

Resource Sharing Through Multi-Round Matchings

Yohai Trabelsi¹, Abhijin Adiga², Sarit Kraus¹, S. S. Ravi^{2,3}, Daniel J. Rosenkrantz^{2,3}

¹ Department of Computer Science, Bar-Ilan University, Ramat Gan, Israel

² Biocomplexity Institute and Initiative, Univ. of Virginia, Charlottesville, VA, USA

³ Dept. of Computer Science, University at Albany – SUNY, Albany, NY, USA

yohai.trabelsi@gmail.com, abhijin@virginia.edu, sarit@cs.biu.ac.il, ssravi0@gmail.com, drosenkrantz@gmail.com

Abstract

Applications such as employees sharing office spaces over a workweek can be modeled as problems where agents are matched to resources over multiple rounds. Agents' requirements limit the set of compatible resources and the rounds in which they want to be matched. Viewing such an application as a multi-round matching problem on a bipartite compatibility graph between agents and resources, we show that a solution (i.e., a set of matchings, with one matching per round) can be found efficiently if one exists. To cope with situations where a solution does not exist, we consider two extensions. In the first extension, a benefit function is defined for each agent and the objective is to find a multi-round matching to maximize the total benefit. For a general class of benefit functions satisfying certain properties (including diminishing returns), we show that this multi-round matching problem is efficiently solvable. This class includes utilitarian and Rawlsian welfare functions. For another benefit function, we show that the maximization problem is NP-hard. In the second extension, the objective is to generate advice to each agent (i.e., a subset of requirements to be relaxed) subject to a budget constraint so that the agent can be matched. We show that this budget-constrained advice generation problem is NP-hard. For this problem, we develop an integer linear programming formulation as well as a heuristic based on local search. We experimentally evaluate our algorithms on synthetic networks and apply them to two real-world situations: shared office spaces and matching courses to classrooms.

1 Introduction

We consider resource allocation problems that arise in practical applications such as hot desking or shared work spaces (Varone and Beffa 2019; Cai and Khan 2010), classroom scheduling (Phillips et al. 2015), matching customers with taxicabs (Karamanis, Anastasiadis, and Angeloudis 2021; Kucharski and Cats 2020), and matching agricultural equipment with farms (Gilbert 2018; Rakhra and Singh 2020). In such scenarios, many agents (individuals, cohorts, farms, or in general, entities) are competing for a limited number of time-shared resources. In our formulation, an agent can be matched to at most one resource in any time slot, but might want to be matched in more than one time slot. Agents may have some *restrictions* that limit the set

of resources to which they can be matched or possible time slots in which they can be matched. Any resource whose specifications do not meet an agent's restrictions is *incompatible* with that agent.

In light of the COVID-19 pandemic, such resource allocation problems have become as important as ever. For example, social-distancing requirements (such as maintaining a six-foot separation between individuals) led to a dramatic decrease in the number of individuals who can occupy an enclosed space, whether it is a classroom (District Management Group 2020; Enriching Students 2020), workspace (Parker 2020) or visitation room (Wong 2021).

We model this resource allocation problem as a k -round matching problem on a bipartite graph, where the two node sets represent the agents and resources respectively. We assume that the rounds are numbered 1 through k , and each round matches a subset of agents with a subset of resources. Each agent specifies the set of permissible rounds in which it can participate and the desired number of rounds (or matchings) in which it needs to be assigned a resource. Consider for example a classroom scheduling for a workweek ($k = 5$). Each lecturer who wants to schedule class sessions for her courses specifies the number of sessions she would like to schedule and the possible weekdays. There may also be additional requirements for classrooms (e.g., room size, location, computer lab). As a result, some agent-resource pairs become incompatible. The PRINCIPAL's objective is to find a set of at most k matchings satisfying all the requirements, if one exists. In the classroom scheduling example, the existence of such a set means that there exists a solution where each lecturer receives the desired number of sessions on the desired days.

We refer to this as the **multi-round matching problem** (MRM). We consider two additional problem formulations to cope with the situation when a solution satisfying all the requirements does not exist.

In the first formulation, an agent receives a benefit (or reward) that depends on the number of rounds where it is matched, and the objective is to find a solution that maximizes the sum of the benefits over all agents. We refer to this as the **multi-round matching for total benefit maximization** problem (abbreviated as MAXTB-MRM). When applying this formulation to the classroom scheduling example, it is possible to find a solution in which the number of

assigned sessions is maximized (by maximizing the corresponding utility function). Alternatively, we can also find a solution in which the minimal assignment ratio for a lecturer (i.e., the number of assigned sessions divided by the number requested) is maximized.

In the second formulation, the objective is to generate suitable advice to each agent to relax its resource requirements so that it becomes compatible with previously incompatible resources. However, the agent incurs a cost to relax its restrictions (e.g., social distancing requirements not met, increase in travel time). So, the suggested advice must satisfy the budget constraints that are chosen by the agents. In other words, the goal of this problem is to suggest relaxations to agents' requirements, subject to budget constraints, so that in the new bipartite graph (obtained after relaxing the requirements), all the agents' requirements are satisfied. In scheduling classrooms, for example, the PRINCIPAL may require that some lecturers waive a subset of their restrictions (e.g., having a far-away room) so that they can get the number of sessions requested. We refer to this to as the **advice generation for multi-round matching** problem (abbreviated as AG-MRM).

Summary of contributions:

(a) An efficient algorithm for MAXTB-MRM. For a general class of benefit functions that satisfy certain properties including monotonicity (i.e., the function is non-decreasing with increase in the number of rounds) and diminishing returns (i.e., increase in the value of the function becomes smaller as the number of rounds is increased), we show that the MAXTB-MRM problem can be solved efficiently by a reduction to the maximum weighted matching problem. A simple example of such a function (where each agent receives a benefit of 1 for each round in which it is matched) represents a **utilitarian social welfare** function (Viner 1949). Our efficient algorithm for this problem yields as a corollary an efficient algorithm for the MRM problem mentioned above. Our algorithm can also be used for a more complex benefit function that models a **Rawlsian social welfare function** (Rawls 1999; Stark 2020), where the goal is to maximize the minimum satisfaction ratio (i.e., the ratio of the number of rounds assigned to an agent to the requested number of rounds) over all the agents.

(b) Maximizing the number of satisfied agents. Given a multi-round matching, we say that an agent is **satisfied** if the matching satisfies all the requirements of the agent. The objective of finding a multi-round matching that satisfies the largest number of agents can be modeled as the problem of maximizing the total benefit by specifying a simple benefit function for each agent. However, such a benefit function *doesn't* have the diminishing returns property. We show that this optimization problem (denoted by MAXSA-MRM) is **NP-hard**, even when each agent needs to be matched in at most three rounds. This result points out the importance of the diminishing returns property in obtaining an efficient algorithm for maximizing the total benefit.

(c) Advice generation. We show that AG-MRM is **NP-hard** even when there is only one agent who needs to be matched in two or more rounds. Recall that the AG-MRM problem

Problem	Results
MAXTB-MRM	Efficient algorithm for a general class of benefit functions. An efficient algorithm for the MRM problem is a corollary.
MAXSA-MRM	NP-hard (reduction from Minimum Vertex Cover for cubic graphs)
AG-MRM AG-MAXSA-MRM	NP-hard (reduction from Minimum Set Cover). ILP and a local search heuristic for AG-MAXSA-MRM

Table 1: Overview of problems and main results

requires that each agent must be satisfied (i.e., assigned the desired number of rounds) in the new compatibility graph (obtained by relaxing the suggested restrictions). It is interesting to note that the hardness of AG-MRM is due to the advice generation part and not the computation of matching. (Without the advice generation part, the AG-MRM problem corresponds to MRM on a given compatibility graph, which is efficiently solvable as mentioned above.) The hardness of AG-MRM directly implies the hardness of the problem (denoted by AG-MAXSA-MRM) where the generated advice must lead to a matching that satisfies the maximum number of agents. We present two solution approaches for the AG-MAXSA-MRM problem: (i) an integer linear program (ILP-AG-MAXSA-MRM) to find an optimal solution and (ii) a pruned local search heuristic (PS-AG-MAXSA-MRM) that uses our algorithm for MRM to generate solutions of good quality.

(d) Experimental results. We present a back-to-the-lab desk-sharing study that has been conducted in a lab at Bar-Ilan University to facilitate lab personnel intending to return to the work place during the COVID-19 epidemic. This study applies our algorithms to guide policies for returning to work. In addition, we present an experimental evaluation of our algorithms on several synthetic data sets as well as on a data set for matching courses to classrooms.

Table 1 shows the list of problems considered in our work and our main results. Due to space limitations, many proofs are omitted; they can be founded in an expanded version (Trabelsi et al. 2022c).

2 Related Work

Resource allocation in multi-agent systems is a well-studied area (e.g., (Chevalere et al. 2006; Gorodetski, Karsaev, and Konushy 2003; Dolgov and Durfee 2006)). The general focus of this work is on topics such as how agents express their requirements, algorithms for allocating resources and evaluating the quality of the resulting allocations. Some complexity and approximability issues in this context are discussed in Nguyen et al. (2013). Zahedi et al. (2020) study the allocation of tasks to agents so that the task allocator can answer queries on counterfactual allocations.

Babaei et al. (2015) presented a survey of approaches to solving timetabling problems that assign university courses to spaces and time slots, avoiding the violation of hard constraints and trying to satisfy soft constraints as much as pos-

sible. A few papers proposed approaches to examination timetabling (e.g., Leite et al. (2019), Elley (2006)). Other work addressed timetabling for high schools (e.g., Fonseca et al. (2016) and Tan et al. (2021)). However, the constraints allowed in many variations of the timetabling problem are much more complex than those in our setting.

Variants of multi-round matching have been considered in game theoretic settings. Anshelevich et al. (2013) present an online matching algorithm where incoming agents are matched in batches. Liu (2020) considers matching a set of long-lived players (akin to resources in our setting) to short-lived players in time-slots. These two works focus on mechanism design to achieve stability and maximize social welfare. Gollapudi et al. (2020) and also Caragiannis and Narang (2022) consider repeated matching with the objective of achieving envy-freeness. However, the latter only considers the case of perfect matching, where agents are matched exactly once in each round. Sühr et al. (2019) considered optimizing fairness in repeated matchings motivated by applications to the ride-hailing problem.

Zanker et al. (2010) discuss the design and evaluation of recommendation systems that allow users to specify soft constraints regarding products of interest. A good discussion on the design of constraint-based recommendation systems appears in Felfernig et al. (2011).

Zhou and Han (2019) propose an approach for a graph-based recommendation system that forms groups of agents with similar preferences to allocate resources. Trabelsi et al. (2022a, 2022b) discussed the problem of advising an agent to modify its preferences so that it can be assigned a resource in a maximum matching. In that work, however, the advice is given to a single agent and only one matching is generated.

To our knowledge, the problems studied in our paper, namely finding multi-round matchings that optimize general benefit functions and advising agents to modify their preferences so that they can be assigned resources in such matchings, have not been addressed in the literature.

3 Definitions and Problem Formulation

3.1 Preliminaries

Agents, resources and matching rounds Let $X = \{x_1, x_2, \dots, x_n\}$ be a set of n agents and $Y = \{y_1, y_2, \dots, y_m\}$ a set of m resources. The objective of the system or the PRINCIPAL is to generate k matchings M_1, \dots, M_k of agents to resources, which we henceforth refer to as a **k-round matching**. However, each agent has certain requirements that need to be satisfied. Firstly, each agent x_i specifies a set $K_i \subseteq \{1, 2, \dots, k\}$ and wants to be matched in ρ_i rounds from K_i .

Restrictions and restrictions graph In addition to round preferences, agents may have restrictions due to which they cannot be matched to certain (or all) resources. We model this using an XY bipartite graph with multiple labels on edges, with each label representing a restriction. We refer to this graph as the **restrictions graph** and denote it as $G_R(X, Y, E_R)$, where E_R is the edge set. Each edge

$e \in E_R$ is associated with a set Γ_e of agent-specific labels representing restrictions.

Compatibility graph Let \mathcal{C}_i denote the set of edge labels or restrictions associated with agent x_i . Suppose a set $C \subseteq \mathcal{C}_i$ of restrictions makes resource y_j *incompatible* with x_i , we draw a labeled edge $\{x_i, y_j\}$ with labels in C . We say that C is the **restrictions set** for the edge $\{x_i, y_j\}$. To make x_i compatible with y_j , *all* the restrictions in C must be removed. There can be hard constraints due to which y_j can never be compatible with x_i even if all the restrictions are removed, in which case, x_i and y_j are not adjacent in G_R . Let $E \subseteq E_R$ be the set of all edges for which the restrictions set is empty. The corresponding subgraph $G(X, Y, E)$ is called the **compatibility graph**. An example showing a restrictions graph, a compatibility graph and multi-round matching appears as Figure 1. This example (Figure 1) is motivated by the *Lab-Space* application considered in this work. There are lab members (agents) who need to be assigned rooms (resources) – one person per room – on five days (i.e., $k = 5$) with some preferences for space and days. Restrictions on each edge are induced by the mismatch between agent preferences and resource properties. For example, x_1 requires a big room, while y_2 is small. Therefore, we add a restriction (or label) big_1 to the edge $\{x_1, y_2\}$. Figure 1 shows the restrictions graph for this problem followed by a multi-round matching solution.

Costs for restrictions removal For each restriction c_i^t , there is a positive cost ψ_i^t associated with removing c_i^t . To remove any set of labels, we use an additive cost model, i.e., for any $C \subseteq \mathcal{C}_i$, the cost incurred by agent x_i for removing all the restrictions in C is $\sum_{c_i^t \in C} \psi_i^t$. An agent x_i is **satisfied** if it is matched in ρ_i rounds belonging to the set K_i . Now, we formally define the multi-round matching and advice generation problems.

3.2 Multi-Round Matching

We begin with a definition of the basic multi-round matching problem (MRM) and discuss other versions.

Multi-Round Matching (MRM)

Instance: A compatibility graph $G(X, Y, E)$, number of rounds of matching k ; for each agent x_i , the set of permissible rounds $K_i \subseteq \{1, \dots, k\}$, the number of rounds $\rho_i \leq |K_i|$ in which x_i wants to be matched.

Requirement: Find a k -round matching that satisfies all the agents (i.e., one that meets the requirements of all the agents), if one exists.

Since a solution may not exist for a given instance of the MRM problem, we consider an extension that uses a **benefit** (or utility) function for each agent. For each agent x_i , the benefit function $\mu_i : \{0, 1, \dots, \rho_i\} \rightarrow \mathbb{R}^+$, where \mathbb{R}^+ is the set of nonnegative real numbers, gives a benefit value $\mu_i(\ell)$ when the number of rounds assigned to x_i is ℓ , $0 \leq \ell \leq \rho_i$. If a k -round matching assigns $\ell_i \leq \rho_i$ rounds to agent x_i , then the total benefit to all the agents is given by $\sum_{i=1}^n \mu_i(\ell_i)$. We can now define a more general version of MRM.

Multi-round matching to maximize total benefit (MAXTB-MRM)

Matching rounds $k = 5$; $\rho_i = 3$ and $\beta_i = 1$, $1 \leq i \leq 4$.
 $K_1 = K_2 = \{1, \dots, 5\}$, $K_3 = K_4 = \{1, 3, 5\}$.
 Attributes: room size (**big** or **small**), WiFi connectivity (**wifi**),
 and cabinet requirement (**cab**). Cost of removing each label is 1.

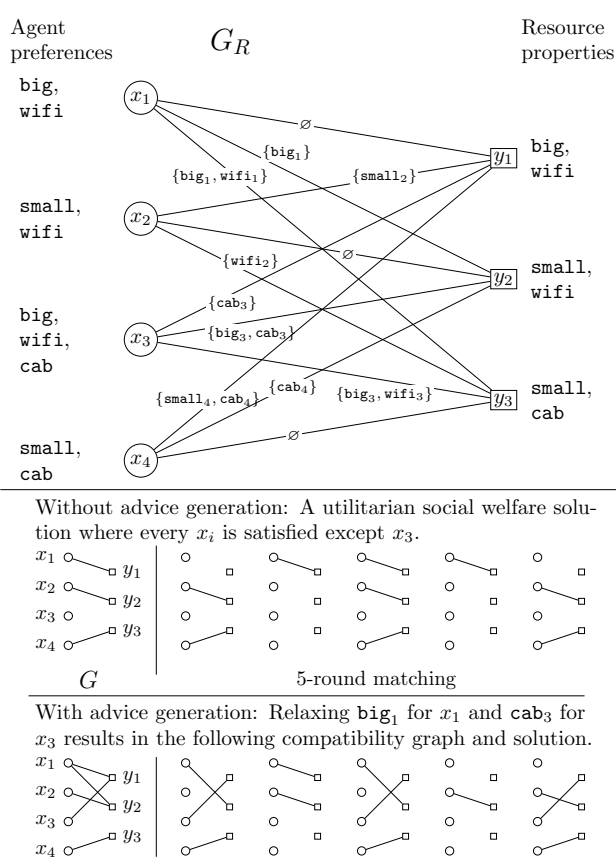


Figure 1: An example: A restrictions graph G_R is induced by agent preferences and resource attributes. The first multi-round matching solution corresponds to the compatibility graph without the advice generation component. Not all agents are satisfied. The second solution satisfies all agents.

Instance: A compatibility graph $G(X, Y, E)$, number of rounds of matching k ; for each agent x_i , the set of permissible rounds $K_i \subseteq \{1, \dots, k\}$, the number of rounds $\rho_i \leq |K_i|$ in which x_i wants to be matched and the benefit function μ_i .

Requirement: Find a k -round matching that maximizes the total benefit over all the agents.

As will be shown in Section 4, when the benefit functions satisfy some properties (including monotonicity and diminishing returns), the MAXTB-MRM problem can be solved efficiently. One such benefit function allows us to obtain an efficient algorithm for the MRM problem. Another benefit function allows us to define the problem of finding a k -round matching that maximizes the number of satisfied agents (abbreviated as MAXSA-MRM). However, this benefit function does not satisfy the diminishing returns property and so our efficient algorithm for MAXTB-MRM can-

not be used for this problem. In fact, we show in Section 4 that the MAXSA-MRM problem is **NP**-hard.

3.3 Advice Generation

We define the multi-round matching advice generation problems, namely AG-MRM and AG-MAXSA-MRM.

Advice Generation for Multi-Round Matching (AG-MRM)

Instance: A restrictions graph $G_R(X, Y, E_R)$, number of rounds of matching k ; for each agent x_i , the set of permissible rounds $K_i \subseteq \{1, \dots, k\}$, the number of times x_i wants to be matched $\rho_i \leq |K_i|$, and budget β_i ; for each label $c_i^t \in C_i$, a cost ψ_i^t of removing that label. (Recall that C_i is the set of labels on the edges incident on agent x_i in G_R .)

Requirement: For each agent x_i , is there a subset $C'_i \subseteq C_i$ such that the following conditions hold? (i) The cost of removing all the labels in C'_i for agent x_i is at most β_i , and (ii) in the resulting compatibility graph, there exists a k -round matching such that all agents are satisfied. If so, find such a k -round matching.

In Figure 1, the advice generation component is illustrated in the bottom-most panel. In this case, we note that there is a solution to the AG-MRM problem.

In general, since a solution may not exist for a given AG-MRM instance, it is natural to consider the version where the goal is to maximize the number of satisfied agents. We denote this version by AG-MAXSA-MRM.

Note: For convenience, we defined optimization versions of problems MAXSA-MRM, AG-MRM and AG-MAXSA-MRM above. In subsequent sections, we show that these problems are **NP**-hard. It can be seen that the decision versions of these three problems are in **NP**; hence, these versions are **NP**-complete.

4 Maximizing Total Benefit

4.1 An Efficient Algorithm for MAXTB-MRM

In this section, we present our algorithm for the MAXTB-MRM problem that finds a k -round matching that maximizes the total benefit over all the agents. This algorithm requires the benefit function for each agent to satisfy some properties. We now specify these properties.

Valid benefit functions: For each agent x_i , $1 \leq i \leq n$, let $\mu_i(\ell)$ denote the non-negative *benefit* that the agent x_i receives if matched in ℓ rounds, for $\ell = 0, 1, 2, \dots, \rho_i$. Let $\delta_i(\ell) = \mu_i(\ell) - \mu_i(\ell - 1)$ for $\ell = 1, 2, \dots, \rho_i$. We say that the benefit function μ_i is **valid** if satisfies *all* of the following four properties: (P1) $\mu_i(0) = 0$; (P2) μ_i is monotone non-decreasing in ℓ ; (P3) μ_i has the *diminishing returns property*, that is, $\mu_i(\ell) - \mu_i(\ell - 1) \leq \mu_i(\ell - 1) - \mu_i(\ell - 2)$ for $2 \leq \ell \leq \rho_i$; and (P4) $\delta_i(\ell) \leq 1$, for $\ell = 1, 2, \dots, \rho_i$. Note that $\mu_i(\cdot)$ satisfies property P3 iff $\delta_i(\cdot)$ is monotone non-increasing in ℓ . Property P4 can be satisfied by normalizing the increments in the benefit values.

Algorithm for MAXTB-MRM: Recall that the goal of this problem is to find a multi-round matching that maximizes the total benefit over all the agents. The basic idea behind our algorithm is to reduce the problem to the maximum

Algorithm 1: ALG-MAXTB-MRM

Input: A compatibility graph $G(X, Y, E)$, for each agent x_i , permissible rounds K_i , the maximum number of rounds ρ_i , and benefit function $\mu_i(\cdot)$.

Output: A matching solution, \mathcal{M}^* .

- 1: $X', Y', E' \leftarrow \emptyset, G' = (X', Y', E')$
 - 2: For each agent $x_i \in X$ and each $h \in K_i$, add a node denoted by x_i^h to X' .
 - 3: For each resource $y_j \in Y$, add k nodes denoted by y_j^h , where $h = 1, 2, \dots, k$, to Y' (type-1 resource nodes).
 - 4: For every agent x_i , add ρ_i nodes, denoted by $y_{i,p}^2$, where $p = 1, 2, \dots, \rho_i$ (type-2).
 - 5: For every agent x_i , add $|K_i| - \rho_i$ nodes, denoted by $y_{i,p}^3$, where $p = 1, 2, \dots, |K_i| - \rho_i$ (type-3).
 - 6: For every edge $\{x_i, y_j\} \in E$ and each $h \in K_i$, add edge $\{x_i^h, y_j^h\}$ to E' with weight $w(x_i^h, y_j^h) = 1$.
 - 7: For each $h \in K_i$ and $p = 1, \dots, \rho_i$, add edge $\{x_i^h, y_{i,p}^2\}$ with weight $w(x_i^h, y_{i,p}^2) = 1 - \delta_i(p)$.
 - 8: For each $h \in K_i$ and $p = 1, \dots, |K_i| - \rho_i$, add edge $\{x_i^h, y_{i,p}^3\}$ with weight $w(x_i^h, y_{i,p}^3) = k + 1$.
 - 9: Find a maximum weighted matching M^* in G'
 - 10: $M_1, M_2, \dots, M_k \leftarrow \emptyset, \mathcal{M}^* = \langle M_1, M_2, \dots, M_k \rangle$
 - 11: For each edge $e = \{x_i^h, y_j^h\} \in M^*$ where y_j^h is a type-1 resource, add the edge $\{x_i, y_j\}$ to M_h .
 - 12: **return** \mathcal{M}^*
-

weighted matching problem on a larger graph. The steps of this construction are shown in ALG-MAXTB-MRM. Given a multi-round solution \mathcal{M} that assigns ℓ_i rounds to agent x_i , $1 \leq i \leq n$, let $\mathbb{B}(\mathcal{M}) = \sum_{i=1}^n \mu_i(\ell_i)$ denote the total benefit due to \mathcal{M} .

Now, we establish the following result.

Theorem 4.1. *Algorithm ALG-MAXTB-MRM produces an optimal solution if every benefit function $\mu_i(\cdot)$ is valid.*

We prove the above theorem through a sequence of definitions and lemmas.

Definition 4.2. A matching M of G' is **saturated** if M includes (i) every node of X' and (ii) every node of Y' that represents a type-3 resource.

Lemma 4.3. *Given any maximum weight matching M for G' , a saturated maximum weight matching M' for G' with the same weight as M can be constructed.*

Proof sketch: We first argue that every maximum weight matching of G' includes all type-3 resource nodes (since each edge incident on those nodes has a large weight). We then argue that other unmatched nodes of X' can be added suitably.

Let the quantity λ be defined by

$$\lambda = \sum_{i=1}^n k(|K_i| - \rho_i) + \sum_{i=1}^n \sum_{\ell=1}^{\rho_i} (1 - \delta_i(\ell)).$$

We use λ throughout the remainder of this section. Note that λ depends only on the parameters of the problem; it does not depend on the algorithm.

Lemma 4.4. *Suppose there is a saturated maximum weight matching M for G' with weight W . Then, a solution to the MAXTB-MRM problem instance with benefit $W - \lambda$ can be constructed if every benefit function $\mu_i(\cdot)$ is valid.*

Proof. Given M , we compute a collection of matchings $\mathcal{M} = \{M_1, M_2, \dots, M_k\}$ for G as follows. For each edge $e = \{x_i^h, y_j^h\} \in M$ where y_j^h is a type-1 resource, add the edge $\{x_i, y_j\}$ to M_h . First, we will show that each $M_h \in \mathcal{M}$ is a matching in G . Suppose M_h is not a matching; then, there exists some x_i or y_j with a degree of at least two in M_h . We will prove it for the first case. The second case is similar. Suppose agent x_i is adjacent to two resources y_j' and y_j'' . This implies that in M , x_i^h is adjacent to both $y_j'^h$ and $y_j''^h$, contradicting the fact that M is a matching in G' . Hence, \mathcal{M} is a valid solution to the MAXTB-MRM problem.

We now derive an expression for the weight W of M . Since every type-3 resource is matched in M , the corresponding edges contribute $\sum_{i=1}^n k(|K_i| - \rho_i)$ to W . For an agent x_i , let the number of edges to type-1 resources be $\gamma_i \leq \rho_i$. Each such edge is of weight 1. The remaining $\rho_i - \gamma_i$ edges correspond to type-2 resources. Since each benefit function, μ_i satisfies the diminishing returns property, the incremental benefits $\delta_i(\cdot)$ are monotone non-increasing in the number of matchings. Thus, we may assume without loss of generality that the top $\rho_i - \gamma_i$ edges in terms of weight are in the maximum weighted matching M . This contributes $\sum_{\ell=1}^{\rho_i - \gamma_i} (1 - \delta_i(\rho_i - \ell + 1))$ to W . Combining the terms corresponding to type-1 and type-2 resources, the contribution of agent x_i to W is $\gamma_i + \sum_{\ell=1}^{\rho_i - \gamma_i} (1 - \delta_i(\rho_i - \ell + 1)) = \sum_{\ell=1}^{\gamma_i} 1 + \sum_{\ell=\gamma_i+1}^{\rho_i} (1 - \delta_i(\ell)) = \sum_{\ell=1}^{\rho_i} (1 - \delta_i(\ell)) + \sum_{\ell=1}^{\gamma_i} \delta_i(\ell)$. Summing this over all the agents, we get $W = \sum_{i=1}^n k(|K_i| - \rho_i) + \sum_{i=1}^n \sum_{\ell=1}^{\rho_i} (1 - \delta_i(\ell)) + \sum_{i=1}^n \sum_{\ell=1}^{\gamma_i} \delta_i(\ell) = \lambda + \sum_{i=1}^n \mu_i(\gamma_i)$. Since $\sum_{i=1}^n \mu_i(\gamma_i) = \mathbb{B}(\mathcal{M})$, the lemma follows. \square

Lemma 4.5. *Suppose there is a solution to the MAXTB-MRM problem instance with benefit Q . Then, there a saturated matching M for G' with weight $Q + \lambda$ can be constructed.*

Proof idea: This proof uses an analysis similar to the one used in the proof of Lemma 4.4.

Lemma 4.6. *There is an optimal solution to the MaxTB-MRM problem instance with benefit Q^* if and only if there is a maximum weight saturated matching M^* for G' with weight $Q^* + \lambda$.*

Proof idea: This is a consequence of Lemmas 4.4 and 4.5.

Theorem 4.1 is a direct consequence of Lemma 4.6.

We now estimate the running time of ALG-MAXTB-MRM.

Proposition 4.7. *Algorithm ALG-MAXTB-MRM runs in time $O(k^{3/2}(n + |E|)\sqrt{n + m})$, where n is the number of agents, m is the number of resources, k is the number of rounds and $|E|$ is the number of edges in the compatibility graph $G(X, Y, E)$.*

4.2 Maximizing Utilitarian Social Welfare

Here, we present a valid benefit function that models utilitarian social welfare. For each agent x_i , let $\mu_i(\ell) = \ell$, for $\ell = 0, 1, \dots, \rho_i$. It is easy to verify that this is a valid benefit function. Hence, Algorithm ALG-MAXTB-MRM can be used to solve the MAXTB-MRM problem with these benefit functions. An optimal solution in this case maximizes the total number of rounds assigned to all the agents, with agent x_i assigned at most ρ_i rounds, $1 \leq i \leq n$.

This benefit function can also be used to show that the MRM problem (where the goal is to check whether there is a solution that satisfies all the agents) can be solved efficiently.

Proposition 4.8. *MRM problem can be solved in polynomial time.*

Proof. For each agent x_i , let the benefit function μ_i be defined by $\mu_i(\ell) = \ell$, for $0 \leq \ell \leq \rho_i$. We will show that for this benefit function, algorithm ALG-MAXTB-MRM produces a solution with total benefit $\sum_{i=1}^n \rho_i$ iff there is a solution to the MRM instance.

Suppose there is a solution to the MRM instance. We can assume (by deleting, if necessary, some rounds in which agents are matched) that each agent x_i is assigned exactly ρ_i rounds. Therefore, the total benefit over all the agents is $\sum_{i=1}^n \rho_i$. Since the maximum benefit that agent x_i can get is ρ_i , this sum also represents the maximum possible total benefit. So, the total benefit of the solution produced by algorithm ALG-MAXTB-MRM is $\sum_{i=1}^n \rho_i$.

For the converse, suppose the maximum benefit produced by the algorithm is equal to $\sum_{i=1}^n \rho_i$. It can be seen that ALG-MAXTB-MRM assigns for each agent x_i , at most ρ_i rounds. (This is due to the type-3 resource nodes for which all the incident edges have large weights.) Since the total benefit is $\sum_{i=1}^n \rho_i$, each agent x_i must be assigned exactly ρ_i rounds. Thus, the MRM instance has a solution. \square

In subsequent sections, we will refer to the algorithm for MRM mentioned in the above proof as ALG-MRM.

4.3 Maximizing Rawlsian Social Welfare

Let $\gamma_i(\mathcal{M})$ denote the number of rounds assigned to agent x_i in a multi-round solution \mathcal{M} . The minimum satisfaction ratio for \mathcal{M} is defined as $\min_i \{\gamma_i(\mathcal{M})/\rho_i\}$. Consider the problem of finding a multi-round matching that *maximizes* the minimum satisfaction ratio. While the maximization objective of this problem seems different from the utilitarian welfare function, we have the following result.

Theorem 4.9. *There exists a reward function μ_i for each agent x_i such that maximizing the total benefit under this function maximizes the Rawlsian social welfare function, i.e., it maximizes the minimum satisfaction ratio over all agents. Further, this reward function is valid and therefore, an optimal solution to the MAXTB-MRM problem under this benefit function can be computed in polynomial time.*

To prove the theorem, we first show how the benefit function is constructed.

1. Let $F = \{k_1/k_2 \mid k_1, k_2 \in \{1, \dots, k\} \text{ and } k_1 \leq k_2\} \cup \{0\}$.

2. Let $\pi : F \rightarrow \{1, \dots, |F|\}$ give the index of each element in F when sorted in descending order.
3. For each $q \in F$, let $\xi(q) = (nk)^{\pi(q)}/(nk)^{|F|}$, where n is the number of agents and k is the number of rounds. Note that each $\xi(q) \in (0, 1]$.
4. The incremental benefit $\delta_i(\ell)$ for an agent x_i for the ℓ th matching is defined as $\delta_i(\ell) = \xi((\ell - 1)/\rho_i)$. Thus, the benefit function μ_i for agent x_i is given by $\mu_i(0) = 0$ and $\mu_i(\ell) = \mu_i(\ell - 1) + \delta_i(\ell)$ for $1 \leq \ell \leq \rho_i$.

It can be seen that μ_i satisfies properties P1, P2, and P4 of a valid benefit function. We now show that it also satisfies P3, the diminishing returns property.

Lemma 4.10. *For each agent x_i , the incremental benefit $\delta_i(\ell)$ is monotone non-increasing in ℓ . Therefore, $\mu_i(\cdot)$ satisfies the diminishing returns property.*

Proof. By definition, for $\ell = 1, \dots, \rho_i - 1$, $\delta_i(\ell)/\delta_i(\ell+1) = \xi((\ell - 1)/\rho_i)/\xi(\ell/\rho_i) \geq nk > 1$, by noting that $\pi(\ell/\rho_i) - \pi((\ell + 1)/\rho_i) \geq 1$. Hence, the lemma holds. \square

To complete the proof of Theorem 4.9, we must show that any solution that maximizes the total benefit for the above benefit function also maximizes the minimum satisfaction ratio. This proof appears in (Trabelsi et al. 2022c).

4.4 Hardness of MAXSA-MRM

Here we consider the MAXSA-MRM problem, where the goal is to find a k -round matching to maximize the number of satisfied agents. A benefit function that models this problem is as follows. For each agent x_i , let $\mu_i(\ell) = 0$ for $0 \leq \ell \leq \rho_i - 1$ and $\mu_i(\rho_i) = 1$. This function can be seen to satisfy properties P1, P2, and P4 of a valid benefit function. However, it does *not* satisfy P3, the diminishing returns property, since $\delta(\rho_i - 1) = 0$ while $\delta(\rho_i) = 1$. Thus, this is not a valid benefit function. This difference is enough to change the complexity of the problem.

Theorem 4.11. *MAXSA-MRM is NP-hard even when the number of rounds (k) is 3.*

Proof idea: We use a reduction from the Minimum Vertex Cover Problem for Cubic graphs which is known to be NP-complete (Garey and Johnson 1979).

5 Advice Generation Problems

Here, we first point out that the AG-MRM and AG-MAXSA-MRM problems are NP-hard. We present two solution approaches for the AG-MAXSA-MRM problem, namely an integer linear program (ILP) formulation and a pruned-search-based optimization using ALG-MRM.

5.1 Complexity Results

Theorem 5.1. *The AG-MRM problem is NP-hard when there is just one agent x_i for which $\rho_i > 1$.*

Proof idea. We use a reduction from the Minimum Set Cover (MSC) problem which is known to be NP-complete (Garey and Johnson 1979).

Any solution to the AG-MRM problem must satisfy *all* the agents. Thus, the hardness of AG-MRM also yields the following:

Corollary 5.2. AG-MAXSA-MRM is *NP-hard*.

5.2 Solution Approaches for AG-MAXSA-MRM

(a) ILP formulation for AG-MAXSA-MRM: A complete ILP formulation for solving the problem AG-MAXSA-MRM can be found at (Trabelsi et al. 2022c).

(b) ALG-MRM guided optimization: We describe a local search heuristic PS-AG-MAXSA-MRM for the MS-AG-MRM problem. It consists of three parts: (1) identification of relevant candidate sets of relaxations for each agent and edges in the restrictions graph (and hence the name **pruned-search**), (2) the valuation of an identified set of relaxations for all agents and (3) a local search algorithm.

In part (1), for each agent x_i , we first identify the sets of restrictions $\Gamma_i = \{C'_i \subseteq C_i \mid \text{Cost of removing all the labels in } C'_i \text{ for agent } x_i \text{ is } \leq \beta_i\}$ to be considered for removal during the search. We choose a subcollection $\Gamma_i^s \subseteq \Gamma_i$ of sets based on the following criteria: (a) only Pareto optimal sets of constraints with respect to the budget are considered, i.e., if $C'_i \in \Gamma_i^s$, then there is no $C''_i \in C_i$ such that the cost of relaxation of $C'_i \cup \{C''_i\}$ is less than or equal to β_i ; (b) each $C'_i \in \Gamma_i^s$ is associated with a set of edges E'_i that will be added to the compatibility graph if C'_i is relaxed; if C'_i and C''_i are such that $E''_i \subseteq E'_i$, then, we retain only C'_i ; (c) if two sets of constraints cause the addition of the same set of edges to the compatibility graph, only one of them is considered.

In part (2), given a collection of relaxation sets, one for each unsatisfied agent, we apply ALG-MRM from Section 4 on the induced compatibility graph. The number of agents that are satisfied in the k -round solution obtained serves as the valuation for the collection of relaxation sets. In part (3), we use simulated annealing local search algorithm (Aarts and Korst 1989) for the actual search.

6 Experimental Results

We applied the algorithms developed in this work to several datasets. Evaluations were based on relevant benefit functions and computing time. We used at least 10 replicates for each experiment when needed (e.g., random assignment of attributes and while using PS-AG-MAXSA-MRM, which is a stochastic optimization algorithm). Error bars are provided in such instances. We considered two novel real-world applications (discussed below). We first describe the attributes of the datasets; the data and the code for running the experiments are available at <https://github.com/yohayt/Resource-Sharing-Through-Multi-Round-Matchings>.

Lab-Space application. We conducted a back-to-the-lab desk sharing study in a lab at Bar-Ilan University (Ramat Gan, Israel) to facilitate lab members intending to return to in-person work during the COVID-19 epidemic. A survey was used to collect the preferences of members. The lab has 14 offices and 31 members. Six of the offices can accommodate two students each and the remaining eight can accommodate one student each. The agent restrictions are: vicinity

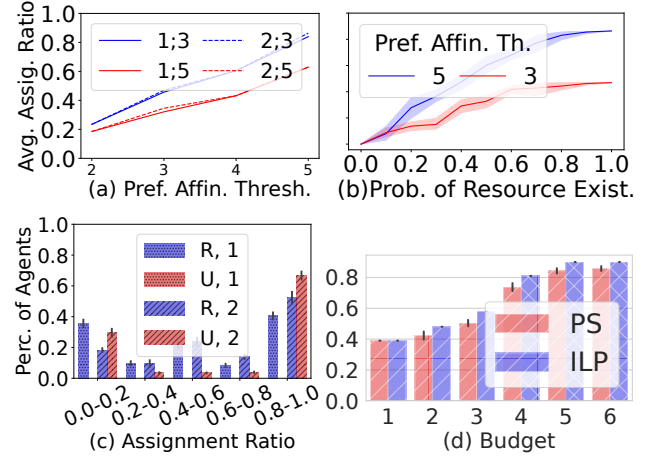


Figure 2: *Lab-Space* experiment results: (a) Average assignment ratio for different preference affinity thresholds. Legend: Maximum Students per Room, preference affinity threshold. (b) Average assignment ratio when the number of resources is restricted; (c) Comparison of different welfare functions for resource probability 0.5, and preference affinity and day affinity thresholds of 5 and 2 respectively. Legend: Rawlsian/Utilitarian, Maximum Students in Room; and (d) Evaluation of the advice generation algorithm for increasing budget for one student per room. y is the fraction of fully satisfied agents.

to the advisor’s room, presence of a cabinet, WiFi connectivity, ambient noise, vicinity to the kitchen, room size, and vicinity to the bathroom. All agent preferences are generated from the survey as a function of the affinity (from 1 (low) to 5 (high)) they provide for each attribute. These affinities are also used as costs for attribute removal.

Course-Classroom dataset. This dataset also comes from Bar-Ilan University; it is for the year 2018–2019. There are 144 classrooms and 153 courses. In the experiments, we focused on two hours on Tuesday and used all the courses that are scheduled in this time slot and all available classes. There were 142 courses and all the 144 classrooms were available for them. Each classroom has two attributes: its *capacity* and the *region* to which it belongs. Although the problem of assigning classrooms to courses is quite common (Phillips et al. 2015), we did not find any publicly available dataset that could be used here. The number of rounds is chosen randomly between 1 and 3 as this data was not available. Also, to generate K_i , ρ_i rounds are sampled uniformly followed by choosing the remaining days with probability 0.5. We used 10 replicates for each experiment.

The generation of the restrictions and compatibility graphs from the agent preferences and resource attributes is described in detail in (Trabelsi et al. 2022c).

Lab-Space multi-round matching. Through this study, the head of the lab (PRINCIPAL) wanted answers to the following questions: Is it possible to accommodate only one person per office? Is it possible to maximize the number

of relevant assignments? Is it possible to achieve all this by satisfying as many preferences of lab members as possible? The number of rounds is 5 (one per weekday). We defined a preference affinity threshold. If the agent's survey score for affinity (1-5) is at least as much as the threshold, then we add this preference as a constraint. Note that the compatibility graph generated for threshold τ is a subgraph of the one generated for $\tau + 1$. The permissible sets of rounds K_i are generated the same way. A day affinity threshold is set. If the agent's preference for a weekday is lower than this threshold, then we do not add it to K_i . In Figure 2(a), we have results for maximizing the utilitarian social welfare for various compatibility graphs generated by adding constraints based on the agent's affinity to each preference. We see that adhering strictly to the preferences of agents gives a very low average assignment ratio γ_i/ρ_i , where γ_i is the number of rounds assigned to agent x_i (who requested ρ_i rounds). We considered two scenarios for large rooms: one or two students. We note that accommodating two students significantly improves the average assignment ratio. Adhering to the preferred days of students does not seem to have many costs, particularly when the preference affinity threshold is high. In Figure 2(b), we demonstrate how critical the current set of resources is for the functioning of the lab. We withheld only a portion of the resources (by sampling) to satisfy agent preferences. However, with just 60% of the resources, it is possible to get an average assignment ratio of 0.6, albeit by ignoring most of the agent preferences. In Figure 2(c), we compare utilitarian social welfare with Rawlsian social welfare. We note that in the utilitarian welfare, many agents end up with a lower assignment ratio compared to Rawlsian, while the total number of matchings across rounds is the same. (A theorem presented in (Trabelsi et al. 2022c) states that for valid and monotone strictly increasing benefit functions, an optimal solution for total benefit also optimizes the total number of rounds assigned to agents.) Hence, from the perspective of fairness in agent satisfaction, the Rawlsian reward function performs better.

Lab-Space advice generation. From the results in Figure 2(d), we observe that there is a sharp increase in the number of agents satisfied for a budget of 4. In the data, there is a strong preference for WiFi (an average score of 4.3 out of 5), while more than 40% of the rooms have poor connectivity. When most of the agents relax this preference, they have access to many rooms. This partially explains the sharp increase in the number of matchings. We also observe that the performance of PS-AG-MAXSA-MRM is close to that of optimum (given by ILP-MAX-AG-MRM).

Course-Classroom multi-round matching. Here, we considered two scenarios: having classes (i) five days a week and (ii) six days a week. A six-day week can accommodate more courses while on the other hand, facilities have to be kept open for an additional day. The results of maximizing the utilitarian multi-round matching are in Figure 3(a). We observe that as ρ_i is increased, the average assignment ratio decreases slowly. The difference between five and six days a week is not significant. Figure 3(b) compares utilitarian and Rawlsian reward functions. Unlike the *Lab-Space* re-

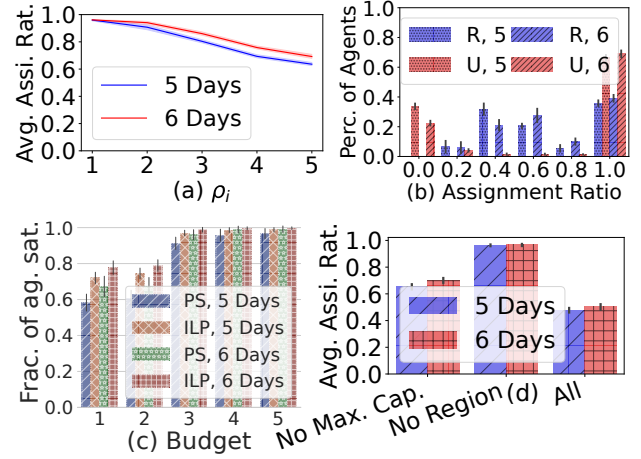


Figure 3: *Course-Classroom* experimental results: (a) Average assignment ratio vs. number of rounds required per agent; (b) Comparison of Rawlsian (R) and Utilitarian (U) welfare functions. (c) Evaluation of advice generation algorithm for increasing budget. (d) Importance of attributes. Days in legends is the number of working days per week.

sults, here we clearly see that the minimum assignment ratio is higher in the former case.

Course-Classroom advice generation In Figure 3(c) we observe that only for low budgets, the difference between five-day week and six-day week solutions is significant. In the same regime, we see a significant difference between PS-AG-MAXSA-MRM and ILP-MAX-AG-MRM. By inspecting the solution sets, we observed that relaxing the region attribute is most effective in increasing the number of satisfied agents, while minimum capacity is least effective. To study the importance of different attributes, we conducted experiments where a single chosen attribute was omitted from the restrictions list. Figure 3(d) indicates that the region attribute has the most impact on the assignments ratio.

Scaling to larger networks. To analyze the performance of the of the advice generation algorithms with respect to the size of the network, we experimented with complete bipartite graphs of various sizes. The results are in (Trabelsi et al. 2022c). In these experiments, we note that not only is the solution obtained using the PS-AG-MAXSA-MRM close to ILP-MAX-AG-MRM, it is orders of magnitude faster.

7 Directions for Future Work

We conclude by mentioning a few directions for future work. One direction is to consider other models for producing compatibility graphs from restrictions graphs. (For example, an edge may be added to the compatibility graph if at least one label on that edge is removed.) It will also be interesting to consider the multi-round matching problems in an online setting, where the set of agents varies over time and compatibility is round-dependent. Finally, it is of interest to investigate approximation algorithms with provable performance guarantees for MAXSA-MRM and AG-MAXSA-MRM.

Acknowledgments

We thank the AAAI 2023 reviewers for their feedback. This research has been partially supported by the Israel Science Foundation under grant 1958/20, the EU Project TAILOR under grant 952215, University of Virginia Strategic Investment Fund award number SIF160, and the US National Science Foundation grant OAC-1916805 (CINES).

References

- Aarts, E.; and Korst, J. 1989. *Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing*. New York, NY: Wiley.
- Anshelevich, E.; Chhabra, M.; Das, S.; and Gerrior, M. 2013. On the social welfare of mechanisms for repeated batch matching. In *AAAI*, 60–66.
- Babaei, H.; Karimpour, J.; and Hadidi, A. 2015. A survey of approaches for university course timetabling problem. *Computers & Industrial Engineering*, 86: 43–59.
- Cai, H.; and Khan, S. 2010. The common first year studio in a hot-desking age: An explorative study on the studio environment and learning. *J. Education in the Built Environment*, 5(2): 39–64.
- Caragiannis, I.; and Narang, S. 2022. Repeatedly Matching Items to Agents Fairly and Efficiently. *arXiv preprint arXiv:2207.01589*.
- Chevaleyre, Y.; Dunne, P. E.; Endriss, U.; Lang, J.; Lemaître, M.; Maudet, N.; Padget, J. A.; Phelps, S.; Rodríguez-Aguilar, J. A.; and Sousa, P. 2006. Issues in Multiagent Resource Allocation. *Informatica (Slovenia)*, 30(1): 3–31.
- District Management Group. 2020. COVID-19: Creating Elementary School Schedules to Support Social Distancing DMGroup Research Briefs. <https://f.hubspotusercontent00.net/hubfs/3412255/Resources%20Page%20Files%20-%20Public/DMGroup-School-Restart-Research-Brief-Elementary-School-Schedules-and-Social-Distancing%5F2020-7.pdf>. Accessed: 2023-03-21.
- Dolgov, D. A.; and Durfee, E. H. 2006. Resource Allocation Among Agents with MDP-Induced Preferences. *J. Artif. Intell. Res.*, 27: 505–549.
- Eley, M. 2006. Ant algorithms for the exam timetabling problem. In *International Conference on the Practice and Theory of Automated Timetabling*, 364–382. Springer.
- Enriching Students. 2020. School Schedules During COVID-19—A Guide to assist schools. <https://www.enrichingstudents.com/school-schedules-during-covid-19/>. Accessed: 2023-03-21.
- Felfernig, A.; Friedrich, G.; Jannach, D.; and Zanker, M. 2011. Developing Constraint-based Recommenders. In *Recommender Systems Handbook*, 187–215. Springer.
- Fonseca, G. H.; Santos, H. G.; and Carrano, E. G. 2016. Integrating matheuristics and metaheuristics for timetabling. *Computers & Operations Research*, 74: 108–117.
- Garey, M. R.; and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman and Co.
- Gilbert, F. 2018. A Guide to Sharing Farm Equipment. <https://projects.sare.org/wp-content/uploads/Sharing-Guide-2018-%5F-Web.pdf>. Accessed: 2023-03-21.
- Gollapudi, S.; Kollias, K.; and Plaut, B. 2020. Almost envy-free repeated matching in two-sided markets. In *Int. Conf. on Web and Internet Economics*, 3–16.
- Gorodetski, V. I.; Karsaev, O.; and Konushy, V. 2003. Multi-agent System for Resource Allocation and Scheduling. In *Proc. CEEMAS*, 236–246.
- Karamanis, R.; Anastasiadis, E.; and Angeloudis, M., P. Stettler. 2021. Assignment and pricing of shared rides in ride-sourcing using combinatorial double auctions. *IEEE Trans. Intelligent Transportation Systems*, 22(9): 5648–5659.
- Kucharski, R.; and Cats, O. 2020. Exact matching of attractive shared rides (ExMAS) for system-wide strategic evaluations. *Transportation Research Part B: Methodological*, 139: 285–310.
- Leite, N.; Melício, F.; and Rosa, A. C. 2019. A fast simulated annealing algorithm for the examination timetabling problem. *Expert Systems with Applications*, 122: 137–151.
- Liu, C. 2020. Stability in repeated matching markets. *arXiv preprint arXiv:2007.03794*.
- Nguyen, T. T.; Roos, M.; and Rothe, J. 2013. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *AMAI*, 68(1-3): 65–90.
- Parker, L. D. 2020. The COVID-19 office in transition: cost, efficiency and the social responsibility business case. *Accounting, Auditing & Accountability Journal*, 33(8): 1943–1967.
- Phillips, A. E.; Waterer, H.; Ehr Gott, M.; and Ryan, D. M. 2015. Integer programming methods for large-scale practical classroom assignment problems. *Computers & Operations Research*, 53: 42–53.
- Rakhra, M.; and Singh, R. 2020. Internet Based Resource Sharing Platform development For Agriculture Machinery and Tools in Punjab, India. In *Proc. 8th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO)*, 636–642.
- Rawls, J. 1999. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Stark, O. 2020. An economics-based rationale for the Rawlsian social welfare program. University of Tübingen Working Papers in Business and Economics, No. 137.
- Sühr, T.; Biega, A. J.; Zehlke, M.; Gummadi, K. P.; and Chakraborty, A. 2019. Two-sided fairness for repeated matchings in two-sided markets: A case study of a ride-hailing platform. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 3082–3092.

- Tan, J. S.; Goh, S. L.; Kendall, G.; and Sabar, N. R. 2021. A survey of the state-of-the-art of optimisation methodologies in school timetabling problems. *Expert Systems with Applications*, 165: 113943.
- Trabelsi, Y.; Adiga, A.; Kraus, S.; and Ravi, S. S. 2022a. Maximizing Resource Allocation Likelihood with Minimum Compromise. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1738–1740.
- Trabelsi, Y.; Adiga, A.; Kraus, S.; and Ravi, S. S. 2022b. Resource Allocation to Agents with Restrictions: Maximizing Likelihood with Minimum Compromise. In *Proc. European Conference on Multi-Agent Systems (EUMAS)*, 403–420.
- Trabelsi, Y.; Adiga, A.; Kraus, S.; Ravi, S. S.; and Rosenkrantz, D. J. 2022c. Resource Sharing Through Multi-Round Matchings. *arXiv preprint arXiv:2211.17199*.
- Varone, S.; and Beffa, C. 2019. Dataset on a problem of assigning activities to children, with various optimization constraints. *Data in brief*, 25: 104168.
- Viner, J. 1949. Bentham and JS Mill: The utilitarian background. *The American Economic Review*, 39(2): 360–382.
- Wong, M. 2021. Self-scheduler for dental students booking consultations with faculty during the COVID-19 pandemic. *Journal of Dental Education*, 85: 1984–1985.
- Zahedi, Z.; Sengupta, S.; and Kambhampati, S. 2020. 'Why not give this work to them?' Explaining AI-Moderated Task-Allocation Outcomes using Negotiation Trees. *CoRR*, abs/2002.01640.
- Zanker, M.; Jessenitschnig, M.; and Schmid, W. 2010. Preference reasoning with soft constraints in constraint-based recommender systems. *Constraints: An Int. J.*, 15(4): 574–595.
- Zhou, W.; and Han, W. 2019. Personalized Recommendation via User Preference Matching. *Information Processing and Management*, 56(3): 955–968.