Collision-Free MPC for Legged Robots in Static and Dynamic Scenes

Conference Paper**Author(s):**

Gaertner, Magnus; Bjelonic, Marko; Farshidian, Farbod; [Hutter, Marco](#)

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000491764>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/ICRA48506.2021.9561326>

Funding acknowledgement:

166232 - Data-driven control approaches for advanced legged locomotion (SNF)

188596 - Perceptive Dynamic Locomotion on Rough Terrain (SNF)

852044 - Learning Mobility for Real Legged Robots (EC)

Collision-Free MPC for Legged Robots in Static and Dynamic Scenes

Magnus Gaertner, Marko Bjelonic, Farbod Farshidian, Marco Hutter

Abstract—We present a model predictive controller (MPC) that automatically discovers collision-free locomotion while simultaneously taking into account the system dynamics, friction constraints, and kinematic limitations. A relaxed barrier function is added to the optimization’s cost function, leading to collision avoidance behavior without increasing the problem’s computational complexity. Our holistic approach does not require any heuristics and enables legged robots to find whole-body motions in the presence of static and dynamic obstacles. We use a dynamically generated euclidean signed distance field for static collision checking. Collision checking for dynamic obstacles is modeled with moving cylinders, increasing the responsiveness to fast-moving agents. Furthermore, we include a Kalman filter motion prediction for moving obstacles into our receding horizon planning, enabling the robot to anticipate possible future collisions. Our experiments demonstrate collision-free motions on a quadrupedal robot in challenging indoor environments. The robot handles complex scenes like overhanging obstacles and dynamic agents by exploring motions at the robot’s dynamic and kinematic limits.

I. INTRODUCTION

Legged robots are fast becoming more common in industrial environments [1], and it is only a matter of time until we see more of these robots in urban environments [2]. A necessity for applications in these environments is compliance with *static* environments and *dynamic* agents such as pedestrians moving side by side with the robot. Also, in remote operations where latency to the operator impedes fully supervised teleoperation, an automatic motion planner improves safety and autonomy.

The potential advantages of legged robots (as shown in Fig. 1) in mobility and versatility compared to wheeled robots come at the cost of higher complexity for the locomotion task. We propose a whole-body Model Predictive Control (MPC) that simultaneously solves torso and feet motions while being responsive to the environment. Thus, kinematic reachability, dynamic stability, and obstacle avoidance are solved as a whole, leading to an optimal behavior. Moreover, the formulation automatically discovers complex and dynamic motions cumbersome to handcraft through heuristics. The robot successfully anticipates possible future collisions and generates avoidance maneuvers automatically. We show that casting collision avoidance into an optimization problem increases the computational cost of the MPC negligibly. But more importantly, it avoids laboriously handcrafting new behaviors.

Fig. 1. Our quadrupedal robot, *ANYmal*, is navigating collision-free through narrow passages in an office environment. First column: Collision-free locomotion while traversing lateral and rotational through a crowded environment. Top right: SDF and sphere decomposition of the robot used to capture the interactions with the environment. Bottom right: Observed trajectory in a narrow corridor while giving a coarse direction input.

A. Related Work

In the following, we categorize existing approaches to collision-free motion planning by their map representation and collision avoidance algorithm and bring them into the context of legged locomotion planning.

1) *Environment representation*: The environment’s representation is a crucial element, and multiple volumetric map representations are available, e.g., octree [3], height-map [4], and euclidean signed distance field (ESDF). ESDFs seem to become increasingly crucial in legged robotics as they can deal with the sensor’s noise in an inherent fashion [5], [6]. Directly storing the distance measurements in the representation enables averaging multiple measurements, thereby filtering the sensor’s noise. *Voxblox* [7] and *FIESTA* [8] are open-source programs that incrementally update an ESDF.

Besides volumetric approaches that are typically aiming for a globally consistent map, local representations are meant to query only local geometric information. One of such representations is *NanoMap* [9]. Here distances are directly queried from laser scans. It introduces less computational overhead as no explicit map representation is build. On the downside, querying takes longer than on a distance field. Another approach is to model individual objects of

All authors are with the Robotic Systems Lab, ETH Zürich, 8092 Zürich, Switzerland. firstname.lastname@mavt.ethz.ch

¹Supplementary video: https://youtu.be/_wkqCVz3gdg

the scene explicitly. For example, the work in [10] extracts line segments and circles from 2D laser scans and tracks their respective positions and velocities. Typically, objects are approximated by geometric primitives such as planes, spheres, ellipsoids, or polytopes.

2) *Optimization-based collision avoidance*: A trajectory optimization method like *Covariant Hamiltonian Optimization for Motion Planning (CHOMP)* [11] is widely used for robot motion planning in the presence of obstacles and other constraints. The trajectory costs are invariant to the time parameterization of the trajectory and use pre-computed SDFs. In contrast, the work in [12], [13] introduces a collision-free trajectory planner, *TrajOpt*. It proposes a sequential convex optimization-based motion planner. Moreover, the algorithm penalizes collisions with a hinge loss. Instead of SDF, a convex-convex collision detection is employed to compute minimal translation to avoid collision and directly considers continuous-time safety. The paper in [14] presents a novel method for reformulating non-differentiable collision avoidance constraints into smooth nonlinear constraints.

The no-collisions constraint can be implemented by constraining the robot's position to be in a safe region. Another approach is to augment the objective with a repulsive potential, typically proportional to the inverse square distance. The former is shown by [15] efficiently as a pairwise distance constraint, and the work in [16] exemplifies the latter together with a constraint on the relative velocity.

The implementation in [17] computes the distance between the robot and the environment by discriminating between dynamic obstacles represented by orientable ellipsoids and static obstacles represented by a 2D occupancy grid map. Similarly, the authors in [18] use orientable ellipsoids for moving obstacles and surface normals for static obstacles. A soft constraint with a shared slack variable among all obstacles is introduced; sharing the slack reduces the computational cost when dealing with multiple obstacles.

The recent work in [19] shows collision-free planning of end-effector trajectories of a mobile manipulator. The authors add soft inequality constraints with relaxed barrier functions to the MPC cost, and a set of collision spheres to approximate the robot.

3) *Collision-free locomotion planning*: The task of finding the continuous-time motion of legged robots, i.e., the trajectories of the torso and the feet, can be achieved by decomposing the problem or by treating the whole problem as a single task. Decomposed approaches separate the feet and torso planning into two different stages, a common practice to make the collision-free locomotion planning more tractable in terms of the optimization's dimensionality. The full set of possible solutions, however, can not be discovered by such a method, since the individual components do not take into account the full set of physical constraints.

In [5], the CHOMP algorithm generates the shortest collision-free trajectory of the foot's swing trajectory, and the work in [6] extends the work to collision-free planning of the torso trajectory. The authors generate a SDF through a heightmap, and overhanging obstacles are handled through

a second 2.5D map of the ceiling. In both approaches, the authors decompose the problem into separate torso and feet trajectory planning. The coordination of each task's solution is one of the main challenges, and heuristics are needed that are cumbersome to handcraft.

By decomposing the locomotion and navigation planner, the planning problem results in similar challenges since a path follower needs to be designed that guides the locomotion planner onto the collision-free path. Recent advancements in robotics competitions are relying on such a decomposition [20], [21]. However, these approaches cannot explore highly dynamic motions at their limits, such as using the system's friction and kinematic limits to escape moving obstacles.

B. Contributions

We present a MPC for legged robots with a novel way to incorporate static and dynamic environments. We treat the whole locomotion planning problem as a holistic optimization, thus simultaneously taking into account the scene, dynamic feasibility, and kinematic reachability. To avoid obstacles, a soft inequality constraint is added as a cost to the switched system. The environment representation is divided into a *static* and *dynamic* scene to increase the responsiveness to fast-moving obstacles. The former uses a dynamically generated ESDF for static collision checking, and the latter is based on a moving cylinder representation per obstacle. Our framework extends the capabilities of legged robots in the following ways:

- The holistic approach automatically discovers complex, collision-free, and dynamic motions without needless heuristics with minimum computational overhead.
- Our whole-body approach enables the robot to explore the full six degrees of freedom (DOF) of the torso to handle overhanging obstacles.
- Our formulation allows to include the prediction of moving obstacles into our receding horizon planning and enables the robot to anticipate possible future collisions before they take place.
- The MPC improves the safety and autonomy in narrow passages.

II. PROBLEM FORMULATION

At each iteration of the MPC, we solve the following optimization problem based on the latest state measurement $\mathbf{x}_s$ at time t_s . We then apply the optimized control policy until a new MPC update arrives.

$$\underset{\mathbf{u}(\cdot)}{\text{minimize}} \quad \int_{t_s}^{t_f} l(\mathbf{x}(t), \mathbf{u}(t), t) dt, \quad (1a)$$

$$\text{subject to} \quad \mathbf{x}(t_s) = \mathbf{x}_s, \quad (1b)$$

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, t), \quad (1c)$$

$$\mathbf{g}(\mathbf{x}, \mathbf{u}, t) = \mathbf{0} \quad (1d)$$

$$\mathbf{h}(\mathbf{x}, \mathbf{u}, t) \geq \mathbf{0}, \quad (1e)$$

where $\mathbf{x}(t)$ is the state and $\mathbf{u}(t)$ is the input at time t . l is a time-varying running cost. The goal is to find a continuous

control policy that minimizes this cost subject to the initial condition (1b), system dynamics (1c), and general equality (1d) and inequality (1e) constraints.

In this paper, we employ the Sequential Linear Quadratic (SLQ) algorithm, which is a Differential Dynamic Programming (DDP) based algorithm for continuous-time systems [22]. As opposed to Nonlinear Programming (NLP) this approach does not require transcription of the continuous state and input but chooses an adaptive step-size given the desired accuracy. For MPC formulation, we rely on a real-time iteration scheme where for each MPC cycle, one iteration of SLQ optimization is performed [23]. In particular we use the MPC-SLQ approach [24] with feedback policy [25].

A. Blind MPC

In this section, we briefly describe the previous MPC formulation [24], [25] on which we build our addition for collision avoidance described in Sec. III.

1) *System dynamics*: We apply our approach to the kinodynamic model of a quadrupedal robot, which describes the dynamics of a single free-floating body along with the full kinematics for each leg. The state $\mathbf{x} \in \mathbb{R}^{24}$ and input $\mathbf{u} \in \mathbb{R}^{24}$ are defined as

$$\begin{aligned} \mathbf{x} &= (\boldsymbol{\theta}, \mathbf{p}, \boldsymbol{\omega}, \mathbf{v}, \mathbf{q}), \\ \mathbf{u} &= (\boldsymbol{\lambda}_{ee}^1, \boldsymbol{\lambda}_{ee}^2, \boldsymbol{\lambda}_{ee}^3, \boldsymbol{\lambda}_{ee}^4, \mathbf{v}_{joint}), \end{aligned} \quad (2)$$

and the equations of motion are given by

$$\begin{cases} \dot{\boldsymbol{\theta}} = \mathbf{T}(\boldsymbol{\theta})\boldsymbol{\omega} \\ \dot{\mathbf{p}} = {}_W\mathbf{R}_B(\boldsymbol{\theta})\mathbf{v} \\ \dot{\boldsymbol{\omega}} = \mathbf{I}^{-1} \left(-\boldsymbol{\omega} \times \mathbf{I}\boldsymbol{\omega} + \sum_{i=1}^4 \mathbf{r}_{ee}^i \times \boldsymbol{\lambda}_{ee}^i \right) \\ \dot{\mathbf{v}} = \mathbf{g}(\boldsymbol{\theta}) + \frac{1}{m} \sum_{i=1}^4 \boldsymbol{\lambda}_{ee}^i \\ \dot{\mathbf{q}} = \mathbf{v}_{joint}, \end{cases} \quad (3)$$

where ${}_W\mathbf{R}_B$ and $\mathbf{T}$ are the rotation matrix from the base to the world frame, and the transformation matrix from base local angular velocities to the Euler angles derivatives in the world frame. $\mathbf{g}$ is the gravitational acceleration in body frame, $\mathbf{I}$, and m are the moment of inertia about the Center of Mass (CoM) and the total mass, respectively. The inertia is assumed to be constant and taken at the default configuration of the robot. $\mathbf{r}_{ee}^i$ is the position of the foot i with respect to (w.r.t.) CoM. $\boldsymbol{\theta}$ is the orientation of the base in Euler angles², $\mathbf{p}$ is the position of the CoM in world frame, $\boldsymbol{\omega}$ is the angular rate, and $\mathbf{v}$ is the linear velocity of the CoM. $\mathbf{q} \in \mathbb{R}^{12}$ is the joint positions. The inputs of the model are the joint velocities $\mathbf{v}_{joint} \in \mathbb{R}^{12}$ and end-effector contact forces $\{\boldsymbol{\lambda}_{ee}^i\}_{i=1}^4$ in the body frame.

2) *Cost function*: The intermediate cost function is formulated as a quadratic function

$$l(\mathbf{x}, \mathbf{u}, t) = \frac{1}{2} \|\mathbf{x} - \mathbf{x}_d\|_{\mathbf{Q}}^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0\|_{\mathbf{R}}^2, \quad (4)$$

²The reference frame is chosen, such that a gimbal lock appears in the unlikely "backflip" position.

Fig. 2. Integration of the static and dynamic environment. Two map servers asynchronously update the environment representations from new sensor data. The updates on the representation are sent to the MPC node, where they are cached. On each iteration of MPC, the most recent cached version is used in the optimization. The augmentation for the collision-free MPC is depicted in blue.

where $\mathbf{R}$ is positive definite and $\mathbf{Q}$ are positive semi-definite weighting matrices. $\mathbf{x}_d$ is the desired state computed from a user-defined velocity command and a default configuration for the joints. $\mathbf{u}_0$ is the equilibrium input vector consisting of the statically stabilizing contact forces and zero joint velocities. Finally, the planning horizon is chosen one second i.e., $t_f - t_s = 1$.

3) *System constraints*: The constraints capture the different modes of each leg at any given point in time. We assume that the mode sequence is a predefined function of time, and we denote the set of all closed contacts at a given time by $\mathcal{C}$. The resulting constraints are

$$\begin{cases} \mathbf{v}_{ee}^i(\mathbf{x}, \mathbf{u}) = \mathbf{0}, & \boldsymbol{\lambda}_{ee}^i \in \mathcal{Q}(\hat{\mathbf{n}}, \mu_c), \quad \text{if } i \in \mathcal{C}, \\ \mathbf{v}_{ee}^i(\mathbf{x}, \mathbf{u}) \cdot \hat{\mathbf{n}} = c(t), & \boldsymbol{\lambda}_{ee}^i = \mathbf{0}, \quad \text{if } i \notin \mathcal{C}, \end{cases} \quad (5)$$

where $\mathbf{v}_{ee}^i$ is the end-effector velocity in the world frame, these constraints ensure that stance legs remain on the ground and swing legs follow a predefined curve, $c(t)$, in the direction of the local surface normal $\hat{\mathbf{n}} \in \mathbb{R}^3$ to avoid foot scuffing. Furthermore, we enforce zero contact force for swing legs and a friction cone constraint, $\mathcal{Q}$, for stance legs defined by the surface normal and friction coefficient $\mu_c = 0.7$.

III. COLLISION-FREE MPC

In this section, we present the collision-free MPC formulation that incorporates an additional constraint to achieve collision avoidance. For collision checking, we use a sphere decomposition of the robot and query the distance to the individual spheres' closest obstacle. In the rest of the section, we first present our collision avoidance formulation. We then discuss two different representations of the environment for static and dynamic scenes.

A. Collision Avoidance

Obstacle avoidance is handled as a soft constraint and added to the previously blind MPC controller by augmenting the objective with the following penalty function

$$l_{\text{aug}}(\mathbf{x}, \mathbf{u}, t) = l(\mathbf{x}, \mathbf{u}, t) + l_c(\mathbf{x}, t), \quad (6)$$

Fig. 3. Left and center: Sphere decomposition of the robot. Right: A Squared-hinge penalty function is used as a barrier function in the objective. It is parametrized by the scaling parameter μ and a safety margin ϵ . A distance $d_i < 0$ indicates a collision.

where l_{aug} is our augmented intermediate cost function, l the previous intermediate cost from (4) and l_c our additional penalty function. For l_c we use a simple, one-sided quadratic barrier function

$$l_c(\mathbf{x}, t) = \sum_{i=1}^p \mu \cdot \frac{1}{2} (\max\{0, \epsilon - d_i(\mathbf{x}, t)\})^2 \quad (7)$$

where p is the number of spheres in the sphere decomposition of the body, and d_i is the distance of the sphere i to the closest obstacle. μ and ϵ are the scaling and safety-margin parameters illustrated in Fig. 3. We used a safety margin of 10 cm and 15 cm for experiments in simulation and real-world, respectively.

In this work, we decompose the robot's main body by six spheres of uniform size ordered in a two-by-three, largely overlapping grid (see Fig. 3). Moreover, two additional spheres are added as collision bodies to account for the front handle and the back's lidar cage.

To deal with the occasional conflict between the collision avoidance and the position tracking objectives, we clip the position tracking error $\|\mathbf{p} - \mathbf{p}_d\|$ to a maximum threshold of 15 cm. Doing so, we prevent a sudden pull towards the commanded trajectory once the robot negotiates the blocking obstacle.

Finally, as the SLQ algorithm requires the Hessian of the penalty w.r.t. the robot state to be positive semi-definite, we use a *Gauss-Newton* approximation to ensure this property.

$$\frac{\partial^2 l_c}{\partial \mathbf{x}^2} = \sum_{i=1}^p \mathcal{H}(\epsilon - d_i) \left(\mu \frac{\partial d_i}{\partial \mathbf{x}} \frac{\partial d_i}{\partial \mathbf{x}}^T + \frac{\partial l_c}{\partial d_i} \frac{\partial^2 d_i}{\partial \mathbf{x}^2} \right), \quad (8a)$$

$$\frac{\partial^2 l_c}{\partial \mathbf{x}^2} \approx \sum_{i=1}^p \mathcal{H}(\epsilon - d_i) \left(\mu \frac{\partial d_i}{\partial \mathbf{x}} \frac{\partial d_i}{\partial \mathbf{x}}^T \right), \quad (8b)$$

where $\mathcal{H} : \mathbb{R} \rightarrow \{0, 1\}$ is the Heaviside step function.

B. Environment Representation

To have the fidelity of a static environment representation and a notion of speed for dynamic obstacles, we combine two environments in a compound collision environment (CCE). The CCE forwards all environmental information, distance and gradient, from the single environment that returns the smallest distance to the MPC formulation.

The environment representations are updated asynchronously in separate nodes and updates are sent to the MPC node. Our setup is depicted in Fig. 2.

1) *Static environment*: We use *Voxblox* [7] to generate an ESDF on a separate Robot Operating System (ROS) node, denoted as *Voxblox Server*. The *Voxblox Server* repeatedly sends updates of the computed ESDF to the MPC node. Upon receiving a new map, we pre-compute the gradient for each voxel of the distance-field and cache it alongside its distance to make the queries within the MPC optimization faster. This caching functionality is taken from [19], although we clip the norm of this gradient to be within its theoretical limits $[0, 1]$ to compensate for errors introduced by finite-difference approximation and the voxelization of the distance field³. As the computation time of the distance field update is cubic in the voxel size, we need to trade-off between update rate and resolution. We use an update rate of 4 Hz and a voxel size of 10 cm.

At query time, the distance is computed using the first-order approximation from the closest voxel center. Gradients are approximated by the pre-computed gradient of the closest voxel.

2) *Dynamic environment*: Although the use of an ESDF representation seems best suited for our requirements, it has one limitation that is hard to overcome - it is inherently static. Consequently, we can not incorporate predictions of the state in the future into the receding time horizon of the controller. Instead, we use a simplistic, cylindrical representation for dynamic obstacles. By introducing an object-based representation for moving obstacles, we can track their trajectories and predict their motions. For this, we make use of an open-source ROS package *obstacle_detector* [10] that takes ordered, two-dimensional laser-scans. We use their *obstacle_publisher* auxiliary node to publish a set of virtually detected circles for the experiments in simulation. The state of an obstacle is the two-dimensional position, velocity, and radius. We use a constant-velocity model to predict the future location of these moving obstacles and integrate it into our receding horizon MPC formulation.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

We show that our formulation discovers complex behaviors leveraging the full dynamic and kinematic capabilities of the robot. Experiments have been carried out both in simulation and real-world on our quadrupedal robot, ANYmal [26]. In this section, we first study the computational overhead incurred by our perceptive formulation compared to the blind version. We then evaluate the performance of our controller in different environments with static and dynamic obstacles.

A. Computational Cost

Benchmarking our controller against a blind MPC reveals that the computation time is only slightly increased. While operating at a frequency of 20 Hz, our collision-free

³In a static environment, the distance to any object will not change more than the distance directly traveled towards it. In particular, the norm is exactly one, whenever there is exactly one closest object.

Fig. 4. Collection of different emerging behaviors during real-world experiments. At each snapshot the commanded (green) and optimized (red) trajectories are shown. First row, lateral commands: The robot adjusts its yaw and position to avoid obstacles along the commanded trajectory. Middle row, rotational commands: The center of rotation is adjusted dynamically to both follow the desired command closely but not collide with the environment. Bottom row: The robot followed more diverse inputs while avoiding obstacles.

Fig. 5. Computational comparison of the collision-free and blind MPC. The serial computation time of the solver is split up into its main parts. Bars show the average of over 1000 solver iterations. Black error bars indicate one standard deviation. On average, the solving time increased only marginally by 0.93 ms. For this benchmarking, we use AMD Ryzen 5 2600x, 16Gb DDR4 RAM, 64bit.

MPC formulation introduces a computational overhead in the backward pass and in searching a strategy amounting to a total increase of 6% or 0.93 ms. At the same time, it breaks down the complexity of the control structure and engineering effort required for connecting inherently-dependent but independently-designed path navigation and tracking modules. More detailed timings are shown in Fig. 5.

B. Whole-Body Planning and Look-Ahead

Our formulation shows its potential best while locomoting through narrow passages and more complex maneuvers such as ducking and dynamic collision avoidance. Thereby, the

Fig. 6. ANYmal ducking under the overhanging obstacle. While commanded to go forward, the robot automatically adjusts its base height and orientation to avoid a collision with the overhanging obstacle. Afterwards, the robot adjusts to its nominal posture.

whole body optimization leverages the inter-dependencies between robot dynamics and collision-free trajectories. We carried out two evaluations in simulation, for ducking under overhanging obstacles (see Fig. 6) and avoidance of dynamic obstacles (see Fig. 7). For those experiments, the robot was given only a constant forward velocity command and no explicit final position was given or updated to avoid the obstacles. Moreover we carried out several experiments in real world (see Fig. 4 and the attached video).

1) *Overhanging obstacles*: The robot automatically ducks under overhanging obstacles and returns to its nominal posture once the obstacle is passed. In contrast to the 2.5D height map in [6], our environment representation based on the SDF can handle the full 3D space, including overhanging obstacles.

2) *Dynamic obstacles*: As mentioned before, our collision-free MPC allows the integration of the prediction of moving obstacles into its receding horizon planning. To showcase this property, we devise two tests where ANYmal should pass through two approaching obstacles from opposite sides. In both cases, the robot is commanded forward with

Fig. 7. Sequences of dynamic obstacle avoidance maneuvers. Top, left to right: Obstacles approaching at 0.5 m/s. The robot is trotting in place waiting for the obstacles to pass, although continuously commanded to go ahead. Note that at (2a) - (4a), the robot even moves backward. Bottom, left to right: Obstacles approaching at 0.2 m/s. The robot accelerates (2b) to pass through the obstacles to then decelerate (4b), (5b) to be again aligned with the desired trajectory.

Fig. 8. Top: Distance of the robot to the closest obstacle. We observe that the distance stays well above zero. Bottom: Distance penalty associated with the current state. For distances above 0.15 m the cost is zero. A Moving average is applied over a time window of 1.6 s.

constant velocity. In the first case, the objects are moving faster (0.5 m/s) than in the second case (0.2 m/s), see Fig. 7. In both cases, the robot automatically anticipates the timing when it is possible to surpass the moving obstacles without collision. We can demonstrate that the robot prefers to surpass these obstacles when the dynamic obstacles move at a slower speed. At higher speeds, the robot trots in place till the obstacles have passed before continuing its path. This highlights that in this formulation, depending on the scene and system's limits (in this case, friction cone limits and the actuation cost), different behavior can emerge within the same optimization problem.

C. Hardware Experiments

Several real-world experiments have shown that our planner performs comparably w.r.t. the simulation results on real hardware. We observed motions like sidestepping, aligning to walls, and backward movement in order to execute a commanded turn. Fig. 4 gives a qualitative overview of the observed motions. While navigating through artificially narrowed down office floors (Fig. 1) for over five minutes, we had only one minor collision with a very slim obstacle

that was not properly mapped. Fig. 8 shows the minimum distance and its associated cost for this experiment. Encountered map inconsistencies caused by an odometry drift were partially solved by a map update rate of 4 Hz and trusting more on the recently received sensor data. The collision-free MPC increases the safety factor of the operation and helps the operator go through cluttered space.

V. CONCLUSIONS

We introduce a collision-free MPC for legged robots that enables our four-legged robot, ANYmal, to navigate cluttered environments and deal with dynamic obstacles. In contrast to existing approaches, the problem is solved as a holistic optimization problem that considers the collision-free space, dynamic feasibility, and kinematic constraints over a receding horizon. With this formulation, we can include the prediction of moving obstacles, which enables the robot to avoid future collisions before they occur. To achieve collision avoidance, we add a relaxed barrier function to the switched system's MPC formulation without increasing the problem's computational complexity. ANYmal manages to automatically discover complex motions, avoiding cluttered scenes with overhanging and moving obstacles. Also, the real-world tests verify that the algorithm can cope with noisy data from exteroceptive sensors. All in all, the collision-free MPC improves the safety and autonomy of the robot.

In future work, we plan to incorporate collision-free planning of the swing feet trajectory to negotiate rough terrain. Furthermore, we aim to employ more sophisticated prediction strategies [27] for moving obstacles based on a history of object motion making a better prediction of their future path beyond the first-order model used in this work.

ACKNOWLEDGMENTS

This research was supported by the Swiss National Science Foundation (SNSF) through project 166232, 188596, the National Centre of Competence in Research Robotics (NCCR Robotics), and the European Union's Horizon2020 (grant agreement No. 852044). Moreover, this work has been conducted as part of ANYmal Research, a community to advance legged robotics.

We appreciate the insightful discussion and support from Marco Tranzatto, ETH Zürich.

REFERENCES

- [1] C. D. Bellicoso, M. Bjelonic, L. Wellhausen, K. Holtmann, F. Günther, M. Tranzatto, P. Fankhauser, and M. Hutter, "Advances in real-world applications for legged robots," *Journal of Field Robotics*, vol. 35, no. 8, pp. 1311–1326, 2018.
- [2] BBC News, "Coronavirus: Robot dog enforces social distancing in singapore park," <https://www.bbc.com/news/av/technology-52619568>, BBC News, 2020, [Online; last accessed 09-March-2021].
- [3] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3d mapping framework based on octrees," *Autonomous robots*, vol. 34, no. 3, pp. 189–206, 2013.
- [4] P. Fankhauser, M. Bloesch, and M. Hutter, "Probabilistic terrain mapping for mobile robots with uncertain localization," *IEEE Robotics and Automation Letters (RA-L)*, vol. 3, no. 4, pp. 3019–3026, 2018.
- [5] P. Fankhauser, M. Bjelonic, C. Dario Bellicoso, T. Miki, and M. Hutter, "Robust rough-terrain locomotion with a quadrupedal robot," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 5761–5768.

- [6] R. Buchanan, L. Wellhausen, M. Bjelonic, T. Bandyopadhyay, N. Kottege, and M. Hutter, "Perceptive whole-body planning for multilegged robots in confined spaces," *Journal of Field Robotics*, 2020.
- [7] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board map planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [8] L. Han, F. Gao, B. Zhou, and S. Shen, "Fiesta: Fast incremental euclidean distance fields for online motion planning of aerial robots," *arXiv preprint arXiv:1903.02144*, 2019.
- [9] P. R. Florence, J. Carter, J. Ware, and R. Tedrake, "Nanomap: Fast, uncertainty-aware proximity queries with lazy search over local 3d data," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 7631–7638.
- [10] M. Przybyła, "Detection and tracking of 2d geometric obstacles from lrf data," in *2017 11th International Workshop on Robot Motion and Control (RoMoCo)*. IEEE, 2017, pp. 135–141.
- [11] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *IEEE International Conference on Robotics and Automation*, 2009, pp. 489–494.
- [12] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization," in *Robotics: science and systems*, vol. 9, no. 1. Citeseer, 2013, pp. 1–10.
- [13] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [14] X. Zhang, A. Liniger, and F. Borrelli, "Optimization-based collision avoidance," *IEEE Transactions on Control Systems Technology*, pp. 1–12, 2020.
- [15] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with centroidal dynamics and full kinematics," in *2014 IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2014, pp. 295–302.
- [16] A. C. Woods, H. M. Lay, and Q. P. Ha, "A novel extended potential field controller for use on aerial robots," in *2016 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2016, pp. 286–291.
- [17] B. Brito, B. Floor, L. Ferranti, and J. Alonso-Mora, "Model predictive contouring control for collision avoidance in unstructured dynamic environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4459–4466, 2019.
- [18] M. Castillo-Lopez, S. A. Sajadi-Alamdari, J. L. Sanchez-Lopez, M. A. Olivares-Mendez, and H. Voos, "Model predictive control for aerial collision avoidance in dynamic environments," in *2018 26th Mediterranean Conference on Control and Automation (MED)*. IEEE, 2018, pp. 1–6.
- [19] J. Pankert and M. Hutter, "Perceptive model predictive control for continuous mobile manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6177–6184, 2020.
- [20] M. Hutter, R. Diethelm, S. Bachmann, P. Fankhauser, C. Gehring, V. Tsounis, A. Lauber, F. Guenther, M. Bjelonic, L. Isler, *et al.*, "Towards a generic solution for inspection of industrial sites," in *Field and Service Robotics*. Springer, 2018, pp. 575–589.
- [21] A. Bouman, M. F. Ginting, N. Alatur, M. Palieri, D. D. Fan, T. Touma, T. Pailevanian, S.-K. Kim, K. Otsu, J. Burdick, *et al.*, "Autonomous spot: Long-range autonomous exploration of extreme environments with legged locomotion," *submitted to IEEE Robotics and Automation Letters*, 2020.
- [22] F. Farshidian, M. Neunert, A. W. Winkler, G. Rey, and J. Buchli, "An efficient optimal planning and control framework for quadrupedal locomotion," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 93–100.
- [23] M. Diehl, H. Bock, and J. Schlöder, "A real-time iteration scheme for nonlinear optimization in optimal feedback control," *SIAM Journal on Control and Optimization*, vol. 43, no. 5, pp. 1714–1736, 2005.
- [24] F. Farshidian, E. Jelavic, A. Satapathy, M. Gifftthaler, and J. Buchli, "Real-time motion planning of legged robots: A model predictive control approach," in *International Conference on Humanoid Robotics (Humanoids)*. IEEE-RAS, 2017, pp. 577–584.
- [25] R. Grandia, F. Farshidian, R. Ranftl, and M. Hutter, "Feedback mpc for torque-controlled legged robots," in *International Conference on Intelligent Robots and Systems (IROS)*. IEEE/RSJ, 2019, pp. 4730–4737.
- [26] M. Hutter, C. Gehring, A. Lauber, F. Gunther, C. D. Bellicoso, V. Tsounis, P. Fankhauser, R. Diethelm, S. Bachmann, M. Blösch, *et al.*, "Anymal-toward legged robots for harsh environments," *Advanced Robotics*, vol. 31, no. 17, pp. 918–931, 2017.
- [27] D. Hoeller, L. Wellhausen, F. Farshidian, and M. Hutter, "Learning a state representation and navigation in cluttered and dynamic environments," *arXiv preprint arXiv:2103.04351*, 2021.