

Anytime Heuristic for Weighted Matching Through Altruism-Inspired Behavior

Panayiotis Danassis, Aris Filos-Ratsikas, Boi Faltings

Artificial Intelligence Laboratory (LIA), École Polytechnique Fédérale de Lausanne (EPFL)
{panayiotis.danassis, aris.filosratsikas, boi.faltings}@epfl.ch

ABSTRACT

We present a novel anytime heuristic (ALMA), inspired by the human principle of altruism, for solving the assignment problem. ALMA is decentralized, completely uncoupled, and requires no communication between the participants. We prove an upper bound on the convergence speed that is polynomial in the desired number of resources and competing agents per resource; crucially, in the realistic case where the aforementioned quantities are bounded independently of the total number of agents/resources, the convergence time remains *constant* as the total problem size increases.

We have evaluated ALMA under three test cases: (i) an anti-coordination scenario where agents with similar preferences compete over the same set of actions, (ii) a resource allocation scenario in an urban environment, under a constant-time constraint, and finally, (iii) an on-line matching scenario using *real* passenger-taxi data. In all of the cases, ALMA was able to reach high social welfare, while being orders of magnitude faster than the centralized, optimal algorithm. The latter allows our algorithm to scale to realistic scenarios with hundreds of thousands of agents, e.g., vehicle coordination in urban environments.

KEYWORDS

Coordination and Cooperation; Resource Allocation; Multi-agent Learning

1 INTRODUCTION

One of the most relevant problems in multi-agent systems (MAS) is finding an optimal allocation between agents. This pertains to role allocation (e.g., team formation for autonomous robots [18]), task assignment (e.g., employees of a factory, taxi-passenger matching [33]), resource allocation (e.g., parking spaces and/or charging stations for autonomous vehicles [15]), etc. What follows is *applicable to any such scenario*, but for concreteness we will refer to the allocation of a set of resources to a set of agents, a setting known as the *assignment problem*, one of the most fundamental combinatorial optimization problems [23].

When designing algorithms for assignment problems, a significant challenge emerges from the nature of real-world applications, which is often distributed and information-restrictive. For the former part, a variety of decentralized algorithms have been developed [6, 17, 19, 34], all of which, though, require polynomial in the problem size number of messages. However, inter-agent interactions often repeat no more than a few hundreds of times. Moreover, sharing plans and preferences creates high overhead, and there is often a lack of responsiveness and/or communication between the participants [29]. Achieving fast convergence and high efficiency in such information-restrictive settings is extremely challenging. Yet,

humans are able to routinely and robustly coordinate in similar everyday scenarios. One driving factor that facilitates human cooperation is the principle of *altruism* [7, 16, 25]. Inspired by human behavior, the proposed heuristic is modeled on the principle of altruism. This results to fast convergence to highly efficient allocations, without any communication between the agents.

A distinctive characteristic of ALMA is that agents make decisions locally, based on (i) the contest for resources that *they* are interested in, (ii) the agents that are interested in the *same* resources. If each agent is interested in only a *subset* of the total resources, ALMA converges in time polynomial in the maximum size of the subsets; not the total number of resources. In particular, if the size of each subset is a constant fraction of the total number of resources, then the convergence time is *constant*, in the sense that it does not grow with the problem size. The same is not true for other algorithms (e.g., the optimal centralized solution) which require time polynomial in the *total* number of agents/resources, even if the aforementioned condition holds. The condition holds by default in many real-world applications; agents have only local knowledge of the world, there is typically a cost associated with acquiring a resource, or agents are simply only interested in resources in their vicinity (e.g., urban environments). This is important, as the proposed approach avoids having to artificially split the problem in subproblems (e.g. by placing bounds or spatial constraints) and solve those separately, in order to make it tractable. Instead, ALMA utilizes a natural domain characteristic, instead of an artificial optimization technique (i.e., artificial bounds). Coupled to the convergence time, the decentralized nature of ALMA makes it applicable to large-scale, real-world applications (e.g., IoT devices, intelligent infrastructure, autonomous vehicles, etc.).

1.1 Our Results

Our main contributions in this paper are:

(1) We introduce a novel, anytime **AL**truistic **MA**tching heuristic (**ALMA**) for solving the assignment problem. ALMA is decentralized, completely uncoupled (i.e., each agent is only aware of his own history of action/reward pairs [31]), and requires no communication between the agents.

(2) We prove that if we bound the maximum number of resources an agent is interested in, and the maximum number of agents competing for a resource, the expected number of steps for any agent to converge is independent of the total problem size. Thus, we do not require to artificially split the problem, or similar techniques, to render it manageable.

(3) We provide a thorough empirical evaluation of ALMA on both synthetic and real data. In particular, we have evaluated ALMA under three test cases: (i) an anti-coordination scenario where agents with similar preferences compete over the same set of actions, (ii)

a resource allocation scenario in an urban environment, under a constant-time constraint, and finally, (iii) an on-line matching scenario using *real* passenger-taxi data. In all of the cases, ALMA achieves high social welfare (total satisfaction of the agents) as compared to the optimal solution, as well as various other algorithms.

1.2 Related Work

The assignment problem consists of finding a maximum weight matching in a weighted bipartite graph and it is one of the best-studied combinatorial optimization problems in the literature. The first polynomial time algorithm (with respect to the total number of nodes, and edges) was introduced by Jacobi in the 19th century [4, 26], and was succeeded by many classical algorithms [3, 14, 23] with the *Hungarian* algorithm of Kuhn [1955] being the most prominent one (see [30] for an overview). The problem can also be solved via linear programming [12], as its LP formulation relaxation admits integral optimal solutions [27]. In Section 3.3, we will apply ALMA on a non-bipartite setting, which corresponds to the more general *maximum weight matching* problem on general graphs. To compute the optimal in this case, we will use the *blossom algorithm* of [13] (see [22]).

In reality, a centralized coordinator is not always available, and if so, it has to know the utilities of all the participants, which is often not feasible. In the literature of the assignment problem, there also exist several decentralized algorithms (e.g., [6, 17, 19, 34] which are the decentralized versions of the aforementioned well-known centralized algorithms). However, these algorithms require polynomial computational time and polynomial number of messages (such as cost matrices [19], pricing information [34], or a basis of the LP [6], etc.). Yet, agent interactions often repeat no more than a few hundreds of times. To the best of our knowledge, a decentralized algorithm that requires no message exchange (i.e., no communication network) between the participants, and achieves high efficiency, like ALMA does, has not appeared in the literature before. Let us stress the importance of such a heuristic: as autonomous agents proliferate, and their number and diversity continue to rise, differences between the agents in terms of origin, communication protocols, or the existence of sub-optimal, legacy agents will bring forth the need to collaborate without any form of explicit communication [29]. Finally, communication between participants creates high overhead as well.

ALMA is inspired by the decentralized allocation algorithm of [10]. Using such a simple learning rule which only requires environmental feedback, allows our approach to scale to hundreds of thousands of agents. Moreover, it does not require global knowledge of utilities; only local knowledge of personal utilities (in fact, we require knowledge of pairwise differences which are far easier to estimate).

2 ALTRUISTIC MATCHING HEURISTIC

In this section, we define ALMA and prove its convergence properties. We begin with the definition of the assignment problem, and its interpretation in our setting.

2.1 The Assignment Problem

The assignment problem consists of finding a maximum weight matching in a weighted bipartite graph, $\mathcal{G} = \{\mathcal{N} \cup \mathcal{R}, \mathcal{E}\}$. In the studied scenario, $\mathcal{N} = \{1, \dots, N\}$ agents compete to acquire $\mathcal{R} = \{1, \dots, R\}$ resources. We assume that each agent n is interested in a subset of the total resources, i.e., $\mathcal{R}^n \subseteq \mathcal{R}$. The weight of an edge $(n, r) \in \mathcal{E}$ represents the utility $(u_n(r))$ agent n receives by acquiring resource r . Each agent can acquire at most one resource, and each resource can be assigned to at most one agent. The goal is to maximize the social welfare (sum of utilities), i.e., $\max_{x \geq 0} \sum_{(n,r) \in \mathcal{E}} u_n(r) x_{n,r}$, subject to $\sum_{r|(n,r) \in \mathcal{E}} x_{n,r} = 1, \forall n \in \mathcal{N}$, and $\sum_{n|(n,r) \in \mathcal{E}} x_{n,r} = 1, \forall r \in \mathcal{R}$.

2.2 Learning Rule

This section describes the proposed heuristic (**ALMA: ALtruistic MAtching** heuristic) for weighted matching. We make the following two assumptions: First, we assume (possibly noisy) knowledge of personal preferences by each agent. Second, we assume that agents can observe feedback from their environment. This is used to inform collisions and detect free resources. It could be achieved by the use of visual, auditory, olfactory sensors etc., or by any other means of feedback from the resource (e.g., by sending an occupancy message). Note here that these messages would be between the requesting agent and the resource, not between the participating agents themselves, and that it suffices to send only 1 bit of information (e.g., 0, 1 for occupied / free respectively).

ALMA learns the right action through repeated trials as follows. Each agent sorts his available resources (possibly $\mathcal{R}^n \subseteq \mathcal{R}$) in decreasing order of utility $(r_1, r_2, \dots, r_i, r_{i+1}, \dots, r_{R^n})$. The set of available actions is denoted as $\mathcal{A} = \{Y, A_{r_1}, \dots, A_{r_{R^n}}\}$, where Y refers to yielding, and A_r refers to accessing resource r . Each agent has a strategy (g_n) that points to a resource and it is initialized to the most preferred one. As long as an agent has not acquired a resource yet, at every time-step, there are two possible scenarios. If $g_n = A_r$ (strategy points to resource r), then agent n attempts to acquire that resource. If there is a collision, the colliding parties back-off with some probability. Otherwise, if $g_n = Y$, the agent chooses a resource r for monitoring. If the resource is free, he sets $g_n \leftarrow A_r$. Alg. 1 presents the pseudo-code of ALMA, which is followed by every agent individually. The back-off probability and the next resource to monitor are computed individually and locally based on the current resource and each agent's utilities, as will be explained in the following section. Finally, note that if the available resources change over time, the agents simply need to sort again the currently available ones.

2.3 Back-off Probability & Resource Selection

Let \mathcal{R} be totally ordered in decreasing utility under $<_n, \forall n \in \mathcal{N}$. If more than one agent compete for resource r_i (step 4 of Alg. 1), each of them will back-off with probability that depends on their utility loss of switching to their respective remaining resources. The loss is given by Eq. 1.

$$loss_n^i = \frac{\sum_{j=i+1}^k u_n(r_i) - u_n(r_j)}{k - i} \quad (1)$$

where $k \in \{i+1, \dots, R^n\}$ denotes the number of remaining resources to be considered. For $k = i+1$, the formula only takes into account the utility loss of switching to the immediate next best resource, while for $k = R^n$ it takes into account the average utility loss of switching to all of the remaining resources. In the remainder of the paper we assume $k = i+1$, i.e., $loss_n^i = u_n(r_i) - u_n(r_{i+1})$. The actual back-off probability can be computed with any monotonically decreasing function f on $loss_n$, i.e., $P_n(r_i, <_n) = f_n(loss_n^i)$. In the evaluation section, we have used two such functions, a linear (Eq. 2), and the logistic function (Eq. 3). The parameter ϵ places a threshold on the minimum / maximum back-off probability for the linear curve, while γ determines the steepness of the logistic curve.

$$f(loss) = \begin{cases} 1 - \epsilon, & \text{if } loss \leq \epsilon \\ \epsilon, & \text{if } 1 - loss \leq \epsilon \\ 1 - loss, & \text{otherwise} \end{cases} \quad (2)$$

$$f(loss) = \frac{1}{1 + e^{-\gamma(0.5 - loss)}} \quad (3)$$

Using the aforescribed rule, agents that do not have good alternatives will be less likely to back-off and vice versa. The ones that do back-off select an alternative resource and examine its availability. The resource selection is performed in sequential order, i.e., $S_n(r_{\text{prev}}, <_n) = r_{\text{prev}+1}$, where r_{prev} denotes the resource selected by that agent in the previous round. We also examined the possibility of using a weighted or uniformly at random selection, but achieved inferior results.

2.4 Altruism-Inspired Behavior

ALMA is inspired by the human principle of altruism. We would expect an altruistic person to give up a resource either to someone who values it more, if that resulted in an improvement of the well-being of society [7], or simply to be nice to others [28]. Such behavior is especially common in situations where the backing-off subject has equally good alternatives. For example, in human pick-up teams, each player typically attempts to fill his most preferred position. If there is a collision, a colliding player might back-off because his teammate is more competent in that role, or because he has an equally good alternative, or simply to be polite; the player backs-off now and assumes that role at some future game. From an alternative viewpoint, following such an altruistic convention leads to a faster convergence which outweighs the loss in utility. Such conventions allow humans to routinely and robustly coordinate in large scale and under dynamic and unpredictable demand. Behavioral conventions are a fundamental part of human societies [21], yet they have not appeared meaningfully in empirical modeling of multi-agent systems. Inspired by human behavior, ALMA attempts to reproduce these simple rules in an artificial setting.

2.5 Convergence

Agents who have not acquired a resource ($g_n = Y$) will not claim an occupied one. Additionally, every time a collision happens, there is a positive probability that some agents will back-off. As a result, the system will converge. The following theorem proves that the expected convergence time is logarithmic in the number of agents N and quadratic in the number of resources R .

Algorithm 1 ALMA: Altruistic Matching Heuristic.

Require: Sort resources ($\mathcal{R}^n \subseteq \mathcal{R}$) in decreasing order of utility $r_1, r_2, \dots, r_i, r_{i+1}, \dots, r_{R^n}$.

Require: Initialize $g_n \leftarrow A_{r_1}$, and $r_{\text{prev}} \leftarrow r_1$.

```

1: procedure ALMA
2:   if  $g_n = A_r$  then
3:     Agent  $n$  attempts to acquire resource  $r$ . Set  $r_{\text{prev}} \leftarrow r$ .
4:     if Collision( $r$ ) then
5:       back-off (set  $g_n \leftarrow Y$ ) with probability  $P_n(r, <_n)$ .
6:     else ( $g_n = Y$ )
7:       Agent  $n$  monitors  $r \leftarrow S_n(r_{\text{prev}}, <_n)$ . Set  $r_{\text{prev}} \leftarrow r$ .
8:       if Free( $r$ ) then set  $g_n \leftarrow A_r$ .
```

THEOREM 2.1. For N agents and R resources, the expected number of steps until the system of agents following Alg. 1 converges to a complete matching is bounded by (4), where $p^* = f(loss^*)$, and $loss^*$ is given by Eq. 5.

$$O\left(R \frac{2-p^*}{2(1-p^*)} \left(\frac{1}{p^*} \log N + R\right)\right) \quad (4)$$

$$loss^* = \arg \min_{loss_n^r} \left(\min_{r \in \mathcal{R}, n \in \mathcal{N}} (loss_n^r), 1 - \max_{r \in \mathcal{R}, n \in \mathcal{N}} (loss_n^r) \right) \quad (5)$$

PROOF. To improve readability, we will only provide a sketch of the proof. Please see [11] for the complete version. The proof is based on [9, 10].

We first assume that every agent, on every collision, backs-off with the same constant probability p . We start with the case of having N agents competing for 1 resource and model our system as a discrete time Markov chain. Intuitively, this Markov chain describes the number of individuals in a decreasing population, but with two caveats: the goal (absorbing state) is to reach a point where only one individual remains, and if we reach zero, we restart. We prove that the expected number of steps until we reach a state where either 1 or 0 agents compete for that resource is $O\left(\frac{1}{p} \log N\right)$.

Moreover, we prove that with high probability, $\Omega\left(\frac{2(1-p)}{2-p}\right)$, only 1 agent will remain (contrary to reaching 0 and restarting the process of claiming the resource), no matter the initial number of agents. Having proven that, we move to the general case of N agents competing for R resources.

At any time, at most N agents can compete for each resource. We call this period a round. During a round, the number of agents competing for a specific resource monotonically decreases, since that resource is perceived as occupied by non-competing agents. Let the round end when either 1 or 0 agents compete for the resource. This will require $O\left(\frac{1}{p} \log N\right)$ steps. If all agents backed-off, it will take on average R steps until at least one of them finds a free resource. We call this period a break. In the worst case, the system will oscillate between a round and a break. According to the above, one oscillation requires in expectation $O\left(\frac{1}{p} \log N + R\right)$ steps. If $R = 1$, as mentioned in the previous paragraph, in expectation there will be $\frac{2-p}{2(1-p)}$ oscillations. For $R > 1$ the expected number of

oscillations is bounded by $O\left(R \frac{2-p}{2(1-p)}\right)$. Thus, we conclude that if all the agents back-off with the same constant probability p , the expected number of steps until the system converges to a complete matching is $O\left(R \frac{2-p}{2(1-p)} \left(\frac{1}{p} \log N + R\right)\right)$.

Next, we drop the constant probability assumption. Intuitively, the worst case scenario corresponds to either all agents having a small back-off probability, thus they keep on competing for the same resource, or all of them having a high back-off probability, thus the process will keep on restarting. These two scenarios correspond to the inner $\left(\frac{1}{p}\right)$ and outer $\left(\frac{2-p}{2(1-p)}\right)$ probability terms of bound (4) respectively. Let $p^* = f(\text{loss}^*)$ be the worst between the smallest or highest back-off probability any agent $n \in \mathcal{N}$ can exhibit, i.e., having loss^* given by Eq. 5. Using p^* instead of the constant p , we bound the expected convergence time according to bound (4). \square

It is worth noting that the back-off probability p^* in bound (4) does not significantly affect the convergence time. For example, using Eq. 2 with a quite small $\epsilon = 0.01$, the resulting quantities would be at most $100R \log N$, and $50R^2$. Most importantly, though, this is a rather loose bound (e.g., agents would rarely back-off with probabilities as extreme as p^*).

Apart from the convergence of the whole system, we are interested in the expected number of steps any individual agent would require in order to acquire a resource. In real-world scenarios, there is typically a cost associated with acquiring a resource. For example, a taxi driver would not be willing to drive to the other end of the city to pick up a low fare passenger. As a result, each agent is typically interested in a subset of the total resources, i.e., $\mathcal{R}^n \subset \mathcal{R}$, thus at each resource there is a bounded number of competing agents. Let R^n denote the maximum number of resources agent n is interested in, and N^r denote the maximum number of agents competing for resource r . By bounding these two quantities (i.e., we consider R^n and N^r to be constant functions of N, R), Corollary 2.2 proves that the expected number of steps any individual agent requires in order to claim a resource is independent of the total problem size (i.e., N , and R), or, in other words, that the convergence time is *constant* in these quantities.

COROLLARY 2.2. *Let $R^n = |\mathcal{R}^n|$, such that $\forall r \in \mathcal{R}^n : u_n(r) > 0$, and $N^r = |\mathcal{N}^r|$, such that $\forall n \in \mathcal{N}^r : u_n(r) > 0$. The expected number of steps until an agent $n \in \mathcal{N}$ following Alg. 1 successfully acquires a resource is bounded by (6), where $p_n^* = f(\text{loss}^*)$ and loss^* is given by Eq. 7, independent of the total problem size N, R .*

$$O\left(\max_{n' \in \cup_{r \in \mathcal{R}^n} \mathcal{N}^r} R^{n'} \frac{2 - p_n^*}{2(1 - p_n^*)} \left(\frac{1}{p_n^*} \log(\max_{r \in \mathcal{R}^n} N^r) + \max_{n' \in \cup_{r \in \mathcal{R}^n} \mathcal{N}^r} R^{n'}\right)\right) \quad (6)$$

$$\text{loss}^* = \arg \min_{\text{loss}_n^r} \left(\min_{r \in \mathcal{R}^n, n \in \mathcal{N}^r} (\text{loss}_n^r), 1 - \max_{r \in \mathcal{R}^n, n \in \mathcal{N}^r} (\text{loss}_n^r) \right) \quad (7)$$

PROOF. The expected number of steps until an agent $n \in \mathcal{N}$ successfully acquires a resource is upper bounded by the total convergence time of the sub-system he belongs to, i.e., the sub-system consisting of the sets of \mathcal{R}^n resources and $\cup_{r \in \mathcal{R}^n} \mathcal{N}^r$ agents. In such scenario, at most $\max_{r \in \mathcal{R}^n} N^r$ agents can compete for any resource. Using Theorem 2.1 for $\max_{r \in \mathcal{R}^n} N^r$ agents, $\max_{n' \in \cup_{r \in \mathcal{R}^n} \mathcal{N}^r} R^{n'}$ resources, and worst-case loss^* given by any agent in $\cup_{r \in \mathcal{R}^n} \mathcal{N}^r$ (i.e.,

Eq 7) results in the desired bound. Note that agents do not compete for already claimed resources (step 8 of Alg. 1), thus the convergence of an agent does not require the convergence of agents of overlapping sub-systems. \square

3 EVALUATION

In this section we evaluate ALMA under various test cases. For the first two, we focus on convergence time and relative difference in social welfare (SW), i.e., $(\text{achieved} - \text{optimal})/\text{optimal}$. In every reported metric, except for the social welfare, we report the average value out of 128 runs of the same problem instance. Error bars represent one standard deviation (SD) of uncertainty. As a measure of social welfare, we report the cumulative regret of the aforementioned 128 runs, i.e., for $i \in [1, 128]$ runs, the reported relative difference in social welfare is $(\sum_i \text{achieved} - \sum_i \text{optimal})/\sum_i \text{optimal}$. This was done to improve visualization of the results in smaller problem sizes, where really small differences result in high SD bars (e.g., if $\text{achieved} = 1 \times 10^{-5}$, and $\text{optimal} = 2 \times 10^{-5}$, the relative difference would be -50% for practically the same matching). The optimal matchings were computed using the Hungarian algorithm¹. The third test case is an on-line setting, thus we report the achieved SW (not the relative difference to the optimal), and the empirical competitive ratio (average out of 128 runs, as before). All the simulations were run on 2x Intel Xeon E5-2680 with 256 GB RAM. In Section 3.1 we use the logistic function (Eq. 3) with $\gamma = 2$, while in Sections 3.2 & 3.3 we use the linear function (Eq. 2) with $\epsilon = 0.1$.

It is important to stress that our goal is not to improve the convergence speed of a centralized, or decentralized algorithm. Rather, the computation time comparisons of Sections 3.1 & 3.2 are meant to ground the actual speed of ALMA, and argue in favor of its applicability on large-scale, real-world scenarios. Given the nature of the problem, we elected to use a specialized algorithm to compute the optimal solution, rather than a general LP-based technique (e.g., the Simplex method). Specifically, we opted to use the Hungarian algorithm which, first, has proven polynomial worst case bound, and second, as our simulations will demonstrate, can handle sufficiently large problems.

3.1 Test Case #1: Uniform, and Noisy Common Preferences

3.1.1 Setting. As a first test case, we cover the extreme scenarios. The first pertains to an anti-coordination scenario in which agents with similar preferences compete over the same set of actions. We call this ‘noisy common preferences’ and model the utilities as: $\forall n, n' \in \mathcal{N}, |u_n(r) - u_{n'}(r)| \leq \text{noise}$, where the noise is sampled from a zero-mean Gaussian distribution, i.e., $\text{noise} \sim \mathcal{N}(0, \sigma^2)$. In the second scenario the utilities are initialized uniformly at random for each agent and resource.

3.1.2 Results. Please see the full version [11] for a detailed account of the results (including comparisons with the greedy and random algorithms), along with the relevant graphs. The convergence time for the system of agents is linear to the number of resources R . ALMA requires approximately 4 to 6 orders of magnitude less computation time, and scales more gracefully, than the

¹We used Kevin L. Stern’s $O(N^3)$ implementation: <https://github.com/KevinStern/>.

centralized Hungarian algorithm. Note also that in real-world applications we would have to take into account communication time, communication reliability protocols, etc., which create additional overhead for the Hungarian or any other algorithm for the assignment problem. The relative difference in social welfare reaches asymptotically zero as R increases. For a small number of resources, ALMA achieves the worst social welfare, approximately 11% worse than the optimal. Intuitively this is because when we have a small number of choices, a single wrong matching can have a significant impact to the final social welfare, while as the number of resources grow, the impact of an erroneous matching is mitigated. For 16384 resources we lose less than 2.5% of the optimal.

3.2 Test Case #2: Resource Allocation in a Cartesian Map with Manhattan Distances

Adopting a simple rule allows the applicability of ALMA to large scale multi-agent systems. In this section we will analyze such a scenario. Specifically we are interested in resource allocation in urban environments (e.g., parking spots / charging stations for autonomous vehicles, taxi - passenger matchings, etc.). The aforementioned problems become ever more relevant due to rapid urbanization, and the natural lack of coordination in the usage of resources [32]. The latter result in the degradation of response (e.g., waiting time) and quality metrics in large cities [32].

3.2.1 Setting. Let us consider a Cartesian map representing a city on which are randomly distributed vehicles and charging stations, as depicted in Fig. 1h. The utility received by a vehicle n for using a charging station r is proportional to the inverse of their distance, i.e., $u_n(r) = 1/d_{n,r}$. Since we are in an urban environment, let $d_{n,r}$ denote the Manhattan distance. Typically, there is a cost each agent is willing to pay to drive to a resource, thus there is a cut-off distance, upon which the utility of acquiring the resource is zero (or possibly negative). This is a typical scenario encountered in resource allocation in urban environments, where there are spatial constraints and local interactions.

The way such problems are typically tackled, is by dividing the map to sub-regions, and solving each individual sub-problem. For example, Singapore is divided into 83 zones based on postal codes [8], and taxi drivers' policies are optimized according to those [24, 33]. On the other hand, not placing bounds means that the current solutions will not scale. To the best of our knowledge, we are the first to propose an anytime heuristic for resource allocation in urban environments that can scale in constant time without the need to artificially split the problem. Instead, ALMA exploits the two typical characteristics of an urban environment: the anonymity in interactions and homogeneity in supply and demand [32] (e.g., assigning any of two equidistant charging stations to a vehicle, would typically result to the same utility). This results in a simple learning rule which, as we will demonstrate in this section, can scale to hundreds of thousands of agents.

3.2.2 Convergence Time. To demonstrate the latter, we placed a bound on the maximum number of resources each agent is interested in, and on the maximum number of agents competing for a resource, specifically $R^n = N^r \in \{8, 16, 32, 64, 128\}$. According to Corollary 2.2, bounding these two quantities should result in

convergence in constant time, regardless of the total problem size (R, N). The latter is corroborated by Fig. 1a, which shows that the average number of time-steps until an agent successfully claims a resource remains constant as we increase the total problem size. Same is true for the system's convergence time (Fig. 1b), which caps as R increases. The small increase is due to outliers, as Fig. 1b reports the convergence time of the last agent. This results to approximately 7 orders of magnitude less computation time than the centralized Hungarian algorithm (Fig. 1c), and this number would grow boundlessly as we increase the total problem size. Moreover, as mentioned, in an actual implementation, any algorithm for the assignment problem would face additional overhead due to communication time, reliability protocols, etc.

3.2.3 Efficiency. Along with the constant convergence time, ALMA is able to reach high quality matchings, achieving less than 7.5% worse social welfare (SW) than the optimal (Fig. 1d). The latter refers to the small bound of $R^n = 8$. As observed in Section 3.1, with a small number of choices, a single wrong matching can have a significant impact to the final social welfare. By increasing the bound to a more realistic number (e.g., $R^n = 32$), we achieve less than 2.5% worse SW. In general, for $R > 2$ resources and different values of R^n , ALMA achieves between 1.9 – 12% loss in SW, while the greedy approach achieves 8.0 – 24% and the random 7.3 – 43.4%. The behavior of the graphs depicted in Fig. 1d for $R^n \in \{8, 16, 32, 64, 128\}$ indicate that, as the problem size (R) increases, the social welfare reaches its lowest value at $R = 2 \times R^n$. To investigate the latter, we have included a graph for increasing R^n (instead of constant to the problem size), specifically $R^n = R/2$. ALMA achieves a constant loss in social welfare (approximately 11%). The greedy approach achieves loss of 14%, while the random solution degrades towards 44% loss.

Compared to Test Case #1, this is a significantly harder problem for a decentralized algorithm with no communication and no global knowledge of the resources. The set of resources each agent is interested in is a proper subset of the set of the total resources, i.e., $\mathcal{R}^n \subsetneq \mathcal{R}$ (or could be $R < N$). Furthermore, the lack of communication between the participants, and the stochastic nature of the algorithm can lead to deadlocks, e.g., in Fig. 1h, if vehicle 2 acquires resource 1, then vehicle 1 does not have an available resource in range. Nonetheless, ALMA results in an almost complete matching. Fig. 1e, depicts the percentage of 'winners' (i.e., agents that have successfully claimed a resource r such that $u_n(r) > 0$). The aforementioned percentage refers to the total population (N) and not the maximum possible matchings (potentially $< N$). As depicted, the percentage of 'winners' is more than 90%, reaching up to 97.8% for $R^n = 128$. We also employed ALMA in larger simulations with up to 131072 agents, and equal resources. As seen in Fig. 1e, the percentage of winners remains stable at around 98%. Even though the size of the problem prohibited us from running the Hungarian algorithm (or an out-of-the-box LP solver) and validating the quality of the achieved matching, the fact that the percentage of winners remains the same suggests that the relative difference in SW will continue on the same trend as in Fig. 1d. Moreover, the average steps per agent to claim a resource remains, as proven, constant (Fig. 1a). The latter validate the applicability of ALMA in large scale applications with hundreds of thousands of agents.

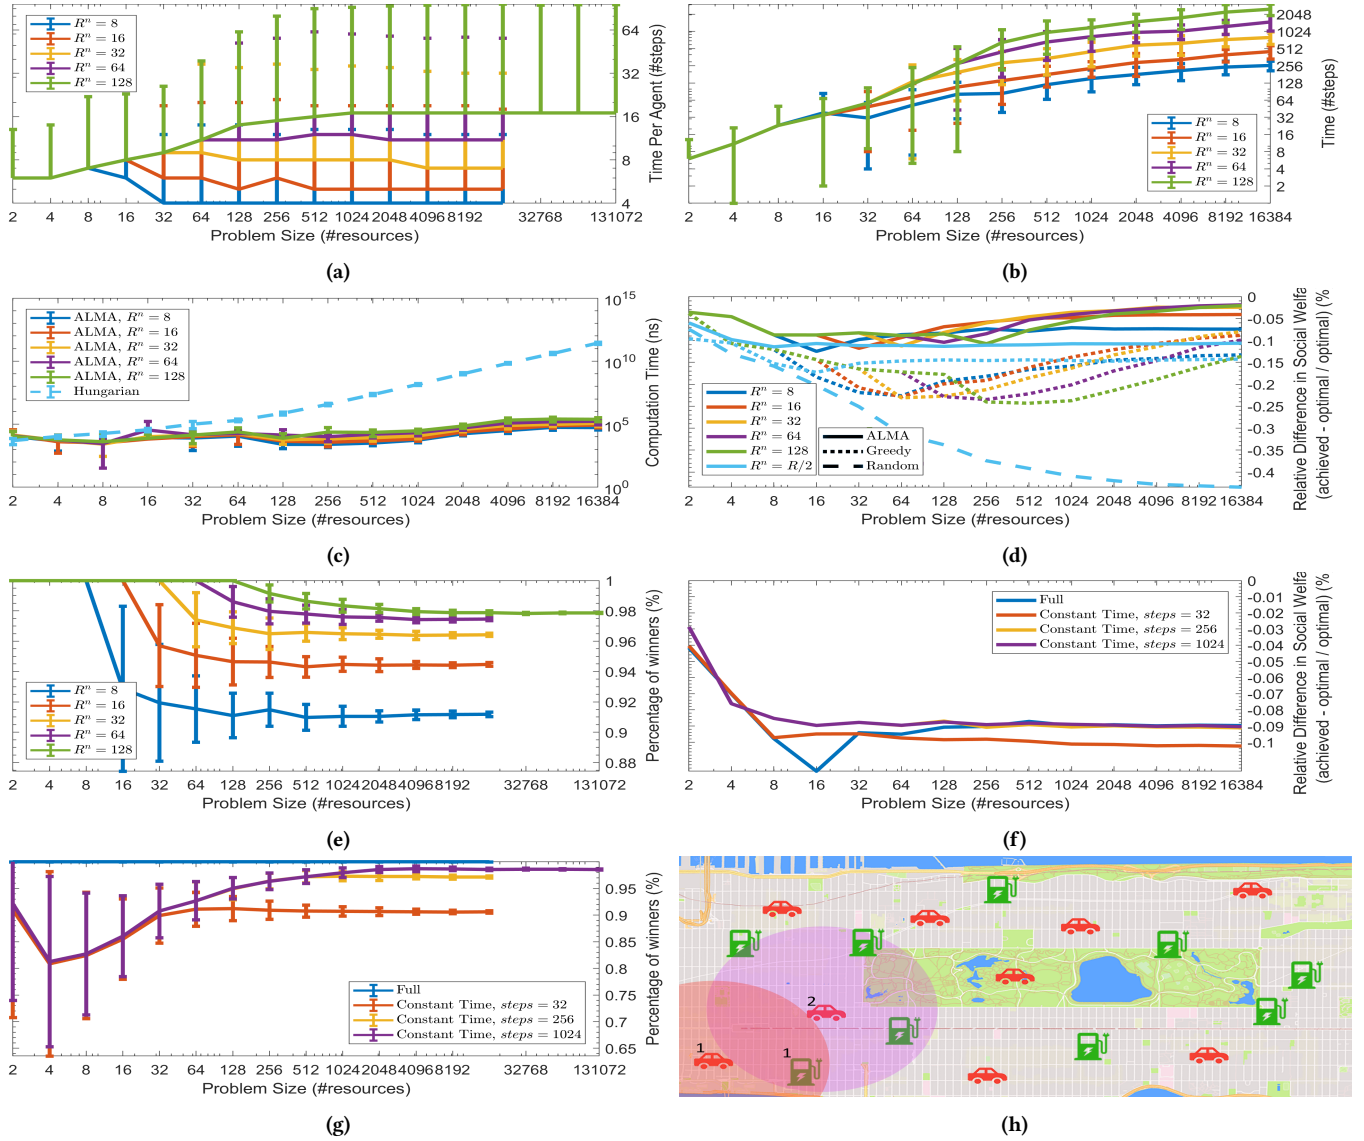


Figure 1: Left to right, top to bottom: (1a) Average time (#steps) for an agent to acquire a resource, (1b) Total convergence time (#steps), (1c) Computation time (ns), (1d) Relative difference in SW (%), (1e) Percentage of ‘winners’, (1f) Relative difference in SW (%) in interrupted execution, (1g) Percentage of ‘winners’ in interrupted execution. The aforementioned are for increasing number of resources, and $N = R$. Fig. (1a), (1b), and (1c) are in double log scale, while the rest are in single log scale. Fig. (1h) presents an example of the studied resource allocation scenario in an urban environment. We assume grid length of $\sqrt{4 \times N}$.

3.2.4 Anytime Property. In the real world, agents are required to run in real time, which imposes time constraints. ALMA can be used as an anytime heuristic as well. To demonstrate the latter, we compare four configurations: the ‘full’ one, which is allowed to run until the systems converges, and three ‘constant time’ versions which are given a time budget of 32, 256, and 1024 time-steps. In this scenario, we do not impose a bound on R^n , N^r , but we assume a cut-off distance, upon which the utility is zero. The cut-off distance was set to 0.25 of the maximum possible distance, i.e., as the problem size grows, so do the R^n , N^r . On par with Test Case #1, the full

version converges in linear time. As depicted in Fig. 1f, the achieved SW is less than 9% worse than the optimal. The inferior results in terms of SW compared to Test Case #1 are because this is a significantly harder problem due to the aforementioned deadlocks. On the other hand though, ALMA benefits from the spatial constraints of the problem. The average number of time-steps an individual agent needs to successfully claim a resource is significantly smaller, which suggest that we can enforce computation time constraints. Restricting to only 32, 256, and 1024 time-steps, results in 1.25%, 0.12%, and 0.03% worse SW than the unrestricted version, respectively.

Even in larger simulations with up to 131072 agents, the percentage of winners (Fig. 1g) remains stable at 98.6%, which suggests that the relative difference in SW will continue on the same trend as in Fig. 1f (we do not suggest that this is the case in any domain. For example, in the noisy common preferences domain of Test Case #1, the quality of the achieved matching decreases boundlessly as we decrease the allotted time. Nevertheless, the afordescribed domain is a realistic one, with a variety of real-world applications). Finally, the repeated nature of such problems suggests that even in the case of a deadlock, the agent which failed to win a resource, will do so in some subsequent round.

3.3 Test Case #3: On-line Taxi Request Match

In this section we present a motivating test case involving ride-sharing, via on-line taxi request matching, using *real* data of taxi rides in New York City. Ride-sharing (or carpooling), offers great potential in congestion relief and environmental sustainability in urban environments. In the past few years, several commercially successful ride-sharing companies have surfaced (e.g., Uber, Lyft, etc.), giving rise to a new incarnation of ride-sharing: dynamic ride-sharing, where passengers are matched in real-time. Ride-sharing, though, results to some passenger disruption due to loss in flexibility, security concerns, etc. Compensation comes in the form of monetary incentives, as it allows passengers to share the travel expenses, and thus reduce the cost. Ride-sharing companies account for a plethora of factors, like current demand, prime time pricing, the cost of the route without any ride-sharing, the likelihood of a match, etc. Yet, a fundamental factor of the cost of any shared ride, no matter if it is a company or a locally-organized car sharing scheme, is the traveled distance.

In this test case, we attempt to maximize the total distanced saved, by matching taxi requests of high overlap. Fig. 2 provides an illustrative example. There are two passengers (depicted as yellow and red) with high overlap routes. Each can drive on their own to their respective destinations (dashed yellow and red line respectively), or share a ride (green line) and reduce travel costs.

Dynamic ride-sharing is an inherently *on-line* setting, as a matching algorithm is unaware of the requests that will appear in the future and needs to make decisions for the requests before they ‘expire’ (a similar setting was studied in [2]). ALMA is highly befitting for such a scenario, as it involves large-scale matchings under dynamic demand, it is highly decentralized, and partially observable.

3.3.1 Setting. We use a dataset ² of all taxi requests (ρ) in New York City during one week ({01-01-16 0:00 - 01-07-16 23:59}, 34077 requests in total). The data include pickup and drop-off times, and geolocations. Requests appear (become open) at their respective pickup time, and wait k_ρ time-steps to find a match. Let a time-step be one minute. After k_ρ time-steps we call request ρ , critical. If a critical request is not matched, we assume they drive off to their destination in a single passenger ride. Let *open*, *critical* denote the sets of open, and critical requests respectively, and let *current* = *open* \cup *critical*. To compute k_ρ we assume the following: There is a minimum *minW*, and a maximum *maxW* waiting time set by the

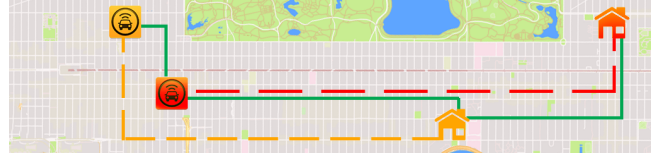


Figure 2: Example of the studied taxi request matching scenario.

ride-sharing company, i.e., $\min W \leq k_\rho \leq \max W, \forall \rho$. Moreover, since each passenger specifies his destination in advance, we can compute the trip time (l_ρ). Assuming people are willing to wait for a time that is proportional to their trip time, let $k_\rho = q \times l_\rho$, where $q \in [0, 1]$. The parameters *minW*, *maxW*, and q can be set by the ride-sharing company. We report results on different values for all of the above parameters. For each pair ρ_1, ρ_2 of requests, we compute the driving distance ($d_{\rho_1, \rho_2} = \min$ of all possible combinations of driving between ρ_1, ρ_2 ’s pickup and drop-off locations) that would be traveled if ρ_1, ρ_2 are matched, i.e., if they share the same taxi. Subsequently, the utility of matching ρ_1 to ρ_2 (distance saved) is $u_{\rho_1}(\rho_2) = d_{\rho_1, \rho_2}$ (km).

Given the on-line nature of the setting, it might be beneficial to use the following *non-myopic heuristic*: avoid matching low utility pairs, as long as the requests are not critical, since more valuable pairs might be presented in the future. Thus, if $u_{\rho_1}(\rho_2) < d_{\min}$, and $\rho_1, \rho_2 \notin \text{critical}$, we do not match ρ_1, ρ_2 . In what follows, we select for each algorithm and for each simulation the value $d_{\min} \in \{0, 500, 1000, 1500, 2000, 2500\}$ that results in the highest score. To compute the actual trip time, and driving distance, we have used the Open Source Routing Machine ³, which computes shortest paths in road networks.

3.3.2 Benchmarks. Each request runs ALMA independently. ALMA waits until the request becomes critical, and then matches it by running Alg. 1, where $\mathcal{N} = \text{critical}$, and $\mathcal{R} = \text{current}$. In this non-bipartite scenario, if an agent is matched under his dual role as a resource, he is immediately removed. As we explained earlier, it is infeasible for an on-line algorithm to compute the optimal matching over the whole period of time. Instead, we consider *just-in-time* and *in batches* optimal solutions. Specifically, we compare to the following [1, 2]:

- **Just-in-time Max Weight Matching (JiTMWM):** Waits until a request becomes critical and then computes a maximum-weight matching of all the current requests, i.e $\mathcal{N} = \mathcal{R} = \text{current}$.
- **Batching Max Weight Matching (BMWM):** Waits x time-steps and then computes a maximum-weight matching of all the current requests, i.e $\mathcal{N} = \mathcal{R} = \text{current}$.
- **Batching Greedy (BG):** Waits x time-steps and then greedily matches current requests, i.e $\mathcal{N} = \mathcal{R} = \text{current}$ (ties are broken randomly). Unmatched open requests are removed. For batch size $x = 1$ we get the simple greedy approach where requests are matched as soon as they appear.

We have also computed the off-line optimal matching (see [11]) and report the *empirical competitive ratio*, i.e., the ratio of the social

²<https://www.kaggle.com/debanjanpaul/new-york-city-taxi-trip-distance-matrix/>

³<http://project-osrm.org/>

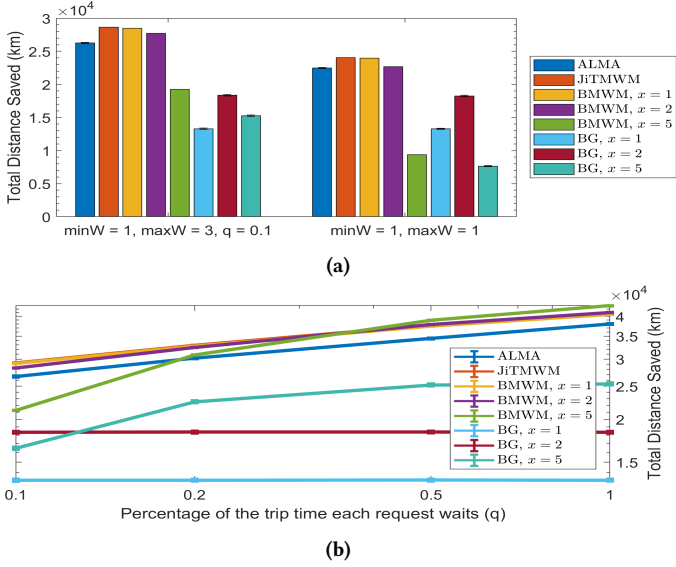


Figure 3: Total distance saved (km) for various values of $\min W$, $\max W$, q . (3a, left) Pragmatic scenario, (3a, right) Requests become critical in just one time-step, (3b) Various levels of waiting time (q) assuming no bounds, i.e., $\min W = 0$, $\max W = \infty$, (double log. scale).

welfare of the on-line algorithm over the welfare of the optimal, as is common in the literature of competitive analysis [5].

3.3.3 Efficiency. Starting with the social welfare, Fig. 3 presents the total distance saved (km) for various values of $\min W$, $\max W$, and q . ALMA loses 8.3% of SW in the pragmatic scenario of Fig. 3a (left), and 6.5% when the requests become critical in just one time-step (Fig. 3a (right)). If no bounds are placed on the minimum and maximum waiting time (i.e., $\min W = 0$, $\max W = \infty$), ALMA exhibits loss of 8 – 11.5%, for various values of q (Fig. 3b). The above are compared to the best performing benchmark on each scenario (JiTMWM, or BMWWM, $x = 5$). Moreover, it significantly outperforms every greedy approach. In the first scenario the BGs lose between 35.8 – 53.5%, in the second between 24 – 68.2%, and in the third between 31.5 – 69%.

Once more, it is worth noting that ALMA requires just a broadcast of a single bit to indicate the occupancy of a resource, while the compared approaches require either *message exchange* for sharing the utility table, or the use of a *centralized authority*. For example, the greedy solution would require message exchange to communicate users’ preferences and resolve collisions in a decentralized setting, and every batching approach would require a common centralized synchronization clock.

Table 1 presents the empirical competitive ratio for the first day of the week of the employed dataset. As we can see, even in the extreme, unlikely scenarios where we assume that people would be willing to wait for more than 10 or 20 minutes (which correspond to large values of q), ALMA achieves high relative efficiency, compared to the off-line (infeasible) benchmark. These scenarios are favorable to the off-line optimal, because requests stay longer in the system

Table 1: The Empirical Competitive Ratio of ALMA, for different values of $(\min W, \max W, q)$. The full table for all the employed on-line heuristics can be found in [11].

$(1, 3, 0.1)$	$(1, 1, -)$	$(0, \infty, 0.1)$	$(0, \infty, 0.2)$	$(0, \infty, 0.5)$	$(0, \infty, 1.0)$
0.79	0.85	0.78	0.75	0.69	0.67

and therefore the algorithm takes more advantage of its foreseeing capabilities (hence, the drop in the competitive ratios is observed in *all* of the employed algorithms). In particular, ALMA achieves an empirical competitive ratio of 0.67 for $q = 1$ and even better ratios for more realistic scenarios (as large as 0.85). The just-in-time and batch versions of the maximum weight matching perform slightly better, but this is to be expected, as they compute the maximum weight matching on the graphs of the current requests. In spite of the unpredictability of the on-line setting, and the dynamic nature of the demand, ALMA is consistently able to exhibit high performance, in all of the employed scenarios.

4 CONCLUSION

Algorithms for solving the assignment problem, whether centralized or distributed, have runtime that increases with the total problem size, even if agents are interested in a small number of resources. Thus, they can only handle problems of some bounded size. Moreover, they require a significant amount of inter-agent communication. Humans on the other hand are routinely called upon to coordinate in large scale in their everyday lives, and are able to fast and robustly match with resources under dynamic and unpredictable demand. One driving factor that facilitates human co-operation is the principle of altruism. Inspired by human behavior, we have introduced a novel anytime heuristic (ALMA) for weighted matching in both bipartite and non-bipartite graphs. ALMA is decentralized and requires agents to only receive partial feedback of success or failure in acquiring a resource. Furthermore, the running time of the heuristic is constant in the total problem size, under reasonable assumptions on the preference domain of the agents. As autonomous agents proliferate (e.g., IoT devices, intelligent infrastructure, autonomous vehicles, etc.), having robust algorithms that can scale to hundreds of thousands of agents is of utmost importance.

The presented results provide an empirical proof of the high quality of the achieved solution in a variety of scenarios, including both synthetic and *real* data, time constraints and on-line settings. Furthermore, both the proven theoretical bound (which guarantees constant convergence time), and the computation speed comparison (which grounds ALMA to a proven fast centralized algorithm), argue for its applicability to large scale, real world problems. As future work, it would be interesting to identify meaningful domains in which ALMA can provide provable worst-case performance guarantees, as well as to empirically evaluate its performance on other real datasets, corresponding to important real-world, large-scale problems.

REFERENCES

- [1] Niels Agatz, Alan L Erera, Martin WP Savelsbergh, and Xing Wang. 2011. Dynamic ride-sharing: A simulation study in metro Atlanta. *Procedia-Social and Behavioral Sciences* 17 (2011), 532–550.
- [2] Itai Ashlagi, Maximilien Burq, Chinmoy Dutta, Patrick Jaillet, Amin Saberi, and Chris Sholley. 2018. Maximum Weight Online Matching with Deadlines. *arXiv preprint arXiv:1808.03526* (2018).
- [3] DP Bertsekas. 1979. A distributed algorithm for the assignment problem. *Laboratory for Information and Decision Systems Working Paper, Massachusetts Institute of Technology, Cambridge, MA* (1979).
- [4] CW Borchardt and CGJ Jacobi. 1865. De investigando ordine systematis aequationum differentialium vulgarium cujuscunque. *Journal für die reine und angewandte Mathematik* 64 (1865), 297–320.
- [5] Allan Borodin and Ran El-Yaniv. 2005. *Online computation and competitive analysis*. Cambridge university press.
- [6] Mathias Bürger, Giuseppe Notarstefano, Francesco Bullo, and Frank Allgöwer. 2012. A distributed simplex algorithm for degenerate linear programs and multi-agent assignments. *Automatica* 48, 9 (2012), 2298–2304.
- [7] Gary Charness and Matthew Rabin. 2002. Understanding Social Preferences with Simple Tests*. *The Quarterly Journal of Economics* 117, 3 (2002), 817–869. <https://doi.org/10.1162/003355302760193904>
- [8] S. F. Cheng and T. D. Nguyen. 2011. TaxiSim: A Multiagent Simulation Platform for Evaluating Taxi Fleet Operations. In *2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, Vol. 2. 14–21. <https://doi.org/10.1109/WI-IAT.2011.138>
- [9] Ludek Cigler and Boi Faltings. 2011. Reaching correlated equilibria through multi-agent learning. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 509–516.
- [10] Panayiotis Danassis and Boi Faltings. 2019. Courtesy as a Means to Coordinate. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems.
- [11] Panayiotis Danassis, Aris Filos-Ratsikas, and Boi Faltings. 2019. Anytime Heuristic for Weighted Matching Through Altruism-Inspired Behavior. *arXiv:1902.09359* (2019). <https://arxiv.org/abs/1902.09359>
- [12] George B Dantzig. 1990. *Origins of the simplex method*. ACM.
- [13] Jack Edmonds. 1965. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of research of the National Bureau of Standards B* 69, 125-130 (1965), 55–56.
- [14] Jack Edmonds and Richard M. Karp. 1972. Theoretical Improvements in Algorithmic Efficiency for Network Flow Problems. *J. ACM* 19, 2 (April 1972), 248–264. <https://doi.org/10.1145/321694.321699>
- [15] Yanfeng Geng and Christos G Cassandras. 2013. New “smart parking” system based on resource allocation and reservations. *IEEE Transactions on Intelligent Transportation Systems* 14, 3 (2013), 1129–1139.
- [16] Herbert Gintis. 2000. Strong reciprocity and human sociality. *Journal of theoretical biology* 206, 2 (2000), 169–179.
- [17] Stefano Giordani, Marin Lujak, and Francesco Martinelli. 2010. A distributed algorithm for the multi-robot task allocation problem. In *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*. Springer, 721–730.
- [18] Tyler Gunn and John Anderson. 2013. Dynamic Heterogeneous Team Formation for Robotic Urban Search and Rescue. *Procedia Computer Science* 19 (2013), 22 – 31. <https://doi.org/10.1016/j.procs.2013.06.009> The 4th International Conference on Ambient Systems, Networks and Technologies (ANT 2013), the 3rd International Conference on Sustainable Energy Information Technology (SEIT-2013).
- [19] S. Ismail and L. Sun. 2017. Decentralized hungarian-based approach for fast and scalable task allocation. In *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*. 23–28. <https://doi.org/10.1109/ICUAS.2017.7991447>
- [20] Harold W Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics (NRL)* 2, 1-2 (1955), 83–97.
- [21] David Lewis. 2008. *Convention: A philosophical study*. John Wiley & Sons.
- [22] László Lovász and Michael D Plummer. 2009. *Matching theory*. Vol. 367. American Mathematical Soc.
- [23] James Munkres. 1957. Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics* 5, 1 (1957), 32–38.
- [24] Duc Thien Nguyen, Akshat Kumar, and Hoong Chuin Lau. 2017. Collective Multiagent Sequential Decision Making Under Uncertainty. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [25] Martin A Nowak and Karl Sigmund. 2005. Evolution of indirect reciprocity. *Nature* 437, 7063 (2005), 1291.
- [26] François Ollivier. 2009. Looking for the order of a system of arbitrary ordinary differential equations. *Applicable Algebra in Engineering, Communication and Computing* 20, 1 (2009), 7–32.
- [27] Christos H. Papadimitriou and Kenneth Steiglitz. 1982. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- [28] Jay Simon. 2016. On the existence of altruistic value and utility functions. *Theory and Decision* 81, 3 (2016), 371–391.
- [29] Peter Stone, Gal A. Kaminka, Sarit Kraus, and Jeffrey S. Rosenschein. 2010. Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination. In *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence*.
- [30] Hsin-Hao Su. 2015. Algorithms for Fundamental Problems in Computer Networks. (2015).
- [31] Mohammad Sadegh Talebi. 2013. Uncoupled Learning Rules for Seeking Equilibria in Repeated Plays: An Overview. *CoRR abs/1310.5660* (2013). <http://arxiv.org/abs/1310.5660>
- [32] Pradeep Varakantham. 2016. Sequential decision making for improving efficiency in urban environments. (2016).
- [33] Pradeep Reddy Varakantham, Shih-Fen Cheng, Geoff Gordon, and Asrar Ahmed. 2012. Decision support for agent populations in uncertain and congested environments. (2012).
- [34] Michael M Zavlanos, Leonid Spesivtsev, and George J Pappas. 2008. A distributed auction algorithm for the assignment problem. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*. IEEE, 1212–1217.