A certified-complete bimanual manipulation planner

Lertkultanon, Puttichai; Pham, Quang-Cuong

2018

Lertkultanon, P., & Pham, Q.-C. (2018). A certified-complete bimanual manipulation planner. IEEE Transactions on Automation Science and Engineering, 15(3), 1355-1368. doi:10.1109/TASE.2018.2791478

<https://hdl.handle.net/10356/105544>

<https://doi.org/10.1109/TASE.2018.2791478>

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at:
<https://doi.org/10.1109/TASE.2018.2791478>

Downloaded on 26 Apr 2023 07:35:01 SGT

A Certified-Complete Bimanual Manipulation Planner

Puttichai Lertkultanon and Quang-Cuong Pham, *Member, IEEE*

Abstract—Planning motions for two robot arms to move an object collaboratively is a difficult problem, mainly because of the closed-chain constraint, which arises whenever two robot hands simultaneously grasp a single rigid object. In this paper, we propose a manipulation planning algorithm to bring an object from an initial stable placement (position and orientation of the object on a support surface) towards a goal stable placement. The key specificity of our algorithm is that it is certified-complete: for a given object and a given environment, we provide a certificate that the algorithm will find a solution to *any* bimanual manipulation query in that environment whenever one exists. Moreover, the certificate is constructive: at run-time, it can be used to quickly find a solution to a given query. The algorithm is tested in software and hardware on a number of large pieces of furniture.

Note to Practitioners—This paper presents an algorithm to solve a difficult class of bimanual manipulation planning problems where a movable object can be moved only when grasped by two robots. This problem arises naturally when manipulating a large and/or heavy object such as a piece of furniture and is therefore essential to industrial automation. The algorithm first pre-computes a certificate, a set of robot motions to move the object between different placement classes that helps guarantee that the algorithm will find a solution to *any* planning query whenever one exists. This certificate is then used to quickly construct a solution trajectory to a planning query and can be reused under the same environment. The algorithm has been empirically verified through software and hardware experiments on a number of large pieces of furniture. An open-source implementation is provided.

Index Terms—Bimanual Manipulation, Certified-Completeness

NOMENCLATURE

c	Composite configuration
q	Robot configuration
T	4×4 homogeneous transformation matrix
g	Grasp parameter vector, see Section III-A
$\mathcal{G}$	Grasp configuration set
$\mathcal{P}$	Placement configuration set
π_p	Projection function from a composite configuration space into $SE(3)$
σ	Path of the object moving on a support surface; $\sigma : [0, 1] \rightarrow SE(2)$
$\mathcal{F}$	Single-transfer connected set, see Definition 10

Fig. 1: (a) A human can effortlessly manipulate a large piece of furniture. (b) Unlike humans, bimanual manipulation is extremely challenging for robots due to, for example, closed-chain constraints.

$\mathcal{F}$	Collection of single-transfer connected sets, see Section IV-A
$\mathcal{E}$	Projection via π_p of the interior of a single-transfer connected set; $\mathcal{E} = \pi_p(\text{int}(\mathcal{F}))$
$\mathcal{T}$	Connected component of object configurations on a support surface; $\mathcal{T} \subset SE(2)$

I. INTRODUCTION

Large or heavy objects are best manipulated using two hands. Humans are good at bimanual manipulation: think of how we can, for example, effortlessly manipulate a large piece of furniture (Fig. 1(a)). By contrast, bimanual manipulation is still challenging for robots, mainly because of the closed-chain constraint, which arises whenever two robot hands simultaneously grasp a single rigid object (Fig. 1(b)). This constraint poses significant challenges for manipulation planning since it (i) reduces the dimension of the configuration space [1], and (ii) restricts the range of motion of each robot arm [2]. Thus, while unimanual manipulation planning is a relatively established research field with solid theoretical foundations and a number of working demonstrations (see e.g., [3], [4], [5], [6] and references therein), results in bimanual manipulation planning are still scarce, see Section II for a review.

This paper specifically considers the harder class of problem instances where the manipulated object can be moved *only* when grasped with both hands, which is the case for large or heavy objects. We propose a manipulation algorithm to bring the object from an initial stable placement (position and orientation of the object on the support surface) towards a goal stable placement. The algorithm works at two levels: task-planning and motion-planning. At the task-planning level, a sequence of stable intermediate placements that allows changes of placement classes is found. At the motion-planning level, the motions of the two arms (when they carry the object between two intermediate placements or when they move

Manuscript received May 2, 2017; revised August 20 and November 10, 2017; accepted December 26, 2017. This work was supported in part by NTUitive Gap Fund NGF-2016-01-028 and SMART Innovation Grant NG000074-ENG.

Puttichai Lertkultanon was with the School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore. He is now with Mujin, Inc., Sumida-ku, Tokyo, Japan (email: L.Puttichai@gmail.com).

Quang-Cuong Pham is with the School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore (email: cuong.pham@normalesup.org).

freely while the object is at an intermediate placement) are determined.

The proposed algorithm possesses a desirable property called *certified-completeness*: for a given object and a given environment, the algorithm computes a certificate that it will find a solution to *any* bimanual manipulation query in that environment whenever one exists. Furthermore, the computed certificate helps construct manipulation solutions quickly as it can be used as part of solutions. The algorithm is tested in software on a number of realistic pieces of furniture. We also conducted a hardware experiment to demonstrate that the algorithm can solve real-world problems. An implementation is openly available at <https://gitlab.com/puttichai/pymanip>.

To the best of our knowledge, there currently exists no complete or certified-complete *bimanual* manipulation planners¹. This is because, in addition to the motion-planning level, manipulation planners include the task-planning level, whose completeness properties are difficult to formalize and to prove. Our algorithm is, therefore, the first to exhibit certified-completeness property. Such property is crucial for automation, where time is a valuable asset, since it essentially helps eliminate the need to spend computation time searching for non-existent manipulation paths. Computed certificates also help the algorithm to find shorter manipulation paths since the robots will only bring the object to different placements only if necessary.

The rest of the paper is organized as follows. In Section II, we review related works in manipulation planning. In Section III, we introduce the background of manipulation planning and give an overview of the proposed bimanual manipulation planner. In Section IV and Section V, we discuss main components of the proposed planner and introduce the notion of certificate. In Section VI, we present software and hardware experiments to validate our approach. Finally, in Section VII, we discuss the advantages and limitations of the approach and sketch some directions for future work.

II. RELATED WORKS

A. Bimanual Manipulation Planning

In a pioneering work [9], the authors considered the problems with the constraint that the movable object could only be moved when grasped by both arms, which is similar to ours. The proposed solutions were based on discretization of the configuration space. Their applicability were therefore limited to low-dimensional problems.

Apart from [9], most existing works on bimanual manipulation fall into one of the following paradigms:

a) Passing an object from one hand to another [10], [11], [12], [13]. These motions can be seen as a way to increase the workspace volume of the system or to help solve single-arm manipulation problems more easily.

b) Focusing on control of interactions (robot-robot, robot-object, or object-object). Examples include assembly manipulation [14], [15], [16], [17] and objects handling [18].

c) Discussing only reaching motions [19], [20], [21] or non-cooperative tasks [22], [23]. This case is most similar to the usual multi-arm path planning where robot arms move *independently but coordinatively* to reach their goals without colliding with one another.

d) No regrasping. In [24], [25], the start and goal configurations are closed-chain configurations. This case is more related to closed-chain motion planning.

e) Others. For example, in [26], the authors solved a manipulation problem by sequentially generating a sequence of object contact states, object poses, and manipulator contact points. They, however, did not take into account robot kinematics. In [27], the authors discussed a problem of a bimanual robot manipulating a foldable chair. The task was solved via chair state discretization.

A more thorough survey of work on bimanual manipulation can be found in [28].

The problem we are interested in, however, is a combination of various subproblems, including c) and d) mentioned above as well as the regrasping problem.

Although the problem itself is similar in nature to the one tackled in [9], the setting here is more practical. Our approach can deal with problems with high degrees-of-freedom (DOF). Furthermore, while the planners presented in [9] did not have any performance guarantee, our proposed planner is certified-complete.

B. Completeness Results in Manipulation Planning

Some completeness results have been presented for a broader class of problems—*multi-modal motion planning*—of which manipulation planning can be viewed as a member. In multi-modal motion planning, the configuration space has multi-modal structure and each mode limits possible motions to a submanifold. For the manipulation planning problem that we consider here, the available *family of modes* of motions consist of *transit* and *transfer*; each family comprises infinitely many modes. For example, each mode in the transit family corresponds to one object placement.

In [29], the authors presented a probabilistically complete multi-modal planner that, however, applies to only problems with *finite* modes. A more relevant result was published in [30] where the authors presented a planner, Random-MMP, that can deal with infinite number of modes. Despite the probabilistic completeness guarantee, using Random-MMP directly to solve a pick-and-place task poses critical disadvantages. This is mainly due to unsupervised mode switches: consider for example when the object is at rest on a supporting surface without being grasped by any robot (i.e., the configuration is in a transit mode), Random-MMP will proceed by simply sampling an adjacent mode, which is essentially sampling *any* grasp (transfer mode). Without utilizing knowledge of how grasps and placements correlates, as is done, e.g., in [6], the planner can be very slow since it indeed needs to randomly sample a correct order of a correct combinations of grasps and placements before it can eventually reach the goal.

¹Note, however, that proofs of completeness have been obtained for some classes of motion planning algorithms, under more or less restrictive and verifiable assumptions, see, e.g., [7], [8]. See also Section II-B for a brief review on completeness results

Apart from the results for multi-modal motion planning, there has also been a recent result in unimanual manipulation planning [31]. The authors presented an asymptotically optimal manipulation planner (the probabilistic completeness of which follows from its asymptotic optimality). The manipulation planner works by building a roadmap for each sampled grasp or placement by using an asymptotically optimal probabilistic roadmap (PRM*) [32]. However, due to the complex geometry of bimanual manipulation systems that arises from closed-chain constraints, the manipulation planner cannot directly generalize to solve bimanual manipulation queries. Also, to solve a complex manipulation query that requires many grasp-ungrasp operations, such as the one presented later in Section VI (see Fig. 7), PRM-based planners will need to build a large number of roadmaps hence enormous space complexity. Furthermore, the completeness of the manipulation planner relies on the notion of transition and contact robustness, which may not be easily verifiable.

In this work, we introduce a more practical notion of completeness, namely certified-completeness. A certified-complete planner computes a certificate that guarantees existence of solutions to any feasible manipulation query. Although the computation of certificates itself is not complete, i.e., there currently exists no theoretical guarantee if such computation will be successful, we show that it is in fact practical to compute such certificates, as presented in Section VI for a number of realistic cases.

Note also that there is also another somewhat related line of research, in which the focus is on computation of space disconnection certificate (see, e.g., [33]).

C. Regrasping

Generally speaking, regrasping is a grasp-changing operation. Here we are interested in the case when manipulators are equipped with parallel jaw grippers, which are the most common and robust grippers in industries. Unlike multi-fingered hands, which can perform in-hand regrasping with general objects, a robot equipped with a parallel gripper has to rely on a *support surface*, on which the object can be placed stably while ungrasped, to change the grasp².

Works on regrasping utilize the knowledge that to realize any regrasping motions, the robot(s) must place the object down on the support surface. The system configuration has naturally to satisfy two criteria: 1) the robot(s) must be grasping the object and 2) the object must be at a stable position. The set of configurations satisfying the aforementioned criteria, denoted as $\mathcal{G} \cap \mathcal{P}$, and connectivity between its different connected components play significant roles in solving regrasping problems.

So far, researchers have been focusing only on developing regrasping algorithms for unimanual manipulation systems. Pioneering works on regrasping problems, including [35], [36], [37], characterized the set $\mathcal{G} \cap \mathcal{P}$ by means of discretization.

Their methods are therefore limited to objects with low geometric complexity. However, the authors of [35] also proposed an interesting notion of Grasp-Placement Table, based on the discretization of $\mathcal{G} \cap \mathcal{P}$, which captured the connectivity of $\mathcal{G} \cap \mathcal{P}$. Recent works on regrasping such as [5], [6], [38] also employed some kinds of graphs to represent the connectivity. Another line of works on regrasping focuses on improving implementations/execution of existing unimanual regrasping algorithms [39], [40].

The set $\mathcal{G} \cap \mathcal{P}$ can, in fact, be grouped into a finite number of subsets, called grasp classes and placement classes [6]. Utilizing this fact, the authors of [6] introduced a high-level Grasp-Placement Graph, which showed potential connectivity between different connected components of $\mathcal{G} \cap \mathcal{P}$. They proposed a manipulation planner that, with the guidance from the graph, explored the configuration space efficiently and systematically.

One possible way to solve a bimanual manipulation planning problem is then to extend the high-level Grasp-Placement Graph [6], originally proposed for unimanual systems, to bimanual cases. However, the combinatorial complexity associated with grasp classes grows much too high, making this approach not suitable even in the case when the object has a moderate number of grasp classes. For example, consider a unimanual setting. Suppose the start and goal placements have m grasps in common but no transfer path directly connecting the two placement classes exists. The planner will have to explore *exhaustively* all m paths connecting placements start and goal in the graph before considering any manipulation path with some intermediate placements. In a bimanual setting, the planner will have to explore all $\mathcal{O}(m^2)$ possibilities in case no direct transfer path exists between the start and goal placements. Suppose there are tens of common grasp classes between the start and goal placements, this means that the planner needs to explore already hundreds of possibilities before trying to plan a manipulation paths with one intermediate placements.

Therefore, we take a different angle to approach bimanual regrasping problems, see Section III-B for overview of the approach. Instead of exploring connectivities between a large number of connected components of $\mathcal{G} \cap \mathcal{P}$, we explore the connectivities between different placement classes since for most objects, the total number of placement classes is significantly lower than that of grasp classes (see, e.g., Table I). This approach helps reduce drastically the combinatorial complexity mentioned earlier.

D. Motion Planning with Closed Kinematic Chains

Closed-chain motion planning is by itself a difficult and challenging problem. Efficient path planners such as Rapidly-exploring Random Trees (RRTs) [41] or their variants cannot be directly applied to solve such problems since the probability of a randomly sampled configuration satisfying the closed-chain constraint is essentially null [1]. This is because valid closed-chain configurations forms a set of manifolds of dimension lower than that of the ambient space. To cope with this issue, various methods have been devised to sample

²In fact, in-hand regrasping with parallel grippers can also be made possible (with minor helps from the environment) under some restricted conditions, e.g., the gripper is made specifically to grasp a specific type of objects [34] and/or all the contact geometries and other properties are precisely modelled.

closed-chain configurations and to interpolate closed-chain trajectories.

Random gradient descent was used in [1] to move a randomly sampled configuration towards the constraint manifold. In [42], the authors proposed breaking the closed kinematic chain into two subchains. A configuration of one subchain is sampled randomly while a configuration of the other is computed so as to close the chain. This method was further improved in [43]. More recent work samples configurations on a tangent space of the constraint manifold [44], [45], [46].

We take a different approach to closed-chain motion planning. Essentially, to interpolate a trajectory between closed-chain configurations, our planner first interpolates a trajectory for the movable object. The trajectory is then tracked by the two robots. We describe our closed-chain motion planner as well as our rationale in Section V.

III. BACKGROUND AND OVERVIEW OF THE BIMANUAL MANIPULATION PLANNING ALGORITHM

A. Background

This Section presents definitions and fundamentals of bimanual manipulation planning built based on previous works [3], [6].

Consider the 3D space where the bimanual manipulation system is located, called world. The world, $\mathcal{W}$, consists of two robots $\mathcal{R}_1$ and $\mathcal{R}_2$, a movable object $\mathcal{O}$, and the environment $\mathcal{E}$. Each robot is equipped with a parallel jaw gripper. The environment also includes *support surface(s)* on which the object is allowed to rest.

Let $\mathcal{C}_{\mathcal{R}_1}$ and $\mathcal{C}_{\mathcal{R}_2}$ be the configuration spaces of the two robots and $\mathcal{C}_{\mathcal{O}} \subseteq SE(3)$ the configuration space of the object. The composite configuration $\mathcal{C}$ is defined as the Cartesian products of the three aforementioned spaces. Each composite configuration $c \in \mathcal{C}$ can then be written as $c = (q_1, q_2, T)$, where $q_1 \in \mathcal{C}_{\mathcal{R}_1}$, $q_2 \in \mathcal{C}_{\mathcal{R}_2}$ and $T \in \mathcal{C}_{\mathcal{O}}$.

We equip with the composite configuration space a metric d defined as a linear combination of Euclidean distance between robot configurations and the distance between the object transformation matrices. In particular, $d(c_a, c_b) = \alpha (\|q_{1a} - q_{2a}\|_2 + \|q_{1b} - q_{2b}\|_2) + (1 - \alpha)w(T_a, T_b)$, where $\alpha \in [0, 1]$ and w is the weighted sum of the minimal geodesic distance between two rotations [47] and the Euclidean distance between two displacements.

We now define a *grasp* and a *placement* as follows.

Definition 1. A *grasp* is a relation between the object pose and the grippers' poses.

One can represent a grasp by, e.g., a pair of relative transformations between each robot gripper and the object. Note that from the definition, any pair of relative transformations can be a grasp. However, the object can be moved only when being grasped by a *valid grasp*. The set of all valid grasps are to be determined by the users, either explicitly (e.g., as a set of grasps) or implicitly (e.g., as conditions to be satisfied by the grippers). Note also that there can be many pair of robot configurations (q_1, q_2) corresponding to exactly the same grasp due to multiplicity of inverse kinematic (IK) solutions associated with the same grippers' poses.

Fig. 2: A parallel gripper with its local frame. The lateral direction is orthogonal to both finger surfaces. The sliding direction is parallel to both finger surfaces and is defined such that the approaching direction is pointing out of the gripper.

The set of all valid grasps can be *parameterized* by a set of parameters [6], which is finite but not necessarily unique. Consider for example an object composed entirely of boxes³ and a gripper shown in Fig. 2. Grasp parameters may be defined as follows [6]. l is an integer indicating the index of the link (box) that the gripper is grasping. a is an integer indicating how the gripper is approaching the object. Assuming, without loss of generality, that each box is aligned with its local coordinate frame. The integer a may be a number from 1 to 6, where if $a = 1$, the gripper's approaching direction is aligned with the $+x$ -axis of the box's local frame; if $a = 2$, the gripper's approaching direction is aligned with the $+y$ -axis, etc. b is an integer indicating the axis of the box's local frame that the gripper's sliding direction is aligned with. And the last parameter⁴ δ is a real number indicating the position of the gripper along the sliding direction. For example, if the gripper is grasping the box at the middle, we may assign $\delta = 0$, and δ increases (or decreases) when the gripper *slides* along the sliding direction. Using this notion, a grasp for the i^{th} robot may be written as a vector $g_i = [l_i \ a_i \ b_i \ \delta_i]^T$ and therefore a bimanual grasp may be written as $g = [g_1^T \ g_2^T]^T$.

Definition 2. A *placement* refers to an object transformation at which the object is in contact with a support surface.

A placement is said to be *stable* if when not in contact with any robot, the object remains stationary.

The set of all stable placements can be seen as a subset of $SE(2)$. They therefore can be parameterized by three parameters x, y , and θ , where x and y represent the position of some nominal point of the object with respect to the support surface, and θ represents the rotation around an axis passing through the point (x, y) and perpendicular to the surface.

With the above definitions of grasps and placements, we can now define a *collision-free* configuration and a *feasible* configuration.

Definition 3. A composite configuration is said to be *collision-free* if there are no collisions in the world except the ones induced by valid grasps and ones induced by placements.

Definition 4. A composite configuration $c = (q_1, q_2, T)$ is said to be *non-singular* if the Jacobians of the two robots,

³This is not a particularly impractical setting since a piece of furniture tends to be geometric and thus can be approximately, if not exactly, represented by boxes.

⁴Note that we can have more grasp parameters. For example, the gripper could tilt around the (virtual) axis connecting the two finger tips. We may introduce another parameter θ to indicate the tilting angle.

$J_{\mathcal{R}_1}(q_1)$ and $J_{\mathcal{R}_2}(q_2)$, have maximal rank. Otherwise, the configuration is said to be *singular*.

Definition 5. A composite configuration is said to be *feasible* if it is collision-free and non-singular and at least one of the following holds: 1) The robots are grasping the object with a valid grasp; 2) The object is at a stable placement.

Let us now consider intrinsic structure of $\mathcal{C}$. For convenience, we define a function $\pi_p : \mathcal{C} \rightarrow \mathcal{C}_O$ that projects a composite configuration $c = (q_1, q_2, T)$ into $SE(3)$ such that $\pi_p(c) = T$. There are two types of subsets of $\mathcal{C}$ induced by valid grasps and stable placements.

Definition 6. Grasp configuration set, $\mathcal{G}$, is the set of feasible composite configurations where the robots are grasping the object with a valid grasp.

Definition 7. Placement configuration set, $\mathcal{P}$, is the set of feasible composite configurations such that

- 1) $\forall c \in \mathcal{P} \pi_p(c)$ is a stable placement and
- 2) $\forall c \in \mathcal{P} \exists c' \in \mathcal{G} \pi_p(c') = \pi_p(c)$.

The second requirement of the placement configuration set is to ensure that for any placement configuration $c \in \mathcal{P}$, its corresponding placement is always reachable by some grasps.

Both $\mathcal{G}$ and $\mathcal{P}$ can be partitioned into a finite number of *grasp classes* and *placement classes*, respectively [6]. From the grasp parameters we introduced earlier, we define a grasp class as a subset of $\mathcal{G}$ whose configurations have the same grasp parameters l (link index) and a (approaching direction). For example, if the object is a box, there will be 6 grasp classes in total. Now consider partitioning of $\mathcal{P}$. Let $\mathcal{H}$ be the convex hull of the object. All stable placements can be grouped based on which surface of $\mathcal{H}$ is in contact with the support surface. Therefore, a placement class is defined as a subset of $\mathcal{P}$ where at each configuration, the same face of $\mathcal{H}$ is in contact with the support surface. For convenience, we will also say that two object transformations are *in the same placement class* if at both transformations, the same face of the convex hull $\mathcal{H}$ is in contact with the support surface.

There are two types of physically realizable *single-mode* paths: transit and transfer. A transit path is a path in $\mathcal{P}$ where the placement remains unchanged throughout while a transfer path is a path in $\mathcal{G}$ where the grasp remains unchanged throughout. A manipulation path is defined as an alternating sequence of single-mode paths. To plan a manipulation path, a *manipulation query* must be provided to a planner. A manipulation query is defined as follows.

Definition 8. A manipulation query, or simply query, Q , is a set of information provided to a manipulation planner to solve for a manipulation trajectory. A query consists of at least a pair of stable placements, T_s and T_g , which are the start and goal object transformations.

A query is said to be *feasible* if $T_s, T_g \in \pi_p(\mathcal{P})$.

Then a manipulation planning problem can be stated as follows.

Problem 1. Given the description of the world and a query $Q = (T_s, T_g)$, find a manipulation trajectory that brings the

object from T_s to T_g .

B. Overview of the Proposed Bimanual Manipulation Planning Algorithm

We propose the following approach to solving a bimanual manipulation query:

- Step 1** Identify the placement classes of T_s and T_g as $\mathcal{P}_s$ and $\mathcal{P}_g$, respectively.
- Step 2** Generate TypeA trajectory, within the placement class $\mathcal{P}_s$, to move the object from T_s to some T'_s .
- Step 3** Generate TypeB trajectory to bring the object from T'_s to some T'_g in the placement class $\mathcal{P}_g$.
- Step 4** Generate TypeA trajectory, within the placement class $\mathcal{P}_g$, to move the object from T'_g to T_g .

A solution to a query will be a sequence of TypeA trajectories, which connect configurations in the same placement class, and TypeB trajectories, which connect configurations from different placement classes.

In the above steps, T'_s (respectively T'_g) is an object transformation that can serve as an initial (respectively goal) transformation of the to-be-generated TypeB trajectory. Note also that in some cases, one may need to generate TypeB trajectories to move the object to, and between, some *intermediate* placements since a direct connection between $\mathcal{P}_s$ and $\mathcal{P}_g$ may not exist or cannot be found. The procedure can be done by repeating Step 2 and Step 3 until the goal placement $\mathcal{P}_g$ is reached.

The completeness of the above approach depends on the completeness of TypeA and TypeB trajectory generation methods. In the remaining of this Section, we give brief overviews of generation of both TypeA and TypeB trajectories, as well as the main algorithm.

1) TypeA: To plan TypeA trajectories, we argue that we can consider the set $\mathcal{T}_i \subset SE(2)$ of object configurations (see Section IV for more details) instead of examining a placement class $\mathcal{P}_i$, which is a subset of the high-dimensional $\mathcal{C}$.

Given two object configurations T_1 and T_2 in the same connected component of $\mathcal{T}_i$, we first generate an object path σ as if the object could move freely by itself on a support surface. Then we present a procedure to generate a TypeA trajectory that moves the object along σ . We prove that, given a valid object path, such a TypeA trajectory always exists and that our procedure will terminate with a solution in finite time.

2) TypeB: Since the motions of the system are severely constrained by closed kinematic chains, randomly generating closed-chain queries, where the start and goal configurations are in different placement classes, has slim chances of the queries being solvable. To resolve this issue, we propose a heuristic to generate closed-chain queries in such a way that, by our intuition, does not require a large range of robot motions to solve them.

Now suppose that one has a TypeB trajectory $M^B : [0, 1] \rightarrow \mathcal{C}$ connecting two placement classes $\mathcal{P}_i$ and $\mathcal{P}_j$, i.e., $M^B(0) \in \mathcal{P}_i$ and $M^B(1) \in \mathcal{P}_j$. Observe that provided that the world does not change, whenever one needs to connect configurations $c_1 \in \mathcal{P}_i$ and $c_2 \in \mathcal{P}_j$, one can reuse the trajectory M^B by planning two TypeA trajectories, M_1^A and

Fig. 3: (a) Placement classes and their verified (TypeB) connectivities. Each $M_i^B, i \in \{1, 2, \dots, 5\}$ is a transfer (TypeB) trajectory connecting two different placement classes. We call the set $\{M_1^B, M_2^B, \dots, M_5^B\}$ a *certificate*. (b) To construct a solution to a query that starts in $\mathcal{P}_1$ and ends in $\mathcal{P}_5$, one simply plans a set of TypeA trajectories, M_1^A, M_2^A, M_3^A , and M_4^A , as shown in dashed lines, to bridge the TypeB trajectories. A solution trajectory is then $M = M_1^A * M_2^B * M_2^A * M_3^B * M_3^A * M_5^B * M_4^A$.

M_2^A , where M_1^A connects c_1 and $M_2^B(0)$ and M_2^A connects $M_2^B(1)$ and c_2 . The composition (as defined in [6]) of the three trajectories, i.e., $M = M_1^A * M_2^B * M_2^A$, then serves as a solution. Since TypeB trajectories can be reused as discussed above, they have to be computed only once and the procedure may as well be offline. This inspires us to introduce a notion of a *certificate*, which is a set of useful TypeB trajectories. Once computed, a solution (if any) to any given bimanual manipulation query can then be constructed from the certificate in the aforementioned manner.

3) *Main Algorithm*: First, we generate a certificate $\mathcal{M}$. This step needs to be done only once per problem setting. Given a query $Q = (T_s, T_g)$, we then extract a placement sequence $\mathcal{P}_1 \rightarrow \mathcal{P}_2 \rightarrow \dots \rightarrow \mathcal{P}_n$, where $\mathcal{P}_1 = \mathcal{P}_s$ and $\mathcal{P}_n = \mathcal{P}_g$, along with their corresponding transfer trajectories $M_1^B, M_2^B, \dots, M_{n-1}^B$, where M_i^B is a TypeB trajectory connecting the i^{th} placement in the sequence to the next. Next, we generate a TypeA trajectory connecting M_i^B and M_{i+1}^B for every i . Finally, a solution to the query is constructed by concatenating all the trajectory (using the composition operation). Fig. 3 illustrates the proposed algorithm.

IV. GENERATING TRAJECTORIES WITHIN A PLACEMENT CLASS

In this Section, we investigate the existence of manipulation paths connecting two composite configurations in the same placement class (TypeA). Note that since a TypeA trajectory

lies entirely within a placement class, each composite configuration along such a trajectory is then a placement configuration and therefore the object will always be on a support surface.

The goal of this Section is to assert that given an object path $\sigma : [0, 1] \rightarrow SE(2)$, as if the object could move freely by itself on a support surface, there exists a finite-length⁵ manipulation path realizing the object motions along σ . In other words, the projection via π_p of the manipulation path is σ .

As the proof of existence itself is nonconstructive, we further propose an algorithm that, given an object path σ together with a certain set of assumptions, will return a manipulation path associated with σ in finite time.

A. Existence of TypeA Paths

First we introduce the notion of *single-transfer connectedness* as follows.

Definition 9. Two composite configurations c_1 and c_2 are single-transfer connected if there exists a transfer path whose terminal configurations are c_1 and c_2 .

Definition 10. A single-transfer connected set $\mathcal{F}$ is a set in which any two composite configurations are single-transfer connected. If such a set is maximal in the sense that for any point $c \in \partial\mathcal{F}$, every neighborhood of c consists of both configurations that are single-transfer connected and not single-transfer connected with c , we call it a single-transfer connected component.

Let $\mathcal{F}$ be a collection of all single-transfer connected components in $\mathcal{G} \cap \mathcal{P}_i$. Since $\mathcal{G} \cap \mathcal{P}_i$ contains no singular configuration, any $c \in \mathcal{G} \cap \mathcal{P}_i$ must be in some single-transfer connected component. That is, $\mathcal{G} \cap \mathcal{P}_i = \bigcup_{\mathcal{F} \in \mathcal{F}} \mathcal{F}$. Define $\mathcal{T}_i$ by $\mathcal{T}_i = \pi_p(\text{int}(\mathcal{G} \cap \mathcal{P}_i))$. We have the following proposition.

Proposition 1. Let $\sigma : [0, 1] \rightarrow SE(2)$ be a path lying in a connected component of $\mathcal{T}_i$. Then σ is a projection via π_p of some finite-length manipulation path.

Proof. First note that $\mathcal{T}_i = \bigcup_{\mathcal{F} \in \mathcal{F}} \pi_p(\text{int}(\mathcal{F}))$. Let each projection $\pi_p(\text{int}(\mathcal{F}))$ be denoted by $\mathcal{E}$ and $\mathcal{E}$ the collection of such sets. Since π_p is an open map, each $\mathcal{E}$ is open. Therefore, $\mathcal{E}$ is an open covering of $\mathcal{T}_i$.

Since σ lies entirely in some connected component of $\mathcal{T}_i$, there exists a subcollection $\mathcal{E}'$ of $\mathcal{E}$ that covers σ . Let $\mathcal{I}$ be the collection of open intervals where each interval, $\mathcal{I}$, corresponds to a domain of the path σ such that the path segment $\sigma(\mathcal{I})$ lies entirely in some open set $\mathcal{E}$.

Now we have that $\mathcal{I}$ is an open covering of $[0, 1]$. Since $[0, 1]$ is compact [48], there exists a finite subcollection of $\mathcal{I}$ that also covers $[0, 1]$. This means that the path σ consists of a finite number of segments where each segment lies entirely in an open set $\mathcal{E}$ and hence is a projection of a transfer path. Therefore, we can conclude that the path σ is a projection of a finite-length manipulation path. ■

⁵We define the length of a manipulation path as in [6]. Generally speaking, the length is proportional to the number of *necessary* regrasping operations along the manipulation path.

However, since the proof of compactness of $[0, 1]$ is not constructive [49], the above proposition does not give us a way to construct a finite-length manipulation path associated to a given object path σ . Note also that since the proof of Reduction Property⁶ given in [50] also relied on the Heine-Borel covering theorem, it also does not provide a practical way to construct a manipulation path.

To explicitly construct an algorithm that, given an object path σ , computes in finite time an associated finite-length manipulation path, we need a set of additional assumptions. The idea behind the construction of the algorithm is that from the uncountable collection $\mathcal{E}$, we need to be able to extract from $\mathcal{E}$ a *countable* (possibly infinite) subcollection that still covers the given path σ . Then from the countable subcollection, we can then iterate through combinations of its members until we find one that covers σ . The Heine-Borel covering theorem helps guarantee that these iterations will eventually terminate in finite time.

Before we proceed to state assumptions, we present the following result. Define the set $\mathcal{F}(g)$ as the union of all element $\mathcal{F}$ of $\mathcal{F}$ where the grasp associated with any composite configuration $c \in \mathcal{F}(g)$ is specified by the bimanual grasp parameter vector g (see Section III). Since there may exist multiple IK solutions associating with one grasp, we may categorize the set $\mathcal{F}(g)$ further into a number of subsets according to classes of the associated IK solution [51], [19]. We write $\mathcal{F}(g, k), k \in K$ to refer to the set $\mathcal{F}$ with a specific grasp g and that any $c \in \mathcal{F}(g, k)$ has the IK solution in the same class as other configurations. Note the according to [51], the index set K is bounded.

Consider a set $\mathcal{E}(g, k)$, defined as the projection via π_p of $\mathcal{F}(g, k)$.

Lemma 1. *An object path $\sigma : [0, 1] \rightarrow SE(2)$ lies entirely in a connected component of $\mathcal{E}(g, k)$ if and only if it is a projection of a transfer path.*

Proof. The result follows directly from the definition of a single-transfer connected component. ■

B. Assumptions

Now we present a set of assumptions as follows.

Assumption 1. *For any object path $\sigma : [0, 1] \rightarrow SE(2)$. The intersection between σ and the set $\mathcal{E}(g, k)$, for any grasp g and IK class index k , consists of finitely many path segments and the domain of the path parameter for each segment is computable.*

In usual manipulation planning settings, environments are relatively controlled such that they should not contain physical obstacles of extremely odd geometries that would eventually result in the set $\mathcal{E}(g, k)$ being divided into infinitely many connected components. Furthermore, the robot singularity set is not likely to divide the feasible configuration space into infinitely many connected components as well. This is true, for example, for a class of *generic* manipulators whose singularity

sets consist of finite *smooth manifolds* [52]. However, the above assumption is still necessary to ensure that each connected component of $\mathcal{E}(g, k)$ is well-behaved, in the sense that a finite number of components would not result in infinitely many segments.

The second assumption is stated as follows.

Assumption 2. *Given an object path σ contained in a connected component of $\mathcal{E}(g)$, where $g = [g_1^\top g_2^\top]^\top$, $g_1 = [l_1 a_1 b_1 \delta_1]^\top$, and $g_2 = [l_2 a_2 b_2 \delta_2]^\top$. There exists a lower bound $\epsilon > 0$, which may depend on σ , such that all $\mathcal{E}(g')$ also contain σ , where $g' = [g_1'^\top g_2'^\top]^\top$, $g_1' = [l_1 a_1 b_1 (\delta_1 + \Delta_1)]^\top$, $g_2' = [l_2 a_2 b_2 (\delta_2 + \Delta_2)]^\top$, and $0 < \Delta_1, \Delta_2 \leq \epsilon$.*

This means that if the robots can grasp the object with the bimanual grasp g and then trace the object path σ , the robots can also grasp the object and trace the same object path with some *nearby grasps*.

Now suppose, without loss of generality, that the domain of the grasp parameter δ is normalized to $(0, 1)$ for each grasp class. Next, we define $\mathcal{B}(d)$, with $d = (d_1, d_2) \in (0, 1) \times (0, 1)$, as the collection of all $\mathcal{E}(g, k)$ such that the grasp parameters $\delta_1 = d_1$ and $\delta_2 = d_2$. Consider the set $\mathcal{A}$ defined by

$$\mathcal{A} = \{\mathcal{B}(h_1), \mathcal{B}(h_2), \mathcal{B}(h_3), \dots\}, \quad (1)$$

where $\{h_1, h_2, h_3, \dots\}$ is a two-dimensional low-discrepancy sequence, such as the sequence introduced in [53]. This set $\mathcal{A}$ is then basically an enumeration of $\mathcal{B}(d)$ at values d from such a sequence. Thanks to Assumption 2, we have that there exists an integer N such that the subset $\mathcal{A}'$ of the first N terms of $\mathcal{A}$ covers σ .

Assumption 3. *The connectivity of a set $\mathcal{T}_i$ for all i can be determined empirically, e.g., by discretization.*

The above discretization can be easily done on the three parameters (x, y, θ) parameterizing the placement. Ranges of the parameters x and y are determined by the user while θ ranges from 0 to 2π . After obtaining the set of discretized coordinates, one tests at each point if the object is collision-free and graspable by the robots. Fig. 4(a) shows the scene in which we tested the connectivity of $\mathcal{T}_i$. The set $\mathcal{T}_i$ is visualized in Fig. 4(b) by being superimposed into the scene.

The idea here is that once the connectivity of the set $\mathcal{T}_i$ is determined, we can treat different connected components of $\mathcal{T}_i$ (if any) as different placement classes when computing a certificate. Therefore, we shall suppose in the sequel that each $\mathcal{T}_i$ consists of one connected component.

C. Algorithm

Consider next the algorithm listed in Algorithm 1. To compute a TypeA motion to move the object along a path segment $\sigma(s), s \in [t, t']$ (in `ComputeCCMotion` line 7), one starts with the initial IK solution of the robot grasping the object at $\sigma(t)$. Note that the grasp as well as the IK solution can be determined uniquely from the element A . Since $[t, t'] \subset (a, b)$, it is possible for one to use a differential IK algorithm [54] to

⁶Reduction Property states that two configurations in the same connected component of $\mathcal{G} \cap \mathcal{P}$ are connected by some manipulation path.

Fig. 4: (a) The chair at a transformation in the placement class of interest. (b) The set $\mathcal{T}_i$ plotted out in 3D, shown in green. Each 3D point represents (x, y, θ) at which the object can be grasp by the robots. The vertical axis is θ , varying from 0 to 2π . One can see from the figure that $\mathcal{T}_i$ is indeed connected.

solve for remaining IK solutions along the path (according to Lemma 1).

Based on assumptions presented in the previous Section, we have the following proposition.

Proposition 2. *Given an object path σ lying entirely in $\mathcal{T}_i$, Algorithm 1, with a complete motion planner deployed as a local planner, will terminate in finite time with a finite-length manipulation trajectory whose projection via π_p is σ .*

Proof. Consider first the function Planner in Algorithm 1. Since the set $\mathcal{A}$ provides a countable open covers of the path σ , and since the closed interval $[0, 1]$ of the path parameter is compact, there exists a finite subcover of σ . Therefore, only a finite number of iterations is required before the while loop (in Planner line 2) terminates. In each iteration of the while loop, one call to the function CheckCover is made. From Assumption 1 of finite intersections with σ , one would require finite time to verify intersections of each element A of $\mathcal{A}'$ with the path σ . Therefore, CheckCover always terminates in finite time.

Since $\mathcal{A}'$ is a finite subcover of the path σ , the while loop in ComputeCCMotion will eventually terminate. This concludes the proof. ■

Note here that instead of a complete motion planner, a *probabilistically* complete motion planner can be used in Algorithm 1, ComputeCCMotion, line 7. In this case, the overall planner will be *probabilistically* certified-complete (see Section VII-A for a more detailed discussion).

V. GENERATING TRAJECTORIES BETWEEN DIFFERENT PLACEMENT CLASSES

It follows from Proposition 1 that if we have one TypeB trajectory that starts in some placement class $\mathcal{P}_i$ and ends in some other placement class $\mathcal{P}_j$, any pair of composite configurations in $\mathcal{P}_i \cup \mathcal{P}_j$ are also manipulation path-connected. In the case when there are only two placement classes available,

---

Algorithm 1: In-placement manipulation planner

Planner(σ):

```

1  $\mathcal{A}' \leftarrow \{\}, j \leftarrow 1$ 
2 while True do
3   Append the  $j^{\text{th}}$  element of  $\mathcal{A}$  to  $\mathcal{A}'$ ,  $j++$ 
4   covered  $\leftarrow$  CheckCover( $\sigma, \mathcal{A}'$ )
5   if covered then
6     break
7 return ComputeCCMotion( $\sigma, \mathcal{A}'$ )
```

ComputeCCMotion($\sigma, \mathcal{A}'$):

```

1  $t \leftarrow 0$ , reached  $\leftarrow$  False
2 traj  $\leftarrow$  an empty trajectory
3 while not reached do
4   Select an element  $A$  from  $\mathcal{A}'$  that contains  $\sigma(t)$ .
   Let  $(a, b)$  be the domain of  $\sigma$  covered by  $A$ 
   such that  $t \in (a, b)$ 
5   Select an element  $A'$  from  $\mathcal{A}'$  that contains
    $\sigma(b)$ . Let  $(a', b')$  be the domain of  $\sigma$  covered
   by  $A'$  such that  $b \in (a', b')$ .
6   Select  $t' \in (a', b)$ 
7   Compute a closed-chain motion for
    $\sigma(s)$ ,  $s \in [t, t']$  and other necessary transit
   motions for regrasp
8   Append the compute trajectories to traj
9   if  $t' == 1$  then
10     break
11    $t \leftarrow t'$ 
12 return traj
```

---

any TypeB path between the two placement classes then guarantees the existence of a solution to any feasible manipulation query. With n_p placement classes available, one only needs a minimum of $n_p - 1$ TypeB trajectories between different pairs of placement classes in order to guarantee the existence of a solution to any feasible query. Therefore, we define the notion of a *certificate* as follows.

Definition 11. *A certificate is a set of transfer paths that spans all the placement classes.*

One can think of placement classes as nodes in a graph. A certificate is then analogous to a set of edges that contains a spanning tree's edges. Although $n_p - 1$ transfer paths are sufficient to guarantee the existence of solutions to any manipulation query, the more transfer paths one has (between distinct pairs of placement classes) can contribute to higher quality of solutions since the system may need to visit a fewer number of intermediate placement classes before reaching the desired placement class.

Since the process of computing a certificate needs to be done only once per problem setting, we suppose that this computation can be done offline and the computation time is not a limiting factor. Therefore, one may aim at generating all $n_p C_2 = n_p(n_p - 1)/2$ transfer paths connecting all possible different pairs of placement classes.

Given a pair of placement classes $\mathcal{P}_i$ and $\mathcal{P}_j$, we divide the process of generating a TypeB trajectory into two main

parts: 1) generating a closed-chain query; and 2) solving a closed-chain query.

A. Generating a Closed-Chain Query

Randomly sampling two object transformations, one from each placement class, may have a relatively low probability that the resulting closed-chain query is solvable. This is mainly due to the fact that the closed-chain constraint greatly reduces the range of motions of the system. To deal with this issue, we propose a heuristic to help generate closed-chain queries that are likely solvable. The idea behind this is that since the bimanual manipulation system can exhibit a very limited range of motions, queries should be generated such that they intuitively do not require a large range of robot motions to solve them.

Recall that object transformations in any placement class can be parameterized by three parameters (Section IV). The problem of generating closed-chain queries is then boiled down to how one generates the three parameters for the start and goal transformations. We first define a *manipulation point* (x_m, y_m) on the support surface. The position parameters (x, y) of the transformations to be generated will be assigned to be this point. Doing so greatly simplifies query generation while not drastically reduce the possibilities of the queries generated since the motion range of the system is already very limited. The manipulation point may be assigned to be on the middle line passing between the two robots (the green line in Fig. 5), roughly speaking, to maximize the reachability of the two robots.

Then the rotation parameter θ can be computed as follows. Let f_j denote the face of $\mathcal{H}$ corresponding to placement class $\mathcal{P}_j$ (the goal placement class). Let $\mathbf{n}_j$ be a normal vector pointing outwards from the face f_j (see the red arrow in Fig. 5(a)). The desired value of θ is such that the projection of $\mathbf{n}_j$ onto the support surface is parallel to the line l . The reason behind this is that once the object is arranged as mentioned, the expected closed-chain motion to move the object from placement $\mathcal{P}_i$ to $\mathcal{P}_j$ will be a relatively easy flipping motion, as shown by the blue arrow in Fig. 5(a). After the start transformation has been computed, the goal transformation can be computed accordingly (Fig. 5(b)).

Apart from the two transformations T_s and T_g , we may also include into the query a grasp g together with associated IK solutions at T_s and/or T_g . Generating a grasp g is straightforward since it can be sampled from a set of grasps available at both transformations.

Computing associated IK solutions at T_s and/or T_g , however, is non-trivial when the query is to be solved via a bidirectional planner. This is because given two composite configurations in the same grasp class, there is no known way to completely determine if they belong to the same connected component of $\mathcal{G}$. Generally, one set of IK solutions of robots grasping the object corresponds to one connected component of $\mathcal{G}$ called self-motion manifold [51]. To choose IK solutions of robots grasping the object at both T_s and/or T_g , we rely on an ad hoc heuristic since, to our knowledge, there is currently no known generalized way to do so.

Fig. 5: (a) The start transformation of the chair. The chair is placed around the center line (green) such that the normal vector (the dotted arrow) of the face f_j is pointing in the direction of the center line. By arranging the chair as shown, the chair is only needed to be flipped in the direction depicted by the curved arrow. (b) The goal transformation of the chair.

Finally, note that this two-stage approach in generating and solving queries may proceed in iterations. If a generated query is not solvable within the given time, one can generate a new query by defining a new manipulation point, e.g., by adding some small perturbation to the point.

B. Solving a Closed-Chain Query

Our closed-chain motion planner (CCPlanner) is adapted from a bidirectional RRT planner [41]. In particular, we build two trees $\mathcal{T}_a$ and $\mathcal{T}_b$, each one is a data structure storing vertices. A vertex $\mathcal{V}$ keeps information of a composite configuration, its parent on the tree, as well as a closed-chain trajectory connecting itself and its parent. The algorithm for CCPlanner is listed in Algorithm 2.

Algorithm 2: Closed-chain motion planner

```

CCPlanner( $T_s, T_g, (q_{1s}, q_{2s}), (q_{1g}, q_{2g})$ ):
1  $\mathcal{T}_a \leftarrow \text{InitializeTree}(T_s, (q_{1s}, q_{2s}))$ 
2  $\mathcal{T}_b \leftarrow \text{InitializeTree}(T_g, (q_{1g}, q_{2g}))$ 
3 while time is not exhausted do
4    $T \leftarrow \text{SampleSE3}()$ 
5   if Extend( $\mathcal{T}_a, T$ ) then
6     if Connect( $\mathcal{T}_a, \mathcal{T}_b$ ) then
7       return ExtractTrajectory( $\mathcal{T}_a, \mathcal{T}_b$ )
8   Swap( $\mathcal{T}_a, \mathcal{T}_b$ )
9 return None

Extend( $\mathcal{T}, T$ ):
1 for  $\mathcal{V}_{\text{near}}$  in KNN( $\mathcal{T}, T$ ) do
2    $T_{\text{new}} \leftarrow \text{Threshold}(\mathcal{V}_{\text{near}}, T)$ 
3    $\mathcal{P} \leftarrow \text{Interpolate}(\mathcal{V}_{\text{near}}.\text{config}, T_{\text{new}})$ 
4   if  $\mathcal{P}$  is not None then
5      $\mathcal{V}_{\text{new}} \leftarrow \text{Vertex}(\mathcal{C}, \mathcal{P})$ 
6      $\mathcal{T}.\text{Add}(\mathcal{V}_{\text{near}}, \mathcal{V}_{\text{new}})$ 
7   return True
8 return False

```

CCPlanner accepts the start and goal transformations, T_s and T_g and IK solutions at T_s and T_g as its inputs. It starts by initializing two trees, i.e., creating root vertices

Fig. 6: Objects used in our experiments of certificate computation (not to scale).

TABLE I: Computation Time for Computing Certificates and Solving Queries for the Objects in Fig. 6.

			Object 1	Object 2	Object 3	Object 4	Object 5	Object 6
# placement classes			4	3	6	6	8	6
# grasp classes			1849	2601	2401	3364	1089	784
Computing Certificates	total computation time	avg. (s.)	415.97	281.18	867.71	1030.44	1374.27	574.54
		std. (s.)	141.49	98.48	212.54	288.06	342.11	120.29
	*planning (s.)		154.36	94.07	604.19	556.92	582.54	199.63
	*grasp check (s.)		101.69	76.27	202.90	227.51	157.90	154.11
	*shortcutting (s.)		14.51	7.29	32.76	33.55	34.83	23.58
	# TypeB trajectories (avg.)		4.00	2.00	11.02	10.84	11.98	7.40
	success rate		1.0	1.0	1.0	1.0	1.0	1.0
Solving Queries	total computation time	avg. (s.)	103.65	45.08	127.45	234.54	109.99	185.08
		std. (s.)	18.04	14.93	29.93	55.96	24.37	31.55
	# TypeB trajectories (avg.)		2.0	2.0	2.0	2.0	2.0	2.0

*Row 3 – 5 of “Computing Certificate” category are planning time, grasp checking time, and shortcutting time averaged over successful closed-chain queries.

storing configuration information. In each planning iteration, CCPlanner randomly samples an object transformation matrix T . Then the planner tries to extend a tree towards it. Upon a successful extension the planner tries to connect two trees together. If the connection attempt is successful, the closed-chain path is extracted from the tree and returned. Some key functions are described in details below.

- **SampleSE3:** A transformation matrix (an element of $SE(3)$) is generated by separately sampling rotational and translational parts. A rotation matrix is uniformly sampled from $SO(3)$ via the method proposed in [55]. A translation vector is sampled uniformly from the user defined range. Then the planner checks if IK solutions exist for both robots. The function iterates until it finds a reachable transformation matrix⁷.

- **Extend:** To extend a tree towards a given transformation T , we search in the tree the set of k vertices whose transformation matrices are nearest to T (via KNN with k defined by the user). The distance metric used is defined in Section III. Then for each vertex, we generate a new transformation T_{new} in the direction from $\mathcal{V}_{\text{new}}.\text{config}.T$ to T such that the distance between T_{new} and $\mathcal{V}_{\text{new}}.\text{config}.T$ does not exceed some pre-defined step size. Then we construct a closed-chain trajectory

⁷Since the arrangement of the two robots is supposedly in such a way that facilitates the manipulation procedure, one can suppose that the volume of the set of transformation matrices that are reachable by the two robots relative to the subset of $SE(3)$ from which a transformation is sampled is non-zero. Therefore, the probability that the generated transformation matrix can satisfy the closure constraint is positive.

connecting the two transformation via *Interpolate*. If the trajectory generation is successful, a new vertex, $\mathcal{V}_{\text{new}}$, is created and added to the tree.

- **Interpolate:** To generate a closed-chain motion, we generate an $SE(3)$ trajectory first by using the method in [47] for the rotational part and using polynomial interpolation for the translational part. Then the trajectory is discretized into small time steps. IK solutions of the two robots are computed at each time step using the differential IK method [54]. In case the differential IK runs into a singularity, the $SE(3)$ trajectory is rejected, and the connection attempt is considered unsuccessful.

Apart from being less complicated implementation-wise, this method gives an exact parameterization of the object trajectory. One can then incorporate various types of constraints into the object trajectory by means of time-parameterization to obtain time-optimal trajectory with respect to the constraints (see [56] for more details on time-optimal path parameterization (TOPP)). Examples of applicable constraints are velocity and acceleration limits for rigid body motions [57] and dynamic grasp stability.

Nevertheless, the user can utilize their trajectory generation method of choice. For example, if singularity avoidance is a major concern, one may use a specialized method to directly deal with singularities (see, e.g., [44]). In case an exact parameterization of a closed-chain trajectory is available, one may also use the TOPP method for redundantly-actuated systems [58] (which is the case for bimanual systems) to

enforce the aforementioned constraints along the trajectory.

- **Connect**: After a successful tree extension, the planner will attempt to connect the newly added vertex to the other tree. The details of procedure are mainly similar to **Extend**, except that this function does not include **Threshold**.

VI. EXPERIMENTAL RESULTS

The planner and all related functions were implemented in Python. We used OpenRAVE [59] as a simulation environment. The robots were two identical 6-DOF industrial manipulators Denso VS-060. Each one was equipped with an industrial-grade 2-finger Robotiq gripper 85. The planning environment is as shown in Fig. 4. All simulations were run on a 2.4 GHz desktop.

A. Computing Certificates

To validate our certificate issuing planner, we ran the planner to compute certificates for a set of objects. All objects, listed in Fig. 6, were furniture pieces, which were relatively large such that they needed two robots grasping in order to move them. For each object, we repeated certificate generation for 50 times. When computing each transfer path, we also had the following additional computations: 1) grasp equilibrium checking 2) closed-chain trajectory shortcutting (200 iterations). Statistics collected from the runs are reported in Table I in “Computing Certificate” category. Here are a few things we would like to point out:

- The planner may spend up to around 45% of the total time generating and planning unsuccessful queries. This is because there currently exists no definite method to check that the generated IK solutions associated with the start and goal transformations belong to the same connected component (self-motion manifold).
- Due to the large number of grasp classes, solving a bimanual manipulation query via, for example, the three-dimensional extension of the high-level Grasp-Placement Graph [6] could potentially be rendered infeasible. Consider Object 2 (Fig. 6(b)), for example, which has 3 placement classes and 51 *unimanual* grasp classes. While the two-dimensional Grasp-Placement Graph has 104 vertices and 2028 edges, its three-dimensional counterpart contains over 3,000 vertices with over 6 million edges.
- For each run, the certificate computation procedure is considered successful if the computed transfer paths span all the placement classes. Although the number of successfully planned TypeB paths varied slightly among different runs, the resulting set of TypeB paths still spanned all placement classes in every run.
- In the current implementation, we check grasp equilibrium by solving a linear program at each discretized time step along a closed-chain trajectory. This approach is, however, time-consuming and restrictive in that it only guarantees static equilibrium⁸. One possible improvement is by formulating contact constraints in terms of inequalities in path

⁸Trajectories with static equilibrium are only guaranteed to be executable at *arbitrarily slow* speed.

Fig. 7: Key object transformations along the computed bimanual manipulation trajectory solving the given query.

parameters and its derivatives [60] by utilizing cone double-description (CDD) method [61]. Then one can *time-optimally* parameterize the trajectory such that it moves as fast as possible while respecting all the constraints (see [56] and references therein for more details). Another possible improvement is to directly plan dynamically feasible motions using a planner based on Admissible Velocity Propagation [62].

B. Solving Bimanual Queries

First, for each object listed in Fig. 6, we hand-picked two placement classes that do not have any direct connection via a TypeB trajectory (information provided by a certificate) and generated a pair of object transformations T_s and T_g from each of the placement classes. Then we repeated solving each query $Q = (T_s, T_g)$ for 50 times. Statistics collected from the runs are reported in Table I in “Solving Queries” category. One of the main factors that causes variations in the running time is the geometry of each object. Larger objects, for example, have *narrower* free space to navigate on the support surface. Furthermore, with larger objects, it is also more difficult for robots to move around and change grasps.

A solution manipulation path to any of the above queries needs to be at least of length 5. However, it is not the case here since a TypeA trajectory, which connects two TypeB trajectories from a certificate, will always contain regrasp operations. This is because the grasps used in the two TypeB trajectories are always different. Therefore, if a planner with an extension of the high-level Grasp-Placement Graph is used, it will spend a considerable amount of time invalidating manipulation paths of shorter length, hence not practical. A general manipulation planner such as a Random-MMP [30] is not likely to terminate with a solution within a reasonable amount of time as well since it needs to sample correctly a relative long sequence of transit and transfer.

Apart from the simulation results, we also successfully carried out a hardware experiment. We constructed a query for Object 2 (Fig. 6(b)) such that there is no direct TypeB

trajectory connecting the two placement classes. The scene with the start and goal transformations of the object provided by the query are shown in Fig. 7(a) and Fig. 7(f), respectively. The computed solution to this query consists of three TypeA trajectories ((a)→(b), (c)→(d), and (e)→(f)) and two TypeB trajectories ((b)→(c) and (d)→(e)).

We used a leader-follower scheme to perform the closed-chain motions on the hardware platform. The leader robot was purely position-controlled while the follower robot had an additional compliance so as to account for model inaccuracies, see [2] for details. The high grip force exerted by the grippers, in combination with this compliant control scheme, prevents any slippage during the execution of the closed-chain motions. The video of the robots executing the motion solving the query can be found at <https://youtu.be/4DcMwr2xxrQ>.

VII. DISCUSSION AND CONCLUSION

A. Discussion

1) *Completeness Property*: The proposed algorithm has a property so-called *certified-completeness*. A computed certificate certifies that the algorithm will find a solution to any bimanual manipulation query, in the environment in which the certificate is computed, whenever one exists. This completeness is, however, up to the completeness of the motion planner deployed as a local planner. To see this, consider the computation of TypeA trajectories according to Algorithm 1. In each iteration of the function `ComputeCCMotion` a segment of the object path σ that lies in a single-transfer connected set is identified. A transfer trajectory that moves the object along the segment of σ can then be computed accordingly by, for example, computing IK solutions for the two robots associated with discretized object configurations along the segment⁹. This closed-chain motion computation will always succeed since the object path segment lies entirely in a single-transfer connected set. However, one also needs to compute transit trajectories to connect consecutive transfer trajectories, i.e., to change the robot grasps. Therefore, the completeness of the overall algorithm will be affected by the completeness of the planner deployed to compute these transit trajectories. When a probabilistically complete motion planner is used, the proposed manipulation planning algorithm will be *probabilistically certified-complete*.

2) *Complexity of Algorithm*: We analyze computational complexity of the certificate computation phase and the query phase in terms of the number of motion planning queries that needs to be solved. Consider the certificate computation phase. Let n_p be the number of all placement classes, $S(i, j)$ the event that a closed-chain query for moving the object between placement classes i and j , $i \neq j$, is successfully solved. Suppose there exists a probability p such that $P[S(i, j)] \geq p$, for all $i \neq j$. The value of p may be estimated empirically. Let X be a random variable representing the number of consecutive failure before the first successfully solved closed-chain query. From direction calculation, one has $E[X] \leq 1/p - 1$. Therefore, for any pair of placement classes, the number of closed-chain

queries one is expected to solve until encountering the first feasible query is $O(1/p)$. In the worst case scenario, one needs to consider all possible $n_p(n_p - 1)/2$ pairs of placement classes in order to compute a certificate. Therefore, the number of closed-chain queries one needs to solve to compute a certificate is $O(n_p^2/p)$. In the query phase, TypeA trajectories are created via Algorithm 1. Since there are n_p placement classes, the number of TypeA trajectory queries that needs to be solved is $O(n_p)$.

Time complexity of motion planning when solving each open-chain or closed-chain query is affected largely by constraints such as collision avoidance (with both physical obstacles and virtual obstacles such as joint limits). These constraints reduces the configuration space expansiveness in a non-trivial manner and therefore, according to the analysis in [63], lead to longer running time. In addition, more restrictive constraints could lead to higher number of regrasp operations required in a TypeA trajectory. This, in turn, results in more transit trajectory queries to be solved.

Complexity may be analyzed in greater details by taking into account complexity of each primitive procedure in motion planning such as nearest neighbor search, collision checking, or IK computation as was done in [32]. However, these results are theoretical. In practice, there are a number of factors that affect the actual running time of the planning algorithm, for example, the constants omitted in those complexity expressions. Careful implementation can also play a crucial role in reducing running time. As an example, consider transit trajectory generation procedure. When the two n -DOF robots change their grasps, the motion planning problem can be formulated as either one $2n$ -DOF problem or two sequential n -DOF problems. This poses a trade-off between motion planning time and motion execution time and it is therefore up to the user to strike a balance between the two.

B. Limitations and Future Works

Various improvements can be made to the planning process. For example, currently when computing TypeB trajectories (see Section V), the generated closed-chain query will fail when the start and goal configurations live in different single-transfer connected components. To improve success rate of certificate computation, one may deploy a closed-chain planner presented in [2], which can plan motions across different connected components, to solve such queries.

Furthermore, the proposed algorithm depends on the accuracy of its inputs, for example, the 3D model of the object, the relative transformation between the robots, the initial object pose, environment geometry, etc. Any discrepancies between the actual environment and the data are accounted for by the compliant control algorithm used in motion execution. However, as these discrepancies can propagate and be intensified throughout the motion execution process, they may eventually lead to failure of execution or cause damages to the manipulated object. One important future direction is then to investigate how the knowledge of data uncertainty can be incorporated into planning so as to plan more robust trajectories.

⁹Methods such as the differential IK method [54] can be used to ensure continuity of IK solutions.

C. Conclusion

In this paper, we first present a set of definitions and fundamentals of bimanual manipulation planning. In order to solve a bimanual manipulation query, it is essential for a planner to obtain information of connectivities between different connected components of the composite configuration space. We propose an algorithm that constructs a manipulation solution by generating and concatenating two types of trajectories: TypeA trajectories, which connect configurations in the same placement class, and TypeB trajectories, which connect configurations from different placement classes. The proposed algorithm is certified-complete. We provide a method to compute a *certificate* for a given object and environment. A certificate, once obtained, guarantees that the algorithm will find a solution to any feasible bimanual manipulation query for the object in that setting in finite time. Information contained in a certificate can be used to construct a solution trajectory. Simulation and experimental results illustrate the validity and capability of our algorithm to plan bimanual manipulation motions for various practical objects.

REFERENCES

- [1] J. H. Yakey, S. M. LaValle, and L. E. Kavraki, "Randomized path planning for linkages with closed kinematic chains," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 951–958, 2001.
- [2] Z. Xian, P. Lertkultanon, and Q.-C. Pham, "Closed-chain manipulation of large objects by multi-arm robotic systems," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 1832–1839, 2017.
- [3] T. Siméon, J.-P. Laumond, J. Cortés, and A. Sahbani, "Manipulation planning with probabilistic roadmaps," *The International Journal of Robotics Research*, vol. 23, no. 7–8, pp. 729–746, 2004.
- [4] M. Gupta, J. Muller, and G. S. Sukhatme, "Using manipulation primitives for object sorting in cluttered environments," *IEEE Transactions on Automation Science and Engineering*, vol. 12, no. 2, pp. 608–614, 2015.
- [5] W. Wan, M. M. Mason, R. Fukui, and Y. Kuniyoshi, "Improving regrasp algorithms to analyze the utility of work surfaces in a workcell," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, 2015, pp. 4326–4333.
- [6] P. Lertkultanon and Q.-C. Pham, "A single-query manipulation planner," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 198–205, 2015.
- [7] S. M. LaValle, *Planning algorithms*. Cambridge University Press, 2006.
- [8] S. Caron, Q.-C. Pham, and Y. Nakamura, "Completeness of randomized kinodynamic planners with state-based steering," *Robotics and Autonomous Systems*, vol. 89, pp. 85–94, 2017.
- [9] Y. Koga and J.-C. Latombe, "Experiments in dual-arm manipulation planning," in *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*, 1992, pp. 2238–2245.
- [10] J.-P. Saut, M. Gharbi, J. Cortés, D. Sidobre, and T. Siméon, "Planning pick-and-place tasks with two-hand regrasping," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, 2010.
- [11] B. Balaguer and S. Carpin, "Bimanual regrasping from unimanual machine learning," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, 2012, pp. 3264–3270.
- [12] K. Harada, T. Tsuji, and J.-P. Laumond, "A manipulation motion planner for dual-arm industrial manipulators," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, 2014, pp. 928–934.
- [13] W. Wan and K. Harada, "Developing and comparing single-arm and dual-arm regrasp," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 234–250, 2016.
- [14] M. J. Hwang, D. Y. Lee, and S. Y. Chung, "Motion planning of bimanual robot for assembly," in *Systems, Man and Cybernetics, 2007. ISIC. IEEE International Conference on*, 2007, pp. 240–245.
- [15] D. Surdilovic, Y. Yakut, T.-M. Nguyen, X. B. Pham, A. Vick, and R. Martin-Martin, "Compliance control with dual-arm humanoid robots: Design, planning and programming," in *Humanoid Robots (Humanoids), 2010 10th IEEE-RAS International Conference on*, 2010, pp. 275–281.
- [16] R. L. A. Shauri and K. Nonami, "Assembly manipulation of small objects by dual-arm manipulator," *Assembly Automation*, vol. 31, no. 3, pp. 263–274, 2011.
- [17] A. Edsinger and C. C. Kemp, "Two arms are better than one: A behavior based control system for assistive bimanual manipulation," in *Recent Progress in Robotics: Viable Robotic Service to Human*, ser. Lecture Notes in Control and Information Sciences. Springer Berlin Heidelberg, 2008, vol. 370, pp. 345–355.
- [18] D. Kruse, J. T. Wen, and R. J. Radke, "A sensor-based dual-arm tele-robotic-system," *IEEE Transactions on Automation Science and Engineering*, vol. 12, no. 1, pp. 4–18, 2015.
- [19] M. Gharbi, J. Cortés, and T. Siméon, "Roadmap composition for multi-arm systems path planning," in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, 2009, pp. 2471–2476.
- [20] N. Vahrenkamp, D. Berenson, T. Asfour, J. J. Kuffner, and R. Dillmann, "Humanoid motion planning for dual-arm manipulation and re-grasping task," in *Intelligent Robots and Systems (IROS), 2009 IEEE/RSJ International Conference on*, 2009, pp. 2464–2470.
- [21] N. Vahrenkamp, M. Do, T. Asfour, and R. Dillmann, "Integrated grasp and motion planning," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 2010, pp. 2883–2888.
- [22] S. K. Agrawal and S. Shirumalla, "Planning motions of a dual-arm free-floating manipulator keeping the base inertially fixed," *Mechanism and Machine Theory*, vol. 30, no. 1, pp. 59–70, 1995.
- [23] F. Zacharias, D. Leidner, F. Schmidt, C. Borst, and G. Hirzinger, "Exploiting structure in two-armed manipulation tasks for humanoid robots," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, 2010, pp. 5446–5452.
- [24] M. Gharbi, J. Cortés, and T. Siméon, "A sampling-based path planner for dual-arm manipulation," in *Advanced Intelligent Mechatronics, 2008. AIM 2008. IEEE/ASME International Conference on*, 2008, pp. 383–388.
- [25] B. Cohen, S. Chitta, and M. Likhachev, "Single- and dual-arm motion planning with heuristic search," *The International Journal of Robotics Research*, vol. 33, no. 2, pp. 305–320, 2013.
- [26] G. Lee, T. Lozano-Pérez, and L. P. Kaelbling, "Hierarchical planning for multi-contact non-prehensile manipulation," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, 2015, pp. 264–271.
- [27] M. Pflueger and G. S. Sukhatme, "Multi-step planning for robotic manipulation," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, 2015, pp. 2496–2501.
- [28] C. Smith, Y. Karayiannidis, L. Nalpantidis, X. Gratal, P. Qi, D. V. Dimarogonas, and D. Kragic, "Dual arm manipulation—a survey," *Robotics and Autonomous Systems*, vol. 60, no. 10, pp. 1340–1353, 2012.
- [29] 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, June 2010.
- [30] K. Hauser and V. Ng-Thow-Hing, "Randomized multi-modal motion planning for a humanoid robot manipulation task," *The International Journal of Robotics Research*, vol. 30, no. 6, pp. 678–698, 2011.
- [31] P. S. Schmitt, W. Neubauer, W. Feiten, K. M. Wurm, G. V. Wichert, and S. Burgard, "Optimal, sampling-based manipulation planning," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, 2017, pp. 3426–3432.
- [32] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [33] T. Bretl, S. Lall, J.-C. Latombe, and S. Rock, *Multi-Step Motion Planning for Free-Climbing Robots*, ser. Springer Tracts in Advanced Robotics. Springer Berlin Heidelberg, 2005, vol. 17, pp. 59–74.
- [34] N. Chavan-Daffe, M. T. Mason, H. Staab, G. Rossano, and A. Rodriguez, "Two-phase gripper to reorient and grasp," in *Automation Science and Engineering (CASE), 2015 IEEE International Conference on*, 2015, pp. 1249–1255.
- [35] P. Tournassoud, T. Lozano-Pérez, and E. Mazer, "Regrasping," in *Robotics and Automation. Proceedings. 1987 IEEE International Conference on*, vol. 4, 1987, pp. 1924–1928.
- [36] F. Röhrdanz and F. M. Wahl, "Generating and evaluating regrasp operations," in *Robotics and Automation, 1997. Proceedings., 1997 IEEE International Conference on*, vol. 3, 1997, pp. 2013–2018.
- [37] S. A. Stoeter, S. Voss, N. P. Papanikolopoulos, and H. Mosemann, "Planning of regrasp operations," in *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, vol. 1, 1999.

- [38] J. Mirabel and F. Lamiroux, "Manipulation planning: addressing the crossed foliation issue," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, 2017, pp. 4032–4037.
- [39] W. Wan, E. C. H. Cheung, J. Pan, and K. Harada, "Optimizing the parameter of tilting surfaces in robotic workcells," in *Automation Science and Engineering (CASE), 2015 IEEE International Conference on*, 2015, pp. 593–599.
- [40] C. Cao, W. Wan, J. Pan, and K. Harada, "Analyzing the utility of a support pin in sequential robotic manipulation," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, 2016, pp. 5499–5504.
- [41] S. M. Lavalle and J. J. Kuffner, "Rapidly-exploring random trees: Progress and prospects," in *Algorithmic and Computational Robotics: New Directions*, B. R. Donald, K. M. Lynch, and D. Rus, Eds. A. K. Peters, 2001, pp. 293–308.
- [42] L. Han and N. M. Amato, "A kinematic-based probabilistic roadmap method for closed chain systems," in *International Workshop on the Algorithmic Foundations of Robotics*, 2000, pp. 233–246.
- [43] J. Cortés, T. Siméon, and J.-P. Laumond, "A random loop generator for planning the motions of closed kinematic chains using PRM methods," in *Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on*, 2002.
- [44] O. Bohigas, M. E. Henderson, L. Ros, M. Manubens, and J. M. Porta, "Planning singularity-free paths on closed-chain manipulators," *IEEE Transactions on Robotics*, vol. 29, no. 4, pp. 888–898, 2013.
- [45] L. Jaillet and J. M. Porta, "Path planning under kinematic constraints by rapidly exploring manifolds," *IEEE Transactions on Robotics*, vol. 29, pp. 105–117, 2013.
- [46] B. Kim, T. T. Um, C. Suh, and F. C. Park, "Tangent bundle rrt: A randomized algorithm for constrained motion planning," *Robotica*, vol. 34, no. 1, pp. 202–225, 2016.
- [47] F. C. Park and B. Ravani, "Smooth invariant interpolation of rotations," *ACM Transactions on Graphics*, vol. 16, no. 3, pp. 277–295, 1997.
- [48] J. R. Munkres, *Topology*, 2nd ed. Prentice Hall, 2000.
- [49] J. Cederquist and S. Negri, *A constructive proof of the Heine-Borel covering theorem for formal reals*, ser. Lecture Notes in Computer Science (LNCS). Berlin, Heidelberg: Springer Berlin Heidelberg, 1996, vol. 1158, pp. 62–75. [Online]. Available: http://dx.doi.org/10.1007/3-540-61780-9_62
- [50] R. Alami, J.-P. Laumond, and T. Siméon, "Two manipulation planning algorithms," in *WAFR Proceedings of the workshop on Algorithmic foundations of robotics*, 1994, pp. 109–125.
- [51] J. W. Burdick, "On the inverse kinematics of redundant manipulators: Characterization of the self-motion manifolds," in *Advanced Robotics: 1989*. Springer Berlin Heidelberg, 1989, ch. 3.
- [52] D. K. Pai and M. C. Leu, "Genericity and singularities of robot manipulators," *IEEE Transactions on Robotics and Automation*, vol. 8, no. 5, pp. 545–559, 1992.
- [53] S. R. Lindemann and S. M. LaValle, "Incremental low-discrepancy lattice methods for motion planning," in *Robotics and Automation, 2003. Proceedings. ICRA '03. IEEE International Conference on*, 2003, pp. 2920–2927.
- [54] B. Siciliano, L. Sciacivco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*, 1st ed., ser. Advanced Textbooks in Control and Signal Processing. Springer London, 2009, ch. Differential Kinematics and Statics, pp. 105–160.
- [55] J. J. Kuffner, "Effective sampling and distance metrics for 3D rigid body path planning," in *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, vol. 4, 2004, pp. 3993–3998.
- [56] Q.-C. Pham, "A general, fast, and robust implementation of the time-optimal path parameterization algorithm," *IEEE Transactions on Robotics*, vol. 30, pp. 1533–1540, 2014.
- [57] H. Nguyen and Q.-C. Pham, "Time-optimal path parameterization of rigid-body motions: applications to spacecraft reorientation," *Journal of Guidance, Control, and Dynamics*, vol. 39, no. 7, pp. 1667–1671, 2016.
- [58] Q.-C. Pham and O. Stasse, "Time-optimal path parameterization for redundantly actuated robots: a numerical integration approach," *IEEE/ASME Transactions on Mechatronics*, vol. 20, no. 6, pp. 3257–3263, 2015.
- [59] R. Diankov, "Automated construction of robotic manipulation programs," Ph.D. dissertation, Carnegie Mellon University, Robotics Institute, August 2010, available at http://www.programmingvision.com/rosen_diankov_thesis.pdf.
- [60] S. Caron, Q.-C. Pham, and Y. Nakamura, "Leveraging cone double description for multi-contact stability of humanoid robots with applications to statics and dynamics," in *Proceedings of Robotics: Science and Systems*, 2015.
- [61] K. Fukuda and A. Prodon, "Double description method revisited," in *Combinatorics and Computer Science*, ser. Lecture Notes in Computer Science (LNCS). Springer, Berlin, Heidelberg, 1996, vol. 1120, pp. 91–111.
- [62] Q.-C. Pham, S. Caron, P. Lertkultanon, and Y. Nakamura, "Admissible velocity propagation: Beyond quasi-static path planning for high-dimensional robots," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 44–67, 2016.
- [63] D. Hsu, J.-C. Latombe, and R. Motwani, "Path planning in expansive configuration spaces," in *Robotics and Automation, 1997. Proceedings. ICRA '97. IEEE International Conference on*, vol. 3, 1997, pp. 2719–2726.

Puttichai Lertkultanon was born in Bangkok, Thailand. He received the BE degree (First-Class Honours) in Electrical Engineering from Chulalongkorn University, Bangkok, Thailand, in 2013. He obtained a PhD in Robotics from Nanyang Technological University, Singapore, in 2017. He joined Mujin Inc., Tokyo, Japan as a Motion Planning Engineer in 2017.

Quang-Cuong Pham was born in Hanoi, Vietnam and grew up in Vietnam and then in France. He graduated from École Normale Supérieure, France, in 2007. He obtained a PhD in Neuroscience from Université Paris VI and Collège de France in 2009. In 2010, he was a Visiting Researcher at the University of São Paulo, Brazil. From 2011 to 2013, he was as a Fellow of the Japan Society for the Promotion of Science (JSPS), studying Robotics at the University of Tokyo. He joined NTU, Singapore as an Assistant Professor in 2013.