Asynchronous Task Plan Refinement for Multi-Robot Task and Motion Planning

Yoonchang Sung[1], Rahul Shome[2], and Peter Stone[3]

Abstract— **This paper explores general multi-robot task and motion planning, where multiple robots in close proximity manipulate objects while satisfying constraints and a given goal. In particular, we formulate the plan refinement problem—which, given a task plan, finds valid assignments of variables corresponding to solution trajectories—as a hybrid constraint satisfaction problem. The proposed algorithm follows several design principles that yield the following features: (1) efficient solution finding due to sequential heuristics and implicit time and roadmap representations, and (2) maximized feasible solution space obtained by introducing minimally necessary coordination-induced constraints and not relying on prevalent simplifications that exist in the literature. The evaluation results demonstrate the planning efficiency of the proposed algorithm, outperforming the synchronous approach in terms of makespan.**

I. INTRODUCTION

Developing multi-robot systems to achieve a desired goal while interacting with objects in the world requires integrated reasoning about task sequencing, task allocation, and motion planning. Task and motion planning (TAMP [1]) jointly addresses the search for a sequence of discrete symbolic actions, the selection of which object to manipulate, and the assignment of continuous values to actions, determining how to execute those actions. However, the TAMP literature has predominantly focused on single-robot problems.

Another closely related topic is multi-robot motion planning [2], [3], which aims to find collision-free paths for multiple robots. In this context, objects are not considered for manipulation but rather are treated as obstacles. Additionally, multi-robot motion planning typically addresses individual motion planning problems, unlike TAMP where a sequence of motion planning problems is considered. The objective of this work is to develop a general-purpose multi-robot TAMP (MR-TAMP) framework that inherits challenges from both of these perspectives.

In existing MR-TAMP research, two prevalent simplifications are the *pre-discretization* of the search space [4], [5] and *synchronous actions* [6]–[9], where robots simultaneously initiate and complete action execution. While these assumptions simplify algorithm design, they can significantly diminish the space of feasible solutions, potentially preventing the solution of certain feasible problems and reducing the diversity of available solution paths.

In this work, our goal is to formulate MR-TAMP problems that maximize the feasible solution space by avoiding both of these simplifications. This approach can be viewed as an extension of the TAMP formulation to MR-TAMP, introducing only the necessary constraints arising from multi-robot coordination. The formulation essentially represents a *hybrid* constraint satisfaction problem (H-CSP [1], [10]), incorporating both discrete and continuous variables.

To achieve this goal, we address a specific aspect in this work, referred to as the *refinement* problem. When a task plan is provided, specifying the sequences of object manipulations for all robots, the objective of the refinement problem is to assign values to all continuous variables that meet the constraints, in order to find solution paths that the robots can execute. This direction shows promise, as we can seamlessly harness state-of-the-art multi-agent task planners from the AI planning community [11], [12] when developing the complete framework in the future.
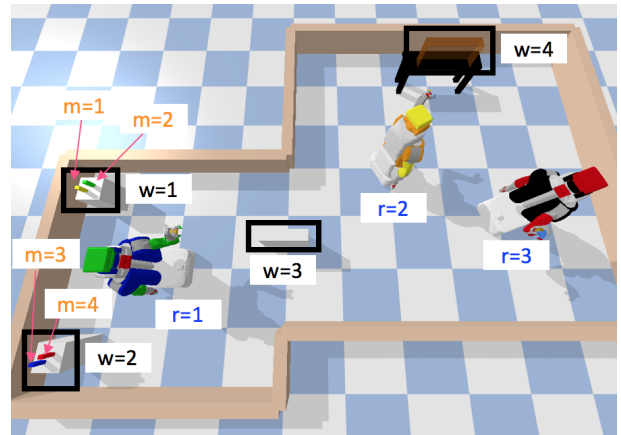

Fig. 1: Example MR-TAMP task showing the initial state. Robots ($\{r\}_{r=1}^3$), movable objects ($\{m\}_{m=1}^4$), and workspace regions ($\{w\}_{w=1}^4$) are depicted in the figure, while fixed objects, corresponding to walls, shelves, table, and cabinet, are not shown. The goal is to move all movable objects from their initial locations to the cabinet (*i.e.*, workspace region 4). The given task involves three robots such that robot 1 moves all movable objects to workspace region 3, and robots 2 and 3 move them to workspace region 4.

Figure 1 illustrates the type of task we address, wherein multiple mobile manipulators operate in close proximity, involving multi-step manipulations such as picking up and placing multiple objects. The process of solving the proposed refinement problem, which aims to satisfy the given task plan and constraints, reveals specific grasp poses, placements, motions, and action scheduling for the robots.

[1]The University of Texas at Austin, Department of Computer Science yooncs8@cs.utexas.edu
[2]The Australian National University, School of Computing rahul.shome@anu.edu.au
[3]The University of Texas at Austin, Department of Computer Science and Sony AI pstone@cs.utexas.edu

Our main contributions can be summarized as follows: (1) the introduction of a general problem formulation for MR-TAMP that is inherently asynchronous and does not require complex scheduling, (2) the identification of fundamental challenges raised by this problem, and (3) the proposal of a search algorithm that incorporates promising heuristics.

II. THE MR-TAMP REFINEMENT PROBLEM

A. Notations and assumptions

Consider R robots, indexed as $\{r\}_{r=1}^{R}$, manipulating objects to achieve a goal in a 3D workspace. The workspace consists of M movable objects, such as cups and plates, indexed as $\{m\}_{m=1}^{M}$ and F fixed objects, such as tables and shelves, indexed as $\{f\}_{f=1}^{F}$. We denote W workspace regions as $\{w\}_{w=1}^{W}$, where movable objects can be placed, such as the surface of the table and the space on the shelf, inspired by the work [13].

While our framework is not necessarily restricted to homogeneous robots (*i.e.*, robots with the same shapes, degrees of freedom, and abilities to move and manipulate), in this paper, we consider homogeneous robots for the sake of notational convenience. Each robot r operates in a d-dimensional configuration space whose configuration is represented as $q_r \in \mathcal{C}_r \subset \mathbb{R}^d$. The pose of a movable object m is denoted as $p_m \in \mathcal{P}_m \subset SE(3)$. Then, the composite configuration space for all robots and movable objects becomes $\mathcal{C} = \prod_{r=1}^{R} \mathcal{C}_r \times \prod_{m=1}^{M} \mathcal{P}_m$. We denote the free space of the composite configuration space as $\mathcal{C}^{\mathrm{F}}$, which represents all possible configurations of robots and movable objects that are positioned stably and do not collide with each other and with fixed objects. Correspondingly, the obstacle space is defined as $\mathcal{C}^{\mathrm{O}} = \mathcal{C} \setminus \mathcal{C}^{\mathrm{F}}$.

We assume quasi-static dynamics in the world, which implies that movable objects remain stable after being manipulated by robots. Additionally, we assume that each movable object can be manipulated by a single robot. Furthermore, we assume deterministic transition effects, full observability, and lossless communication among robots. While our focus in this work is on pick-and-place tasks where geometric constraints are of major concern, our ultimate aim is to position this work as a foundational framework in MR-TAMP that can effectively address a wider range of practical challenges in the future, including those that relax the assumptions mentioned in this paragraph.

B. Mode-based abstract actions

We employ the notion of a *mode* [14]–[17], denoted by σ, which specifies a constraint submanifold of $\mathcal{C}^{\mathrm{F}}$, to define actions. These modes are determined by the contact points between the robot and the movable object (*e.g.*, robot r grasping movable object m), while the remaining objects remain stationary. We consider two types of modes: a *transit mode* σ^{S} where a robot moves with an empty hand, and a *transfer mode* σ^{F} where a robot moves while holding a movable object. The transition between two adjacent modes can be facilitated through a *transition configuration*, which represents the robot's grasping or placing configuration.

We define the *abstract action* based on these two modes. Let the abstract action be $a = \{\sigma^{\mathrm{S}}(r, m, w, w'), \sigma^{\mathrm{F}}(r, m, w, w')\}$. $\sigma^{\mathrm{S}}(r, m, w, w')$ indicates that robot r moves from workspace region w to another workspace region w' with an empty hand in order to grasp movable object m located in w'. $\sigma^{\mathrm{F}}(r, m, w, w')$ indicates that robot r, while already grasping movable object m in workspace region w, moves and places it in another workspace region w'. These actions are still abstract because continuous parameters, such as robot configurations $\{q_r\}_{r=1}^{R}$ and object poses $\{p_m\}_{m=1}^{M}$, are not yet specified. Abstract actions may encompass both arm and base motions, as illustrated in Figure 1.

C. H-CSP for refinement

We formulate the refinement of abstract actions into fully specified actions that robots can execute as an H-CSP problem. This problem involves assigning values from the domains of variables while ensuring that the assigned values do not violate any constraints. The variable set is defined as $\mathcal{V} = \{\{v_r^q\}_{r=1}^{R}, \{v_r^g\}_{r=1}^{R}, \{v_m^p\}_{m=1}^{M}\}$, where v_r^q is a transition configuration variable for robot r, v_r^g is a grasp variable for robot r, and v_m^p is a pose variable for movable object m. The domains for these variables are defined as follows: for v_r^q, $\mathcal{D}_r^q = \mathcal{C}_r$; for v_r^g, $\mathcal{D}_r^g = \cup_{m=1}^{M} \mathcal{G}_{r,m}$; and for v_m^p, $\mathcal{D}_m^p = \mathcal{P}_m$. Here, $\mathcal{G}_{r,m} \ni g_{r,m} = (r, m, \gamma_{r,m})$ indicates that robot r grasps movable object m with a relative transformation $\gamma_{r,m}$ between the pose of the robot r's end-effector and the pose of the object p_m. The abstract actions are associated with variables as goal variables, where $\sigma^{\mathrm{S}}(r, m, w, w')$ includes v_r^q and v_r^g, while $\sigma^{\mathrm{F}}(r, m, w, w')$ includes v_r^q and v_m^p.

We present the mode-specific constraints, which are parameterized, that the assigned values must satisfy as follows:

- $\texttt{Motion}\big(q_r, q_r'; G_r = (V_r, E_r)\big)$ represents a reachability constraint for robot r from a start configuration q_r to a goal configuration q_r'. This constraint is verified by applying existing motion planners. In our case, we utilize probabilistic roadmap (PRM [18]) planners, which are a well-known sampling-based motion planner. These planners generate a roadmap data structure $G_r = (V_r, E_r)$, where the vertex set V_r consists of robot r's configurations, and the edge set E_r consists of edges (often straight lines in C_r) that connect pairs of vertices in V_r. The PRM planners assist in finding a path that connects the start configuration q_r and the goal configuration q_r'.
- $\texttt{CFree}\big(\{q_r\}_{r=1}^{R}, \{p_m\}_{m=1}^{M}, \{f\}_{f=1}^{F}\big)$ ensures that there are no pairwise collisions among robots at $\{q_r\}_{r=1}^{R}$ configurations, movable objects at $\{p_m\}_{m=1}^{M}$ poses, and fixed objects $\{f\}_{f=1}^{F}$. Eventually, the tensor product of individual roadmaps $\{G_r\}_{r=1}^{R}$, as shown in the $\texttt{Motion}$ constraint, must find R paths corresponding to R robots that satisfy this $\texttt{CFree}$ constraint.
- $\texttt{Kin}(q_r, g_{r,m}, p_m)$ guarantees the existence of a kinematic solution for grasping movable object m at pose p_m with grasp $g_{r,m}$ and robot r's configuration q_r.
- $\texttt{Grasp}(g_{r,m}, p_m)$ represents a graspability constraint, indicating that a movable object at pose p_m can be grasped

with grasp $g_{r,m}$.

- $\texttt{Hold}(r, m)$ ensures that movable object m is securely attached to the hand of robot r. When this constraint is activated, it affects other constraints in the following manner. In the $\texttt{CFree}$ constraint, the pose p_m of movable object m is no longer considered directly, but can be computed based on grasp $g_{r,m}$ and robot r's configuration q_r. Additionally, the collision detection between robot r and movable object m is no longer considered in the $\texttt{CFree}$ constraint. Furthermore, it prevents the activation of $\texttt{Grasp}$ constraints for other robots besides r, ensuring that the same movable object cannot be grasped by multiple robots while it is already being held.
- $\texttt{Contain}(m, w)$ constrains that movable object m is stably placed within workspace region w.

Among the constraints, $\texttt{Motion}$, $\texttt{CFree}$, and $\texttt{Kin}$ are always applied to both types of abstract actions, σ^{S} and σ^{F}. $\texttt{Grasp}$ and $\texttt{Contain}$ constraints serve as goals within the abstract actions. The constraints applied for each type of abstract action are presented as follows:

- σ^{S}: $\texttt{Motion}$, $\texttt{CFree}$, $\texttt{Kin}$, and $\texttt{Grasp}$.
- σ^{F}: $\texttt{Motion}$, $\texttt{CFree}$, $\texttt{Kin}$, $\texttt{Hold}$, and $\texttt{Contain}$.

Note that we do not introduce a constraint enforcing the synchronous start and end of abstract actions for all robots, characterizing the synchronous approach. Therefore, our formulation strictly generalizes the synchronous formulation.

D. The proposed problem

In this work, we address a partial problem where *ground* abstract actions for all robots are provided, which means that the arguments r, m, w, and w' are grounded in all instances of σ^{S} and σ^{F}, as well as the ordering among abstract actions. However, we still need to assign values to the variables of the corresponding abstract actions that satisfy the constraints specified in Section II-C. This particular approach is referred to as the *sequence-before-satisfy* strategy in the TAMP literature [1], and our focus is on addressing the satisfy part, or refinement, assuming that sequencing is given.

Specifically, we are provided with a tuple $\left\langle \{a_r^{A_r}\}_{r=1}^R, \prec \right\rangle$, where $a_r^{A_r}$ represents a set of abstract actions for robot r, and A_r is an index set specific to robot r, allowing robots to have different cardinalities of abstract actions. $\prec$ is a set of ordering constraints that determine the sequencing of the provided abstract actions.

It is important to note that these ordering constraints can apply not only to abstract actions of the same robot but also to abstract actions of different robots. For instance, if movable object m is initially placed in workspace region w, then the refinement of $\sigma^{\text{S}}(r, m, w', w'')$ for robot r cannot be carried out until another robot r' executes $\sigma^{\text{S}}(r', m, w, w')$, as the movable object m is not yet located within workspace region w'.

Furthermore, $\prec$ does not specify the ordering between every pair of abstract actions from $\{a_r^{A_r}\}_{r=1}^R$. $\prec$ is *minimally* given in the sense that it only specifies the sequence of workspace regions where each movable object is placed. Any

orderings that require geometric reasoning are not included and must be determined by solving the refinement problem. For instance, suppose workspace region w has limited space. In that case, robot r can only feasibly place movable object m in workspace region w (*e.g.*, $\sigma^{\text{F}}(r, m, w', w)$) after another robot r' removes another movable object m' from the same workspace region (*e.g.*, $\sigma^{\text{F}}(r', m', w, w')$), creating empty space in workspace region w.

Let $s_0 = \left((q_r)_{r=1}^R, (p_m)_{m=1}^M \right)$ represent the initial state, specifying the initial configurations of all robots and the initial poses of all movable objects. The refinement problem is then defined as follows: given a tuple $\left\langle \{a_r^{A_r}\}_{r=1}^R, \prec, s_0 \right\rangle$, the goal is to find valid assignments of variables defined in Section II-C for all abstract actions $\{a_r^{A_r}\}_{r=1}^R$, potentially introducing additional ordering constraints while respecting the given ordering constraints $\prec$ and the mode-specific constraints.

III. ALGORITHM

Solving the proposed problem while respecting all the constraints simultaneously is highly challenging, as even a single robot TAMP problem is known to be intractable (*i.e.*, PSPACE-hard [19]). Additionally, explicitly constructing a composite roadmap from $\{G_r\}_{r=1}^R$ is computationally expensive, especially considering the exponential increase in the number of samples required by the motion planner (such as PRM in our case) to cover the composite configuration space for all robots (*i.e.*, $\prod_{r=1}^R C_r$). Moreover, the path for an abstract action and its length can only be determined after it has been computed by the motion planner, making it difficult to anticipate in advance when a robot will place a movable object. Consequently, it is challenging to identify when the $\texttt{CFree}$ constraints are affected by the $\texttt{Hold}$ constraints without evaluating all the relevant $\texttt{Motion}$ constraints.

A. Overall framework

We propose a heuristic-based search algorithm to efficiently solve the refinement problem, incorporating the following four principles.

(1) **Least commitment**: We follow the *least commitment* principle [20], avoiding the introduction of additional ordering constraints unless absolutely necessary. This approach increases the size of the feasible solution space, leading to a more diverse set of solutions.

(2) **Sequential heuristics**: Instead of solving the problem in one step, we decompose it into a sequence of subproblems. We relax the problem by neglecting some of the mode-specific constraints, creating a relaxed problem that serves as a necessary condition for the subsequent problem in the sequence. The first subproblem is the most relaxed, and as we progress through the sequence, the neglected constraints are reintroduced incrementally. Additionally, the relaxed problem provides heuristics for guiding the search in the next subproblem. This decomposition approach is appealing because it can efficiently find a solution if one exists or effectively detect infeasibility in the early stage

of the sequence. The flow chart illustrating this process is depicted in Figure 2.

(3) **Implicit time representation**: Unlike many existing multi-robot task planning or TAMP approaches that explicitly represent time for temporal planning, our formulation and algorithm do not require explicit time representation. This approach avoids the complexity of introducing a scheduling problem and aligns with the observation made by Boutilier and Brafman [21] that explicit time representation is not always necessary. In our approach, time is implicitly revealed as a byproduct of solving the refinement problem.

(4) **Implicit composite roadmap construction**: As mentioned at the beginning of this section, explicit construction of a composite roadmap from $\{G_r\}_{r=1}^R$ is impractical. Instead, we employ the concept of implicit composite roadmap construction, referred to as *subdimensional expansion* in the literature [22]–[24]. This approach involves generating individual roadmaps for each robot independently, ignoring collisions with other robots. These individual roadmaps are then combined in a manner that takes into account robot-robot collisions. The resulting composite roadmap consists only of explored vertices and edges.
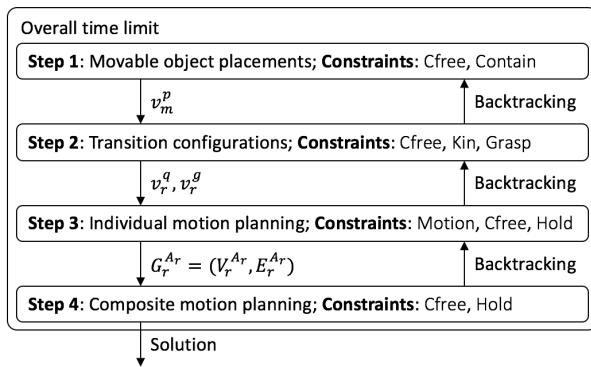

Fig. 2: The overall framework.

In the following subsections, we present each component of the algorithm depicted in Figure 2.

B. Movable object placements

In this step, we relax most of the mode-specific constraints and retain only the CFree and Contain constraints in all abstract actions. Moreover, in the CFree constraint, we disregard the robot configurations from the argument, resulting in $\mathrm{CFree}\left(\{p_m\}_{m=1}^M, \{f\}_{f=1}^F\right)$. This step can be seen as the teleportation of movable objects from one workspace region to another, excluding any robot involvement. The objective is to find valid assignments for all the relevant pose variables v_m^p present in the given abstract actions $\{a_r^{A_r}\}_{r=1}^R$ to satisfy the Contain constraint and, if necessary, introduce additional ordering constraints to resolve the CFree constraint with other movable objects.

This subproblem can be effectively solved by further decomposing it into multiple workspace region-specific problems since placements in one workspace region are completely independent of those in other regions. Let's consider

a specific workspace region w present in the given abstract actions $\{a_r^{A_r}\}_{r=1}^R$; we apply the same procedure to other relevant workspaces. For workspace region w, we find a set of sequences that specify the ordering of addition (or placement) and removal operations for each relevant movable object. Note that this sequence set can be derived entirely from $\langle \{a_r^{A_r}\}_{r=1}^R, \prec \rangle$, and that each sequence consists of an alternating sequence of addition and removal operations.

From the sequence set in workspace region w, we determine pose variable assignments for the subset of sequences that involve addition operations. We employ a sampling strategy by uniformly drawing a predetermined number of placement samples in workspace region w for each movable object in the subset.

To avoid unnecessary introduction of additional ordering constraints, we make the following observation: If we can solve the most constrained problem, where the CFree constraint is applied to all movable objects already located in workspace region w and those that will be added, no additional ordering constraints will be necessary. Only when the CFree constraint is violated for some movable objects, do we introduce additional ordering constraints, ensuring that one movable object is added after another is removed. This observation is based on the idea that in spacious workspace regions, the CFree constraint is mostly satisfied without the need for additional ordering constraints. However, in tiny workspace regions, many ordering constraints may be required.

In this step, for each abstract action under consideration, we store information about the movable objects from $\{m\}_{m=1}^M$ and fixed objects from $\{f\}_{f=1}^F$ involved in evaluating collisions with the corresponding movable object. This cached information will be utilized in the subsequent steps.

If no valid assignments can be found even after evaluating all possible combinations of the predetermined number of placement samples, we have two options. First, we can stop the process and declare the problem as infeasible. In this case, the next steps do not need to be attempted, as they rely on finding valid assignments in this subproblem. Alternatively, we can choose to draw more samples until a predetermined time limit is reached.

C. Transition configurations

After obtaining valid assignments for all relevant pose variables associated with abstract actions $\{a_r^{A_r}\}_{r=1}^R$, our next step is to find valid assignments for all relevant transition configuration variables v_r^q and grasp variables v_r^g. However, in this process, we continue to disregard certain mode-specific constraints, such as Motion and Hold, as well as the presence of other robots. Instead, we focus on considering the CFree, Kin, and Grasp constraints. This step aims to identify feasible transition configurations and grasps for all given abstract actions $\{a_r^{A_r}\}_{r=1}^R$ that are compatible with the movable object poses obtained in the previous step.

We no longer need to take the workspace region-specific approach as in the previous step. Instead, we address this subproblem for each pair of abstract actions of the same

robot consisting of σ^{S} and σ^{F} sequenced by the ordering constraints $\prec$. Let's consider the sequential abstract actions corresponding to robot r, denoted as $\sigma^{\text{S}}(r, m, w, w')$ and $\sigma^{\text{F}}(r, m, w', w'')$. These abstract actions indicate that robot r grasps movable object m in workspace region w' and moves to workspace region w'' to place the object there. The same rule is applied to all other pairs of abstract actions of the same robot sequenced by the ordering constraints $\prec$.

Instead of considering $\{q_r\}_{r=1}^R$ as arguments in the CFree constraint, we only consider robot r's configuration q_r, ignoring other robots. As for the remaining object-related arguments, we retrieve the collision information cached in the previous step, which indicates which objects must be considered for collision checking. Since collisions among objects have already been confirmed in the previous step, we only assess collisions between robot r and the relevant objects using the CFree constraint.

Since the Grasp constraint is associated with the mode σ^{S}, we first find a valid assignment for the grasp variable v_r^g corresponding to the abstract action σ^{S} by sampling a predetermined number of grasps. Once a valid grasp is found, we compute q_r with respect to the grasp $g_{r,m}$ using the Kin constraint. In the case of a mobile manipulator, as used in our experiments, computing q_r involves determining a base pose and subsequently solving an inverse kinematic problem (*i.e.*, Kin) to verify reachability to grasp $g_{r,m}$ [25]. This computed q_r is for the abstract action σ^{S}. Similarly, the same grasp $g_{r,m}$ is used to find another q_r' for the corresponding abstract action σ^{F}. The computed configurations q_r and q_r' are then used in their respective CFree constraints to ensure collision-free transition configurations.

If valid transition configurations can be found for all the abstract actions $\{a_r^{A_r}\}_{r=1}^R$ from the set of possible grasp samples, we can proceed to the next step. However, if valid transition configurations cannot be found, we have three options. First, we can choose to stop the process as explained in the previous step, indicating that a solution cannot be found. Second, we can backtrack to the previous step and explore unevaluated combinations of placement samples to potentially find valid transition configurations. To improve efficiency, we can also inform the previous step about the cause of failure, allowing suitable ordering constraints to be added and prevent the same failures in future attempts. Lastly, we can increase the number of grasp samples and reevaluate this step to improve the chances of finding valid transition configurations.

D. Individual motion planning

Even after obtaining feasible transition configurations, as mentioned in the fourth principle, solving for paths of all robots simultaneously by explicitly constructing a composite roadmap is a challenging task. To address this complexity, we leverage the discrete RRT (dRRT [23], [24]) algorithm, which is built upon the subdimensional expansion concept. The dRRT algorithm is specifically designed for solving *single-modal* motion planning problems involving multiple robots. In our algorithm, we extend the capabilities of dRRT

in two aspects: (1) individual motion planning is generalized to multi-modal motion planning, considering multiple abstract actions, and (2) our algorithm accommodates robots holding objects, which affects the collision-checking process.

In this step, we focus on considering the Motion and Hold constraints, given feasible transition configurations. During individual motion planning, we still disregard the presence of other robots. Furthermore, we assume that all movable objects, except for the one held by the corresponding robot, have been placed in their respective workspace regions, as determined in the movable object placement step. As a result, the CFree constraint still includes the same arguments as in the previous transition configuration step. However, the Hold constraint allows for collision between the robot and the movable object it holds.

Unlike the previous steps, we decompose this subproblem into multiple individual motion planning problems. Specifically, we can find a sequence of abstract actions for each robot from $\langle \{a_r^{A_r}\}_{r=1}^R, \prec \rangle$ and apply PRM to each abstract action in the sequence. In this case, the transition configurations serve as start and goal configurations, and we generate a predetermined number of samples in the respective configuration space C_r. Throughout this process, we apply the CFree and Hold constraints as mentioned before. This subproblem can be seen as the verification of reachability from the start transition configuration of the first abstract action to the goal transition configuration of the last abstract action.

If valid individual paths can be found for all robots, we can proceed to the last step. Otherwise, we have the same three options as in the transition configuration step.

E. Composite motion planning

We are now ready to consider all the intact mode-specific constraints introduced in Section II-C by merging the individual paths obtained from the previous step. This step involves constructing a tensor-product roadmap from individual roadmaps $\{G_r^{A_r} = (V_r^{A_r}, E_r^{A_r})\}_{r=1}^R$, where A_r is the abstract action index set for robot r. We denote the resulting tensor-product roadmap as $G = (V, E)$. In G, the set of vertices V is the Cartesian product of the vertices from $\{G_r^{A_r}\}_{r=1}^R$, represented as $V = \{(v_1, ..., v_r, ..., v_R) | \forall r \; v_r \in V_r^{A_r}\}$. The set of edges E is defined as $E = \big\{\big((v_1, ..., v_r, ..., v_R), (v_1', ..., v_r', ..., v_R')\big) | \forall i \; \exists (v_r, v_r') \; \big((v_r, v_r') \in E_r^{A_r} \lor v_r = v_r'\big)\big\}$. Note that in E, the condition $v_r = v_r'$ allows some robots to remain stationary. However, since robot-robot collisions and collisions between robots and movable objects held by other robots were not considered in the CFree constraint in the previous steps, some edges in E may contain collision paths among robots.

Due to limited space, we provide a brief explanation of how dRRT works and how we modify it for our problem. For detailed explanations, please refer to the works [23], [24]. dRRT is based on RRT [26] and serves as the underlying framework for constructing the composite search graph G. dRRT incrementally builds G by sampling configurations in the composite configuration space $\prod_{r=1}^R C_r$ and connecting

them using an oracle function that searches for neighboring vertices. The oracle function finds the nearest neighbor vertex v_r and another neighbor vertex v'_r within the individual roadmap $G_r^{A_r}$ for a given sampled configuration. During the composite search, the intact `CFree` and `Hold` constraints, as explained in Section II-C, are used to ensure collision-free and object-holding paths.

During the composite search, when the goal configuration (*i.e.*, transition configuration) of one robot's roadmap is reached, the next roadmap for the same robot is considered. The ordering constraints $\prec$ are taken into account in the composite search, ensuring that no adjacent edges connected to a goal configuration of the corresponding roadmap are used until another robot's roadmap, as determined by $\prec$, is reached.

If the modified dRRT algorithm finds a valid composite path for all robots, we declare that a solution path satisfying the mode-specific constraints and ordering constraints has been found, given the input $\left\langle \{a_r^{A_r}\}_{r=1}^R, \prec, s_0 \right\rangle$. dRRT has its own time limit, and if this limit is exceeded, we backtrack to the previous step. Additionally, we set an overall time limit for the entire process, and if this limit is exceeded, the algorithm terminates with no solution.

IV. EXPERIMENTS

We perform two sets of experiments in PyBullet [27] to evaluate the performance of the proposed algorithm. (1) Ablation study: We analyze the effectiveness of decomposition by comparing planning time with merged hierarchies. (2) Comparison with the synchronous approach: We evaluate the makespan (*i.e.*, the execution time of the last robot) of our algorithm against the synchronous method to highlight our method's ability to discover more effective solutions.

All the experiments are conducted using the task shown in Figure 1. We consider mobile manipulators as our robots with three and seven-dimensional configuration spaces for base motion and arm motion, respectively. Each abstract action consists of a sequence of three motion planning problems: base motion reaching a desired base position, arm motion grasping a target object, and arm motion returning to a home position. Base poses and grasp poses are all sampled, as is typically done in the literature [16], [25]. Due to the limited space, we provide the details of the task, such as specifications of input tuple $\left\langle \{a_r^{A_r}\}_{r=1}^R, \prec \right\rangle$, in the video. As the task contains 15 abstract actions, there are a total of 45 individual motion planning problems to solve the task. We report the results in Table I, where statistics are collected by solving the problem with 25 different random seeds.

Algorithms	Our algorithm	MERGE 1&2	MERGE 1–3
Planning time (s)	324.7 ± 40.2	371.2 ± 54.6	−

Algorithms		Our algorithm	Synchronous
Makespan (simulation steps)		5118.3 ± 148.4	7432.1 ± 211.8

TABLE I: Experimental results. The numbers represent mean and 95% confidence interval. − implies that all instances take longer than 10 minutes to solve.

Ablation study: Since the importance of the decomposition between Steps 3 and 4 is emphasized in dRRT [23], [24], we focus on the importance of decomposition among Steps 1, 2, and 3. The first ablation is to merge Steps 1 and 2 (*i.e.*, MERGE 1&2), and the second one is to merge Steps 1, 2, and 3 (*i.e.*, MERGE 1–3).

The results in the first row of Table I indicate that useful heuristics can be found by decomposition, and thus, a solution is found quickly. MERGE 1–3 takes longer than 10 minutes in all instances due to the generation of unnecessary motion planning problems in Step 3 that do not lead to a solution. We observe some differences between our algorithm and MERGE 1&2, but they are not significant. This implies that, although MERGE 1&2 had to solve many unnecessary inverse kinematic problems, the heuristic found by Step 2 is powerful in solving the rest of the problem, as Steps 3 and 4 consume the majority of planning time.

Comparison with the synchronous approach: In the synchronous approach, all robots either leave and arrive at their corresponding transition configurations at the same time or remain idle during that time period. In tasks where robots manipulate objects in the same workspace regions (*e.g.*, all robots converge at workspace region 3), if the planner does not find feasible transition configurations for all robots, some robots need to remain idle. Moreover, robots 2 and 3 can only start moving to workspace region 3 after robot 1 places an object there.

Makespan results in the second row of Table I support that our asynchronous algorithm is more execution time efficient than the synchronous one, which aligns with the above observations. In any case, the synchronous approach is impractical; if one of the abstract actions requires a robot to move a long distance, all the remaining robots must wait.

V. RELATED WORK

In this section, we briefly review existing MR-TAMP research, in addition to those referred to in the introduction, which rely on pre-discretization or the synchronous approach. Various task types have been investigated, such as assembly [28]–[30] and clutter removal [31]. Challenges that have not been addressed in this work are discussed in the context of MR-TAMP, including decentralized communication [5] and spatial and temporal uncertainty [32].

One distinguishing feature of this work is its implicit time representation, whereas the majority of existing works [29], [30], [33], [34] reason about time explicitly, which incurs the relatively complex overhead of task scheduling.

To solve MR-TAMP problems efficiently, various approximations have been introduced, including state space decomposition [33], [35], [36] and shared space graph [37]. Although incorporating approximations may lead to a loss of feasibility guarantees, it is an interesting avenue for future research.

Optimization-based approaches [29], [38] have made progress in MR-TAMP by leveraging logic-geometric programming [39]. The most recent work [29] in this direction

focuses on the assembly task but still relies on explicit time representations.

VI. Conclusion

In this work, we formulate a general MR-TAMP problem as H-CSP when a task plan is given, which is inherently asynchronous. We propose a refinement planning algorithm driven by design principles and evaluate its efficiency and advantages over the synchronous approach in simulation.

An immediate direction for future work is to develop a partial-order task planner capable of generating the input tuple of abstract actions and ordering constraints to complete the framework. This framework should facilitate bidirectional communication between the task planner and the proposed refinement planner to support full integration.

References

[1] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez, "Integrated task and motion planning," *Annual review of control, robotics, and autonomous systems*, vol. 4, pp. 265–293, 2021.

[2] R. Stern, N. Sturtevant, A. Felner, S. Koenig, H. Ma, T. Walker, J. Li, D. Atzmon, L. Cohen, T. Kumar *et al.*, "Multi-agent pathfinding: Definitions, variants, and benchmarks," in *Proceedings of the International Symposium on Combinatorial Search*, vol. 10, no. 1, 2019, pp. 151–158.

[3] A. Madridano, A. Al-Kaff, D. Martín, and A. De La Escalera, "Trajectory planning for multi-robot systems: Methods and applications," *Expert Systems with Applications*, vol. 173, p. 114660, 2021.

[4] K. Brown, O. Peltzer, M. A. Sehr, M. Schwager, and M. J. Kochenderfer, "Optimal sequential task assignment and path finding for multi-agent robotic assembly planning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 441–447.

[5] Y. Chen, U. Rosolia, and A. D. Ames, "Decentralized task and path planning for multi-robot systems," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4337–4344, 2021.

[6] R. Shome and K. E. Bekris, "Synchronized multi-arm rearrangement guided by mode graphs with capacity constraints," in *Algorithmic Foundations of Robotics XIV: Proceedings of the Fourteenth Workshop on the Algorithmic Foundations of Robotics 14*. Springer, 2021, pp. 243–260.

[7] T. Pan, A. M. Wells, R. Shome, and L. E. Kavraki, "A general task and motion planning framework for multiple manipulators," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 3168–3174.

[8] H. Zhang, S.-H. Chan, J. Zhong, J. Li, S. Koenig, and S. Nikolaidis, "A mip-based approach for multi-robot geometric task-and-motion planning," in *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2022, pp. 2102–2109.

[9] J. Ahn, C. Kim, and C. Nam, "Coordination of two robotic manipulators for object retrieval in clutter," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 1039–1045.

[10] T. Lozano-Pérez and L. P. Kaelbling, "A constraint-based method for solving sequential manipulation planning problems," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 3684–3691.

[11] M. De Weerdt and B. Clement, "Introduction to planning in multiagent systems," *Multiagent and Grid Systems*, vol. 5, no. 4, pp. 345–355, 2009.

[12] A. Torreño, E. Onaindia, A. Komenda, and M. Štolba, "Cooperative multi-agent planning: A survey," *ACM Computing Surveys (CSUR)*, vol. 50, no. 6, pp. 1–32, 2017.

[13] D. Berenson, S. Srinivasa, and J. Kuffner, "Task space regions: A framework for pose-constrained manipulation planning," *The International Journal of Robotics Research*, vol. 30, no. 12, pp. 1435–1460, 2011.

[14] T. Siméon, S. Leroy, and J.-P. Lauumond, "Path coordination for multiple mobile robots: A resolution-complete algorithm," *IEEE transactions on robotics and automation*, vol. 18, no. 1, pp. 42–49, 2002.

[15] K. Hauser and J.-C. Latombe, "Multi-modal motion planning in non-expansive spaces," *The International Journal of Robotics Research*, vol. 29, no. 7, pp. 897–915, 2010.

[16] C. R. Garrett, T. Lozano-Perez, and L. P. Kaelbling, "Ffrob: Leveraging symbolic planning for efficient task and motion planning," *The International Journal of Robotics Research*, vol. 37, no. 1, pp. 104–136, 2018.

[17] W. Vega-Brown and N. Roy, "Asymptotically optimal planning under piecewise-analytic constraints," in *Algorithmic Foundations of Robotics XII*. Springer, 2020, pp. 528–543.

[18] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.

[19] J. Canny, *The complexity of robot motion planning*. MIT press, 1988.

[20] D. S. Weld, "An introduction to least commitment planning," *AI magazine*, vol. 15, no. 4, pp. 27–27, 1994.

[21] C. Boutilier and R. I. Brafman, "Partial-order planning with concurrent interacting actions," *Journal of Artificial Intelligence Research*, vol. 14, pp. 105–136, 2001.

[22] G. Wagner and H. Choset, "Subdimensional expansion for multirobot path planning," *Artificial Intelligence*, vol. 219, pp. 1–24, 2015.

[23] K. Solovey, O. Salzman, and D. Halperin, "Finding a needle in an exponential haystack: Discrete rrt for exploration of implicit roadmaps in multi-robot motion planning," *The International Journal of Robotics Research*, vol. 35, no. 5, pp. 501–513, 2016.

[24] R. Shome, K. Solovey, A. Dobson, D. Halperin, and K. E. Bekris, "drrt*: Scalable and informed asymptotically-optimal multi-robot motion planning," *Autonomous Robots*, vol. 44, no. 3, pp. 443–467, 2020.

[25] R. Diankov, "Automated construction of robotic manipulation programs," 2010.

[26] S. M. LaValle *et al.*, "Rapidly-exploring random trees: A new tool for path planning," 1998.

[27] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," http://pybullet.org, 2016–2021.

[28] R. A. Knepper, T. Layton, J. Romanishin, and D. Rus, "Ikeabot: An autonomous multi-robot coordinated furniture assembly system," in *2013 IEEE International conference on robotics and automation*. IEEE, 2013, pp. 855–862.

[29] V. N. Hartmann, A. Orthey, D. Driess, O. S. Oguz, and M. Toussaint, "Long-horizon multi-robot rearrangement planning for construction assembly," *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 239–252, 2022.

[30] J. Chen, J. Li, Y. Huang, C. Garrett, D. Sun, C. Fan, A. Hofmann, C. Mueller, S. Koenig, and B. C. Williams, "Cooperative task and motion planning for multi-arm assembly systems," *arXiv preprint arXiv:2203.02475*, 2022.

[31] W. N. Tang, S. D. Han, and J. Yu, "Computing high-quality clutter removal solutions for multiple robots," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 7963–7970.

[32] M. Faroni, A. Umbrico, M. Beschi, A. Orlandini, A. Cesta, and N. Pedrocchi, "Optimal task and motion planning and execution for multiagent systems in dynamic environments," *IEEE Transactions on Cybernetics*, 2023.

[33] H. Karami, A. Thomas, and F. Mastrogiovanni, "Task allocation for multi-robot task and motion planning: A case for object picking in cluttered workspaces," in *International Conference of the Italian Association for Artificial Intelligence*. Springer, 2021, pp. 3–17.

[34] A. Messing, G. Neville, S. Chernova, S. Hutchinson, and H. Ravichandar, "Grstaps: Graphically recursive simultaneous task allocation, planning, and scheduling," *The International Journal of Robotics Research*, vol. 41, no. 2, pp. 232–256, 2022.

[35] J. Motes, R. Sandström, H. Lee, S. Thomas, and N. M. Amato, "Multi-robot task and motion planning with subtask dependencies," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3338–3345, 2020.

[36] J. Motes, T. Chen, T. Bretl, M. M. Aguirre, and N. M. Amato, "Hypergraph-based multi-robot task and motion planning," *IEEE Transactions on Robotics*, 2023.

[37] I. Umay, B. Fidan, and W. Melek, "An integrated task and motion planning technique for multi-robot-systems," in *2019 IEEE International Symposium on Robotic and Sensors Environments (ROSE)*. IEEE, 2019, pp. 1–7.

[38] M. Toussaint and M. Lopes, "Multi-bound tree search for logic-geometric programming in cooperative manipulation domains," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4044–4051.

[39] M. Toussaint, "Logic-geometric programming: An optimization-based approach to combined task and motion planning," in *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.