Decidability of Semi-Holonomic Prehensile Task and Motion Planning

Ashwin Deshpande, Leslie Pack Kaelbling, and Tomás Lozano-Pérez

Massachusetts Institute of Technology, Cambridge, Massachusetts, USA
ashwind,lpk,tlp@csail.mit.edu

Abstract. In this paper, we define *semi-holonomic controllability* (SHC) and a general task and motion planning framework. We give a perturbation algorithm that can take a prehensile task and motion planning (PTAMP) domain and create a *jointly-controllable-open* (JC-open) variant with practically identical semantics. We then present a decomposition-based algorithm that computes the reachability set of a problem instance if a controllability criterion is met. Finally, by showing that JC-open domains satisfy the controllability criterion, we can conclude that PTAMP is decidable.

1 Introduction

The last few decades of robotic planning have been dominated by sample-based techniques. Sample-based techniques are very useful tools to quickly find solutions in many domains. However, they suffer from the notable drawback that they cannot prove that a solution does not exist for a particular problem.

The existence of a probabilistically complete algorithm for a planning problem does not settle the question of whether a complete decision procedure, an algorithm that indicates whether a solution does or does not exist for any problem instance, exists. For "classic" motion planning, a holonomic robot among static obstacles, we know that exact algorithms exist for the general case [1, 2]. However, for motion planning in the presence of movable objects, the results are much more limited.

The formal treatment of the problem of planning among movable objects was initiated by Wilfong [3]. When the number of placements and grasps is finite, the problem can be shown to be decidable by building a manipulation graph consisting of a finite number of transfer and transit paths [4]. Decidability for continuous grasps and placements, but involving a single movable object, was shown by Dacre-Wright et al. [5]. More recently, decidability was shown for planning with two objects under restrictive geometries and dynamics [6].

In this paper, we consider a much more general version of planning in the presence of movable obstacles. We allow an arbitrary dimensional world with an arbitrary number of robots, objects, and obstacles, all with semi-algebraic geometries. We also assume that each robot can be holonomically controlled, and each object can be holonomically manipulated. In this manner, we can account for various continuous polynomial dynamics including translations, rotations, stretching, twisting, and morphing. We do restrict our attention to prehensile manipulation, where objects are rigidly attached to appropriate robots during manipulation. We call the resulting class of problems "prehensile task and motion planning" (PTAMP).

In the first section, we define a general task and motion planning framework capable of representing a large variety of planning problems including PTAMP. At the core of the framework is the concept of semi-holonomic controllability (SHC), which accurately describes the intrinsic dynamics of many task and motion planning problems including PTAMP.

Next, we state the central result of the paper: *jointly-controllably-open* (JC-open) domains are decidable. We then give a perturbation algorithm and show that any real-world PTAMP can be rewritten to be JC-open.

Finally, we give a constructive proof of the decidability of JC-open domains. Our algorithm is divided in four parts. First, we describe the decomposition algorithm which decomposes the configuration space into a finite number of manifolds with special properties. Next, we use techniques from differential geometry to calculate the internal controllability of each manifold. Afterwards, for every manifold, we calculate its stratified controllability, i.e. the controllability gained by leaving a manifold and utilizing the controllability of neighboring manifolds. To accomplish this step, we present the convergence condition, which we shows holds for JC-open domains. Finally, we execute a graph search to calculate the reachability set for our initial configuration and test for the existence of a solution.

2 Semi-holonomic task and motion planning framework

We consider an example PTAMP domain with one robot A and several movable objects B_1,...,B_k as shown in figure 1a. The configuration space of the problem is the Cartesian product of individual configuration spaces for the robot and each movable object, i.e. $\mathcal{C}_A \times \mathcal{C}_{B_1} \times ... \times \mathcal{C}_{B_k}$. There are two types of operators: MOVEROBOT, in which A transits around the space, and MANIPULATE-B$_i$, in which A manipulates object B_i while remaining in contact. We are given an initial configuration and set of goal configurations.

Semantically, for MOVEROBOT, we need to be able to modify the dimensions $\mathcal{C}_A$ without affecting any other dimensions. We enforce that each operator exhibits *semi-holonomic controllability* (SHC) in that a subset of dimensions, F, are marked as "free dimensions" and can be holonomically modified by the operator; non-free dimensions must be held constant. In our example, $F_{\text{MoveRobot}}$ and $F_{\text{Manipulate-B}_i}$ are $\{\mathcal{C}_A\}$ and $\{\mathcal{C}_A,\mathcal{C}_{B_i}\}$ respectively.

For each operator, we also set a predicate R which encodes the set of configurations that can be in any valid execution trajectory of the operator. $R_{\text{MoveRobot}}$ includes the set of all configurations in which there are no collisions between the robot, the objects, and the walls. $R_{\text{Manipulate-B}_i}$ is similar to $R_{\text{MoveRobot}}$, but it also eliminates configurations in which A and B_i are not in contact. Note that while R is a subset of the configuration space, any specific operator instance is confined to a cross-section of R corresponding to the free dimensions F. In addition, R can only be tested against the current configuration and does not have a memory; therefore, if the robot starts MANIPULATE-B$_i$ with one grasp, it may end the operation with another grasp as

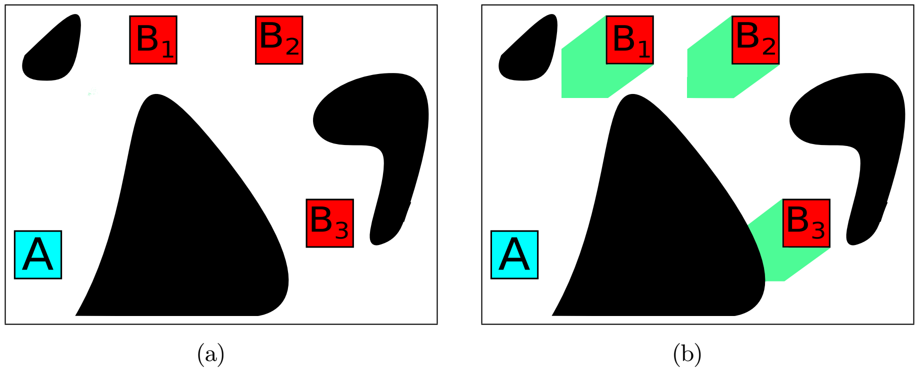

Fig. 1: (a) Basic PTAMP domain with one robot, three movable objects and obstacles. (b) Augmented PTAMP domain with shadows in green.

long as the object and robot remain in manipulation range. This "sliding grasp" phenomenon is corrected by replacing that trajectory with a sequence of transit and transfer operations, which is possible if the system allows the robot a small amount of leeway. So far, the predicates cover the typical geometric constraints of manipulation planning.

However, our framework affords us additional flexibility. For example, consider an extended scenario in which the robot is equipped with a solar panel and the movable objects cast shadows on the ground as shown in figure 1b. The robot can transit as before but cannot be in the shade when it manipulate objects as it requires additional power. In the updated scenario, $F_{\text{MoveRobot}}$, $F_{\text{Manipulate-B}_i}$ and $R_{\text{MoveRobot}}$ remain constant. However, $R_{\text{Manipulate-B}_i}$ must be updated to remove all areas that are in the shade. Therefore, unlike in typical manipulation planning, a shaded area can be traversable by one operator and impassable by another. This highlights the power of our representation as each operator can have its own unique dynamics.

A *domain* is a tuple $\mathcal{M}=(D,O)$ where:

- D is a set of n configuration dimensions of the entire domain, including robots and movable objects, each defined over $\mathbb{R}$. We assume D has been augmented with all the requisite dimensions utilized by dimension theory in embedding any elements of the configuration space that are typically expressed in alternate spaces to Euclidean space, e.g. angles in S^1 to $\mathbb{R}^2$ [2, 7].
- O is a set of operators. Each operator $o=(F,R)$, contains a set of free dimensions $F \subseteq D$ and a predicate R. Each operator is assumed to exhibit SHC (definition 2).
- R is a predicate such that every trajectory of its operator o must be contained within R. R is defined as a semi-algebraic subset of $\mathbb{R}^n$, i.e., there exists a finite set of finite-coefficient rational polynomials f of finite degree such that:

$$R=\left\{x\in\mathbb{R}^n \,|\, \bigcup_i \bigcap_j f_{ij}(x)\otimes 0\right\} \tag{1}$$

where each $\otimes$ is a binary relation in the set $\{>,<,\geq,\leq,=,\neq\}$. Let R^{all} be the set of R for all operators.

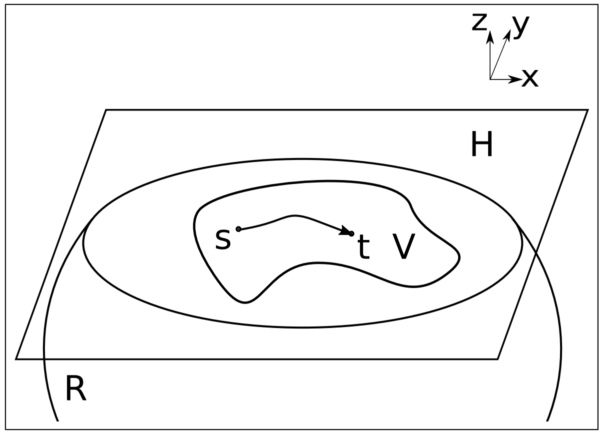

Fig. 2: Illustration of semi-holonomic controllability. In this example, the free dimensions of the operator are $\{x,y\}$, so H is a $\{x,y\}$-section. V shows a connected set on $H \bigcap R$. For any two points (s,t) inside V, there is a path between them without leaving V.

A *problem instance* is a tuple $I = (\mathcal{M}, x_0, G)$ where:

- $\mathcal{M}$ is a domain.
- $x_0 \in \mathbb{R}^n$ is the initial configuration.
- G is a semi-algebraic subset of $\mathbb{R}^n$ that represents the set of goal configurations.

An operator execution trajectory is a continuous trajectory in the configuration space. A sequence of operator execution trajectories is legal only if the terminus of each trajectory serves as the initial state of the subsequent trajectory. A sequence of operator execution trajectories is a solution to a problem instance, $\mathcal{I}$, if starting from x_0 and applying each execution trajectory sequentially, we end on a configuration contained within G. A decision procedure gives a proof that either such a solution does or does not exist in bounded time.

Before describing SHC, we first describe extrusion sets and sections, which extrude spatial subsets along some dimensions:

Definition 1. *Let $F \subseteq D$ be a subset of dimensions, $\overline{F} = (D \setminus F)$, and $V \subseteq \mathbb{R}^{|\overline{F}|}$ be a subset of the span of $\overline{F}$. For a configuration $x \in \mathbb{R}^n$, let $\text{proj}_{\overline{F}}(x) : \mathbb{R}^n \to \mathbb{R}^{|\overline{F}|}$ drop the dimensions F from x. $P(F,V)$ is an F-**extrusion set** with respect to V if:*

$$P(F,V) = \{x \in \mathbb{R}^n \,|\, \text{proj}_{\overline{F}}(x) \in V\}$$

*If V is a singleton set, then we describe the F-extrusion set $P(F,V)$ as an F-**section**.*

We can now define semi-holonomic controllability, which allows for holonomic behavior in a subset of the dimensions as shown in figure 2:

Definition 2. *Let $o = (F, R)$ be an operator, $s \in R$ be a configuration, $H = P(F, \text{proj}_{\overline{F}}(s))$, V be a connected subset of $H \bigcap R$, and $t \in V$ be another configuration. Then, o exhibits* ***semi-holonomic controllability (SHC)*** *if there exists a trajectory of o from s to t while staying inside V for all V, s, and t.*

A domain is SHC if all of its operators are SHC. Note that SHC is a form of factored holonomicity and cannot be used to model general non-holonomic problems.

3 Decidability result

In this section, we give our main decidability result: PTAMP is decidable. Unfortunately, not every problem in our framework is decidable, so we first define the concept of an *open* domain and state a primary result:

Definition 3. *An* ***open*** *domain is one in which for every operator $o = (F, R)$, R is an open set in $\mathbb{R}^n$.*

Theorem 1. *Every problem instance in an open, SHC domain is decidable.*

Our attempt to apply this theorem to PTAMP immediately fails as $R_{\text{MANIPULATEB}_i}$ may have codimension 1. Unfortunately, PTAMP is not open in general. However, we can perturb the domain in a manner similar to Canny to create an open domain with approximately the same dynamics [8]. This involves making all polynomial relations in equation (1) open.

First, we replace all polynomials of the form $f(x) = 0$ and $f(x) \neq 0$ with the logical statements $f(X) \geq 0 \wedge f(X) \leq 0$ and $f(X) < 0 \vee f(X) > 0$ respectively. We also replace $f(X) \geq 0$ and $f(X) > 0$ with $-f(X) \leq 0$ and $-f(X) < 0$ respectively. We then perturb the resulting inequalities by some small, positive number ϵ. Since manifolds with codimension 1 or higher are often semantically meaningful (such as $R_{\text{MANIPULATEB}_i}$), we expand non-strict inequalities by replacing $f(X) \leq 0$ with the extension $\oplus(f(X) < 0, B^n_\epsilon)$ where $\oplus$ is the Minkowski sum and B^n_ϵ is the n-dimensional ball of radius ϵ. Similarly, we contract strict inequalities by replacing $f(X) < 0$ with the contraction $\overline{\oplus(f(X) \geq 0, B^n_{2\epsilon})}$ where the line represents complementation. As semi-algebraic geometry is closed under Minkowski sums and complementation, the resulting system is also semi-algebraic. The perturbation may affect the solvability of a domain. However, for real-world systems, the effect is negligible and the difference between the original and perturbed system cannot be physically measured for small enough ϵ.

However, a potentially significant effect of the perturbation involves identities like trigonometric functions. Typically, an angle θ in S^1 is embedded in $\mathbb{R}^2$ on a pair of real dimensions θ_S and θ_C, corresponding to $\sin(\theta)$ and $\cos(\theta)$, and dynamics are generally polynomial in terms of θ_S and θ_C. The extra degree of freedom is removed by the identity $\theta_S^2 + \theta_C^2 = 1$. We call this trigonometric identity a *semantic-invariant*, since the domain semantics are fundamentally affected under perturbation no matter how small the value of ϵ. We address this issue by relaxing the requirement that domains be open to that domains be *jointly-controllably-open* (JC-open). We first define a set of dimensions to be *jointly-controllable* when they can always be manipulated together:

Definition 4. *Let* $B = \{B_1, B_2, ..., B_k\} \subseteq D$ *be a subset of dimensions.* B *is* ***jointly-controllable*** *if:*

$$\forall o_i \in O. \forall j,k. B_j \in F_i \leftrightarrow B_k \in F_i$$

In our trigonometric example, θ_S is only ever modified when θ_C is modified and vice-versa, so those dimensions are jointly-controllable.

Definition 5. *Let* v *be an* n*-vector. The* nonzero *function returns a subset of* D *such that:*

$$\text{nonzero}(v) = \{d_i \in D \mid v[i] \neq 0\}$$

Definition 6. *Let* $J = \{J_1, J_2, ...\}$ *be a maximal partitioning of dimensions into jointly-controllable sets for a domain. A manifold* M *is* ***JC-Open*** *if for every point* $x \in M$*, there exists a set of basis vectors* $V_x = \{v_x^1, v_x^2, ...\}$ *that span the tangent space of* M *at* x*,* $T_x M$*, such that*

$$\forall v_x^i \in V_x. \exists J_k \in J. \text{nonzero}(v_x) \subseteq J_k$$

Definition 7. *A domain is JC-Open if all predicates in* R^{all} *are JC-Open.*

We can now broaden theorem 1:

Theorem 2. *Every problem instance in a JC-open, SHC domain is decidable.*

Note that every open domain is JC-open with singleton jointly-controllable sets; therefore, theorem 1 follows from theorem 2, and we only provide a proof for theorem 2 in the next section. We can perturb a domain to be JC-open while preserving semantic-invariants by not perturbing any polynomial that is dependent on only the variables in a single jointly-controllable set of dimensions. Such perturbations are unneeded, as the system always exhibits holonomic controllability or no controllability in those dimensions.

In the general variant of PTAMP, we have may several robots and several movable objects. There are two types of operators, one to transit one or more robots and the other for one or more robots to manipulate one or more objects when the robots and objects are in proximity. We make the assumption that any time a group of robots move, all the dimensions in the configuration spaces of the robots can be modified. Similarly, in manipulation operations, all the dimensions in the joint configuration space of the involved robots and objects can be modified. The configuration dimensions for each robot and object therefore constitute a maximal jointly-controllable set. We assume that the predicates defining these operators are semi-algebraic. We also make the assumption that any semantic-invariants arising from dynamics for each robot or object are only dependent upon the dimensions in the configuration space for that robot or object. The only elements of the problem that combine jointly-controllable sets are those that model physical collisions and robot-object proximity. However, both constraints can be rewritten by perturbation, as they are not semantic-invariants as they model purely physical phenomenon. Finally, as PTAMP can be rewritten to be JC-open, we can use the Chow-Rashevsky theorem to approximate a sliding grasp trajectory to a sequence of transit and manipulation operations with standard fixed grasps. Within the scope of the stated conditions, PTAMP is decidable regardless of the dynamics, number of robots and objects, and complexity of the semi-algebraic geometries.

4 Decision procedure

In this section, we describe a decision procedure which determines the existence of a solution for any problem instance in a JC-open domain. The primary intuition behind the entire approach is that since all the operators are SHC, the dynamics of the operators are strongly tied to the Euclidean axes. Therefore, rather than taking the typical differential geometry view that any particular coordinate system is irrelevant, we decompose the configuration space relative to the Euclidean axes into a finite number of manifolds with properties that are subsequently defined. For each manifold, we calculate its controllabilities, the set of directions that can be traversed from a point within the manifold using sequences of operations. These controllabilities can be used to calculate orbits, a foliation of reachable sets within the manifold. We first calculate the internal and exterior controllability for each such manifold. Next, with the aid of a constraint, we calculate the stratified controllability for each manifold by examining which additional controllabilities can be achieved by leveraging the controllabilities of adjacent manifolds. Once the controllabilities of each manifold are computed, the initial configuration is used to calculate the reachability set, which is then intersected against the goal condition to test for satisfiability.

The following sections describe, give pseudocode, and prove key properties of the four phases: decomposition, internal/exterior controllability, stratified controllability, and reachability. The following pseudocode is a roadmap for decision procedure DP.

Algorithm DP — Input: $\mathcal{I}=(\mathcal{M},x_0,G)=((D,O),x_0,G)$ — Output: Solvability

$\mathfrak{A}\leftarrow$ Decomposition(R^{all})
$\mathcal{E},\mathcal{D}\leftarrow$ Internal/Exterior Controllability($\mathfrak{A}$,O)
$\mathcal{S}\leftarrow$ Stratified Controllability($\mathfrak{A}$,$\mathcal{D}$,$\mathcal{E}$)
return Reachability($\mathfrak{A}$,x_0,$\mathcal{D}$,$\mathcal{S}$,G)

4.1 Decomposition phase

The decomposition phase decomposes the configuration space along extrusion sets using algebraic geometry techniques. First, we give an overview of cylindrical algebraic decomposition (CAD), full cylindrical algebraic decomposition (FCAD), and the decomposition phase algorithm. Then, we define geometric correctness conditions for this phase and show that the algorithm satisfies them.

Cylindrical algebraic decomposition (CAD) is an algorithm to decompose semi-algebraic sets into a stratification and was used to prove the decidability of motion planning [1]. We use the notation E_i^j to represent an ordering of a subset of dimensions, $[e_i,e_{i+1},...,e_{j-1},e_j]$. CAD requires two inputs: a decomposition ordering E_1^n (a permutation of D) and the set of predicates Q, which are collectively defined by the polynomial set $Y(E_1^n)$. CAD iteratively processes dimensions backwards over E_1^n. On each projection iteration, CAD identifies two types of events: the intersection set of two polynomials (or a self-intersection) or the set on a polynomial in which the normal is orthogonal to the dimension being processed. On the i-th iteration, CAD extrudes these events along the processed dimension as the set of polynomials $Y(E_1^{n-i})$. This process is continued to create polynomials of the form $Y(E_1^i)$ for $1 \leq i \leq n$. We omit a detailed description of the projection operators as CAD is covered in depth in various other publications [9, 10].

While CAD has traditionally been used to decompose semi-algebraic sets, we are especially interested in CAD for some of the side effects it produces based on the decomposition ordering. The full cylindrical algebraic decomposition (FCAD) algorithm takes a set of predicates as its input, runs CAD over every decomposition ordering, and intersects the results as shown in figure 3. Although there are $O(n!)$ unique decomposition orderings, only the set of extruded dimensions is relevant (not their specific ordering), necessitating only 2^n total projection iterations. We run FCAD once with the predicate set R^{all}. We then run FCAD a second time with the predicate set FCAD(R^{all}) and let $\mathfrak{A}$ be the resulting set of manifolds.

Algorithm Decomposition — Input: R^{all} — Output: $\mathfrak{A}$

return $\mathfrak{A} \leftarrow$ FCAD(FCAD(R^{all}))

Since the geometry of manifolds and the dynamics of operators are linked by their relation to the Euclidean axes, we define characteristic sets and dimensional sets to mathematically express these properties.

Definition 8. *A* ***characteristic set*** *is a set of subsets of D. A* ***dimensional set*** *is a subset of D. Let $\mathcal{C}$ be a characteristic set. As a shorthand, we define $\mathcal{C}^* = \bigcup \mathcal{C}$ to create a dimensional set from $\mathcal{C}$.*

Characteristic sets and dimensional sets can be used to express both the controllability of a manifold and its geometry. In the decomposition phase, we concern ourselves only with the latter. In particular, we are interested in when vectors are orthogonal to the Euclidean axes. The tangent characteristic set of a manifold is a boolean representation of the geometry of the tangent space at each point.

Definition 9. *The* ***tangent characteristic set****, $\mathcal{T}$, of a manifold M at a point x is defined as:*

$$\mathcal{T}_x M = \{B \subseteq D \mid \exists v_x \in T_x M.\text{nonzero}(v_x) = B\}$$

As shown in figures 5a and 5b, in general, the tangent characteristic set of a manifold can vary from point to point. However, for *aligned* manifolds as shown in figures 5c and 5d, the tangent characteristic set is constant. Aligned manifolds are useful as SHC operators have uniform controllability throughout the manifold.

Definition 10. *If $\mathcal{T}_x M = \mathcal{T}_y M$ for all $x, y \in M$, then M is* ***aligned****. We denote the tangent characteristic set of the entire aligned manifold M as $\mathcal{T}M$.*

Proposition 1. *For any manifold M, if $B_1 \in \mathcal{T}M$ and $B_2 \in \mathcal{T}M$, then $B_1 \bigcup B_2 \in \mathcal{T}M$.*

Proof Sketch. The tangent space of a manifold is a locally, finitely generated distribution, so the span of two components of the tangent space must also be in the tangent space. □

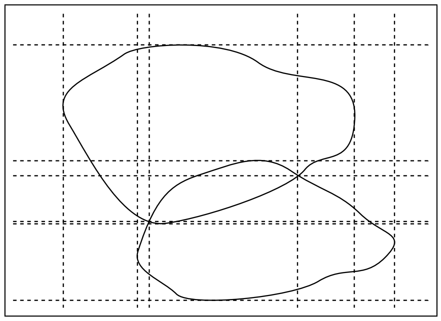

Fig. 3: Example FCAD decomposition. The resulting decomposition is the union of performing CAD with the decomposition orderings $\{x, y\}$ and $\{y,x\}$.

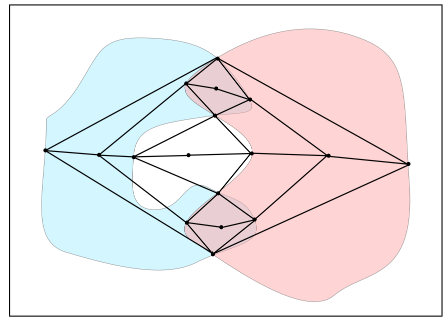

Fig. 4: Example adjacency graph of two predicates. Vertices are maximal, connected regions that are invariant to the polynomials constituting the predicates and edges indicate adjacent regions.

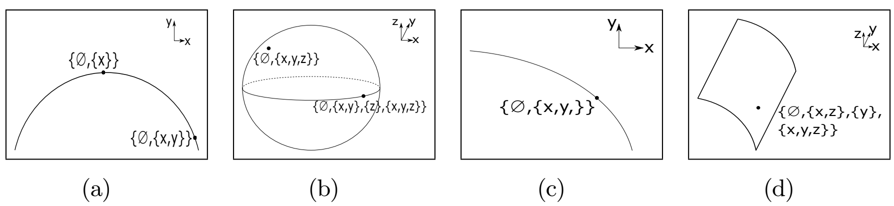

Fig. 5: Tangent characteristic set is marked at given points. (a) and (b) show non-aligned manifolds. (c) and (d) show aligned manifolds.

Next, we define *adjacency graphs* to characterize the adjacency of regions of the configuration space as shown in figure 4.

Definition 11. *Let* $Q = \{Q_1,Q_2,...,Q_k\}$ *be a set of subsets of* $\mathbb{R}^n$. *Let* T *be the set of maximal, connected regions that are sign-invariant to the polynomials constituting* Q. *We define the* ***adjacency graph*** *of* Q *as the graph having nodes* T *and edges* $(t_1,t_2) \in T^2$ *whenever* $t_1 \bigcup t_2$ *is contiguous.*

Definition 12. *Let* $B \subseteq D$ *be a subset of dimensions. A manifold* M *is* B***-uniform*** *to* Q *if* M *is a* B*-extrusion set and every* B*-section through* M *has the same adjacency graph of* Q.

We start our analysis by noting that CAD produces uniform extrusion sets.

Proposition 2. *Assume CAD is run on the set of predicates* Q *with decomposition ordering* E_1^n. *Let* $W_i = \bigcup_{j=1}^{n-i-1} Y(E_1^j)$ *be the set of polynomials created by CAD that are independent of* E_{n-i}^n. W_i *divides the space into a finite number of* E_{n-i}^n*-extrusion sets, each sign-invariant with respect to* W_i. *Each such* E_{n-i}^n*-extrusion set,* M, *is* E_{n-i}^n*-uniform to* Q.

Proof Sketch. Let J_1 and J_2 be two E^n_{n-i}-sections that intersect M. Let S_M, S_1 and S_2 be the projection of M, J_1 and J_2 respectively onto the dimensions E^{n-i-1}_1. Let T be a path from S_1 to S_2 such that $T \subseteq S_M$. For contradiction, we examine the two cases when the adjacency graphs of J_1 and J_2 might differ:

- **A vertex disappears** By the mean-value theorem, a sign-invariant manifold can only appear or disappear between J_1 and J_2 if for all T, there exists some value $S_3 \in T$ such that within the E^n_{n-i} extrusion of S_3, either two polynomials in W_i intersect or a polynomial in W_i is orthogonal to some dimension in E^{n-i-1}_1. The set of S_3 for all T divides S_1 and S_2 and its extrusion must be a member of W_i. Therefore, M is not sign-invariant to W_i, which is a contradiction.
- **An edge disappears** Similarly, by the mean-value theorem, for all T, there exists a set of E^n_{n-i}-sections at which the adjacency between the two predicate-invariant manifolds disappears. At that point, two polynomials forming Q must intersect. Following the same reasoning as before, M is not sign-invariant to W_i, which is a contradiction.

□

The decomposition computed by this procedure has the following properties for each manifold $M \in \mathfrak{A}$:

Proposition 3. *M is sign-invariant to R^{all}.*

Proof Sketch. CAD decomposes the configuration space into cells that are sign-invariant to the input predicate set regardless of decomposition ordering, and the intersection operator preserves sign-invariance. □

Proposition 4. *M is aligned.*

Proof Sketch. If M is n-dimensional, then trivially it is aligned. Otherwise, assume for contradiction that there exists two configurations $x \in M$ and $y \in M$ and a vector $v_x \in T_xM$, such that there does not exist a vector $v_y \in T_yM$ where $\text{nonzero}(v_x) = \text{nonzero}(v_y)$. Since CAD produces regular manifolds, M must be regular and have a constant dimensionality. Therefore, at y there must exist some vector $z_y \in T_yM$ such that $\text{nonzero}(z_y) \notin \mathcal{T}_xM$ and either $\text{nonzero}(v_x) \subset \text{nonzero}(z_y)$ or $\text{nonzero}(v_x) \supset \text{nonzero}(z_y)$. Therefore, either x or y must lie on a polynomial created by CAD, so x and y cannot be in the same cell that is sign-invariant to R^{all}, which is a contradiction. □

Proposition 5. *For every subset of dimensions $B \subseteq D$, M is B-uniform to Q. Furthermore, any pair of B-sections intersecting M, isomorphic manifolds have the same alignment.*

Proof Sketch. As FCAD runs CAD in every decomposition ordering, by proposition 2, after one iteration of FCAD, every manifold in the resulting decomposition, L, is H-uniform for every $H \subseteq D$. Since propositions 3 and 4 apply after even a single application of FCAD, every manifold in L is aligned and sign-invariant to R^{all}. As the subsequent FCAD decomposition is run on the predicates L, by proposition 2, each resulting manifold in the second decomposition is B-uniform with respect to L. Therefore, isomorphic manifolds in adjacency graphs of $\mathfrak{A}$ must have the same alignment. □

4.2 Internal/exterior controllability phase

For the next phase, we calculate both the internal and exterior controllability of every manifold $M \in \mathfrak{A}$. The exterior characteristic set is the set of directions that can be utilized to exit M. The internal characteristic set is the set of directions that can be traversed within M without leaving the manifold and define internal orbits of M.

Definition 13. *Let ϕ_M be the set of operators such that $\forall o_i \in \phi_M . M \subseteq R_i$.*

The exterior characteristic set, $\mathcal{E}$, is easily computed, because we can ignore interaction between operators. We take the union of all the directions in which each operator can individually utilize to leave the manifold.

Definition 14. *Let* Pow *be the powerset function. The* ***exterior characteristic set*** *of M, $\mathcal{E}M$ is:*

$$\mathcal{E}M = \bigcup_{o_i \in \phi_M} \mathrm{Pow}(F_i)$$

The internal characteristic set, $\mathcal{D}$, describes the set of directions that can be traversed within M without leaving the manifold. The calculation is slightly more complex since both the interactions of operators as well as the geometry of M need to be taken into account. First, we define the internal characteristic set of a single operator.

Definition 15. *Let M be an aligned manifold and $o_i \in \phi_M$ be an operator. The* ***internal characteristic set*** *of M with respect to o_i, $\mathcal{D}_{o_i}M$ is:*

$$\mathcal{D}_{o_i}M = \mathrm{Pow}(F_i) \bigcap \mathcal{T}M \tag{2}$$

The internal dimensional set of M, $\mathcal{D}^*M$, with respect to all operators is:

$$\mathcal{D}^*M = \bigcup_{o_i \in \phi_M} \mathcal{D}^*_{o_i}M \tag{3}$$

Therefore, every manifold M has internal orbits that are $\mathcal{D}^*M$-sections.

Algorithm Internal/Exterior Controllability — Input: $\mathfrak{A}$,O — Output: $\mathcal{D}$,$\mathcal{E}$

```
for each M ∈ 𝔄 do
    ℰM = ⋃_{o_i ∈ φ_M} Pow(F_i)
    for each o_i ∈ φ_M do
        𝒟_{o_i} M = Pow(F_i) ⋂ 𝒯M
    𝒟*M = ⋃_{o_i ∈ φ_M} 𝒟*_{o_i} M
return ℰ,𝒟
```

We prove that equation 3 follows from equation 2:

Proposition 6. *For any aligned manifold M, $\mathcal{D}^*M = \bigcup_{o_i \in \phi_M} \bigcup \mathcal{D}_{o_i}M$.*

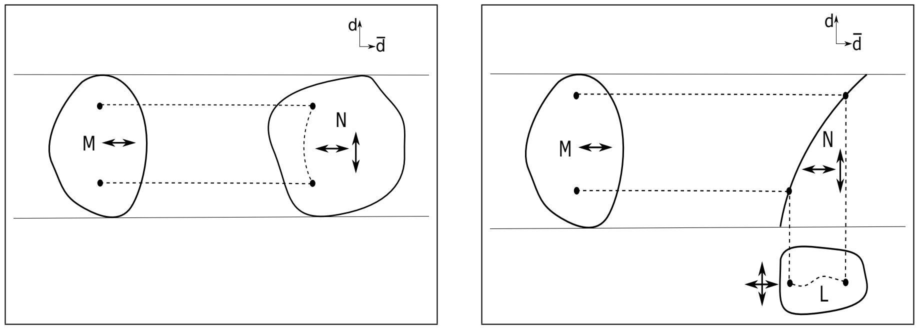

(a)

(b)

Fig. 6: (a) M utilizes the controllability of N to gain controllability in the d dimension. (b) Calculating the stratified controllability of M requires calculating the stratified controllability of N.

Proof Sketch. First, we show that $\bigcup \mathcal{D}_{o_i} M = \mathcal{D}^*_{o_i} M \in \mathcal{D}_{o_i} M$ and therefore orbits within M are $\mathcal{D}^*_{o_i} M$-sections. Let $B_1 \in \mathcal{D}_{o_i} M$ and $B_2 \in \mathcal{D}_{o_i} M$. By construction, $B_1 \bigcup B_2 \subseteq F_i$. By proposition 1, $B_1 \bigcup B_2 \subseteq \mathcal{T} M$. Therefore, $B_1 \bigcup B_2 \subseteq \mathcal{D}_{o_i} M$. $\mathcal{D}^*_{o_i} M$ is just the union of all $B \in \mathcal{D}_{o_i} M$ and so must also be a member of $\mathcal{D}_{o_i} M$.

Let $H = \bigcup_{o_i \in \phi_M} \mathcal{D}^*_{o_i} M$. $H \in \mathcal{T} M$ by proposition 1. The Lie closure of ϕ_M spans every non-empty H-section of M. Therefore, by the Chow-Rashevsky theorem [11], the intersection of those H-sections and M constitute the orbits of M. □

4.3 Stratified controllability phase

In the stratified controllability phase, for each manifold M, we calculate the stratified controllability as the dimensional set $\mathcal{S}^* M$ that can be achieved by leaving M and utilizing the controllability of adjacent manifolds [12] as shown in figure 6a. For every dimension d, we attempt to gain stratified controllability d for M. We travel in a $\overline{d}$-section from M and look for a reachable manifold N that has the ability to travel in the d direction. We then travel in the d direction on N before retracing our steps back to M.

Determining when manifold N allows for movement in the d direction based on its internal controllability is straightforward, but is tricky when the stratified controllability of N itself needs to be calculated to compute the stratified controllability of M as shown in figure 6b. In order to prevent cyclic computations, we state a condition, which holds for JC-Open domains, that allows us to immediately ascertain the controllability of N in the d dimension.

Condition 3 (Convergence condition)**.** *For every manifold $N \in \mathfrak{A}$:*

$$|\mathcal{E}^* N| \leq 1 \vee (\forall d \in D. d \in \mathcal{T}^* N \wedge d \in \mathcal{E}^* N \rightarrow d \in \mathcal{D}^* N \vee d \in \mathcal{S}^* N)$$

Roughly stated, the convergence condition ensures that for every manifold N whose operators allow for travel in more than one direction, for every dimension d, if some tangent vector of N contains a nonzero d component and an operator allows for movement in some direction with a nonzero d component, then N must be controllable in the d direction via either internal controllability or stratified controllability.

Algorithm Stratified Controllability — Input: $\mathfrak{A}$,$\mathcal{D}$,$\mathcal{E}$ — Output: $\mathcal{S}$

```
for each manifold M ∈ 𝔄 do
    for each dimension d ∈ D do
        for each manifold N ∈ 𝔄 that is reachable in some d̄-section do
            if d ∈ ℰ*N then
                Add d to S*M
return 𝒮
```

First, we address when it is possible to traverse between a pair of adjacent manifolds. As CAD produces a stratification, two resulting manifolds can only be adjacent if one is in the stratum of the other.

Definition 16. *Let M_1 and M_2 be adjacent manifolds. Assume M_2 is in the stratum of M_1. Let $B \subseteq D$ be a subset of dimensions. M_1 and M_2 are B-**traversable** if for every point $x \in M_2$, there exists some vector v such that* $\text{nonzero}(v) = B$, $\text{nonzero}(v) \in \mathcal{E}M_1$, $\text{nonzero}(v) \in \mathcal{E}M_2$, *and $x + \epsilon v \in M_1$ for small enough ϵ.*

Proposition 7. *If M_1 and M_2 are B-traversable, then for each configuration $x_1 \in M_1$, there exists a reachable configuration $x_2 \in M_2$ and vice-versa.*

Proof Sketch. Again, assume M_2 is contained in the closure of M_1. Since M_1 and M_2 are B-traversable, by definition 16, every point in M_2 can reach a point in M_1. Let I be the intersection of the closures of M_1 and M_2. The decomposition phase extrudes I in every subset of dimensions. Assume for contradiction that $\hat{x}_1 \in M_1$ cannot reach M_2. Then $\hat{x}_1 \notin P(B, \text{proj}_{\overline{B}}(I))$, so M_1 is not invariant with respect to the polynomials created by FCAD, which is a contradiction. □

We use proposition 7 to construct paths from chains of traversable manifolds.

Proposition 8. *If M can reach manifold N on some d-section, then M can reach N on any d-section that intersects M.*

Proof. By proposition 5, the adjacency graphs and alignment of manifolds of every d-section that intersects M is the same. Therefore, proposition 7 must hold for every manifold-manifold transition in any d-section. □

Proposition 9. *If the convergence condition holds, then we compute the closure of stratified controllability.*

Proof Sketch. Let N be a manifold that might grant M the ability to move in the d dimension. Therefore, $d \in \mathcal{E}^*N$. If $|\mathcal{E}^*N| \leq 1$, then N is unreachable from any $\overline{d}$-section. If $d \notin \mathcal{T}^*N$, then by proposition 5, $d \notin \mathcal{T}^*M$, so M cannot travel in the d dimension with stratified controllability regardless of the controllability of N. Therefore, N must be able to travel in the d direction either internally or via stratified controllability via the convergence condition, which exhausts all cases. □

Proposition 10. *JC-open domains satisfy the convergence condition.*

Proof Sketch. Consider the case when $|\mathcal{E}^*N|>1$, $d\in\mathcal{T}^*N$, and $d\in\mathcal{E}^*N$ as it is the only scenario in which the convergence condition may be violated. Let d be a member of J, a maximal jointly-controllable set of dimensions. Since $d\in\mathcal{E}^*N$, $\forall j\in J.j\in\mathcal{E}^*N$. Since $d\in\mathcal{T}^*N$, by definition 7, we know there must exist a JC-open neighborhood around every point in N. Since at least one vector at every point has a d component and is in the tangent space of the open neighborhood, we can travel in the direction of that vector to effect movement in the d dimension before returning to N. Therefore, $d\in\mathcal{S}^*N$, which satisfies the convergence condition. □

4.4 Reachability phase

In the first three phases, only the domain information is needed; for the final phase, the initial configuration and goal are also taken into account. From the initial configuration x_0, we calculate the reachability set $\mathfrak{R}$. By using a graph search across manifolds, we add the reachabilities of adjacent manifolds until all reachable manifolds are visited. Since we have already computed the internal and stratified controllability of each manifold, there is no need to ever revisit a manifold in order to gain additional reachability. Finally, after $\mathfrak{R}$ has been fully computed, we intersect $\mathfrak{R}$ with the goal condition and test for emptiness to determine the existence of a solution.

We calculate $\mathfrak{R}$ by a graph search on the graph $\mathfrak{G}$. Initially, the vertices in $\mathfrak{G}$ are $\mathfrak{A}$, and there are no edges in $\mathfrak{G}$. However, over the course of the algorithm, we add edges from visited manifolds to reachable manifolds. Through the computation, we let I_M be the initial reachability set for manifold M. Let M_0 be the manifold that contains the initial configuration x_0. We set I_{M_0} to x_0.

When visiting a manifold M with initial reachable set I_M, we construct H_M, the reachability set of M, by take the following steps:

- Initialize H_M to I_M.
- Extrude H_M according to the controllability of M. For any dimension d such that $d\in\mathcal{D}^*M$ or $d\in\mathcal{S}^*M$, we extrude H_M in the dimension d within the confines of M.
- Add edges from M to all reachable neighboring manifolds. For each adjacent manifold N, we test if N can be reached from the reachable region of M by testing if (M,N) is B-traversable for some B and $\mathrm{Cl}(H_M)\bigcap\mathrm{Cl}(N)\neq\emptyset$ where Cl is the closure operator. For each such manifold, we add the edge (M,N) and set $I_N=\mathrm{Cl}(H_M)\bigcap\mathrm{Cl}(N)$.
- Add H_M to $\mathfrak{R}$.

After the graph search terminates, we check for the existence of a solution by testing the emptiness of $\mathfrak{R}\bigcap G$.

Algorithm Reachability — Input: $\mathfrak{A}$,x_0,$\mathcal{D}$,$\mathcal{S}$,G — Output: Solvablility

```
𝔊 ← (𝔄,∅)
I_{M_0} ← x_0
while M ← next visited node in graph search of 𝔊 do
    H_M ← I_M
    for each d ∈ D do
        if d ∈ 𝒟*M or d ∈ 𝒮*M then
            Extrude H_M in dimension d
    for each manifold N where (M,N) is traversable and Cl(H_M) ⋂ Cl(N) ≠ ∅ do
        I_N ← Cl(H_M) ⋂ Cl(N)
        Add edge (M,N)
    Add H_M to ℜ
return ℜ ⋂ G ≠ ∅
```

Proposition 11. *Let H_M be the maximal reachable set within M. For an adjacent and reachable manifold N, if $I_N = \mathrm{Cl}(H_M) \bigcap \mathrm{Cl}(N)$, then H_N constitutes the maximal reachable set within N.*

Proof Sketch. Since H_M is the maximal reachable set within M, I_N constitutes the maximal region that can be reached on the border of M and N. The extrusion of H_N in all dimensions of $\mathcal{D}^*N$ and $\mathcal{S}^*N$ must be reachable and constitute the maximal reachability of of N since we have computed the maximal internal and stratified controllabilities. □

Proposition 12. *The reachability set within any manifold N is independent of the order of the graph search algorithm.*

Proof Sketch. Let N be reached via two different paths from adjacent manifolds M_1 and M_2. As x_0 can be reached from the reachable set of M_1, the reachability set within N when entering from M_1 is a superset of the reachability set within N when entering from M_2. However, the argument can also be reversed. Therefore, the reachability set within N is independent of its predecessor. □

5 Conclusion

In this paper, we defined SHC and JC-open domains. We described how the general PTAMP problem can be perturbed to be JC-open without significant semantic changes. We then gave a decision procedure, DP, for domains that satisfy the convergence condition. Last, we showed that JC-open domains satisfy the convergence condition. Therefore, the general PTAMP problem is decidable.

As the proof of decidability is constructive, through careful bookkeeping, one can recover the exact sequence of actions required to reach the goal if the problem is feasible. While the recovered solution is almost certainly not optimal, it places an upper bound on the cost of some feasible plan. Therefore, for any cost function in which the cost increases strictly-monotonically with respect to the number of actions taken, one can use the method of Cheng et al. [13] to exhaustively search for an optimal solution for PTAMP.

Acknowledgments We gratefully acknowledge support from NSF grants 1420927 and 1523767, from ONR grant N00014-14-1- 0486, and from ARO grant W911NF1410433. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of our sponsors.

References

1. Schwartz, J.T., Sharir, M.: On the piano movers problem. ii. general techniques for computing topological properties of real algebraic manifolds. Advances in applied Mathematics **4**(3) (1983) 298–351
2. Canny, J.: The complexity of robot motion planning. MIT press (1988)
3. Wilfong, G.: Motion planning in the presence of movable obstacles. In: Proceedings of the fourth annual symposium on Computational geometry, ACM (1988) 279–288
4. Alami, R., Laumond, J.P., Siméon, T.: Two manipulation planning algorithms. In: Proceedings of the Workshop on Algorithmic Foundations of Robotics. WAFR, Natick, MA, USA, A. K. Peters, Ltd. (1995) 109–125
5. Dacre-Wright, B., Laumond, J.P., Alami, R.: Motion planning for a robot and a movable object amidst polygonal obstacles. In: Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on, IEEE (1992) 2474–2480
6. Vendittelli, M., Laumond, J.P., Mishra, B.: Decidability of robot manipulation planning: Three disks in the plane. In: WAFR. (2014)
7. Munkres, J.: Topology. Featured Titles for Topology Series. Prentice Hall, Incorporated (2000)
8. Canny, J.: Computing roadmaps of general semi-algebraic sets. The Computer Journal **36**(5) (1993) 504–514
9. Collins, G.E.: Quantifier elimination for real closed fields by cylindrical algebraic decompostion. In: Automata Theory and Formal Languages 2nd GI Conference Kaiserslautern, May 20–23, 1975, Springer (1975) 134–183
10. Hong, H.: An improvement of the projection operator in cylindrical algebraic decomposition. In: Proceedings of the international symposium on Symbolic and algebraic computation, ACM (1990) 261–264
11. Jurdjevic, V.: Geometric Control Theory. Cambridge Studies in Advanced Mathematics. Cambridge University Press (1997)
12. Sussman, H.: Lie brackets, real analyticity and geometric control. Differential Geometric Control Theory **27** (1982) 1–116
13. Cheng, P., Pappas, G., Kumar, V.: Decidability of motion planning with differential constraints. In: Robotics and Automation, 2007 IEEE International Conference on, IEEE (2007) 1826–1831