

Balancing Deployment Costs in Multi-Robot Task Assignment

Nils Wilde, and Javier Alonso-Mora

Abstract—Multi-Robot Task Assignment (MRTA) studies the problem of allocating spatially distributed tasks to a fleet of cooperative robots as well as determining the optimal task sequence for each robot. Common objectives include minimizing the task waiting times and robot tour lengths or maximizing the number of serviced tasks within given time windows. However, this does not consider an equitable distribution of the workload among the fleet. Uneven workloads are often undesirable when few robots may service most tasks while parts of the fleet remain underused. On the other hand, fully balanced workloads may insufficiently consider the total operation cost and thus can deploy robots in a redundant manner. In this paper, we study MRTA from the viewpoint of multi-objective optimization (MOO), formulating the problem of simultaneously minimizing the costs of all individual robot tours. This treatment allows for attaining more balanced solutions than common formulations using the sum or maximum of tour costs. We present a generalist formulation using a scalar objective and establish theoretical guarantees on the attainable trade-offs. Further, we derive an effective heuristic based on a p -norm of tour lengths that is able to find balanced workloads among robots. Our approach is agnostic to the specific choice of MRTA solver and we show how it can be incorporated into two state-of-the-art algorithms. We demonstrate our approach in experiments for offline and online MRTA setups, including servicing tasks as well as pickup and delivery, and highlight its advantages compared to state-of-the-art formulations.

I. INTRODUCTION

Multi-Robot Task Assignment (MRTA) is a fundamental challenge in the deployment of multi-robot systems. Collaboratively, autonomous mobile systems are able to dynamically schedule and plan the completion of complex real-world tasks in applications such as autonomous mobility on demand [1]–[4], pickup and delivery [5]–[7], service and personal care [7] and environmental monitoring [8], [9].

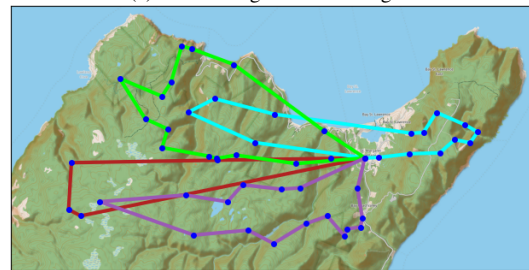
We study the problem of balanced workloads in MRTA, *i.e.*, finding efficient solutions that attain fair distribution of operation cost or tour length. MRTA consists of two core components: task assignment – the decision of which robot is servicing which task, and routing – the tour for each robot in the fleet. Task assignment algorithms are commonly designed to achieve either i) efficient routing *e.g.*, minimizing the total tour length for all robots, ii) efficient service, *e.g.*, minimizing the wait time of tasks or the makespan of the mission, iii)

This research is partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) and from the European Union through ERC, INTERACT, under Grant 101041863. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.


N. Wilde is with the Faculty of Computer Science at Dalhousie University, Halifax, Canada. J. Alonso-Mora is with the Department for Cognitive Robotics, Delft University of Technology, Delft, Netherlands, nils.wilde@dal.ca, j.alonsomora@tudelft.nl.



(a) Minimizing total tour length.



(b) Minimizing maximum tour lengths.



(c) Minimizing a p -norm of tour lengths for $p = 2$.

Fig. 1: Example tours for an environmental monitoring mission with four robots and time windows under different objective functions.

effective service under time constraints such as throughput or maximizing the number of tasks serviced on time [10], or iv) a balanced performance with respect to multiple objectives [7]. However, workload equitability and fairness has found less attention in the literature. Common objectives such as the sum of tour lengths [11] can lead to highly imbalanced tours where few robots service many tasks while others are underutilized. Minimizing the maximum tour length can alleviate this effect, yet may not be resourceful. We illustrate this in Figure 1. The problem is offline task assignment for four robots where all locations must be visited once within a relatively long deadline. As shown in 1a, minimizing the sum of tour length leads to a solution where only two robots are deployed. Minimizing the max tour length (1b) makes use of the full fleet and avoids the imbalance in the workload; however, the red and purple robot have largely redundant tours, making the solution less resourceful. Thus, both approaches are not ideal in practical applications.

Therefore, we study a novel formulation for MRTA. We observe that the problem of minimizing the tour length for a fleet of robots implicitly poses a multi-objective optimization (MOO) problem: the goal is to simultaneously minimize the tour length for *all* robots. We propose a p -norm of tour lengths as an effective heuristic to the MOO formulation. This is highlighted in Figure 1c. Here, all four robots are deployed and find an effective balance of workloads. This reduces the maximum tour length by 42% compared to min-sum in 1a, and the total travel distance by 27% compared to min-max in 1b.

Our approach generalizes the common sum and max of tour lengths and is thus able to compute a wider range of trade-offs of deployment costs. We provide theoretical insights into the expressiveness of the approach, and show-case how the formulation can be adapted in two state-of-the-art task assignment algorithms without increasing their computational complexity. Through simulation experiments, we showcase that the p -norm objective is highly effective in offline MRTA as well as online multi-robot pickup and delivery (MRPD), achieving balanced trade-offs.

a) Related Work: Dynamic Vehicle Routing (DVR) and Multi-robot task assignment (MRTA) are fundamental challenges in the deployment of fleets of mobile robots with active research spanning over multiple decades [12]–[14] until today [8], [15], [16].

With recent advancements in the design of scalable algorithms for large fleets [2], [3] and distributed optimization [11], [17], [18], there is an increasing interest in considering fairness and equitability for multi-robot deployment [8], [19]. Closely related to our work, the authors of [19] study the problem of equitable workloads in the multi Traveling Salesman Problem (mTSP), an offline variant of MRTA. They propose an optimal algorithm of a fairness variant of mTSP as well as a discussion on the relationship of their approach to min-sum, min-max and p -norm objectives. However, in MRTA one often faces additional constraints such as time-windows or partial orderings for pickup and delivery such that optimal solutions where the fairness approach cannot directly be adapted. Therefore, we study how the p -norm objective can be used as an effective heuristic for complex MRTA variants including online settings. Equitable task service is considered in [8] for DVR with moderate workloads. Using a p -norm objective, the maximum waiting times of tasks can be significantly reduced. Our work is closely related to this: we also propose a p -norm objective for MRTA; yet we seek to balance the workload of robots instead of task service times. Further, we address a generalized multi-objective formulation of MRTA, and show how this can be applied to complex variants such as MRPD. The authors of [20] use a squared distance – equivalent to a 2-norm – for MRTA. In our work, we provide a theoretical analysis on how cost functions between robots can lead to multi-objective trade-offs, together with a fundamental Pareto-complete solution, and the p -norm as an effective heuristic. The work of [21] considers DVR with multiple vehicles.

Using *equitable partitioning* the workspace is divided into regions of similar workload, tackling the multi-robot case via parallelization of single vehicle deployment. However, partitioning approaches do not generalize to problem variants such as MRPD, since deliveries may require robots to leave initially assigned partitions.

Practical variants of MRTA are often required to consider multiple objectives. Among the most common objective are the sum of costs among different robots [11], [17], [18], the maximum or makespan of costs [22] and the waiting time of tasks [8]. The authors of [23] explicitly consider simultaneously minimizing the average and maximum cost of robot tours. We show that our generalized MRTA formulation contains this case while the proposed p -norm heuristic is known to balance these cost functions [22]. Trade-offs between task completion time, priorities and energy consumption are considered in a multi-objective formulation in [24], while [7] provides an explicit treatment of multi-objective MRTA for any cost function. While our work is also using a multi-objective formulation, we actually do not consider different cost functions, but instead trade-offs of the same cost function, *e.g.*, tour length, among different robots.

b) Contributions: We make four contributions: i) We present a multi-objective formulation for MRTA to capture a generalized notion of the problem where the operational cost, *e.g.*, tour length of robots, is to be optimized. ii) We provide theoretical insights into how p -norms are a suitable to balance tour lengths for all robots in the fleet, and implicitly trades-off the maximum and total workload. iii) We show how the proposed formulation can be integrated in two different state-of-the-art MRTA solvers. iv) In simulation experiments for offline and online MRTA, we demonstrate that the proposed formulation is highly competitive on several measures including sum of tour cost, maximum tour cost and equitability of workloads, alleviating complementary weaknesses of common baselines.

II. PROBLEM STATEMENT

We consider an MRTA problem: Given an environment encoded as a weighted graph $G = (V, E, d)$, a fleet of m robots needs to service a set of N tasks $\mathcal{T} = \{T_1, \dots, T_N\}$ where each task requires one of the robots to travel to one or more sites, *i.e.*, vertices $s_1, s_2, \dots \in V$, in a fixed order. This formulation is general, it encapsulates simple tasks that require a robot to only visit one location, as well as more complex settings such as pickup and delivery. Further, each location of some task T_j is associated with a deadline t_j^d . We consider a fleet of m robots, and denote the set of robot indices by \mathcal{R} . Let $\pi(i, \mathcal{T}_i)$ denote the tour of the i -th robot servicing tasks $\mathcal{T}_i \subseteq \mathcal{T}$, and let $a(k, T_j, \pi(i, \mathcal{T}_i))$ denote the time robot i arrives on site s_k of task T_j . Finally, a cost function $c(\pi)$ captures the length of each robot's tour.

Commonly, MRTA problems seek to minimize the total length of all robot tours, subject to all tasks being serviced on time by exactly one robot [11]. However, this leads to imbalanced assignments where some robots receive high

workload while others may remain underused or even idle. Thus, we introduce a more general treatment of task assignment where we formulate a multi-objective problem that seeks to minimize the cost of all robot tours simultaneously.

Problem 1 (Generalized MRTA (GMRTA)). Given a graph G , tasks \mathcal{T} , and a fleet of m robots, we seek to find robot tours π_1, \dots, π_m that solve

$$\min_{\pi_1, \dots, \pi_m} \left(c(\pi_1), c(\pi_2), \dots, c(\pi_m) \right) \quad (1a)$$

$$s.t. \cup_{i=1}^m \mathcal{T}_i = \mathcal{T}, \quad (1b)$$

$$\mathcal{T}_i \cap \mathcal{T}_j = \emptyset, \text{ for all } i \neq j \quad (1c)$$

$$a(k, T_j, \pi(i, T_i)) \leq t_j^d. \quad (1d)$$

We observe that (1a) poses a MOO problem since no ordering or scalarization of objectives is given. Constraints (1b) and (1c) ensure that every task is assigned to exactly one robot, constraint (1d) ensures that all tasks are serviced before their respective deadlines.

III. METHOD

We begin by making fundamental observations regarding the relation of the proposed GMRTA problem and common objectives. We then present a scalarized solution approach using p -norms and establish theoretical guarantees on its expressiveness. Further, we show how the p -norm objective can be integrated in two state-of-the-art MRTA solvers.

A. Problem Analysis

A common goal in MRTA is minimizing the length of tours, *i.e.*, $c(\pi)$ captures the travel distance of a robot. However, there are two fundamentally different variants: i) minimizing the sum of tour lengths, *i.e.*,

$$\min \sum_{i \in \mathcal{R}} c(\pi_i), \quad (2)$$

and ii) minimizing the maximum tour length, *i.e.*,

$$\min \max_{i \in \mathcal{R}} c(\pi_i). \quad (3)$$

These two variants are well studied in the offline problem of mTSP [22]; both are a special case of our problem:

Observation 1 (Min-sum and min-max scalarization). The sum and max tour cost objectives are scalarized treatments of the MOO in (1).

Indeed, the two formulations are similar to a linear scalarization and a Chebyshev scalarization, respectively [25]. Yet, they do not have scalarization weights on the objective functions. However, we observe that sum and max tour costs are both capturing extreme solutions that either under-utilize the fleet capabilities, or are inconsiderate of the total cost and thus not resourceful. This is illustrated in Figure 2 for an mTSP instance with 10 robots. In fact, minimizing the sum of tour lengths can lead to highly unbalanced workloads: In the absence of time-windows and capacity constraints, an optimal solution will only use a single robot, as shown in 2a. In contrast, when minimizing the max of tour length, robots

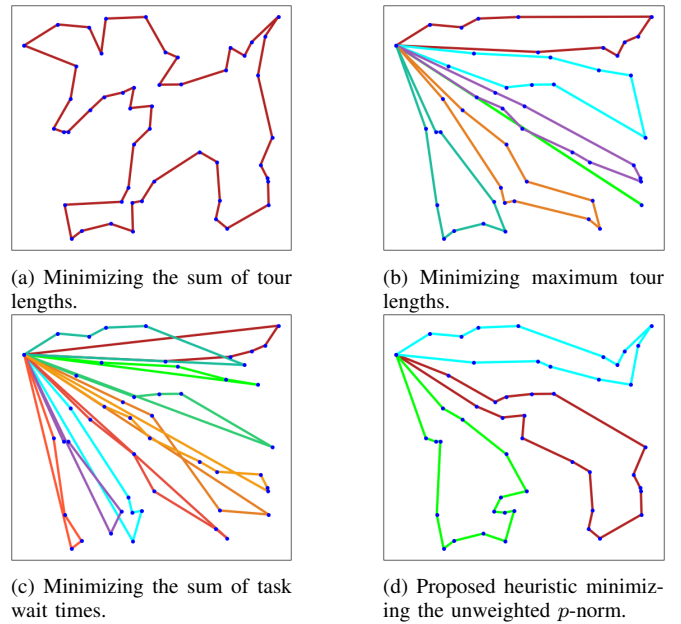


Fig. 2: Euclidean mTSP solutions for different objectives. While sum and max length as well as sum wait time achieve extreme solutions, the proposed objective using a p -norm of tour lengths (here with $p = 2$) attains a more balanced solution.

are often deployed in an almost redundant manner. In 2b, the max tour length objective uses six robots, where the green tour is only visiting only one vertex that is very close to the purple tour. In 2c, we show tours when the waiting time of tasks is minimized. This leads to using all ten robots, yet in a highly redundant manner since the objective does not consider tour lengths. Thus, this approach is more suitable for online MRTA, especially when facing tight task deadlines.

Balancing sum and max tour length is common challenge in the context of mTSP [19], [22], [26], which itself poses a multi-objective optimization problem. In the context of MRTA, we refer to this as MaxSum-MRTA (MS-MRTA). Yet, while closely related, MS-MRTA does not fully capture GMRTA, which we characterize in the following lemma.

Lemma 1 (Shortcomings of MS-MRTA). A Pareto-optimal solution to Problem 1 (GMRTA) is not necessarily a Pareto-optimal solution to MS-MRTA.

Proof. We construct a simple GMRTA instance with two solutions using the graph in Figure 3. In solution A (solid lines), robot 1 services tasks a and c, and robot 2 services only b. In solution B (dashed) robot 1 services only a, and robot 2 services b and c. Thus, we have cost vectors $c^A =$




Fig. 3: Illustration for the proof of Lemma 1. Blue and orange indicate tours for robots 1 and 2, respectively.

$[10, 4]$ and $c^B = [8, 5]$. This corresponds to a MS-MRTA instance with solutions $p^A = [10, 14]$ and $p^B = [8, 13]$.

Clearly, \mathbf{p}^B dominates \mathbf{p}^A , showing that solutions can be Pareto-optimal for GMRTA but not for MS-MRTA. \square

In essence, Lemma 1 establishes that solving MS-MRTA is not fundamentally addressing GMRTA since it is not able to find all Pareto-optimal solutions. However, we show that the inverse statement holds:

Proposition 1 (MRTA contains MS-MRTA). Any Pareto-optimal solution to MS-MRTA is also a Pareto-optimal solution to GMRTA.

Proof. We offer a proof by contradiction, showing that there cannot exist a cost vector for GMRTA that is not Pareto-optimal while the corresponding cost vector of MS-MRTA is Pareto-optimal. Let $\mathbf{c}^A = [c_1^A, \dots, c_m^A]$ be the cost vector for GMRTA, and $\mathbf{p}^A = [\max(\mathbf{c}^A), \sum(\mathbf{c}^A)]$ the corresponding cost vector for MS-MRTA. Assume \mathbf{p}^A is Pareto-optimal. In order for \mathbf{c}^A to not be Pareto-optimal there must exist another cost vector \mathbf{c}^B where $c_i^B \leq c_i^A$ holds for all $i \in \mathcal{R}$, and where $c_j^B < c_j^A$ for some $j \in \mathcal{R}$ [25]. Yet, this would imply that $\sum(\mathbf{c}^B) < \sum(\mathbf{c}^A)$ and $\max(\mathbf{c}^B) \leq \sum(\mathbf{c}^A)$, such that \mathbf{p}^A is dominated by \mathbf{p}^B and thus \mathbf{p}^A was not Pareto-optimal, contradicting the statement. Hence, any Pareto-optimal solution to MS-MRTA has at least one corresponding solution to GMRTA that is also Pareto-optimal. \square

The proposition asserts that GMRTA is a generalization of the well-known trade-off described by MS-MRTA. Next we study solution techniques to GMRTA and their applicability in common MRTA algorithms.

B. Fundamental Solution

A fundamental solution to GMRTA can be attained by using a Chebyshev scalarization, *i.e.*, weighted maximization. This results in solving

$$\min \max_{i \in \mathcal{R}} w_i c(\pi_i), \quad (4)$$

where w_1, \dots, w_n are tunable parameters. Adequately choosing these weights allows for finding any Pareto-optimal solution of GMRTA [25], [27] and is therefore a fundamental solution for GMRTA (*Note: to avoid weakly Pareto-optimal solutions, the equation needs to be augmented with a tie-breaking term $\alpha \sum c(\pi_i)$ for some small $\alpha > 0$ [25]*).

However, this approach comes with two challenges: i) the objective in (4) now has n tunable parameters, making it difficult to identify desirable system behaviour in real-world settings, especially for large fleets, *i.e.*, large n . ii) The min-max optimization cannot be directly integrated in existing MRTA solvers, such as group-based assignment [2].

C. Heuristic Solution

To address these challenges, we propose a heuristic approach based on a p -norm for some $p > 0$, taking the form

$$\min \left(\sum_{i \in \mathcal{R}} (w_i c(\pi_i))^p \right)^{1/p}. \quad (5)$$

Without loss of generality [25], this can be simplified to

$$\min \sum_{i \in \mathcal{R}} (w_i c(\pi_i))^p. \quad (6)$$

This approach approximates any Pareto-optimal solution:

Lemma 2. As $p \rightarrow \infty$, the objective in (6) is a *Pareto-complete* solution to MS-MRTA, *i.e.*, is able to attain any Pareto-optimal solution.

We omit a formal proof since it is an established result that weighted p -norms are approximating any Pareto-optimal solution for large values of p [28] as the p -norm then approximates the maximization in (4).

As a further simplified heuristic, we consider the unweighted case $w_1, \dots, w_n = 1$ and the objective becomes

$$\min \sum_{i \in \mathcal{R}} c(\pi_i)^p. \quad (7)$$

Since the cost functions in our formulation are tour lengths for a fleet of homogeneous robots, this simplification is intuitive: scalarization weights are commonly used to combine incomparable objectives such as trajectory length and risk. Nonetheless, scalarization weights are required to establish theoretical guarantees that any Pareto-optimal solution can be found. Therefore, the unweighted p -norm has no guarantees on approximating any Pareto-optimal solution. However, the reduced number of tunable parameters leads to a practical heuristic, reducing the need for complex tuning of scalarization weights commonly found in MOO approaches [7]. An example is illustrated in Figure 2d for $p = 2$, where it attains a balanced solution using three robots that each cover distinct parts of the environment.

If p were set to 1, (7) is equivalent to the sum of costs from (2), while $p \rightarrow \infty$ approaches the maximum of costs from (3). Thus, different choices of p attain different trade-offs between the sum and max tour length [22] and therefore provides a Pareto-complete solution to MS-MRTA. In practice, p still requires expert tuning. We explore the effect of different values of p in our simulation experiments.

D. Implementation in state-of-the-art algorithms

We briefly show how the proposed objectives from (6) and (7) can be incorporated in two state-of-the-art heuristics, namely a large neighbourhood search [29], [30], and the group assignment algorithm [2], [11].

a) Large Neighbourhood Search (LNS): LNS is a popular approach for tackling complex variants of TSP including task allocation for heterogeneous systems [30], [31], and inspection tasks [32]. Beginning with a random solution, the algorithm quickly explores the search space by repeatedly removing subsets of tasks from current robot tours and reinserting them, using a variety of removal and insertion heuristics. Insertion heuristics rely on costs of subtours, finding the position in the current subtour where insertion of a vertex comes at minimal cost increase. Thus, implementing the p -norm objective only requires raising the insertion cost metrics to a power of p , as in (7). This is applicable

for various insertion heuristics such as cheapest, nearest, furthest and random insertion [29]. Hence, the additional computational effort is within $\mathcal{O}(1)$.

b) Group Assignment: The group assignment algorithm proposed in [2], [11] consists of two phases: First, it computes feasible groups of unassigned tasks that can be allocated to each robot, together with corresponding tentative tours. In the second step, an optimal assignment – with respect to the summed cost of tours – is computed using Integer Linear Programming (ILP). We observe that (7) can be integrated in the first phase: Single robot tours minimizing (7) can be found via common heuristics such as LNS, or exhaustive search for sufficiently small group sizes [2]. The resulting tour cost $c(\pi)^p$ can then be directly used as an input for the assignment ILP, without changing the ILP itself.

IV. NUMERICAL RESULTS

We highlight the efficacy of the proposed p -norm objective in two different task assignment experiments and showcase its strength compared to common formulations.

a) Experiment setup: The first experiment is offline MRTA in euclidean space, as in Figure 2. We solve the problem using LNS with a cheapest insertion heuristic. The second experiment consists of online MRPD – similar to [7] – in an office environment, where robots need to deliver packages between rooms and each robot has a limited capacity of four packages. Here, we employ the centralized group assignment algorithm from [2]. In both experiments tasks have loose deadlines, chosen such that the number of robots is sufficient to deliver all tasks on time. We test setups with 50-200 tasks and 2-8 robots for offline MRTA, and setups with 50-200 tasks and 4-8 robots in online MRPD.

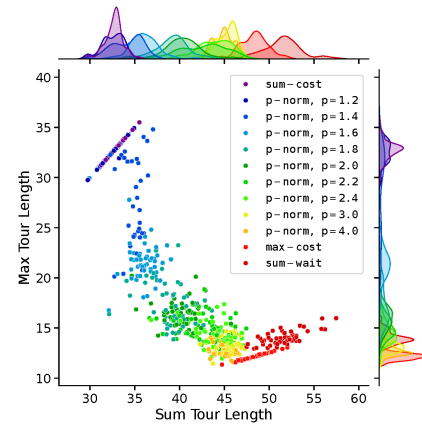
b) Baselines: We compare against three common baselines: minimizing the sum of tour lengths [11], [17], [18], minimizing the maximum tour length [22] and minimizing the summed wait times of tasks [7], [8], labeled as *sum-cost*, *max-cost* and *sum-wait*, respectively. Since we use the group assignment algorithm for our second experiment we cannot directly use a maximum objective. In this case, we approximate *max-cost* with (7) using $p = 5$. The proposed unweighted p -norm objective from (7) is labeled *p-norm*. Empirically, we found that relevant trade-offs are attained for $1 \leq p \leq 4$, for larger values, we closely approach *max-cost*.

c) Performance Measures: We employ three measures: i) the total operational cost captured by the sum of all tour lengths, ii) the total completion time captured by the maximum tour length, and iii) the fairness of robot workloads, captured by the ϵ -fairness measure from [19], defined as

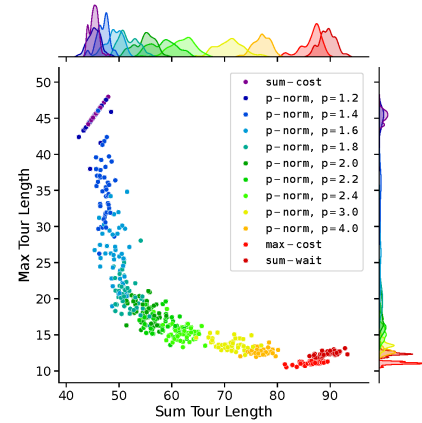
$$\epsilon(c) = \frac{1}{\sqrt{m} - 1} \left(\frac{\|c\|_1}{\|c\|_2} - 1 \right). \quad (8)$$

A. Offline MRTA

Figure 4 shows a qualitative overview of Pareto-optimal trade-offs between sum and max tour length (*i.e.*, MS-



(a) 100 tasks, 4 robots.



(b) 200 tasks, 8 robots.

Fig. 4: Illustration of Pareto-optimal trade-offs for offline MRTA.

MRTA). Different points for each objective correspond to different instances with randomly sampled of task locations.

Naturally, *sum-cost* and *max-cost* achieve the best results on total and maximum tour length, respectively, but *sum-cost* performs poorly on maximum tour length, and *max-cost* performs poorly on total tour length. This highlights that both are attaining rather extreme solutions. Further, *sum-wait* shows a low maximum length, but attains the highest total tour length among all objective functions. In fact, its solutions are dominated by the ones found for *max-cost*. In contrast, the proposed *p-norm* achieves a variety of balanced trade-offs, in particular for the large instance: While $p \geq 2$, *max-cost* and *sum-wait* perform similarly for 100 tasks and 4 robots (a), the gap is substantially larger for the setup with 200 tasks and 8 robots (b). Here, *p-norm* achieves distinct compromises with $p = 2$ and $p = 2.2$, attaining arguably the most practical solutions.

Figure 5 shows the trade-off of maximum and minimum tour length for an example with $m = 2$ robots, *i.e.*, the trade-off between c_1 and c_2 as formulated in Problem 1 but where cost indexes have been re-ordered. Here, *sum-cost* attains an extreme solution, only using one robot, while *max-cost* attains solutions where both robots have costs of ≈ 25 . Again, different values of p allow for attaining various trade-offs. We observe that the Pareto front is not convex. This

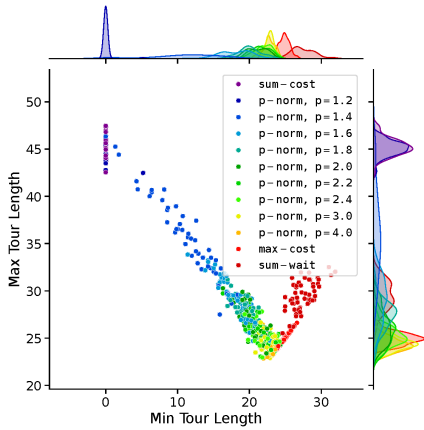


Fig. 5: Illustration of Pareto-optimal trade-offs for offline GMRTA with 2 robots and 200 tasks.

justifies the more complex fundamental solution in (4) based on Chebyshev scalarization (and our approximation thereof with a p -norm) opposed to a weighted sum – weighted sum scalarization only attains solutions on the convex hull of a Pareto front [25], [27]. Lastly, solutions found by sum-wait are dominated and thus not Pareto-optimal; hence, sum-wait is not suitable to address Problem 1.

Numerical results are summarized in Figure 6. Overall, the plots confirm the observations from Figure 4 that p -norm is able to find more balanced trade-offs. Indeed, $p = 2$ attains a close to best performance on total and maximum tour length. Additionally, we observe that with increasing complexity, *i.e.*, number of tasks and number of robots, the gap between sum-cost and max-cost increases, while the proposed p -norm performs consistently close to the respective best results on either measure. The ϵ -fairness is reported in Figure 6c. Here, sum-cost always attains the worst fairness value of 0 since parts of the fleet remain idle, while sum-wait is close to the optimum of 1. While max-cost and p -norm achieve high values ($\epsilon > .9$) for small fleets, the fairness for p -norm drops for larger fleets across all task numbers, and for max-cost for larger fleets and 50-100 tasks. In these cases, not all robots are required, creating an imbalance in workloads. Yet, recalling plot 6a, this comes at the benefit of largely decreased total tour length.

In summary, the different cost functions exhibit distinct characteristics for offline MRTA. While sum-cost, max-cost and sum-wait attain rather extreme solutions for MS-MRTA, p -norm with $p \in [1.6, 2.4]$ finds substantially more balanced and thus practical trade-offs as well as high fairness when all fleet resources are demanded.

B. Online MRPD

We now report results for Experiment 2 on online MRPD. Again, we begin with a qualitative overview of Pareto-optimal trade-offs between sum and max tour length, shown in Figure 7. To improve clarity, we include fewer values for p . Across instances the performance of all objective functions exhibit substantial statistical variance due to the randomness of generated task locations and arrival times. Therefore, in the smaller problem instances with 50 tasks and 4 robots

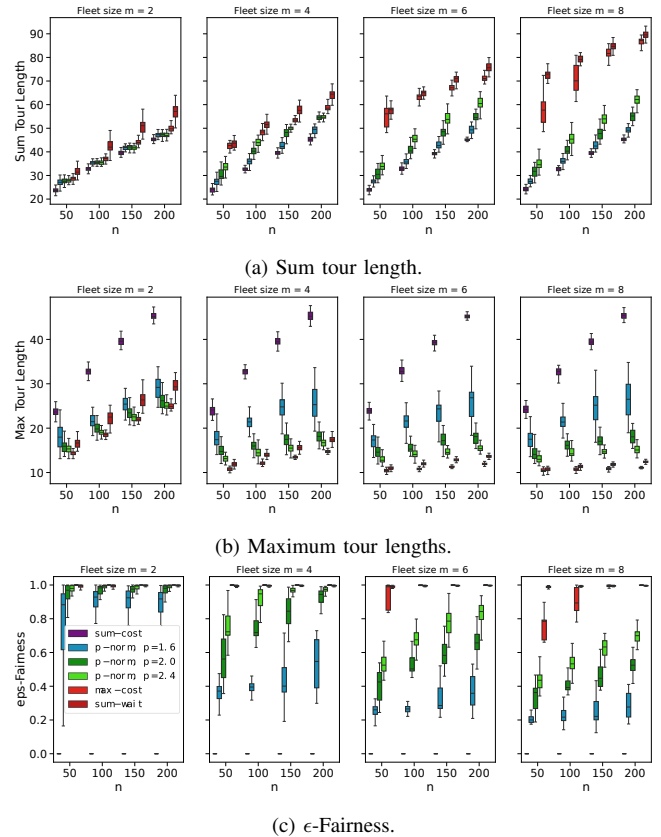


Fig. 6: Results for offline MRTA.

(a), the different objectives lead to similar behaviours with only minor trends. Yet, for the larger instances with 200 task and 8 robots in (b), the different objectives exhibit more distinct characteristics. Here, p -norm – in particular $p = 2$ and $p = 2.4$ – achieves competitive performance on both measures: its distribution over the sum tour length is similar to sum-cost, while its distribution over max tour cost is comparable to max-cost and sum-wait. Further, $p = 2$ and $p = 2.4$ avoid the long distribution tails of sum-cost and max-cost and thus shows a more reliable performance.

Next, we examine the statistical results, shown in Figure 8. First, we notice that for few tasks ($n = 50$), there is no significant difference between the objectives on the sum and max tour lengths since the system is underused and there are only few tasks being requested simultaneously. For larger fleets and growing n , we observe that the gap between max-cost and sum-wait increases; yet with a smaller gap on the sum tour length compared to the first experiment. Nonetheless, p -norm still performs close to sum-cost on the sum tour length, with particular advantages over max-cost and sum-wait for 150 and 200 tasks. It also remains competitive to max-cost on the max tour length for larger fleet sizes, albeit with increased statistical deviation compared to the results in Figure 6. For sum-wait we observe that it is least resourceful and achieves the highest sum tour length, substantially above the results for max-cost in setups with many tasks. Yet, sum-wait shows a strong performance on the max tour length, in some cases even out-

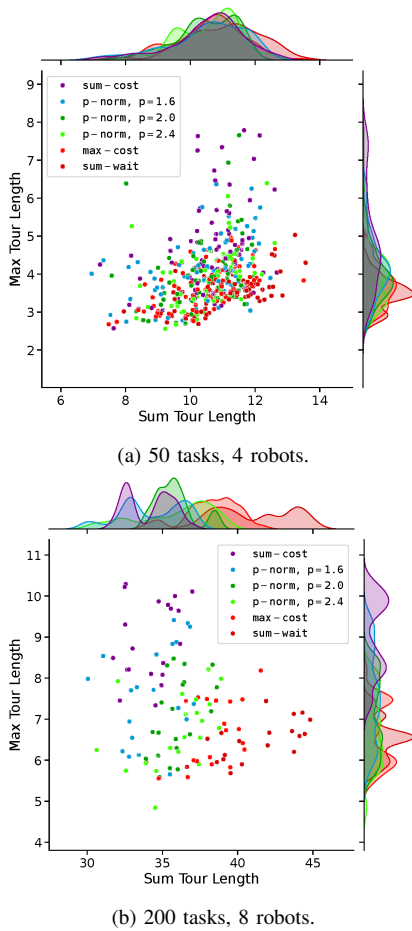


Fig. 7: Illustration of Pareto-optimal trade-offs for online MRPD.

performing *max-cost*. Considering timely task completion ensures that all available robots are active and thus helps balancing the workload between robots. This is also supported by the ϵ -Fairness measure in 8c. Similar to the Experiment 1, *sum-wait* achieves the highest fairness with values close to 1 for large number of tasks. The proposed *p-norm* attains competitive fairness for 100 tasks and more, with $\epsilon > .8$ on average. In comparison, *sum-cost* performs poorly on the fairness measures, matching earlier observations that parts of the robot fleet remain underused.

In conclusion, the second experiment confirms the main trends from earlier. While the increased statistical variance induced by the higher degree of randomness in the online setup narrows the gap on sum and max tour length, our proposed *p-norm* objective is still able to find balanced trade-offs on all measures and thus avoids the weak points of the baselines. Thus, the numerical results have highlighted that our method finds effective solutions to MRTA. Employing the simplified unweighted heuristic allows for a practical approach that avoids complex tuning of scalarization weights commonly found in multi-objective formulations.

V. DISCUSSION

We explored a novel perspective on MRTA addressing the shortcomings of commonly employed objective functions.

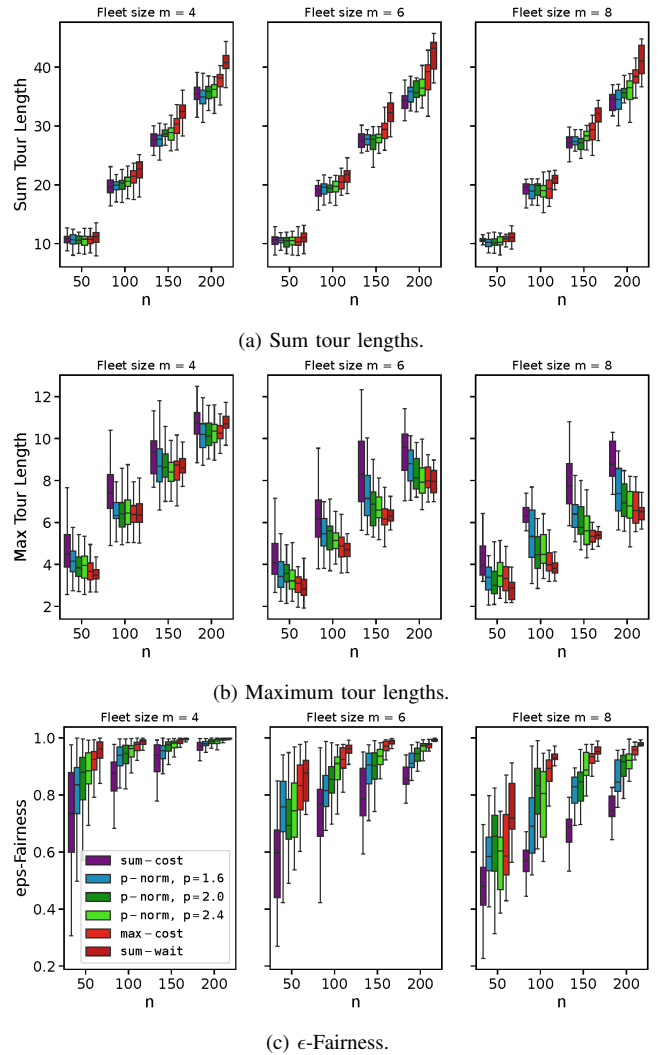


Fig. 8: Results for online MRPD.

We formulated MRTA as a MOO problem that considers the deployment costs of each robot, *e.g.*, tour length, as competing objectives. We characterized the problems expressiveness compared to the well-known trade-off between total and maximum tour length, and provided a fundamental solution. Further, we presented a heuristic solution using a *p-norm* objective, that only requires one tunable parameter and showed how it can be adapted into two state-of-the-art MRTA solvers. In two simulation experiments, we showcased that the proposed approach is suitable to find MRTA solutions with balanced trade-offs between maximum and total tour costs as well as equitable workloads among robots.

a) Limitations: One of the main limitations of the proposed heuristic in (7) is that it requires a set of homogeneous cost functions, *e.g.*, the tour lengths captures the operational cost for a set of identical robots. Thus, the proposed *unweighted p-norm* does not extend to i) heterogeneous teams, and ii) explicitly multi-objective problems with multiple cost functions such as operational cost and service quality. Both cases require scalarization weights as in (6), resulting in more complex parameter tuning. Further, while the weighted maximization in (4) comes with strong theoretical guarantees on

Pareto-completeness, our proposed p -norm heuristic does not have a bound on the error for finite p . The simulation results in this paper cover two problem setups of moderate size using synthetic data. Further evaluation with large instances such as city scale ride-sharing, additional constraints such as battery limits and recharging tours and real-world data sets for the problem setup would further show the practical impact. Lastly, the choice of the best value of p depends on the particular problem instance and thus p remains a tuning parameter. Yet, our simulations suggest that the relevant values of p are found within a small range.

b) Future Work: Future work should consider extending the proposed approach to heterogeneous teams, and including combination of different cost functions; ideally while limiting the number of tunable scalarization parameters. Further, using a p -norm instead of summation of individual robot costs could be relevant in other multi-robot problems such as team-orienteeing, multi-robot navigation and multi-agent path finding (MAPF). Robustness and computationally efficiency of multi-robot systems requires distributed optimization. Thus, future work should explore how the proposed approach can be adapted in decentralized task assignment frameworks. Finally, beyond the scope of multi-robot planning, using p -norms scalarizations is a promising approximation of the weighted maximization, *i.e.*, Chebyshev scalarization. Future work should investigate theoretical bounds on the error of this approach, compared to Pareto-complete methods such as weighted maximization.

REFERENCES

- [1] C.-I. Vasile, J. Tumova, S. Karaman, C. Belta, and D. Rus, "Minimum-violation self motion planning for mobility-on-demand," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1481–1488.
- [2] J. Alonso-Mora, S. Samaranyake, A. Wallar, E. Frazzoli, and D. Rus, "On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment," *Proceedings of the National Academy of Sciences*, vol. 114, no. 3, pp. 462–467, 2017.
- [3] A. Simonetto, J. Monteil, and C. Gambella, "Real-time city-scale ridesharing via linear assignment problems," *Transportation Research Part C: Emerging Technologies*, vol. 101, pp. 208–232, 2019.
- [4] G. Zardini, N. Lanzetti, M. Pavone, and E. Frazzoli, "Analysis and control of autonomous mobility-on-demand systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 633–658, 2022.
- [5] K. Vakil and A. Pierson, "Decentralized drone swaps for online rebalancing of drone delivery tasks," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 16 276–16 282.
- [6] P. Grippa, D. A. Behrens, F. Wall, and C. Bettstetter, "Drone delivery systems: Job assignment and dimensioning," *Autonomous Robots*, vol. 43, no. 2, pp. 261–274, 2019.
- [7] N. Wilde and J. Alonso-Mora, "Statistically distinct plans for multi-objective task assignment," *IEEE Transactions on Robotics*, vol. 40, pp. 2217–2232, 2024.
- [8] A. Botros, B. Gilhuly, N. Wilde, A. Sadeghi, J. Alonso-Mora, and S. L. Smith, "Optimizing task waiting times in dynamic vehicle routing," *IEEE Robotics and Automation Letters*, vol. 8, no. 9, pp. 5520–5527, 2023.
- [9] N. Wilde, A. Sadeghi, and S. L. Smith, "Learning submodular objectives for team environmental monitoring," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 960–967, 2022.
- [10] X. Bai, C. Li, B. Zhang, Z. Wu, and S. S. Ge, "Efficient performance impact algorithms for multirobot task assignment with deadlines," *IEEE Transactions on Industrial Electronics*, vol. 71, no. 11, pp. 14 373–14 382, 2024.
- [11] X. Bai, A. Fielbaum, M. Kronmüller, L. Knoedler, and J. Alonso-Mora, "Group-based distributed auction algorithms for multi-robot task assignment," *IEEE Transactions on Automation Science and Engineering*, vol. 20, no. 2, pp. 1292–1303, 2022.
- [12] H. N. Psarafitis, "Dynamic Vehicle Routing Problems," *Vehicle routing: Methods and studies*, vol. 16, pp. 223–248, 1988.
- [13] —, "A dynamic programming solution to the single vehicle many-to-many immediate request dial-a-ride problem," *Transportation Science*, vol. 14, no. 2, pp. 130–154, 1980.
- [14] D. J. Bertsimas and G. Van Ryzin, "A stochastic and dynamic vehicle routing problem in the euclidean plane," *Operations Research*, vol. 39, no. 4, pp. 601–615, 1991.
- [15] B. H. O. Rios, E. C. Xavier, F. K. Miyazawa, P. Amorim, E. Curcio, and M. J. Santos, "Recent dynamic vehicle routing problems: A survey," *Computers & Industrial Engineering*, vol. 160, p. 107604, 2021.
- [16] N. Soeffker, M. W. Ulmer, and D. C. Mattfeld, "Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review," *European Journal of Operational Research*, 2021.
- [17] H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE transactions on robotics*, vol. 25, no. 4, pp. 912–926, 2009.
- [18] A. Camisa, A. Testa, and G. Notarstefano, "Multi-robot pickup and delivery via distributed resource allocation," *IEEE Transactions on Robotics*, 2022.
- [19] A. S. Bhadoriya, D. Deka, and K. Sundar, "Equitable routing—rethinking the multiple traveling salesman problem," *arXiv preprint arXiv:2404.08157*, 2024.
- [20] J. Alonso-Mora, A. Breitenmoser, M. Ruffi, R. Siegwart, and P. Beardsley, "Image and animation display with multiple mobile robots," *The International Journal of Robotics Research*, vol. 31, no. 6, pp. 753–773, 2012.
- [21] M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo, "Distributed algorithms for environment partitioning in mobile robotic networks," *IEEE Transactions on Automatic Control*, vol. 56, no. 8, pp. 1834–1848, 2011.
- [22] Y. Guo, Z. Ren, and C. Wang, "imtsp: Solving min-max multiple traveling salesman problem with imperative learning," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 10 245–10 252.
- [23] C. Wei, Z. Ji, and B. Cai, "Particle swarm optimization for cooperative multi-robot task allocation: a multi-objective approach," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2530–2537, 2020.
- [24] A. T. Tolmidis and L. Petrou, "Multi-objective optimization for dynamic task allocation in a multi-robot system," *Engineering Applications of Artificial Intelligence*, vol. 26, no. 5–6, pp. 1458–1468, 2013.
- [25] J. Branke, K. Deb, K. Miettinen, and R. Slowiński, *Multiobjective optimization: Interactive and evolutionary approaches*. Springer Science & Business Media, 2008, vol. 5252.
- [26] A. Botros, N. Wilde, A. Sadeghi, J. Alonso-Mora, and S. L. Smith, "Regret-based sampling of pareto fronts for multi-objective robot planning problems," *IEEE Transactions on Robotics*, 2024.
- [27] N. Wilde, S. L. Smith, and J. Alonso-Mora, "Scalarizing multi-objective robot planning problems using weighted maximization," *IEEE Robotics and Automation Letters*, vol. 9, no. 3, pp. 2503–2510, 2024.
- [28] S. Helfrich, K. Prinz, and S. Ruzika, "The weighted p -norm weight set decomposition for multiobjective discrete optimization problems," *Journal of Optimization Theory and Applications*, vol. 202, no. 3, pp. 1187–1216, 2024.
- [29] S. L. Smith and F. Imeson, "GLNS: An effective large neighborhood search heuristic for the generalized traveling salesman problem," *Computers & Operations Research*, vol. 87, pp. 1–19, 2017.
- [30] A. Sadeghi and S. L. Smith, "Heterogeneous task allocation and sequencing via decentralized large neighborhood search," *Unmanned Systems*, vol. 5, no. 02, pp. 79–95, 2017.
- [31] N. Wilde and J. Alonso-Mora, "Designing heterogeneous robot fleets for task allocation and sequencing," in *2023 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. IEEE, 2023, pp. 156–162.
- [32] H. Dhani, K. Yu, T. Williams, V. Vajipey, and P. Tokekar, "Gatsbi: An online gtsb-based algorithm for targeted surface bridge inspection," in *2023 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2023, pp. 1199–1206.