

# Pushing the Boundaries of Asymptotic Optimality in Integrated Task and Motion Planning

Rahul Shome, Daniel Nakhimovich, and Kostas E. Bekris


Rutgers University, New Brunswick, NJ, USA.

**Abstract.** Integrated task and motion planning problems describe a multi-modal state space, which is often abstracted as a set of smooth manifolds that are connected via sets of transition states. One approach to solving such problems is to sample reachable states in each of the manifolds, while simultaneously sampling transition states. Prior work has shown that in order to achieve asymptotically optimal (AO) solutions for such piecewise-smooth task planning problems, it is sufficient to double the connection radius required for AO sampling-based motion planning. This was shown under the assumption that the transition sets themselves are smooth. The current work builds upon this result and demonstrates that it is sufficient to use the same connection radius as for standard AO motion planning. Furthermore, the current work studies the case that the transition sets are non-smooth boundary points of the valid state space, which is frequently the case in practice, such as when a gripper grasps an object. This paper generalizes the notion of clearance that is typically assumed in motion and task planning to include such individual, potentially non-smooth transition states. It is shown that asymptotic optimality is retained under this generalized regime.


## 1 Motivation

Integrated task and motion planning (TAMP) corresponds to simultaneously searching for continuous motions and discrete sequences of actions which resolve the target task when combined. Consider the motivational example of prehensile manipulation planning as in Fig. 1, where a robotic arm must plan both how to move its joints and compute a sequence of discrete grasps and placements. Efficiently solving TAMP problems empowers robots to manipulate objects in the real-world once perception pipelines have provided the location of the detected objects.

Many asymptotically optimal (AO) algorithms for motion planning such as PRM\*, RRT\*, FMT\* [13, 11], have existed for some time now and high level task planning is typically done with an informed tree search. Integrated TAMP, however, is more challenging than its two constituent problems as different task primitives, referred to as *task modes*, often correspond to different configuration spaces of possibly varying dimensionality. Transitioning between different task modes requires sampling the boundaries



**Fig. 1.** A multi-modal planning problem where the robot can perform sequences of continuous motions and discrete actions to achieve a target arrangement.



**Fig. 2.** (Left:) A toy problem with a gripper free to move vertically (along  $x$ ), and close its fingers (expressed by the  $g$  axis). A configuration that touches an object between its fingers lies at the non-differentiable boundary, which is the apex of the triangle in the  $g$ - $x$  graph. (Right:) A task planning problem involving picking the green object, and placing it in the blue region.

between them with a specialized subroutine that is task dependent. For instance, in the case of manipulation planning, sampling the boundary of modes corresponds to identifying different grasps or placements of an object, which, even for simple cases, is not necessarily smooth (see Fig. 2 left). This violates assumptions typically made in sampling-based planning literature [13, 11, 25].

Early work in manipulation TAMP focused on modeling the search space [1, 15]. Recent notable efforts [8] proposed integrated TAMP solutions via multi-modal search strategies, which construct sampling-based roadmaps over individual modes. The key idea constructs an incremental multi-modal PRM [9] over a search tree in the space of task modes. Symbolic planning has also been incorporated into TAMP using action predicates [7], incremental constraints [3], and more recently in an *inside-out* approach by devising a specialized hybrid sampler [23]. The mentioned approaches [8, 9, 7, 23, 3] have been shown to be probabilistically complete, but not necessarily AO.

Other problem instances, such as non-prehensile manipulation [5], or alternative solution frameworks, such as constrained optimization formulations [24], and hybrid approaches [10] using answer set programming have also been studied. There are also hierarchical search strategies, which at a low-level call time-budgeted motion planning subroutines that guide the search over viable actions [12]. Subsequently, application of such compositional techniques [18, 4, 20, 19], that perform task planning using sequences of motion plans, were used in domains of manipulation and rearrangement.

Of specific interest to the current work is the FOBT algorithm [25], which was designed for piecewise-analytic task planning domains, using a PRM\*-like motion planner. Notably, it argued an AO TAMP solution for a roadmap connection radius *twice* of what was argued for motion planning [13]. A contribution of this work is to demonstrate that previous algorithms [25, 9, 18] can be argued to be AO *without inflating the connection radius*. The key insight is that the AO properties of the motion planning roadmaps can be extended to the TAMP domain under a specific set of identified conditions.

Though not a focus in the current work, the arguments presented here should hold for near-optimal roadmaps [22, 21] in the interior of *orbits*, assuming appropriate radius [21] is chosen for transition connections. .

A standard assumption in the sampling-based planning is the existence of clearance, i.e., minimum distance from obstacles. This can be violated in task planning, as shown in Fig. 2 (left), where a target grasp lies on a non-differentiable boundary point with zero-clearance. Another contribution of the current work is to model what happens in such non-smooth boundaries by extending the definition of robust convergence [11, 13]. It is also shown that such countable singular points do not affect the AO properties of integrated task and motion planning.

## 2 Integrated Task and Motion Planning

This section first defines useful terms and notations, then outlines a general TAMP algorithm ALGO that will be analyzed in later sections.

**C-space abstraction:** A robotic system is describable by a configuration  $q$  in a  $d$ -dimensional configuration space  $\mathbb{C} \subset \mathbb{R}^d$ . The robot geometries exist in a workspace  $\mathcal{W} \subset \mathbb{R}^3$ , part of which is occupied by obstacles. This gives rise to an open subset  $\mathbb{C}_{\text{free}} \subset \mathbb{C}$  of the configuration space, which does not result in collisions with obstacles, and the complement obstacle subset  $\mathbb{C}_{\text{obs}} = \mathbb{C} / \mathbb{C}_{\text{free}}$ . The boundary of  $\mathbb{C}_{\text{free}}$  is denoted by  $\partial\mathbb{C}_{\text{free}}$ , and  $\overline{\mathbb{C}_{\text{free}}}$  denotes its closure, i.e.,  $\mathbb{C}_{\text{free}}$  and its boundary.

**Paths:** A parameterized continuous curve  $\pi$  in  $\overline{\mathbb{C}_{\text{free}}}$  is used to denote valid paths of the robot as  $\pi : [0, 1] \rightarrow \overline{\mathbb{C}_{\text{free}}}$ . A path from  $q_0$  to  $q_1$  in  $\overline{\mathbb{C}_{\text{free}}}$  is denoted as  $\pi_{q_0 \rightarrow q_1}$  such that  $\pi(0) = q_0$  and  $\pi(1) = q_1$ . Let  $\mathbb{P}$  be the set of all valid paths. Then, the path cost is defined as a mapping  $\mathbf{C} : \mathbb{P} \rightarrow \mathbb{R}^+$ , which returns a positive measure of a path. The current work considers Euclidean arc-length cost, which is Lipschitz continuous.

**Modes:** This work adopts the language of prior work for task planning [25, 9] and define a finite set of modes  $\mathfrak{M} = \{\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_k\}$ , which correspond to different operational constraints of the robot for different task components. The set of all possible configurations within a given mode is denoted as  $\mathbb{C}^{\mathcal{M}_i}$ ; thus,  $\mathbb{C} = \cup_{\mathcal{M}_i \in \mathfrak{M}} \mathbb{C}^{\mathcal{M}_i}$ . Initially, the discussion will restrict each  $\mathbb{C}^{\mathcal{M}_i}$  to be analytic with smooth boundaries, but this assumption will be later waived in Section 5.

**Orbits:** Dimensionality reducing constraints force  $\mathbb{C}^{\mathcal{M}_i} \subset \mathbb{R}^{d_{\mathcal{M}_i}}$  to be a lower dimensional manifold compared to  $\mathbb{C}$ , i.e.,  $d_{\mathcal{M}_i} < d$ . This arises when the task requires the robot to be constrained for some of its degrees of freedom. Within these modes, define an orbit  $\mathcal{O}_{\mathcal{M}_i}(x)$  as a maximal, path-connected, subset of  $\mathbb{C}^{\mathcal{M}_i}$ , which contains configurations  $x \in \mathbb{C}^{\mathcal{M}_i}$ . Often in manipulation planning, orbits of a mode  $\mathcal{M}_i$  are non-overlapping subsets of a robot's  $\mathbb{C}^{\mathcal{M}_i}$  corresponding to a specific grasp for the grasped object and specific placement of the non-grasped objects.

**Transitions:** In manipulation planning, the configurations where the robot just grasped or placed an object lie on the border of two modes (two specific orbits of those modes). These configurations at the intersection of two orbits, are called transition states. Formally, a configuration  $t \in \mathbb{C}_{\text{free}}$  is a transition state if  $t \in \partial\mathcal{O}_{\mathcal{M}_i} \cap \partial\mathcal{O}_{\mathcal{M}_j}$  for  $i \neq j$ .

**TAMP paths:** For a (non-trivial) task planning problem with a starting configuration of  $q_s$  and goal configuration of  $q_g$  a feasible solution path will traverse multiple orbits over a transition sequence  $T = (t_i)_{i=1}^M$  of length  $M$ . A feasible TAMP path to a task planning query is denoted as  $\Pi = \bigoplus_{i=0}^M \pi_{q_i \rightarrow q_{i+1}}$ ,  $q_i \in (q_s, T, q_g)$ , where  $\bigoplus$  denotes path concatenation. The cost  $\mathbf{C}(\Pi)$  of such a path is taken to be the sum of the costs of paths being concatenated. This formulation implies that traversing a transition has zero (or constant) cost. Such costs are ultimately domain dependent, and cases where the cost function is not locally smooth over transition manifolds are not considered in the current work. The optimal path is denoted as  $\Pi^*$ , and the transition sequence that it traverses as  $T^*$ . Some additional notation will be used for the analysis of orbital roadmap construction: Define  $\mathcal{B}_b(q) \in \mathbb{R}^d$  as an open  $d$ -dim. hyperball centered at  $q$  with radius  $b$ . Let  $\mu$  denote the measure of the  $\mathbb{C}$  space, which corresponds to the Lebesgue measure (generalized notion of volume). Denote the measure of a ball of radius one as  $\mu_1$ , and  $n$  as the number of samples in an orbit.

## 2.1 Asymptotic optimality of Sampling-based Algorithms

This work focuses on algorithms that build roadmaps similar to random geometric graphs (RGG) [17] via sampling. It has been shown that sampling-based roadmaps inherit properties of the underlying RGGs [22].

**Definition 1 (Roadmap).** A roadmap is defined as a graph  $\mathcal{G}_n(V_n, E_n)$ , where  $V_n$  corresponds to the  $n$  points of a sampling sequence  $\mathcal{X}_n$ . Edges  $e(u, v)$  between vertices  $u, v \in V_n$  are added in the edge set  $E_n$ , when:

- a)  $\|u, v\| \leq r_n$ , where  $r_n$  is the connection radius of the roadmap, and
- b) if the geodesic path connecting  $u$  to  $v$  is collision free.


Traditional roadmap construction [14, 13] has focused on the interior of  $\mathbb{C}_{\text{free}}$ . Given any start  $q_0$  and goal configuration  $q_1$  in the interior of  $\mathbb{C}_{\text{free}}$ , asymptotic optimality is defined in terms of the optimal path ( $\pi^*$ ) connecting  $q_0$  to  $q_1$ .

**Definition 2 (Asymptotically Optimal Motion Planning on  $\mathcal{G}_n$ ).** An algorithm that builds a roadmap  $\mathcal{G}_n$  in  $\mathbb{C}_{\text{free}}$ , and returns the shortest path  $\pi_{\mathcal{G}_n}$  connecting a start and goal query point lying in the interior of  $\mathbb{C}_{\text{free}}$ , is asymptotically optimal [11] if:  $\mathbf{C}(\pi_{\mathcal{G}_n}) \leq (1 + \epsilon) \cdot \mathbf{C}(\pi^*) \quad \forall \epsilon > 0, \text{ as } n \rightarrow \infty$ .

Prior work [13, 11] has provided precise bounds on the radius  $r_n$  so that the roadmap  $\mathcal{G}_n$  satisfies this property for all start and goal query points in the interior of  $\mathbb{C}_{\text{free}}$ .

## 2.2 Algorithmic Outline: Forward Search Tree over Orbital Roadmaps

Consider a high-level task planning algorithm ALGO that maintains roadmaps on orbits as well as a high level *forward search tree* representing the connectivity of orbits through transition states. This could be done explicitly or implicitly through factorization [9, 25, 6]. Fig. 3 depicts such an abstraction of the task planning space with roadmaps constructed inside orbits. Consider the multi-modal problem where a robot is tasked to pick-and-place a single object. The planner begins in  $q_s$  in  $\mathcal{O}_0$  where  $\mathcal{M}_0$  corresponds to *transit*. Transitions to the adjacent *transfer* mode  $\mathcal{M}_1$  are achieved through sampled grasps to reach  $\mathcal{O}_1, \mathcal{O}_2$ . Samples in the interior of  $\mathcal{O}_0$  connect  $q_s$  to  $t_1, t_2$ . Then, sampled stable positions with the object reach  $\mathcal{O}_3, \mathcal{O}_4$ .  $t_3$  is an arm configuration that achieves the desired object placement with grasp  $t_1$ . When  $q_g$  is reached in  $\mathcal{O}_4$  with motions connecting  $q_s \rightarrow t_1 \rightarrow t_3 \rightarrow q_g$  the solution can be reported.



**Fig. 3.** Left: A  $\mathbb{C}$ -space split into 2 modes with roadmaps (black) drawn within orbits connected by transition states (pink). The start and goal configurations are drawn in green and red respectively. Right: A high-level orbital graph keeping track of the connections between the roadmaps within orbits (nodes labeled  $\mathcal{O}_i$ ) through the transition states (nodes labeled  $t_j$ ).

**Definition 3 (Asymptotic Optimality of ALGO).** *Algorithm ALGO, which returns a feasible path  $\Pi^n$  after  $n$  iterations, is asymptotically optimal if:  $\mathbf{C}(\Pi^n) \leq (1 + \epsilon) \cdot \mathbf{C}(\Pi^*) \quad \forall \epsilon > 0, \text{ as } n \rightarrow \infty.$*

The conditions for ALGO to achieve AO are the following: (i) It can sample transition states over positive measure subsets of mode boundaries; (ii) The orbital roadmap eventually connects every configuration inside an orbit explored by ALGO; (iii) The number of sampled transitions  $n_t \rightarrow \infty$  as the number of iterations  $n \rightarrow \infty$ ; (iv) For each discovered orbit  $\mathcal{O}_i$  the size of its roadmap  $n_i \rightarrow \infty$  as  $n \rightarrow \infty$ .

Algorithm 1 provides a high-level description of ALGO. The inputs are the initial configuration  $q_s$ , a goal region  $G$ , and two positive parameters: the number of interior orbit samples  $N_m$  and neighboring transitions samples  $N_t$  added to the roadmap per iteration from the expanded orbit. A *forward search-tree*  $\mathbb{T}$  over orbits is built, beginning with the orbit containing  $q_s$ . Each iteration, every orbit must have a non-zero probability of being selected (assured by random node selection) as  $\mathcal{O}_{\text{sel}}$ . The subroutine `expand_roadmap` adds more nodes and edges to an asymptotically optimal roadmap construction algorithm (e.g., PRM\*, or FMT\*) with  $N_m$  more samples. The orbits are not known apriori but  $N_t$  new transition points are uniformly selected from the boundaries with all neighboring modes, and connected to the roadmap  $\mathcal{G}_{\mathcal{O}_{\text{sel}}}$ . Each of the new (empty) orbits are added to the tree  $\mathbb{T}$ , so that they get a chance to be expanded in the future.  $\Pi$  keeps track of the best cost path that reaches the goal region  $G$ .

---

**Algorithm 1:** `ALGO( $q_s, G, N_m, N_t$ )`

---

```

1  $\Pi \leftarrow \emptyset;$ 
2  $\mathbb{T}(V, E);$ 
3  $\mathbb{T}.V \leftarrow \mathbb{T}.V \cup \mathcal{O}(q_s);$ 
4 for  $n$  iters do
5    $\mathcal{O}_{\text{sel}} \leftarrow \text{uniform\_random}(\mathbb{T}.V);$ 
6    $\mathcal{G}_{\mathcal{O}_{\text{sel}}} \leftarrow \text{expand\_roadmap}(\mathcal{G}_{\mathcal{O}_{\text{sel}}}, N_m);$ 
7    $\mathcal{N}_{\mathcal{O}_{\text{sel}}} \leftarrow \text{uniform\_boundary\_sample}(\mathcal{O}_{\text{sel}}, N_t);$ 
8    $\mathcal{G}_{\mathcal{O}_{\text{sel}}} \leftarrow \text{add\_transitions}(\mathcal{G}_{\mathcal{O}_{\text{sel}}}, \mathcal{N}_{\mathcal{O}_{\text{sel}}});$ 
9    $\mathbb{T}.V \leftarrow \mathbb{T}.V \cup \mathcal{N}_{\mathcal{O}_{\text{sel}}};$ 
10   $\mathbb{T}.E \leftarrow \mathbb{T}.E \cup \{(\mathcal{O}_{\text{sel}}, \mathcal{O}_{\text{neighbor}}) \mid \forall \mathcal{O}_{\text{neighbor}} \in \mathcal{N}_{\mathcal{O}_{\text{sel}}}\};$ 
11   $\Pi \leftarrow \text{retrace\_path}(G);$ 
12 return  $\Pi;$ 

```

---

ALGO is introduced as a general version of several existing algorithms that are similar in structure, while differing in the order and exact nature of exploration of the TAMP search space. As such, arguments made in Sec. 4 should apply to algorithms that maintain the key properties of ALGO. The original Multi-modal PRM-based algorithm [8] expands every orbit per iteration. So does more recent work [18]. Random-MMP [9] selects one orbit and samples one neighboring transition per iteration. More recently FOBT [25] samples an orbit per iteration, and if it hasn't been previously explored, sets  $N_m = \Theta(N), N_t = \Theta(N)$  and argues asymptotic optimality when  $N \rightarrow \infty$ . The arguments made for arguing the AO of FOBT will be summarized in the next section.

### 3 Summary of previous results: FOBT

This section summarizes a previous result showing that FOBT is AO to highlight key assumptions for the new result presented later.

Previous work proposed the FOBT algorithm [25]. To show asymptotic optimality, this work uses topological tools to build a geometric construction, which traces a robust task planning solution possessing clearance in the interior of each orbit, while maintaining bounded error from the optimal task planning path. FOBT uses the following rule for the connection radius of the roadmap in each orbit  $\mathcal{O}$ :

$$r_{\mathcal{O}}(n) > 4\left(1 + \frac{1}{d_{\mathcal{O}}}\right)^{\frac{1}{d_{\mathcal{O}}}} \left(\frac{\mu(\mathcal{O})}{\mu_1}\right)^{\frac{1}{d_{\mathcal{O}}}} \left(\frac{\log n}{n}\right)^{\frac{1}{d_{\mathcal{O}}}}, \quad (1)$$

where  $d_{\mathcal{O}}$  is the dimensionality of the orbit,  $\mu$  being the Lebesgue measure or volume, and  $\mu_1$  is the volume of a unit hyperball. *Note that this radius is effectively **twice** the radius sufficient for  $\text{PRM}^*$  [13] for AO motion planning within an orbit.*

**Theorem 1 ([25] Theorem 1).** *Let  $\mathcal{G}_n$  be a geometric graph with  $n$  vertices constructed using the connection radius in Eq. 1 across the orbits of a task planning space. Let  $\mathbf{C}_n^{\text{FOBT}}$  be the minimum cost of a path on  $\mathcal{G}_n$  and  $\mathbf{C}^*$  the optimal solution cost, then:  $\Pr(\{\limsup_{n \rightarrow \infty} \mathbf{C}_n^{\text{FOBT}} = \mathbf{C}^*\}) = 1$ .*

Consider segments of the optimal path  $\pi_{t_i \rightarrow t_{i+1}}^*$ , which traverses an orbit between two transition states (Fig 4) that are a concatenation of:

- (i) One hyperball whose closure contains the start of the segment. (Assuming smooth boundaries, say  $\mathcal{B}_{\delta}(q_{t-})$ , such that  $t_i \in \overline{\mathcal{B}_{\delta}(q_{t-})}$ );
- (ii) One hyperball whose closure contains the end of the segment. (Assuming smooth boundaries, say  $\mathcal{B}_{\delta}(q_{t+})$ , such that  $t_{i+1} \in \overline{\mathcal{B}_{\delta}(q_{t+})}$ );
- (iii) A strongly  $\delta$ -clear path (refer to it as:  $\pi_{q_{t-} \rightarrow q_{t+}}$ ) joining the two hyperballs through the interior of the space, i.e., one that is always some  $\delta > 0$  away from obstacles.

**Assumption 1 (Construction of Segments)** *Similar to motion planning setups [13, 11], the interior of each segment  $\pi^*$  is robustly optimal, i.e., there exists a sequence of strongly clear paths, which are homotopically equivalent to optimal segments.*

**Assumption 2 (Hyperballs Around Transitions)** *The boundaries containing between orbits in neighboring modes are assumed to be smooth (lower dimensional) manifolds, such that it is possible to describe small enough hyperball regions of radius independent of the algorithm. It is additionally assumed that a transition sampler is capable of discovering all such positive volume regions by uniformly sampling the appropriate sub-manifold.*

Restating Equation 20 from [25], the probability of failing to be asymptotically optimal is shown to diminish as follows for any  $\epsilon > 0$

$$\begin{aligned} \Pr(\mathbf{C}_n^{\text{FOBT}} \geq (1 + \epsilon)\mathbf{C}^*) &\leq \\ \Pr(\{\text{Failing to sample in the neighborhood of transitions after } n \text{ samples}\}) &+ \\ \Pr(\{\text{Failing to trace bounded error path inside orbit after } n \text{ samples}\}) & \end{aligned} \quad (2)$$

It was shown that both of the probability terms on the right-hand side of the inequality go to 0 as  $n \rightarrow \infty$ . Given the query points  $q_s$  and  $q_g$ , the output from the algorithm ( $\Pi_{\text{FOBT}}^n$ ) after  $n$  iterations can be described as follows.

**Output of FOBT** ( $q_s, q_g$ ):

$$\Pi_{\text{FOBT}}^n = \bigoplus_{i=0}^M \pi_{q_i \rightarrow q_{i+1}} \text{ such that } \mathbf{C}_{\text{FOBT}}^n = \mathbf{C}(\Pi_{\text{FOBT}}^n) \leq (1 + \epsilon) \mathbf{C}^* \quad (3)$$

$$\text{where } q_i \in (q_s, T^{\text{FOBT}}, q_g), \quad T^{\text{FOBT}} = (t_i)_{i=1}^M. \quad (4)$$

A chief takeaway from the FOBT algorithm is that discovering a  $t_i$  is guaranteed by uniformly sampling the boundaries of modes (or orbits). Note that this property is a consequence of the assumptions about the spaces involved, and not a feature of the algorithm FOBT. It suffices to give this boundary sampling subroutine enough attempts, which tend to infinity.

## 4 New arguments in integrated task and motion planning

Previous work has provided an analysis that sampling-based planning solutions for TAMP are AO as long as they use a connection radius twice as large as that of AO sampling-based motion planners [25]. This section builds a sequence of arguments to show that the same connection radius as that in motion planning is also sufficient for TAMP.

**Summary of Arguments:** Theorem 2 demonstrates the existence of a “robust” solution to the TAMP problem given the current assumptions. Theorem 3 argues that an optimal TAMP solution must comprise of a sequence of optimal paths between orbit transitions. Theorem 4 argues that an AO roadmap-based motion planner will remain AO even for a query between a start and a goal on the boundary of an orbit. Theorem 5 uses Theorems 3 and 4 to argue that ALGO is AO if it can find a robust transition sequence. Theorem 6 details the conditions necessary for a boundary sampler to find such a transition sequence. Finally, bringing it all together, Theorem 7 proves that ALGO is AO.

Define the set of all possible valid finite transition sequences

$$\mathcal{T} = \{T = (t_i)_{i=1}^M \mid 0 < M \leq M_{max}\},$$

where each  $t_i$  is a transition state, and motion planning in a single orbit can connect to an orbit in a neighboring mode through  $t_{i+1}$ .  $M_{max}$  is assumed to be finite as in previous work [8, 25].

With a slight abuse of notation, define the cost of this transition sequence  $\mathbf{C}(T)$  as the least-cost task planning solution that can be obtained over  $T$  between an implicit start and goal state, traced along piecewise robustly optimal segments (similar to Assumption 1).

**Theorem 2 (Robust Optimality of Task Planning).** *For a task planning query with piece-wise analytic constraints, there exists a sequence of hyperballs on the boundaries*

with radius independent of  $n$ , such that any transition sequence passing through these regions has bounded error to the optimal cost. There can be multiple such sequences.

More specifically, there exists a non-empty set of transition sequences  $\mathcal{T}^+ \subset \mathcal{T}$  such that for all small  $\epsilon^+ > 0$

$$\mathbf{C}(T^+) \leq (1 + \epsilon^+)c^*, \quad \forall T^+ \in \mathcal{T}^+,$$

Without loss of generality, set  $T^+ = (t_i)_{i=1}^M$ . Then, there exist  $M$  balls  $\mathcal{B}_\theta(t_i)$  centered around transitions  $t_i$  with radius  $\theta > 0$  such that for all sequences  $T' = (t'_i)_{i=1}^M$  with  $t'_i \in \mathcal{B}_\theta(t_i)$  the cost is similarly bounded as  $\mathbf{C}(T') \leq (1 + \epsilon^+)c^*$ .

Note that each hyperball  $\mathcal{B}_\theta(t_i)$  has a dimensionality identical to the submanifold where  $t_i$  was sampled.

*Proof.* The proof is a direct consequence of Assumption 2, and the guaranteed existence of  $T^{\text{FOBT}}$  as shown in Eq. 4, which means for

$$\mathcal{T}^+ = \{T^{\text{FOBT}}\}, \quad \epsilon^+ = \epsilon \text{ (from FOBT)}$$


the existence of  $\mathcal{T}^+$  is assured. This directly proves the robust optimality of  $\Pi^*$  in the described setup for task planning. This additionally guarantees that the property holds for all small  $\epsilon^+ > 0$ .  $\square$

**Note:** This is a straightforward extension of the arguments presented in the previous work, and such robustness was inherently assumed therein. It should be pointed out that even though the minimal set  $\mathcal{T}^+ = \{T^{\text{FOBT}}\}$  which essentially lies in the neighborhood of  $T^*$  suffices to argue robust optimality, the relaxation of the definition to allow the existence of an arbitrary set  $\mathcal{T}^+$  captures a lot of situations in task planning where the optimal solution is rarely unique. Consider the problem of rearranging objects  $A$  and  $B$ . It is possible, the two solutions that transfer  $A$ , then  $B$  versus transferring  $B$  then  $A$ , are effectively identical in cost, though drastically different in terms of how the transition sequences look.

**Implication:** Theorem 2 indicates that there are positive volumes (the hyperballs around *desired* transition configurations) in the submanifold on which these transitions exist, and are sampled. This allows the cost of discovered solutions connected through these to have (an arbitrarily small) bounded deviation from the optimal cost.

**Theorem 3 (Pairwise-optimal Planning Over Robust Transition Sequence).** *Given an  $\epsilon^+$ -robust transition sequence  $T^+$ , a path  $\Pi$  constructed from optimal orbital segments traversing  $T^+$  maintains the  $\epsilon^+$  cost bound.*

*Proof.* Let  $T^+ \in \mathcal{T}^+$  and let  $(q_s, T^+, q_g)$  be the sequence of transitions with start and goal configurations concatenated at either end. Let  $\pi_{q_i \rightarrow q_{i+1}}$  denote a feasible motion planning solution over a pairwise connection and  $\pi_{q_i \rightarrow q_{i+1}}^*$  be an optimal connection.



**Fig. 4.** The image shows the segment of the trajectory inside an orbit, where the problem is decomposed into connecting to the transition points and motion planning in the interior.



Let  $\Pi^+$  be the path concatenations of  $\pi_{q_i \rightarrow q_{i+1}}$  such that  $\mathbf{C}(\Pi^+) \leq \mathbf{C}(T^+)$ . Let  $\Pi^*$  be the path concatenations of  $\pi_{q_i \rightarrow q_{i+1}}^*$ . The result immediately follows since

$$\begin{aligned} \mathbf{C}(\pi_{q_i \rightarrow q_{i+1}}^*) &\leq \mathbf{C}(\pi_{q_i \rightarrow q_{i+1}}) \quad \forall i, 0 \leq i \leq M \\ \implies \sum_{i=0}^M \mathbf{C}(\pi_{q_i \rightarrow q_{i+1}}^*) &\leq \mathbf{C}(T^+) \leq (1 + \epsilon^+)c^* \end{aligned}$$

This implies that it suffices to reason about the optimality of the pairwise motion planning problems, as long as the set of transitions from  $\mathcal{T}^+$  are sampled.  $\square$

**Theorem 4 (Pairwise-asymptotically Optimal Planning Between Robust Transition Sequence).** Given  $\mathcal{G}_n$  constructed by an asymptotically optimal roadmap-based planner and solution path  $\pi_{t_i \rightarrow t_{i+1}}^n$  found from  $\mathcal{G}_n$  for two transitions  $t_i$  to  $t_{i+1}$ , then

$$\lim_{n \rightarrow \infty} Pr(\{\mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^n) \geq (1 + \epsilon)\mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^*)\}) = 0, \quad \forall \epsilon > 0$$

*Proof.* This needs some additional consideration since the transition configurations lie on the boundary of the orbit, instead of the interior. Traditional sampling-based roadmap methods ([13, 11]) guarantee the following event for interior points, say from  $q_0$  to  $q_1$ . Restating Def 2:  $\lim_{n \rightarrow \infty} Pr(\{\mathbf{C}(\pi_{q_0 \rightarrow q_1}^n) \geq (1 + \epsilon)\mathbf{C}(\pi_{q_0 \rightarrow q_1}^*)\}) = 0, \quad \forall \epsilon > 0$ .

Note that  $q_0$ , and  $q_1$  are not unique in any way, and the above property essentially holds for any two points in the interior of the space. Let the connection radius be  $r_n > 0$ . Given Assumption 1 about the segment connecting  $t_i$  to  $t_{i+1}$ , there is a small enough hyperball that can touch the smooth boundary points (Fig. 5). Construct balls of radius  $\frac{r_n}{2}$ . As  $n$  increases,  $r_n$  decreases and at some point becomes sufficiently small to satisfy Assumption 1. Any point in the interior of such a ball is in the *interior* of the space. Let  $q_0^n$  be the center of such a ball, that gets closer to  $t_i$  as  $n$  grows and  $\frac{r_n}{2}$  shrinks. Any sample in such a ball must be connected to  $t_i$ .


Inspect the event that such a hyperball of radius  $\frac{r_n}{2}$  fails to have a sample in sample set  $\mathcal{X}_n$  for either  $t_i$  or  $t_{i+1}$ :  $\{\mathcal{B}_{\frac{r_n}{2}}(q_0^n) \cap \mathcal{X}_n = \emptyset\}$  and  $\{\mathcal{B}_{\frac{r_n}{2}}(q_1^n) \cap \mathcal{X}_n = \emptyset\}$ .

$$Pr(\{\mathcal{B}_{\frac{r_n}{2}}(q_0^n) \cap \mathcal{X}_n = \emptyset\}) = \left(1 - \frac{\mu(\mathcal{B}_{\frac{r_n}{2}}(q_0^n))}{\mu(\mathcal{O})}\right)^n \quad (5)$$

$$= \left(1 - \frac{\mu_1}{\mu(\mathcal{O})} \left(\frac{r_n}{2}\right)^d\right)^n \leq e^{-\frac{\mu_1}{2^d \mu(\mathcal{O})} n r_n^d} \quad (6)$$

$$\implies \lim_{n \rightarrow \infty} Pr(\{\mathcal{B}_{\frac{r_n}{2}}(q_0^n) \cap \mathcal{X}_n = \emptyset\}) = 0, \quad \text{when } \lim_{n \rightarrow \infty} n r_n^d \rightarrow \infty \quad (7)$$

The same argument holds for  $t_{i+1}$ . It is evident that the connection radii recommended by sampling-based roadmap planners already make this probability go to 0.



**Fig. 5.** Construction at smooth transition boundary.

Let  $q_0$  be the sample that the event  $\{\mathcal{B}_{\frac{r}{2}}(q_0^n) \cap \mathcal{X}_n \neq \emptyset\}$  generated, and similarly  $q_1$  for  $t_{i+1}$ . The exact configurations do not matter as they are guaranteed to be in the interior of the space.

So combining the failure conditions and by using the union bound it is possible to write:

$$Pr(\{\mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^n) \geq (1 + \epsilon)\mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^*)\}) \quad (8)$$

$$\leq Pr(\{\mathcal{B}_{\frac{r}{2}}(q_0^n) \cap \mathcal{X}_n = \emptyset\}) \quad (9)$$

$$+ Pr(\{\mathbf{C}(\pi_{q_0 \rightarrow q_1}^n) \geq (1 + \epsilon)\mathbf{C}(\pi_{q_0 \rightarrow q_1}^*)\}) \quad (10)$$

$$+ Pr(\{\mathcal{B}_{\frac{r}{2}}(q_1^n) \cap \mathcal{X}_n = \emptyset\}) \quad (11)$$

Take the limit on both sides. Since all the right hand side terms go to 0, the probability of the event  $\{\mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^n) \geq (1 + \epsilon)\mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^*)\}$  goes to 0 as well, indicating that this event ceases to happen asymptotically. Note that the underlying connection radius has not been changed.  $\square$

**Theorem 5 (Pairwise-asymptotically Optimal Planning Converges in Cost).** Given  $n_i$  samples in each pairwise motion planning problem  $t_i$  to  $t_{i+1}$  over the robust transition sequence  $T^+ \in \mathcal{T}^+$ , a path  $\Pi^n$  is generated after  $n$  total iterations of the algorithm. For all arbitrarily small  $\epsilon_+ > 0$ , the following holds:

$$\lim_{n \rightarrow \infty} Pr(\{\mathbf{C}(\Pi^n) > (1 + \epsilon_+)\mathbf{C}^*\}) = 0, \quad \forall \epsilon_+ > 0$$

*Proof.* Following the result of guaranteed convergence for each pairwise motion planning problem, it needs to be shown that each segment along  $T^+$  gets enough attempts ( $n_i$ ) to allow convergence. Thus a necessary algorithmic condition is that for each motion plan connecting  $t_i$  to  $t_{i+1}$ , an AO sampling-based roadmap needs  $n_i$  samples such that as  $n \rightarrow \infty$ ,  $n_i \rightarrow \infty$ ,  $\forall \mathcal{O}_i$ .

The rest of the proof follows by combining the pairwise results to compose  $\Pi^n$ . Asymptotically it is known:

$$\mathbf{C}(\Pi^n) = \sum_{i=1}^{M-1} \mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^{n_i}) \leq \sum_{i=1}^{M-1} (1 + \epsilon_i)\mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^*) \quad (12)$$

$$\leq \left(1 + \sum_{i=1}^{M-1} \epsilon_i\right)\mathbf{C}(T^+) \leq \left(1 + \sum_{i=1}^{M-1} \epsilon_i\right)(1 + \epsilon^+)\mathbf{C}^* \leq (1 + \epsilon_+)\mathbf{C}^*. \quad (13)$$

Since each of the *epsilon* terms are arbitrarily small by definition, for all small  $\epsilon_+ > 0$ , there would be some small enough values of  $\epsilon_1, \dots, \epsilon_{M-1}, \epsilon^+$  that satisfy the bound. Similarly the probability is evidently going to 0 since  $M$  is independent of  $n$  and  $n_i$ . For the estimated values of  $\epsilon_1, \dots, \epsilon_{M-1}, \epsilon^+$ , the probability follows from union bound

$$\lim_{n \rightarrow \infty} Pr(\{\mathbf{C}(\Pi^n) > (1 + \epsilon^+)\mathbf{C}^*\}) \quad (14)$$

$$\leq \lim_{n \rightarrow \infty} \sum_{i=1}^{M-1} Pr(\{\mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^{n_i}) > (1 + \epsilon_i)\mathbf{C}(\pi_{t_i \rightarrow t_{i+1}}^*)\}) = 0 \quad (15)$$

It follows that in a robustly optimal task planning problem the solution cost from ALGO can get arbitrarily close to the optimal solution cost if along a discovered transition sequence  $T^+ \in \mathcal{T}^+$  it:

- solves every pairwise transition connection over an orbit using an asymptotically optimal sampling-based motion planner *with the standard AO radius sufficient for motion planning*
- ensures every orbital roadmap ( $n_i$ ) grows infinitely as the total number of high-level iterations ( $n$ ) grows to infinity.

This resolves the second part of Eq. 2. Note that the argument holds over  $(q_s, T^+, q_g)$ .  $\square$

**Theorem 6 (Sampling Sequence of Transitions).** *A forward search tree  $\mathbb{T}$ , which selects an orbit per iteration with probability  $\Theta(\frac{1}{|\mathbb{T}|})$ , and uniformly samples an  $N_t$  expected number of neighboring transitions every iteration, is guaranteed to expand a sequence of transitions that are  $\epsilon^+$ -robust.*

*Proof.* Consider orbits to be selected uniformly at random from the working search tree. Let  $N_t > 0$  be the number of transitions added on each orbit expansion. For each desired pairwise transition  $t_i \rightarrow t_{i+1}$  from an orbit  $\mathcal{O}_i$ , robust optimality guarantees the existence of a positive volume around  $t_{i+1}$ . The probability of sampling in that volume is a small constant  $\epsilon' = \frac{\mu(\mathcal{B}_\epsilon(t_{i+1}))}{\mu_\cap} > 0$  which is independent of  $N_t$  or  $n$ . Here  $\epsilon$  is some small radius describing the region, and  $\mu_\cap$  the volume of the submanifold that is being sampled.

This formulation makes the search tree identical in behavior to the *naive random tree* described in previous work [16]. Reusing the arguments of [16](Theorem 18), by substituting the transition probability with any  $\epsilon' > 0$ , sampling the correct sequence of transitions from some  $T^+$  through the regions defined by Theorem 2 is guaranteed asymptotically. This resolves the first part of Eq. 2.  $\square$

Note that additionally [16](Theorem 3) ensures that every orbit will also be expanded infinitely often, guaranteeing all  $n_i \rightarrow \infty$  in each orbit when every expansion contributes to a positive expected number ( $N_m$ ) of samples to  $n_i$ . In our model,  $N_t$  simply needs to be positive, and can be a constant by the same argument, to ensure coverage of the boundary submanifold.

**Theorem 7 (Asymptotically Optimal Task Planning).** *When applied to a robust task planning problem, solvable by a finite number of mode transitions, ALGO is AO if*

1. *a forward search tree over transition configurations is sampled uniformly at random on the transition submanifold*
2. *the number of samples in each orbit and number of neighboring transitions increases asymptotically*
3. *the roadmap contained in an orbit uses a connection radius of*

$$r_n(\mathcal{O}) \geq \text{AO motion planning radius in each orbit } \mathcal{O} \text{ [13, 11]}$$

*Proof.* This follows immediately from Theorem 5 and 6.  $\square$

## 5 Model for Approaching Boundaries

The previous sections have relied on the assumption that the boundaries of modes are smooth (Fig. 4). This section shows that a relaxing this assumption to allow start/goal points to lie on a non-smooth boundary can still result in AO under certain conditions. The following discussion drops the  $\mathcal{M}$  superscript notation and refers to  $\mathbb{C}_{\text{free}} \subset \mathbb{C}$  as the subspace of obstacle free configurations (of a single mode).

The typical analysis framework for sampling-based motion planning focuses on tiling the interior of  $\mathbb{C}_{\text{free}}$  with hyperballs over solution paths. For smooth boundaries  $\partial\mathbb{C}_{\text{free}}$  it is possible to tile solution paths with balls that touch the boundary; but, any irregularity on the boundary will violate this condition. This is readily demonstrated in Fig. 2.

### 5.1 Cone Condition

To argue results for cases where the boundary is not smooth everywhere, this work borrows certain topological tools for non-smooth boundaries. The proposed framework still makes some assumptions in terms of the underlying space.

**Definition 4 (Cone).** A “ $q$ -cornet” [2]  $\mathcal{H}_b(q, \mathbf{v}, \phi)$  is the intersection of a convex cone with apex at  $q$ , and a hyperball  $\mathcal{B}_b(q)$  of radius  $b$ . The cone is symmetric about vector axis  $\mathbf{v}$  and the “opening” of the cone is parameterized by  $\phi = \frac{\mu(\mathcal{H}_b(q, \mathbf{v}, \phi))}{\mu(\mathcal{B}_b(q))} \in (0, \frac{1}{2}]$ .

With a slight abuse of notation, this work refers to a “ $q$ -cornet” as a “cone”, similar to the underlying literature [2].


**Assumption 3 (Cone Condition)** For every point  $q \in \partial\mathbb{C}_{\text{free}}$  there exist values  $b > 0$ ,  $\phi > 0$  and a vector  $\mathbf{v}$  so that there is a cone  $\mathcal{H}_b(q, \mathbf{v}, \phi) \in \mathbb{C}_{\text{free}}$ .

Note that  $\mathbb{C}_{\text{free}}$  automatically satisfies the (Poincaré) cone condition [2] in its interior since the underlying topology of the interior contains hyperballs at any configuration. This assumption is violated in pathological regions, such as degenerate narrow passages. The following proposition shows that the cones introduced by Assumption 3 have a sufficient intersection with the interior of  $\mathbb{C}_{\text{free}}$  to allow sampling processes to work.

**Proposition 1 (Intersection of Cone and Free Interior).** Given  $q \in \partial\mathbb{C}_{\text{free}}$  and its associated cone  $\mathcal{H}_b(q, \mathbf{v}, \phi)$  from the cone condition, there exists a point  $q'$  and a small enough radius  $b'$  so that  $\mathcal{B}_{b'}(q') \subset \mathcal{H}_b(q, \mathbf{v}, \phi) \cap \mathbb{C}_{\text{free}}$ , i.e., there is a hyperball at the intersection of the cone and the interior of the free configuration space.

*Proof.* Since  $\phi > 0$ ,  $\exists \mathcal{H}_b(q, \mathbf{v}, \phi) \subset \overline{\mathbb{C}_{\text{free}}}$  so that  $\mu(\mathcal{H}_b(q, \mathbf{v}, \phi)) > 0$ . This implies  $\mu(\mathcal{H}_b(q, \mathbf{v}, \phi) \cap \mathbb{C}_{\text{free}}) > 0$ , since  $\mu(\partial\mathbb{C}_{\text{free}}) = 0$ . Given the underlying topology of the space, the positive measure intersection  $\mathcal{H}_b(q, \mathbf{v}, \phi) \cap \mathbb{C}_{\text{free}}$  can contain a small enough hyperball  $\mathcal{B}_{b'}(q') \subset \{\mathcal{H}_b(q, \mathbf{v}, \phi) \cap \mathbb{C}_{\text{free}}\}$  for  $b' > 0$ .

Let  $\vartheta_q$  be the supremum radius  $b'$  of a hyperball at the intersection of the cone  $\mathcal{H}_b(q, \mathbf{v}, \phi)$  and  $\mathbb{C}_{\text{free}}$ . This maximum radius can also be defined for the start  $q_0$  and goal  $q_1$  query points as  $\vartheta_{q_0}$  and  $\vartheta_{q_1}$ , respectively.



**Fig. 6.** The figure describes the cone condition for boundary paths showing cones at the boundary points, and an intersection with the interior of the space.

## 5.2 Robust Clearance for Boundary Paths

The solution paths for the problems considered by this work must connect points that lie on the boundary  $\partial\mathbb{C}_{\text{free}}$ . Such paths will be referred to as “boundary paths”, as shown in Figure 6. The remainder of this section formally defines boundary paths and extends the notion of  $\delta$ -clearance to such paths. Finally, it shows that a motion planning problem that meets this extended notion of clearance for start and goal points can be solved using AO roadmaps.

**Definition 5 (Boundary paths).** *For a boundary path  $\pi$  it is true that  $\pi[0] \in \partial\mathbb{C}_{\text{free}}$  or  $\pi[1] \in \partial\mathbb{C}_{\text{free}}$ .*

Typically, certain clearance properties need to be satisfied for solution paths in order for sampling-based planners to be able to discover them. Consider a sequence of subspaces – parameterized by a decreasing  $\delta > 0$ , which approach the entire  $\mathbb{C}_{\text{free}}$ .


**Definition 6 ( $\delta$ -interior space).** *Given some  $\delta > 0$ , the  $\delta$ -interior space  $\mathbb{C}_\delta \subset \mathbb{C}_{\text{free}}$  consists of all configurations at least  $\delta$  distance away from  $\partial\mathbb{C}_{\text{free}}$ .*

The benefit of the cone condition and resulting proposition is that boundary paths can be decomposed into three segments: (a) one that passes through the cone defined at  $\pi[0] = q_0$ , (b) a second segment that transitions into the interior of  $\mathbb{C}_{\text{free}}$ , and (c) a third segment that connects to the cone defined at  $\pi[1] = q_1$ . Given this idea, the notion of “strong  $\delta$ -clearance for boundary paths” is introduced.

**Definition 7 (Strong  $\delta$ -clearance for Boundary Paths).** *A boundary path  $\pi$  satisfies “strong  $\delta$ -clearance for boundary paths” for some  $\delta > 0$ , if there are path parametrizations  $0 \leq t_- < t_+ \leq 1$  for the path, so that:*

- *the subset of the path  $\pi(0 : t_-)$  from  $q_0 = \pi[0] \in \partial\mathbb{C}_{\text{free}}$  to  $\pi[t_-] \in \mathbb{C}_\delta$  lies entirely in some  $\mathcal{H}_{\hat{b}_0}(q_0, \mathbf{v}_0, \phi_0) \subset \overline{\mathbb{C}_{\text{free}}}$  for  $\hat{b}_0, \phi_0 > 0$ .*
- *the subset of the path  $\pi(t_- : t_+)$  lies in the  $\mathbb{C}_\delta$*
- *the subset of the path  $\pi(t_+ : 1)$  from  $\pi[t_+] \in \mathbb{C}_\delta$  to  $q_1 = \pi[1] \in \partial\mathbb{C}_{\text{free}}$  lies entirely in some  $\mathcal{H}_{\hat{b}_1}(q_1, \mathbf{v}_1, \phi_1) \subset \overline{\mathbb{C}_{\text{free}}}$ , for  $\hat{b}_1, \phi_1 > 0$ .*

The construction of strongly  $\delta$ -clear boundary paths is feasible by Proposition 1 for some  $0 < \delta \leq \frac{1}{2} \min(\vartheta_{q_0}, \vartheta_{q_1})$ , when  $q_0$  and  $q_1$  are connected through  $\mathbb{C}_{\text{free}}$ . In general, any range of  $\delta$  where such a construction is possible can be considered. A view of such a path at one of the ends is shown in Figure 7 (right).



**Fig. 7.** (Left:) Strong  $\delta_\epsilon$ -clearance Convergence, and (Right:) Strong  $\delta$  convergence of some  $\pi_\epsilon$  w.r.t. a  $\pi^*$  that traces the boundary.

Note that the optimal boundary path  $\pi^*$  for a motion planning problem may approach arbitrarily close to obstacle boundaries and hence violate strong  $\delta$ -clearance conditions for boundary paths. In order to model regions of space in the vicinity of  $\pi^*$ , which contain “near-optimal” paths with a cost that gets arbitrarily close to  $C(\pi^*)$ , the notion of  $\delta_\epsilon$ -clear convergence from prior work [11] is adopted here.  $\pi^*$  has to exhibit some “weak clearance” and allow the existence of a sequence of strong  $\delta$ -clearance boundary paths that converge to it.

**Definition 8 (Strong  $\delta_\epsilon$ -clearance Convergence).** *A motion planning problem exhibits “strong  $\delta_\epsilon$ -clearance convergence” for boundary paths, if for all small  $\epsilon > 0$ , there exists some range of clearance values  $\delta \in (0, \delta_\epsilon]$ , such that a “strong  $\delta$ -clear boundary path”  $\pi_\epsilon$  with  $\pi_\epsilon[0] = \pi^*[0]$  and  $\pi_\epsilon[1] = \pi^*[1]$ , has its cost bounded relative to the cost of the optimal path  $\pi^*$  as follows:  $C(\pi_\epsilon) \leq (1 + \epsilon)C(\pi^*)$ .*

**Assumption 4** *Assume the motion planning problem under inspection exhibits “strong  $\delta_\epsilon$ -clearance convergence” (as shown in Fig. 7).*

**Theorem 8 (AO Connection Radius Is Sufficient for Non-smooth Boundaries).** *In the case where  $q_0, q_1 \in \partial C_{\text{free}}$ , the connection radius that guarantees asymptotic optimality of the interior of the space is sufficient for connecting  $q_0$  to  $q_1$ .*

*Proof.* The arguments are identical to the proof of Theorem 4. The probability of failing to connect these boundary points can be split into the probability of failing to connect each boundary point to the interior, and the probability of failing to connect two interior points with bounded error. The only difference occurs in Eq. 5-7, where the ball is replaced by the cone such that the volume now decreases by a constant fraction  $\phi$ . Recall that  $\mathcal{X}_n$  is a set of  $n$  samples and  $\mu_1$  denotes the volume of a ball of radius one.


$$Pr(\{\mathcal{H}_\phi(t_i, \mathbf{v}, r_n) \cap \mathcal{X}_n = \emptyset\}) = \left(1 - \frac{\mu(\mathcal{H}_\phi(t_i, \mathbf{v}, r_n))}{\mu(C_{\text{free}})}\right)^n \quad (16)$$

$$= \left(1 - \frac{\phi \mu_1 r_n^d}{\mu(C_{\text{free}})}\right)^n \leq e^{-\frac{\phi \mu_1}{\mu(C_{\text{free}})} n r_n^d} \quad (17)$$

$$\implies \lim_{n \rightarrow \infty} Pr(\{\mathcal{H}_\phi(t_i, \mathbf{v}, r_n) \cap \mathcal{X}_n = \emptyset\}) = 0, \text{ when } \lim_{n \rightarrow \infty} n r_n^d \rightarrow \infty \quad (18)$$

This bound is still readily satisfied by the connection radius argued in asymptotically optimal roadmap-based methods([13, 11]). It follows from the other arguments that motion planning inside an orbit, and task planning across a robustly optimal sequence of transitions both converge to the optimal cost asymptotically.  $\square$

An instance of such a motion planning problem for increasingly narrow *cones* is demonstrated in Fig. 8. Note that in the context of task planning the non-smoothness of the boundaries necessitate a stronger assumption about the *transition sampler* in task planning. The transition sampler needs to be aware of the precise submanifold to sample, which might no longer be the dimensionality of the boundary, since the apex of the cones might lie on a lower dimensional space. This is reasonable in practice. For instance, a grasp sampler for a parallel gripper will sample transitions that constrain the alignment of the object between the fingers.



**Fig. 8.** Instances of motion planning between two non-smooth boundary points in 2D, using the radius from  $\text{PRM}^*$ , for a red triangular robot that can only translate. The start is shown, and the green arc shows the path traced by the robot’s apex to reach the lower triangular cavity.

## 6 Discussion

This work aims to highlight useful analysis tools for understanding the properties of sampling-based TAMP planners in order to achieve AO. In particular, this paper argues that given reasonable assumptions about the TAMP planning problem an algorithm that can guarantee sufficient sampling of mode transitions as well as orbit interiors will retain the AO properties of the underlying sampling-based motion planner. Another contribution is the relaxation of the smooth boundary assumption widely applied by AO motion and TAMP planners. This work shows that it suffices that such boundary points are countably finite along the path and permit hypercones that open into the interior of the space. Moving away from smooth boundaries is reassuring as smoothness is hard to justify in practical problems, especially in manipulation problems with contact. It is also an issue for many motion planning queries, such as docking at a charging station.

In term of scalability, and practical performance, the algorithmic structure presented in this work is rather general, and if implemented naively would prove inefficient. Devising effective admissible heuristics in TAMP is critical for the quick discovery of high-quality solutions. It is also of interest to study whether the domain of transition sequences permit smoothing operations that are typically used in motion planning. The description of the homotopic properties in the multi-modal space can prove useful tools in improving TAMP solutions. The study of convergence rates and the inspection of finite time properties for AO planners is an important consideration. The design of practically efficient TAMP planners for realistic problems remains an active research area.

**ACKNOWLEDGMENTS:** The authors would like to acknowledge the support of NSF awards IIS:1617744 and CCF:1934924. Any findings reported here do not reflect the opinions of the sponsor.

## References

1. Alami, R., Simeon, T., Laumond, J.P.: A geometrical approach to planning manipulation tasks. the case of discrete placements and grasps (1990)
2. Cholaquidis, A., Cuevas, A., Fraiman, R.: On poincaré cone property. *The Annals of Statistics* **42**(1), 255–284 (2014)
3. Dantam, N.T., Kingston, Z.K., Chaudhuri, S., Kavraki, L.E.: Incremental task and motion planning: A constraint-based approach. In: *Robotics: Science and Systems* (2016)
4. Dobson, A., Bekris, K.E.: Planning Representations and Algorithms for Prehensile Multi-Arm Manipulation. In: *IEEE IROS* (2015)
5. Dogar, M.R., Srinivasa, S.S.: A Framework for Push-Grasping in Clutter. In: *Robotics: Science and Systems* (2011)
6. Garrett, C.R., Lozano-Pérez, T., Kaelbling, L.P.: Sample-based methods for factored task and motion planning. In: *Robotics: Science and Systems* (2017)
7. Garrett, C.R., Lozano-Perez, T., Kaelbling, L.P.: FFRob: Leveraging symbolic planning for efficient task and motion planning. *IJRR* **37**(1), 104–136 (2018)
8. Hauser, K., Latombe, J.C.: Multi-modal motion planning in non-expansive spaces. *The International Journal of Robotics Research* **29**(7), 897–915 (2010)
9. Hauser, K., Ng-Thow-Hing, V.: Randomized Multi-Modal Motion Planning for a Humanoid Robot Manipulation Task. *International Journal of Robotics Research* **30**(6) (2011)
10. Havur, G., Ozbilgin, G., Erdem, E., Patoglu, V.: Geometric rearrangement of multiple movable objects on cluttered surfaces: A hybrid reasoning approach. In: *IEEE ICRA* (2014)
11. Janson, L., Schmerling, E., Clark, A., Pavone, M.: Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *IJRR* (2015)
12. Kaelbling, L.P., Lozano-Pérez, T.: Hierarchical Task and Motion Planning in the Now. In: *IEEE ICRA* (2011)
13. Karaman, S., Frazzoli, E.: Sampling-based Algorithms for Optimal Motion Planning. *IJRR* **30**(7) (2011)
14. Kavraki, L.E., Svestka, P., Latombe, J.C., Overmars, M.: Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces. *IEEE Transactions on Robotics and Automation* **12**(4) (1996)
15. Koga, Y., Latombe, J.C.: On multi-arm manipulation planning. In: *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pp. 945–952. IEEE (1994)
16. Li, Y., Littlefield, Z., Bekris, K.E.: Asymptotically optimal sampling-based kinodynamic planning. *The International Journal of Robotics Research* **35**(5), 528–564 (2016)
17. Penrose, M.: *Random geometric graphs*. 5. Oxford university press (2003)
18. Schmitt, P.S., Neubauer, W., Feiten, W., Wurm, K.M., Wichert, G.V., Burgard, W.: Optimal, sampling-based manipulation planning. In: *ICRA*. IEEE (2017)
19. Shome, R., Bekris, K.E.: Anytime multi-arm task and motion planning for pick-and-place of individual objects via handoffs. In: *IEEE MRS*, pp. 37–43. IEEE (2019)
20. Shome, R., Solovey, K., Yu, J., Halperin, D., Bekris, K.E.: Fast, high-quality dual-arm rearrangement in synchronous, monotone tabletop setups. In: *WAFR* (2018)
21. Solovey, K., Kleinbort, M.: The critical radius in sampling-based motion planning. *IJRR* **39**(2-3), 266–285 (2020)
22. Solovey, K., Salzman, O., Halperin, D.: New perspective on sampling-based motion planning via random geometric graphs. *IJRR* **37**(10), 1117–1133 (2018)
23. Thomason, W., Knepper, R.A.: A unified sampling-based approach to integrated task and motion planning. In: *Proceedings International Symposium on Robotics Research* (2019)
24. Toussaint, M.: Logic-geometric programming: An optimization-based approach to combined task and motion planning. In: *International Joint Conference on Artificial Intelligence* (2015)
25. Vega-Brown, W., Roy, N.: Asymptotically optimal planning under piecewise-analytic constraints. In: *WAFR* (2016)