} is not PO (that is, for all $t > \text{last}_{\mathcal{A}}$, $\sum_{j=1}^t \mathcal{A}_j^t = 1$, and for $t = \text{last}_{\mathcal{A}}$, $\sum_{j=1}^t \mathcal{A}_j^t < 1$). Let \mathcal{A}^* be an allocation that minimizes $\text{last}_{\mathcal{A}}$ among all allocations whose fairness ratio is σ . That is,

$$\mathcal{A}^* = \underset{\mathcal{A}: \text{fairness of } \mathcal{A} \text{ is } \sigma}{\text{argmin}} \text{last}_{\mathcal{A}}.$$

Let $t^* = \text{last}_{\mathcal{A}^*}$. At round t^* , there is some unallocated resource, $x = 1 - \sum_{j=1}^{t^*} \mathcal{A}_j^{t^*}$. Let d_1 be some donor in round t^* . Consider the allocation \mathcal{A}' that is identical to \mathcal{A}^* on every round, except for round t^* , where all the agents except d_1 have the same allocation and $\mathcal{A}'_{d_1} = \mathcal{A}^*_{d_1} + x$. It is easy to see that \mathcal{A}' is a feasible allocation that achieves the same fairness as \mathcal{A}^* , but $\text{last}_{\mathcal{A}'} < \text{last}_{\mathcal{A}^*}$, a contradiction. \square

We have the following corollary:

Corollary 4.4. *Rewriting LP 4.1 with equalities in the first line of constraints gives an LP that achieves the same optimal fairness and also guarantees SDPO.*

4.1 Sound Allocations

Although the allocations of LP 4.1 are optimal for arrivals, there is no guarantee that they are also valid for departures. For example, we can get from $(0.5, 0.5)$ to $(0.6, 0.2, 0.2)$ with one donor, but if one of the agents that is allocated 0.2 departs, we can't go back.³

We need to develop another mechanism in order to handle departures. To help us do so, we observe that the allocation of LP 4.1 for N agents has the following interesting property:

The last $d + 1$ agents (the poorest ones) are allocated at least $\frac{\sigma}{N}$. That is $\mathcal{A}_j^N \geq \frac{\sigma}{N}$, for $j \in [N - d, N]$. For $j \in [N - 2d - 1, N - d - 1]$, we have that $\mathcal{A}_j^N \geq \frac{\sigma}{N-1}$; this is because for these values of j it must hold that $\mathcal{A}_j^N \geq \mathcal{A}_{j+d}^{N-1} \geq \frac{\sigma}{N-1}$. Similarly, for $j \in [N - 3d - 2, N - 2d - 2]$ we have $\mathcal{A}_j^N \geq \frac{\sigma}{N-2}$, and so on.

We call this property *soundness*. Denote

$$\psi_i(t, \sigma, d) = \frac{\sigma}{t - \lfloor \frac{t-i}{d+1} \rfloor}.$$

As σ and d are usually obvious from the context, we simply write $\psi_i(t)$ instead of $\psi_i(t, \sigma, d)$. If t is also obvious, we simply write ψ_i .

³ Although this example might seem unnatural, optimal solutions that possess this anomalous quality actually exist: When $d = 1$, $N = 20$, one possible optimal allocation includes the allocation vectors $A^2 = (0.5, 0.5)$, $A^3 = (0.5 + 2\epsilon, 0.25 - \epsilon, 0.25 - \epsilon)$, for some $\epsilon > 0$.

Definition 4.5 (Sound Allocation). *We say that an allocation \mathcal{A}^t , in the presence of t agents, is (t, σ, d) -sound, if for all $i \in [1, t]$,*

$$\mathcal{A}_i^t \geq \psi_i(t).$$

In other words, let $g = \lfloor \frac{t}{d+1} \rfloor$. An allocation is (t, σ, d) -sound if:

- There are $t \bmod (d + 1)$ agents whose allocation is at least $\frac{\sigma}{t-g}$.
- There are $g(d + 1)$ (different) agents, split in groups of size $d + 1$, such that all agents in group j are allocated at least $\frac{\sigma}{t-j}$, for all $j \in [0, g - 1]$.

A (t, σ, d) -sound allocation is (t, σ, d) -tight when the above inequalities are tight:

Definition 4.6 (Tight Allocation). *We say that an allocation \mathcal{A}^t , in the presence of t agents, is (t, σ, d) -tight, if for all $i \in [1, t]$,*

$$\mathcal{A}_i^t = \psi_i(t).$$

We call the difference between a (t, σ, d) -sound allocation and a (t, σ, d) -tight allocation the *surplus*.

Definition 4.7 (Surplus). *The surplus Φ_t of a (t, σ, d) -sound allocation \mathcal{A}^t is defined as follows.*

$$\Phi_t = \sum_{i=1}^t (\mathcal{A}_i^t - \psi_i(t)).$$

Clearly, if \mathcal{A}^t is (t, σ, d) -sound, and $\Phi_t = 0$ the allocation is *tight*. Also notice that $\Phi_t = 0$ does not imply Pareto optimality: σ can be any real number.

Example 4.8. *The allocation $(\frac{\sigma}{5}, \frac{\sigma}{6}, \frac{\sigma}{6}, \frac{\sigma}{6}, \frac{\sigma}{7}, \frac{\sigma}{7}, \frac{\sigma}{7})$ is $(7, \sigma, 2)$ -tight. On the other hand, the allocation $(\frac{\sigma}{5}, \frac{\sigma}{6}, \frac{\sigma}{6}, \frac{\sigma}{6}, \frac{\sigma}{6}, \frac{\sigma}{7}, \frac{\sigma}{7})$ is $(7, \sigma, 2)$ -sound, and the surplus is $\frac{\sigma}{42}$. Notice that if $\sigma = \frac{1}{2}$, neither allocation is Pareto optimal, as an allocation is PO only if its shares sum to 1. If $\sigma = 1$, neither allocation is feasible.*

Example 4.9. *Table 3 presents an optimal set of allocations for $d = 1, N = 5$. The allocation for $t = 5$ is $(5, 30/37, 1)$ -tight. For $t \in [1, 4]$, the allocation \mathcal{A}^t is $(t, 30/37, 1)$ -sound, but not $(t, 30/37, 1)$ -tight. The allocation \mathcal{A}^2 is also $(2, 32/37, 1)$ -sound.*

No. of Agents	Allocation
1	1
2	42/74, 32/74
3	32/74, 22/74, 20/74
4	24/74, 20/74, 15/74, 15/74
5	20/74, 15/74, 15/74, 12/74, 12/74

Table 3: An optimal allocation for $d = 1, N = 5$.

We notice the following property of sound allocations:

Remark 4.10. *A (t, σ, d) -sound allocation has fairness ratio of at least σ : the number of agents is t , and the smallest share is at least $\frac{\sigma}{t}$.*

5 Solving Departures: The Sound Mechanism

We solve the issue of departures by defining a new mechanism, SOUND. The SOUND mechanism takes as input a (N, σ, d) -sound allocation for N agents, and builds $N - 1$ allocations backwards, such that

1. The allocations are valid (we can go “back and forth”).
2. The fairness ratio is non-decreasing.

We explain how SOUND works in Subsection 5.1 and prove its correctness in Subsection 5.2. In Section 6, we show that giving SOUND the final allocation of LP 4.1 as its input allows us to deal with departures without any loss in fairness. In Section 7 we show how we can use SOUND in unbounded systems to ensure all of the allocations have fairness that is arbitrarily close to optimal.

5.1 The Mechanism

Mechanism 5.1. *The SOUND Mechanism receives as input a (N, σ, d) -sound allocation for N agents, and creates allocation vectors \mathcal{A}^t , for $t \in [1, N - 1]$.*

The allocations are created backwards from $t = N - 1$, as follows.

- $t \in [d+1, N - 1]$: *Given \mathcal{A}^{t+1} , we can compute \mathcal{A}^t : Set $\mathcal{A}_{j+d}^t = \mathcal{A}_j^{t+1}$ for all $j \in [1, t-d]$. Take the $d + 1$ smallest shares of \mathcal{A}^{t+1} and create d new shares, $i \in [1, d]$,*

$$\mathcal{A}_i^t = \psi_i(t).$$

Allocate any surplus equally among the new shares.

- $t \leq d$: $\mathcal{A}_i^t = \frac{1}{t}$ for all $i \in [1, t]$.

When there are t agents in the system, $t \in [1, N]$, and the poorest agent departs (see Remark 5.2), we create \mathcal{A}^{t-1} as described above, by taking the leaving agent’s resource and distributing it among the d next poorest to create the d wealthiest allocations of \mathcal{A}^{t-1} . When an agent arrives, we redistribute the d largest shares to recreate the $d + 1$ smallest shares of \mathcal{A}^{t+1} .

Remark 5.2. *We can always assume that the poorest agent leaves: if agent $j \neq t$ leaves, we allocate $\mathcal{A}_j^t - \mathcal{A}_t^t$ to agent t , and now we can treat j as the poorest agent.*

Remark 5.3. *Assuming that SOUND is able to create the allocations as described, it is immediate by the definition of soundness that all allocations produced by it are (t, σ, d) -sound.*

5.2 Correctness

In order to prove that SOUND is correct it suffices to show that: (1) it is possible to create \mathcal{A}^{t+1} from \mathcal{A}^t using at most d donors, (2) it is possible to create \mathcal{A}^t from \mathcal{A}^{t+1} if any agent leaves, without decreasing the share of any of the remaining agents, (3) the fairness ratio for all allocations is at least as much as that of the input allocation.

Proof outline We first show that, for all $t \in [d+1, N]$, the surplus is non-increasing (Claim 5.4). We use Claim 5.4 to prove Proposition 5.6, which proves (1) and (2). Given that SOUND can indeed create all the allocations in its description, (3) is immediate from Remarks 4.10 and 5.3: every allocation that is produced by SOUND has a fairness ratio of σ .

Claim 5.4. *In the execution of SOUND, for $t \in [d + 1, N - 1]$, the surplus satisfies*

$$\Phi_t \geq \Phi_{t+1}.$$

Proof. Recall that in SOUND, the allocations are created backwards; that is, \mathcal{A}^t is created from \mathcal{A}^{t+1} . When creating allocation \mathcal{A}^t from \mathcal{A}^{t+1} , all shares except for the $d + 1$ smallest remain unchanged, and therefore whatever surplus these agents possess in \mathcal{A}^{t+1} , they also possess in \mathcal{A}^t . We take the $d + 1$ poorest agents and combine their shares to create the d richest shares of \mathcal{A}^t . The difference in surplus between \mathcal{A}^t and \mathcal{A}^{t+1} is exactly the difference between the sums of these shares. If the shares of the $d + 1$ poorest agents are exactly $\frac{\sigma}{t+1}$, it suffices to prove that this is sufficient to create the d new shares, to imply $\Phi_t \geq \Phi_{t+1}$. We can assume this is the case, as if it is not, the surplus that the $d + 1$ poorest agents had is also transferred to the d new shares of \mathcal{A}^t , and this cannot affect the difference. In other words,

$$\Phi_t - \Phi_{t+1} = \frac{\sigma(d+1)}{t+1} - \sum_{i=1}^d \psi_i(t).$$

We show that the right hand expression is non-negative.

$$\begin{aligned} \frac{\sigma(d+1)}{t+1} - \sum_{i=1}^d \psi_i(t) &= \frac{\sigma(d+1)}{t+1} - \sum_{i=1}^d \frac{\sigma}{t - \lfloor \frac{t-i}{d+1} \rfloor} \\ &\geq \frac{\sigma(d+1)}{t+1} - \sum_{i=1}^d \frac{\sigma}{t - \frac{t-d}{d+1}} \\ &\geq \frac{\sigma(d+1)}{t+1} - \frac{d(d+1)\sigma}{dt+d} = 0, \end{aligned} \tag{4}$$

where Inequality (4) is due to Claim 5.5. □

Claim 5.5. *For $t \geq d + 1$,*

$$\sum_{i=1}^d \frac{1}{t - \lfloor \frac{t-i}{d+1} \rfloor} \leq \sum_{i=1}^d \frac{1}{t - \frac{t-d}{d+1}}.$$

The proof of Claim 5.5 is deferred to Appendix A.1.

We show (1) and (2) in the following Proposition.

Proposition 5.6. *Let $\mathcal{A}^t, t \in [1, N]$ be the allocations created by SOUND. Then for any $t \in [1, N - 1]$,*

- *It is possible to create \mathcal{A}^{t+1} from \mathcal{A}^t using at most d donors, and*
- *It is possible to create \mathcal{A}^t from \mathcal{A}^{t+1} if any agent leaves, without decreasing the share of any of the remaining agents.*

Proof. The second part of the proposition implies the first: Notice that \mathcal{A}_i^{t+1} is equal to \mathcal{A}_{i+d}^t , for all $i \in [1, t - d]$. Assume that we have shown we can go from \mathcal{A}^{t+1} to \mathcal{A}^t using only the share of the poorest agent in \mathcal{A}^{t+1} : that share has been split among only d agents. We can use the shares of these d agents and go back to \mathcal{A}^{t+1} from \mathcal{A}^t .

For the second part, we notice that as \mathcal{A}^{t+1} is (t, σ, d) -sound, its smallest $d + 1$ shares are at least $\frac{\sigma}{t+1}$ each. One of these agents departs and we use her share to increase the allocation of the remaining d of them, in order to create d shares that are greater than or equal to $\psi_i(t)$ for $i \in [1, d]$. As $\frac{\sigma}{t+1} \leq \psi_i(t)$ for $i \in [1, d]$, it suffices to show that the sum of the $d + 1$ poorest shares of \mathcal{A}^{t+1} suffice to create the d new shares of \mathcal{A}^t , i.e., $\Phi_t \geq \Phi_{t+1}$, which is true from Claim 5.4. □

6 Applying Sound to Bounded Capacity Systems

In the previous section, we saw that SOUND works correctly whenever it is given a sound allocation as input. Since the last allocation of LP 4.1 is (N, σ, d) -sound, we can give it as input to SOUND. We call this mechanism SOUND – LP. In this section we will prove several properties of the fairness and envy ratios of SOUND – LP, to obtain the following.

Theorem 6.1. *For any $d > 0$, there exists a d -disruptive, PM, SDPO allocation mechanism for bounded-capacity systems that guarantees optimal fairness. Furthermore, the mechanism is $\frac{1}{(d+1) \ln(\frac{d+1}{d})}$ -fair and $\frac{3(d+1)^2}{2d^2}$ -envy-free.*

We saw that SOUND is d -disruptive, PM and SDPO in Section 5. The fact that it receives an allocation from LP 4.1 as its input, and the fairness of the other allocations can only be greater shows that SOUND is optimal. Lemmas 6.2 (Subsection 6.1) and 8.4 (Subsection 8.1) complete the proof of the theorem.

6.1 Bounding the Fairness Ratio

The following lemma gives matching upper and lower bounds on the fairness of SOUND.

Lemma 6.2. *SOUND – LP has fairness ratio σ that satisfies*

$$\frac{1}{(d+1) \ln(\frac{d+1}{d}) - \frac{(d+1)^2}{2dN}} \geq \sigma \geq \frac{1}{(d+1) \ln(\frac{d+1}{d})}.$$

We can actually compute the exact value of the fairness of LP 4.1 and SOUND. We require the following observation.

Observation 6.3. *The allocation for N agents of LP 4.1, \mathcal{A}^N , is (N, σ, d) -tight. (Alternatively, $\Phi_N = 0$.)*

Proof. We saw in Claim 5.4 that in the execution of SOUND, for all $t \in [d+1, N-1]$, $\Phi_t \geq \Phi_{t+1}$. It is easy to verify that the claim holds for LP 4.1 as well. The minimal surplus of LP 4.1 is therefore when there are N agents in the system. The surplus for this allocation must be 0, otherwise all allocations have a positive surplus, in contradiction to the optimality of LP 4.1. \square

As \mathcal{A}^N is (N, σ, d) -tight, it is easy to compute the exact fairness ratio as a function of the number of agents and donors:

Lemma 6.4. *The fairness ratio of SOUND is*

$$\sigma = \frac{1}{\sum_{i=1}^N \frac{1}{N - \lfloor \frac{N-i}{d+1} \rfloor}}.$$

Proof. Since \mathcal{A}^N is (N, σ, d) -tight, the i -th agent is allocated $\frac{\sigma}{N - \lfloor \frac{N-i}{d+1} \rfloor}$. The sum of all the shares must be equal to 1, giving

$$\sum_{i=1}^N \frac{\sigma}{N - \lfloor \frac{N-i}{d+1} \rfloor} = 1.$$

Rearranging gives the lemma. \square

Remark 6.5. We can compute \mathcal{A}^N directly, using Lemma 6.4. Somewhat surprisingly, even though we started with LP 4.1 and used it for the proofs, it turns out that we never actually need to write or solve an LP in order to achieve optimal fairness.

Using Lemma 6.4, we can give lower and upper bounds for σ as a function of d and N . The lower bound is given in Corollary 6.6 and the upper bound in Corollary 6.7.

Corollary 6.6 (to Lemma 6.4). *The fairness ratio σ of a bounded-capacity system with d donors and at most N agents satisfies*

$$\sigma \geq \frac{1}{(d+1) \ln \left(\frac{d+1}{d} \right)}.$$

Proof. Let $N = c(d+1) + b$, where c and b are non-negative integers, and $b \leq d$.

$$\sum_{i=1}^N \frac{1}{N - \lfloor \frac{N-i}{d+1} \rfloor} = (d+1) \sum_{i=N+1-c}^N \frac{1}{i} + \frac{b}{N-c} \quad (5)$$

$$\leq (d+1) \ln \left(\frac{N}{N-c} \right) + \frac{b}{N-c} \quad (6)$$

$$\leq (d+1) \ln \left(\frac{d+1}{d} \right), \quad (7)$$

where Equality (5) is due to Claim 6.8, Inequality (6) is due to Proposition 2.2 and Inequality (7) is due to Claim 6.9.

From Lemma 6.4 and the above,

$$\sigma = \frac{1}{\sum_{i=1}^N \frac{1}{N - \lfloor \frac{N-i}{d+1} \rfloor}} \geq \frac{1}{(d+1) \ln \left(\frac{d+1}{d} \right)}.$$

□

Corollary 6.7 (to Lemma 6.4). *The fairness ratio σ of a bounded-capacity system with d donors and at most N agents satisfies*

$$\sigma \leq \frac{1}{(d+1) \ln \left(\frac{d+1}{d} \right) - \frac{(d+1)^2}{2dN}}.$$

The proof of Corollary 6.7 is deferred to Appendix A.2; the proofs of the following claims are deferred to Appendices A.3 and A.4 respectively.

Claim 6.8. *Let $N = c(d+1) + b$, where c and b are non-negative integers, and $b \leq d$.*

$$\sum_{i=b+1}^N \frac{\sigma}{N - \lfloor \frac{N-i}{d+1} \rfloor} = \sigma(d+1) \sum_{i=N+1-c}^N \frac{1}{i}.$$

Claim 6.9. *Let $N = c(d+1) + b$, where c and b are non-negative integers, and $b \leq d$. Then*

$$(d+1) \ln \left(\frac{N}{N-c} \right) + \frac{b}{N-c} \leq (d+1) \ln \left(\frac{d+1}{d} \right).$$

7 Unbounded Systems

In this section we study the resource division problem in general (unbounded) systems, with no known upper bound on the number of agents that can be served. We prove the following theorem, which is the main result of this paper.

Theorem 7.1. *For any $d > 0$ and $\epsilon > 0$ ($d \in \mathbb{N}, \epsilon \in \mathbb{R}$), there exists a d -disruptive, PM and SDPO mechanism with fairness ratio σ that satisfies:*

$$\sigma \geq \frac{1}{(d+1) \ln\left(\frac{d+1}{d}\right)} - \epsilon.$$

Furthermore, its envy ratio is at most $\frac{3(d+1)^2}{2d^2}$.

For ease of presentation we present our mechanism in two parts. First we present a mechanism for the unbounded case, HARMONIC (Subsection 7.1). HARMONIC has a fairness ratio that increases as agents arrive, and that fairness ratio converges to $\frac{1}{(d+1) \ln\left(\frac{d+1}{d}\right)}$ as the number of agents grows.

Then, we show that when there are t agents in the system, the allocation of HARMONIC (with fairness ratio σ_t) is (σ_t, d, t) -sound. Thus, we can give it as input to SOUND and keep the fairness ratio σ_t when there are fewer than t agents present in the system.

In order to guarantee a fairness ratio of at least $\left((d+1) \ln\left(\frac{d+1}{d}\right)\right)^{-1} - \epsilon$, our main mechanism, SOUND-HARMONIC, chooses an $N = \frac{1}{\epsilon}$ (see Claim 7.7), computes the allocation of HARMONIC when there are N agents in the system, and gives it as input to SOUND. Whenever at most N agents are present, we have the allocation SOUND dictates, and the allocation of HARMONIC otherwise. This guarantees that for all allocations, the fairness ratio is at least $\frac{1}{(d+1) \ln\left(\frac{d+1}{d}\right)} - \epsilon$. Figures 2 and 3 in Appendix C show the fairness of SOUND-HARMONIC as a function of the number of agents in the system, for $d = 1$ and 3.

7.1 The HARMONIC Mechanism

The first component of our main mechanism is the HARMONIC mechanism. Set $\alpha_t = d(t-1)$. HARMONIC sets the i^{th} highest allocation when there are t agents in the system to be⁴

$$\mathcal{A}_i^t = \frac{\gamma_t}{\alpha_t + i - 1}, \quad (8)$$

where γ_t is the normalizing factor, i.e.,

$$\gamma_t = \left(\sum_{i=0}^{t-1} \frac{1}{\alpha_t + i} \right)^{-1}.$$

That is, HARMONIC ensures that the allocation vector is

$$\mathcal{A} = \gamma_t \left(\frac{1}{\alpha_t}, \frac{1}{\alpha_t + 1}, \dots, \frac{1}{\alpha_t + t - 1} \right). \quad (9)$$

Figure 1 shows the allocations of HARMONIC when $d = 1$, for different numbers of agents present in the system.

In the rest of this subsection we show how this simple characterization allows us to handle arrivals and departures, and understand how the fairness ratio of the HARMONIC mechanism behaves.

⁴Note that we are counting here from 1 to t ; we will often sum from 0 to $t-1$.

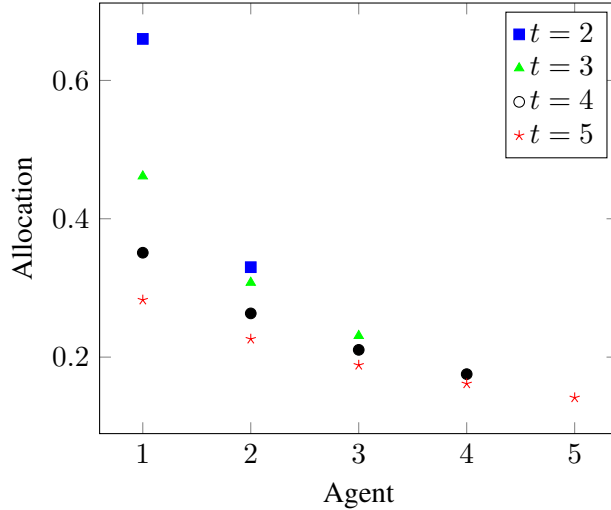


Figure 1: The HARMONIC Mechanism allocations for each agent, when $d = 1$, for different number of agents (t) present in the system.

Arrivals

When a new agent arrives, we take the largest d allocations, combine their resources, create $d + 1$ new allocations, and “top off” the remaining agents (in order to get exactly the allocation of Equation (8)). As before, after each arrival the agents are re-numbered according to their shares (1 has the biggest share).

In order to show that we can get from the t -th allocation to the $(t + 1)$ -th allocation using only d donors we need to show: (1) the shares of the agents that are not donors do not decrease, and (2) the d largest shares suffice to create the $d + 1$ new shares.

(1) is proved by Claim 7.2. (2) follows immediately from (1) and the fact that the shares of both allocations sum up to 1.

Recall that \mathcal{A}_i^t denotes i^{th} highest allocation when there are t agents in the system.

Claim 7.2. *For all $d \geq 1$, $t \geq d + 1$, and $i \in [0, t - 1]$, the allocations of HARMONIC satisfy*

$$\mathcal{A}_i^{t+1} \geq \mathcal{A}_{i+d}^t.$$

The proof of Claim 7.2 is deferred to Appendix A.5.

Departures

when an agent departs, we just reallocate the resource as it should be for $t - 1$ agents: we take the leaving agent’s resource and distribute it among the remaining agents to get the harmonic allocation.

We need to show that (1) the shares of the non-departing agents do not decrease when the poorest agent departs (i.e. we only use her share) and (2) the poorest agent’s share is large enough. As was the case for arrivals, (2) is implied by (1). We prove (1) in Claim 7.3.

Claim 7.3. *For all $d \geq 1$, $t \geq d + 1$, and $0 \leq i \leq t - 1$, the allocations of HARMONIC satisfy*

$$A_i^t \geq A_i^{t+1}.$$

The proof of Claim 7.3 is deferred to Appendix A.6.

Fairness Ratio

Proposition 7.4 and Corollary 7.5 bound the fairness ratio of HARMONIC.

Proposition 7.4. *In HARMONIC with $d \geq 1$ donors, if there are $t \geq 2$ agents in the system, then*

$$\sigma_t \geq \frac{t}{(t-1)(d+1) \left(\ln \frac{(d+1)(t-1)}{d(t-1)-1} \right)}.$$

Proof. The least allocated agent is allocated $\frac{\gamma_t}{\alpha_t + t - 1}$ of the resource. The fairness is therefore

$$\sigma_t = \frac{t\gamma_t}{\alpha_t + t - 1}. \quad (10)$$

Setting $b_t = \alpha_t + t - 1$,

$$\begin{aligned} \sigma_t &= \frac{t\gamma_t}{b_t} \\ &= \frac{t}{b_t} \left(\sum_{x=\alpha_t}^{b_t} \frac{1}{x} \right)^{-1} \\ &\geq \frac{t}{b_t} (\ln(b_t) - \ln(\alpha_t - 1))^{-1} \\ &= \frac{t}{(t-1)(d+1)} \left(\ln \frac{(d+1)(t-1)}{d(t-1)-1} \right)^{-1}, \end{aligned} \quad (11)$$

where Inequality (11) is due to Proposition 2.2. \square

By taking the limit we get the following corollary:

Corollary 7.5. *In the HARMONIC mechanism with $d \geq 1$ donors, if there are t agents in the system, then*

$$\sigma_t \rightarrow \frac{1}{(d+1) \ln \left(\frac{d+1}{d} \right)},$$

as $t \rightarrow \infty$.

7.2 The SOUND-HARMONIC Mechanism

If the allocation of HARMONIC when there are N agents in the system is sound, we can give it as input to the SOUND mechanism and be able to obtain the same fairness ratio when there are fewer than N agents in the system. Claim 7.6 that this is indeed the case by showing that the allocations of HARMONIC when there are N agents in the system conform to the definition of soundness (Definition 4.5). Denote $\psi_i(\sigma_N) = \psi_i(N, \sigma_N, d)$,

Claim 7.6. *Let $\psi(\sigma_N) = (\psi_1(\sigma_N), \psi_2(\sigma_N), \dots, \psi_N(\sigma_N))$ be an allocation, where σ_N is the fairness at round N of HARMONIC and let $B = (B_1, B_2, \dots, B_N)$ be the allocation at round N of HARMONIC. Then for all $i \in [1, N]$,*

$$B_i \geq \psi_i(\sigma_N).$$

Proof. From Equations (9) and (10),

$$\begin{aligned}
\psi_i(\sigma_N) &= \frac{\sigma_N}{N - \lfloor \frac{N-i}{d+1} \rfloor} \\
&= \frac{N\gamma_N}{(d+1)(N-1)(N - \lfloor \frac{N-i}{d+1} \rfloor)} \\
&\leq \frac{N\gamma_N}{dN^2 - Nd + iN - i} \\
&\leq \frac{\gamma_N}{dN - d + i - 1} = B_i,
\end{aligned}$$

as $i \leq N$. □

Corollary 7.5 shows that $\sigma_N \xrightarrow{N \rightarrow \infty} \frac{1}{(d+1)\ln(\frac{d+1}{d})}$. The following claim, whose correctness is easy to verify, shows that in order to guarantee a fairness ratio of $\frac{1}{(d+1)\ln(\frac{d+1}{d})} - \epsilon$, we can choose $N = 1/\epsilon$, compute the allocation of HARMONIC when there are N agents in the system, and give this allocation to SOUND.

Claim 7.7.

$$\frac{t}{(t-1)(d+1) \left(\ln \frac{(d+1)(t-1)}{d(t-1)-1} \right)} \geq \frac{1}{(d+1)\ln\left(\frac{d+1}{d}\right)} - \frac{1}{t-1}.$$

To summarize, we have shown that in order to guarantee a fairness ratio of at least $\sigma = \left((d+1)\ln\left(\frac{d+1}{d}\right) \right)^{-1} - \epsilon$, we let $N = \frac{1}{\epsilon}$, compute the allocation of HARMONIC when there are N agents in the system, and give this as input to SOUND. When there are N agents in the system, the allocation of HARMONIC is (σ, d, N) -sound, and so we can give it as input to SOUND to guarantee the fairness ratio is σ when there are fewer than N agents present in the system. As the fairness of HARMONIC is increasing in the number of agents present in the system, we are guaranteed that when there are more than N agents in the system, the fairness is at least σ as well.

8 Bounding the Envy

In this section, we give an upper bound on the envy of the SOUND- LP and SOUND-HARMONIC mechanisms. In both scenarios, t represents the number of agents currently present in the system. For SOUND- LP , N represents the cap on the number of agents in the system, and for SOUND-HARMONIC, it represents the number of agents in the allocation that is given to SOUND.

For simplicity, throughout this section, we assume that N and t are multiples of $(d+1)$. In particular, $N = a(d+1), t = b(d+1)$. It is not too difficult to verify the proofs hold for all values of N and t .

8.1 Bounding the Envy of SOUND- LP

Before bounding the envy of SOUND- LP , we prove two propositions about the surplus.

Proposition 8.1. *The surplus of SOUND- LP , in the presence of t agents satisfies $\Phi_t \leq \frac{\sigma(d+1)^2}{2td}$.*

Proof. We assume that $N = a(d+1), t = b(d+1), a \geq b > 0$,

We saw in Claim 5.4 that

$$\Phi_k - \Phi_{k+1} = \frac{\sigma(d+1)}{k+1} - \sum_{i=1}^d \psi_i(k).$$

From Observation 6.3, $\Phi_N = 0$, therefore

$$\begin{aligned}\Phi_t &= \Phi_t - \Phi_N = \sum_{k=t}^{N-1} \Phi_k - \Phi_{k+1} \\ &= \sum_{k=t}^{N-1} \left(\frac{\sigma(d+1)}{k+1} - \sum_{i=1}^d \psi_i(k) \right) \\ &= \sum_{k=t}^{N-1} \frac{\sigma(d+1)}{k+1} - \sum_{j=1}^{d(N-t)} \psi_j(t)\end{aligned}\tag{12}$$

$$= \sigma(d+1) \left(\sum_{k=b(d+1)+1}^{a(d+1)} \frac{1}{k} - \sum_{i=bd+1}^{ad} \frac{1}{i} \right)\tag{13}$$

$$\begin{aligned}&\leq \sigma(d+1) \left(\ln(a(d+1)) - \ln(b(d+1)) + \ln(bd) - \ln(ad) + \frac{1}{2bd} \right) \\ &= \frac{\sigma(d+1)}{2bd},\end{aligned}\tag{14}$$

where Equality (12) is because $\psi_i(k+1) = \psi_{i+d}(k)$, Equality (13) is due to Claim 8.2 and Inequality (14) is due to Proposition 2.2. \square

Claim 8.2. *Let $N = a(d+1)$, $t = b(d+1)$, $a \geq b > 0$. Then*

$$\sum_{j=1}^{d(N-t)} \psi_j(t) = \sigma(d+1) \sum_{i=bd+1}^{ad} \frac{1}{i}.$$

The proof of Claim 8.2 is deferred to Appendix A.7.

Proposition 8.3. *Assume that an allocation \mathcal{A}^t of SOUND has a surplus of Φ_t . Then for all $i \in [1, t]$, agent i 's share is at most*

$$\psi_i(t) + \frac{\Phi_t - \Phi_n}{d}.$$

Proof. Any new surplus produced by SOUND is split equally among d agents in each round. Therefore no single agent can receive more than a $1/d$ share of the total new surplus created between round n and t . \square

Lemma 8.4. *Let \mathcal{A}^t be the allocation of SOUND – LP in the presence of t agents. Then ξ , the envy of \mathcal{A}^t satisfies*

$$\xi \leq \frac{3(d+1)^2}{2d^2}.$$

Proof. From Proposition 8.1, $\Phi_t \leq \frac{\sigma(d+1)}{2td}$. Recall that the envy ratio is at most the ratio between the biggest share over the smallest share. The smallest share is simply ψ_t . From Observation 6.3, $\Phi_N = 0$, and so, from Proposition 8.3 we conclude that the biggest share is $\psi_1 + \Phi_t/d$. For $t \leq d$ there is no envy. Therefore the envy ratio is at most:

$$\xi \leq \frac{\psi_1 + \Phi_t/d}{\psi_t} \leq \frac{\frac{\sigma}{td} + \frac{\sigma(d+1)}{2td^2}}{\frac{\sigma}{t(d+1)}} = \frac{(3d+1)(d+1)}{2d^2} \leq \frac{3(d+1)^2}{2d^2}.$$

\square

8.2 Bounding the Envy of SOUND-HARMONIC

In this subsection we bound the envy ratio of the SOUND-HARMONIC mechanism. It is easy to obtain a tight bound for the envy of the HARMONIC part of the mechanism, and we do so. However, we do not attempt to provide a tight bound on the envy of the SOUND part.

Theorem 8.5. *The envy ratio ξ of the SOUND-HARMONIC mechanism with parameter N is at most $\frac{3(d+1)^2}{2d^2}$ during the SOUND part of the mechanism and at most $\frac{d+1}{d}$ during the HARMONIC part of the mechanism.*

Proof. When there are more than N agents, we are in the HARMONIC part of our mechanism. The wealthiest agent has $\frac{\gamma_t}{\alpha_t} = \frac{\gamma_t}{d(t-1)}$ while the poorest agent has $\frac{\gamma_t}{\alpha_{t+t-1}} = \frac{\gamma_t}{(d+1)(t-1)}$. Thus the envy is exactly $\frac{d+1}{d}$.

When there are $t \leq N$ agents, Lemma 8.6 holds. \square

Lemma 8.6. *Let \mathcal{A}^t be the allocation of SOUND-HARMONIC in the presence of $t \leq N$ agents. Then ξ , the envy of \mathcal{A}^t satisfies*

$$\xi \leq \frac{3(d+1)^2}{2d^2}.$$

The proof is virtually identical to the proof of Lemma 8.4 and is omitted. We note that instead of Observation 6.3, we require Claim 8.7, and instead of Proposition 8.1, we require Proposition 8.8.

Claim 8.7. *For the N^{th} allocation of HARMONIC, it holds that $\Phi_N = O(\frac{1}{N^2})$.*

Proof. Assume, for simplicity, that $N = a(d+1), t = b(d+1)$.

$$B_t = \frac{\gamma_N}{(d+1)(cd+b-1)},$$

and

$$\psi_t = \frac{c(d+1)\gamma_N}{(d+1)(cd+c-1)(cd+b)} = \frac{c\gamma_N}{(cd+c-1)(cd+b)}.$$

It is easy to verify that

$$\begin{aligned} B_t - \psi_t &= \frac{\gamma_N(c-b)}{(d+1)(cd+c-1)(cd+b-1)(cd+b)} \\ &\leq \frac{(c-b)}{(cd+c-1)(cd+b-1)(cd+b)} = O\left(\frac{1}{N^2}\right). \end{aligned}$$

\square

Proposition 8.8. *The surplus of the SOUND mechanism, with input the N^{th} allocation of HARMONIC, in the presence of t agents, satisfies*

$$\Phi_t \leq \frac{\sigma(d+1)^2}{2td} + O\left(\frac{1}{N^2}\right).$$

The proof of Proposition 8.8 is identical to the proof of Proposition 8.1, except that instead of $\Phi_N = 0$, we set Φ_N to be $O(\frac{1}{N^2})$.

Acknowledgements

The authors would like to thank the anonymous reviewers for their useful feedback.

References

- [1] ALON, N. 1987. Splitting necklaces. *Advances in Mathematics* 63, 3, 247–253.
- [2] BHATTACHARYA, A. A., CULLER, D., FRIEDMAN, E., GHODSI, A., SHENKER, S., AND STOICA, I. 2013. Hierarchical scheduling for diverse datacenter workloads. In *Proceedings of the 4th Annual Symposium on Cloud Computing*. SOCC '13. ACM, New York, NY, USA, 4:1–4:15.
- [3] BRAMS, S. J. AND TAYLOR, A. D. 1995. An envy-free cake division protocol. *American Mathematical Monthly*, 9–18.
- [4] BRAMS, S. J. AND TAYLOR, A. D. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- [5] DOLEV, D., FEITELSON, D. G., HALPERN, J. Y., KUPFERMAN, R., AND LINIAL, N. 2012. No justified complaints: On fair sharing of multiple resources. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. ITCS '12. ACM, New York, NY, USA, 68–75.
- [6] DUBINS, L. E. AND SPANIER, E. H. 1961. How to cut a cake fairly. *American mathematical monthly*, 1–17.
- [7] FELDMAN, M., FIAT, A., LEONARDI, S., AND SANKOWSKI, P. 2012. Revenue maximizing envy-free multi-unit auctions with budgets. In *ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012*. 532–549.
- [8] FRIEDMAN, E., GHODSI, A., AND PSOMAS, C.-A. 2014. Strategyproof allocation of discrete jobs on multiple machines. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 529–546.
- [9] GHODSI, A., ZAHARIA, M., HINDMAN, B., KONWINSKI, A., SHENKER, S., AND STOICA, I. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*. NSDI'11. USENIX Association, Berkeley, CA, USA, 24–24.
- [10] GUTMAN, A. AND NISAN, N. 2012. Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 719–728.
- [11] KALAI, E. AND SMORODINSKY, M. 1975. Other solutions to nash's bargaining problem. *Econometrica: Journal of the Econometric Society*, 513–518.
- [12] KASH, I., PROCACCIA, A. D., AND SHAH, N. 2013. No agent left behind: Dynamic fair division of multiple resources. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 351–358.
- [13] MILOJIČIĆ, D. S., DOUGLIS, F., PAINDAVEINE, Y., WHEELER, R., AND ZHOU, S. 2000. Process migration. *ACM Computing Surveys (CSUR)* 32, 3, 241–299.
- [14] PARKES, D. C., PROCACCIA, A. D., AND SHAH, N. 2012. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC '12. ACM, New York, NY, USA, 808–825.

- [15] PAZNER, E. A. AND SCHMEIDLER, D. 1978. Egalitarian equivalent allocations: A new concept of economic equity. *The Quarterly Journal of Economics*, 671–687.
- [16] STEINHAUS, H. 1948. The problem of fair division. *Econometrica* 16, 1.
- [17] STROMQUIST, W. 1980. How to cut a cake fairly. *American Mathematical Monthly*, 640–644.
- [18] WALSH, T. 2011. Online cake cutting. In *Algorithmic Decision Theory*. Springer, 292–305.
- [19] WANG, W., LI, B., AND LIANG, B. 2014. Dominant resource fairness in cloud computing systems with heterogeneous servers. In *INFOCOM, 2014 Proceedings IEEE*. IEEE, 583–591.

A Missing Proofs

A.1 Proof of Claim 5.5

Claim 5.5. For $t \geq d + 1$,

$$\sum_{i=1}^d \frac{1}{t - \lfloor \frac{t-i}{d+1} \rfloor} \leq \sum_{i=1}^d \frac{1}{t - \frac{t-d}{d+1}}.$$

Proof. Notice that $\lfloor \frac{t-i}{d+1} \rfloor$ can only take two values for $i \in [1, d]$. Let $p \in [1, d]$ be such that $(t-p) \bmod (d+1) = 0$. The value $\frac{t-p}{d+1}$ appears p times, and the value $\frac{t-p}{d+1} - 1$ appears $d-p$ times.

If $d = 1$ or there does not exist such a p then the inequality holds immediately. Otherwise, such a p exists and then,

$$\begin{aligned} \sum_{i=1}^d \frac{1}{t - \lfloor \frac{t-i}{d+1} \rfloor} &= \frac{p}{t - \frac{t-p}{d+1}} + \frac{d-p}{t - \frac{t-p}{d+1} + 1} \\ &= (d+1) \left(\frac{p}{dt+p} + \frac{d-p}{dt+p+d+1} \right) \\ &\leq \frac{d+1}{t+1} \\ &= \sum_{i=1}^d \frac{1}{t - \frac{t-d}{d+1}}. \end{aligned} \tag{15}$$

where Inequality (15) is because for $p = d$,

$$\frac{p}{dt+p} + \frac{d-p}{dt+p+d+1} = \frac{1}{t+1},$$

and differentiating w.r.t. p , it is easy to verify that $\frac{p}{dt+p} + \frac{d-p}{dt+p+d+1}$ is monotonically increasing in the interval $[1, d]$. \square

A.2 Proof of Corollary 6.7

Corollary 6.7 (to Lemma 6.4). *The fairness ratio σ of a bounded-capacity system with d donors and at most N agents satisfies*

$$\sigma \leq \frac{1}{(d+1) \ln \left(\frac{d+1}{d} \right) - \frac{(d+1)^2}{2dN}}.$$

For simplicity, we prove Corollary 6.7 only for when N is a multiple of $(d + 1)$.

Proof. Assume $N = c(d + 1)$, $c > 0$, then from Claim 6.8,

$$\sigma = \frac{1}{\sum_{i=1}^N \frac{1}{N - \lfloor \frac{N-i}{d+1} \rfloor}} = \frac{1}{(d+1) \sum_{i=N+1-c}^N \frac{1}{i}}.$$

We can use Proposition 2.2 to obtain a lower bound for $\sum_{i=N+1-c}^N \frac{1}{i}$:

$$\sum_{i=N+1-c}^N \frac{1}{i} \geq \ln \left(\frac{d+1}{d} \right) - \frac{d+1}{2dN},$$

which gives the corollary. \square

We note that there is another way to upper bound the fairness ratio, by taking the dual of LP 4.1 and providing a feasible solution for it. This approach gives a slightly better upper bound of $\frac{1}{(d+1) \ln \left(\frac{d+1}{d} \right)} + \frac{2}{N}$. As it is a much more cumbersome proof, and the gain from using it is minimal, we omit it from the paper.

A.3 Proof of Claim 6.8

Claim 6.8. Let $N = c(d + 1) + b$, where c and b are non-negative integers, and $b \leq d$.

$$\sum_{i=b+1}^N \frac{\sigma}{N - \lfloor \frac{N-i}{d+1} \rfloor} = \sigma(d+1) \sum_{i=N+1-c}^N \frac{1}{i}.$$

Proof. There are $d + 1$ values of i for which $\lfloor \frac{N-i}{d+1} \rfloor$ gives the same value. Therefore, there are c distinct, consecutive values of $\lfloor \frac{N-i}{d+1} \rfloor$ in the sum, with each one appearing $d + 1$ times: the first one is $c - 1$ and the last one is 0. Therefore,

$$\sum_{i=b+1}^{c(d+1)+b} \frac{\sigma}{N - \lfloor \frac{N-i}{d+1} \rfloor} = \sigma(d+1) \sum_{j=0}^{c-1} \frac{1}{N-j}.$$

Substituting $i = N - j$ completes the claim. \square

A.4 Proof of Claim 6.9

Claim 6.9. Let $N = c(d + 1) + b$, where c and b are non-negative integers, and $b \leq d$. Then

$$(d+1) \ln \left(\frac{N}{N-c} \right) + \frac{b}{N-c} \leq (d+1) \ln \left(\frac{d+1}{d} \right).$$

Proof. We show that $e^{((d+1) \ln \frac{N}{N-c} + \frac{b}{N-c})} \leq \left(\frac{d+1}{d} \right)^{d+1}$, by showing that

$$e^{\frac{b}{N-c}} \leq \frac{\left(\frac{d+1}{d} \right)^{d+1}}{e^{(d+1) \ln \frac{N}{N-c}}} = \frac{\left(\frac{d+1}{d} \right)^{d+1}}{\left(\frac{N}{N-c} \right)^{d+1}} = \left(\frac{(d+1)(N-c)}{dN} \right)^{d+1}.$$

Taking the $(d+1)^{th}$ root, we have

$$\begin{aligned}
e^{\frac{b}{(cd+b)(d+1)}} &\leq \frac{1}{1 - \frac{b}{(cd+b)(d+1)}} \\
&= \frac{(cd+b)(d+1)}{d(c(d+1)+b)} \\
&= \frac{(d+1)(N-c)}{dN}.
\end{aligned} \tag{16}$$

Inequality (16) is because $e^x \leq \frac{1}{1-x}$ for $x < 1$. □

A.5 Proof of Claim 7.2

Claim 7.2. For all $d \geq 1$, $t \geq d+1$, and $i \in [0, t-1]$, the allocations of HARMONIC satisfy

$$\mathcal{A}_i^{t+1} \geq \mathcal{A}_{i+d}^t.$$

Proof. From Equation (8), $\mathcal{A}_i^t = \frac{\gamma^t}{d(t-1)+i-1}$. We therefore need to show, for $d \geq 1$, $t \geq d+1$, and $0 \leq i \leq t-1$

$$\frac{\gamma_{t+1}}{dt+i} \geq \frac{\gamma_t}{d(t-1)+i+d}.$$

The denominators are identical, thus it suffices to prove that $\gamma_{t+1} \geq \gamma_t$. As $\gamma_t = \left(\sum_{i=0}^{t-1} \frac{1}{d(t-1)+i} \right)^{-1}$, this is the same as showing that

$$\sum_{i=0}^{t-1} \frac{1}{d(t-1)+i} \geq \sum_{i=0}^t \frac{1}{dt+i}.$$

Since the sums have overlapping elements this holds if and only if

$$\sum_{i=-d}^{-1} \frac{1}{dt+i} \geq \sum_{i=t-d}^t \frac{1}{dt+i}$$

We rearrange so that the expression becomes

$$\begin{aligned}
\sum_{i=1}^d \left(\frac{1}{dt-i} - \frac{1}{dt+t-i} \right) &\geq \frac{1}{dt+t}. \\
\sum_{i=1}^d \left(\frac{1}{dt-i} - \frac{1}{dt+t-i} \right) &= \sum_{i=1}^d \left(\frac{t}{(dt-i)(dt+t-i)} \right) \\
&\geq d \left(\frac{t}{(dt-1)(dt+t-1)} \right) \\
&= \frac{dt}{(dt)^2 - 2dt + dt^2 - t + 1} \\
&= \frac{1}{dt+t-2-1/d+1/dt} \\
&\geq \frac{1}{dt+t},
\end{aligned}$$

where the last inequality holds for all $d, t \geq 1$. □

A.6 Proof of Claim 7.3

Claim 7.3. For all $d \geq 1$, $t \geq d + 1$, and $0 \leq i \leq t - 1$, the allocations of HARMONIC satisfy

$$A_i^t \geq A_i^{t+1}.$$

Proof. In other words, as $A_i^t = \frac{\gamma_t}{d(t-1)+i-1}$ (Equation (8)), we need to prove that for all $d \geq 1$, $k \geq d + 1$, and $0 \leq j \leq k - 1$,

$$\frac{\gamma_{k+1}}{dk+j} \leq \frac{\gamma_k}{d(k-1)+j}.$$

As $\gamma_k = \left(\sum_{i=0}^{k-1} \frac{1}{d(k-1)+i} \right)^{-1}$ We need to show that

$$\gamma_k \geq \left(1 - \frac{d}{dk+j} \right) \gamma_{k+1}$$

$$\gamma_k \geq \gamma_{k+1} - \frac{d}{dk+j} \gamma_{k+1}$$

We've seen in the proof of Claim 7.2 that this is equivalent to:

$$\frac{1}{dk+k} \geq \sum_{i=0}^d \left(\frac{1}{dk-i} - \frac{1}{dk+k-i} \right) - \frac{d}{dk+j} \gamma_{k+1}$$

$$\frac{1}{dk+k} + \frac{d}{dk+j} \gamma_{k+1} \geq \sum_{i=0}^d \left(\frac{1}{dk-i} - \frac{1}{dk+k-i} \right)$$

We've seen that $\gamma_{k+1} \geq \gamma_k$, and since $j \leq k - 1$ we have that $\frac{d}{dk+j} \geq \frac{d}{dk+k}$. Thus, we only need to show that

$$\frac{1}{dk+k} + \frac{d}{dk+k} \gamma_k \geq \sum_{i=0}^d \left(\frac{1}{dk-i} - \frac{1}{dk+k-i} \right)$$

We rewrite this as Claim A.1. □

Claim A.1. The following holds for all $d \geq 1$, $k \geq d + 2$, and $0 \leq i \leq k - 1$.

$$\left(\frac{1}{dk+k} \right) \left(1 + \sum_{i=0}^{k-1} \frac{d}{d(k-1)+i} \right) \geq \sum_{i=0}^d \left(\frac{1}{dk-i} - \frac{1}{dk+k-i} \right)$$

Proof. We only prove the claim for $k \geq 7$, for simplicity. It is easy to verify that the claim holds for all $k \geq d + 2$.

$$\begin{aligned}
\sum_{i=1}^d \left(\frac{1}{dk-i} - \frac{1}{dk+k-i} \right) &= \sum_{i=1}^d \left(\frac{k}{(dk-i)(dk+k-i)} \right) \\
&\leq \frac{dk}{(dk-d)(dk+k-d)} \\
&\leq \frac{dk}{(dk)^2 - 2d^2k + dk^2 - dk + d^2} \\
&\leq \frac{1}{dk+k-2d-1} \\
&\leq \frac{1}{(d+1)(k-2)} \\
&= \frac{1}{k(d+1)} \left(1 + \frac{2}{k-2} \right) \\
&\leq \left(\frac{1}{dk+k} \right) \left(1 + \frac{kd}{(d+1)(k-1)} \right) \\
&\leq \left(\frac{1}{dk+k} \right) \left(1 + \sum_{i=0}^{k-1} \frac{d}{d(k-1)+i} \right),
\end{aligned} \tag{17}$$

for all $d \geq 1, k \geq 7$. (we require $k \geq 7$ for Inequality (17)). \square

A.7 Proof of Claim 8.2

Claim 8.2. Let $N = a(d+1), t = b(d+1), a \geq b > 0$. Then

$$\sum_{j=1}^{d(N-t)} \psi_j(t) = \sigma(d+1) \sum_{i=bd+1}^{ad} \frac{1}{i}.$$

Proof. Let $c = a - b$. There are $d+1$ values of j for which $\lfloor \frac{t-j}{d+1} \rfloor$ gives the same value. Therefore, there are cd distinct, consecutive values in the sum, with each one appearing $d+1$ times, and the first one being $\frac{\sigma}{t-b+1}$. Therefore,

$$\begin{aligned}
\sum_{j=1}^{dc(d+1)} \psi_j(t) &= \sigma(d+1) \sum_{j=1}^{cd} \frac{1}{t-b+j} \\
&= \sigma(d+1) \sum_{i=t-b+1}^{t-b+cd} \frac{1}{i} \\
&= \sigma(d+1) \sum_{i=bd+1}^{ad} \frac{1}{i},
\end{aligned}$$

substituting $i = t - b + j$. \square

B Optimizing Envy

We can use a variation of LP 4.1 to find the allocations that give optimal envy ratio. The catch is that envy cannot be a variable; the constraints will not be linear any more. The following variation works: fix some envy ϵ and check if the following set of constraints is feasible:

Linear Program B.1.

$$\begin{aligned} \sum_{j=1}^t a_{t,j} &= 1, & t \in [1, N] \\ \frac{a_{t,i}}{a_{t,j}} &\geq \epsilon, & t \in [1, N], i, j \in [1, t] \\ a_{t+1,j} &\geq a_{t,j+d}, & t \in [d, N-1], j \in [1, t-d] \end{aligned}$$

We can use binary search over ϵ and compute the optimal envy for a given N in polynomial time.

C Fairness of the SOUND-HARMONIC Mechanism




Figure 2: Fairness ratio for $d = 1$ and $N = 10$

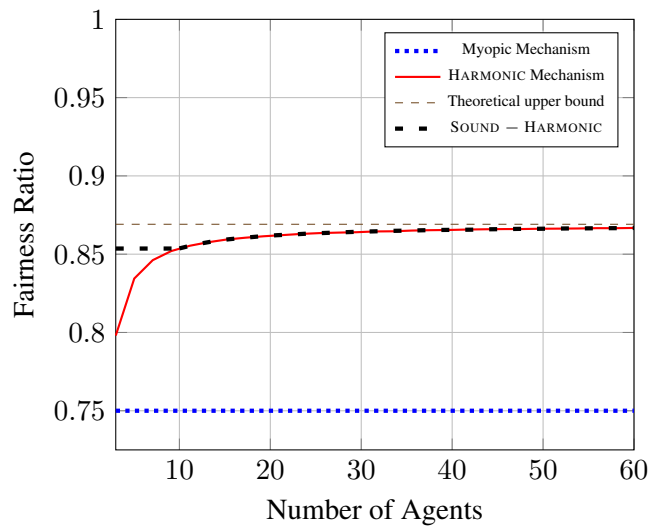


Figure 3: Fairness ratio for $d = 3$ and $N = 10$