

Motion Planning for a Climbing Robot with Stochastic Grasps

Stephanie Newdick¹, Nitin Ongole¹, Tony G. Chen²,
Edward Schmerling¹, Mark R. Cutkosky², Marco Pavone¹

Abstract—ReachBot is a robot that uses extendable and retractable booms as limbs to move around unpredictable environments such as martian caves. Each boom is capped by a microspine gripper designed for grasping rocky surfaces. Motion planning for ReachBot must be versatile to accommodate variable terrain features and robust to mitigate risks from the stochastic nature of grasping with spines. In this paper, we introduce a graph traversal algorithm to select a discrete sequence of grasps based on available terrain features suitable for grasping. This discrete plan is complemented by a decoupled motion planner that considers the alternating phases of body movement and end-effector movement, using a combination of sampling-based planning and sequential convex programming to optimize individual phases. We use our motion planner to plan a trajectory across a simulated 2D cave environment with at least 90% probability of success and demonstrate improved robustness over a baseline trajectory. Finally, we use a simplified prototype to verify a body movement trajectory generated by our motion planning algorithm.

I. INTRODUCTION

There is a growing interest in planetary operations that require robots capable of locomotion and mobile manipulation in a variety of terrains. For example, NASA's interest in exploring caves, cliffs, and other rocky terrain on Mars [1, 2] motivates the need for space robots that combine sparse-anchored mobility with high-wrench (force and moment) manipulation. However, there is a key technology gap in existing solutions: small robots are typically restricted to a small reach and limited wrench capability. Conversely, large robots, particularly rigid-link articulated-arm robots, are hampered by high mass and complexity, which scale poorly with increased reach.

ReachBot is a new robot concept for mobile manipulation in environments that require climbing, particularly with sparse anchor points. As debuted in our past works [3, 4], ReachBot consists of a small body that uses multiple extendable booms as appendages to reach to nearby or distant rock surfaces and grip them. These booms are ideal for space applications [5]: they are lightweight and compact when rolled up, while still being strong when deployed and capable of extending many times the span of the robot body (Fig. 1).

With a large reachable workspace, ReachBot can target distant anchor points to pass over local irregularities, leveraging the advantages of a large robot (e.g., RoboSimian, Tri-ATHLETE [6]–[8]) but with reduced weight and complexity. Climbing robots of similar weight (e.g., LEMUR [9]) have limited reach, preventing them from traversing environments




Fig. 1: A simplified, planar, cable-based ReachBot prototype moves under force control (from 1 \rightarrow 2) while keeping contact forces within stochastic limits at four anchor points.

with sparse anchor points or obstructions such as rockfall and crevasses. ReachBot's large workspace allows for versatility in anchor point selection, a notable advantage when navigating these environments. The dynamical model and simulation presented in past work demonstrated ReachBot's feasibility [3], but did not provide a high-level motion plan for ReachBot to explore and adapt to unique terrain geometries. In this paper we present a planning approach that leverages ReachBot's versatile design while accounting for the uncertain terrain encountered in a martian cave.

ReachBot's motion planner must address multiple challenges that arise from its mobility paradigm. It uses microspine grippers to create attachments to surrounding basaltic surfaces, allowing it to manipulate its body by applying forces [4]. The stochastic nature of grasping rock with spines introduces two complications. First, the non-zero failure probability of each new grasp highlights the importance of making as few regrasps as possible, requiring a motion planner that looks ahead at least a few moves before grasping a particular anchor point. Second, the stochasticity of the limit surface in force space for each attachment introduces uncertainty into the planning that we must consider to ensure robust mobility. Furthermore, the dynamic and kinematic constraints create a highly nonlinear system.

Statement of Contributions: This paper addresses the need for a grasp-aware motion planner to generate robust trajectories for an extendable-boom climbing robot through three main contributions. (1) We propose ReachBot Planner (RBP), a framework for a non-gaited climbing robot that splits the motion planning problem into three decomposed segments: footstep planning, body movement, and end-effector movement. (2) We introduce a convex metric for stochastic grasp success and use this metric to quantify and optimize the robustness of the planned trajectories. (3) We leverage the convexity of this metric and employ sequential convex programming to locally optimize trajectories between discrete

¹Stephanie Newdick, Nitin Ongole, Edward Schmerling, and Marco Pavone are with the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305. {snewdick, nitin99, schmrlnrg, pavone}@stanford.edu.

²Tony G. Chen and Mark R. Cutkosky are with the Department of Mechanical Engineering, Stanford University, Stanford, CA 94305. {agchen, cutkosky}@stanford.edu.

footsteps, then validate the trajectories in simulation and on a 2D prototype. While this approach is formulated around ReachBot, we believe it is broadly applicable to cable-supported robots and other multi-limbed climbing robots with uncertain grasp strength (e.g. [6, 9, 10]).

Paper Organization: The rest of this paper is organized as follows. In Section II, we discuss existing motion planning approaches for legged, cable-suspended, and climbing robots. In Section III, we introduce our robotic formulation and gripper model and present metrics for robustness and feasibility of the system. In Section IV, we present a hierarchical motion plan wherein we first determine a discrete sequence of footsteps, then use sampling-based planning and sequential convex programming to optimize each continuous movement segment with respect to the metric of robustness. In Section V, we compare our strategy against a naive baseline in simulation and validate our simulation results on a hardware prototype.

II. RELATED WORK

Motion planning for legged or climbing robots requires a combination of discrete footstep planning and continuous trajectory optimization. Many recent approaches use mixed-integer convex optimization methods to find the footstep and continuous trajectory simultaneously [11, 12]. Additionally, several papers have demonstrated the ability of nonlinear solvers to provide coupled motion plans [13]. However, such methods have proven to be prohibitively slow for robots with highly nonlinear dynamics like ReachBot [14].

Alternatively, decoupled planners for multi-contact problems offer sub-optimal but faster solutions. Decoupled planners can be split into two categories: motion-before-contact and contact-before-motion. The former approach uses continuous planning to find a centroid trajectory, then determines contact placement. It is effective for robots using gaited mobility or for robots locomoting on flat or otherwise predictable terrain where contact points will always be available [10, 15]–[17]. Conversely, a contact-before-motion approach is required for robots such as ReachBot that navigate uncertain terrain with irregular and unpredictable geometry.

Contact-before-motion planners generate a discrete plan by considering adjacent contacts, often using graph search methods. After a sequence of contacts is determined, a continuous trajectory is planned to string the sequence together using optimal control or sampling-based planning [18]. Contact-before-motion planners typically use optimal or locally optimal control methods such as sequential convex programming (SCP) to minimize control effort [19].

Capuchin [20] and LEMUR-3 [9] offer two examples of non-linear systems designed for climbing in irregular terrain for which contact-before-motion planners have been demonstrated. The motion planner for Capuchin generates a footstep path by traversing a graph of stances where every feasible transition is an edge on the graph; LEMUR’s sequence of stances is chosen by a human operator. Both robots use sampling-based methods to compute feasible paths between two stances.

Although many robots rely on contact interactions that are either unknown or inherently stochastic (like ReachBot’s




Fig. 2: (Left) ReachBot in a 4-stance with all four booms attached. The axial forces F_i for each boom i must counteract the weight of the body, Mg . (Right) ReachBot in a 3-stance where one end-effector is detached. The weight of the gripper provides an external torque ($mg\Delta x$) on the body which must be supported by the active shoulder joint and the other three booms.

spiny grippers), few existing motion planners explicitly account for this uncertainty. One such planner models gripper uncertainty by augmenting friction cone constraints for a hexapod climbing robot [21]. However, this solution’s reliance on a predefined gait would sacrifice the versatility of ReachBot’s large workspace. Our contribution is therefore to present a contact-before-motion planner for a multi-contact robot with nonlinear dynamics that minimizes the likelihood of grasp failure throughout the trajectory.

III. PROBLEM FORMULATION

In this section, we reintroduce a formulation of ReachBot that will be used throughout this work. We then introduce a grasp limit surface model that represents the stochasticity of grasping with microspines as Gaussian uncertainty. By parameterizing ReachBot’s probability of success in any configuration, this grasp model is crucial to the development of a robust motion plan. Finally, we define a notion of transition feasibility used to link sequential footsteps.

This paper assumes a 2D simplification of ReachBot to coincide with past work in modeling and prototype development [3, 4]. Planar ReachBot includes a robot body with four shoulder mechanisms, each having two controllable joints: (1) a prismatic joint composed of a motor that extends and retracts a deployable boom (see F_i in Fig. 2) and (2) a revolute shoulder joint that pivots the boom in the plane. As in our past work, ReachBot moves by alternating between body movement and end-effector movement. A *stance* is an assignment of end-effectors to anchor points. During body movement, shown in Fig. 2 (left), all four of ReachBot’s end-effectors remain anchored in a “4-stance” while the body repositions. During end-effector movement, ReachBot must hold itself in a “3-stance” while it moves its free end-effector to a new anchor point as shown in Fig. 2 (right).

A. Gripper model

The contact model presented in this section determines the allowable pulling forces (magnitude and direction) for a given grasp. Microspine grippers have been demonstrated to sustain large loads on a variety of rock-climbing robots [9, 22, 23]. Past analyses of spiny grippers developed models to generate a stochastic wrench limit surface for a given gripper and surface topography [24]–[26]. We are currently developing models to construct a 6D limit surface for ReachBot’s

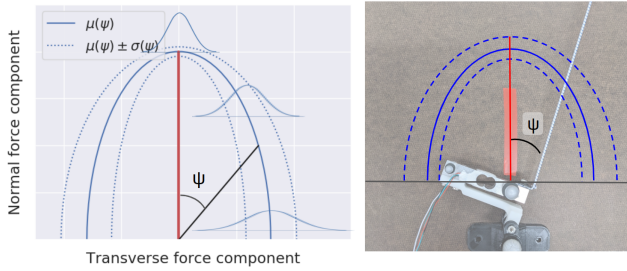


Fig. 3: (Left) Slice of stochastic half-ellipsoid used to model the limit surface of a grasp. Direction of pulling force is parameterized by ψ , where $\psi = 0$ corresponds to a purely normal force, and $\psi = \pm\pi/2$ denotes lateral pulling. The expected limit is given by $\mu(\psi)$, and uncertainty is taken to be Gaussian, with standard deviation $\sigma(\psi)$ varying linearly with ψ . (Right) Limit surface ellipsoid overlaid on prototype anchor. We determine the limiting force as a function of ψ , then compare it to the pulling force measured by a strain gauge.

gripper based on grasp angle and the size and distribution of asperities in the rock. Under the assumption that booms can pull but not push, we approximate this limit surface as a stochastic half ellipsoid parameterized by empirically-determined semi-major and semi-minor axes [25].

Fig. 3 shows a 2D slice of an example ellipsoidal limit surface. The force that the grasp can sustain is a function of the pull angle ψ . The spine-surface interaction is itself a probabilistic function [24, 26], and therefore the expected sustainable gripper force at a given angle ψ is also stochastic. We assume that this uncertainty can be modeled as Gaussian, with the standard deviation increasing linearly as a function of ψ . This is consistent with Gaussian models of asperity slope variation on basaltic surfaces [27]. Asperity slope will be the largest source of uncertainty for ReachBot’s gripper design, so we assume an overall Gaussian uncertainty model.

B. Robustness metrics

ReachBot’s robustness to failure depends on its configuration, defined by stance, pose, and contact forces. In this section, we define a metric for configuration robustness based on the likelihood of failure of each individual grasp.

We consider a generalized ReachBot with n booms so the formulation can be applied to a 3-stance, 4-stance, or future n -stance 3D configuration. Using a Gaussian model of uncertainty, the maximum sustainable force by grasp i is given by

$$f_{\{i\},\max} \sim \mathcal{N}(\mu_{\{i\}}(\psi_{\{i\}}), \sigma_{\{i\}}^2(\psi_{\{i\}})), \quad (1)$$

where $\mathcal{N}(\mu_{\{i\}}, \sigma_{\{i\}})$ represents the stochastic pull force limit for grasp i . This distribution is defined by local rock properties and the pull angle $\psi_{\{i\}}$ for each grasp i . We define the probability of success for the n -stance configuration as the joint probability that all forces $f_{\{i\}}$ are within their respective limits. Since the wrist joints are passive, $f_{\{i\},\max}$ for each grasp is fully determined by the anchor point coordinates and the pose and orientation of ReachBot. Therefore, the probability of each grasp’s success is conditionally independent given the robot’s pose and stance, and we can

write the joint probability as

$$P(f_{\{1:n\}} \leq f_{\{1:n\},\max} | \psi_{\{1:n\}}) = \prod_{i=1}^n P(f_{\{i\}} \leq f_{\{i\},\max} | \psi_{\{i\}}). \quad (2)$$

By negating the force terms and applying the Gaussian cumulative distribution function $g(\cdot)$, we can rewrite the probability of success as

$$P(-f_{\{1:n\},\max} \leq -f_{\{1:n\}} | \psi_{\{1:n\}}) = \prod_{i=1}^n g\left(\frac{-f_{\{i\}} + \mu_{\{i\}}(\psi_{\{i\}})}{\sigma_{\{i\}}(\psi_{\{i\}})}\right). \quad (3)$$

We take the log of both sides to define a convex robustness function

$$r(s, u) = \sum_{i=1}^n \log \left[g\left(\frac{-f_{\{i\}} + \mu_{\{i\}}(\psi_{\{i\}})}{\sigma_{\{i\}}(\psi_{\{i\}})}\right) \right], \quad (4)$$

where $r(s, u)$ represents the negative log of the probability of success for a stance with a configuration parameterized by state vector s and control vector u .

For the current robot prototype, the gripper will reach failure before enough pulling force can be applied to break the boom in tension. Therefore, the robustness of a configuration depends only on the probability of grasp failure. In future work, our gripper will also be able to apply modest pushing force on the wall, so robustness will additionally depend on potential failure due to exceeding the Coulomb friction cone, or the compression and bending limits of the boom.

Even with adequate force sensing capability and grasp limit surface modeling, the stochastic nature of spine contact on unknown rock surfaces introduces a non-zero chance of failure with every new grasp. Moreover, there is some chance that the rock could fracture. Therefore, to increase the robustness of a multi-footstep trajectory, we must additionally aim to minimize the overall number of regrasps.

C. Transition feasibility

To define transition feasibility, or the feasibility of moving between two stances, we must first define the feasibility of a pose. For a given stance, ReachBot’s pose (position and orientation) is feasible if it meets the following requirements:

- It meets kinematic constraints. For our prototype, the booms can retract to a minimum of 1cm and extend to a maximum of 2m. The shoulder joints have a range of motion of $\pm 45^\circ$ in either direction from equilibrium.
- No part of the robot body or booms is colliding with itself or with the environment.
- There exists a combination of control inputs where ReachBot resists external forces (e.g. gravity) and the configuration’s robustness (4) is greater than some minimum value. For the rest of this paper, we define a minimum allowable probability of success of 90%, so

$$r_{\min} = -\log(0.90). \quad (5)$$

A transition between two 4-stances is feasible if there exists a pose that is (1) feasible in the starting 4-stance, (2) feasible in the goal 4-stance, (3) feasible in the common 3-stance while holding the free boom statically at its starting anchor point, and (4) feasible in the common 3-stance while holding the free boom statically at its final anchor point. In other words, for a stance transition to be feasible, there must

exist a pose that ReachBot can hold throughout the entire maneuver while maintaining robustness of at least r_{\min} .

IV. ROBUST MOTION PLANNING

In this section, we present ReachBot Planner decomposed in three pieces: footstep planning, body movement, and end-effector movement. First, we present a discrete footstep planner that uses graph search to define a sequence of stances linked by feasible transitions. From this discrete plan, we break ReachBot's motion down into a sequence of alternating body movement and end-effector movement phases. For both of these continuous phases, we generate a feasible trajectory using sampling-based techniques, then locally optimize these trajectories using sequential convex programming. We present representative samples of a full path generated by RBP for a simulated 2D environment where the planner has a priori knowledge of the randomly-generated anchor point locations and parameters of the limit surfaces' ellipsoids.

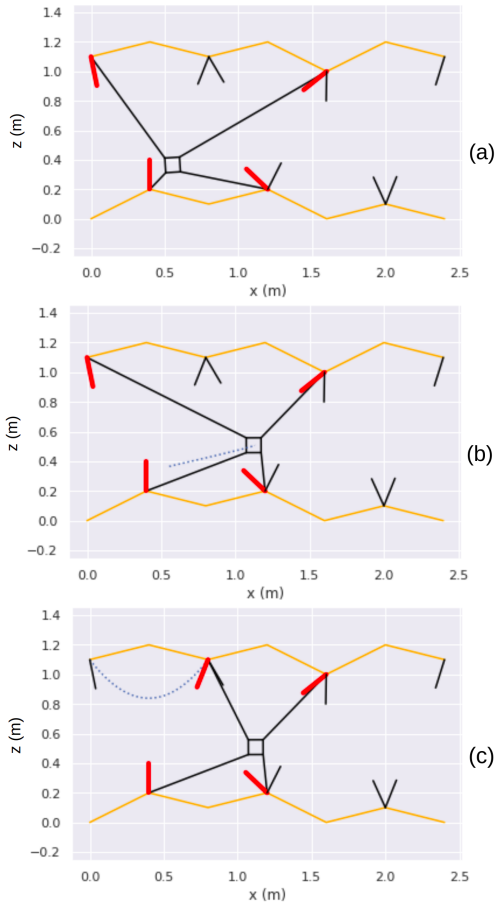


Fig. 4: Snapshots of ReachBot's configuration at discrete points in its footstep plan. The solid orange lines represent the cave walls, and the black line segments represent the line at $\psi = 0$ (see Fig. 3) for each respective anchor point limit surface. This line is red if that anchor point is being grasped. Figure (a) shows a snapshot of a 4-stance before body movement. From (a) to (b), ReachBot undergoes a body movement phase to assume a configuration where it can stably detach and move its top left end-effector. From (b) to (c), ReachBot moves its end-effector between adjacent 4-stances. The blue dotted line signifies the trajectory of either movement.

A. Footstep planning

The overall goal of footstep planning is to find a path of feasible poses and stance transitions to move the body from its starting position to its goal position. We assume for now that we have full knowledge of the available grasp points, including their position, orientation, and limit surface parameters throughout the environment.

We construct a graph wherein each 4-stance is a vertex, and each edge connecting adjacent vertices represents the common 3-stance. We use A* to traverse the graph with distance-to-goal heuristic to encourage larger steps and therefore fewer regrasps. To speed up graph traversal, we only populate edges where there is a feasible transition between the two vertices. The limiting torques that a 3-stance must resist during a transition occur immediately after detachment and immediately before initiating the next grasp. These two scenarios represent the maximum horizontal extension of the moving boom in either direction, and therefore the maximum external moments exerted on the body. We leverage the convexity of (4) in addition to affine static equilibrium constraints to search for a feasible transition pose. In particular, we use Hooke-Jeeves gradient-free optimization to search around randomly sampled poses for a pose from which ReachBot can resist both limiting torques. If such a point is found, the transition is feasible, and we continue searching the graph in that direction. Fig. 4 shows an example of two steps ReachBot takes during its path from start to goal.

B. Body movement phase

In the body movement phase, ReachBot moves within a 4-stance to prepare for the next stance transition. The footstep path guarantees that each pose at the beginning and end of a required body movement is feasible, but we further require that every intermediate pose must be statically feasible. To this end, we generate a discretized trajectory of statically feasible poses by sampling locally from a straight-line trajectory, verifying at each point that a combination of forces exists to hold that pose statically with at least 90% probability of success. We use this trajectory to warm start the SCP solver, which maximizes the success probability by sequentially approximating the non-convex body movement problem with the following convex optimization problem:

$$\min_{\substack{s_k, u_k \\ \forall k \in \{1..N\}}} \sum_{k=1}^N \left(-\alpha r(s_k, u_k) + \lambda_s \|s_k^{(w)} - s_k^{(w-1)}\| + \lambda_u \|u_k^{(w)} - u_k^{(w-1)}\| \right) \quad (6)$$

$$\begin{aligned} \text{s.t. } & \|s_{k+1} - A(s_k, u_k)\| \leq \beta^{(w)} & \forall k \in \{1..N-1\} \\ & s_1 = s_{\text{start}} \\ & s_N = s_{\text{goal}} \\ & \theta_{\min} \leq \theta_{k\{i\}} \leq \theta_{\max} & \forall k \in \{1..N\}, i \in \{1..4\} \\ & b_{\min} \leq b_{k\{i\}} \leq b_{\max} & \forall k \in \{1..N\}, i \in \{1..4\} \\ & f_{\min} \leq f_{k\{i\}} \leq f_{\max} & \forall k \in \{1..N\}, i \in \{1..4\} \\ & r_{\min} \leq r(s_k, u_k) & \forall k \in \{1..N\} \end{aligned}$$

where the optimization variables are the state vector

$$s_k = [x_k, y_k, \phi_k, \dot{x}_k, \dot{y}_k, \dot{\phi}_k], \quad (7)$$

which defines the body position (x, y) , orientation ϕ , linear velocities (\dot{x}, \dot{y}) and angular velocity $\dot{\phi}$ at timestep k , and the control vector

$$u_k = [f_{k\{1\}}, f_{k\{2\}}, f_{k\{3\}}, f_{k\{4\}}], \quad (8)$$

which defines the control input for each boom, $f_{k\{i\}}$, equal to the prismatic force applied to boom i at timestep k . We use $(\cdot)^{(w)}$ to denote the value of (\cdot) at SCP iteration w ; where not explicitly noted all variables and derived quantities correspond to SCP iteration w .

To optimize for robustness, we minimize $-r(s, u)$ at all timesteps in the trajectory (6). In each iteration of SCP, we penalize divergence from the local convexification using trust parameters for the state λ_s and control λ_u .

The function A represents ReachBot's dynamics linearized around the previous SCP iterate. Because the seed trajectory $(s^{(0)}, u^{(0)})$ is constructed from a string of static intermediate poses, we expect it to be dynamically infeasible. We use the parameter $\beta^{(w)}$, which is reduced every iteration, to allow infeasible dynamics in early iterations of SCP but constrain that the final trajectory has feasible dynamics.


The start and goal states s_{start} and s_{goal} for each body movement phase are determined by the footstep planner. We apply kinematic and control constraints as they were introduced in Section III-C, where $b_{k\{i\}}$ is the length of boom i at timestep k and $\theta_{k\{i\}}$ is the angle from rotational equilibrium of boom i and timestep k . We additionally constrain a lower bound on robustness as defined in Section III-C. Fig. 5(a) shows the performance of RBP on the example body movement segment in Fig. 4. After optimization, the controls are within actuator limits and overall smaller in magnitude, as seen in Fig. 5(a, left). A smaller pulling magnitude decreases the probability of exceeding a gripper's limit surface, as evidenced by the increased probability of success shown in Fig. 5(a, right). Both the seed and optimized trajectory are feasible with a minimum probability of success above 90%.

C. End-effector movement phase


In the end-effector movement phase, ReachBot remains in a transitional 3-stance while it moves its free end-effector from one anchor point to another defined by the footstep planner in Section IV-A. During end-effector movement, the cantilevered boom exerts an external wrench on the robot body, requiring additional compensation to keep ReachBot stationary. By keeping movement slow enough, we can ignore inertial effects and assert that the external wrench is due entirely to the cantilever. With this assertion, the external wrench is independent of the exact path taken by the end-effector, so we arbitrarily define a collision-free parabolic trajectory between the two anchor points. We generate a seed path of feasible body poses as in Section IV-B, then use SCP to maximize robustness throughout the trajectory.

The full optimization problem is nearly identical to that in IV-B with only slight modifications. First, the robustness $r(s_k, u_k)$ only considers the ellipses of the three attached points. Second, the control vector must include torsional input at the shoulder joint in addition to prismatic forces

$$u_k = [f_{k\{i1\}}, f_{k\{i2\}}, f_{k\{i3\}}, f_{k\{j\}}, t_{k\{j\}}],$$



(a) Analysis of the body movement between Fig. 4(a) and Fig. 4(b).



(b) Analysis of the end-effector movement between Fig. 4(b) and Fig. 4(c).

Fig. 5: (Left) applied forces (one color per boom); (right) probability of success throughout the trajectory. In general, the SCP-optimized trajectory (solid line) encourages less control effort than the seed trajectory (dotted line), decreasing the probability of exceeding the gripper limit surface. This corresponds to an increased probability of success in the optimized trajectory.

where $f_{k\{i1\}}, f_{k\{i2\}}, f_{k\{i3\}}$ are forces applied at the three anchored booms, and $f_{k\{j\}}$ and $t_{k\{j\}}$ represent the prismatic force and rotational torque, respectively, applied to moving boom j at timestep k . Lastly, an external wrench term imparted by the cantilevered boom and end-effector is included in the nonlinear dynamics function A .

Fig. 5(b) shows a similar performance increase after optimizing end-effector movement as for body movement. We again consider an example segment from Fig. 4. The smaller actuator output shown in Fig. 5(b, left) induces an increased probability of success shown in Fig. 5(b, right). Both the seed and optimized trajectory are feasible with a minimum probability of success above 90%.

V. RESULTS

In this section, we introduce a naive footstep planner that neglects limit surface constraints. We compare this naive baseline to RBP in terms of success rate and computation times. We then demonstrate body movement on a hardware platform to verify robustness throughout trajectory.

A. Simulation Results

We construct a naive footstep planner that imitates the graph traversal algorithm presented in Section IV-A but without considering the feasibility of the poses due to gripper limitations. Fig. 6 shows the results of 100 paths generated by both footstep planners for environments with randomized anchor point placement and limit surface magnitude and orientation. To minimize distance-to-goal, the naive planner greedily selects the farthest anchor point in reach, resulting in a plan with fewer regrasps. In general, fewer regrasps corresponds to fewer chances for failure. However, the overwhelming superiority of RBP highlights the importance of considering limit surfaces: a path with more regrasps that ensures higher robustness per grasp performs better overall than simply minimizing number of regrasps.

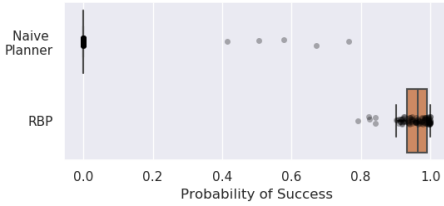


Fig. 6: Probability of success of footstep planners over 100 randomized trials. The naive baseline planner does not account for limit surfaces, leading to a low success rate due to gripper failure, though it does succeed in some instantiations. RBP returns footstep paths with over 90% probability of success in 86 trials out of 100 and paths with over 90% success rate in more than half of the trials.

Table I provides the average compute time of RBP’s footstep planner over 100 randomized iterations as compared with the naive planner. As the naive planner produces so few feasible footstep paths, we cannot compare computation time of the continuous portions of the trajectory to a baseline. The computation time for the initial seed trajectory as well as additional time to perform SCP optimization are listed in Table I for individual body and end-effector movement segments. The total computation time for the multistep trajectory plan is, for our unoptimized implementation, on the order of minutes owing to the highly nonlinear nature of the optimization and our corresponding conservative trust region definition. We note, however, that the computation time is compatible (even accounting for replanning) with the intended quasistatic deployment of ReachBot where this duration is comparable to the implementation of a single body movement segment.

	RBP time (s)	Naive time (s)
footstep plan	66.19	0.0123
seed for body mvmt	0.0839	N/A
SCP for body mvmt	239.78	N/A
seed for end-eff. mvmt	0.0417	N/A
SCP for end-eff. mvmt	398.70	N/A

TABLE I: Planning time for each segment of the motion plan.

B. Experimental Results

In order to validate the trajectory produced by RBP, we built a critical function prototype with a small robot body and elastic cords in place of ReachBot’s rigid booms. The purpose of this prototype is to verify robustness by measuring the force magnitude and pulling angle at each anchor throughout a trajectory. Since the cords are rigidly attached to the anchors, we only use this setup to test body movement trajectories. There are four motors driving four independent winches controlling both the lengths and tensions of the cords. Each anchor is equipped with a single axis load cell mounted on a hinge which allows it to freely rotate. The full experimental setup is shown in Fig. 1, with a close-up view of an anchor in Fig. 3 (Right).

To validate simulation results, we execute a body movement trajectory on our hardware prototype. Compliance in the elastic cords adds artificial transient dynamics, but we use closed-loop force control to track the trajectory, so the forces on the anchors during motion are representative of a rigid-boom ReachBot. Using the force and angle at each anchor,

we calculate the probability of success of each representative grasp based on the predefined limit surface parameters. The results are shown in Figure 7 after being passed through a low-pass filter to reduce noise. Notable trends include the increasing probability of success at the lower right anchor as the boom angle approaches the ellipse normal direction ($\psi = 0$). Similarly, we see the probability of success for the lower left anchor decrease as the boom angle diverts from the ellipse’s normal.

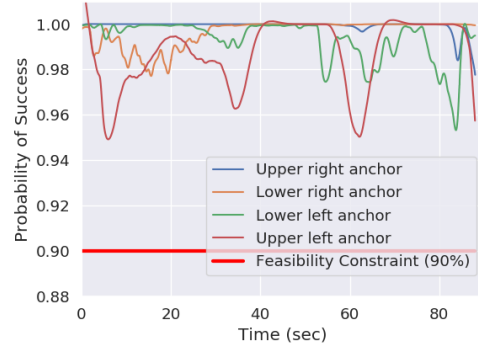


Fig. 7: Probability of grasp maintenance for each anchor on our hardware platform while implementing a single body movement trajectory. Probability is a function of force as defined in (2).

VI. CONCLUSION

In this paper, we presented ReachBot Planner (RBP), a motion planner for a climbing robot that incorporates stochastic limit surfaces via a convex metric for robustness based on an ellipsoidal limit surface model. First, RBP determines a sequence of feasible stances, then it generates a feasible seed trajectory for each body and end-effector movement segment that it optimizes uses SCP to maximize probability of success throughout the trajectory. We demonstrated on hardware that the RBP-generated trajectory keeps the forces of all four grasps within their respective limit surfaces during body movement, and showed in simulation that a naive planner results in an unacceptably low probability of success, suggesting the necessity of a planner such as RBP. Furthermore, the stark difference between performance of a naive planner as compared to RBP implies high sensitivity to the limit surface model, encouraging further work in modeling for microspine grippers.

Our eventual goal is to design, build, and field-test a 3D ReachBot prototype. Extending RBP to 3D should be straightforward as the convex formulation of the robustness metric is preserved in 3D. However, full knowledge of available grasp points is unrealistic in a real-world setting, so we are concurrently developing perception strategies to quantify grasp sites in an unfamiliar environment. Additionally, we will develop a metric for robustness to dynamic failure and perform experiments on a representative hardware platform before deploying it on a 3D robot.

ACKNOWLEDGEMENTS

Support for this work was provided by NASA under the NIAC program and by the Air Force under an STTR award with Altius Space Machines. S. Newdick is supported by NASA NSTGRO and T. G. Chen is supported by NSF GRFP.

REFERENCES

- [1] National Research Council, "Decadal survey vision and voyages for planetary science in the decade 2013–2022," National Academy Press, Tech. Rep., 2011.
- [2] M. G. A. Lap tre, J. G. O'Rourke, L. K. Schaefer, K. L. Siebach, C. Spalding, S. M. Tikoo, and R. D. Wordsworth, "Probing space to understand earth," *Nature Reviews Earth & Environment*, vol. 1, pp. 170–181, 2020.
- [3] S. Schneider, A. Bylard, T. G. Chen, P. Wang, M. R. Cutkosky, and M. Pavone, "ReachBot: A small robot for large mobile manipulation tasks," in *IEEE Aerospace Conference*, 2022, in press.
- [4] T. G. Chen, B. Miller, C. Winston, S. Schneider, A. Bylard, M. Pavone, and M. R. Cutkosky, "ReachBot: A small robot with exceptional reach for rough terrain," in *Proc. IEEE Conf. on Robotics and Automation*, 2022, press.
- [5] C. Leclerc, L. Wilson, M. A. Bessa, and S. Pellegrino, "Characterization of ultra-thin composite triangular rollable and collapsible booms," in *AIAA SPACE Conferences & Exposition*, 2017.
- [6] P. Hebert, M. Bajracharya, J. Ma, N. Hudson, A. Aydemir, J. Reid, C. Bergh, J. Borders, M. Frost, M. Hagman, J. Leichty, P. Backes, B. Kennedy, P. Karplus, B. Satzinger, K. Byl, K. Shankar, and J. Burdick, "Mobile manipulation and mobility as manipulation—Design and algorithms of RoboSimian," *Journal of Field Robotics*, vol. 32, no. 2, pp. 255–274, 2015.
- [7] B. H. Wilcox, "ATHLETE: A limbed vehicle for solar system exploration," in *IEEE Aerospace Conference*, 2012.
- [8] M. Heverly, J. Matthews, M. Frost, and C. McQuin, "Development of the Tri-ATHLETE lunar vehicle prototype," in *Aerospace Mechanisms Symposium*, 2010.
- [9] A. Parness, N. Abcouwer, C. Fuller, N. Wiltsie, J. Nash, and B. Kennedy, "LEMUR 3: A limbed climbing robot for extreme terrain mobility in space," in *Proc. IEEE Conf. on Robotics and Automation*, 2017.
- [10] A. Capua, A. Shapiro, and S. Shoval, "Spiderbot: a cable-suspended walking robot," *Mechanism and Machine Theory*, vol. 82, pp. 56–70, 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094114X14001839>
- [11] B. Aceituno-Cabezas, C. Mastalli, H. Dai, M. Focchi, A. Radulescu, D. G. Caldwell, J. Cappelletto, J. C. Grieco, G. Fernandez-Lopez, and C. Semini, "Simultaneous contact, gait, and motion planning for robust multilegged locomotion via mixed-integer convex optimization," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2531–2538, 2017.
- [12] M. S. Ahn, H. Chae, and D. W. Hong, "Stable, autonomous, unknown terrain locomotion for quadrupeds based on visual feedback and mixed-integer convex optimization," in *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2018.
- [13] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli, "Gait and trajectory optimization for legged systems through phase-based end-effector parameterization," *IEEE Robotics and Automation Letters*, vol. 3, pp. 1560–1567, 2018.
- [14] X. Lin, J. Zhang, J. Shen, G. Fernandez, and D. W. Hong, "Optimization based motion planning for multi-limbed vertical climbing robots," in *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2019.
- [15] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *Int. Journal of Robotics Research*, vol. 33, no. 1, pp. 69–81, 2014.
- [16] A. Laurenzi, E. M. Hoffman, and N. G. Tsagarakis, "Quadrupedal walking motion and footstep placement through linear model predictive control," in *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2018.
- [17] A. Sintov, T. Avramovich, and A. Shapiro, "Design and motion planning of an autonomous climbing robot with claws," *Robotics and Autonomous Systems*, vol. 59, no. 11, pp. 1008–1019, 2011.
- [18] K. Hauser, T. Bretl, and J.-C. Latombe, "Non-gaited humanoid locomotion planning," in *IEEE RAS Int. Conf. on Humanoid Robots*, 2005.
- [19] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *Int. Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [20] R. Zhang and J.-C. Latombe, "Capuchin: A free-climbing robot," *Int. Journal of Advanced Robotic Systems*, vol. 10, no. 4, 2013.
- [21] Y. Shirai, X. Lin, Y. Tanaka, A. Mehta, and D. Hong, "Risk-aware motion planning for a limbed robot with stochastic gripping forces using nonlinear programming," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 4994–5001, 2020.
- [22] A. T. Asbeck, S. Kim, M. R. Cutkosky, W. R. Provancher, and M. Lanzetta, "Scaling hard vertical surfaces with compliant microspine arrays," *Int. Journal of Robotics Research*, vol. 25, no. 12, pp. 1165–1179, 2006.
- [23] W. Ruotolo, F. S. Roig, and M. R. Cutkosky, "Load-sharing in soft and spiny paws for a large climbing robot," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1439–1446, 2019.
- [24] H. Jiang, S. Wang, and M. R. Cutkosky, "Stochastic models of compliant spine arrays for rough surface grasping," *Int. Journal of Robotics Research*, vol. 37, no. 7, pp. 669–687, 2018.
- [25] S. Wang, H. Jiang, T. Myung Huh, D. Sun, W. Ruotolo, M. Miller, W. R. T. Roderick, H. S. Stuart, and M. R. Cutkosky, "Spinyhand: Contact load sharing for a human-scale climbing robot," *Journal of Mechanisms and Robotics*, vol. 11, no. 3, 2019.
- [26] S. Iacoponi, M. Calisti, and C. Laschi, "Simulation and analysis of microspines interlocking behavior on rocky surfaces: an in-depth study of the isolated spine," *Journal of Mechanisms and Robotics*, vol. 12, no. 6, 2020.
- [27] H. Jiang, "Surface grasping with bio-inspired, opposed adhesion," Ph.D. dissertation, Stanford University, 2017.