

Floating-base manipulation on zero-perturbation manifolds

Brian A. Bittner, Jason Reid, Kevin C. Wolfe

Abstract—To achieve high-dexterity motion planning on floating-base systems, the base dynamics induced by arm motions must be treated carefully. In general, it is a significant challenge to establish a fixed-base frame during tasking due to forces and torques on the base that arise directly from arm motions (e.g. arm drag in low Reynolds environments and arm momentum in high Reynolds environments). While thrusters can in theory be used to regulate the vehicle pose, it is often insufficient to establish a stable pose for precise tasking, whether the cause be due to underactuation, modeling inaccuracy, suboptimal control parameters, or insufficient power. We propose a solution that asks the thrusters to do less high bandwidth perturbation correction by planning arm motions that induce zero perturbation on the base. We are able to cast our motion planner as a nonholonomic rapidly-exploring random tree (RRT) by representing the floating-base dynamics as pfaffian constraints on joint velocity. These constraints guide the manipulators to move on zero-perturbation manifolds (which inhabit a subspace of the tangent space of the internal configuration space). To invoke this representation (termed a *perturbation map*) we assume the body velocity (perturbation) of the base to be a joint-defined linear mapping of joint velocity and describe situations where this assumption is realistic (including underwater, aerial, and orbital environments). The core insight of this work is that when perturbation of the floating-base has affine structure with respect to joint velocity, it provides the system a class of kinematic reduction that permits the use of sample-based motion planners (specifically a nonholonomic RRT). We show that this allows rapid, exploration-gearred motion planning for high degree of freedom systems in obstacle rich environments, even on floating-base systems with nontrivial dynamics.

I. INTRODUCTION

Floating-base manipulation platforms offer the potential to perform critical tasking in underwater [1], [2], [3], [4], aerial [5], [6], [7], and orbital [8], [9] environments. While manipulation has succeeded in high complexity environments [10], [11], [12], floating-base systems have yet to achieve the same ability to precisely manipulate objects in similar environments. The inability to physically anchor any link of the kinematic chain can induce complex dynamics on the ungrounded robot. These dynamics present differential constraints that are often challenging for the motion planner to accommodate when pursuing precision tasking. If the floating-base manipulator assumes a kinematic operating environment, perturbations on the base can cause imprecision in movement that is unpredictable and unsafe, especially in obstacle rich environments. By considering perturbation dynamics, one can pursue more precise positioning of the end effector(s).

Typically in floating-base manipulation, the dynamics are represented as a second-order mechanical system [13], [14],

Johns Hopkins University Applied Physics Laboratory, Laurel, MD, USA.
brian.bittner@jhuapl.edu

[15]. Such a representation stipulates that given a path in the joint configuration, the rate along that path has a variational effect on the evolution of the robot in its workspace. These timing-based concerns induce substantial complexity on the motion planning problem; these planning problems are classified as kinodynamic [16]. Generally, kinodynamic path planning problems are formulated as trajectory optimization problems. These approaches typically approximate the Hamilton-Jacobi-Bellman equation, which recursively formulates the optimal control law for a physical system. Such exploitative methods excel at finding local refinements to seeded trajectories that leverage the agile capabilities of the system [17], [18], [19], as long as the underlying model maintains high fidelity to the realized system and environment. A separate trajectory optimization approach involves policy gradient methods [20], [21], which are performant at rapidly finding collision free paths, but again are typically used to find local refinements to approaches (due to the exploitative nature of gradient descent).¹ Additionally, this class of solutions suffer from poor sampling scalability with the dimension of the configuration – termed the curse of dimensionality. This notion of scalability is encountered in many practical scenarios, such as adding joints to improve the reachable configurations of a robot in key task domains (e.g. pick and place in a cluttered environment). Exploitative trajectory optimization methods thus tend to be insufficient for accomplishing many critical tasks due to their locality and sample inefficiency. However, as is typically the case in mobility architectures, a higher level planner might successfully hand off a performant, coarse plan intended for final refinement by the trajectory optimizer.

The main contribution of this work is the justification, design, and implementation of a framework that decouples perturbations during floating-base manipulation within a sampling-based motion planner (specifically an RRT [22]). We show that a kinematic reduction of the floating-base dynamics can be applied to systems through mechanical design, environment selection, and thruster-aided pose control. This reduces the motion planning problem from kinodynamic to kinematic, making sampling-based motion planning algorithms accessible. Empirically, such methods extend to high dimensional problems with practical sample efficiency [23]. Sample efficiency for large dimensions is critical, since as described earlier, higher numbers of joints will be required for complex tasking in cluttered environments. Gaining this property (sample efficiency for platforms with significant

¹While stochastic formulations are designed to specifically avoid local minima, we emphasize that the nature of a policy gradient search is far more localized than sampling-based path planning methods.




Fig. 1. A thirteen-link Purcell swimmer with knowledge of its kinematic chain via (7) (left) is unable to track (red proximal to end effectors) the shapes (blue) sequentially attempted right to left due to perturbation induced on the base (red proximal to base) during tracking. The thirteen-link Purcell swimmer with knowledge of its dynamics as well as its kinematic chain written as pfaffian constraints via (8) (right) is able to succeed given the same task. This tracking result is explained in detail in Section IV. The performant tracker leverages a surface in the joint velocity space which we term a zero-perturbation manifold (ZPM). This surface is the core element of a new planning framework covered in Section V. Accompanying video in Supplementary Materials.

internal complexity – required for many critical application domains) while respecting the dynamics is the core capability that we pursue through this work. Using the method detailed in Section V, we structurally expect (A) higher precision end effector tracking and placement than planners with no dynamics knowledge as well as (B) rapid computability of valid plans in many-joint, obstacle rich scenarios. We show in obstacle rich scenarios that for emplacement and extraction tasks, the proposed method can shift induced perturbation to a trivial moment during the plan, allowing the ZPM to facilitate precise engagement with a target configuration.

The class of kinematic systems we assume stipulates that the body velocity is a joint-defined linear mapping of joint velocity. These pfaffian constraints can be used at any configuration to identify a subspace of joint velocities where instantaneous joint motion yields zero perturbation (a sufficient condition being that the dimension of the internal configuration is greater or equal to the dimension of the position space). These pfaffian constraints are termed a *perturbation map*, and this map provides a nonheuristic manifold that can be used to guide the system to rapidly find collision-free plans. Next, we provide background on kinematic representations of floating-base systems in application-oriented physical scenarios, followed by discussion of the key properties of the ZPM in Section III along with examples of end-effector tracking and planning solutions that leverage this structure.

II. BACKGROUND: KINEMATIC REDUCTIONS OF LAGRANGIAN SYSTEMS

In this work, we seek to compute plans for floating-base systems when their motion model can be written as a first-order (kinematic) system expressed as pfaffian constraints connecting velocity on the shape space (joint space) B to velocity on the position space G . These constraints are expressed as

$$\dot{x}_b = P(\theta)\dot{\theta} \quad (1)$$

yielding a linear mapping, P (as a function of n joints $\theta \in B$) from joint velocity $\dot{\theta} \in TB$ to body velocity $\dot{x}_b \in TG$. This structure will be shown to be an exact representation of the physics in conventional orbital environments and scenarios governed by Rayleigh dissipation. We'll show that violations of this physical structure can be related to intuitive phenomena that can be designed out of the problem through mechanical design, environment selection, and control.

A. Obtaining Environmental Symmetry

Gravity directly violates environmental symmetry, inducing position and orientation dependent forces on the floating-base. This will, for instance, cause the robot to move when the joints are still, clearly violating our assumed model at $\dot{\theta} = 0$. Gravity is however negligible for orbital robots as well as neutrally buoyant underwater systems². In less convenient scenarios, one might supply a model-based control law to compensate for gravity with thrusters, through which one could regain environmental symmetry under the dynamics resulting from the thruster-governed control law.

When we can assert symmetry of the dynamics over G for Lagrangian systems, there exists a powerful reframing of the dynamics that provides three distinct equations written for the generalized momentum, body velocity, and joint acceleration of the system. This result from the field of geometric mechanics is known as the reduced Lagrangian:

$$\dot{x}_b = P(\theta)\dot{\theta} + I(\theta)p \quad (2)$$

$$\dot{p} = f_1(\theta, \dot{\theta}) \quad (3)$$

$$\ddot{\theta} = f_2(\theta, \dot{\theta}, p) \quad (4)$$

for which p provides the generalized momentum and f_1 and f_2 are nonlinear functions whose specific structure is available in [24].

²Some platforms such as RE2s sapien class system are designed to be neutrally buoyant to avoid dynamical complexity.

B. Mitigating Interaction Between Inertia and Friction

When the generalized momentum starts at zero (as is the imagined case for many floating-base tasks), the dynamics of the body are governed exclusively by conservation of zero net angular and linear momentum. The resulting dynamics take the form of (1) and is referred to as the *mechanical connection* [25]. A similar phenomena occurs when the motion model becomes dominated by Rayleigh dissipation. Here, friction-governed forces applied along the body yield a viscous interaction with the fluid, again identical to the form in (1) and termed the *viscous connection* [26].

Coupling of forces driven by inertia and friction result in second-order behaviors. Friction governed systems can be designed to avoid chattering or high speed maneuvering. They can also increase the damping that governs their interaction with the environment, which could be through mass reduction, substrate selection, or modifications to thruster control laws (e.g. tuning the derivative gain of a PID controller). By keeping $|p| \ll 1$, dissipative systems can ensure that $|P(\theta)\dot{\theta}| \gg |I(\theta)p|$, such that the system remains in a principally kinematic state. Zero momentum conservation systems (e.g. orbital manipulators) can be further designed to avoid momentum shifting events such as contact or thrust generation. Since both of these events add or reduce momentum from the system, they will drive the $|f_1| > 0$ and excite the latter term, $I(\theta)p$, in (2) (making the system no longer principally kinematic).

C. Approximating Nonviscous Friction

In general, friction can certainly act in non-viscous ways, violating the structure of (1). One can think of the viability of the Rayleigh dissipation model as the viability of a first-order Taylor series approximation of the more general friction-governed motion model $\dot{x}_b|\theta = f(\dot{\theta})$. However, it has been observed in practice that the assumption of a Rayleigh dissipation model provides a strong fit across a surprisingly broad class of robots [27], [28].

III. ZERO PERTURBATION MANIFOLDS

We now shift focus to explaining how the representations motivated in Section II populate a manifold that can be exploited for motion planning that induces zero perturbation on the floating-base. The linear mapping $P(\theta)$ has a null space computable at all shapes θ

$$\text{null}(P(\theta)) \subset TB \quad (5)$$

which inhabits a subspace of the shape velocity space TB . This null space provides a basis in the joint velocity space for exploration from a shape that will infinitesimally invoke zero perturbation on the base. This null space changes as a function of θ , and so during a plan, the allowable motions available to the joints are subject to change. When the null space changes continuously with respect to shape, the null space populates a connected submanifold of joint velocity space. We know this since the notion of connectedness is preserved under continuous mappings. However, the null space is provably not continuous in this sense at singular P where

the rank of the nullspace increases over that of non-singular P . Thus, the zero perturbation manifold (ZPM), $\text{null}(P)$, has a basis of size $\dim(B) - \dim(G)$ and continuously changes over its connected regions. It is however a larger manifold that is not continuous, e.g. at locations of singular P .

Continuity of the null space is however important for generating C^1 smooth arm motions. If a null space basis vector disappears while moving a joint in that direction, it will appear as a nonsmooth path in joint space (if persisting along the zero perturbation manifold). This manifold is not unlike leveraging the null space of components of the manipulator Jacobian to stabilize elevation, roll and pitch when moving a weight along a table, as is done in work that in-part inspired this approach [29], [30].

In classical manipulation scenarios, additional joints provide exploitable redundancy. Likewise, each additional joint increases the dimension of the zero perturbation manifold. Despite this, the probability of the null space directly aligning with two randomly selected shapes is approximately zero:

$$\forall \delta \in B : \mathbf{p}(\text{null}(P(\theta))^T \text{null}(P(\theta + \delta)) = 1) \approx 0. \quad (6)$$

where \mathbf{p} provides our single-case use of notation for probability. While it is easy to grow a tree along the null space from a known joint location, it can be hard to grow the tree along the null space toward a known joint location. One might think of an Ackermann steering vehicle (a system whose maneuvers are governed by pfaffian constraints) that is placed arbitrarily close to a final location, then asked to arrive at that location directly, without knowledge of Lie brackets or "parallel parking" maneuvers.

IV. LOCAL END EFFECTOR TRACKING ALONG THE ZERO PERTURBATION MANIFOLD

Here we leverage the exploitable redundancy property discussed in Section III to show that a many jointed, bimanual swimming robot can complete a welding task with a closed form solution. We compare two platforms. A baseline platform has access to its kinematic chain, so induced disturbances on the body cannot be accounted for:

$$\dot{x}_{ee} = J(\theta)\dot{\theta}. \quad (7)$$

It takes a pseudoinverse of the Jacobian to compute $\dot{\theta}$ given a desired velocity for the end effector \dot{x}_{ee} . The second platform has access to both its kinematic chain and the model of the viscous fluidic model that it inhabits:

$$\begin{bmatrix} \dot{x}_{ee} \\ \dot{x}_b \end{bmatrix} = \begin{bmatrix} J(\theta) \\ P(\theta) \end{bmatrix} \dot{\theta}. \quad (8)$$

where \dot{x}_b can be set to zero.

This toy system is a simple adaptation of the Purcell swimming model [31], [32], where we magnify the ratio of across-to-along body forces by a factor of 5, mimicking a more fin-like relationship with the water to excite larger perturbations. In the trial, \dot{x}_{ee} is set by the residual between the location of the hand (with respect to the base) and the desired hand location (with respect to the base). The desired hand location is a function of time that traces each

shape (from right to left)). The closest hand is driven to track the desired hand location, the other being free to use its neighboring links as a tail-like appendage if necessary to stabilize the base. Fig. 1 shows that the thirteen-link swimmer was capable of exactly tracing the desired shapes with a closed form solution when leveraging the pfaffian form of the dynamics. Meanwhile, the platform embedded with only kinematic-chain knowledge both drifted and was unable to draw shapes resembling those desired.

V. RRT ON ZERO PERTURBATION MANIFOLDS

In Sections II and III we overviewed a key class of floating-based dynamics and its relevance for distinction of a zero-perturbation manifold that occupies TB . Here, we will show how such a manifold can be used effectively within sampling-based motion planners for rapid exploration based planning on high degree of freedom floating-base robots in obstacle rich settings. Equation (6) shows us that connecting two points in the shape space can be challenging. This suggests that objects like probabilistic roadmaps (PRMs) will be structurally challenging to generate, since the interconnectedness of nodes requires direct confrontation of this issue. Turning our attention to single-query planning approaches, RRT offers the ability to grow quickly from a known location. As we have established, knowing the final location does not make it easy to find a plan that arrives exactly at the final location. One might think of a car placed at a close distance but incorrect orientation to a desired pose. While Lie brackets or primitive-based motion planning may offer a solution for the car, it is not obvious that such approaches are computationally tractable in the floating-base manipulation setting.

We observe three options, which include growing the RRT from an initial location to a desired location, growing a bi-RRT from both locations towards each other, and growing from the desired location toward the initial location. We accept the third option as suboptimal but highly practical since it affords precision at the desired location. An RRT will be designed to rapidly grow from a location of interest toward the robots initial configuration, all the while restricting itself to the zero perturbation manifold. The optimal path selected from this tree is the core contribution of this work. The robot will start by connecting itself from initial configuration to the ZPM. Once the joints reach the ZPM, the arms stop while the thrusters mitigate any induced initial perturbation. Post regulation, the floating-base system is afforded a zero perturbation path, through clutter, to the object of interest.

Here we will outline our approach for floating-base manipulation along zero perturbation manifolds. Algorithm 1 is provided a start and goal internal configuration for the robots (θ_s, θ_g) . As described above, we will start at the goal location, where there is a higher need for precision (such as in emplacement and extraction tasking). We initialize a tree T with this node and query a sample function at each iteration. This will provide a random internal configuration or the start configuration (with probability 0.5 in this work). We then attempt to grow the tree toward this

Algorithm 1 $RRTZPM(\theta_s, \theta_g)$

```

1:  $T.init(\theta_g)$ 
2: while  $TimeRemaining()$  do
3:    $\theta_{sample} \leftarrow sample(i, \theta_s)$ 
4:    $\theta_{nearest} \leftarrow nearest(T, \theta_{sample})$ 
5:    $\theta_{extend} \leftarrow ExtendAlongZPM(T, \theta_{nearest}, \theta_{sample})$ 
6:   if  $\|\theta_{extend} - \theta_s\| < \alpha$  then
7:     return  $Smooth(ExtractPath(T, \theta_s))$ 
8:   end if
9:   return null
10: end while

```

Algorithm 2 $ExtendAlongZPM(T, \theta_{near}, \theta_{sample})$

```

1:  $\theta_{extend} \leftarrow \theta_{near}$ ;  $spacing \leftarrow 0$ 
2: while  $Timeout()$  or  $Collision(\theta_{extend})$  do
3:    $\delta\theta \leftarrow \theta_{extend} - \theta_{sample}$ 
4:    $u, s, v \leftarrow svd(P(\theta_{extend}))$ 
5:    $P_{null} \leftarrow v[dim(G) : ]$ 
6:   if  $P_{null} \perp \delta\theta$  then
7:     return  $\theta_{extend}$ 
8:   else
9:      $\theta_{extend} \leftarrow \theta_{extend} + unit(\delta\theta P_{null}^T P_{null}) dt$ 
10:     $spacing \leftarrow spacing + dt$ 
11:    if  $spacing > L$  then
12:       $T.addVertex(\theta_{near}, \theta_{extend})$ 
13:       $\theta_{near} \leftarrow \theta_{extend}$ ;  $spacing \leftarrow 0$ 
14:    end if
15:  end if
16: end while
17: return  $\theta_{extend}$ 

```

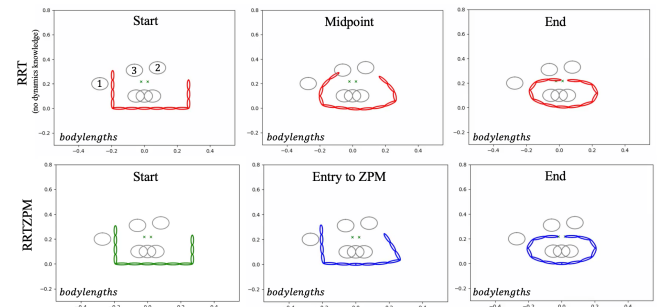


Fig. 2. Using a classical RRT (top), the thirteen-link Purcell swimmer is not able to accurately obtain a desired pose during emplacement (taken from 1 of 30 trials detailed in Fig. 3). We show snapshots of the swimmer at the beginning (first column), midpoint (second column), and end (final column), approaching hand locations (pairs of green x's) while avoiding obstacles (grey circles). Using RRTZPM (bottom), the swimmer is able to reach the desired endpoint. We show snapshots of the behavior, with columns carrying the same timestamps with the exception of the middle, which marks the transition point to the ZPM. Accompanying video in Supplementary Materials.




Fig. 3. This Fig. show the displacement of the Purcell swimmer in distance (left Δd in bodylengths) and orientation (right $\Delta\psi$ in radians) over 30 trials for emplacement with one obstacle (top row, obstacle numbers identified in Fig. 2), two obstacles (middle row), and three obstacles (bottom row). The mean (solid) and bounds (transparent band) of performance are plotted for the RRTZPM (blue, Algorithm 1) motivated in this work as well as the classical RRT (red) which does not incorporate any dynamics knowledge.

sampled configuration, while adhering to the zero perturbation manifold. The computation details of this process are available in Algorithm 2, which involves computation of the null space via singular value decomposition (SVD) and projection of the error between goal and current point along the computed subspace of the internal configuration that carries zero eigenvalues. If we avoid singular configurations (as is assumed in this algorithm), then the null space is consistently of $\text{rank}(\text{null}(P)) = \text{dim}(B) - \text{dim}(G)$ and can be indexed directly as the corresponding bottom rows of V . We can project the growth direction onto this manifold (see line 9 of Algorithm 2 where *unit* produces a unit vector), and stop extending in the event of a timeout, collision, or perpendicular orientation of the null space and remaining error. During extension, we grow the tree by adding a node after such a length has been extended from a prior node in the tree. Once the search has timed out or achieved sufficient proximity to the goal configuration, the algorithm returns a smoothed path. In this work, smoothing involves attempting to find a shortcut between the start and end point of the path by direct projection onto the ZPM, then recursively repeating the procedure on two randomly split segments of the path (until those segments reduce to two points). This smoothing process is repeated three times in this work.

VI. RESULTS

Here we will ask the thirteen-link Purcell swimmer to plan in constrained environments while executing an emplacement task. This task will require that both hands reach a particular location in the workspace. Through expectation that the base does not move, we stipulate a shape that will achieve the desired end effector locations and provide this desired shape as a goal to the planner. As mentioned before, the planner will

grow from the goal towards the starting location. We identify that there are a variety of viable starting configurations from which to approach the emplacement task, and encode this within the planner. When trying to connect to starting shapes θ_s , we will provide a randomly selected, generalized U -shape for the swimmer.³ Any such shape can translate in the coordinate perpendicular to the base link to exit or enter the starting configuration, making it a viable entry point to start the emplacement plan from. Thus, we claim to be solving a broader planning problem (starting from a disengaged, distal position with any U -shaped initial joint configuration) even though we only explicitly solve the non-trivial component of the planning problem (once engaged in an obstacle rich environment).

We start with one obstacle and provide the RRTZPM 30 separate trials to find connections to U -shapes along the zero perturbation manifolds. We then ask the RRT to connect the same two configurations. In each trial we allow the planner to replan 4 times after 10s intervals if a collision free path is not found. Paths are termed collision free if they would not cause a collision with a static base. In every trial both planners succeed during the first 10s interval. Fig. 3 shows the displacement as a function of a fraction of path length, spanning the start and finish of the attempt. It is clear that the proposed method accrues a brief initial displacement to access the zero perturbation manifold (distinguishing the effect of the properties following from (6)). However, once on the ZPM, the platform is capable of achieving the desired configuration with zero displacement. Meanwhile, the classical RRT accrues error, critically during the final steps of engagement of the desired internal configuration. This classical approach persistently accrues error as it approaches the final location, exactly when the proposed method accrues zero error by structure.

We see this relationship persist as a second and third obstacle are added. Notably, during trials with 3 obstacles, the RRT unsuccessfully attempts to find a collision free path in 20 of 30 trials. For all trials, the proposed method experiences an initial perturbation of the base followed by an unperturbed approach to the final configuration. We observe this difference in performance in Fig. 2, where we can see the drift over the duration of the trial for the classical RRT on the top row, as well as the lack of drift for RRTZPM.

In obstacle-free environments we compare this approach to an iterative time varying linear quadratic regulator (TVLQR [33]). This provides the opportunity to observe the exploration-driven solutions available via RRTZPM with respect to this exploitation-driven approximation of the Hamilton-Jacobi-Bellman equation. The iterative TVLQR is seeded with a direct route from initial to final shape for the linear system, $\begin{bmatrix} \dot{x}_b \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} P(\theta) \\ I \end{bmatrix} \dot{\theta}$, where we take the shape velocity $\dot{\theta}$ as the control input. We discretize each trajectory at 50 samples. For $n = \text{dim}(B)$ and $m = \text{dim}(G)$ take $Q =$

³A U includes a negative quarter rotation of a joint left of the base and a positive quarter rotation of a joint right of the base.

VII. DISCUSSION AND CONCLUSION

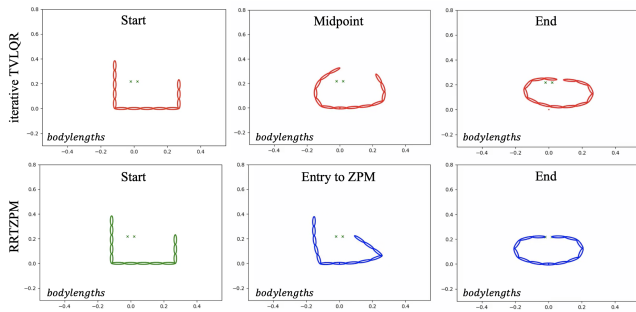


Fig. 4. Using iterative TVLQR (top), the thirteen-link Purcell swimmer is not able to accurately obtain a desired pose during emplacement (taken from 1 of 30 trials in Fig. 5). We show snapshots of the swimmer at the beginning (first column), midpoint (second column), and end (final column), approaching hand locations (pairs of green x's) while avoiding obstacles (grey circles). Using RRTZPM (bottom), the swimmer is able to reach the desired endpoint. We show snapshots of the behavior, with columns carrying the same timestamps with the exception of the middle, which marks the transition point to the ZPM. Accompanying video in Supplementary Materials.




Fig. 5. This Fig. shows the displacement of the Purcell swimmer in distance (left Δd in bodylengths) and orientation (right $\Delta\psi$ in radians) over 30 trials with no obstacles. The mean (solid) and bounds (transparent band) of performance are plotted for RRTZPM (blue, Algorithm 1) motivated in this work as well as the iterative TVLQR which is subject to the locality of exploitation-driven solutions.

$\begin{bmatrix} I^{m \times m} & 0 \\ 0 & 0 \end{bmatrix}$, $Q_f = \begin{bmatrix} I^{m \times m} & 0 \\ 0 & I^{n \times n} \end{bmatrix}$, and $R = Ie^{-2}$ where Q and R are cost matrices and Q_f is the final element of Q . By iteratively supplying the output of a TVLQR solution to itself, we provide the opportunity to further refine the solution with the context of the dynamics centered at the most recent solution. At the end of 20 iterations we have a trajectory that has been locally refined to find the desired joint location while avoiding perturbations. We ask the swimmer to achieve the previously identified emplacement shape configuration from randomly selected U -shapes. Fig. 4 shows that over 30 trials we observe the RRTZPM is capable of accessing the ZPM, although it accessed later in the maneuver than in Fig 3 when it had the choice of starting U -shape locations. The iterative TVLQR critically accrues error as it arrives at the emplacement location, repeating the precision mitigating behavior seen in the RRT in Fig. 3. We can see impact of this error at the emplacement location in Fig. 5, where the error accrued through iterative TVLQR renders the platform unable to complete the emplacement task.

The active use of the other arm to decouple the generated drag force in Fig. 1 reminded us of the use of tails in many animals [34], [35]. This may provide inspiration for design of systems in similar dynamic environments. Overall, the end effector tracking results of Fig. 1 are not entirely surprising (prior work exists in null space projection for swimming robots [4]). The example reinforces what is known in manipulation, that redundancy offers a greater variety of options to achieve the desired result. Observing this property in a floating-base end effector tracking provided reason to pursue embedding this capability in a global planner.

The emplacement performance of the RRTZPM method in Fig. 3 and Fig. 2 shows us that it was possible to embed this property in a high degree of freedom platform with three obstacles. It is clear from comparison to the RRT planner (that neglects dynamics) that the ZPM is providing significant value. The practical relevance of this result is made clear in Fig. 2, where the RRT platform can be seen colliding and drifting from the desired location. Such precision is unacceptable for even coarse emplacement tasking.

This relationship persists when we relate the motivated method to iterative TVLQR in an obstacle-free environment. The exploitation driven algorithm can locally refine the initial trajectory, but is unable to find a more global solution that permits acceptable precision for the emplacement task. This supports the perceived value of our method which can respect dynamics while finding inventive solutions rapidly.

The measured perturbation rejection on these simulated systems can be impacted largely by the time-step of the integration. For real systems, feasibility of precisely tracking shape evolution along the ZPM could impact the quality of the overall plan, leading to interest in tools that can estimate or improve the accuracy of motion models in fully kinematic [36] and perturbed [37] settings.

The results suggest that the proposed methods provide a novel class of algorithms that can address the weaknesses of classical sampling-based motion planners and optimization-based approaches. We showed that we were able to maintain the "sampling-scalability at high dimension" feature of RRTs while providing them the dynamical context to avoid precision-related failures. We showed that the proposed method maintained the exploration-gearred inventiveness of RRTs, permitting viable solutions where exploitation-driven optimization schemes will fail.

Our examples are restricted to a planar swimming simulated system. Further questions will remain pertinent in extensions of these results to other systems, simulated or real. How viable are pfaffian constraint approximations for systems that do not exist exactly in high or low Reynolds domains? How many degrees of freedom are typically needed to achieve similar results in three-dimensional scenarios? How easy is it to execute trajectories along this manifold with conventional hardware? In future work we plan to explore open questions directly on hardware in underwater and orbital settings.

REFERENCES

- [1] J. C. Kinsey, R. M. Eustice, and L. L. Whitcomb, "A survey of underwater vehicle navigation: Recent advances and new challenges," in *IFAC Conference of Manoeuvring and Control of Marine Craft*, vol. 88. Lisbon, 2006, pp. 1–12.
- [2] E. Simetti, G. Casalino, S. Torelli, A. Sperinde, and A. Turetta, "Floating underwater manipulation: Developed control methodology and experimental validation within the trident project," *Journal of Field Robotics*, vol. 31, no. 3, pp. 364–385, 2014.
- [3] P. Cieslak, P. Ridao, and M. Giergiel, "Autonomous underwater panel operation by girona500 uvms: A practical approach to autonomous underwater manipulation," in *International Conference on Robotics and Automation*. IEEE, 2015, pp. 529–536.
- [4] G. Brantner and O. Khatib, "Controlling ocean one: Human–robot collaboration for deep-sea manipulation," *Journal of Field Robotics*, vol. 38, no. 1, pp. 28–51, 2021.
- [5] F. Ruggiero, V. Lippiello, and A. Ollero, "Aerial manipulation: A literature review," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1957–1964, 2018.
- [6] M. Kobilarov, "Nonlinear trajectory control of multi-body aerial manipulators," *Journal of Intelligent & Robotic Systems*, vol. 73, no. 1, pp. 679–692, 2014.
- [7] M. Zhao, K. Okada, and M. Inaba, "Versatile articulated aerial robot dragon: Aerial manipulation and grasping by vectorable thrust control," *The International Journal of Robotics Research*, p. 02783649221112446, 2022.
- [8] P. Kazanzides, B. P. Vagvolgyi, W. Pryor, A. Deguet, S. Leonard, and L. L. Whitcomb, "Teleoperation and visualization interfaces for remote intervention in space," *Frontiers in Robotics and AI*, vol. 8, 2021.
- [9] E. Papadopoulos, F. Aghili, O. Ma, and R. Lampariello, "Robotic manipulation and capture in space: A survey," *Frontiers in Robotics and AI*, p. 228, 2021.
- [10] E. Krotkov, D. Hackett, L. Jackel, M. Perschbacher, J. Pippine, J. Strauss, G. Pratt, and C. Orłowski, "The darpa robotics challenge finals: Results and perspectives," in *The DARPA Robotics Challenge Finals: Humanoid Robots To The Rescue*. Springer, 2018, pp. 1–26.
- [11] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, p. eaau4984, 2019.
- [12] D. McConachie, A. Dobson, M. Ruan, and D. Berenson, "Manipulating deformable objects by interleaving prediction, planning, and control," *The International Journal of Robotics Research*, vol. 39, no. 8, pp. 957–982, 2020.
- [13] S. Dubowsky and E. Papadopoulos, "The kinematics, dynamics, and control of free-flying and free-floating space robotic systems," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 5, pp. 531–543, 1993.
- [14] M. Mistry, J. Buchli, and S. Schaal, "Inverse dynamics control of floating base systems using orthogonal decomposition," in *International Conference on Robotics and Automation*. IEEE, 2010, pp. 3406–3412.
- [15] L. Righetti, J. Buchli, M. Mistry, and S. Schaal, "Inverse dynamics control of floating-base robots with external constraints: A unified view," in *International Conference on Robotics and Automation*. IEEE, 2011, pp. 1085–1090.
- [16] S. M. LaValle and J. J. Kuffner Jr, "Randomized kinodynamic planning," *The International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.
- [17] S. Kleff, A. Meduri, R. Budhiraja, N. Mansard, and L. Righetti, "High-frequency nonlinear model predictive control of a manipulator," in *International Conference on Robotics and Automation*. IEEE, 2021, pp. 7330–7336.
- [18] T. Rybus, K. Seweryn, and J. Z. Sasiadek, "Control system for free-floating space manipulator based on nonlinear model predictive control (nmpe)," *Journal of Intelligent & Robotic Systems*, vol. 85, no. 3, pp. 491–509, 2017.
- [19] D. Lee, H. Seo, D. Kim, and H. J. Kim, "Aerial manipulation using model predictive control for opening a hinged door," in *International Conference on Robotics and Automation*. IEEE, 2020, pp. 1237–1242.
- [20] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *International Conference on Robotics and Automation*. IEEE, 2009, pp. 489–494.
- [21] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "Stomp: Stochastic trajectory optimization for motion planning," in *International Conference on Robotics and Automation*. IEEE, 2011, pp. 4569–4574.
- [22] S. M. LaValle, J. J. Kuffner, B. Donald, *et al.*, "Rapidly-exploring random trees: Progress and prospects," *Algorithmic and computational robotics: new directions*, vol. 5, pp. 293–308, 2001.
- [23] M. Elbanhawi and M. Simic, "Sampling-based robot motion planning: A review," *IEEE Access*, vol. 2, pp. 56–77, 2014.
- [24] J. Ostrowski and J. Burdick, "The geometric mechanics of undulatory robotic locomotion," *The International Journal of Robotics Research*, vol. 17, no. 7, pp. 683–701, 1998.
- [25] J. E. Marsden and J. Ostrowski, "Symmetries in motion: Geometric foundations of motion control," 1998.
- [26] S. D. Kelly and R. M. Murray, "Geometric phases and robotic locomotion," *Journal of Robotic Systems*, vol. 12, no. 6, pp. 417–431, 1995.
- [27] H. C. Astley, J. R. Mendelson III, J. Dai, C. Gong, B. Chong, J. M. Rieser, P. E. Schiebel, S. S. Sharpe, R. L. Hatton, H. Choset, *et al.*, "Surprising simplicities and syntheses in limbless self-propulsion in sand," *Journal of Experimental Biology*, vol. 223, no. 5, p. jeb103564, 2020.
- [28] D. Zhao, B. Bittner, G. Clifton, N. Gravish, and S. Revzen, "Walking is like slithering: A unifying, data-driven view of locomotion," *Proceedings of the National Academy of Sciences*, vol. 119, no. 37, p. e2113222119, 2022.
- [29] D. Berenson, S. S. Srinivasa, D. Ferguson, and J. J. Kuffner, "Manipulation planning on constraint manifolds," in *International Conference on Robotics and Automation*. IEEE, 2009, pp. 625–632.
- [30] C. Suh, T. T. Um, B. Kim, H. Noh, M. Kim, and F. C. Park, "Tangent space rrt: A randomized planning algorithm on constraint manifolds," in *International Conference on Robotics and Automation*. IEEE, 2011, pp. 4968–4973.
- [31] L. E. Becker, S. A. Koehler, and H. A. Stone, "On self-propulsion of micro-machines at low reynolds number: Purcell's three-link swimmer," *Journal of Fluid Mechanics*, vol. 490, pp. 15–35, 2003.
- [32] R. L. Hatton and H. Choset, "Geometric swimming at low and high reynolds numbers," *IEEE Transactions on Robotics*, vol. 29, no. 3, pp. 615–624, 2013.
- [33] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, vol. 38, no. 1, pp. 3–20, 2002.
- [34] A. Jusufi, D. I. Goldman, S. Revzen, and R. J. Full, "Active tails enhance arboreal acrobatics in geckos," *Proceedings of the National Academy of Sciences*, vol. 105, no. 11, pp. 4215–4219, 2008.
- [35] T. Libby, T. Y. Moore, E. Chang-Siu, D. Li, D. J. Cohen, A. Jusufi, and R. J. Full, "Tail-assisted pitch control in lizards, robots and dinosaurs," *Nature*, vol. 481, no. 7380, pp. 181–184, 2012.
- [36] B. Bittner, R. L. Hatton, and S. Revzen, "Geometrically optimal gaits: a data-driven approach," *Nonlinear Dynamics*, vol. 94, pp. 1933–1948, 2018.
- [37] M. D. Kvalheim, B. Bittner, and S. Revzen, "Gait modeling and optimization for the perturbed stokes regime," *Nonlinear Dynamics*, vol. 97, pp. 2249–2270, 2019.