

Section Patterns: Efficiently Solving Narrow Passage Problems in Multilevel Motion Planning

Andreas Orthey¹ and Marc Toussaint^{1,2}

Abstract—Sampling-based planning methods often become inefficient due to narrow passages. Narrow passages induce a higher runtime, because the chance to sample them becomes vanishingly small. In recent work, we showed that narrow passages can be approached by relaxing the problem using admissible lower-dimensional projections of the state space. Those relaxations often increase the volume of narrow passages under projection. Solving the relaxed problem is often efficient and produces an admissible heuristic we can exploit. However, given a base path, i.e. a solution to a relaxed problem, there are currently no tailored methods to efficiently exploit the base path. To efficiently exploit the base path and thereby its admissible heuristic, we develop section patterns, which are solution strategies to efficiently exploit base paths in particular around narrow passages. To coordinate section patterns, we develop the pattern dance algorithm, which efficiently coordinates section patterns to reactively traverse narrow passages. We combine the pattern dance algorithm with previously developed multilevel planning algorithms and benchmark them on challenging planning problems like the Bugtrap, the double L-shape, an egress problem and on four pregrasp scenarios for a 37 degrees of freedom shadow hand mounted on a KUKA LWR robot. Our results confirm that section patterns are useful to efficiently solve high-dimensional narrow passage motion planning problems.

I. INTRODUCTION

Sampling-based motion planning algorithms are a successful paradigm to automate robotic tasks [52]. However, sampling-based algorithms do not perform well when the state space of the robot contains narrow passages [59, 94, 38, 80], which are low-measure regions which have to be traversed to reach a goal. Narrow passages are often occurring in tasks which are particularly important in robotic applications, like grasping, peg-in-hole, egress/ingress or long-horizon planning problems [24, 35].

In previous work, we and other research teams have shown that we can often efficiently solve high-dimensional planning problems by using admissible lower-dimensional projections of the state space, a topic we refer to as multilevel motion planning [22, 5, 68, 77, 100]. When using a multilevel motion planning framework, we can often use solutions to simplified planning stages as admissible heuristics for the original problem [71, 1]. To efficiently exploit those admissible

¹Max Planck Institute for Intelligent Systems, Stuttgart, Germany. Marc Toussaint thanks the MPI-IS for the Max Planck Fellowship. aorthey@is.mpg.de

²Technical University of Berlin, Germany

©2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.




Fig. 1: Efficient exploitation of admissible heuristics (stemming from solution to relaxed problem) using the triple step pattern. The triple step pattern is one of four section patterns we advocate to efficiently exploit admissible heuristics near narrow passages.

heuristics, we can use biased sampling methods [68, 76], which we can combine with classical planning algorithms like the rapidly-exploring random tree algorithm [64], the probabilistic roadmap planner [66], its optimal star versions [68] or the fast marching tree planner [76]. However, while showing promising runtimes, those algorithms are prone to get trapped when run on problems involving narrow passages.

In this work, we address narrow passages in multilevel motion planning problems by developing section patterns. Section patterns are methods to explicitly address problematic situations that occur when we exploit solutions to relaxed problems. We introduce four section patterns. First, we introduce the Manhattan pattern, which we use to compute solution paths which actuate the minimal amount of joints to reach a goal region, which is advantageous for high dimensional

systems [13, 68]. Second, we introduce the Wriggle pattern, which we use to make small random walk steps to traverse a narrow passage. Third, we introduce the Tunnel pattern, which we use to steer around small infeasible regions. Fourth, we introduce the Triple step pattern, which we use to backtrack in case the algorithm gets stuck. In Fig. 1 the Triple step pattern is showcased for a 37-degree of freedom (dof) robotic hand. We execute the pattern when a collision occurs (1). We first backstep (2), then sidestep (3) and finally we make a forward step (4) to reach a goal position. The details of this and the other patterns will be detailed later in this paper.

To coordinate the execution of the four section patterns, we develop a novel algorithm we call *pattern dance*. The pattern dance algorithm applies the section patterns sequentially by trying first a pattern which is easy to compute (Manhattan pattern) and reverting to the more complex pattern like Wriggle or Tunnel only if needed. If all those patterns fail, we revert to the Triple step pattern, which is the most computationally demanding pattern. We embed this pattern dance algorithm into four multilevel planners [68], namely the quotient space RRT (QRRT) [64], the quotient space roadmap planner (QMP) [66] and its optimal versions QRRT* and QMP* [68].

Our contributions are as follows.

- 1) We develop section patterns to efficiently exploit base space paths (solutions to relaxed problems).
- 2) To coordinate sections patterns, we develop the pattern dance algorithm.
- 3) We combine the pattern dance algorithm with four multilevel planners (QRRT, QRRT*, QMP, QMP*) and compare against 36 planners from the open motion planning library (OMPL) and a previous sidestepping algorithm [68] on 7 challenging scenarios.

II. RELATED WORK

Let us review the literature by focusing on two topics. First, we focus on generating admissible heuristics [20] for motion planning problems involving continuous domains [52]. We discuss sources of admissible heuristics like constraint relaxations, lazy search, informed trees and past experience. Second, given an admissible heuristic, we review methods to efficiently exploit the heuristic either using path section approaches, local minima avoidance or narrow passage handling.

A. Generating Admissible Heuristics

Motion planning [52] is a well studied topic which has been successfully applied to a wide range of problem domains [63]. One of the most promising paradigms to solve motion planning problems are (asymptotically optimal) sampling-based planners [45, 84, 83, 6, 25]. However, these planners might become inefficient in state spaces which are too high-dimensional [68], contain intricate constraints [42] or narrow passages [54]. We can, however, often solve such problems efficiently, if we use admissible heuristics [1].

We believe there are three large sources of admissible heuristics. First, we can compute admissible heuristics as solutions to relaxed problems [71]. Early instances of this idea to motion planning can be found in the constraint relaxation

frameworks by Ferbach and Barraquand [22], Sekhavat et al. [87] and Bayazit et al. [5]. Newer instances of this idea are putting the focus on different aspects like the specific type of projection [92, 30] or the type of lower-dimensional space [66, 9]. We refer to all those frameworks under the collective term multilevel motion planning [68]. We can apply multilevel frameworks both to holonomic [76, 77] and nonholonomic planning problems [100, 68]. To create multilevel abstraction, we can often remove links from a robot [5, 107], shrink links [3, 80] or approximate a robot by simpler geometries, either exact [66, 31] or approximate [10, 78, 95]. While most methods use prespecified levels of abstraction, we can also use workspace information to compute abstractions on the fly [104, 57], adaptively switch between abstractions [91] or learn useful abstractions for specific instances [9]. Our approach is similar, in that we also use a multilevel motion planning framework [68]. However, our work is complementary, in that we focus specifically on computing path sections in the presence of narrow passages in the state space.

A second source of admissible heuristics are lazy search [8, 33] and informed sets [26, 44]. Instead of using relaxations, we can compute lazy paths (paths not checked for collisions), either forward from the start [36] or backwards from the goal [90], to create an efficient heuristic which we can exploit using dedicated algorithms [27]. Once a solution exists, we can also exploit informed sets, sets which exclude all states with provable higher cost-to-go [26, 27]. Those methods are particularly important, since edge evaluations is one of the bottlenecks in motion planning [47]. It therefore makes sense to develop heuristics which evaluate edges as late as possible [61, 37].

Third, inspired by pattern database approaches in discrete search [14, 21, 39], we can also construct admissible heuristics by using past experience. We can achieve this by either precomputing motion primitives, like steering functions or controllers like linear quadratic regulators [82, 81]. Or, we can store previous solution paths directly and use them as heuristics in new environments [17, 75]. Our work is complementary in that we assume a heuristic given and we focus on exploiting this heuristic as efficiently as possible.

B. Exploiting Admissible Heuristic

Given an admissible heuristic, we can optimally exploit it by discretizing the state space [23] and by using the A* algorithm [34, 71, 1]. However, discretizing the state space usually does not scale well to higher dimensional state spaces [11, 72, 28] and performance would be sensitive to the resolution used [19]. To avoid discretization, we found three categories of work which use continuous methods to exploit admissible heuristics.

First, we can use biased sampling methods. A straightforward way would be to represent the heuristic value of a state by the radius of a hypersphere around the state [55]. We could then exploit this hypersphere using dynamic domain sampling [103]. Using such a scheme, we would expand states with higher heuristic values more often. Depending on the exact type of heuristic function used, we would obtain sampling

distributions which would increase the probability to sample states which are near to restricted workspace geometries [97, 102], to state space obstacles [2] or to narrow passage [38]. Those sampling distributions could also be learned over time to improve sampling [58, 41]. Our approach is similar in that we also use sampling-based methods. We differ, however, in that we concentrate on designing efficient patterns complementary to biased sampling methods.

Given a solution to a relaxed problem, we can often use this solution as a guide path heuristic [107, 95] to quickly find a solution in the original state space. Using the parlance of fiber bundles, we call this the find section problem [68]. This problem requires a relaxed solution (a base path), which we can find by computing workspace regions [73], by using workspace graphs [16, 96] or by using a simpler robot geometry [95]. In more complex environments, it is often advantageous to use multiple base paths [101, 16] which decompose the original problem into smaller subproblems [74, 7, 65]. To exploit a base path, we can often use restriction sampling [70, 66], which is highly efficient in high-dimensional state spaces, where uniform sampling would most likely fail to find solutions in a reasonable time [31]. Apart from biasing sampling, we can also explicitly search over the set of states which project onto the base path [107], which we call the path restriction. To find paths over path restrictions, we previously developed a sidestepping approach [68], where we propagate states along the path restriction and execute sidesteps when collision occur. However, as we show in Sec. V, sidesteps are often not beneficial for narrow passages. While we also search over path restrictions, we differ by developing dedicated patterns to more efficiently traverse narrow passages.

Path section approaches and other heuristic search methods often fail because they reach local minima. We define a local minimum as a region in state space where the heuristic is not or only weakly correlated with the true cost-to-go [99]. To address local minima, we can choose one of two approaches. First, we could preemptively avoid local minima. If the environment is static, we can learn minima regions and use this information to update the heuristic function [99]. Second, we could try to escape local minima. There exist several methods to escape local minima like deflating the heuristic value of states close to obstacles [18] or increasing the search resolution to prevent evaluation of closeby states [19]. A related idea is to utilize Tabu search [29] to prevent sampling in previously visited regions.

It is important to make the distinction between local minima which trap the planner and regions which might look like local minima but which a planner can actually traverse. We call such regions narrow passages [85]. To verify the existence of narrow passages in low-dimensional state spaces, we can use exact infeasibility proofs [86, 4], for example using geometrical shapes like alpha complices [62] or cell decomposition methods [106]. Because many state spaces have a local product structure, we can often use configuration space slices [56, 88] to efficiently test for infeasibility [98]. If the problem is feasible, we could then use the geometrical shapes to enumerate narrow passages [60]. To exploit narrow passages, we could bias sampling to the most constricted areas [102, 94]. We differ to those




Fig. 2: Left: Fiber bundle $\mathbb{R}^3 / \mathbb{R}^2$ with base space B , total space X , fiber space F , mappings π , π_F and fiber F_b over base element b . Right: Path restriction $r(p)$ over base path p . Adapted from [68].

approaches by not explicitly modeling narrow passages or local minima, but we instead develop reactive measures to escape minima and to traverse narrow passages. We thereby avoid spending time on irrelevant narrow passages.

III. BACKGROUND

Let us describe the necessary background to follow the exposition of our algorithm in Sec. IV and Sec. V. We start by explaining multilevel motion planning, i.e. planning with sequences of relaxed subproblems. While several formulations exist, we believe the framework of fiber bundles [68] to be a good way to concisely model multilevel abstractions and describe our algorithms. We then describe the concepts of lift, path restriction and path section which are particularly important. Finally, we describe the notion of admissible heuristics, which is one of the fundamental concepts to exploit solutions to relaxed problems [71].

A. Optimal Motion Planning

Let X be the state space of the robot. To each state space we associate a constraint function $c: X \rightarrow \{0, 1\}$ which evaluates to 0 if a state is constraint-free and to 1 otherwise. We use the constraint function to define the free state space $X_{\text{free}} = \{x \in X \mid c(x) = 0\}$. Together with an initial configuration $x_I \in X_{\text{free}}$ and a goal configuration $x_G \in X_{\text{free}}$, we define an optimal motion planning problem [45, 83, 6] as the tuple $(X_{\text{free}}, x_I, x_G, c)$, whereby our task is to develop an algorithm which computes a path from x_I to x_G while staying in X_{free} and minimizing the cost functional c . In this work, we use a minimal-length cost functional, but other costs are also possible like minimal energy or maximum clearance.

B. Multilevel Motion Planning

Since high-dimensional motion planning problems are often too computationally expensive to solve, we use a sequence of relaxed problems which we refer to as multilevel abstractions [68]. Given a state space X , let us denote a multilevel abstraction as the tuple (X_1, \dots, X_K, g) with $X_K = X$. To

each state space X_k , we associate a constraint function c_k and a projection π_k from X_k to X_{k-1} . We say that the projection π_k is admissible (w.r.t. the constraint functions), if $c_{k-1}(\pi_k(x)) \leq c_k(x)$ for any x in X_k . With admissibility, we basically guarantee that solutions are preserved under projections [64]. If we would allow inadmissible projections, we would potentially sacrifice solutions and thereby sacrifice (probabilistic) completeness.

C. Fiber Bundle Formulation

When working with multilevel abstraction, we quickly stumble upon situations where we lack the appropriate vocabulary to describe solution strategies. As a remedy, we describe multilevel abstractions using the framework of fiber bundles [89, 40, 53]. A fiber bundle is a tuple $(X_k; X_{k-1}; F_k; \pi_k; \pi_k)$ consisting of a total space X_k , a base space X_{k-1} , a fiber space F_k , a projection mapping π_k from total to base space and a fiber projection mapping π_k from total to fiber space. We assume the projection mapping π_k to be admissible. With a fiber bundle, we model product spaces which locally decompose as $X_k = X_{k-1} \times F_k$. The total space X_k is a union of fiber spaces which are parameterized by the base space X_{k-1} . If the level k is unimportant for the task at hand, we often refer to a fiber bundle as the tuple $(X; B; F; \pi; \pi)$ with X being the total, B the base, F the fiber space and π, π the base and fiber projection, respectively. We visualize a prototypical fiber bundle in Fig. 2 (left). For more details and motivation, we refer to our prior work [68]. For the purpose of this paper, we focus on the three concepts of lift, path restriction and path section, which we explain next.

D. Lift

Let $(X; B; F; \pi; \pi)$ be a fiber bundle and let $b \in B$ be a base space element. We often like to project the element b back to the total space X . We call this operation a lift [79, 68]. We define a lift as a mapping $LIFT : B \rightarrow X$. To uniquely select an element in X , we will overload this function as a mapping $LIFT : B \times F \rightarrow X$ by providing a fiber space element f in F . If X is a product space, we define the lift as $LIFT(b; f) = (b; f)$ [68].

E. Path Restrictions

Let $p : I \rightarrow B$ with $I = [0; 1]$ be a path on the base space (a base path). Given a base path, one of the most central sets which we use in this work are path restrictions. A path restriction is the set $r(p) = \{x \in X \mid \exists t \in I, p(t) = \pi(x)\}$, whereby $\pi : X \rightarrow B$ is the projection from X to B . We visualize this situation in Fig. 2 (right), where we show the image of a base path on the disk-shaped base space and its associated path restriction on the total space.

F. Path Sections

Given a path restriction, we are often interested in finding paths which are lying inside the path restriction. We call them path sections [89]. A (smooth) path section w.r.t. a base path

p is a continuous mapping s from base space B to total space X such that $\pi(s(u)) = u$ for any u in the image of p [53]. This means, for each base path element, we select a unique state from the path restriction—in a continuous manner.

G. Admissible Heuristics

Our motivation to introduce path restrictions and path sections comes from the role they play in exploiting admissible heuristics. Given a goal state x_G , an admissible heuristic $h(x)$ for a state x in X is a lower-bound on the true cost-to-go (or value) function $h^*(x)$, which we define as the cost of the optimal path from x to x_G through X_{free} . Formally, we write this condition as $h(x) \leq h^*(x)$ [71, 1, 64].

Given an admissible heuristic, we can try to reach the goal x_G by using locally optimal decisions [34]. If we are at a state x , we can make an optimal decision by doing a two-step approach. First, we compute the f -value of all its neighbors, which is the sum of its heuristic value and its cost-to-come from the start state. We then expand the state (node) with the lowest f -value, because, under the admissible heuristic, it is our best guess to efficiently reach the goal [71].

However, in a continuous domain, we cannot straightforwardly compute all neighboring states. Instead, we imagine computing a small ϵ -neighborhood around the state. To compute heuristic values, we project the complete neighborhood down onto the base space. To reach the goal, our best guess is to make a step into the direction of the current minimal-cost base path. The states which we would expand in that way are exactly the states on the path restriction. By searching a path section over this path restriction, we efficiently exploit the admissible heuristic given by the base path.

IV. FIND SECTIONS USING PATTERN DANCE

Our goal is to develop an algorithm which solves the find section problem, the problem of finding a path section over a given path restriction. After we state the problem, we discuss how the problem fits into the more general framework of motion planning using multilevel abstractions [68]. Finally, we discuss the pattern dance algorithm, which coordinates four section patterns to efficiently find feasible path sections.

A. Find Section Problem

Let $(X; B; F; \pi; \pi)$ be a fiber bundle on X (possibly in a sequence of fiber bundles) and let $p : I \rightarrow B$ be a base path on B starting at (x_I) and ending at (x_G) . Given the base path p and its path restriction $r(p) \subset X$, our goal is to develop an algorithm to find a feasible path section, i.e. a path lying in the intersection of the path restriction $r(p)$ and the free state space X_{free} connecting x_I to x_G . We call this problem the *find section problem*.

To illustrate the find section problem, we visualize it in Fig. 3. The figure shows a base path p on B (bottom) and its restriction $r(p)$ on X (top). Our goal is to connect x_I to x_G while staying inside $r(p)$. To efficiently solve the find section problem, we often need to track information along the path restriction. To track this information, we introduce the notion




Fig. 3: Path restriction $r(p)$ on a total space X over a base path p from base space B , together with initial state x_I , goal state x_G , projection π and head pointer with head pointer H , consisting of state x and location l .

Algorithm 1 MultilevelPlanner($x_I; x_G; X_1; \dots; X_K$)

Input: Initial state x_I , goal state x_G , state spaces $X_1; \dots; X_K$

Output: Graphs $G_1; \dots; G_K$

```

1: Let  $\mathbf{X}$  be a PRIORITY QUEUE
2: for  $k = 1$  to  $K$  do
3:   FINDSECTION( $X_k$ )
4:    $\mathbf{X}$ .PUSH( $X_k$ ; IMPORTANCE( $X_k$ ))
5:   while : PTC( $X_k$ ) do
6:      $X_{\text{select}} = \mathbf{X}$ .POP
7:     GROW( $X_{\text{select}}$ )
8:      $\mathbf{X}$ .PUSH( $X_{\text{select}}$ ; IMPORTANCE( $X_{\text{select}}$ ))
9:   end while
10: end for
11: return GRAPHS( $X_1; \dots; X_K$ )

```

of a head pointer H as the tuple $H = (x; l; r)$ consisting of a path restriction $r(p) : X$ over a base path p in B , a current state x in $r(p)$ and a location $l \in [0; 1]$ defining the position along the base path. We think of the head pointer as a ruler which we move forward along the path restriction towards the goal state. In pseudocode, we refer to the current state as STATE(H) and its location as LOCATION(H).

B. Find Sections in Multilevel Planning

The find section problem is a subproblem of the more general multilevel motion planning problem (see Sec. III-B). In previous works, we proposed to solve multilevel planning problems using a dedicated multilevel planner [68]. To clarify the role of finding sections, we describe this multilevel planner in Alg. 1. We initialize this algorithm with an initial state x_I , a goal state x_G and a sequence of bundle spaces $X_1; \dots; X_K$. To search for a feasible path, we first initialize a priority queue (Line 1), then we iteratively explore the bundle spaces (Line 2) by first trying to solve the find section problem (Line 3),

Algorithm 2 FindSection(X_k)

```

1: if EXISTS( $X_{k-1}$ ) then
2:    $\rho = \text{BASEPATH}(G_{k-1})$ 
3:    $r = \text{RESTRICTION}(\rho)$ 
4:    $H = \text{HEADPOINTER}(x_I; \text{location} = 0; r)$ 
5:   PATTERN DANCE( $H$ )
6: end if

```

then pushing the k -th bundle space into the priority queue (Line 4). We compute the importance of a bundle space by the sampling density of its associated graph [68] as

$$\text{IMPORTANCE}(X_k) = \frac{1}{j|V_k|^{1-n_k} + 1} \quad (1)$$

with $j|V_k|$ being the number of nodes in the graph G_k on X_k and n_k is the dimensionality of X_k . We then go into a while loop which terminates if a planner terminate condition (PTC) of the k -th space is not fulfilled (Line 5). A PTC can be a timelimit, an iteration limit or a desired cost. We then pop the space with the highest importance from the queue (Line 6), execute one grow iteration for the selected bundle space (Line 7) and push the space back to the queue thereby updating its importance (Line 8). The planner terminates if the PTC of all bundle spaces is false and returns the graphs of all computed levels (Line 11). From those graphs, we can then compute the (optimal) solution path using a discrete A* search [34] (if one exists). All multilevel planners share this high-level structure. Multilevel planners differ by how the GROW function is implemented.

We previously developed four multilevel planners. First, the quotient-space roadmap planner (QMP), in which we implement GROW as a probabilistic roadmap (PRM) step [46]. Second, the quotient-space rapidly-exploring random tree (QRRT), in which we implement GROW as an RRT step [48]. Finally, we use the two asymptotically optimal versions QRRT* and QMP*, in which we implement a step of RRT* and PRM* [45], respectively. The algorithms also differ in how we compute the distance metric and how we implement sampling inside the grow function, as we detail in our previous publication [68].

The main contribution of our paper, the pattern dance algorithm, is an efficient method to solve the find section problem. The integration into the multilevel planner is shown in the FINDSECTION method in Alg. 2. First, we check if there exists a base space (Line 1). We then compute a base path p from the underlying graph or tree on the base space (Line 2). We then build a path restriction r from p (Line 3) and create a head on the path restriction (Line 4). We then call the pattern dance algorithm with the head as input.

C. Pattern Dance Algorithm

We depict the pseudocode of the pattern dance algorithm in Alg. 3. The input is a head over the path restriction and a recursion depth (initially set to zero). Inside the pattern dance algorithm, we coordinate the execution of four section patterns. The rationale behind the coordination is to try less

Algorithm 3 PatternDance(H ; depth = 0)

Parameters: Maximum branching factor B_{\max} , maximum depth factor D_{\max} , base space step size χ_{k-1}

```

1: if MANHATTANPATTERN( $H$ ) then
2:   return true
3: end if
4: if depth  $D_{\max}$  then
5:   return false
6: end if
7: if WRIGGLEPATTERN( $H$ ) or TUNNELPATTERN( $H$ ) then
8:   return PATTERNDANCE( $H$ , depth+1)
9: end if
10:  $l$  LOCATION( $H$ ) +  $\chi_{k-1}$ 
11:  $X_B$  BASEPATHAT( $p$ ;  $l$ )
12: for  $j \in [1; B_{\max}]$  do
13:    $X_F$  SAMPLEFIBER( $X_B$ )
14:    $X$  LIFT( $X_B$ ;  $X_F$ )
15:    $X_H$  STATE( $H$ )
16:   if ISVALID( $X$ ) and :CHECKMOTION( $X_H$ ;  $X$ ) then
17:     if TRIPLESTEPATTERN( $H$ ;  $X$ ) then
18:       return PATTERNDANCE( $H$ , depth+1)
19:     end if
20:   end if
21: end for

```

complex patterns first while we can successfully move the head forward along the path restriction. Only if no progress is made, we revert to more and more complex patterns to resolve the situation. We found this to be an efficient strategy to quickly find sections.

Those four section patterns are detailed in Sec. V and either move the head forward by controlling the lowest amount of joints possible (MANHATTANPATTERN), execute random walk steps with forward bias (WRIGGLEPATTERN), try to overcome small barriers using steps outside the path restriction (TUNNELPATTERN) or use a dedicated backtracking procedure (TRIPLESTEPATTERN) to efficiently find feasible path sections.

Before going into detail, we provide a brief summary and motivation. The algorithm iterates through all four patterns, starting with the computationally cheapest MANHATTANPATTERN (Line 1). If the pattern succeeds, we successfully return (Line 2). Otherwise, we check if we reached the maximum recursion depth (Line 4) and return with failure (Line 5).

If the depth is below the maximum depth, we continue by executing first the WRIGGLEPATTERN and the TUNNELPATTERN (Line 7). If one pattern successfully terminates, we recursively call the pattern dance algorithm and we increase the recursion depth (Line 8). If no pattern successfully terminates, we backtrack using the TRIPLESTEPATTERN. To execute the triple step pattern, we first interpolate a single step forward along the base path (Line 10, 11). We then attempt to find a valid fiber space element for a maximum of B_{\max} attempts (Line 12). This is done by first sampling a fiber state over the given base state (Line 13). We then lift the state to the path restriction (Line 14) to obtain a state x . If this state is valid and we *cannot* reach it from the head state (Line 16),

Parameter	Description	Values used
D_{\max}	Maximum depth of pattern dance	3
B_{\max}	Maximum branching of pattern dance	500
S_{\max}	Maximum sampling attempts	100
$\delta_{X_{k-1}}$	Step size on base space	$0.01\mu_{X_{k-1}}$
δ_{F_k}	Step size on fiber space	$0.01\mu_{F_k}$

TABLE I: Parameters used in algorithm. The variable χ refers to the measure (volume) of the state space X .

we execute the triple step pattern with target x (Line 17). If we successfully executed the pattern, we call the pattern dance algorithm again recursively. Note that the small forward step of χ_{k-1} (Line 10) is an essential component of our algorithm. If we would sample directly over the head base state, we often would sample symmetrical local minima (as an example, see state p_i^l in Fig. 8). We found this to be particularly important for higher dimensional state spaces, where we often encounter infinitely many symmetrical local minima (consider the set of horizontal rotations of the cylinder before entering the opening in the Bugtrap scenario in Sec. VI).

To implement the section patterns and the pattern dance algorithm, we use the open motion planning library (OMPL) [93]. The algorithms are freely available and part of our multilevel motion planning extension of OMPL [68]. All code can be downloaded over github¹. All parameters used in the algorithms are shown in Table I, including the values we use for the evaluations. The values for B_{\max} ; S_{\max} ; D_{\max} are chosen as large as possible to still give good performance on our hardware.

V. SECTION PATTERNS

The pattern dance algorithm relies on four section patterns, to which we like to provide more detail and motivation. Each of those section patterns is a particular approach to efficiently traverse narrow passages and escape local minima, whereby a local minimum is defined as a region where the heuristic cost is only weakly correlated with the true cost-to-go [99]. Each section pattern takes as input a head pointer and tries to move this head pointer forward along the path restriction. Please also consult Fig. 3 for visualization of the terminology used.

A. Manhattan Pattern

Our first section pattern to propagate the head pointer H is the Manhattan (MH) pattern. With the MH pattern, we interpolate a path between the head state and the goal state along the path restriction. To interpolate, we first interpolate along the base path while keeping the fiber element fixed. Once we reach the end of the base path, we interpolate along the fiber space to the goal state. This method is motivated by our desire to actuate the smallest number of joints at the same time, which is advantageous for high-dimensional systems [13].

We detail the MH pattern in Alg. 4. We take as input a head pointer H over a path restriction r with base path p .

¹<https://github.com/aorthey/MotionExplorer/> and <https://github.com/aorthey/ompl/>.

Algorithm 4 ManhattanPattern(H)

Parameters: Base space step size x_{k-1}

- 1: x_H STATE(H)
- 2: x_F PROJECTFIBER(x_H) \cdot $F(x_H)$
- 3: l LOCATION(H)
- 4: S ;
- 5: **while** $l < \text{LENGTH}(\rho)$ **do**
- 6: x_B BASEPATHAT($\rho; l$) \cdot State $\rho(l)$ on base path
- 7: x LIFT($x_B; x_F$)
- 8: S $S \cup \{x\}$
- 9: l $l + x_{k-1}$
- 10: **end while**
- 11: S $S \cup \{x_{GG}\}$
- 12: H CHECKMOTION(S) \cdot Return Last Valid
- 13: **return** HASREACHEDGOAL(H)




Fig. 4: Left: Rectangular rigid robot which has to traverse a narrow passage from a green start to a red goal state. Right: The geometry of its state spaces (darker colors are closer to start state).

We first project the head state onto the fiber (Line 1-2) by using the fiber projection F . We then take the location of the head pointer along the base path (Line 3) and step along the base path in increments of x_{k-1} (Line 5-10) and add the states to the path S (Line 4). This is done by computing the next base state (Line 6), lifting the base state into the total space (Line 7) and adding it to the path (Line 8). Once we reached the end of the base path, we add the goal state to the section (Line 11). The resulting path S is schematically shown in Fig. 3. Finally, we evaluate the path by moving along until a constraint violation occurs or we reached the goal state (Line 12). The function CHECKMOTION returns the last valid state which we use to update the head H . We then return true if the head has reached the goal and false otherwise.




Fig. 5: Two path restrictions for the rectangular rigid robot near a narrow passage. Left: Robot moves along a straight base path. Right: Robot moves along a slanted base path.

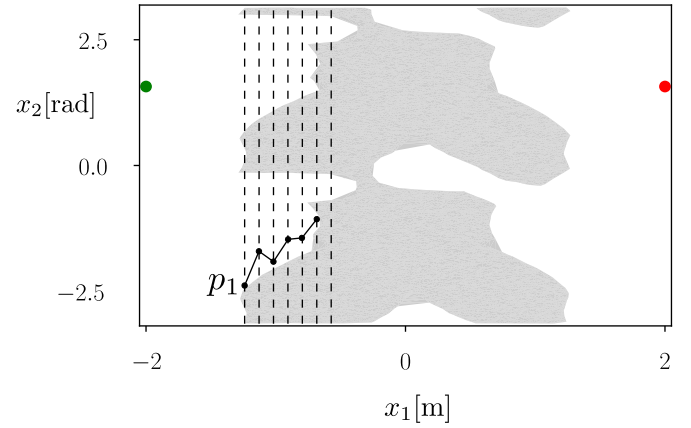


Fig. 6: Wriggle pattern to traverse a narrow passage: Given a feasible state p_1 , we make coordinated random walk steps along the fibers of the path restriction. The distance between fibers is determined by the base space step size parameter x_{k-1} .

B. Interlude: The Geometry near Narrow Passages

The next three section patterns are tailor-made solutions to either traverse a narrow passage or to escape a local minimum. To motivate those patterns, we first study the geometry of state spaces near narrow passages. We use a simple toy example of a rigid rectangular body moving in the 2D plane. The state space of this rigid body is the special Euclidean group $SE(2)$, consisting of position and orientation. We assume that the body is located near to a narrow passages as shown in Fig. 4 (left). We will further assume that our task is to move the rigid body through the narrow passage, from a start state (green) to a goal state (red). We will represent a state as $(x_0; x_1; x_2) \in SE(2)$, with $x_0; x_1$ being vertical and horizontal displacement and x_2 the orientation. We visualize a subset of the state space in Fig. 4 (right), whereby points in collision are colored from dark red (low x_1 value, close to start) to bright blue (high x_1 value, close to goal).

To generate path restrictions, we first use a relaxation of the problem onto a circular disk as shown in Fig. 4 (Left). We model this relaxation using the fiber bundle $SE(2) \times \mathbb{R}^2$ with base space \mathbb{R}^2 and total space $SE(2)$ [66]. Let us assume a base path $\rho: I \rightarrow \mathbb{R}^2$ for the disk to be given. This path induces a two-dimensional path restriction in $SE(2)$, two of which we visualize in Fig. 5. The left figure shows a path restriction for a base path going straight through the passage, as shown in Fig. 4. The right figure shows a path restriction for a base path which goes slanted through the passage. Both are also slices through the state space geometry shown in Fig. 4 (right). From Fig. 5, we observe that there are at least three failure cases. Either, we reach a local minimum, we collide with constraints near a narrow passage or we get stuck in front of a small but infeasible region. For each case, we develop a dedicated section pattern to either advance or backtrack.

Algorithm 5 WrigglePattern(H)

Parameters: Base space step size x_{k-1} , fiber space step size F_k , maximum samples S_{\max}

```

1:  $l \leftarrow \text{LOCATION}(H) + x_{k-1}$ 
2: steps  $\leftarrow 0$ 
3: while  $l < \text{LENGTH}(\rho)$  do
4:    $X_B \leftarrow \text{BASEPATHAT}(\rho; l)$ 
5:    $X_H \leftarrow \text{STATE}(H)$ 
6:    $X_{F_H} \leftarrow \text{PROJECTFIBER}(X_H)$ 
7:   ctr  $\leftarrow 0$ 
8:   while ctr  $< S_{\max}$  do
9:      $X_F \leftarrow \text{SAMPLEUNIFORMNEAR}(X_{F_H}; F_k)$ 
10:     $X \leftarrow \text{LIFT}(X_B; X_F)$ 
11:    if ISVALID( $X$ ) then
12:      if CHECKMOTION( $X_H; X$ ) then
13:         $G_k \leftarrow G_k \cup \{X_H; X\}$ 
14:        UPDATEHEAD( $H; X$ )
15:        steps  $\leftarrow$  steps + 1
16:        break
17:      end if
18:    end if
19:    ctr  $\leftarrow$  ctr + 1
20:  end while
21:  if ctr  $\geq S_{\max}$  then
22:    break
23:  end if
24: end while
25: return steps  $> 0$ 

```

C. Wriggle Pattern

If we reach a local minimum, the triple step pattern is a way to backtrack to a narrow passage. However, we often might execute the triple step pattern prematurely, because we bumped into constraints near or in a narrow passage. To circumvent those situations, we use the wriggle pattern. With the wriggle pattern, we make coordinated random steps along the fibers of the path restriction and accept a step if it is valid, which is similar to retraction-based sampling [105]. We visualize this pattern in Fig. 6.

We show the pseudocode in Alg. 5. We start by making one x_{k-1} step forward from the head (Line 1). Until we have not reached the end (Line 3), we get the base state at location l (Line 4), and get the fiber element of the head state (Line 6). We then sample for S_{\max} rounds (Line 8) by sampling a fiber state in the F_k proximity of the head fiber state (Line 9). We then lift the base and fiber state (Line 10) and check if the state is valid (Line 11). If the state is valid, we check if the motion from the head to the new state is feasible (Line 12-17). We terminate if we could not expand the state (Line 21-23) or reach the end. We then return true if we made at least one step (Line 25).

D. Tunnel Pattern

While the wriggle pattern locally explores the neighborhood *inside* the path restriction, we often encounter situations where we find it advantageous to momentarily step *outside* the

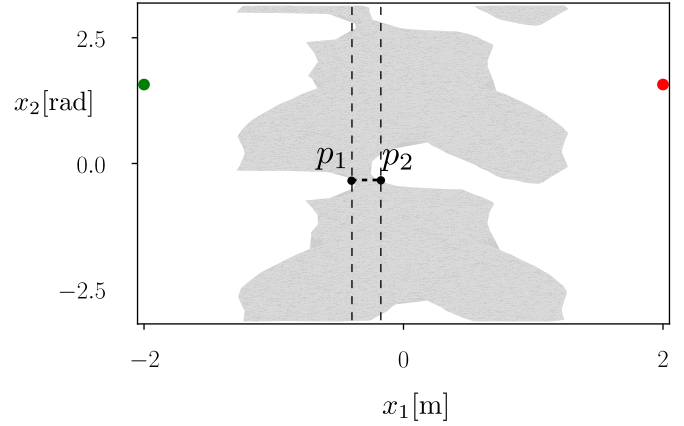


Fig. 7: Tunnel pattern to traverse a narrow passage: Given two feasible states p_1 and p_2 , we connect them by momentarily leaving the path restriction to circumnavigate the infeasible region between them.

path restriction to overcome an infeasible region. From the perspective of the path restriction, we “tunnel” through the infeasible region, which we therefore refer to as the tunnel pattern. With the tunnel pattern, we assume to be located at a local minimum p_1 as shown in Fig. 7. To resolve this situation, we try to find the next valid state p_2 while keeping the fiber element constant. We then try to connect p_1 to p_2 by sampling valid states in a smoothly increasing neighborhood of the base space and a constant neighborhood in fiber space. While p_2 is not reached, we accept new states if they decrease the distance to p_2 .

We show the pseudocode in Alg. 6. We first search for a tunnel ending state X_{End} at base path location l_{End} (Line 1). To find the tunnel ending, we step forward along the base path without changing the fiber until we find a valid state. We then try to connect the head state X_H to the tunnel ending state X_{End} . We use a while loop to move along the relevant base path segment from the head location l to the tunnel end location l_{End} (Line 6). We first check if we can connect the head state to the tunnel end state (Line 7). If true, we add a new edge into the graph (Line 8), set the head to the tunnel ending state (Line 9) and return true (Line 10). Otherwise, we step forward along the base path with step size x_{k-1} (Line 12) and query the base state at l (Line 13). Instead of using the base state exactly, we use a smoothly increasing neighborhood parameter δ . The value of δ depends on the counter ctr and smoothly interpolates between 0 and $10 \cdot x_{k-1}$ using an Hermite polynomial [15] (Line 14). We then attempt to make a step towards the tunnel ending for a maximum of S_{\max} attempts (Line 16). We do this by sampling a base space element (Line 17) and a fiber element (Line 18). We then lift the state (Line 19) and check for validity (Line 20). If the new state is valid, its distance is closer to the tunnel ending and we can connect it to the head state (Line 22), we add a new edge to the graph (Line 23), set the head state to the new state (Line 24) and continue forward (Line 25). If we fail to find a better sample for S_{\max} attempts, we return false (Line 30-32). We also return false if we reach the base path location l_{End} without having a valid connection

Algorithm 6 TunnelPattern(H)

Parameters: Base space step size x_{k-1} , fiber space step size F_k , maximum samples S_{\max}

```

1: ( $X_{\text{End}}; l_{\text{End}}$ ) TUNNELEND( $H$ )
2:  $X_H$  STATE( $H$ )
3:  $X_{F_H}$  PROJECTFIBER( $X_H$ )
4:  $d_{\text{best}}$  DISTANCE( $X_H; X_{\text{End}}$ )
5:  $l$  LOCATION( $H$ )
6: while  $l \neq l_{\text{End}}$  do
7:   if CHECKMOTION( $X_H; X_{\text{End}}$ ) then
8:      $G_k \leftarrow G_k \cup \{X_H; X_{\text{End}}\}$ 
9:     UPDATEHEAD( $H; X_{\text{End}}$ )
10:    return true
11:  end if
12:   $l \leftarrow l + x_{k-1}$ 
13:   $X_B$  BASEPATHAT( $p; l$ )
14:  SMOOTHPARAMETER( $0; 10 \cdot x_{k-1}; S_{\max}$ )
15:  ctr  $\leftarrow 0$ 
16:  while ctr <  $S_{\max}$  do
17:     $X_B$  SAMPLEUNIFORMNEAR( $X_B; \text{ctr}$ )
18:     $X_F$  SAMPLEUNIFORMNEAR( $X_{F_H}; F_k$ )
19:     $X$  LIFT( $X_B; X_F$ )
20:    if ISVALID( $X$ ) then
21:       $d$  DISTANCE( $X; X_{\text{End}}$ )
22:      if  $d < d_{\text{best}}$  and CHECKMOTION( $X_H; X$ ) then
23:         $G_k \leftarrow G_k \cup \{X_H; X\}$ 
24:         $X_H \leftarrow X$ 
25:        break
26:      end if
27:    end if
28:    ctr  $\leftarrow \text{ctr} + 1$ 
29:  end while
30:  if ctr  $\geq S_{\max}$  then
31:    return false
32:  end if
33: end while
34: return false

```

(Line 34).

E. Triple Step Pattern

To escape a local minimum, we develop the triple step pattern. With the triple step pattern, we connect two states on the path restriction using a triple backtracking step.

The idea of the triple step pattern is to connect two states on (or near) the same fiber. Before explaining the pattern in detail, we first visualize the pattern in Fig. 8. You can see a rectangular rigid body in the plane, which is currently at state p_1 (Fig. 8 (a)) and which we like to move to state p_4 (Fig. 8 (d)). To connect those states, we first move backwards along the path restriction from p_1 to another state p_2 (Fig. 8 (b)) while moving from p_4 to another state p_3 (Fig. 8 (c)), respectively. We move backwards until we can connect p_2 and p_3 by a straight line segment. In that case we execute a backstep from p_1 to p_2 , a sidestep (along the fiber marked) from p_2 to p_3 and a forward step from p_3 to p_4 . Note that p_4

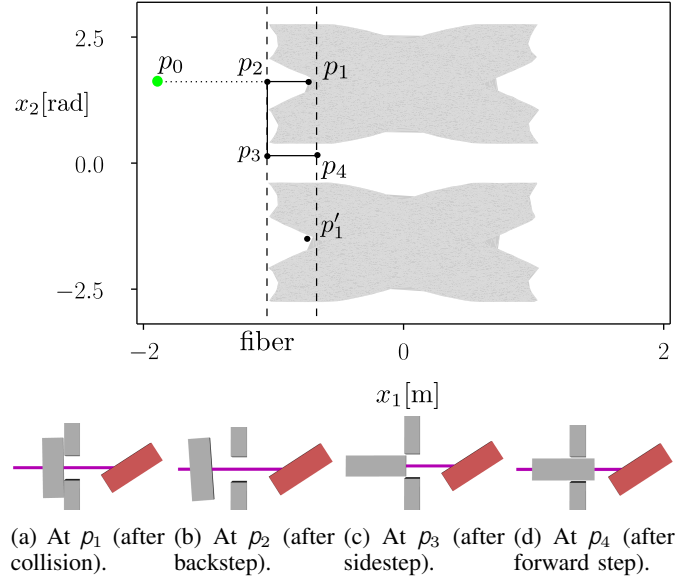


Fig. 8: Triple step pattern to traverse a narrow passage: We start at a state p_1 (a), backstep to a state p_2 (b), sidestep along the fiber to p_3 (c) and then step forward to reach a state p_4 (d).

is slightly moved forward such that we avoid situations where we backtrack to a symmetric local minimum like p'_1 which would not improve our location along the path restriction.

We show the pseudocode for the triple step pattern in Alg. 7. Our goal is to connect the head state to the given state X . We first compute a midpoint on the fiber space (Line 5) (to minimize the number of CHECKMOTION calls [61]). We then move backwards along the base path while we are greater than the parameter x_{k-1} (Line 6-7). For each location, we interpolate a base state (Line 8), lift the state using the fiber midpoint (Line 9) and check if this state is valid. If it is valid, we compute intermediate states x_1 and x_2 (Line 11, 12) and check if the motion between them is feasible (Line 13). If that is true, we additionally check if the backward and forward steps are feasible (Line 14, 15). If that is true, we add those edges to the graph (Line 16-18) and update the head to our new state X (Line 19). In that case we return true (Line 20). If we fail to find such a triple step, we terminate once we reach the beginning of the base path location and return false (Line 27).

VI. EVALUATIONS

To evaluate our pattern dance algorithm, we integrate it into the multilevel planner QRRT, QRRT*, QMP and QMP*, as we discussed in Sec. IV-B. We then conduct two comparisons. First, we compare our planner to 36 available planning algorithms in the Open motion planning library (OMPL) [63] on 7 challenging environments as shown in Fig. 9. For each algorithm, we use the abbreviated name. For a full list of algorithms with full names and associated publication, see [68] and the OMPL documentation [93]. Second, we compare the multilevel planner with the pattern dance algorithm to an




Fig. 9: Scenarios for evaluations. The task is to move the robot from the start state (green) to the goal state (red). Top Row (left to right): Bugtrap (6-dof), Double L Shape (6-dof) (goal configuration not shown) and Chain Egress (10-dof). Bottom Row: Overhand, Underhand, Single-Finger and Double-Finger Pregrasp (each 37-dof) (start configurations not shown).

older version of the same multilevel planner, where we use a recursive sidestepping algorithm to quickly find sections [68].

A. Evaluation Metric

To evaluate, we use a 8GB RAM 4-core 2.5GHz laptop running Ubuntu 16.04. For each experiment, we use a minimum length cost (for planner which support cost functions) and we let each planner run 10 times with a cut-off time limit of 60 seconds. We then report on the average runtime over those 10 runs. We show the results in Table II.

Concerning the results, there are two notes of caution. First, we let each OMPL planner run out-of-the-box without any parameter tuning. Further tuning of parameters could potentially improve results significantly. Second, due to the high number of planner and scenarios, we let each planner run only 10 times and take the average. However, averaging over 10 runs might exhibit more variance and thereby create more outliers.

B. 06-dof Bugtrap

For the first evaluation, we use the Bugtrap scenario [54] (Fig. 9a). The lowest runtime we found in the literature is 22.17s for a version of the Selective-Retraction-RRT [54, 105]. However, this runtime is not directly comparable due to

different hardware, implementation, parameters and operating systems. To relax the problem, we use an inscribed sphere at the center of the cylindrical bug as shown in Fig. 10g and Fig. 10d.

We show the results of our evaluation in Fig. II. The best performing planner is QMP (3rd planner in table) with 0.51s followed by QMP* (4) with 0.90s and QRRT (1) with 4.45s. We also see good performance of the BiTRRT (13) planner [43] with 11.54s. We note that the QRRT* (2) algorithm requires 24.87s, which we believe to be caused by the additional burden of rewiring the tree [84, 68].

C. 06-dof Double L shape

In the next evaluation, we like to show that the section patterns are not specific to the cylindrical geometry, but are more widely applicable to other rigid bodies. As demonstration, we use the double L-shape scenario [97], where two L-shape bodies are connected to each other as shown in Fig. 9b. The task is to move through a vertical wall with a small quadratic hole. We use a two-level relaxation by using an inscribed sphere as shown in Fig. 10h and Fig. 10e. To make our method more robust against base paths too close to obstacles, we increase the size of the sphere slightly to increase clearance from obstacles.

Algorithm 7 TripleStepPattern($H; X$)

Parameters: Base space step size X_{k-1}

```

1:  $X_H$  STATE( $H$ )
2:  $l$  LOCATION( $H$ )
3:  $X_{F_1}$  PROJECTFIBER( $H$ )
4:  $X_{F_2}$  PROJECTFIBER( $X$ )
5:  $X_{F_m}$  STEER( $X_{F_1}; X_{F_2}; 0.5$ ) . Midpoint
6: while  $l > X_{k-1}$  do
7:    $l = l - X_{k-1}$ 
8:    $X_B$  BASEPATHAT( $p; l$ )
9:    $X_{mid}$  LIFT( $X_B; X_{F_m}$ )
10:  if ISVALID( $X_{mid}$ ) then
11:     $X_1$  LIFT( $X_B; X_{F_1}$ )
12:     $X_2$  LIFT( $X_B; X_{F_2}$ )
13:    if CHECKMOTION( $X_1; X_2$ ) then
14:      if CHECKMOTION( $X_H; X_1$ ) then
15:        if CHECKMOTION( $X_2; X$ ) then
16:           $G_k$   $G_k [fX_H; X_1g$ 
17:           $G_k$   $G_k [fX_1; X_2g$ 
18:           $G_k$   $G_k [fX_2; Xg$ 
19:          UPDATEHEAD( $h; X$ )
20:          return true
21:        end if
22:      end if
23:    break . End While Loop
24:  end if
25: end while
26: return false

```

Our evaluation shows that QMP performs best with 1.27s followed by QMP* (1.63s), QRRT (1.86s) and QRRT* (2.00s). The next best planner from OMPL is LBKPIECE1 (38) with 49.79s.

D. 10-dof Chain Egress

In the third evaluation, we like to increase the complexity by considering an articulated chain (10-dof) as shown in Fig. 9c. The task is to remove the chain from a pipe, a typical egress scenario. Note that for such systems, we can find analytical feasible path sections if we assume the base path of the head to be curvature constrained [67]. However, we will not make such assumption in this paper.

To relax the problem, we use an inscribed sphere in the head of the chain as shown in Fig. 10i and Fig. 10f. As in the case of the double L-shape, we slightly increase the size of the sphere to make our method more robust against base paths too close to obstacles.

In our evaluations, we show that QRRT performs best with 0.55s followed by QRRT* (0.56s). The next best planners are TRRT (11) (0.81s), QMP (1.91), BiTRRT (12) (4.57s) and QMP* with 7.29s. Note that there are 12 OMPL planner which cannot address this problem, because they do not support compound state spaces or do not have dedicated projection functions for such spaces.




Fig. 10: Multilevel abstraction using simplified models.

E. 37-dof Pre-Grasp

For the next evaluations, we compute (pre-)grasping paths for a ShadowHand mounted on a KUKA LWR robot. The tasks are to compute an overhand grasp on a ball (Fig. 9d), an underhand grasp on a metal piece (Fig. 9e), a single-finger precision grasp on a mug (Fig. 9f) and a double-finger precision grasp on a scissor (Fig. 9g). The starting state for all scenarios is an upright position of the arm with hand being open, as shown in Fig. 10a. To relax the problem, we use a three-level abstraction by first removing three fingers (Fig. 10b) and subsequently removing the thumb (Fig. 10c) of the hand.

Our evaluations show the following results. First, for the Ball scenario, we see that QMP and QMP* perform best with 0.86s. The next best planner is the OMPL planner BiRLRT (15) [57] with 1.52s, QRRT with 2.01s and RRTConnect (6) with 1.70s. We note that also the planner PDST (35) [51], RLRT (14) [57] and KPIECE1 (36) [92] perform competitively with 3.25s, 3.68s and 6.27s, respectively. The planner QRRT* does not perform well on this problem instance with 25.35s, due to similar problems as on the Bugtrap scenario. Second, for the underhand grasp on the metal piece, we see that QMP* performs best with 1.94s followed by RRTConnect (6) with 8.16s and QMP with 18.98s. We will address the discrepancy between QMP and QMP* further in Sec. VII. Third, for the

Runtime in seconds (10 run average)		06D Bugtrap	06D Double Lshape	10D Chain Egress	37D ShadowHand Ball	37D ShadowHand Metal	37D ShadowHand Mug	37D ShadowHand Scissor
1	QRRT (ours)	4.45	1.86	0.55	2.01	35.63	19.80	60.00
2	QRRT* (ours)	24.87	2.00	0.56	25.35	43.95	60.00	60.00
3	QMP (ours)	0.51	1.27	1.91	0.86	18.98	1.20	14.52
4	QMP* (ours)	0.90	1.63	7.29	0.86	1.94	1.63	37.27
5	RRT	60.00	60.00	49.77	60.00	60.00	60.00	60.00
6	RRTConnect	60.00	60.00	60.00	<i>1.70</i>	<i>8.16</i>	57.38	60.00
7	RRT#	60.00	60.00	45.43	60.00	60.00	60.00	60.00
8	RRT*	60.00	60.00	51.74	60.00	60.00	60.00	60.00
9	RRTXstatic	60.00	60.00	50.49	60.00	60.00	60.00	60.00
10	LazyRRT	60.00	60.00	55.56	60.00	60.00	60.00	60.00
11	TRRT	60.00	60.00	<i>0.81</i>	42.08	60.00	60.00	60.00
12	BiTRRT	11.54	54.30	<i>4.57</i>	60.00	60.00	60.00	60.00
13	LBTRRT	60.00	60.00	60.00	60.00	60.00	60.00	60.00
14	RLRT	60.00	60.00	51.39	<i>3.68</i>	28.47	60.00	60.00
15	BiRLRT	60.00	57.40	60.00	<i>1.52</i>	25.60	60.00	60.00
16	pRRT	60.00	60.00	49.41	60.00	60.00	60.00	60.00
17	FMT	60.00	60.00	60.00	60.00	60.00	60.00	60.00
18	BFMT	60.00	50.34	60.00	60.00	60.00	60.00	60.00
19	PRM	60.00	56.47	60.00	37.25	52.72	60.00	60.00
20	PRM*	60.00	57.80	60.00	34.24	50.04	60.00	60.00
21	LazyPRM	60.00	60.00	60.00	60.00	60.00	60.00	60.00
22	LazyPRM*	60.00	60.00	60.00	54.06	60.00	60.00	60.00
23	SPARS	60.00	59.73	60.00	60.00	60.00	60.00	60.00
24	SPARStwo	60.00	54.69	60.00	60.00	60.00	60.00	60.00
25	SST	60.00	60.00	60.00	60.00	60.00	60.00	60.00
26	EST	60.00	60.00	50.46	24.96	45.64	60.00	60.00
27	BiEST	60.00	60.00	59.85	29.79	33.36	60.00	60.00
28	InformedRRT*	60.00	60.00	-	60.00	60.00	60.00	60.00
29	SORRT*	60.00	60.00	-	60.00	60.00	60.00	60.00
30	kBIT*	60.00	60.00	-	34.17	46.44	60.00	60.00
31	kABIT*	60.00	60.00	-	50.28	44.56	60.00	60.00
32	AIT*	60.00	60.00	-	55.35	60.00	60.00	60.00
33	STRIDE	60.00	60.00	-	29.58	48.98	60.00	60.00
34	ProjEST	60.00	60.00	-	47.77	60.00	60.00	60.00
35	PDST	60.00	60.00	-	<i>3.25</i>	54.42	60.00	60.00
36	KPIECE1	60.00	60.00	-	<i>6.27</i>	32.48	60.00	60.00
37	BKPIECE1	60.00	60.00	-	52.35	60.00	60.00	60.00
38	LBKPIECE1	60.00	49.79	-	60.00	60.00	60.00	60.00
39	SBL	60.00	50.30	-	60.00	60.00	60.00	60.00
40	CForest	60.00	60.00	-	60.00	60.00	60.00	60.00

TABLE II: Runtime (s) of motion planner on the scenarios from Fig. 9, each averaged over 10 runs with cut-off time limit of 60s. An entry --- means that planner does not support the particular state space.

single-finger precision grasp on the mug, we observe that QMP performs best with 1.20s followed by QMP* with 1.63s. While QRRT performs significantly worse (19:80s), QRRT* was not able to solve this problem (60:00s). Fourth, for the double-finger precision grasp on the scissor, we observe that QMP performs best with 14.52s followed by QMP* with 37:27s.

Runtime in seconds (10 run average)		06D Bugtrap	06D Double Lshape	10D Chain Egress	37D ShadowHand Ball	37D ShadowHand Metal	37D ShadowHand Mug	37D ShadowHand Scissor
1	QMP (ours)	0.51	1.27	1.91	0.86	18.98	1.20	14.52
2	QMP (SideStepping)	60.00	26.08	60.00	1.07	55.37	6 ^a	60.00
3	QMP* (ours)	0.90	1.63	7.29	0.86	1.94	1.63	37.27
4	QMP* (SideStepping)	60.00	30.11	60.00	1.76	60.00	12 ^a	60.00
5	QRRT (ours)	4.45	1.86	0.55	2.01	35.63	19.80	60.00
6	QRRT (SideStepping)	60.00	27.72	9.14	18.65	60.00	44 ^a	60.00
7	QRRT* (ours)	24.87	2.00	0.56	25.35	43.95	60.00	60.00
8	QRRT* (SideStepping)	60.00	60.00	16.42	42.33	54.05	48 ^a	60.00

TABLE III: Comparison of multilevel planners with sidestepping [68] versus multilevel planner with our pattern dance algorithm.

^a Taken from [68].

No other planner is able to solve this problem. We will further discuss the high runtime of both QMP and QMP* in detail in Sec. VII.

VII. LIMITATIONS AND DISCUSSION

While our evaluations support the usage of section patterns for narrow passage planning problems, we also like to point out two limitations of our approach. To each limitation, we will discuss possible ways to eventually address and resolve the limitation.




Fig. 11: Limitations of Section Pattern Approach. Base path does not admit a feasible path section. See text for clarification.

A. Increased runtime on Metal and Scissor Scenario

The first limitation is the increased runtime of our planner on the 37D ShadowHand Scissor and the Metal scenario. We distinguish between two subproblems. First, we observe that QRRT and QRRT* have a runtime of 60s on the Scissor scenario. Both scenarios, however, are ingress scenarios, where the planner needs to find a narrow passage on the base space

Run	1	2	3	4	5	6	7	8	9	10
37D ShadowHand Metal Scenario										
QMP	1.53	1.11	1.20	0.99	1.06	60.00	60.00	2.93	1.02	60.00
QMP*	0.98	1.15	0.93	1.23	2.73	1.13	1.03	7.61	0.98	1.65
37D ShadowHand Scissor Scenario										
QMP	1.45	1.50	2.14	2.17	60.00	60.00	2.44	7.49	1.51	6.51
QMP*	60.00	60.00	2.22	60.00	6.27	60.00	60.00	60.00	1.92	2.30

TABLE IV: Runtime (s) for QMP and QMP* on each run. Average runtimes are 18.98s/1.94s (QMP/QMP*) for the Metal scenarios and 14.52s/37.27s for the Scissor scenario.

to enter the goal region, which is challenging for RRT-like algorithms [48] and could be addressed using a bidirectional version of QRRT.

Second, we observe that QMP and QMP* require 14:52s and 37:27s to solve the Scissor scenario and that QMP requires 18:98s to solve the Metal scenario. To explain this rather large increase in runtime, we have a closer look at the individual runtimes, which we show in Table IV. We can observe that both planner exhibit one of two outputs. Either, they quickly return a solution (usually less than 3s, always less than 10s) or they fail and time out at 60s (three/two times for QMP, zero/six times for QMP*). To us, this indicates that both algorithms might be sensitive to the base space path. If the base path is not smooth enough, has kinks in it or is too close to obstacles, then we might not be able to solve it with the pattern dance algorithm. We could address this problem in the future by either additional smoothing of the base space path [100], by introducing conservative heuristics [12] or by switching to a different relaxed model [91].

B. Base path does not admit a feasible section

While all multilevel planner are probabilistically complete, we often need the pattern dance algorithm to efficiently solve a problem. However, we might encounter scenarios, where the base path does not admit a feasible path section. Such a situation is shown in Fig. 11. The scenario depicts an X-shape robot, which has to traverse a shape-sorter box with different openings, which we relax by inscribing a sphere (right). Planning for the spherical robot might produce a base path going through the wrong hole. Such a base path does not admit a feasible path section, meaning there are no paths along the path restriction of the base path to traverse towards the goal. While multilevel planner are probabilistically complete and would eventually resolve the situation, we would not be able to solve this situation using our pattern dance algorithm. To address such situations, we could either compute several base paths [65, 32, 101, 69, 7, 74] and consider them as a multi-arm bandit problem over path restrictions [49] or we could automatically choose an alternative relaxation using either a meta-heuristic [9] or a brute-force search [64].

VIII. CONCLUSION

We developed the pattern dance algorithm, which takes as input a base space path and efficiently searches for a feasible

section in its path restriction using four dedicated section patterns, which we named Manhattan, Wriggle, Tunnel and Triple step. We showed in evaluations, that our pattern dance algorithm successfully coordinates section patterns and outperforms a similar sidestepping algorithm [68]. We then showed that multilevel motion planning algorithms using our pattern dance algorithm outperform classical planner from the OMPL library on challenging narrow passage scenarios including the Bugtrap, chain egress and precision grasping. With some exceptions, we often observed runtime improvements by one to two orders of magnitudes.

While we demonstrated to efficiently solve narrow passage problems, we also pointed out two limitations. First, we observe an increased runtime in some planning instances. We could address this problem by either optimizing the base path [107], by improved neighborhood modeling [50] or by learning the section patterns themselves [41]. Second, we cannot handle cases where the base path does not admit a path section. We could address this problem by computing multiple base paths [65, 69, 101] or using more informed graph restriction sampling methods [64].

Despite limitations, we believe to have contributed a novel solution method which we can use to efficiently find sections over base path restrictions. We believe our method to be a promising tool to further probe, understand and efficiently exploit high-dimensional state spaces.

IX. ACKNOWLEDGEMENT

Marc Toussaint thanks the Max Planck Institute for Intelligent Systems for the Max Planck Fellowship.



Andreas Orthey is a postdoctoral researcher in computational robotics at the TU Berlin funded by the Max Planck Institute for Intelligent Systems (MPI-IS). Previously, he has been a Research Fellow with the Alexander von Humboldt Foundation (AvH) at the University of Stuttgart, the Japan Society for the Promotion of Science (JSPS) at the AIST in Tsukuba, Japan, a Postdoctoral Researcher at the Worcester Polytechnic Institute (WPI), MA, USA and a Doctoral Candidate at the LAAS-CNRS in Toulouse, France. He holds a PhD Degree from INP

Toulouse and a Master's Degree with Honours from the TU Berlin. His research interest lies in optimization and planning for complex and high-dimensional robotic systems.



Marc Toussaint is professor for Intelligent Systems at TU Berlin since March 2020 and Max Planck Fellow at the MPI for Intelligent Systems since November 2018. In 2017/18 he spend a year as visiting scholar at MIT, before that some months with Amazon Robotics, and was professor for Machine Learning and Robotics at the University of Stuttgart since 2012. In his view, a key in understanding and creating intelligence is the interplay of learning and reasoning, where learning becomes the enabler for strongly generalizing reasoning and acting in our physical world. His research therefore bridges between AI planning, machine learning, and robotics. His work was awarded best paper at R:SS'18 and ICMLA'07, and runner up at R:SS'12 and UAI'08.

REFERENCES

- [1] S. Aine, S. Swaminathan, V. Narayanan, V. Hwang, and M. Likhachev, "Multi-heuristic A*," *International Journal of Robotics Research* vol. 35, no. 1-3, pp. 224–243, 2016.
- [2] N. M. Amato, O. B. Bayazit, L. K. Dale, C. Jones, and D. Vallejo, "OBPRM: An obstacle-based PRM for 3D workspaces," *Workshop on the Algorithmic Foundations of Robotics* 1998, pp. 155–168.
- [3] B. Baginski, "Local motion planning for manipulators based on shrinking and growing geometry models," *IEEE International Conference on Robotics and Automation* Citeseer, 1996, pp. 3303–3308.
- [4] J. Basch, L. J. Guibas, D. Hsu, and A. T. Nguyen, "Disconnection proofs for motion planning," *IEEE International Conference on Robotics and Automation* vol. 2. IEEE, 2001, pp. 1765–1772.
- [5] O. B. Bayazit, D. Xie, and N. M. Amato, "Iterative relaxation of constraints: a framework for improving automated motion planning," in *IEEE International Conference on Intelligent Robots and Systems* 2005, pp. 3433–3440.
- [6] K. E. Bekris and R. Shome, "Asymptotically optimal sampling-based planners," *arXiv preprint arXiv:1911.04044* 2020.
- [7] S. Bhattacharya and R. Ghrist, "Path homotopy invariants and their application to optimal trajectory planning," *Annals of Mathematics and Artificial Intelligence*, vol. 84, no. 3-4, pp. 139–160, 2018.
- [8] R. Bohlin and L. E. Kavraki, "Path planning using lazy PRM," *IEEE International Conference on Robotics and Automation* vol. 1. IEEE, 2000, pp. 521–528.
- [9] M. Brandao and I. Havoutis, "Learning sequences of approximations for hierarchical motion planning," in *International Conference on Automated Planning and Scheduling* vol. 30, 2020, pp. 508–516.
- [10] O. Brock and L. E. Kavraki, "Decomposition-based motion planning: A framework for real-time motion planning in high-dimensional configuration spaces," *IEEE International Conference on Robotics and Automation* vol. 2. IEEE, 2001, pp. 1469–1474.
- [11] H.-J. Bungartz and M. Griebel, "Sparse grids," *Acta numerica* vol. 13, pp. 147–269, 2004.
- [12] I. Chatterjee, M. Likhachev, A. Khadke, and M. Veloso, "Speeding up search-based motion planning via conservative heuristics," *International Conference on Automated Planning and Scheduling* vol. 29, no. 1, 2019, pp. 674–679.
- [13] J. Corès, L. Jaillet, and T. Sifon, "Disassembly path planning for complex articulated objects," *IEEE Transactions on Robotics* vol. 24, no. 2, pp. 475–481, 2008.
- [14] J. C. Culberson and J. Schaeffer, "Pattern databases," *Computational Intelligence* vol. 14, no. 3, pp. 318–334, 1998.
- [15] C. De Boer, K. Hollig, and M. Sabin, "High accuracy geometric hermite interpolation," *University of Wisconsin-Madison Department of Computer Sciences, Tech. Rep.*, 1987.
- [16] J. Denny, R. Sandam, A. Bregger, and N. M. Amato, "Dynamic region-biased rapidly-exploring random trees," *Algorithmic Foundations of Robotics XII* Springer, 2020, pp. 640–655.
- [17] D. Driess, J.-S. Ha, and M. Toussaint, "Deep visual reasoning: Learning to predict action sequences for task and motion planning from an initial scene image," in *Robotics: Science and Systems* 2020.
- [18] W. Du, S.-K. Kim, O. Salzman, and M. Likhachev, "Escaping local minima in search-based planning using soft duplicate detection," in *IEEE International Conference on Intelligent Robots and Systems* IEEE, 2019, pp. 2365–2371.
- [19] W. Du, F. Islam, and M. Likhachev, "Multi-resolution A*," *arXiv preprint arXiv:2004.06684* 2020.
- [20] S. Edelkamp and S. Schroedter, *Heuristic search: theory and applications*. Elsevier, 2011.
- [21] S. Edelkamp, P. Kissmann, and Torralba, "Symbolic A* search with pattern databases and the merge-and-shrink abstraction," *European Conference on Artificial Intelligence* 2012, pp. 306–311.
- [22] P. Ferbach and J. Barraquand, "A method of progressive constraints for manipulation planning," *IEEE Transactions on Robotics* vol. 13, no. 4, pp. 473–485, 1997.
- [23] D. Ferguson, M. Likhachev, and A. Stentz, "A guide to heuristic-based path planning," in *International Conference on Automated Planning and Scheduling* 2005, pp. 9–18.
- [24] M. Fu, A. Kuntz, O. Salzman, and R. Alterovitz, "Toward asymptotically-optimal inspection planning via efficient near-optimal graph search," in *Robotics: Science and Systems* June 2019.
- [25] J. D. Gammell and M. P. Strub, "A survey of asymptotically optimal sampling-based motion planning methods," *arXiv preprint arXiv:2009.10484* 2020.
- [26] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Informed RRT*," *Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic*, *IEEE International Conference on Intelligent Robots and Systems* IEEE, 2014, pp. 2997–3004.
- [27] J. D. Gammell, T. D. Barfoot, and S. S. Srinivasa, "Batch informed trees (BIT*): Informed asymptotically optimal anytime search," *International Journal of Robotics Research* vol. 39, no. 5, pp. 543–567, 2020.
- [28] M. B. Giles, "Multilevel Monte Carlo methods," *Acta Numerica* vol. 24, pp. 259–328, 2015.
- [29] F. Glover and M. Laguna, "Tabu search," *Handbook of combinatorial optimization* Springer, 1998, pp. 2093–2229.
- [30] K. Gochev, A. Safonova, and M. Likhachev, "Planning with adaptive dimensionality for mobile manipulation," *IEEE International Conference on Robotics and Automation* 2012, pp. 2944–2951.
- [31] M. X. Grey, A. D. Ames, and C. K. Liu, "Footstep and motion planning in semi-unstructured environments using randomized possibility graphs," in *IEEE International Conference on Robotics and Automation* 2017, pp. 4747–4753.
- [32] J.-S. Ha, S.-S. Park, and H.-L. Choi, "Topology-guided path integral approach for stochastic optimal control in cluttered environments," *IEEE Robotics and Autonomous Systems* vol. 113, pp. 81–93, 2019.
- [33] N. Haghtalab, S. Mackenzie, A. D. Procaccia, O. Salzman, and S. S. Srinivasa, "The provable virtue of laziness in motion planning," *arXiv preprint arXiv:1710.04101* 2017.
- [34] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE transactions on Systems Science and Cybernetics* vol. 4, no. 2, pp. 100–107, 1968.
- [35] V. N. Hartmann, O. S. Oguz, D. Driess, M. Toussaint, and A. Menges, "Robust task and motion planning for long-horizon architectural construction planning," in *IEEE International Conference on Intelligent Robots and Systems* 2020.
- [36] K. Hauser, "Fast interpolation and time-optimization with contact," *International Journal of Robotics Research* vol. 33, no. 9, pp. 1231–1250, 2014.
- [37] B. Hou, S. Choudhury, G. Lee, A. Mandalika, and S. S. Srinivasa, "Posterior sampling for anytime motion planning on graphs with expensive-to-evaluate edges," *arXiv preprint arXiv:2002.11853* 2020.
- [38] D. Hsu, T. Jiang, J. Reif, and Z. Sun, "The bridge test for sampling narrow passages with probabilistic roadmap planners," *IEEE International Conference on Robotics and Automation* vol. 3. IEEE, 2003, pp. 4420–4426.
- [39] S. Hu and N. R. Sturtevant, "Direction-optimizing breadth-first search with external memory storage," in *International Joint Conference on Artificial Intelligence*, 2019, pp. 1258–1264.
- [40] D. Husemoller, *Fibre bundles* Springer, 1966, vol. 5.
- [41] B. Ichter, J. Harrison, and M. Pavone, "Learning sampling distributions for robot motion planning," in *IEEE International Conference on Robotics and Automation* IEEE, 2018, pp. 7087–7094.
- [42] L. Jaillet and J. M. Porta, "Path planning under kinematic constraints by rapidly exploring manifolds," *IEEE Transactions on Robotics* vol. 29, no. 1, pp. 105–117, 2012.
- [43] L. Jaillet, J. Corès, and T. Sifon, "Sampling-based path planning on configuration-space costmaps," *IEEE Transactions on Robotics* vol. 26, no. 4, pp. 635–646, 2010.
- [44] S. S. Joshi, S. Hutchinson, and P. Tsiotras, "Time-informed exploration for robot motion planning," *arXiv preprint arXiv:2004.05241* 2020.
- [45] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *International Journal of Robotics Research* vol. 30, no. 7, pp. 846–894, 2011.
- [46] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics* vol. 12, no. 4, pp. 566–580, 1996.
- [47] M. Kleinbort, O. Salzman, and D. Halperin, "Collision detection or nearest-neighbor search? On the computational bottleneck in sampling-based motion planning," in *Algorithmic Foundations of Robotics XII* Springer, 2020, pp. 624–639.
- [48] J. J. Kuffner and S. M. LaValle, "RRT-connect: An efficient approach to single-query path planning," *IEEE International Conference on Robotics and Automation* vol. 2, 2000, pp. 995–1001.
- [49] H. Kurniawati and D. Hsu, "Workspace-based connectivity oracle: An adaptive sampling strategy for PRM planning," *Algorithmic Foundations of Robotics VII* Springer, 2008, pp. 35–51.
- [50] B. Lacevic and D. Osmanovic, "Improved C-space exploration and path planning for robotic manipulators using distance information," in *IEEE International Conference on Robotics and Automation* 2020.
- [51] A. M. Ladd and L. E. Kavraki, "Fast tree-based exploration of

- state space for robots with dynamics,” in *Algorithmic Foundations of Robotics VI*. Springer, 2004.
- [52] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
- [53] J. M. Lee, *Introduction to Smooth Manifolds*. New York, NY: Springer New York, 2003.
- [54] J. Lee, O. Kwon, L. Zhang, and S.-e. Yoon, “SR-RRT: Selective retraction-based RRT planner,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 2543–2550.
- [55] Z. Littlefield and K. E. Bekris, “Efficient and asymptotically optimal kinodynamic motion planning via dominance-informed regions,” in *IEEE International Conference on Intelligent Robots and Systems*. IEEE, 2018, pp. 1–9.
- [56] T. Lozano-Perez, “A simple motion-planning algorithm for general robot manipulators,” *IEEE Journal on Robotics and Automation*, vol. 3, no. 3, pp. 224–238, 1987.
- [57] R. Luna, M. Moll, J. Badger, and L. E. Kavraki, “A scalable motion planner for high-dimensional kinematic systems,” *International Journal of Robotics Research*, vol. 39, no. 4, pp. 361–388, 2020.
- [58] Y. Luo, H. Bai, D. Hsu, and W. S. Lee, “Importance sampling for online planning under uncertainty,” *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 162–181, 2019.
- [59] J. Mainprice, N. Ratliff, M. Toussaint, and S. Schaal, “An interior point method solving motion planning problems with narrow passages,” in *IEEE International Conference on Robot and Human Interactive Communication*. IEEE, 2020, pp. 547–552.
- [60] M. Manak, “Voronoi-based detection of pockets in proteins defined by large and small probes,” *Journal of computational chemistry*, vol. 40, no. 19, pp. 1758–1771, 2019.
- [61] A. Mandalika, S. Choudhury, O. Salzman, and S. Srinivasa, “Generalized lazy search for robot motion planning: Interleaving search and edge evaluation via event-based toggles,” in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 29, no. 1, 2019, pp. 745–753.
- [62] Z. McCarthy, T. Bretl, and S. Hutchinson, “Proving path non-existence using sampling and alpha shapes,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 2563–2569.
- [63] M. Moll, I. A. Şucan, and L. E. Kavraki, “Benchmarking motion planning algorithms: An extensible infrastructure for analysis and visualization,” *IEEE Robotics and Automation Magazine*, vol. 22, no. 3, pp. 96–102, September 2015.
- [64] A. Orthey and M. Toussaint, “Rapidly-exploring quotient-space trees: Motion planning using sequential simplifications,” *International Symposium of Robotics Research*, 2019.
- [65] —, “Visualizing local minima in multi-robot motion planning using multilevel morse theory,” in *Workshop on the Algorithmic Foundations of Robotics*, 2020.
- [66] A. Orthey, A. Escande, and E. Yoshida, “Quotient-space motion planning,” in *IEEE International Conference on Intelligent Robots and Systems*. IEEE, 2018, pp. 8089–8096.
- [67] A. Orthey, O. Roussel, O. Stasse, and M. Taïx, “Motion planning in irreducible path spaces,” *IEEE Robotics and Autonomous Systems*, vol. 109, pp. 97–108, 2018.
- [68] A. Orthey, S. Akbar, and M. Toussaint, “Multilevel motion planning: A fiber bundle formulation,” 2020, arXiv:2007.09435 [cs.RO].
- [69] T. Osa, “Multimodal trajectory optimization for motion planning,” *International Journal of Robotics Research*, vol. 39, no. 8, pp. 983–1001, 2020.
- [70] L. Palmieri, S. Koenig, and K. O. Arras, “RRT-based nonholonomic motion planning using any-angle path biasing,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2016, pp. 2775–2781.
- [71] J. Pearl, “Heuristics: Intelligent search strategies for computer problem solving,” *Addison Wesley*, 1984.
- [72] S. M. Persson and I. Sharf, “Sampling-based A* algorithm for robot path-planning,” *The International Journal of Robotics Research*, vol. 33, no. 13, pp. 1683–1708, 2014.
- [73] E. Plaku, L. E. Kavraki, and M. Y. Vardi, “Discrete search leading continuous exploration for kinodynamic motion planning,” in *Robotics: Science and Systems*, 2007, pp. 326–333.
- [74] F. T. Pokorny, M. Hawasly, and S. Ramamoorthy, “Topological trajectory classification with filtrations of simplicial complexes and persistent homology,” *International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 204–223, 2016.
- [75] A. H. Qureshi, Y. Miao, A. Simeonov, and M. C. Yip, “Motion planning networks: Bridging the gap between learning-based and classical motion planners,” *IEEE Transactions on Robotics*, 2020.
- [76] W. Reid, R. Fitch, A. H. Göktoğan, and S. Sukkarieh, “Sampling-based hierarchical motion planning for a reconfigurable wheel-on-leg planetary analogue exploration rover,” *Journal of Field Robotics*, 2019.
- [77] W. Reid, R. Fitch, A. H. Göktoğan, and S. Sukkarieh, “Motion planning for reconfigurable mobile robots using hierarchical fast marching trees,” in *Algorithmic Foundations of Robotics XII*. Springer, 2020, pp. 656–671.
- [78] M. Rickert, A. Sieverling, and O. Brock, “Balancing exploration and exploitation in sampling-based motion planning,” *IEEE Transactions on Robotics*, vol. 30, no. 6, pp. 1305–1317, 2014.
- [79] J. Röwekämper, G. Tipaldi, and W. Burgard, “Learning to guide random tree planners in high dimensional spaces,” in *IEEE International Conference on Intelligent Robots and Systems*, 2013.
- [80] M. Saha, J.-C. Latombe, Y.-C. Chang, and F. Prinz, “Finding narrow passages with probabilistic roadmap: The small-step retraction method,” *Autonomous Robots*, vol. 19, no. 3, pp. 301–319, 2005.
- [81] B. Sakkak, L. Bascetta, G. Ferretti, and M. Prandini, “An admissible heuristic to improve convergence in kinodynamic planners using motion primitives,” *IEEE Control Systems Letters*, vol. 4, no. 1, pp. 175–180, 2019.
- [82] —, “Sampling-based optimal kinodynamic planning with motion primitives,” *Autonomous Robots*, vol. 43, no. 7, pp. 1715–1732, 2019.
- [83] O. Salzman, “Sampling-based robot motion planning,” *Communications of the ACM*, vol. 62, no. 10, pp. 54–63, 2019.
- [84] O. Salzman and D. Halperin, “Asymptotically near-optimal RRT for fast, high-quality motion planning,” *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 473–483, 2016.
- [85] O. Salzman, M. Hemmer, and D. Halperin, “On the power of manifold samples in exploring configuration spaces and the dimensionality of narrow passages,” in *Algorithmic Foundations of Robotics X*, E. Frazzoli, T. Lozano-Perez, N. Roy, and D. Rus, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 313–329.
- [86] A. Schweikard and F. Schwarzler, “Detecting geometric infeasibility,” *Artificial Intelligence*, vol. 105, no. 1-2, pp. 139–159, 1998.
- [87] S. Sekhavat, P. Svestka, J.-P. Laumond, and M. H. Overmars, “Multilevel path planning for nonholonomic robots using semiholonomic subsystems,” *International Journal of Robotics Research*, vol. 17, no. 8, pp. 840–857, 1998.
- [88] A. Sintov, S. Macenski, A. Borum, and T. Bretl, “Motion planning for dual-arm manipulation of elastic rods,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6065–6072, 2020.
- [89] N. E. Steenrod, *The topology of fibre bundles*. Princeton Univ. Press, 1951.
- [90] M. P. Strub and J. D. Gammell, “Adaptively informed trees (AIT*): Fast asymptotically optimal path planning through adaptive heuristics,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2020.
- [91] B. M. K. Styler and R. Simmons, “Plan-time multi-model switching for motion planning,” in *International Conference on Automated Planning and Scheduling*, 2017.
- [92] I. A. Şucan and L. E. Kavraki, “A sampling-based tree planner for systems with complex dynamics,” *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 116–131, 2011.
- [93] I. A. Şucan, M. Moll, and L. Kavraki, “The open motion planning library,” *IEEE Robotics and Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [94] J. Szkandera, I. Kolingerová, and M. Maňák, “Narrow passage problem solution for motion planning,” in *Computational Science – ICCS 2020*, V. V. Krzhizhanovskaya, G. Závodszy, M. H. Lees, J. J. Dongarra, P. M. A. Sloot, S. Brissos, and J. Teixeira, Eds. Springer International Publishing, 2020, pp. 459–470.
- [95] S. Tonneau, A. D. Prete, J. Pettré, C. Park, D. Manocha, and N. Mansard, “An Efficient Acyclic Contact Planner for Multiplied Robots,” *IEEE Transactions on Robotics*, vol. 34, no. 3, pp. 586–601, June 2018.
- [96] D. Uwacu, R. Rex, B. Wang, S. Thomas, and N. M. Amato, “Annotated-skeleton biased motion planning for faster relevant region discovery,” *arXiv preprint arXiv:2003.02176*, 2020.
- [97] J. P. Van den Berg and M. H. Overmars, “Using workspace information as a guide to non-uniform sampling in probabilistic roadmap planners,” *International Journal of Robotics Research*, vol. 24, no. 12, pp. 1055–1071, 2005.
- [98] A. Varava, J. F. Carvalho, F. T. Pokorny, and D. Kragic, “Free space of rigid objects: Caging, path non-existence, and narrow passage detection,” in *International Journal of Robotics Research*. Springer, 2020.
- [99] S. Vats, V. Narayanan, and M. Likhachev, “Learning to avoid local minima in planning for static environments,” in *International Conference*

on *Automated Planning and Scheduling*, 2017.

- [100] E. Vidal, M. Moll, N. Palomeras, J. D. Hernández, M. Carreras, and L. E. Kavraki, "Online multilayered motion planning with dynamic constraints for autonomous underwater vehicles," in *IEEE International Conference on Robotics and Automation*. IEEE, 2019, pp. 8936–8942.
- [101] V. Vonásek and R. Pěnička, "Sampling-based motion planning of 3d solid objects guided by multiple approximate solutions," in *IEEE International Conference on Intelligent Robots and Systems*. IEEE, 2019, pp. 1480–1487.
- [102] Y. Yang and O. Brock, "Efficient motion planning based on disassembly," in *Robotics: Science and Systems*, Cambridge, USA, June 2005.
- [103] A. Yershova and S. M. LaValle, "Motion planning for highly constrained spaces," *Robot Motion and Control*, vol. 396, p. 297, 2009.
- [104] E. Yoshida, "Humanoid motion planning using multi-level dof exploitation based on randomized method," in *IEEE International Conference on Intelligent Robots and Systems*. IEEE, 2005, pp. 3378–3383.
- [105] L. Zhang and D. Manocha, "An efficient retraction-based RRT planner," in *IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 3743–3750.
- [106] L. Zhang, Y. J. Kim, and D. Manocha, "A simple path non-existence algorithm using c-obstacle query," in *Algorithmic Foundation of Robotics VII*. Springer, 2008, pp. 269–284.
- [107] L. Zhang, J. Pan, and D. Manocha, "Motion planning of human-like robots using constrained coordination," in *IEEE International Conference on Humanoid Robots*, 2009, pp. 188–195.