Short Papers

Navigation for Legged Mobility: Dynamic Climbing

Max P. Austin, Mario Y. Harper, Jason M. Brown, Emmanuel G. Collins, and Jonathan E. Clark

***Abstract*—Autonomous navigation through unstructured terrains has been most effectively demonstrated by animals, who utilize a large set of locomotive styles to move through their native habitats. While legged robots have recently demonstrated several of these locomotion modalities (such as walking, running, jumping, and climbing vertical walls), motion planners have yet to be able to leverage these unique mobility characteristics. In this article, we outline some of the specific motion planning challenges faced when attempting to plan for legged systems with dynamic gaits, with specific instances of these demonstrated by the dynamic climbing platform TAILS. Using a unique implementation of sampling-based model predictive optimization, we demonstrate the ability to motion plan around obstacles on vertical walls and experimentally demonstrate this on TAILS by navigating through traditionally difficult narrow gap problems.**

***Index Terms*—Climbing robots, legged locomotion, path planning.**

I. Introduction

Legged robotic platforms using both quasi-static (capable of continuously enforcing stability criterions during movement) and dynamic behaviors (operating at speeds that intrinsically reduce the control authority and primarily rely on conservation of momentum for stability) have demonstrated a vast array of distinct locomotion modalities, specific examples include walking, running, jumping, crawling, and wall climbing [1]–[4]. While these modalities provide a wide variety of navigation options, current motion planning techniques have yet to leverage these unique modes. While having enabled impressive performance, motion planners generally map legged locomotion to the limited capabilities of traditional wheeled and tracked vehicles [5] or constrain the behaviors to quasi-static or enforceably stable motions [6]–[9].

These algorithms used in long horizon motion planning (navigation) are typically based on sampling-based methods [10], heuristic search approaches [11], [12], or a combination of these. Sampling-based algorithms have shown great versatility in handling kinodynamic constraints [13], allowing them to plan for nonholonomic motions and unique constraints, but they generally cannot handle the constrained computation times of dynamic legged locomotion (on the order of 10 ms), which the heuristic-based planning methods such as those based on the A* algorithm enable. However, heuristic search planners require understanding of robot behavior to formulate an effective cost function that will guide the search toward the goal, something that is difficult due to the complexity of legged systems employing dynamic behaviors.

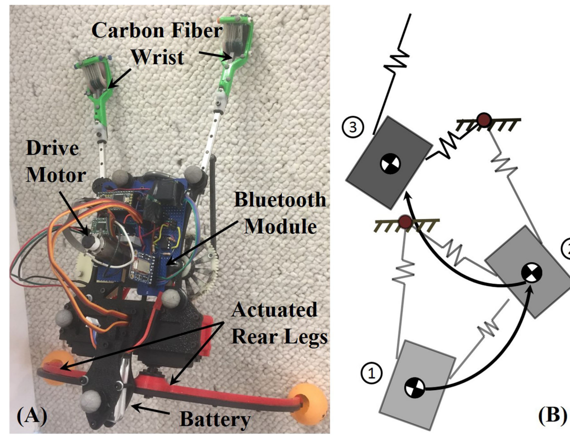

Fig. 1. (a) TAILS: one of the fastest legged climbing robots. (b) Example of Full–Goldman style dynamic climbing utilized by TAILS showing pendular COM trajectory and heading angle deviation (exaggerated for clarity).

The complex motion characteristics of legged locomotion, especially dynamic locomotion that has brief flight phases, present some fundamental opportunities and challenges to motion planning that are largely yet to be addressed. These challenges include motion impediment versus foot placement constraints, postural changes (specifically enforced body shape changes), unusual (both holonomic and nonholonomic) motion constraints, and discrete state estimation, and control authority.

The challenges described above, inherent to planning for a wide variety of legged locomotion, are demonstrated explicitly by dynamic vertical climbing, which imposes very unusual foot contact and motion constraints. Navigation within vertical climbing has been limited to quasi-static foot-step planning [14]–[16], primarily because of the limited maneuverability of dynamic climbers. Recently, however, maneuverability has been demonstrated on the dynamic climbing platform TAILS, shown in Fig. 1, which is a member of the family of the fastest legged climbing robots (capable of climbing upward at over 1.5 body lengths per second). TAILS, using a single drive motor to actuate both front legs, is able to instantiate the animal inspired Full–Goldman climbing dynamics [17]. The Full–Goldman model captures the pendularlike oscillations [see Fig. 1(b)] characteristic of rapidly climbing animals. This style of climbing naturally provides passive self-stabilizing of heading angle using entirely feedforward control. Maneuverability

Manuscript received April 1, 2019; accepted November 16, 2019. This work was supported in part by Cooperative Agreement DAAD 19-01-2-0012 and in part by the National Science Foundation under Grant CMMI-1351524. This paper was recommended for publication by Associate Editor R. Hatton and Editor E. Yoshida upon evaluation of the reviewers' comments. *(Corresponding author: Max P. Austin.)*

M. P. Austin, M. Y. Harper, J. M. Brown, and J. E. Clark are with the FAMU/FSU College of Engineering, Tallahassee, FL 32310-6046 USA (e-mail: mpa12c@my.fsu.edu; myh15@fsu.edu; jmb10t@my.fsu.edu; jeclark@fsu.edu).

E. G. Collins is with the J. B. Speed School of Engineering, University of Louisville, Louisville, KY 40292 USA (e-mail: ecollins@eng.fsu.edu).

This article has supplementary downloadable material available at http://ieeexplore.ieee.org, provided by the authors. The material consists of a video. The size of the video is 6.87 MB. Contact mpa12c@my.fsu.edu and jeclark@fsu.edu for more details.

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TRO.2019.2958207

was achieved via upward strafing and dynamic downward locomotion using the added rear leg actuation, which enables independent relative body roll and pitch control [18]. Thus, creating the minimal set of behaviors necessary to navigate through an arbitrary obstacle field. However, no existing motion planner is capable of utilizing the full motion capabilities and handling the unique motion constraints imposed by a platform like TAILS.

Inspired by this, we describe a hybrid heuristic and sampling-based planning method, utilizing a unique implementation of sampling-based model predictive optimization (SBMPO) [19], to motion plan, for the first time, for a dynamic legged platform on a vertical wall. In Section II, we detail fundamental challenges for the planning of dynamic legged robots. Next, in Section III, we describe TAILS and how these challenges are manifested on the platform. Our algorithmic advances to address these challenges are given in Section IV. In Sections V and VI, we demonstrate the robot maneuvering through a set of difficult navigation problems. Finally, Section VII concludes this article.

II. Navigation for Gaited Legged Locomotion

Dynamic gaited locomotion (continuous locomotion that is instantaneously unstable but cyclically stable) has been utilized by both biology and robotics to generate fast efficient motions, but impose numerous challenges to motion planning. This section defines a set of characteristic challenges many legged systems present toward navigating utilizing the system's full host of behaviors. Some of these have been individually addressed, but this set has yet to be completely addressed.

A. Discrete Control Authority

While moving quasi-statically, legged platforms have the ability, like traditional vehicles, to instantaneously change their plan or behavior. However, when moving dynamically, legged platforms generally experience very brief discrete stance and flight phases (where all the feet are no longer in contact with the surfaces). The platform can only significantly adjust its behavior during the short contact period. Discontinuity in control authority suggests the use of discrete measurement and control for planning. Since legged systems often have periodic body dynamics and frequent and irregular perturbations when moving rapidly, state estimation may require special consideration. Additionally, impacts from discrete foot touchdown events can cause significant perturbations to body motion.

B. Postural Change

During locomotion, the effective shape of the body is persistently changing as a result of either bending legs for locomotion, an active change to the body shape (i.e., crouching down), and/or natural body oscillations (which tend to be more significant for dynamic motions). Collision detection is impacted by all of these cases, but active posture control inherent to legged platforms has the potential to be a boon to navigation in cluttered environments as demonstrated by cockroaches using their ability to contort their bodies to fit through narrow gaps [20]. Similar postural changes can be used to improve energy efficiency and speed, such as choosing a high stepping gait to rapidly traverse shallow water.

C. Platform Specific Motion Constraints

While all vehicles have some specific motion constraints, such as how wheeled vehicles generally lack the capability to move laterally, legged robots exhibit more heading specific behaviors, including asymmetries in locomotion based on direction of travel (e.g., going forward versus backward). For example, due to the direction of their claws, some cats are capable of climbing up a tree but not down, and therefore must chose a different behavior to get down, such as jumping. Additionally, dynamic legged robots have a tendency to experience significant transient periods when changing directions or switching locomotion modes.

D. Motion Impediment Versus Foot Placement Constraints

Consider a robot walking across a bridge made up of two parallel, wooden planks, the robot's body is freely traversing over the open air between the planks but the feet may not be placed within the gap. Contrary to this suppose a robot with waterproof legs is traversing through shallow water allowing its feet to be submerged but incapable of submerging its body. These two examples highlight situations where considering the foot and the body separately (thinking about motion impediments versus foot placement constraints) could enable planning in more complex environments. In legged locomotion the propulsion mechanisms (the feet) are discretely placed and decentralized from a primarily passive body. The relative motion between these objects suggests that the feet and body can be treated as separate, but coupled, systems with distinct constraints.

III. Platform

A. TAILS: Design and Control

TAILS, shown in Fig. 1(a) [18], is a climbing robot, which demonstrates rapid (40 cm/s) ascension of vertical walls and instantiates the Full–Goldman climbing dynamics. It does this by linearly retracting its forefeet through a sinusoidal extension trajectory locked 180° out of phase using a single brushed dc motor. Directional adhesion [21] to the substrate is achieved through microspine arrays contained within the forefeet, which facilitate the rapid attachment and detachment necessary for dynamic climbing. The forefeet are connected to the body through a spring-loaded sliding wrist mechanism, which mitigates peaks in the wall reaction forces.

To enable autonomous operation, some adaptations were made to TAILS. The sliding wrists, previously made of 3-D printed acrylonitrile butadiene styrene (ABS) plastic, have been replaced with carbon fiber shafts in order to manage the higher strain seen by strafing and downward climbing gaits. An Adafruit Bluefruit LE bluetooth module was added to enable offboard communication with an external path planning computer. Additionally, a three-cell lithium polymer battery was added to facilitate untethered locomotion.

Low-level control over the robot's climbing maneuvers, which produces the collection of behaviors that the motion planner can select for navigation, is divided into a feedforward forelimb controller that maintains a fixed retraction frequency and a kinematic body orientation controller. The body orientation control (described in more detail in [18]) prescribes the servo angle of the rear legs to set pitch [which allows for transposition between body size in plane to out of plane (OOP)] and roll (which can induce strafing with a constant nonzero value or downward climbing with a sinusoidal trajectory).

B. Planning Challenges Exhibited by TAILS

While TAILS is the first dynamic legged robot capable of maneuvering on a vertical wall, its physical design and control impose several constraints that make planning navigation difficult. Specific instances of the four broad classes of challenges outlined in Section II are given as follows.

1) Attachment Constraint Challenges: For TAILS' feet, the use of directional adhesion imparts both a surface and heading angle constraint, as attachment can only be achieved on sufficiently rough substrates and when loaded by gravitational forces. Because of this, dynamical climbing naturally stabilizes to and requires the feet remain above the center of mass (COM), which bounds the acceptable heading angles to a narrow range with the robot oriented essentially upright.

2) Challenge of Heading Angle Constrained Behaviors: With the fixed body heading angle, maneuverability is achieved using upward climbing, downward climbing, and diagonal strafing. Upward climbing has a broad range of velocities with minimal transient behavior, whereas downward climbing is significantly more sensitive to parameter variation and has longer initial transient periods, which result in a constrained range of speeds. Because of the orientation constraints, lateral motion is constrained to strafing, which progresses upward and laterally (with upward mobility still dominating).

3) Challenge of Posture and Dynamic Coupling: The effective size of the robot varies with the speed, climbing behavior, and prescribed posture. In previous results, it has been found that the body swing magnitude is much higher at low retraction frequencies and is attenuated at high frequencies, which causes the effective lateral size of the body to increase or decrease [18]. A specific example of this is the increased pitch required for strafing (nearly double vertical climbing), as shown in the video attachment.

4) State Estimation Challenges: During standard locomotion, the body uses pendular dynamic swings, which adds complexity to position tracking. While the average position change of a representative locomotion template shown in Fig. 1(b) is almost purely vertical, the instantaneous position of the robot's body deviates throughout the step from the prescribed linear path toward a goal.

IV. Algorithm

A perfectly well-behaved robot with known and consistent dynamics will have many options for types of motion planners as there is a clearly established relationship between the control space and the resulting motions. SBMPO was chosen as it is particularly adept at complex situations (such as the holonomic and nonholonomic constraints imposed by TAILS) by planning directly in the control space using a propagation relationship to map controls to whole-body velocities, which alleviates the need for an invertible model.

A. Sampling-Based Model Predictive Optimization

SBMPO uses kinematic, dynamic, or power models of a system to predict behavior (a model-centric approach) and takes advantage of both heuristic-based and sampling-based search methods to intelligently plan an optimal trajectory (within sampling resolution). This algorithm is based on lifelong-planning A* (LPA*) and shares its guarantees on completeness and optimality. Upon termination, SBMPO will produce a sequence of nodes representing a minimal cost trajectory among those represented by the graph.

The algorithm is comprised of three largely independent components that encompass all the problem specific components (see Fig. 2): an optimizer, sampler, and model. The optimization is conducted by using LPA* [12], which operates on a unique graph of the output space, generated in SBMPO by inserting samples of the control input into the prediction model. The model-based components, i.e., the prediction model, cost function, and heuristic function, allow the algorithm great flexibility in dealing with the specific constraints of TAILS.

Details of the SBMPO algorithm are found in prior publications [22], [23]. The general steps of SBMPO, specialized to TAILS, can be summarized as follows.

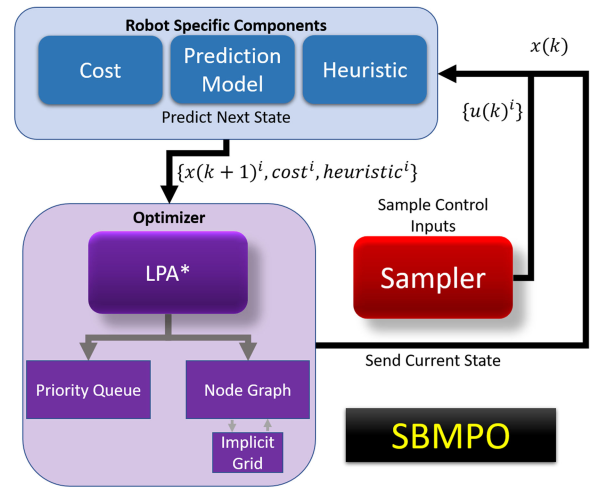

Fig. 2. SBMPO is a motion planning paradigm that uses sampling and cost-to-goal heuristic estimates. Similar to an rapidly-explored random trees (RRT)-based algorithm, a tree is generated by sampling, allowing direct use of kinematic and dynamic models to inform the planner.

1) A set $\{u(k)\}$ of feasible control inputs is generated by sampling. For TAILS, the vector u is comprised of three elements, offset (impacting strafing angle), frequency (impacting velocity), and pitch (impacting stability and velocity).
2) For TAILS, the state x is comprised of the global Cartesian position. The set of corresponding next states $\{x(k+1)\}$ in the graph is generated by propagating $\{u(k)\}$ through the prediction model that, for TAILS, is a deep feedforward network. Each edge of the directed graph represents the sampled input at a given time (say k^*) and leads to a node representing the state of the robot at time k^*+1.
3) LPA* is used to search the graph for the optimal route. All A*-type algorithms leverage a heuristic to rapidly reduce the computation time needed for optimization; the unique constraints of TAILS greatly benefited from a specially tailored heuristic function, as described in Section IV-C.

SBMPO is favorable for motion planning on TAILS as it works with both linear and nonlinear models, samples in the input (control) space, and predicts using kinematic or dynamic models. Additionally, SBMPO plans feasible trajectories directly (not requiring post processing of a path to generate a trajectory) and efficiently (by using heuristics). Finally, it has been demonstrated to be effective in numerous applications on real-world systems with diverse motion constraints, including underwater vehicles [24], legged robots [25], [26], and spacecraft [27].

B. Learning the Model

Specific adaptation of SBMPO for the TAILS robot required the creation of an accurate motion prediction model. Although SBMPO can directly integrate an analytically determined dynamic model, to achieve efficient computations, it typically utilizes an enhanced kinematic models (i.e., kinematic models preceded by integrators) and considers constraints from the robot's dynamic model [27]–[29].

Dynamic legged robotics incur significant intrastride variability in COM position. These robots also have nontrivial analytic models of body motion, where high-fidelity simulations often rely on computationally expensive propagation techniques. We employed a data-driven approach to determine this motion prediction model, a deep feedforward neural network was trained from experimental data.

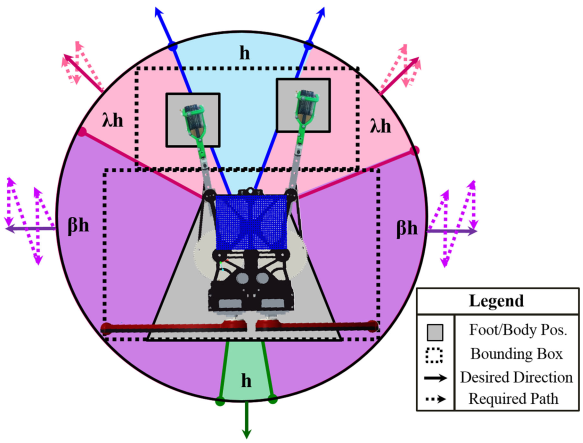

Fig. 3. TAILS mapping of directional heuristic scaling. Separate bounding boxes are used on the forelimbs and the body. The blue region is defined by the limits of robot strafing, and the traversal cost to any point therein can be reasonably approximated by Euclidean distance (h). As the angle increases, the relative estimated cost can no longer be expressed by a simple Euclidean distance and is necessarily inflated. Arrows show planned paths at different angles. The robot requires more distance and time to traverse horizontally due to the fixed heading angle.

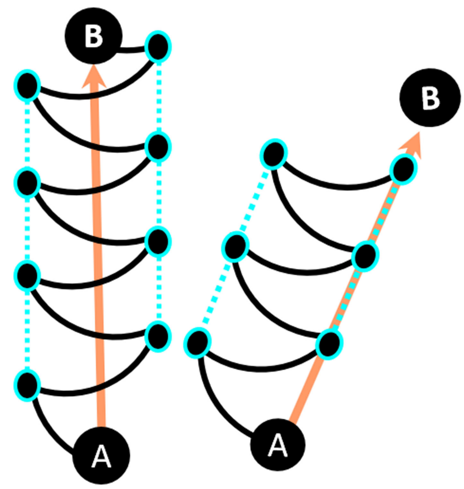

Fig. 4. Example COM trajectory of the robot (black) is divided into steps based on critical points of the body swing (cyan circles). The approximated position (brown) is centered between the critical points in vertical climbing and set to the lower bound in strafing.

In order to address the complex constraints inherent to the TAILS platform, several adaptations to the software were necessary. Particularly, an accurate motion prediction model needed to be created. In order to create this, a series of motion-data collection experiments were conducted to determine the prediction model relating the control parameters (frequency, offset, and pitch) to upward velocity V_Z and lateral velocity V_X. Incrementing through several combinations of the three control parameters, steady-state behavior was recorded for 180 motion characterization experiments. The details of these experiments are given in Section V-A.

The network was structured with two hidden layers of 15 and 8 units, respectively, and was intentionally designed to be noncomplex with the desire to be adaptable online. This simple structure has the capacity to learn without needing significant training time, and it can accommodate additional inputs and outputs without the need for increased computational resources. The stochastic gradient descent optimizer was used with a tanh activation function for both hidden layers. Learning rates were set to 0.005 and batch size was 16.

C. *Directional Heuristic and Obstacle Definitions*

Heuristic search methods such as A* require some form of approximation of the cost-to-attain-goal in order to direct the search. Directing the search allows the algorithm to find an optimal path faster than an uninformed search method. If properly chosen, SBMPO, by virtue of the heuristic, is computationally faster than most other forms of graph search and exploration algorithms [11]. The ideal heuristic function estimates the cost to goal as closely as possible without exceeding the actual cost.

Because of the unique motion characteristics of TAILS, a new cost function and optimistic heuristic function were required. Since TAILS does not have the capacity to change heading angle and can only move side to side through diagonal strafing maneuvers, lateral maneuvers are much more distance costly than others (see Fig. 3).

The heuristic functions must also take into account the relative ease of vertical mobility and the increased costs of horizontal movement. A distance-based heuristic was defined based on the Euclidean distance between the COM of TAILS and the goal coordinate. This heuristic was then inflated using the tuning parameters λ and β ($\lambda = 3$ and $\beta = 8$) based on direction of the goal region relative to the robot (shown in Fig. 3).

Another substantive change that TAILS required was in the treatment of obstacles. Some obstacles constrain the motion of the whole body, much like a traditional robot, whereas others only constrain the feet. As an example, a small void or a narrow crack in the wall does not limit traversal for the hanging body but it certainly presents a constraint for placing the feet. The body pitch also impacts the size of the body OOP, some obstacles require the robot to hug the wall as closely as possible (in the cases of crawling under an overhang).

We have chosen to treat the robot as two separate objects (shown in dotted rectangles in Fig. 3) each with their own constraints and interactions. The feet are treated in a traditional manner with obstacle detection preventing exploration of regions overlapping with the defined foot area. The body is defined with a 3-D constraint and is allowed to interpret when an obstacle does not constrain motion.

D. *Discrete Control and Dynamic-Based State Estimation*

An important consideration when developing models for dynamic legged robots is the employed gait, a cyclic pattern of leg motions, used to maneuver and maintain stable locomotion. This type of locomotion requires a fundamental shift in planning methodology from a time-based model to that of a stride-based model [5] due to the discrete nature of the motion. As discussed in Section III-B4, during a single stride, naturally occurring body oscillations (such as roll, pitch, and yaw) occur that adds significant volatility to robot state information. By limiting the data collection and control to a single point during a stride, it has been found [5] that data relevant for motion planning are preserved and less experimental data are needed to determine an accurate motion model.

For TAILS, the pendular nature of the Full–Goldman dynamics produces yaw oscillations large enough to be leveraged for defining stride information. In order to track the robot during path execution, the position is approximated using the critical points of body swing, as shown in Fig. 4. Steps are determined through a direction change in the periodic motion of the robot's yaw angle, with the boundary points

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

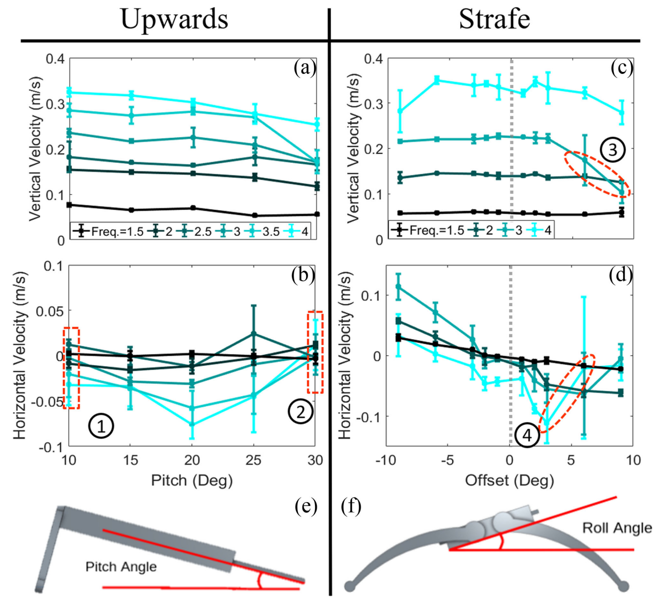

Fig. 5. Experimental average velocity data from the upward and strafing parameter sweeps in the vertical and horizontal directions. Examples of rear servo induced body pitch (e) and body roll offset (f) based on the quasi-static model are also shown.

TABLE I
PARAMETERS USED IN CLIMBING CHARACTERIZATION PARAMETER SWEEP

Parameter	Range: Upward	Range: Strafing	Units
Driving Frequency	1.5,2,**2.5**,**3**,**3.5**,**4**	**1.5**, **2**, 3, 4	*hz*
Pitch	**10**,15,20,25,**30**	**25**	*deg*
Roll Offset	**0**	± 1,2,**3**,**6**,**9**	*deg*

Negative roll offset angles correspond to the robot rolled to the right and therefore strafing to the right, positive body roll is toward the left.

being defined for use in constructing variable sized bounding boxes. When the robot is using an upward or downward gait, the magnitude of the body yaw is approximately symmetric about the zero. In this case, the coordinates of the approximated COM position are chosen to be the instantaneous vertical position and the horizontal position is averaged between the position current and prior step. During strafing, due to the asymmetry of the yaw sinusoid, we chose to use the lower boundary of the body swing as the approximated position, as this boundary is closer to the rest position of the body.

V. EXPERIMENTAL SETUP AND RESULTS

A. Network Training

1) Training Data Acquisition: In order to train the neural network, two separate parameter sweeps were run, one each for upward climbing [varying driving frequency and body pitch shown in Fig. 5(e)] and strafing behaviors [varying driving frequency and roll angle offset shown in Fig. 5(f)], with the ranges shown in Table I. The roll offset range was selected to produce both large lateral velocities and lateral velocities for fine positioning, with only a single pitch of 25° tested based on its reliability during initial experimentation. Downward climbing utilized a frequency of 1 Hz to insure gait robustness. A Nelder–Mead optimization was run to retune the control parameters from [18] Fig. 5(a)–(d) to account for the robot modifications. Experimental COM data were recorded using a Vicon motion capture system at 300 Hz for the entire randomized set of tests, which were repeated three times and run off of a power supply held at a constant 15 V.

The resulting average velocities from the parameter sweep are shown in Fig. 5. The vertical velocity seen in Fig. 5(a) depends, as expected, primarily on driving frequency. There is some decrease in speed and increase in variability with higher body pitch. The coupling between the posture and climbing dynamics is particularly noticeable in the horizontal velocity. In upward climbing, as shown in Fig. 5(b), the horizontal velocity shows significant undesired strafing particularly in the high-frequency gaits with midrange pitches, whereas low and high pitches held lateral velocities near zero with minor deviation at all tested frequencies [boxes 1 and 2 in Fig. 5(b)]. Examining the strafing experiments, with vertical and lateral velocity shown in Fig. 5(c) and (d), respectively, it can be seen that the offset magnitude is directly correlated to lateral velocity and therefore the amount of strafing generated. However, there is significantly more variability in climbing speed when compared to upward tests, especially at higher frequencies with positive offset angles. As shown by 3 in Fig. 5, there is greater deviation in vertical velocity at a frequency of 3.5 Hz, whereas 4 shows the significant variability in the horizontal velocity at 4 Hz. While not shown, the vertical velocity in downward gaits demonstrated significantly more variability than either the upward or strafing behavior. Due to these factors, a limited subset of gaits were selected to train the network, as shown in Table I (bolded text). These gaits allowed a range of vertical velocities from 7.7 to 32.4 cm/s and a range of average heading angles between 25.9° left and 27.4° right.

2) Neural Network Results: The network was trained and attained a 98.9% goodness of fit in predicting motions based on the R^2 value. In conducting hardware validation of this learned model, it was seen that additional dynamics were affecting the robot's motion. In particular, transition effects from or to a downward climbing command were noted to occasionally induce significant amounts of error.

B. Path Planning: Methods

In order to test the physical instantiation of the legged climbing motion planner, a set of obstacle courses were designed for the robot to navigate. These courses were designed to simulate real-world climbing environments and are categorized into three sets based on physical characteristics: 1) climbable substrates, 2) unclimbable obstacles, and 3) narrow gaps. Climbable substrates are surfaces with sufficient asperity size and strength to facilitate the directional adhesion property of microspines, such as cinderblocks, stone aggregate, or in our case, carpet. Unclimbable obstacles are either substrates with high fragility or insufficient roughness to allow attachment (which constitute forefoot obstacles) or discrete surface changes such as extrusions or ledges (which constitute body obstacles). In these experiments, we only use forefoot obstacles for simplicity. Narrow gaps are regions of climbable substrate between close, unclimbable obstacles or exterior obstacles, which require the robot to change its behavior or morphology to negotiate.

A set of difficult motion planning problems were chosen to define the obstacle courses shown in Fig. 6. Due to the mobility constraints of the climber, the standard narrow gap problem has three permutations based on orientation: vertical [see Fig. 6(a)], horizontal [see Fig. 6(b)], and the OOP narrow gap [see Fig. 6(c)]. Additionally, a discrete obstacle is tested [see Fig. 6(d)], which is representative of a simple unclimbable obstacle, such as a window in the real world. Each course was manually entered into the planner and tested at least three times.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

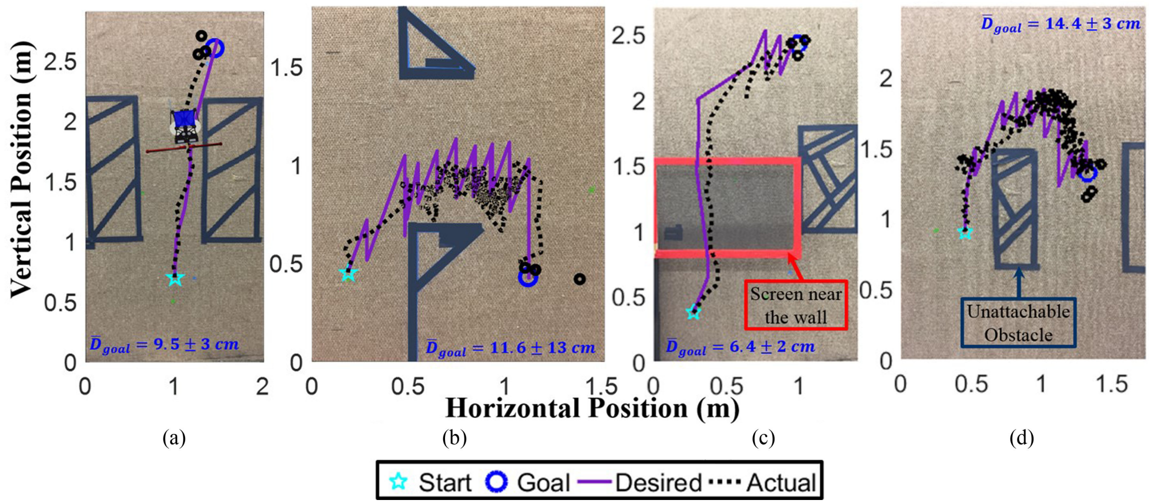

Fig. 6. Results from physical experimentation. The SBPMO desired path (purple) from start node (star) to the goal node (blue circle). The robot's swing-filtered body path is shown in black with final positions from three trials in black circles. The average and standard deviation distance from the goal is presented in blue text. The three versions of the narrow gap problem are shown in (a)–(c) and (d) shows the discrete obstacle. (a) Vertical Gap. (b) Horizontal Gap. (c) OOP Gap. (d) Box Obstacle.

TABLE II
COMPARISON OF COMPUTATION TIME FOR GAP TRAVERSAL PLANNING

Case	Naive Heuristic	Directional Heuristic
Vertical	0.0015s	0.0011s
Horizontal	188.7042s	0.7464s
Out of Plane	failed	1.0099s

VI. Path Planning: Results and Discussion

A. Heuristic Comparison

With a naive heuristic of Euclidean distance, the computation time for a 2-m trajectory to the right costs 188.7 s and explored 192 711 nodes in the planning space. The addition of the custom directional heuristic function reduced the planning time to 0.75 s with expansions of only 3197 nodes. The comparison of the heuristics is further shown in Table II. These results suggest that the directional inflation factor appears to facilitate rapid computation of trajectories.

B. Path Plans and Execution

The planned paths for each case (solid purple lines in Fig. 6) were then tested in a feedforward manner on TAILS. A characteristic experimental path is shown as the dotted lines where the effective paths were determined using the stride-based COM path filtering described in Section IV-D. Final robot positions from other trials are shown as black circles. For the vertical narrow gap [see Fig. 6(a)] that has unclimbable obstacles outside of a 50 cm wide and 120 cm tall gap, the plan was able to primarily utilize the consistent upward gaits while avoiding the obstacles, even though the body swung over the obstacles.

For the horizontal narrow gap [see Fig. 6(b)] that has unclimbable obstacles around a 40 cm wide and 80 cm tall gap, the plan generated resulted in frequent lateral strafing to downward climbing motions. While able to navigate horizontally, these repeated changes cause a buildup in unmodeled transient effects causing reduced accuracy in reaching the goal.

For the OOP narrow gap case [see Fig. 6(c)], which has a mesh screen (outlined by red on the figure) close enough to the wall that only vertical gaits can be implemented under it and an unclimbable obstacle at the boundary, the planner was able to choose the correct gaits to navigate to the goal. The planned path (see video attachment) strafes as much as possible toward the goal initially, switches to a flat upward gait to fit through the gap, and then strafes and navigates to the goal.

Finally, the vertical obstacle [see Fig. 6(d)], which has a 33 cm wide and 90 cm tall obstacle between the start and goal locations, the planner ended up choosing to go over the obstacle (which allowed more room for horizontal strafing maneuvers) before going down to the goal. This path was able to be completed effectively; however, similar to the horizontal gap, lateral motions caused variance in the accuracy of the final position.

With the improved heuristic, unique coupled bounding box obstacle detection, and discrete control, our planner was able to generate achievable motion plans for each of the cases, although significant transients impacted the robot's performance of lateral motions. For the experimental cases without significant lateral motions [see Fig. 6(a) and (c)], TAILS finished within 10 cm of the goal, whereas for cases with significant constrained lateral motion the average distance to goal was high with standard deviations up to 13 cm. In order to test the planner over extended horizons, the algorithm was tasked with designing a path on a real building, as shown in Fig. 7. The resulting path would require 257.8 s to traverse, navigating from near the ground to a perch near one of the upper windows. This was not tested experimentally due to battery limits and concerns over the reliability of the attachment mechanism with the substrate.

The significant unmodeled transients and system drift seen in the lateral experiments could likely be handled by replanning and trajectory correction techniques. However, replanning with this unique mobility will require careful considerations of the same locomotion constraints, the original planner was required to utilize as well as the ability to determine if the system can effectively merge to the previously defined path. If replanning was determined to be necessary, it could be done by

Fig. 7. Theoretically feasible path generated for climbing on a real-world building (FSU's Doak Campbell Stadium).

use of a parallel-SBMPO [30] developed specifically to allow real-time replanning on rapid motion robots.

VII. Conclusion

This article presented a novel planning solution for the unique mobility challenges exhibited by legged locomotion, in particular the dynamic climbing behaviors of TAILS. To do this, a set of fundamental planning challenges for legged locomotion were articulated and related to their equivalent representation on the TAILS platform. TAILS exhibited special attachment constraints due to the directional adhesion mechanism, a nonholonomic heading angle constraint, coupling between the robots size and locomotive behavior, and a style of locomotion that is challenging for state estimation. These were addressed in the planner through sampling only the valid control inputs, applying a unique directional heuristic, and by treating the body as a coupled set of bounding boxes.

Motion plans were developed and tested experimentally for a set of difficult planning problems adapted for climbing. In our case, the narrow passage problem has three permutations, the vertical, horizontal, and OOP directions, each of which pose different planning challenges. The vertical problem showed that tracking feet and body separately allowed navigation through very thin gaps, the use of a directional heuristic allowed efficient computation of a path through the horizontal gap, and the OOP problem highlighted the planners ability to account for changes in the bounding box geometry with gait. Together these demonstrated, for the first time, motion planned for a dynamic legged robot on a vertical wall.

With these advances, we are one step closer to enabling motion planners to fully utilize the unique mobility features provided by legged robotics. Some areas for improvement include more explicit modeling of the significant transient behavior, increasing the number of potential locomotion modalities, and improving the collision detection/planning of the discrete feet. A potential motion planning regime that considers long-term and short-term patterns of motion (using an long-short term memory (LSTM) neural network) could likely model both the persistent steady-state motions as well as the short-duration transients, whereas replanning may account for variability in the motion. With further refinement, from both the locomotion development/modeling and planner development, TAILS should be able to execute planning maneuvers, untethered, on large-scale real-world buildings.

Acknowledgment

The U.S. Government is authorized to reproduce and distribute reprints for government purposes not withstanding any copyright notation thereon.

References

[1] U. Saranli, M. Buehler, and D. E. Koditschek, "RHex: A simple and highly mobile hexapod robot," *Int. J. Robot. Res.*, vol. 20, no. 7, pp. 616–631, 2001.

[2] H.-W. Park, P. M. Wensing, and S. Kim, "High-speed bounding with the MIT Cheetah 2: Control design and experiments," *Int. J. Robot. Res.*, vol. 36, no. 2, pp. 167–192, 2017.

[3] D. J. Blackman *et al.*, "Leg design for running and jumping dynamics," in *Proc. IEEE Int. Conf. Robot. Biomimetics*, 2017, pp. 2617–2623.

[4] B. D. Miller, P. R. Rivera, J. D. Dickson, and J. E. Clark, "Running up a wall: The role and challenges of dynamic climbing in enhancing multimodal legged systems," *Bioinspiration Biomimetics*, vol. 10, no. 2, 2015, Art. no. 025005. [Online]. Available: http://stacks.iop.org/1748-3190/10/i=2/a=025005

[5] M. Harper, J. Pace, N. Gupta, C. Ordonez, and E. G. Collins, "Kinematic modeling of a RHex-type robot using a neural network," *Proc. SPIE*, vol. 10195, 2017, Art. no. 1019507.

[6] M. Wermelinger, P. Fankhauser, R. Diethelm, P. Krüsi, R. Siegwart, and M. Hutter, "Navigation planning for legged robots in challenging terrain," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 1184–1189.

[7] S. Feng, E. Whitman, X. Xinjilefu, and C. G. Atkeson, "Optimization-based full body control for the DARPA robotics challenge," *J. Field Robot.*, vol. 32, no. 2, pp. 293–312, 2015.

[8] M. Johnson *et al.*, "Team IHMC's lessons learned from the DARPA robotics challenge trials," *J. Field Robot.*, vol. 32, no. 2, pp. 192–208, 2015.

[9] R. Tedrake, "LQR-trees: Feedback motion planning on sparse randomized trees," in *Proc. Robot., Sci. Syst.*, 2009.

[10] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. J. Robot. Res.*, vol. 30, no. 7, pp. 846–894, 2011.

[11] S. Aine, S. Swaminathan, V. Narayanan, V. Hwang, and M. Likhachev, "Multi-heuristic A*," *Int. J. Robot. Res.*, vol. 35, no. 1–3, pp. 224–243, 2016.

[12] S. Koenig, M. Likhachev, and D. Furcy, "Lifelong planning A*," *Artif. Intell.*, vol. 155, no. 1, pp. 93–146, 2004.

[13] S. M. LaValle and J. J. Kuffner, Jr, "Randomized kinodynamic planning," *Int. J. Robot. Res.*, vol. 20, no. 5, pp. 378–400, 2001.

[14] T. Bretl, S. Rock, J.-C. Latombe, B. Kennedy, and H. Aghazarian, "Free-climbing with a multi-use robot," in *Proc. Exp. Robot. IX*, 2006, pp. 449–458.

[15] M. P. Murphy, C. Kute, Y. Mengüç, and M. Sitti, "Waalbot II: Adhesion recovery and improved performance of a climbing robot using fibrillar adhesives," *Int. J. Robot. Res.*, vol. 30, no. 1, pp. 118–133, 2011.

[16] T. Bretl, "Motion planning of multi-limbed robots subject to equilibrium constraints: The free-climbing robot problem," *Int. J. Robot. Res.*, vol. 25, no. 4, pp. 317–342, 2006.

[17] D. I. Goldman, T. S. Chen, D. M. Dudek, and R. J. Full, "Dynamics of rapid vertical climbing in cockroaches reveals a template," *J. Exp. Biol.*, vol. 209, no. 15, pp. 2990–3000, 2006.

[18] J. M. Brown, M. P. Austin, B. Kanwar, T. E. Jonas, and J. E. Clark, "Maneuverability in dynamic vertical climbing," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 4340–4347.

[19] D. Dunlap, C. Caldwell, and E. Collins, "Nonlinear model predictive control using sampling and goal-directed optimization," in *Proc. Multi-Conf. Syst. Control*, Yokohama, Japan, Sep. 8–9, 2010, pp. 1349–1356.

[20] K. Jayaram and R. J. Full, "Cockroaches traverse crevices, crawl rapidly in confined spaces, and inspire a soft, legged robot," *Proc. Nat. Acad. Sci. USA*, vol. 113, no. 8, pp. E950–E957, 2016.

[21] S. Kim, A. T. Asbeck, M. R. Cutkosky, and W. R. Provancher, "SpinybotII: Climbing hard walls with compliant microspines," in *Proc. 12th Int. Conf. Adv. Robot.*, 2005, pp. 601–606.

[22] D. D. Dunlap, *et al.*, "Motion planning for mobile robots via sampling-based model predictive optimization," in *Recent Advances in Mobile Robotics*. Rijeka, Croatia: InTech, 2011.

[23] O. Chuy, E. Collins, D. Dunlap, and A. Sharma, "Sampling-based direct trajectory generation using the minimum time cost function," in *Proc. Int. Symp. Exp. Robot.*, 2013, pp. 651–666.

[24] C. V. Caldwell, D. D. Dunlap, and E. G. Collins, Jr, "Application of sampling based model predictive control to an autonomous underwater vehicle," *Ship Sci. Technol.*, vol. 4, no. 7, pp. 55–63, Jul. 2010.

[25] C. Ordonez, N. Gupta, E. G. Collins, J. Clark, and A. M. Johnson, "Power modeling of the XRL hexapedal robot and its application to energy efficient motion planning," in *Proc. Int. Conf. Climbing Walking Robots*, Baltimore, MD, USA, Jul. 2012, pp. 689–696.

[26] M. Y. Harper, J. V. Nicholson, E. G. Collins, J. Pusey, and J. E. Clark, "Energy efficient navigation for running legged robots," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2019, pp. 6770–6776.

[27] G. D. Francis, E. G. Collins, Jr, O. Chuy, and A. Sharma, "Sampling-based trajectory generation for autonomous spacecraft rendezvous and docking," in *Proc. AIAA Guid., Navig., Control Conf.*, 2013, p. 4549.

[28] O. Chuy, E. Collins, W. Yu, and C. Ordonez, "Power modeling of a skid steered, wheeled robotic ground vehicle," in *Proc. IEEE Conf. Robot. Autom.*, Kobe, Japan, May 2009, pp. 4118–4123.

[29] C. Ordonez, N. Gupta, O. Chuy, and E. Collins, "Momentum based traversal of mobility challenges for autonomous ground vehicles," in *Proc. IEEE Int. Conf. Robot. Autom.*, Karlsruhe, Germany, May 6–10, 2013, pp. 752–759.

[30] M. Y. Harper, C. Ordonez, E. G. Collins, and G. Erlebacher, "Parallel approach to motion planning in uncertain environments," *Proc. SPIE*, vol. 10640, 2018, Art. no. 106400H. [Online]. Available: https://doi.org/10.1117/12.2304433