Multi-Robot Task Allocation with Time Window and Ordering Constraints

Elina Suslova and Pooyan Fazli

Abstract— The multi-robot task allocation problem comprises task assignment, coalition formation, task scheduling, and routing. We extend the distributed constraint optimization problem (DCOP) formalism to allocate tasks to a team of robots. The tasks have time window and ordering constraints. Each robot creates a simple temporal network to maintain the tasks in its schedule. The proposed layered framework, called L-DCOP, forms efficient coalitions among robots to accomplish the tasks more efficiently as a result of their collective abilities. We conduct extensive experiments to assess the performance of the proposed algorithm and compare it against a benchmark auction-based approach. The results show that L-DCOP increases the task completion rate and task completion frequency by 1.7% and 10.1%, respectively, and reduces the task execution time by 52.5% on average.

I. INTRODUCTION

Multi-robot systems provide robust, flexible, and efficient solutions for tackling real-world applications, such as search and rescue [1], patrol and monitoring [2], [3], and distributed servicing tasks [4], [5]. Multi-robot task allocation (MRTA) is a challenging problem that involves task assignment, coalition formation, task scheduling, and routing. MRTA aims to recruit the best single or multiple robots to accomplish the tasks while optimizing the performance metrics. In variants of the MRTA problem, limitations are imposed on both robots and tasks, such as capability; capacity; and temporal, spatial, hard, and soft constraints [6]. Constraints reduce the possible solution set for the task allocation problem. The goal of this work is to address the MRTA problem with time window and ordering constraints on the tasks.

Tasks and robots are spatially distributed in the environment. In order to complete as many tasks as possible while satisfying their time window and ordering constraints, each robot must be provided with a schedule to complete the maximum number of tasks in the minimum time. Moreover, by forming teams or coalitions [7], [8], robots can achieve such tasks more efficiently as a result of their collective abilities.

There are various reasons for forming coalitions among robots to complete the tasks [1]. First, the task workload may be high, so a single robot may not be able to perform the given task within the specified time window. Second, coalition among robots leads to faster completion of tasks; therefore, the robots will have enough time to attempt other tasks in the environment. Third, the travel distance to the task may be too long for a single robot to reach the task in time and complete it individually within the specified time window. Hence, it is critical that the processes of coalition formation are managed effectively among robots.

The distributed constraint optimization problem (DCOP) is a powerful framework for modeling many real-world problems involving collaborative multi-agent systems [9]. In DCOP, agents coordinate through local communication to choose values in a distributed manner that optimize the team's global objective function. Despite the vast literature on DCOP, we are unaware of any work that has extended these models to handle tasks with time window and ordering constraints. Thus, the contributions of the paper are as follows:

- We present a layered iterative DCOP framework, called L-DCOP, to address the multi-robot task assignment and scheduling problem with time window and ordering constraints.
- The proposed framework facilitates coalition formation among robots to maximize the task completion rate and frequency while minimizing the average task execution time.
- We empirically evaluate the performance of L-DCOP against a benchmark auction-based method and show that our algorithm increases the task completion rate and task completion frequency by 1.7% and 10.1%, respectively, and reduces the task execution time by 52.5% on average.

II. BACKGROUND AND STATE OF THE ART

Centralized approaches [10] are able to find optimal solutions for the MRTA problem. However, MRTA is an NPhard problem [11], and as the number of robots or tasks increases, the problem becomes intractable. For this reason, approximated or greedy methods [12] are used extensively to solve the MRTA problem. Centralized methods suffer from a single point of failure, poor scalability, and need to generate a new global solution every time the attributes of the environment, tasks, and robots change. On the other hand, decentralized approaches are more robust to unreliable communication and robot failures and can repair the solution locally, however, they do not guarantee optimality. Below, we will discuss the two common classes of decentralized methods addressing the MRTA problem.

A. Market-based Approaches

Market- and auction-based approaches [13] are popular methods for distributed task allocation in multi-robot systems. An auctioneer announces a task, and each robot uses

Elina Suslova and Pooyan Fazli are with the Department of Computer Science, San Francisco State University, San Francisco, CA 94132, USA {esuslova,pooyan}@sfsu.edu

its local information to compute a bid, which is an estimate of the robot's expected cost or utility of performing the task. The auctioneer collects the bids and selects the winning robot that will be responsible for executing the task. McIntire et al. [14] introduced an iterated sequential single-item auction algorithm to allocate tasks with ordering constraints. Nunes et al. [15] extended the previous work to accommodate tasks with both time window and ordering constraints. Despite the vast literature on market-based approaches, we are unaware of any work that has extended these methods to integrate coalition formation to handle the task allocation problem with time window and ordering constraints.

B. DCOP-based Approaches

DCOP algorithms rely on local message passing between robots to find solutions but do not, to our knowledge, currently handle tasks with time window and ordering constraints. Farinelli et al. [16] used the DCOP framework and the max-sum algorithm to coordinate low-power embedded devices in a decentralized manner. Ramchurn et al. [1] used DCOP for task allocation in the search and rescue domain. They introduced fast max-sum, which is a more efficient and robust variation of the max-sum algorithm, to solve the task allocation problem and facilitate coalition formation among various types of agents: ambulance, fire brigade, and police agents. As with a real-world disaster scenario, tasks have deadlines and are spatially distributed, and agents must synchronize their arrival time at victim locations. In addition, previous work [17], [18] improved the DCOP model to reduce the communication overhead among the agents in the rescue domain.

III. PROBLEM STATEMENT

Let $\mathcal{R} = \{r_1, r_2, ..., r_{|\mathcal{R}|}\}$ be a finite set of robots and $\mathcal{K} = \{k_1, k_2, ..., k_{|\mathcal{K}|}\}$ be a finite set of tasks with time window and ordering constraints. The time window constraint specifies the time interval within which a task needs to be performed. In particular, the time window determines the earliest start time (EST) and the latest finish time (LFT) of the task. For example, suppose a task must be performed between 9:00 am and 2:00 pm. Each task may have one or more predecessor and successor tasks. More precisely, executing a task is possible only if all the predecessor tasks are completed first. $k_i \prec k_j$ denotes that k_i precedes k_j , or k_i must be completed before any robot can start executing k_i . Each task has a location (L) associated with it, and a robot must be present at the location to execute the task. The initial locations of robots and tasks are chosen randomly on the map. Moreover, each task has a duration (D) and a type (τ) . Each robot can perform only one task at a time, and robots are heterogeneous in the sense that each robot can execute only predefined types of tasks. Robots can perform each task individually or form teams or coalitions. When multiple robots $C_{k_i} \subseteq \mathcal{R}$ collaborate on one task k_i , the task duration D_{k_i} is divided by the number of robots, $\frac{D_{k_i}}{|C_{k_i}|}$, task duration D_{k_i} is divided by the number of robots, $\frac{D_i}{|C_i|}$ and robots in the coalition contribute to the task equally. Furthermore, robots in a coalition do not have to work on

a task simultaneously; they are allowed to perform the task partially and then move to their next assigned task, leaving the rest to other robots in the coalition.

In brief, there are tasks in the environment whose locations, durations, types, time windows, and orders are given to the set of robots. Each robot's duty is to go to its assigned task's location, perform the task completely or partially, and then move to the next assigned task on the map. Tasks are allocated and scheduled before the execution is started.

In order to assess the performance of the algorithms, we define three evaluation metrics: 1) task completion rate (TCR), 2) task completion frequency (TCF), and 3) average task execution time (ATET). The evaluation metrics are defined as below:

1)
$$\text{TCR} = \frac{|\hat{\mathcal{K}}|}{|\mathcal{K}|} \times 100,$$

2) $\text{TCF} = \frac{|\hat{\mathcal{K}}|}{m},$
3) $\text{ATET} = \frac{1}{|\mathcal{K}|} \sum_{i=1}^{|\mathcal{K}|} \text{ET}_{k_i},$

where $|\hat{\mathcal{K}}|$ is the number of completed tasks, $|\mathcal{K}|$ is the total number of tasks, *m* or *makespan* is the latest finish time of the last task, and ET_{k_i} is the execution time of task k_i . We believe that TCF is a better metric than makespan to assess the performance in the MRTA problem because different solutions do not always complete an equal number of tasks. One algorithm may schedule and complete more tasks, resulting in a higher makespan, and another algorithm may complete fewer tasks, resulting in a lower makespan. The goal is to maximize the task completion rate (TCR) and task completion frequency (TCF) and minimize the average task execution time (ATET).

IV. PROPOSED APPROACH

We present L-DCOP, a layered iterative DCOP framework, to form coalitions of robots to accomplish a set of tasks that are constrained by time windows and have dependencies with other tasks. The overall scheme of the approach is shown in Figures 1–4.

A. Precedence Graph

Tasks \mathcal{K} are given to the set of robots \mathcal{R} through a precedence graph. The precedence graph $\mathcal{G}_{\mathcal{P}} = (\mathcal{K}, \mathcal{E}_{P})$ is a directed acyclic graph (DAG) with nodes \mathcal{K} corresponding to tasks and edges \mathcal{E}_{P} representing the ordering constraints between the tasks. A directed edge $e_{ij} \in \mathcal{E}_{P}$ indicates that task k_i should be completed before performing task k_j or in other words $k_i \prec k_j$. Figure 1 shows a sample precedence graph with eight tasks. The precedence graph is divided into layers, such that there is no ordering constraint between the tasks in each layer. Hence, each layer contains a set of tasks that can be executed independently. Initially, tasks in the first layer are allocated through iterations of the DCOP formulation detailed in the following sections. After assigning the tasks in the first layer, they are removed from



and duration (D). Edges show the ordering constraints. represent the tasks \mathcal{K} (shown by squares).

Fig. 1: Precedence graph. Each box represents a task Fig. 2: Factor graphs. The variable nodes \mathcal{X} represent the with ID, earliest start time (EST), latest finish time (LFT), robots \mathcal{R} (shown by circles), and the function nodes \mathcal{F}



Fig. 3: The simple temporal network (STN) formation for robot 2. The STN contains the time points (ST: start time, FT: finish time), duration (D), travel time (TT), and time window and ordering constraints of the four tasks assigned to robot 2. Task 1 is executed, and then there is a choice of completing task 4 or task 5 first. Task 7 is done at the end of the schedule.



Fig. 4: Final schedules for robot 1 to 4. \mathcal{R} is the set of robots, and \mathcal{K} is the set of tasks.

the precedence graph, and the tasks in the next layer are allocated through DCOP. This process continues until all the tasks in the precedence graph are assigned.

B. Distributed Constraint Optimization Problem

Distributed constraint optimization problem (DCOP) provides a powerful framework to model multi-robot collaboration and coordination problems [9]. DCOP generalizes the distributed constraint satisfaction problem (DisCSP). In DisCSP, the constraints are all hard, meaning that the solution

must satisfy all of them. However, real-world problems often contain soft constraints too, which need to be satisfied as much as possible. DCOP handles both types of constraints.

A distributed constraint optimization problem (DCOP) is formally defined by a tuple $\langle \mathcal{R}, \mathcal{X}, \mathcal{D}, \mathcal{F} \rangle$, where:

- $\mathcal{R} = \{r_1, \dots, r_{|\mathcal{R}|}\}$ is a finite set of robots.
- $\mathcal{X} = \{x_1, \dots, x_{|\mathcal{X}|}\}$ is a finite set of variables, where $|\mathcal{X}| = |\mathcal{R}|$ in our problem.
- $\mathcal{D} = \{\mathcal{D}_{x_1}, \dots, \mathcal{D}_{x_{|\mathcal{X}|}}\}$ is a set of finite domains for the variables in \mathcal{X} , with \mathcal{D}_{x_i} being the domain of variable

 x_i .

F = {*f*₁,...,*f*_{|*F*|}} is a finite set of cost functions, where |*F*| = |*K*| in our problem. Each cost function is defined over a set of variables: *f_i* : ∏<sub>*x*∈*X*_{*f_i} <i>D_x* → ℝ₀⁺ ∪ {+∞}, where infeasible assignments have +∞ utility and *X*_{*f_i*} ⊆ *X*.
</sub></sub>

To follow the DCOP formulation, each variable $x_i \in \mathcal{X}$ is assigned to a robot $r_i \in \mathcal{R}$, who has the sole responsibility for the variable's value. In our problem, each robot controls exactly one variable. The domain \mathcal{D}_{x_i} of each variable x_i consists of the tasks that the corresponding robot r_i is capable of doing. Each task k_j is represented as a cost function f_j . The function f_j shows the cost of accomplishing the corresponding task with different numbers of robots $\mathcal{C}_{k_j} \subseteq \mathcal{R}$, from no robot to the coalition of all the robots capable of doing task k_j . Each robot knows only about the functions in which it is involved. We compute the cost function f_j for task k_j as follows:

$$f_j = \alpha \times \max_{r_i \in \mathcal{C}_{k_j}} (m(r_i, k_j)) + (1 - \alpha) \times \sum_{r_i \in \mathcal{C}_{k_j}} tt(\mathbf{L}_{r_i}, \mathbf{L}_{k_j}),$$
(1)

where $m(r_i, k_j)$ is the latest finish time of the last task in robot r_i 's schedule if task k_j were to be assigned to the robot, $tt(L_{r_i}, L_{k_j})$ is the time it takes for robot r_i to travel from its current location to task k_j 's location, and α is a hyperparameter set to 0.7 in our experiments. The objective is to find a complete value assignment for the variables of \mathcal{X} (denoted by **x**) that minimizes the following global cost function:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{j=1}^{|\mathcal{F}|} f_j(\mathbf{x}_{f_j}), \qquad (2)$$

where \mathbf{x}_{f_j} is the partial value assignment for the variables of $\mathcal{X}_{f_i} \subseteq \mathcal{X}$.

C. DCOP Representation: Factor Graph

There are different ways to represent a DCOP problem. The factor graph [19] is a bipartite cyclic graph chosen to represent the problem. In the factor graph $\mathcal{G}_{\varphi} = (\mathcal{X} + \mathcal{F}, \mathcal{E}_{\varphi})$, the variable nodes \mathcal{X} represent the robots (shown by circles in Figure 2), and the function nodes \mathcal{F} represent the tasks (shown by squares in Figure 2). An undirected edge $e_{ij} \in \mathcal{E}_{\varphi}$ between a variable node x_i and a function node f_j indicates that robot r_i is capable of performing task k_j . Each robot has a local view of the factor graph that includes only its immediate neighbors. Figure 2 shows three factor graphs for the three layers of the precedence graph in Figure 1.

D. Solving DCOP: Max-Sum Algorithm

Solving DCOP exactly is NP-hard [20], and for this reason approximate methods such as max-sum [16] are used to solve the optimization problem. The max-sum algorithm is an incomplete inference-based method that iteratively performs message passing on the factor graph corresponding to the DCOP. In each iteration, every variable node sends messages to all the function nodes that it connects to (Equation 3), and every function node sends messages to all the variable nodes that it is connected to in the graph (Equation 4).

• The message from variable x_i to function f_j is

$$q_{i \to j}(x_i) = \beta_{ij} + \sum_{k \in \mathcal{M}_i \setminus j} s_{k \to i}(x_i), \qquad (3)$$

where $q_{i \to j}$ is the message sent from variable node x_i to function node f_j , β_{ij} is a scalar chosen such that $\sum_{x_i} q_{i \to j}(x_i) = 0$, and \mathcal{M}_i is the set of indices of all the function nodes connected to variable node x_i in the factor graph.

• The message from function f_i to variable x_i is

$$s_{j \to i}(x_i) = \min_{\mathbf{x}_{f_j} \setminus i} [f_j(\mathbf{x}_{f_j}) + \sum_{k \in \mathcal{N}_j \setminus i} q_{k \to j}(x_k)], \quad (4)$$

where $s_{j \to i}$ is the message sent from function node f_j to variable node x_i , \mathcal{N}_j is the set of indices of all the variable nodes connected to function node f_j , and $\mathbf{x}_{f_j} \setminus i \equiv \{x_k \mid k \in \mathcal{N}_j \setminus i\}.$

This process continues until the messages converge or a fixed number of iterations is reached. Each robot r_i then selects the best task by aggregating the cost values received from its neighboring robots through the adjacent function nodes:

$$x_i^* = \underset{x_i}{\operatorname{argmin}} \sum_{j \in \mathcal{M}_i} s_{j \to i}(x_i).$$
(5)

It may be the case that there are more tasks than robots in a layer of the precedence graph, or a task is not assigned to any robot. The max-sum algorithm is run on each layer for a fixed number of iterations so that all the tasks in the layer are processed. The algorithm processes and allocates the tasks in the precedence graph layer by layer. Max-sum has a time complexity of $O(d^l)$ [9], where $d = \max_{\mathcal{D}_{x_i} \in \mathcal{D}} |\mathcal{D}_{x_i}|$ is the size of the largest domain (i.e., the largest number of tasks a robot may perform), and $l = \max_{r_i \in \mathcal{R}} |N_{r_i}|$ is the largest number of neighboring robots.

E. Managing Schedules: Simple Temporal Networks

To maintain a schedule for the tasks, each robot creates a simple temporal network (STN) [21]. An STN is a graph $\mathcal{G}_{\sigma} = (\mathcal{T}, \mathcal{E}_{\sigma})$ in which nodes represent the start time (ST) or finish time (FT) of the tasks, and edges show the duration (D) of a task or travel time (TT) between two tasks. The start time of a task k_i should be scheduled in the interval $[\text{EST}_{k_i}, \text{LFT}_{k_i} - \frac{D_{k_i}}{|\mathcal{C}_{k_i}|}]$, and the finish time should be scheduled in the interval $[\text{EST}_{k_i} + \frac{D_{k_i}}{|\mathcal{C}_{k_i}|}, \text{LFT}_{k_i}]$.

Each robot r_i uses its STN to calculate the cost of a task k_j if the task were to be assigned to the robot. The robot's corresponding variable node x_i uses this value to initiate the message with the function node f_j :

$$\alpha \times m(r_i, k_j) + (1 - \alpha) \times tt(\mathbf{L}_{r_i}, \mathbf{L}_{k_j}).$$
(6)

Whenever a new task is assigned to a robot, the robot checks every possible position in the corresponding layer of the STN to insert the task. To find a solution and a valid schedule for the STN, all the tasks should be assigned start-time and finish-time points that fulfill the tasks' time window and ordering constraints and lead to the lowest makespan. The Floyd-Warshall algorithm is used to solve the STN in $O(n^3)$ polynomial time, where n is the number of time points in the network. Figure 3 demonstrates the gradual formation of the complete STN for robot 2 for the sample precedence graph of Figure 1. Figure 4 shows the final schedule for each robot after solving their associated STNs.

V. EXPERIMENTS AND RESULTS

A. Experimental Setup

We generate random precedence graphs using the method presented by Melancon et al. [22]. The map used in the experiments is a 100×100 grid in a 2D coordinate plane. We set our scheduling time frame (TF) to be 5000 steps, meaning that all the scheduled tasks must be completed within a fixed number of steps. Tasks are generated with a set of parameters, including location (L), time duration (D), earliest start time (EST), latest finish time (LFT), and type (τ). The *time percentile* (TP) detailed below determines the EST of a task in the scheduling time frame. Tasks are distributed randomly throughout the map. For each task, the time duration is sampled uniformly from one of the three different integer intervals. Short task durations are sampled from SH = [100, 300], long task durations are sampled from LO = [800, 1000], and other task durations are sampled from the full interval FU = [100, 1000]. The EST and LFT of each task are sampled uniformly from the following intervals:

$$EST = [0, TF \times TP], \tag{7}$$

$$LFT = [EST + D, TF].$$
(8)

We define three different time percentiles (TP) for the EST of tasks. In the first case, we assign the EST of all tasks to be scheduled within the initial 25% of the scheduling time frame. This setting provides the most flexible time window constraints, giving sufficient time for each task to be performed and allowing robots to schedule and complete more tasks. In the second case, we assign the EST of all tasks to be within the initial 50% of the scheduling time frame. This setting provides more restricted time window constraints for the robots compared to the previous case. In the third case, we assign the EST of all tasks to be within the initial 75% of the scheduling time frame. Using this setting, robots face the most restricted temporal constraints for the tasks.

We define two types of tasks in the environment and randomly assign a type to each task to identify whether the task can be performed by a certain robot. Each robot is randomly assigned to perform one of the two types of tasks or both. We make sure that there is at least one robot for each type of task. Our experiments also examine another parameter, the number of robots involved in the MRTA problem. For each experiment, we use 4, 6, or 8 robots to complete the tasks. In addition, we consider a different number of tasks, i.e., 10, 15, and 25 in the experiments.

In summary, we conducted the experiments with the following settings:

- number of robots: 4, 6, and 8
- number of precedence graphs per setting: 25
- number of tasks per precedence graph: $10,\,15,\,\mathrm{and}\;25$
- scheduling time frame: 5000 steps
- EST percentile of the tasks: 25%, 50%, and 75%
- task duration: SH = [100, 300], LO = [800, 1000], and FU = [100, 1000]

We compared L-DCOP with an auction-based approach from the literature called prioritized iterated auction (PIA) [14], [15]. PIA handles both time window and ordering constraints imposed on the tasks, but can only be used in noncoalition scenarios. Ours is the first work to extend coalition formation to handle tasks with time window and ordering constraints.

B. Results

We conducted extensive experiments to assess the performance of the proposed algorithm. Figures 5, 6, and 7 show the results of the experiments under different settings, averaged over 25 randomly generated precedence graphs. The summary of the results are as follow:

1) Results for tasks with an EST within the first 25% of the scheduling time frame: For 10 and 15 tasks with time duration intervals of [100, 300], [100, 1000], and [800, 1000], L-DCOP allocates more or an equal number of tasks compared to PIA and outperforms PIA under task completion frequency (TCF) and average task execution time (ATET) in 100% of the cases.

For 25 tasks, L-DCOP allocates more or an equal number of tasks compared to PIA in 100% of the cases and outperforms PIA under TCF in $\approx 89\%$ (8 out of 9) and under ATET in 100% of the cases.

2) Results for tasks with an EST within the first 50% of the scheduling time frame: For 10, 15, and 25 tasks with a task duration in the intervals of [100, 300], [100, 1000], and [800, 1000], L-DCOP allocates more or an equal number of tasks compared to PIA and outperforms PIA under TCF and ATET in 100% of the cases.

3) Results for tasks with an EST within the first 75% of the scheduling time frame: For 10 tasks with a task duration in the intervals of [100, 300], [100, 1000], and [800, 1000], L-DCOP schedules more or an equal number of tasks compared to PIA in \approx 78% (7 out of 9) of the cases and outperforms PIA under TCF and ATET in 100% of the cases.

For 15 tasks, L-DCOP schedules more or an equal number of tasks compared to PIA and outperforms PIA under TCF and ATET in 100% of the cases.

For 25 tasks, L-DCOP schedules more or an equal number of tasks compared to PIA in 100% of the cases and outperforms PIA under TCF in $\approx 89\%$ (8 out of 9) and under ATET in 100% of the cases.



Fig. 5: Results for tasks with an EST within the first 25% of the scheduling time frame averaged over 25 randomly generated precedence graphs. Task duration: SH = [100, 300], LO = [800, 1000], and FU = [100, 1000].

C. Discussion

Task durations affect the overall performance of the proposed algorithm. Shorter task durations (e.g., [100, 300]) lead to coalitions becoming less effective, in particular for smaller numbers of robots (4 robots), because each robot performs a relatively short subtask compared to the time it spends traveling to the task. For coalition formation, it is more favorable to have longer task durations, since the robots' combined travel time is justified by the amount of time each robot spends performing part of the task. Overall, the formation of coalitions in L-DCOP significantly reduces the average task execution time (ATET), and the experiments show that L-DCOP outperforms PIA under ATET in 100% of the cases.

The time duration interval [800, 1000] is challenging for smaller robot teams (4 robots), in particular when the number of tasks is high (25 tasks). Since the time duration per task is high, it takes longer for each task to be completed by a single robot, and not all tasks can be allocated due to time window and ordering constraints. L-DCOP performs better than PIA in this case, in particular with a larger number of robots (6 and 8 robots), since a coalition of robots can finish a task in a shorter time allowing more tasks to be scheduled and completed in the future.

In addition, we examined the earliest start time (EST) parameter in the experiments. If the EST of tasks is scheduled early (i.e., within the initial 25% of the scheduling time frame), the time window of the tasks will be longer and more flexible on average. Hence, robots schedule and complete more tasks in this setting. When the EST of tasks is assigned to be in the later percentiles of the scheduling time frame (50% and 75%), the time window will be shorter and less flexible on average. Under these settings, it is more challenging for robots to schedule and complete the tasks.

In summary, the results show that L-DCOP increases the task completion rate and task completion frequency by 1.7% and 10.1%, respectively, and reduces the task execution time by 52.5% on average.

VI. CONCLUSION AND FUTURE WORK

We presented L-DCOP, the first DCOP formulation of the multi-robot task allocation problem with time window and ordering constraints. The proposed method forms efficient coalitions among robots to maximize the task completion rate and frequency while minimizing the task execution time.



Fig. 6: Results for tasks with an EST within the first 50% of the scheduling time frame averaged over 25 randomly generated precedence graphs. Task duration: SH = [100, 300], LO = [800, 1000], and FU = [100, 1000].

The method outperforms an auction-based approach under different evaluation metrics. For future work, we plan to extend this work in various directions:

- Scalability: Scalability is a major challenge in DCOPS due to communication overhead among the robots. We plan to work on improved variations of the max-sum algorithm to reduce this overhead [23], [1].
- **Heterogeneity:** The algorithm should be able to handle various forms of heterogeneity, such as different motion or sensing capabilities of the robots.
- **Dynamic environments:** The method should be able to adapt to changes in the environment without having to be completely re-run. For example, robots or tasks can be added to or removed from the environment. Dynamic DCOPs [24] can potentially address this issue.
- Uncertainty: Handling uncertainties, such as changing costs and task durations, is another challenge to address in the MRTA problem.
- **Robustness:** The algorithm should be robust to failures [25], such as robot action failure and unreliable communication.
- Physical limitations: The method should be able to

handle energy- or battery- constrained robots.

• **Boundedness:** The algorithm should provide provable bounds on the solution quality for the MRTA problem.

REFERENCES

- S. Ramchurn, A. Farinelli, K. Macarthur, M. Polukarov, and N. Jennings, "Decentralised coordination in robocup rescue," *The Computer Journal*, vol. 53, no. 9, pp. 1447–1461, 2010.
- [2] P. Fazli, A. Davoodi, and A. K. Mackworth, "Multi-robot repeated area coverage," *Autonomous Robots*, vol. 34, no. 4, p. 251–276, 2013.
- [3] P. Fazli and A. K. Mackworth, "The effects of communication and visual range on multi-robot repeated boundary coverage," in *IEEE International Symposium on Safety, Security, and Rescue Robotics* (SSRR), 2012, pp. 1–8.
- [4] J. Ji, P. Fazli, S. Liu, T. Pereira, D. Lu, J. Liu, M. Veloso, and X. Chen, "Help me! sharing of instructions between remote and heterogeneous robots," in *International Conference on Social Robotics (ICSR)*, 2016, pp. 786–795.
- [5] M. Veloso, J. Biswas, B. Coltin, and S. Rosenthal, "Cobots: Robust symbiotic autonomous mobile service robots," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2015, pp. 4423–4429.
- [6] E. Nunes, M. Manner, H. Mitiche, and M. Gini, "A taxonomy for task allocation problems with temporal and ordering constraints," *Robotics* and Autonomous Systems, vol. 90, pp. 55–70, 2017.
- [7] J. A. Adams *et al.*, "Coalition formation for task allocation: Theory and algorithms," *Autonomous Agents and Multi-Agent Systems*, vol. 22, no. 2, pp. 225–248, 2011.



Fig. 7: Results for tasks with an EST within the first 75% of the scheduling time frame averaged over 25 randomly generated precedence graphs. Task duration: SH = [100, 300], LO = [800, 1000], and FU = [100, 1000].

- [8] L. E. Parker and F. Tang, "Building multirobot coalitions through automated task solution synthesis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1289–1305, 2006.
- [9] F. Fioretto, E. Pontelli, and W. Yeoh, "Distributed constraint optimization problems and applications: A survey," *Journal of Artificial Intelligence Research*, vol. 61, no. 1, pp. 623–698, 2018.
- [10] S. D. Ramchurn, M. Polukarov, A. Farinelli, C. Truong, and N. R. Jennings, "Coalition formation with spatial and temporal constraints," in *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2010, pp. 1181–1188.
- [11] B. P. Gerkey and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *The International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.
- [12] M. C. Gombolay, R. Wilcox, and J. A. Shah, "Fast scheduling of multirobot teams with temporospatial constraints," in *Robotics: Science and Systems (RSS)*, 2013, pp. 49–56.
- [13] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: A survey and analysis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257–1270, 2006.
- [14] M. McIntire, E. Nunes, and M. Gini, "Iterated multi-robot auctions for precedence-constrained task scheduling," in *International Conference* on Autonomous Agents and Multiagent Systems (AAMAS), 2016, pp. 1078–1086.
- [15] E. Nunes, M. McIntire, and M. Gini, "Decentralized allocation of tasks with temporal and precedence constraints to a team of robots," in *IEEE International Conference on Simulation, Modeling, and Programming* for Autonomous Robots (SIMPAR), 2016, pp. 197–202.
- [16] A. Farinelli, A. Rogers, A. Petcu, and N. R. Jennings, "Decentralised coordination of low-power embedded devices using the max-sum algorithm," in *International Conference on Autonomous Agents and*

Multiagent systems (AAMAS), 2008, pp. 639-646.

- [17] M. Pujol-Gonzalez, J. Cerquides, A. Farinelli, P. Meseguer, and J. A. Rodriguez-Aguilar, "Efficient inter-team task allocation in robocup rescue," in *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015, pp. 413–421.
- [18] P. Scerri, A. Farinelli, S. Okamoto, and M. Tambe, "Allocating tasks in extreme teams," in *International Conference on Autonomous Agents* and *Multiagent Systems (AAMAS)*, 2005, pp. 727–734.
- [19] F. R. Kschischang, B. J. Frey, and H. . Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 498–519, 2001.
- [20] P. J. Modi, W.-M. Shen, M. Tambe, and M. Yokoo, "Adopt: asynchronous distributed constraint optimization with quality guarantees," *Artificial Intelligence*, vol. 161, no. 1, pp. 149–180, 2005.
- [21] R. Dechter, I. Meiri, and J. Pearl, "Temporal constraint networks," *Artificial intelligence*, vol. 49, no. 1-3, pp. 61–95, 1991.
- [22] G. Melançon, I. Dutour, and M. Bousquet-Mélou, "Random generation of directed acyclic graphs," *Electronic Notes in Discrete Mathematics*, vol. 10, pp. 202–207, 2001.
- [23] Y. Kim and V. Lesser, "Improved max-sum algorithm for dcop with n-ary constraints," in *International Conference on Autonomous Agents* and Multiagent Systems (AAMAS), 2013, pp. 191–198.
- [24] K. D. Hoang, F. Fioretto, P. Hou, M. Yokoo, W. Yeoh, and R. Zivan, "Proactive dynamic distributed constraint optimization," in *International Conference on Autonomous Agents and Multiagent Systems* (AAMAS), 2016, p. 597–605.
- [25] P. Fazli, A. Davoodi, P. Pasquier, and A. K. Mackworth, "Complete and robust cooperative robot area coverage with limited range," in *IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), 2010, pp. 5577–5582.