

Generating Continuous Motion and Force Plans in Real-Time for Legged Mobile Manipulation

Parker Ewen¹, Jean-Pierre Sleiman², Yuxin Chen¹, Wei-Chun Lu¹, Marco Hutter², Ram Vasudevan¹

Abstract—Manipulators can be added to legged robots, allowing them to interact with and change their environment. Legged mobile manipulation planners must consider how contact forces generated by these manipulators affect the system. Current planning strategies either treat these forces as immutable during planning or are unable to optimize over these contact forces while operating in real-time. This paper presents the Stability and Task Oriented Receding-Horizon Motion and Manipulation Autonomous Planner (STORMMAP) that is able to generate continuous plans for the robot’s motion and manipulation force trajectories that ensure dynamic feasibility and stability of the platform, and incentivizes accomplishing manipulation and motion tasks specified by a user. A variety of simulated experiments on a quadruped with a manipulator mounted to its torso demonstrate the versatility of STORMMAP. In contrast to existing state of the art methods, the approach described in this paper generates continuous plans in under ten milliseconds, an order of magnitude faster than previous strategies.

I. INTRODUCTION

Legged robots offer unique advantages when compared to wheeled or aerial robots when autonomously performing tasks in complex terrains while accommodating non-trivial payloads. To improve the utility of legged robots, additional appendages, such as arms, can be integrated into the platform, thereby providing manipulation capabilities. The addition of manipulators can enable legged platforms to interact with and alter their environments; for example, to clear obstacles which may be otherwise impassable.

In fact, manipulation capabilities have been implemented on a number of existing legged platforms. The torque-controllable quadruped Anymal [1] was outfitted with a Kinova Jaco six DoF manipulator [2] to accomplish a variety of tasks [3]. The hydraulically-actuated quadruped (HyQ) robot [4][5] and the Boston Dynamics’ SpotMini also have manipulation capabilities. For humanoid robots such as Digit [6], Toro [7], and Boston Dynamic’s Atlas, arms are an inherent part of their humanoid physiology.

The addition of manipulation appendages introduces additional forces that must be accounted for during planning. For instance, if these interaction forces are unaccounted for while assessing the stability of the platform, the forces that enable manipulation may unduly contribute to instability and subsequent failure.

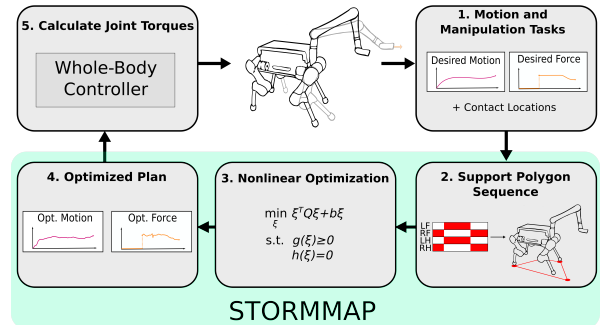


Fig. 1: Block diagram for the real-time dynamic legged mobile manipulation planner STORMMAP. (1) The user specifies desired motion and manipulation tasks for the robot to complete. (2) A sequence of support polygons is generated for one gait cycle to be used for the ZMP stability criterion. (3) A motion trajectory and manipulation force plan is computed using nonlinear optimization. (4) This plan accomplishes the motion and manipulation tasks if it can do so while maintaining dynamic feasibility and stability. (5) The robot’s controller then calculates the joint torques and contact forces at the feet to achieve the generated plan.

The existing literature in planning for legged robotic systems with manipulators can be divided into two categories. The first set of approaches, which we refer to as static techniques, are focused on planning the platform’s motion while assuming that the forces at the end-effector are unalterable during planning. Static planning either accounts for this constant expected manipulation force in the dynamics of the platform [8], [9] or treats the new contact created by the manipulator as an additional point affecting the structure of the robot’s support [10]–[12]. These static methods generally either result in motion plans which sacrifice dynamic feasibility or are unable to compute a feasible motion plan while satisfying the imposed constraints.

The second set of approaches, which we refer to as dynamic techniques, focus on leveraging the additional appendages to improve the stability of the overall system by solving for both the contact locations and contact forces of these new appendages. Such strategies have been used to construct plans to traverse complex terrain [13]–[15] and often leverage mixed-integer optimization to plan foot- and hand-holds [16], [17]. By planning for the contact locations and forces alongside the motion trajectories, these dynamic strategies can use the manipulator to stabilize the robot while completing manipulation tasks. However, dynamic strategies suffer from either long computational times associated with the complex nature of planning for the contact locations, or utilize coarse time discretizations to speed up computations.

The contribution of this paper is the introduction of a real-time dynamic planning strategy, the Stability and Task Oriented Receding-Horizon Motion and Manipulation Au-

¹Authors are with Robotics and Optimization for Analysis of Human Motion Lab, University of Michigan, USA

²Authors are with the Robotic Systems Laboratory, ETH Zürich, Zürich 8092, Switzerland

This work is supported by the National Science Foundation Career Award 1751093, the Ford Motor Company via the Ford-UM Alliance under award N022977, and the Office of Naval Research under Award Number N00014-18-1-2575.

onomous Planner (STORMMAP) for legged mobile manipulation. Plans are generated for a time horizon which is 80 times longer than their computational time, allowing these plans to account for disturbances and to be generated in a receding-horizon fashion. STORMMAP jointly plans for both the motion trajectory of the robot and the contact force at a manipulator's end-effector given a contact location. To the best of our knowledge, this is the fastest dynamic planning strategy in the literature based on planning speed by an order of magnitude [13]–[15], [17], [18].

The paper is organized as follows: We first introduce the robot model and contextualize the use of a legged mobile manipulation planner in Section II. In Section III we formulate the parameterization used to reduce the computational complexity of the optimization problem. Details regarding the nonlinear optimization problem used for STORMMAP are laid out in Section IV. We demonstrate the versatility of STORMMAP on the Anymal platform [1] which is equipped with a Kinova Jaco 6 DoF arm [2] at its front. The planner is validated in a variety of experiments described in Section V, which consist of both manipulation tasks and instability-inducing scenarios. We include a baseline motion planner for comparison that does not consider the manipulation forces when planning. Finally, we offer concluding remarks in Section VI.

II. PROBLEM SETUP

This section introduces the model of the robot and motivates the use of a legged mobile manipulation planner to generate motion trajectories and manipulation force plans. Let the generalized coordinates, \mathbf{q} , be a vector of continuous functions of time that describe the configuration of a robot, \mathbf{u} be a vector of differentiable functions of time, and $\dot{\mathbf{u}}$ be their time derivatives, each with the following forms:

$$\mathbf{q} = \begin{bmatrix} \mathbf{q}_{b_P} \\ \mathbf{q}_{b_R} \\ \mathbf{q}_j \end{bmatrix} \quad \mathbf{u} = \begin{bmatrix} \mathbf{v}_b \\ \boldsymbol{\omega}_b \\ \dot{\mathbf{q}}_j \end{bmatrix} \quad \dot{\mathbf{u}} = \begin{bmatrix} \mathbf{a}_b \\ \dot{\boldsymbol{\omega}}_b \\ \ddot{\mathbf{q}}_j \end{bmatrix} \quad (1)$$

\mathbf{q}_{b_P} and \mathbf{q}_{b_R} are functions of time describing the robot's center-of-mass (CoM) position and body rotation, respectively, \mathbf{v}_b and $\boldsymbol{\omega}_b$ are functions of time describing the robot's linear and angular velocity, respectively, \mathbf{a}_b and $\dot{\boldsymbol{\omega}}_b$ are functions of time describing the robot's linear and angular accelerations, respectively, and \mathbf{q}_j are functions of time describing the n_j actuated joint positions of the robot. We assume that position of the robot, \mathbf{q}_{b_P} , is parameterized using Cartesian coordinates.

The high-fidelity model for a legged robot with a manipulator derived using the projected Newton-Euler equation with the constraint compliant Lagrange formulation [19, Eq. 4.55a] can be written as follows:

$$\mathbf{M}(\mathbf{q}(t))\dot{\mathbf{u}}(t) + \mathbf{b}(\mathbf{q}(t), \mathbf{u}(t)) = \mathbf{S}^T \boldsymbol{\tau}(t) + \mathbf{J}(t)^T \mathbf{f}(t) \quad (2)$$

where at time t , $\mathbf{M}(\mathbf{q}(t))$ is the mass matrix, $\mathbf{b}(\mathbf{q}(t), \mathbf{u}(t))$ is the vector of Coriolis, centrifugal and gravity terms, $\boldsymbol{\tau}(t)$ are the torques acting in the direction of the generalized coordinates, \mathbf{S} is the selection matrix for the actuated joints,

$\mathbf{f}(t)$ are the contact forces acting on the robot, and $\mathbf{J}(t)$ is the contact Jacobian matrix. We make the following assumption about this model:

Assumption 1. *There exists a controller which uses (2) to calculate the joint torques and contact forces for the robot to achieve some desired base acceleration, \mathbf{a}_b , and a desired manipulation force, \mathbf{f}_m , that is a subset of the set of contact forces, \mathbf{f} .*

Various techniques have been proposed in the literature to construct a controller that satisfies this assumption [1], [3], [5], [9], [17], [20], [21].

The goal of our planner is to generate trajectories for \mathbf{q}_{b_P} , \mathbf{v}_b , \mathbf{a}_b , and \mathbf{f}_m over a time interval $[0, T]$ that are dynamically feasible, maintain stability, and can accomplish a pre-defined motion or manipulation task, if possible. We formulate this planner using the following nonlinear optimization problem:

$$\min_{\mathbf{q}_{b_P}, \mathbf{v}_b, \mathbf{a}_b, \mathbf{f}_m} \text{Cost}(\mathbf{q}_{b_P}, \mathbf{v}_b, \mathbf{a}_b, \mathbf{f}_m) \quad (3a)$$

$$\text{s.t. Motion Spline Junctions}(\mathbf{q}_{b_P}, \mathbf{v}_b) \quad (3b)$$

$$\text{Initial Spline Point}(\mathbf{q}_{b_P}, \mathbf{v}_b) \quad (3c)$$

$$\text{Friction Pyramid}(\mathbf{a}_b, \mathbf{f}_m) \quad (3d)$$

$$\text{ZMP Stability}(\mathbf{q}_{b_P}, \mathbf{a}_b, \mathbf{f}_m) \quad (3e)$$

$$\text{Force Limits}(\mathbf{f}_m) \quad (3f)$$

$$\text{Free Motion Direction}(\mathbf{f}_m) \quad (3g)$$

We achieve a smooth motion trajectory by constraining the position and velocity of the robot to be continuous (3b) and constrain the initial position and velocity of the optimized trajectory to be equivalent to the measured value at the start of the optimization (3c). The friction pyramid models the relationship between the forces at the contact locations and the accelerations of the robot to ensure dynamic feasibility of the generated plans (3d). The stability of the platform is maintained in the presence of manipulation forces using the ZMP stability model [22] that includes these forces (3e). Lastly, the manipulation force direction is specified using (3f) and are also prohibited in free motion directions (3g). The cost is used to incentivize the completion of tasks.

In practice, this optimization problem is solved by discretizing the functions that appear as decision variables at pre-specified sampling times. The accuracy of the approximation is dependent on the chosen time discretization, with a finer discretization leading to a more accurate approximation but more parameters, thereby increasing the solution space of the optimization and the computational time required to solve the problem. We overcome this curse of dimensionality by following the example set in [1], [3] and parameterize the plan using quintic splines. We reformulate (3) using the spline parameterization which results in a nonlinear optimization problem whose average computational time is under 10 milliseconds. The speed of the optimization is such that we can use STORMMAP to plan in a receding horizon fashion as the planning horizon is longer than the time it takes to compute a plan.

III. SPLINE PARAMETERIZATION

This section describes the parameterization used to encode the motion and manipulation force trajectories as quintic splines. The time-dependent CoM trajectory can be parameterized using a set of quintic spline parameters. Spline parameters for the motion trajectory are denoted by α while those for manipulation force are denoted by β . For instance, the spline parameterization for the x direction of the motion trajectory is

$$x(t) = \mathbf{t}^\top(t)\alpha_x, \quad \dot{x}(t) = \dot{\mathbf{t}}^\top(t)\alpha_x, \quad \ddot{x}(t) = \ddot{\mathbf{t}}^\top(t)\alpha_x, \quad (4)$$

where $\alpha_x = [\alpha_{x,5}, \alpha_{x,4}, \alpha_{x,3}, \alpha_{x,2}, \alpha_{x,1}, \alpha_{x,0}]^\top$ and $\mathbf{t} = [t^5, t^4, t^3, t^2, t, 1]^\top$. The same parameterization is used for the manipulation force at the end effector. We can describe the full motion or manipulation force trajectory as:

$$\begin{bmatrix} x(t) \\ y(t) \\ z(t) \end{bmatrix} = T(t) \begin{bmatrix} \alpha_x \\ \alpha_y \\ \alpha_z \end{bmatrix} \quad (5)$$

where $T(t) = \text{diag}(\mathbf{t}^\top(t), \mathbf{t}^\top(t), \mathbf{t}^\top(t))$.

Next, we account for the hybrid dynamics of legged robotic systems. To account for the contact transitions at the feet during walking, we consider the following:

Assumption 2. *Given a user-specified desired velocity, which is assumed constant over the motion trajectory, a set of foot contact locations can be computed using the linear inverted pendulum model [23]. This contact sequence is calculated for one gait cycle with a velocity-dependant time horizon. A new support polygon is generated for each lift-off and touch-down event of the feet, and using a dynamic trotting gait, this yields a sequence of five support polygons over a gait cycle.*

A single piecewise quintic spline is used per support polygon in the generated sequence described in Assumption 2. This means that for each optimization, we solve for the parameters of five splines associated with the motion trajectory. These splines are indexed using i and each has a domain $[t_0^{(i)}, t_f^{(i)}]$. We concatenate the spline parameters for an entire motion trajectory composed of a series of piecewise splines into a single vector α .

To expand on Assumption 2, we form a desired motion trajectory given the current position of the robot and the user-specified constant desired velocity. We evaluate this desired motion trajectory over the entire optimization horizon as follows:

Definition 3. *Let π_α denote the desired position trajectory of the robot, with $\dot{\pi}_\alpha$ and $\ddot{\pi}_\alpha$ the desired velocity and acceleration trajectories, respectively. We evaluate π_α by first setting $\pi_\alpha(0)$ to the measured position of the robot and calculating $\pi_\alpha(T)$ by integrating the constant desired velocity from Assumption 2 over the optimization horizon $[0, T]$. The desired velocity trajectory is set to the constant user-specified velocity and the desired acceleration trajectory is set to zero.*

To account for the contact transitions at the manipulator, we make the following assumption:

Assumption 4. *The contact locations and times of contact for the manipulator are known a priori to the planner. Contact occurs when the manipulator has at least one direction of constrained motion, and contact is broken when all directions at the manipulator are in free motion.*

Note that the contact locations and times of contact for the manipulator can be either specified by the user or by an external algorithm. When contact occurs or is broken, we call this a *contact transition*. We denote the number of contact transitions that occur over the optimization horizon as n_m . Similar to the motion trajectory, we construct a sequence of force splines to describe the force across a time horizon. We use one force spline to describe the force before a transition, indexed as j with the domain of each spline equal to $[t_0^{(j)}, t_f^{(j)})$, and one force spline for the force after the transition, indexed by $j + 1$. This allows for a discontinuity in the force at the transition. As in the case of the motion trajectory, we concatenate all of the spline parameters for the manipulation force trajectory into a single vector that we denote by β .

When possible, we want the generated manipulation force plan to achieve some user-specified manipulation task. To define such a task we assume the following:

Assumption 5. *For a manipulation task, we are given a desired manipulation force and contact location.*

We also define the set of *force events* which represent the time-ordered set of desired manipulation tasks:

Definition 6. *Let a force event $\pi_\beta^{(j)}$ denote the desired manipulation force over a range $[t_0^{(j)}, t_f^{(j)})$ which occurs before the j^{th} contact transition. The set of force events is the set of all such force events arranged sequentially, denoted as $\pi_\beta = \{\pi_\beta^{(j)} | j \in J\}$, such that, combined, they span the range of the optimization horizon $[0, T]$. Here, $J \in 1, \dots, n_m + 1$.*

Using the spline parameterization for the motion trajectory and manipulation force, we rewrite the optimization problem in (3) to optimize over the set of spline parameters.

IV. OPTIMIZATION

This section describes the formulation of the cost functions and constraints of the nonlinear optimization problem (3). The vector of motion and force spline parameters mapping to the state trajectories is the optimization variable. In practice, it is challenging to validate a constraint over a continuous time interval. We therefore discretize the time based on the following assumption:

Assumption 7. *Take each piecewise motion and force spline and sample them at six evenly spaced time intervals, including the initial and final point of the spline.*

We denote the set of sample times as S_t . This discretization is needed, as we cannot enforce constraints over an

entire trajectory. The fineness of this discretization makes it unlikely that the generated continuous trajectories will violate a constraint between time samples. This discretization does not increase the solution space of problem (3) due to the spline parameterization, even though the same number of constraints are used for both the original and parameterized optimizations.

The remainder of this section details the formulation of the nonlinear optimization problem used to generate motion and manipulation force plans. The costs of the optimization problem in (3) are described in Section IV-A, while the equality and inequality constraints are formulated in Sections IV-B and IV-C, respectively.

A. Cost

The cost in the optimization problem (3) is assumed to be a positive definite quadratic form. In particular, the quadratic cost, Q , and linear cost, b , take the following form:

$$Q = \text{diag}(Q_\alpha^{(1)} + Q_{\alpha_{im}}, \dots, Q_\alpha^{(5)}, Q_\beta^{(1)} + Q_{\beta_{im}}, \dots, Q_\beta^{(n_m)})$$

$$b^\top = [b_\alpha^{(1)\top} + b_{\alpha_{im}}^\top, \dots, b_\alpha^{(5)\top}, b_\beta^{(1)\top} + b_{\beta_{im}}^\top, \dots, b_\beta^{(n_m)\top}] \quad (6)$$

where $Q_\alpha^{(i)} \in \mathbb{R}^{6 \times 6}$ and $b_\alpha^{(i)} \in \mathbb{R}^{6 \times 1}$ are the quadratic and linear costs for the i^{th} motion spline, and $Q_\beta^{(j)} \in \mathbb{R}^{6 \times 6}$ and $b_\beta^{(j)} \in \mathbb{R}^{6 \times 1}$ are the quadratic and linear costs for the j^{th} force spline. The costs $Q_{\alpha_{im}}^{(1)}$ and $b_{\alpha_{im}}^{(1)}$ minimize the difference between the initial acceleration of the motion trajectory and the measured acceleration, while $Q_{\beta_{im}}^{(1)}$ and $b_{\beta_{im}}^{(1)}$ minimize the difference between the initial force and measured force.

The cost per motion spline is a linear combination of several time-dependent costs and one time-independent cost. We compute the quadratic and linear cost terms for each motion spline as:

$$Q_\alpha^{(i)} = Q_{acc}^{(i)} + \sum_{t_s \in S_t} (Q_{\alpha_\pi}^{(i)}(t_s) + Q_{\alpha_d}^{(i)}(t_s))$$

$$b_\alpha^{(i)} = \sum_{t_s \in S_t} (b_{\alpha_\pi}^{(i)}(t_s) + b_{\alpha_d}^{(i)}(t_s)) \quad (7)$$

where $Q_{\alpha_\pi}^{(i)}$ is a cost to incentivize a desired task, $Q_{\alpha_d}^{(i)}$ penalizes deviations between subsequent plans and $Q_{acc}^{(i)}$ penalizes large accelerations over the spline. An equivalent calculation to (7) is used for the cost per force spline, with individual costs denoted using β . The specific elements of (7) are more formally defined in the remainder of this subsection.

1) *Desired Motion and Manipulation Tasks:* We want the generated motion and manipulation force plans to achieve desired tasks, when feasible. To incentivize these tasks, we use the desired motion trajectory, π_α , and its derivatives (Definition 3), along with the set of force events, π_β , representing the desired manipulation tasks (Definition 6). We minimize the difference between values sampled from the motion and manipulation force plans and their respective desired trajectories at each sampling time t_s in a least-squares manner. An example for the motion splines follows, where i

is the i^{th} motion spline corresponding to the sampling time:

$$\alpha^{(i)\top} Q_{\alpha_\pi}^{(i)}(t_s) \alpha^{(i)} + b_{\alpha_\pi}^{(i)}(t_s) \alpha^{(i)} = \left\| T(t_s) \alpha^{(i)} - \pi_t(t_s) \right\|_2 + \left\| \dot{T}(t_s) \alpha^{(i)} - \dot{\pi}_t(t_s) \right\|_2 + \left\| \ddot{T}(t_s) \alpha^{(i)} - \ddot{\pi}_t(t_s) \right\|_2 \quad (8)$$

An equivalent formulation is used to calculate the costs for the j^{th} force spline, $Q_{\beta_\pi}^{(j)}(t_s)$ and $b_{\beta_\pi}^{(j)}(t_s)$.

2) *Deviation from Previous Plans:* Because the optimization problem we are solving is nonlinear, each time we solve it, we initialize the solver using the solution found in the previous iteration. Note, we solve the optimization problem using an SQP solver. Given t_d , the time which has elapsed since the start of the previous plan, and a sampled time, t_s , we minimize the difference between the current plan and the previous successfully computed plan using least squares:

$$\alpha^{(i)\top} Q_{\alpha_d}^{(i)}(t_s) \alpha^{(i)} + b_{\alpha_d}^{(i)}(t_s) \alpha^{(i)} = \left\| T(t_s) \alpha^{(i)} - T(t_s + t_d) \alpha_{prev}^{(i)} \right\|_2 + \left\| \dot{T}(t_s) \alpha^{(i)} - \dot{T}(t_s + t_d) \alpha_{prev}^{(i)} \right\|_2 + \left\| \ddot{T}(t_s) \alpha^{(i)} - \ddot{T}(t_s + t_d) \alpha_{prev}^{(i)} \right\|_2 \quad (9)$$

An equivalent formulation exists for the force spline costs $Q_{\beta_d}^{(j)}(t_s)$ and $b_{\beta_d}^{(j)}(t_s)$.

3) *Initial and Measured State Values:* Large differences between the initial acceleration and manipulation force of the optimized plan and measured accelerations, a_b^* , and manipulation forces, f_m^* , can cause jumps that may be infeasible to achieve. Minimizing the difference between the initial and measured values can prevent these jumps:

$$\alpha^{(1)\top} Q_{\alpha_{im}} \alpha^{(1)\top} + b_{\alpha_{im}}(t_s) \alpha^{(1)} = \left\| \dot{T}(0) \alpha^{(1)} - a_b^* \right\|_2 \quad (10)$$

$$\beta^{(1)\top} Q_{\beta_{im}} \beta^{(1)\top} + b_{\beta_{im}}(t_s) \beta^{(1)} = \left\| T(0) \beta^{(1)} - f_m^* \right\|_2 \quad (11)$$

These quadratic and linear costs are used in (6) for the first motion and force spline only.

4) *Minimize Acceleration:* We want the generated motion trajectory to be as close to the user-specified trajectory, which has a constant velocity, as possible. Thus, we minimize the accelerations of the generated plan. Using the method adopted from [24] and implemented in [25], [26], we derive the quadratic cost by squaring and integrating the acceleration over the time duration of each spline and solve for the central term $Q_{acc}^{(i)} = \int_{t_0}^{t_f} \ddot{t}^\top \ddot{t} dt$. This implementation minimizes the cost over the entire motion spline. For a more thorough treatment of this cost, see [25], [26]. Note, this additional cost is only used for the motion splines.

B. Equality Constraints

Three equality constraints are used in the optimization problem (3). They ensure the piecewise motion splines are connected and continuous in their first derivative (3b), that the optimized initial position and velocity are equal to the measured values (3c), and that manipulation forces cannot be applied in a direction of free motion (3g). To ensure feasibility of the optimization problem, we include slack variables for each constraint.




Fig. 2: Using the baseline planner, Anymal cannot push a table while walking forwards. STORMMAP is able to accomplish both the manipulation and motion task. To make the task harder a random, unknown disturbance is applied to Anymal perpendicular to the direction of motion. Since forces in the y direction were not constrained for this task, STORMMAP mitigates this disturbance by applying an opposing force to the table. Forces shown are acting on the robot.

1) *Motion Spline Junctions*: Recall from Assumption 2 that a sequence of piecewise motion splines are used to account for the hybrid nature of the system. To ensure these splines connect and are continuous in the first derivative, we formulate the following equality for (3b):

$$\begin{bmatrix} T(t_f^{(i)}) & T(t_0^{(i+1)}) \\ \dot{T}(t_f^{(i)}) & \dot{T}(t_0^{(i+1)}) \end{bmatrix} \begin{bmatrix} \alpha^{(i)} \\ -\alpha^{(i+1)} \end{bmatrix} = \mathbf{0}_{2 \times 1} \quad (12)$$

2) *Initial Spline Point*: If the initial position and velocity of the optimized motion trajectory are not equivalent to the measured values, this would cause large initial accelerations during implementation. To prevent this, we ensure that the initial point of the first motion spline at time zero is equal to the measured position of the robot, r^* . Likewise, we ensure the initial velocity is set to the measured velocity, v^* , which leads to the constraint (3c):

$$\begin{bmatrix} T(0) \\ \dot{T}(0) \end{bmatrix} \alpha^{(1)} = \begin{bmatrix} r^* \\ v^* \end{bmatrix} \quad (13)$$

This constraint is only for the first point of the optimized motion trajectory, and thus this constraint only applies to the first motion spline.

3) *Free Motion Direction*: To maintain dynamic feasibility, we ensure forces cannot be applied in directions for which there is free motion by setting these forces equal to zero. This is done by specifying the direction of free motion through a selection matrix $S^{(j)} = \text{diag}(S_x^{(j)}, S_y^{(j)}, S_z^{(j)})$, where (3g) is then:

$$S^{(j)} \beta^{(j)} = \mathbf{0}_{6 \times 1}, \quad (14)$$

where each component $S_i^{(j)} \in \mathbb{R}^{6 \times 6}$ is defined as follows:

$$S_i^{(j)} = \begin{cases} \mathbb{I}_{6 \times 6} & \text{for free-motion} \\ \mathbf{0}_{6 \times 6} & \text{otherwise} \end{cases} \quad (15)$$

C. Inequality Constraints

The inequality constraints ensure that the generated motion trajectory and manipulation plans remain dynamically feasible and stable. This is done using the friction pyramid model (3d) and ZMP stability criterion (3e). Additionally, we prevent the manipulator from applying a force in the direction of free motion (3g).

1) *Friction Pyramid*: Feasibility of the motion and manipulation plan requires the robot's contact forces can generate the desired accelerations without slipping. The friction pyramid model [27] used in (3d), a linearized approximation of the friction cone model that also includes the forces at the manipulator's end-effector, can describe this condition:

$$\begin{bmatrix} 1 & 0 & -\mu \\ -1 & 0 & -\mu \\ 0 & 1 & -\mu \\ 0 & -1 & -\mu \end{bmatrix} \mathbf{a}_b(t_s) + \begin{bmatrix} 1 & 0 & -\mu \\ 1 & 0 & -\mu \\ 0 & 1 & -\mu \\ 0 & 1 & -\mu \end{bmatrix} \mathbf{f}_m(t_s) \leq \begin{bmatrix} \mu mg \\ \mu mg \\ \mu mg \\ \mu mg \end{bmatrix} \quad (16)$$

where m is the whole-body mass of the robot, g is the gravitational constant $-9.81 \frac{m}{s^2}$, and μ is the coefficient of friction. Forces at the manipulator can be used to increase the feasible accelerations of the base if the tangential forces at the feet cannot accommodate these accelerations. This constraint applies to each sample time $t_s \in S_t$.

2) *ZMP Stability*: One method of ensuring stability of a legged system is by using the ZMP stability criterion, which simplifies the robot model to that of an inverted pendulum [22]. This constraint requires that the ZMP remain within the support polygon for all time. Similar to the method employed by [8], [9], we include the manipulation force into the ZMP derivation and evaluate the constraint at each sample time $t_s \in S_t$:

$$\mathbf{r}_{zmp}(t_s) = \frac{\mathbf{n}(t_s) \times \boldsymbol{\tau}(t_s)}{\mathbf{n}(t_s) \cdot \mathbf{f}(t_s)} \quad (17)$$

$$\mathbf{f}(t_s) = m(\mathbf{g} - \mathbf{a}_b(t_s)) + \mathbf{f}_m(t_s) \quad (18)$$

$$\boldsymbol{\tau}(t_s) = \mathbf{q}_{bp}(t_s) \times m(\mathbf{g} - \mathbf{a}_b(t_s)) \dots + (\mathbf{q}_{bp}(t_s) + C_{Ic} \mathbf{r}_{c-m}) \times \mathbf{f}_m(t_s) \quad (19)$$

where, at time t_s , $\mathbf{n}(t_s)$ is the normal to the ground, $\mathbf{f}(t_s)$ the force exerted by the robot onto the environment (18), and $\boldsymbol{\tau}(t_s)$ is the resultant moment (19), m the robot's whole-body mass, $\mathbf{g} = [0, 0, g]^T$, $C_{Ic}(t_s)$ the rotation matrix between the inertial frame and the body frame, and $\mathbf{r}_{c-m}(t_s)$ the position of the CoM to the manipulation force center in the body frame which is assumed constant over the planning horizon.

Let $\mathbf{d}(t_s) = [\mathbf{a}(t_s), \mathbf{b}(t_s), \mathbf{0}_{4 \times 1}]$ and $\mathbf{c}(t_s)$ represent the vectorized line variables of the support polygon at sample time t_s . The ZMP constraint (3e) is satisfied if:

$$\mathbf{d}(t_s) \mathbf{r}_{zmp}(t_s) + \mathbf{c}(t_s) \leq \mathbf{0}_{4 \times 1} \quad (20)$$

We can see from (19) that this constraint is nonlinear.

3) *Force Limit*: Force direction is specified using hard bounds on the force to specify the manipulator is required to pull or push on the environment. Requiring the force to be either greater than or less than zero dictates whether the manipulation force is pushing or pulling on the environment, respectively. For each $t_s \in S_t$, one can represent (3f) as:

$$\mathbf{f}_{lim}^- \leq \mathbf{f}_m(t_s) \leq \mathbf{f}_{lim}^+. \quad (21)$$

V. SIMULATED EXPERIMENTS

This section demonstrates the performance of the algorithm proposed in this paper in simulation on the Anymal robot. We describe the implementation of STORMMAP in Section V-A. Section V-B considers an experiment in which simultaneous motion and manipulation tasks are specified. Section V-C considers an instability-prone scenario in which the robot must rely upon the manipulator to remain stable. Finally, Section V-D describes the computational time of STORMMAP across these experiments. These two experiments and several others are shown in the related video ¹.

A. Implementation

To solve the nonlinear optimization problem in (3), this paper relies on sequential quadratic programming (SQP), active-set method. This method was included within the Anymal software and relies upon the Quadprog++ [28] library. A series of experiments are run to show object manipulation and stabilizing behaviours using the Anymal platform [1] with a 6 DOF Kinova Jaco [2] robotic arm mounted on its base. The controller to realize the plans generated by the solution to (3) uses the robot model in (2) with hierarchical null space projection [3] to control the joint torques of the robot. We compare a subset of the experiments with a baseline planner used on Anymal [26], [29]. This planner does not consider manipulation forces when planning the motion trajectory. Experiments are run using the Gazebo simulator on a desktop computer with an i5-4590 processor with 16GB of RAM. Lastly, because we solve the planning problem in a receding horizon fashion, we use the previously successful plan as an initial guess for the SQP solver.

B. Manipulation Task

To test the manipulation capabilities of the STORMMAP, the robot is required to move a table across a room as illustrated in Fig. 2. A lateral force of 50N is required to push the table, and the coefficient of friction between the hard plastic feet of the robot and the hardwood floor is set at $\mu = 0.3$. Note that the planner is given this coefficient of friction. To make this task more difficult we apply a disturbance to the robot's torso. This disturbance is perpendicular to the robot's motion and is applied 4 seconds after it begins pushing. The direction and magnitude of the disturbance are unknown to the planner. Results for the table pushing experiment using STORMMAP are depicted in Fig. 2b. From the force graph in Fig. 2c we see a disturbance force of 15N applied to the base in the y-direction, shown

in red. During this disturbance, STORMMAP modifies the manipulation force plan to counteract this disturbance force, and continues to apply this force to bring the robot back in-line with the table. Using the baseline Anymal motion planner for this experiment, the robot is able to push the table forward initially, but is unable to maintain an appropriate forward velocity (Fig. 2a). After several seconds of pushing, the manipulator reaches the limit of its reachable space and is unable to push the table further.

C. Stability-Oriented Scenario

While STORMMAP can be used to accomplish manipulation tasks, it can also provide stability to the robot in otherwise challenging scenarios. In this experiment, the robot grabs onto a railing and walks across a slippery floor at a constant desired speed. During the experiment, the manipulator can slide along the railing and apply forces tangential to but not along the railing. STORMMAP optimizes for the manipulation forces at the end-effector in both the y and z directions. The force in the y direction keeps the ZMP within the support polygon while reducing the necessary sway of the base. The force in the z direction also changes the position of the ZMP and keeps it within the support polygon. Additionally, the robot is also able to increase its effective weight by pushing upward on the railing, allowing the robot to exert larger tangential forces at the feet to accelerate forward without having its feet slip on the surface. No forces are applied in the x direction, as we have specified this as a direction of free motion in the constraint in (14).

D. Time Performance

The planning horizon was approximately 0.8 seconds across every test. The optimization took on average 5 ms to compute motion trajectories and manipulation force plans, and rarely exceeding a planning time of 10 ms.

The speed of the optimization is such that the STORMMAP is able to run in a receding-horizon fashion. Once a plan is successfully computed, it is stored and used for the robot's controller and the computation of a new plan starts. STORMMAP is, to the best of our knowledge, an order of magnitude faster compared with other planners from the literature based on planning speed [13]–[15], [17], [18].

VI. DISCUSSION AND CONCLUSION

This paper describes a legged mobile manipulation planner, STORMMAP, which can compute plans for both the motion trajectory of the robot as well as the manipulation forces it exerts. This is done faster than previous dynamic planning strategies and can run in real-time. By formulating the planner as a nonlinear optimization problem with pre-specified contact locations for the manipulator's end-effector, we can plan dynamically feasible motions and forces while ensuring the stability of the platform.

In the future we will include a parallel planner to optimize for the contact locations of the manipulator. Additionally, we will test STORMMAP on a physical robot.

¹Supplementary video can be found here: <https://youtu.be/A9cBz8uuJAs>

REFERENCES

- [1] M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, *et al.*, “Anymal-a highly mobile and dynamic quadrupedal robot,” in *International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2016, pp. 38–44.
- [2] A. Campeau-Lecours, H. Lamontagne, S. Latour, P. Fauteux, V. Maheu, F. Boucher, C. Deguire, and L.-J. C. L’Ecuyer, “Kinova modular robot arms for service robotics applications,” in *Rapid Automation: Concepts, Methodologies, Tools, and Applications*, IGI Global, 2019, pp. 693–719.
- [3] C. D. Bellicoso, K. Krämer, M. Stäubli, D. Sako, F. Jenelten, M. Bjelonic, and M. Hutter, “Alma-articulated locomotion and manipulation for a torque-controllable robot,” in *International Conference on Robotics and Automation (ICRA)*, IEEE, 2019, pp. 8477–8483.
- [4] C. Semini, N. G. Tsagarakis, E. Guglielmino, M. Focchi, F. Cannella, and D. G. Caldwell, “Design of hyq—a hydraulically and electrically actuated quadruped robot,” *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, vol. 225, no. 6, pp. 831–849, 2011.
- [5] B. U. Rehman, M. Focchi, J. Lee, H. Dallali, D. G. Caldwell, and C. Semini, “Towards a multi-legged mobile manipulator,” in *International Conference on Robotics and Automation (ICRA)*, IEEE, 2016, pp. 3618–3624.
- [6] J. Hurst, “Walk this way: To be useful around people, robots need to learn how to move like we do,” *IEEE Spectrum*, vol. 56, no. 3, pp. 30–51, 2019.
- [7] J. Engelsberger, A. Werner, C. Ott, B. Henze, M. A. Roa, G. Garofalo, R. Burger, A. Beyer, O. Eiberger, K. Schmid, and A. Albu-Schäffer, “Overview of the torque-controlled humanoid robot toro,” in *International Conference on Humanoid Robots*, IEEE, 2014, pp. 916–923.
- [8] T. Takubo, K. Inoue, and T. Arai, “Pushing an object considering the hand reflect forces by humanoid robot in dynamic walking,” in *International Conference on Robotics and Automation (ICRA)*, IEEE, 2005, pp. 1706–1711.
- [9] B. J. Stephens and C. G. Atkeson, “Dynamic balance force control for compliant humanoid robots,” in *International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2010, pp. 1248–1255.
- [10] K. Inomata and Y. Uchimura, “3dzmp-based control of a humanoid robot with reaction forces at 3-dimensional contact points,” in *International Workshop on Advanced Motion Control (AMC)*, IEEE, 2010, pp. 402–407.
- [11] K. Harada, S. Kajita, K. Kaneko, and H. Hirukawa, “Zmp analysis for arm/leg coordination,” in *International Conference on Intelligent Robots and Systems (IROS)*, IEEE, vol. 1, 2003, pp. 75–81.
- [12] K. Harada, S. Kajita, K. Kaneko, and H. Hirukawa, “Pushing manipulation by humanoid considering two-kinds of zmps,” in *International Conference on Robotics and Automation (ICRA)*, vol. 2, 2003, 1627–1632 vol.2.
- [13] Y. Lin and D. Berenson, “Humanoid navigation planning in large unstructured environments using traversability - based segmentation,” in *International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2018, pp. 7375–7382.
- [14] Y.-C. Lin, L. Righetti, and D. Berenson, *Robust humanoid contact planning with learned zero- and one-step capturability prediction*, 2019.
- [15] K. Bouyarmane and A. Kheddar, “Humanoid robot locomotion and manipulation step planning,” *Advanced Robotics*, vol. 26, no. 10, pp. 1099–1126, 2012.
- [16] R. Deits and R. Tedrake, “Footstep planning on uneven terrain with mixed-integer convex optimization,” in *International Conference on Humanoid Robots*, IEEE, 2014, pp. 279–286.
- [17] B. Ponton, A. Herzog, S. Schaal, and L. Righetti, *A convex model of momentum dynamics for multi-contact motion generation*, 2016.
- [18] H. Dai, A. Valenzuela, and R. Tedrake, “Whole-body motion planning with centroidal dynamics and full kinematics,” in *International Conference on Humanoid Robots*, IEEE, 2014, pp. 295–302.
- [19] H. Hahn, *Rigid Body Dynamics of Mechanisms: 1 Theoretical Basis*. Springer Science & Business Media, 2013.
- [20] T. Apgar, P. Clary, K. Green, A. Fern, and J. W. Hurst, “Fast online trajectory optimization for the bipedal robot cassie,” in *Robotics: Science and Systems*, 2018.
- [21] C. Mastalli, I. Havoutis, M. Focchi, D. G. Caldwell, and C. Semini, “Motion planning for quadrupedal locomotion: Coupled planning, terrain mapping, and whole-body control,” *IEEE Transactions on Robotics*, 1–14, 2020.
- [22] M. Vukobratović and J. Stepanenko, “On the stability of anthropomorphic systems,” *Mathematical biosciences*, vol. 15, no. 1-2, pp. 1–37, 1972.
- [23] C. Gehring, S. Coros, M. Hutter, C. Dario Bellicoso, H. Heijnen, R. Diethelm, M. Bloesch, P. Fankhauser, J. Hwangbo, M. Hoepfner, and R. Siegwart, “Practice makes perfect: An optimization-based approach to controlling agile motions for a quadruped robot,” *Robotics Automation Magazine*, vol. 23, no. 1, pp. 34–43, 2016.
- [24] M. Kalakrishnan, J. Buchli, P. Pastor, M. Mistry, and S. Schaal, “Fast, robust quadruped locomotion over challenging terrain,” in *International Conference on Robotics and Automation (ICRA)*, IEEE, 2010, pp. 2665–2670.
- [25] C. D. Bellicoso, F. Jenelten, P. Fankhauser, C. Gehring, J. Hwangbo, and M. Hutter, “Dynamic locomotion and whole-body control for quadrupedal robots,” in *International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2017, pp. 3359–3365.
- [26] C. D. Bellicoso, F. Jenelten, C. Gehring, and M. Hutter, “Dynamic locomotion through online nonlinear motion optimization for quadrupedal robots,” *Robotics and Automation Letters*, vol. 3, no. 3, pp. 2261–2268, 2018.
- [27] J. C. Trinkle, J.-S. Pang, S. Sudarsky, and G. Lo, “On dynamic multi-rigid-body contact problems with coulomb friction,” *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, vol. 77, no. 4, pp. 267–279, 1997.
- [28] L. D. Gaspero, *Quadprog++*, 1998.
- [29] D. Bellicoso, “Optimization-based planning and control for multi-limbed walking robots,” Ph.D. dissertation, ETH Zurich, 2019.