

Task-Driven Manipulation with Reconfigurable Parallel Robots

Daniel Morton¹, Mark Cutkosky¹, and Marco Pavone²

Abstract—ReachBot, a proposed robotic platform, employs extendable booms as limbs for mobility in challenging environments, such as martian caves. When attached to the environment, ReachBot acts as a parallel robot, with re-configuration driven by the ability to detach and re-place the booms. This ability enables manipulation-focused scientific objectives: for instance, through operating tools, or handling and transporting samples. To achieve these capabilities, we develop a two-part solution, optimizing for robustness against task uncertainty and stochastic failure modes. First, we present a mixed-integer stance planner to determine the positioning of ReachBot’s booms to maximize the task wrench space about the nominal point(s). Second, we present a convex tension planner to determine boom tensions for the desired task wrenches, accounting for the probabilistic nature of microspine grasping. We demonstrate improvements in key robustness metrics from the field of dexterous manipulation, and show a large increase in the volume of the manipulation workspace. Finally, we employ Monte-Carlo simulation to validate the robustness of these methods, demonstrating good performance across a range of randomized tasks and environments, and generalization to cable-driven morphologies. We make our code available at our project webpage, https://stanfordasl.github.io/reachbot_manipulation/

I. INTRODUCTION

Scientific research missions in space have significantly advanced our understanding of Earth, the solar system, and the broader cosmos. Recently, caves and lava tubes on celestial bodies like the Moon and Mars have emerged as areas of interest due to their distinct geological and astrobiological characteristics, offering potential insights into the history of the solar system [1], [2]. Further exploration of these regions may uncover details about their past and future habitability, but this exploration is limited by the capabilities of existing robotic platforms [3], and their inability to access hard-to-reach locations. While passive observation of the environment is suitable for some missions, direct interaction and manipulation unlock a broader range of scientific possibilities, particularly when unconstrained by a small workspace. For instance, a robot with broad mobility and manipulation capabilities could extract geological samples from challenging regions, or assemble preliminary infrastructure for a habitat, in preparation for human arrival.

To address this need, ReachBot is a proposed robotic concept for enhanced mobility in challenging environments

This work was supported by NASA under the NIAC program (Grant #80NSSC22K0766) and a NASA Space Technology Graduate Research Opportunity. D. Morton was supported by NSF GRFP.

¹Daniel Morton and Mark Cutkosky are with the Department of Mechanical Engineering, Stanford University, Stanford, CA 94305. {dmorton, cutkosky}@stanford.edu

²Marco Pavone is with the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305. pavone@stanford.edu



Fig. 1. ReachBot performs a sample-extraction manipulation task in a cave environment. By optimally planning (i) where to place the booms and (ii) the tensions in each boom, we ensure that this task is completed even with uncertain pose and wrench estimation, and stochastic failure modes.

[4]–[8]. Using deployable booms as reconfigurable prismatic joints, ReachBot can extend and traverse across large regions, accessing hard-to-reach areas of scientific interest (Fig. 1). However, accessing these areas is only the first step of a larger science mission with manipulation-based objectives. Building on the sample extraction example, this task will require (i) aligning a camera to analyze a region while maintaining a steady view, (ii) applying a force or torque to extract the sample, through grasping or drilling, and (iii) transporting the sample to a retrieval location. During this process, ReachBot must also address the challenges of (i) the stochastic nature of grasping a rock with microspines, (ii) uncertainty in the required set of wrenches and poses, and (iii) vibrations and disturbances introduced during execution.

A. Related Work

For this manipulation application, with multiple booms attached to the environment, ReachBot is most directly related to cable-driven parallel robots (CDPRs) and cable-suspended robots (CSRs) [9]. Of these, ReachBot is comparable to fully-constrained 6D (position and orientation) CDPRs (requiring at least 7 cables) – as opposed to CDPRs which only control position with fewer than 7 cables. Representative examples of fully-constrained and under-constrained CDPRs are [10] and [11], respectively.

Generally, these CDPRs are not reconfigurable, meaning that the positions of the cable attachment points are fixed. There are a few notable examples of reconfigurable CDPRs [12], yet often this reconfigurability is either limited to a single translational axis, or it is a process that is not performed on a task-by-task basis. SpiderBot [13] is an example

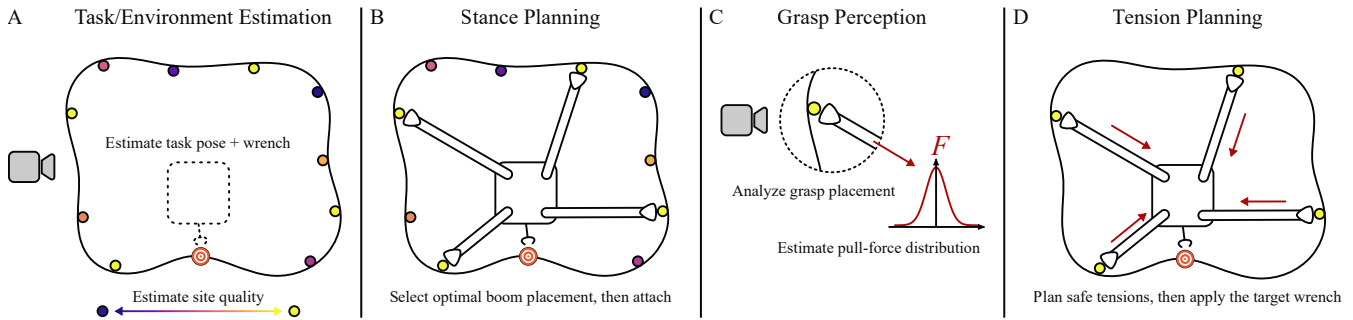


Fig. 2. **Concept of operations.** (A): Upon approaching an area of scientific interest, ReachBot’s perception system identifies where the task needs to be performed, estimates a task wrench, identifies candidate grasp sites in the area, and estimates the quality of each site. (B): The *stance planner* considers the grasp sites and determines the optimal placement of each boom, for robust task execution. (C): After each boom is attached, the perception system can identify the quality of each grasp and determine the expected pull-force distribution for each site. (D): Using this grasp quality information, the *tension planner* determines the tensions on each boom to safely achieve the target task wrench.

of a highly reconfigurable CSR, and is likely the best direct comparison to ReachBot. However, (i) SpiderBot is not fully-constrained, having only 4 cables, and (ii) the authors only study mobility planning rather than using it for manipulation. Others have considered a task-specific optimization for CDPRs [14] with similar objectives such as maximizing the tension-closure workspace volume, or minimizing the cable force. However, this analysis is limited to a 3-DOF universal joint and does not consider reconfigurability.

Due to ReachBot’s similarities to dexterous manipulation (trading fingers that push for booms that pull) [8], we can also look at grasp optimization, much of which builds on the work of [15]. Of particular interest are methods for planning over a set of discrete candidate grasp points [16] and computing optimal contact forces [17], which are highly applicable to the planning methods presented here.

Our past work [6], [8] has also explored a similar motion planning problem: given a starting and ending position of the ReachBot body and potential grasp sites, plan a dynamically feasible trajectory and a sequence of stances for minimizing control effort. However, this motion-focused approach does not consider the use of ReachBot for specific manipulation tasks, which typically require a different stance and tensioning than what is produced by the motion planner.

Statement of Contributions: This paper presents a manipulation planning method for ReachBot, optimizing for task robustness under uncertain parameters and stochastic failure modes. We establish a two-part architecture: a *stance planner* and a *tension planner*, each inspired by work in the field of dexterous grasping. We present an analysis of how these planners improve performance on relevant tasks, and increase the size of the manipulation workspace. These methods are compared against a baseline method in a Monte-Carlo-based simulation. We also present insight into how these methods perform on similar cable-driven robot architectures. Additional media and code is available at our project webpage, https://stanfordasl.github.io/reachbot_manipulation/

II. PLANNING FRAMEWORK

For ReachBot to execute a manipulation task, we consider a multi-stage scenario (Fig. 2) where ReachBot will first

observe the environment, select the grasp sites to place its booms, and then tension its booms to complete the task. Here, we define a “task” as a sequence of pose/wrench pairs: a pose in $SE(3)$, and a wrench in \mathbb{R}^6 (a combination of a force and torque). A single-pose task in the context of ReachBot may imply rejecting disturbances while performing imaging of the environment, and a multi-pose task may involve a pick-and-place task, such as placing sensors or instruments.

To account for these stages, and the different information obtained by the perception system at each stage, we structure the manipulation planner as a two-part system. The *stance planner* uses the perception of the environment (candidate grasp sites and estimated quality of each site) and the task (estimated positioning of the robot and the required wrench/wrenches) to plan a ReachBot stance which can complete the task while being robust against uncertainty and disturbances. Then, the *tension planner* uses the perception of the grasps (their expected pull force distribution based on the microspine gripper model [8]) to safely tension the booms for the desired task without causing a grasp failure.

A. Assumptions

To guide the experiments, we assume the following:

- 1) ReachBot is located in an environment where it can span the full width of the space with its booms, such as a lava tube or cave.
- 2) We consider an 8-boom configuration, which has been previously shown to satisfy capability constraints while minimizing mass and interference [5].
- 3) Manipulation tasks are performed with an end effector on the base of the robot.
- 4) All booms are primarily loaded in tension but can support a small bending moment at the shoulder joint. When comparing to CDPRs, we consider cables as pure-tension members.
- 5) Booms and cables are massless compared to the body.

We also use the following specifications in our experiments: a ReachBot mass of 10 kg, a maximum gripper pull force of 30N, a maximum boom shoulder moment of 1 Nm, and an environment subject to Martian gravity ($g = 3.71m/s^2$).




Fig. 3. **Relationship between the ReachBot stance, wrench space, and task polytope.** The X/Y force space, a subset of the wrench space for a 2D ReachBot, is plotted, along with the task polytope and the task ellipsoid. (A): With no specified task, this ReachBot stance can resist arbitrary disturbance forces about the origin of the force space, with magnitudes up to the radius of the ball. (B): With a specified task ellipsoid or polytope, this stance can successfully execute the task even with some uncertainty, since the center of the ellipsoid/polytope is contained in the hull. (C): The task-ellipsoid/polytope metric can also indicate that a task is feasible, even when ReachBot is in a stance which is notably *not* in force closure (i.e. the origin of the space is not contained in the interior of the hull).

Note that this method can generalize to a parallel robot with any number of cables or booms, even if we have fewer than 7 (the minimum number required for closure in \mathbb{R}^3 with unidirectional forces). This method also generalizes to a manipulation task performed at the end of an arm or boom: to do so simply requires a transformation of the resultant wrench back to the base of the robot.

B. Stance planner

1) *Preliminaries:* To select the ReachBot stance that maximizes the probability of a successful manipulation task, we leverage tools from dexterous manipulation: namely, wrench spaces, task ellipsoids, and task polytopes [15], [18], [19].

Classically, a grasp wrench space is a subset of \mathbb{R}^6 corresponding to the set of all possible wrenches (forces and torques) that can be applied by a set of point contact forces in a dexterous grasp. We generalize this notion to our setting by considering the set of all possible wrenches that can be applied by the ReachBot body, given its current stance and the maximum force that can be applied along each boom. We use the L_∞ definition from [15], with a modification to describe the unidirectional tensile force from a cable or boom, rather than a cone or pyramid of possible forces from a point contact with friction (Eq. 1). Here, \oplus denotes the Minkowski sum between the sets of applied wrenches by each boom (each being a line through wrench space), and w_i denotes the maximum wrench from boom i .

$$W_{L_\infty} = \text{ConvexHull} \left(\bigoplus_{i=1}^n \{\mathbf{0}, w_i\} \right) \quad (1)$$

Within this wrench space, we can define a task ellipsoid: a set of expected wrenches during the manipulation task,

parameterized by an ellipsoid in \mathbb{R}^6 [18]. The sizing of this ellipsoid can be selected to best reflect the task: for instance, in our sample extraction task example (Fig. 1), lifting the sample will require forces distributed primarily along the z direction, and we also anticipate some small disturbance forces and torques along other directions. The corresponding ellipsoid, therefore, would be primarily distributed about z .

However, the task ellipsoid is difficult to efficiently optimize for. Instead, we use a task polytope, an l_1 approximation to the ellipsoid via its principal axes [19]. We parameterize this task polytope (Eq. 2) by a set of task basis wrenches B_1, \dots, B_p (each in \mathbb{R}^6) and weights $\sigma_1, \dots, \sigma_p$. The basis B dictates the principal components of the task (typically left as an identity), and the weights σ dictate the relative size of the polytope along each basis direction (typically set to the standard deviations of the task in each wrench component).

$$\mathcal{P}_{task} = \text{ConvexHull}(\{\sigma_1 \cdot B_1, \dots, \sigma_p \cdot B_p\}) \quad (2)$$

Using a task polytope also guarantees that all wrenches within its positive spanning set are possible with the given ReachBot stance. Therefore, by selecting an appropriate set of parameters B and σ , we can guarantee that the task wrench and any uncertain disturbance wrenches are contained in the task polytope and thus can be achieved by the ReachBot stance. This method's independence from force closure enables it to work even when the task is in a difficult location (Fig. 3 C), or if there are a limited number of grasp sites available. This also leads to favorable worst-case outcomes, even in the event of grasp failure and subsequent closure loss.

2) *Method:* Given a pose and a task ellipsoid defining the distribution of wrenches, we optimize based on the smallest-

magnitude disturbance wrench within the task polytope which can be resisted by the ReachBot stance (Eq. 3). In general, this objective seeks to maximize the size of the wrench space about the task ellipsoid: an “inverse Chebyshev ball” problem, where we modify the polyhedron to best fit a ball or ellipsoid. Empirically, optimizing via the task polytope correlates very strongly with the true task ellipsoid.

We also enforce the following constraints:

- (4) *Unique boom assignment*: Each boom can be assigned to at most one grasp site.
- (5) *Unique site assignment*: Each grasp site can be assigned to at most one boom.
- (6) *Wrench space scaling*: The size of the wrench space along each task basis wrench is non-negative.
- (7) *Boom non-interference*: The directions of the booms are each restricted to lie in an outward-facing second-order cone, to prevent self-intersection with the robot body.
- (8) *Tension limits*: Each boom’s tension is positive and limited by the maximum gripper force, the quality of the grasp site, and the Boolean assignment variable.
- (9) *Force-wrench space relationship*: This constraint relates the boom forces to the size of the wrench space about the nominal desired wrench.

The stance planner is a mixed-integer convex program (MICP), and can be expressed as follows. Here, A is the Boolean boom-to-site assignment matrix, s defines the maximum distance in the wrench space along each basis direction, D is the precomputed set of possible directions between booms and sites, N and θ parameterize the direction and size of the boom reachability cones, T is the set of boom tensions, t_{\max} is the maximum boom tension, q is an estimated quality parameter for each site, w is a resultant wrench component, w_{des} is the desired wrench, and B is the set of task basis vectors. \odot denotes the Hadamard (element-wise) product, and (i, j, k) index over all booms, sites, and task basis vectors, respectively.

$$\text{maximize } \min(\sigma^{-1} \odot s) \quad (3)$$

$$\text{subject to } A\mathbf{1}_n \leq \mathbf{1}_m \quad (4)$$

$$A^T \mathbf{1}_m \leq \mathbf{1}_n \quad (5)$$

$$s \geq \mathbf{0} \quad (6)$$

$$D_{ij} \cdot N_i A_{ij} \geq \cos(\theta) A_{ij} \quad \forall i, j \quad (7)$$

$$\mathbf{0} \leq T_{ijk} \leq t_{\max} q_j A_{ij} \quad \forall i, j, k \quad (8)$$

$$\sum_{i=1}^m \sum_{j=1}^n w_{ijk} = w_{\text{des}} + s_k B_k \quad \forall k \quad (9)$$

3) *Handling pose uncertainty*: The task polytope objective inherently optimizes for robustness against wrench uncertainty, yet pose uncertainty is also inherent to this planning problem. For instance, we may estimate a nominal pre-grasp pose for a pick-and-place task, but after approaching the object, it may be better to slightly adjust the pose to best grasp it.

When this occurs, the off-nominal pose will shift the directions of each boom, changing the resultant wrench on

the robot’s body. Often, these uncertain poses can lead to a dramatic shift in the facets of the nominal wrench space, leading to a potentially much smaller task ellipsoid, or in the worst case, geometric infeasibility of the task. Note that this sensitivity of the wrench space to pose uncertainty is dependent on the configuration of the robot and thus cannot be known prior to evaluating the stance planner.

Unfortunately, expanding the domain of the problem to account for uncertainty in boom directions leads to computational intractability within reasonable planning times. Despite this, we propose a simple method of handling pose uncertainty that integrates easily with our existing wrench-uncertainty-focused planner.

We rely on the empirical observations that (i) orientation errors are the primary contributor to geometric infeasibility of a target wrench, and (ii) stances robust to torque uncertainty are also robust against orientation uncertainty. Using these observations, we can re-scale the torque component of the task polytope, increasing the weighting along these dimensions as a function of orientation uncertainty. Note that when orientation uncertainty is 0, we leave this weighting, λ as 1. A simple heuristic that accounts for this is $\lambda = 1 + \|\sigma_F\| \|\Delta\Theta\| / \|\sigma_\tau\|$, where $\|\Delta\Theta\|$ is the estimated average magnitude of the angular error, and $(\|\sigma_F\|, \|\sigma_\tau\|)$ are the norms of the force and torque components of the nominal task polytope. As an example, if we expect rotational errors up to 10° in any axis for Task B in Fig. 4, this scaling factor λ equals 4.3, increasing the torque weighting and bringing the task polytope closer to the ball metric shown in Task A.

4) *Multi-pose extension*: If the manipulation task calls for a sequence of wrenches applied at different poses, we expand the dimensionality of the problem proportionally to the number of poses. For instance, a pick-and-place task may be defined by pose/wrench pairs at four key points: pre-grasp, post-grasp, pre-place, and post-place. Additional pose/wrench pairs can also be added at intermediate positions for long-horizon trajectories. Given this, we then optimize for the smallest-magnitude disturbance wrench across all of these pose/wrench pairs in the trajectory.

C. Tension planner

Once the stance of ReachBot is fixed, we then determine the tensions in each boom to apply the desired task wrench. For our objective, we employ the probability of success metric from our previous work [6] (Eq. 10), based on an estimated force distribution which is generated through a learned model of the grasp quality [8].

This convex program (CP) is subject to two constraints:

(11) *Tension limits*: Each boom tension is positive and limited by the maximum gripper force and the quality of the grasp site

(12) *Task execution*: The resultant wrench from all booms equals the desired task wrench.

This problem can be expressed as follows, where Φ represents the Gaussian CDF, μ_i, σ_i represent the estimated pull force normal distribution parameters for the i^{th} grasp, and t is the boom tension variable. As with the stance

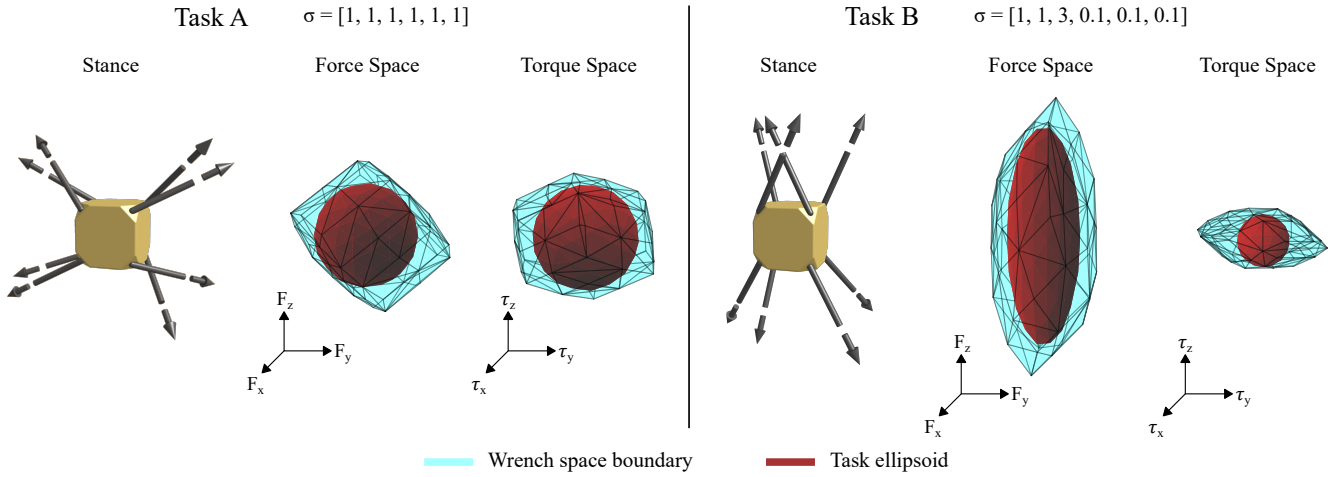


Fig. 4. **Impact of the task-based stance optimization on the ReachBot wrench space.** Here, we show the optimal ReachBot stances and corresponding wrench space (decomposed into the spaces of pure forces and torques) for two different tasks. By changing the task definition from a ball (left) to an ellipsoid biased towards forces along the z direction (right), the *stance planner* adjusts the wrench space to maximize the size of the inscribed task ellipsoid. This enables ReachBot to be robust against uncertainty in the wrench or set of wrenches experienced during manipulation.

planner, t_{\max} is the maximum boom tension, q is the site quality parameter, w_i is the resultant wrench from the i^{th} boom, and w_{des} is desired wrench:

$$\text{maximize} \quad \sum_{i=1}^n \log \left[\Phi \left(\frac{\mu_i - t_i}{\sigma_i} \right) \right] \quad (10)$$

$$\text{subject to} \quad \mathbf{0} \leq t \leq t_{\max} \cdot q \quad (11)$$

$$\sum_{i=1}^m w_i = w_{\text{des}} \quad (12)$$

III. EXPERIMENTS AND RESULTS

A. Naïve baseline

As a point of comparison for the stance planner, we also develop a naïve baseline method, which extends out the cables (Fig. 5) and guarantees that the nominal wrench can be achieved. This method retains many of the same constraints as in the stance planner, but makes no assumptions about robustness to off-nominal wrenches or poses. Notation follows from the stance planner.

$$\begin{aligned} &\text{maximize} \quad \sum_i \sum_j A_{ij} D_{ij} \cdot N_{ij} \\ &\text{subject to} \quad A \mathbf{1}_n \leq \mathbf{1}_m \\ &\quad \quad \quad A^T \mathbf{1}_m \leq \mathbf{1}_n \\ &\quad \quad \quad D_{ij} \cdot N_i A_{ij} \geq \cos(\theta) A_{ij} \quad \forall i, j \\ &\quad \quad \quad \mathbf{0} \leq T_{ij} \leq t_{\max} q_j A_{ij} \quad \forall i, j \\ &\quad \quad \quad \sum_{i=1}^m \sum_{j=1}^n w_{ij} = w_{\text{des}} \end{aligned}$$

B. Computation times

The stance planner, solved with MOSEK via the CVXPY interface, typically finds the optimal solution within 3 seconds for small problems, and around 30 seconds for large problems (see Table I). Here, problem size is dictated by the number of feasible sites for each boom – a function of the total number of sites, and the cone direction constraint. This compute time is reasonable given ReachBot’s speed.




Fig. 5. **Comparison of two possible ReachBot stances for a set of possible sites.** Left: a *naïve* configuration, simply directing the booms outwards from the body. This configuration has good control over the applied force, but is very weak in applying torque, since each boom has no effective lever arm on the body. Right: an *optimal* configuration, with good control over both applied force and torque.

The tension planner, solved with ECOS via the CVXPY interface, typically solves in about 1 millisecond. This is always a fixed problem size, and can be efficiently re-solved online for various applied wrenches.

The naïve planner, also solved with MOSEK, is still a MICP like the stance planner, but its reduced dimensionality allows for significantly faster computation times.

All timings are based on a desktop computer with an Intel Xeon E5-2643 CPU at 3.40GHz, with 64 GB RAM.

TABLE I
AVERAGE COMPUTE TIMES (SECONDS)

	Stance planner		Tension planner
	Optimal	Naïve	
Sites 10	2.24	0.204	0.0014
Sites 20	8.10	0.244	0.0014
Sites 30	29.3	0.292	0.0014




Fig. 6. **Impact of stance optimization on workspace size.** Left: manipulation workspace for cable-driven and boom-driven parallel robots, in a naïve stance. Right: the manipulation workspace for each robot, but in an optimal stance. This stance optimization leads to significant increases in the workspace volume, and reduces the downsides of a lower-degree-of-freedom cable-driven architecture.

C. Evaluation of optimal stance planning on workspace size

The capability of applying a target wrench is only one part of the manipulation problem: another key requirement is the workspace size. Our previous ReachBot works have already highlighted the large workspace size for ReachBot [4], but this concerns a *reachability* workspace, rather than a *manipulation* workspace. While ReachBot may be able to access a large space with the existing motion planner, there are no guarantees that it can apply the wrenches needed for a manipulation task. In general, we need to see that ReachBot can both access a location (the location is in the *Static Equilibrium Workspace*) and apply the desired wrench (the location is in the *Wrench Closure Workspace*, for that specific wrench) [20].

Therefore, to evaluate the impact of our planner on the workspace, we compare the manipulation workspaces between a *naïve* stance and an *optimal* stance (See Fig. 5 for a comparison of these stances). Additionally, we look at how the workspace varies between a cable-driven and a boom-driven ReachBot, to help guide the system design as well as demonstrate the applicability of this method to other reconfigurable parallel robots.

As seen in Fig. 6, in the naïve stance, the cable-driven ReachBot is significantly hindered by the suboptimal cable placement, primarily due to the lack of torque authority in this stance. As purely tensile members, cables must rely solely on the configuration’s geometry to achieve a target wrench, whereas booms can passively resist small torques through the reaction torques at their shoulder joints. This can be seen via the torque component of the robot Jacobian: the geometry of a cable-driven robot may lead to this being rank-deficient, while the added degrees of freedom from a boom-driven robot can mitigate this issue. This leads to a relatively large workspace for booms (Fig. 6), even with a suboptimal stance.

However, if we use the stance planner, this leads to a dramatic improvement in the workspace volume. Both the boom-driven and cable-driven robots see an improvement, though this increase is much more significant for the cable-driven ReachBot. So, while cables are in general less capable than booms, smart planning of the cable positioning mitigates

these downsides and leads to comparable performance.

D. Validation in probabilistic simulation

Our optimization objectives for the stance planner and the tension planner are designed to maximize robustness and safety, but to validate these, we must evaluate the “true” simulated task robustness of these plans. Specifically, we want to show that by using these proposed methods, ReachBot can: (i) apply the required wrench (or set of wrenches) for the manipulation task, even when these can differ from the original plan, (ii) maintain stability even under unmodeled vibrational effects or disturbances, and (iii) execute the task in a manner which leads to a low chance of grasp failure.

To show this, we performed a Monte-Carlo validation of the task robustness in simulation. For 1000 sampled tasks across 10 randomized environments, we determine if (1) the task is geometrically feasible given the planned ReachBot stance, and if so, (2) tension the ReachBot booms and sample from the stochastic grasp site model to determine if a grasp failure occurred. We denote a failure at step (1) as a geometric failure, and a stochastic failure at step (2). Note that stochastic failure is only evaluated if the problem is geometrically feasible (i.e. the desired wrench or set of wrenches lie within the wrench space at each pose).

As seen in Fig. 7, while the naïve planner is guaranteed to support the nominal task, its lack of uncertainty handling leads to poor performance on off-nominal poses and wrenches. Notably, the main failure mode for the naïve planner is geometric: i.e., given the ReachBot stance, there is no set of boom/cable tensions that will achieve the target wrench. The primary reason for this failure tends to be an inability to produce the required torque for the task, even if this is small. This is especially an issue with cables, which cannot passively resist small torques, unlike booms. This passive torque resistance allows a naïve stance with booms to perform adequately on simple tasks, though as the task complexity increases, performance degrades significantly.

Conversely, our optimal stance planner demonstrates high task performance across both simple and complex tasks, and the stance optimization is particularly critical for these more




Fig. 7. **Task performance results**, evaluated across randomized environments and tasks in Monte-Carlo simulation. We consider a representative single-pose task as holding a grasped object steady while rejecting arbitrary disturbances, and a multi-pose task as picking-and-placing an object between two locations in the environment. A geometric failure indicates that the task wrench is infeasible given the current stance, and a stochastic failure indicates that a commanded tension exceeds a sampled maximum force from a probabilistic model.

complex sequences of wrenches and poses. We also see that stochastic failure is generally always present, due to the nature of microspine grasping, but a good configuration can also mitigate the chance of failure. For example, a naïve stance might evenly distribute tensile loads across all sites for a nominal wrench, but rejecting a disturbance may require a very large tensile force applied to a single boom. An optimal stance under the same disturbance will be able to distribute the added loading across multiple booms, reducing the chance of failure.

IV. CONCLUSION

In this work, we have established a two-part optimization-based planning method for robust manipulation with ReachBot, leveraging concepts from dexterous grasping and cable-driven parallel robots. First, the stance planner selects the placement of ReachBot’s booms through a mixed-integer convex program, adjusting the wrench space to improve a task-polytope-based robustness metric. Second, the tension planner determines the forces in each boom to apply a desired task wrench, solving a fast convex program to minimize the chance of a grasp failure.

Through our experiments, we’ve shown that these optimizations significantly increase the size of the manipulation workspace, and lead to robust task execution even under disturbances or uncertainty in the desired poses or wrenches required for the task. Given the high reliability of these methods, this new capability will motivate new missions and scientific experiments on the Moon, Mars, and beyond.

Future work in this area includes integrating these optimal manipulation stances into the ReachBot motion planner, such that the final stance from the motion plan is the optimal stance for the given manipulation task. Second, construction on new reconfigurable cable-driven ReachBot hardware is underway, and experiments on this platform will further demonstrate the need for these planning methods.

REFERENCES

- [1] R. J. L evill e and S. Datta, “Lava tubes and basaltic caves as astrobiological targets on earth and mars: A review,” *Planetary and Space Science*, vol. 58, no. 4, pp. 592–598, 2010.
- [2] N. J. Fox, “Science 2020-2024: A vision for scientific excellence,” tech. rep., NASA Science Mission Directorate, Sept. 2023.
- [3] N. Jacobstein, “Nasa’s perseverance: Robot laboratory on mars,” *Science Robotics*, vol. 6, no. 52, p. eabh3167, 2021.
- [4] S. Schneider, A. Bylard, T. G. Chen, P. Wang, M. Cutkosky, and M. Pavone, “Reachbot: A small robot for large mobile manipulation tasks,” in *2022 IEEE Aerospace Conference*, IEEE, 2022.
- [5] S. Newdick, T. G. Chen, B. Hockman, E. Schmerling, M. R. Cutkosky, and M. Pavone, “Designing reachbot: System design process with a case study of a martian lava tube mission,” in *2023 IEEE Aerospace Conference*, pp. 1–9, 2023.
- [6] S. Newdick, N. Ongole, T. G. Chen, E. Schmerling, M. R. Cutkosky, and M. Pavone, “Motion planning for a climbing robot with stochastic grasps,” in *IEEE ICRA 2023*, pp. 11838–11844, 2023.
- [7] 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 *IEEE ICRA 2022*, pp. 4517–4523, 2022.
- [8] T. G. Chen, S. Newdick, J. Di, C. Bosio, N. Ongole, M. Lap tre, M. Pavone, and M. R. Cutkosky, “Locomotion as manipulation with reachbot,” *Science Robotics*, vol. 9, no. 89, p. eadi9762, 2024.
- [9] S. Qian, B. Zi, W.-W. Shang, and Q.-S. Xu, “A review on cable-driven parallel robots,” *Chinese Journal of Mechanical Engineering*, vol. 31, p. 66, Aug 2018.
- [10] A. Pott, H. Mutherich, W. Kraus, V. Schmidt, P. Miermeister, and A. Verl, *IPAnema: A family of Cable-Driven Parallel Robots for Industrial Applications*, pp. 119–134. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013.
- [11] L. Cone, “Skycam: An aerial robotic camera system,” *Byte*, vol. 10, no. 10, pp. 122–132, 1985.
- [12] D. Q. Nguyen, *On the Study of Large-Dimension Reconfigurable Cable-Driven Parallel Robots*. Theses, Universit  Montpellier 2, Dec. 2014.
- [13] A. Capua, A. Shapiro, and S. Shoval, “Spiderbot: a cable-suspended walking robot,” *Mechanism and Machine Theory*, vol. 82, pp. 56–70, 2014.
- [14] D. Lau, K. Bhalerao, D. Oetomo, and S. K. Halgamuge, “On the task specific evaluation and optimisation of cable-driven manipulators,” in *Advances in Reconfigurable Mechanisms and Robots I* (J. S. Dai, M. Zoppi, and X. Kong, eds.), (London), pp. 707–716, Springer London, 2012.
- [15] C. Ferrari and J. Canny, “Planning optimal grasps,” in *IEEE ICRA 1992*, pp. 2290–2295 vol.3, 1992.
- [16] A. Cauligi, P. Culbertson, B. Stellato, D. Bertsimas, M. Schwager, and M. Pavone, “Learning mixed-integer convex optimization strategies for robot planning and control,” in *IEEE CDC 2020*, pp. 1698–1705, 2020.
- [17] S. P. Boyd and B. Wegbreit, “Fast computation of optimal contact forces,” *IEEE T-RO*, vol. 23, no. 6, pp. 1117–1132, 2007.
- [18] Z. Li and S. Sastry, “Task-oriented optimal grasping by multifingered robot hands,” *IEEE Journal on Robotics and Automation*, vol. 4, no. 1, pp. 32–44, 1988.
- [19] R. Haschke, J. Steil, I. Steuwer, and H. Ritter, “Task-oriented quality measures for dextrous grasping,” in *IEEE CIRA 2005*, pp. 689–694, 2005.
- [20] M. Gouttefarde, J.-P. Merlet, and D. Daney, “Determination of the wrench-closure workspace of 6-dof parallel cable-driven mechanisms,” in *Advances in Robot Kinematics*, (Dordrecht), pp. 315–322, Springer Netherlands, 2006.