

Perceptive Locomotion through Whole-Body MPC and Optimal Region Selection

THOMAS CORBÈRES¹, CARLOS MASTALLI², WOLFGANG MERKT³, JAEHYUN SHIM¹, IOANNIS HAVOUTIS³, MAURICE FALLON³, NICOLAS MANSARD⁴, THOMAS FLAYOLS⁴, SETHU VIJAYAKUMAR¹, STEVE TONNEAU¹, (MEMBER, IEEE)

¹School of Informatics, University of Edinburgh, UK

²School of Eng. and Physical Sciences, Heriot-Watt University, UK

³Oxford Robotics Institute, University of Oxford, UK

⁴LAAS-CNRS, France

Corresponding author: T. Corbères (e-mail: t.corberes@sms.ed.ac.uk).

This research was supported by (1) the European Commission under the Horizon 2020 project Memory of Motion (MEMMO, project ID: 780684) and (2) ROBOTEX 2.0 (Grants ROBOTEX ANR-10-EQPX-44-01 and TIRREX-ANR-21-ESRE-0015).

ABSTRACT Real-time synthesis of legged locomotion maneuvers in challenging industrial settings is still an open problem, requiring simultaneous determination of footsteps locations several steps ahead while generating whole-body motions close to the robot's limits. State estimation and perception errors impose the practical constraint of fast re-planning motions in a model predictive control (MPC) framework. We first observe that the computational limitation of perceptive locomotion pipelines lies in the combinatorics of contact surface selection. Re-planning contact locations on selected surfaces can be accomplished at MPC frequencies (50-100 Hz). Then, whole-body motion generation typically follows a reference trajectory for the robot base to facilitate convergence. We propose removing this constraint to robustly address unforeseen events such as contact slipping, by leveraging a state-of-the-art whole-body MPC (Croccodyl). Our contributions are integrated into a complete framework for perceptive locomotion, validated under diverse terrain conditions, and demonstrated in challenging trials that push the robot's actuation limits, as well as in the ICRA 2023 quadruped challenge simulation.

INDEX TERMS Legged locomotion, Motion planning, Quadrupedal robots, Robot vision systems

I. INTRODUCTION

Reliable and autonomous locomotion for legged robots in arbitrary environments is a longstanding challenge. The hardware maturity of quadruped robots [1]–[3] motivates the development of a motion synthesis framework for applications including inspections in industrial areas [4]. Synthesising motions in this context requires handling the issues of *contact decision* (where should the robot step?) and *Whole-Body Model Predictive Control* (WB-MPC) of the robot (what motion creates the contact?).

Each contact decision defines high-dimensional, non-linear geometric and dynamic constraints on the WB-MPC that prevent a trivial decoupling of the two issues: How to prove that a contact plan is valid without finding a feasible whole-body motion to achieve it? Even worse, the environments we consider comprise holes and gaps,

introducing a *combinatorics* problem: On which contact surface(s) should the robot step?

A. STATE OF THE ART

The mathematical complexity of the legged locomotion problem in arbitrary environments often leads to an undesired decoupling between contact planning and whole-body control is required. Typically, a contact plan describing the contact locations is first computed, assumed to be feasible, and provided as input to a WB-MPC framework to generate the whole-body motion. As a result, the contact decision must be made using an approximated robot model, under the uncertainty that results from imperfect perception and state estimation. The complexity of the approximated model has, unsurprisingly, a strong correlation with the versatility and computational efficiency of the proposed approach.



FIGURE 1. Industrial staircase descent with onboard perception. Video: <https://youtu.be/bNVxTh0eIXI>.

1) Offline contact planning with full kinematics

Early approaches to contact planning are robust and complete, because they integrate a whole-body kinematic model in the planning phase, with quasi-static feasibility guarantees. These *contact-before-motion* approaches [5]–[7] mix graph-based search with contact-posture sampling to explicitly tackle the problem’s combinatorics. Whole-body feasibility is explicitly checked before validating each new contact. The generality of these approaches is unmatched, but the associated computational cost is high (from minutes to hours), which prevents online re-planning. The computation time can be reduced to a few seconds by constraining the root’s path [8], and then by approximating the robot with low-dimensional abstractions [9], [10]. These abstractions use hand-crafted heuristics for collision avoidance and geometry [11], while dynamic feasibility is often asserted using centroidal dynamics (e.g., [12], [13]). Such approximations proved unreliable for the most challenging scenarios (such as car egress motions [9]), mostly because of the difficulty of approximating collision avoidance constraints. Still, most “2.5D” environments (which can be accurately represented with a heightmap) composed of quasi-flat contact surfaces (i.e., a friction cone containing a gravity vector) can be handled with such constraints [9]. Unsurprisingly, these environments, including rubles, stepping stones and staircases, are the application targets for most contributions in the literature.

2) Online contact planning with reduced dynamics

The combination of reduced dynamic models and simplified collision constraints makes optimisation-based techniques tractable. Optimal control is attractive as it allows us to find solutions robust to uncertainties through the minimisation of selected criteria. Sampling-based approaches, instead, only look at feasibility. To model the combinatorics, the first approach is to relax the problem by modelling the discrete (boolean) variables that represent the contact decisions with

continuous variables, resulting in a formulation that can be readily solved by off-the-shelf nonlinear programming (NLP) solvers [14], [15]. However, there is no guarantee that the system’s dynamic constraints will be satisfied even with reduced models, with contacts potentially planned where no surfaces exist [16]. Alternatively, Linear Complementary Constraints (LCP) can be used to accurately model contact constraints, but they are notoriously difficult to handle by NLP solvers [17]. In both cases learning initial guesses can help the solver converge to a feasible solution [18], [19]. The second approach is to explicitly handle combinatorics using Mixed-Integer Programming (MIP). MIP solvers tackle contact planning problems for bipeds [20] and quadrupeds [21], provided that the underlying optimisation problem is convex, which results in a conservative approximation of the dynamics [22]. Monte Carlo Tree Search (MCTS) has recently been proposed as a promising alternative to MIP that could provide a relevant trade-off between exploration and exploitation [23]. In this work, we choose MIP for planning contacts as it has experimentally led to the most effective results in terms of computation time and reliability [24].

3) Perceptive locomotion with instantaneous decisions

The validity of a method is influenced by whether the environment is fully known or not. Reactive perceptive pipelines exist on the LittleDog robot [25]–[27], though they require high-precision terrain pre-mapping and an external motion capture system. When the environment is not fully known, it is typically modelled as an elevation map by fusing depth sensor information within proprioceptive information [28], [29], which may require specific developments to handle drifts in the state estimation [28]. Recent approaches propose to directly optimise the next contact position, the torso orientation and obstacle avoidance for the foot trajectory based on this input [30]–[32]. As discussed in Section IX, we believe that combining these approaches is a promising research avenue.

4) Whole-body predictive control relying on CoM motions

Because of the nonlinearity induced by any changes to the contact plan, the WB-MPC rarely challenges the step locations, even though the approximations do not guarantee feasibility. The uncertainties resulting from state estimation and environment perception motivate frequent re-computation of the contact plan, which is not possible as their frequency is usually low (about 5 Hz in [16]). Furthermore, the WB-MPC is usually additionally constrained to track a reference trajectory for the Centre Of Mass (COM) or the base [33], [34] to facilitate convergence, but we argue that this tracking is problematic when perturbations such as contact slipping occur. Our conclusion is that the use of reduced models for contact planning necessarily leads to errors in the WB-MPC that result in slipping contacts. In the current state of the art, reduced models appear necessary for satisfying real-time constraints. Mitigating this issue involves allowing to adapt

a contact plan at a higher frequency. However, it involves writing a WB-MPC that robustly accommodates these errors and gives as much freedom as possible when following a contact plan.

B. CONTRIBUTION

This paper extends our published conference paper [24], where we propose a *contact repositioning* module between the contact surface planner and the MPC to adapt to the robot's estimated state and perceived environment uncertainties. To achieve this, we decouple the contact surface selection from the control, but not the computation of the contact position on that surface, which we instead update synchronously with the MPC output and updated state of the robot, at 50 Hz. As such, the contact repositioning module is our primary contribution. This work is thus to be considered as an experimental and theoretical extension in the following manner.

We propose a novel, complete perceptive locomotion architecture comprised of the following features: terrain segmentation, real-time surface selection and footstep planning, free-collision foot-swing planning, and whole-body MPC which considers torque limits and generates local controllers. It relies on four technical contributions:

- (i) a contribution to the decoupling between contact planning and WB-MPC,
- (ii) an empirical demonstration of the added value of WB-MPC in perceptive locomotion,
- (iii) a perception pipeline based on a convex segmentation approach using onboard terrain elevation maps, and
- (iv) exhaustive hardware trials on challenging terrains demonstrating increased capabilities for the ANYmal B robot.

Finally, we extend the plane segmentation algorithm [35], which decomposes potential contact surfaces into a sequence of convex surfaces needed by our contact planner. We use the Visvalingam–Whyatt and Tess2 algorithms to correctly handle overlapping surfaces and conservatively reduce the complexity of the scene, leading to a more versatile and robust environment construction. Our framework results in state-of-the-art locomotion capabilities under a wide range of conditions, robust to strong perturbations including sliding contacts and missed steps during stair climbing.

The reader should note that we have incorporated a feedback WB-MPC [36], [37] into our framework, as opposed to the more conventional approach in our previous work, which involved MPC with reduced dynamics followed by a WBC running at a higher frequency. Constraining only the end-effector trajectories, instead of the CoM motions, allows the WB-MPC to freely accommodate substantial perturbations and perception errors, and to maximise the robot's capabilities by optimising its posture. While a subset of our experimental results are shared by both papers, their contributions are orthogonal.

II. ARCHITECTURE OVERVIEW

An overview of the locomotion pipeline is presented in Fig. 2. Walkable surfaces are described via convex planes extracted from the terrain elevation map at 1 Hz. Given the current state of the robot (position / orientation of the base, position of active contacts) and a desired velocity (joystick input), a mixed-integer program is used to select the convex surfaces on which the next 6 or 8 steps will occur, depending on the gait used. A new surface plan is computed at the beginning of each new phase, which corresponds to approximately 3-5 Hz in our experiments. Given the next contact surfaces, the trajectory of each end-effector is updated at 50 Hz, before each iteration of the whole-body MPC. The control policy is then sent to the robot with a Riccati gain controller. State estimation is performed onboard at 400 Hz by fusing inertial sensors from IMU and odometry provided by the ANYbotics software [38]. Finally, a LIDAR is used to correct the drift of these measurements by analysing fixed points in the environment.

III. DEFINITIONS AND NOTATIONS

In line with the notations used in our previous work [24], the *robot state* is formally described by the:

- Centre Of Mass (COM) position, velocity and acceleration $\mathbf{c}$, $\dot{\mathbf{c}}$ and $\ddot{\mathbf{c}}$, each in $\mathbb{R}^3$;
- base transformation matrix in the world frame;
- 3D position of each end-effector in the world frame;
- gait, i.e., the list of effectors currently in contact, as well as the contacts to be activated and deactivated over the planning horizon.

The *horizon* n is defined as the number of future contact creations considered. In the case of the trotting gait, a horizon $n = 6$ describes three steps, as at each step two contacts are created simultaneously.

At the *Surface Selection* planning stage, motion is decomposed into *contact phases*. Each contact phase is associated with a number of feet in contact and one or more contacts are broken/created at each phase. In the case of a trotting gait with a horizon $n = 6$, it corresponds to three contact phases since two feet move at the same time. For a walking gait, the horizon $n = 8$ contains 8 contact phases since only one foot moves at a time.

The *environment* is the union of $m + 1$ disjoint quasi-flat¹ contact surfaces $\mathcal{S} = \bigcup_{i=0}^m \mathcal{S}^i$. Each set $\mathcal{S}^i$ is a polygon embedded in a 3D plane, i.e.,

$$\mathcal{S}^i := \{\mathbf{p} \in \mathbb{R}^3 | \mathbf{S}^i \mathbf{p} \leq \mathbf{s}^i\}, \quad (1)$$

where $\mathbf{S}^i \in \mathbb{R}^{h \times 3}$ and $\mathbf{s}^i \in \mathbb{R}^h$ are respectively a constant matrix and a vector defining the h half-space bounding surface.

The *contact plan* is described as a list of contact surfaces $\mathcal{S}_k^j \in \mathcal{S}$, $1 \leq j \leq l$ with l being the total number of end-effectors and k being the k -th contact phase.

¹ie such that its associated friction cone contains the gravity vector.




FIGURE 2. Overview of our perceptive locomotion pipeline. Around 1Hz, the perceptive elements *Elevation map* and *Convex Plane Segmentation* extract convex planes from the surrounding environment (Section IV). The *Surface Selection* block, running around 5Hz (depending on the gait chosen), chooses between them the next surfaces of contact (Section V). At the frequency of the MPC, 50 Hz, *Footstep and Collision free-trajectory* elements generate the curve for each moving foot (Section VI). Finally, the *Whole-Body MPC* and *Riccati gain controller* synthesise the motion (Section VII).

IV. PERCEPTION

This section reviews the *Elevation map* and *Convex Plane Segmentation* components of the architecture (Fig. 2).

A. SENSORS

Regarding the proprioceptive sensors, the state is estimated by fusing leg odometry and the Xsens MTi-100 IMU [38]. Additionally, a rotating Hokuyo UTM-30LX-EW lidar sensor is placed at the back of the robot to correct the drift of the state estimation with an iterative closest point (ICP) method [39]. A single depth camera Intel RealSense D435, mounted at the front of the robot, extracts height information from the surrounding area into a point cloud. High-accuracy presets are used on the camera to prioritize the accuracy of the point cloud over speed. This increases the quality of the elevation map and no filters are needed to post-process the point clouds. However, this preset degrades the frequency in our setup to a range between 2 and 5 Hz for point cloud collection.

B. CONTACT SURFACES EXTRACTION

The probabilistic and local mapping method in [28] converts point cloud data into an elevation map locally around the robot's pose. Including proprioceptive localisation from kinematic and inertial measurements produces an estimate of the surroundings as a 2.5D heightmap. Potential surfaces are then extracted with the Plane-Seg algorithm [35] by clustering planar points with similar normals. The planes are extracted without memory and therefore the surfaces can change suddenly with each height map update. The quality and consistency of the surfaces then depend entirely on the quality of the height map, as developed in Sec. IX.

C. REFINEMENT AND MARGIN OF SAFETY

Post-processing of Plane-Seg planes is essential to ensure safe footstep decisions. We have added it for three reasons. First, the complexity of the contact surfaces (i.e., the number

of points) has a significant impact on the computation time of the surface selection algorithm. We propose conservatively approximate surfaces with more than 8 vertices with an 8-vertex polygon. Filtering is also done to remove unreachable surfaces. For example, the plane extraction method could return overlapping surfaces when considering a staircase. Here, the ground surface is often detected below the steps and planning a footstep inside it will obviously result in failures. Finally, a safety margin is applied, around 4 cm on each surface, to avoid putting feet on the edges of the surfaces. This can be harmful in the event of estimation errors. This margin is also useful to avoid knee collisions of the knees with the environment, which are not explicitly accounted for.

1) Vertices number reduction

The number of vertices is first reduced using the conservative Visvalingam–Whyatt line simplification algorithm [40]. It eliminates progressively the points from a line that forms the smallest area with its closest neighbours as described in pseudo-code 1. No hyperparameters other than the final number of points (which is 8) are needed. Fig. 3a shows an example of this reduction.

Algorithm 1 Pseudo-code for Visvalingam–Whyatt line simplification

```

1: List of  $n$  vertices :  $L = [P_0 \dots P_n]$ 
2: while  $\text{len}(L) > n_{\max}$  do
3:   for  $i \in [0, \text{len}(L) - 2]$  do
4:     Compute area of  $[P_i, P_{i+1}, P_{i+2}]$ 
5:   end for
6:   Remove  $P_{i+1}$  corresponding to the smallest area.
7: end while
8: return  $L$ 
  
```

2) Safety margins

On this updated contour, an inner and an outer margin are computed parallel to each edge as shown in Fig. 3b. The

inner margin allows the robot to step into a safe area. The outer margin artificially increases the size of the obstacle to prevent the end-effector from getting too close of the obstacle. This is done to avoid collisions while computing swing-foot trajectories. These two margins are used to avoid stepping on the corner of an obstacle due to state estimation errors and for collision avoidance with the body.




FIGURE 3. Reduction to an 8-vertex polygon using the Visvalingam-Whyatt algorithm on an initial 20-vertex polygon on the left **3a**. Applying inner and outer margins to an 8-vertex polygon **3b**.

3) Convex decomposition

Starting from the lowest surface, the overlapping surfaces are removed and a convex decomposition is performed on the resulting contour using the Tess2 algorithm [41] as shown in Fig. 4a. Smaller areas under 0.03 m^2 are deleted. Algorithm IV-C3 gives the pseudo-code of this post-processing. In addition, a rectangle is added below the robot's position, at the estimated height of the feet. This is to ensure that there is always a surface under the robot's feet if the elevation map has not been built. This surface is treated differently in the process to avoid overlapping it with real obstacles but has been removed from the pseudo-code for clarity.




FIGURE 4. Example of surface processing and convex decomposition. These two figures represent the same 3D scene with 2 surfaces overlapping the ground surface. On the right Fig. 4b, the scene is viewed from a top perspective.

Algorithm 2 Pseudo-code for surface processing

```

1: for all surfaces do
2:   Reduce nb of points with Visvalingam's algorithm.
3:   Compute inner and outer contour.
4: end for
5:  $L_i \leftarrow$  List of surfaces in ascending order of height.
6:  $L_f \leftarrow$  Empty list.
7: while  $L_i$  is not empty do
8:   Get first surface  $S_f$  and remove it from  $L_i$ .
9:    $L_o \leftarrow$  Empty list.
10:  for all surface in  $L_i$  do
11:    if surface.outer intersect with  $S_f$ .inner then
12:      Increment  $L_o$  with outer contour of surface.
13:    end if
14:  end for
15:  if  $L_o$  is not empty then
16:    Convex decomposition between  $S_f$ .inner and  $L_o$ .
17:    Remove small areas.
18:    Add remaining surfaces in  $L_f$ .
19:  else
20:    Add  $S_f$ .inner in  $L_f$ 
21:  end if
22: end while
23: return  $L_f$ 

```

V. SURFACES SELECTION

Given the current state of the robot (active contacts and position, COM location), the environment given as a union of non-intersecting surfaces, as well as a desired target velocity for the robot and a desired gait, our *Surface Selection* (Fig. 2 - yellow) algorithm computes a feasible contact plan, composed of n contact surfaces that the robot should step on for the planning horizon (Sec. III). We set $n = 8$ for a walking gait and $n = 6$ for a trotting gait in our experiments. The surface selection algorithm is executed between 3 and 5 Hz. This means that the contact plan is updated before each new step is made by the robot. The frequency depends on the gait and each optimisation starts at the beginning of each step.

The algorithm is implemented as a Mixed-Integer Program (MIP) [42]. We use the SLIM formulation of this algorithm [16], with adaptations to better match the desired robot behaviour. These adaptations are in [24] and are described here for completeness. In this previous work, the number of contacts optimised was set to 4 for the Solo robot [43]. Here the ANYmal's robot dynamics are slower and the timing between each new contact is longer². This gives more time for computing the contact plan, allowing us to increase the planning horizon. We refer the reader to [24] for empirical justifications for these choices.

We first describe how the potential contact surfaces are pre-filtered to improve the algorithm's computational performance without loss of generality. We then provide the mathematical formulation of the *Surface Selection* algorithm. We conclude this section with the details of the cost function used in this optimisation problem.

²The duration of a step was set to 160 ms on Solo whereas it is set to 600 ms for a walking gait and 300 ms for a trotting gait on ANYmal.




FIGURE 5. (a) CoM extrapolation along the horizon. (b) ROMs of the 4 effectors for the current state. (c) ROMs along the horizon for the front right foot.

A. PRE-SELECTION

We first pre-filter the number of contact surfaces using the robot's range of motion (ROM) to reduce the combinatorics. To do so, a CoM trajectory is extrapolated from the joystick velocity command (Fig. 5a) over the contact phases. The position of each state and the yaw angle (orientation around the z-axis) are integrated from the linear and angular (yaw only) desired velocity. The roll (orientation around the x-axis) and pitch (orientation around the y-axis) angles of the guide are computed as the average slope of the terrain around the robot position, given by solving (Fig. 5a)

$$\min_n \|\mathbf{A}\mathbf{n} - \mathbf{B}\|^2$$

with $\forall i, j \in N_x \times N_y, \mathbf{A}[i + jN_x, :] = [x_i, y_j, 1],$
 $\mathbf{B}[i + jN_x] = \text{elevation}(x_i, y_j),$

where x_i and y_j are respectively the positions on the x and y axis around the robot's position with a user-defined resolution of $(N_x, N_y) \in \mathbb{N}^{+2}$. The elevation function gives the terrain height at the position (x_i, y_j) and can be obtained directly from the heightmap or by evaluating the convex plane corresponding to this 2D position. $\mathbf{n} = [a, b, c] \in \mathbb{R}^3$ is the vector optimised to form the plane of equation $ax + by - z + c = 0$. The experiments use 10×10 as resolution. For each extrapolated state c_j^* (8 states for a walking gait as $n = 8$ and 3 for a trotting gait as $n = 6$ but two feet contacts are created at each phase), the 6D configuration is given by:

$$c_j^* = \begin{bmatrix} x_0 + \int_0^{t_j} (v_x^* \cos(\psi^* t) + v_y^* \sin(\psi^* t)) dt \\ y_0 - \int_0^{t_j} (v_x^* \sin(\psi^* t) - v_y^* \cos(\psi^* t)) dt \\ ax_j + by_j + c + h_{\text{ref}} \\ \arctan(b) \\ -\arctan(a) \\ \psi_0 + \psi^* t_j \end{bmatrix}, \quad (3)$$

where the state is described using a 3D position $[x_j, y_j, z_j]$ and 3 Euler angles. The current extrapolated positions x_j and y_j are used to compute the configuration height. v_x^*, v_y^* are the x-y linear velocities and $\dot{\psi}^*$ the yaw angular velocity coming from the joystick input. $h_{\text{ref}} = 0.48$ is the robot's nominal height. t_j is the step duration, manually defined for each gait.




FIGURE 6. Environment $\mathcal{S}$ with 3 contact surfaces for a walking gait (1 foot moving at a time). Circles: initial position; squares: next steps locations.

For each c_j^* , the ROM of each moving leg is intersected with the surfaces. The ROM is approximated by a convex set and represented by the green surfaces in Fig. 5b. Only the surfaces that intersect this ROM are selected. The distance between two convex sets is computed efficiently with the GJK algorithm [44] from the Pinocchio and HPP-FCL libraries [45]. In our experiments, this usually reduces the number of potential surfaces from 20 to 3 for each moving foot on average. This significantly reduces combinatorics (from about 20^n to 3^n possible combinations).

B. SURFACE SELECTION ALGORITHM AS A MIP

The surface selection module computes a contact plan for the robot that satisfies linearized kinematic constraints [12], [46]. In the following, we recall the mathematical formulation of the problem as a mixed-integer program (MIP). This MIP outputs the contact surfaces selected. It also computes the 3D locations of footsteps as a by-product. However, they are discarded as the target footsteps will be adapted at a higher frequency by the footstep planner.

1) Contact constraint representation

The perception pipeline provides a set of $m + 1 \in \mathbb{N}^+$ disjoint quasi-flat surfaces, which define the environment

$\mathcal{S}$, and the motion is decomposed into n contact phases (Sec. III), as shown in Fig. 6. To constrain a point $\mathbf{p} \in \mathbb{R}^3$ to be in contact, we must write:

$$\exists i, \mathbf{S}^i \mathbf{p} \leq \mathbf{s}^i \Leftrightarrow \mathbf{S}^0 \mathbf{p} \leq \mathbf{s}^0 \vee \dots \vee \mathbf{S}^m \mathbf{p} \leq \mathbf{s}^m. \quad (4)$$

The *or* constraint is classically expressed as an integer constraint using the Big-M formulation [47] as follows. We introduce a vector of binary variables $\mathbf{a} = [a_0, \dots, a_m] \in \{0, 1\}^{m+1}$ and a sufficiently large constant $M \in \mathbb{R}^+, M \gg 0$. (4) is equivalently rewritten as:

$$\forall i, \mathbf{S}^i \mathbf{p} \leq \mathbf{s}^i + M(1 - a_i); \sum_{i=0}^m a_i = 1. \quad (5)$$

Under this formulation if $a_i = 1$, then $\mathbf{p}$ belongs to $\mathcal{S}^i$ (and is thus in contact). Instead, if $a_i = 0$ for a sufficiently large M then the constraint $\mathbf{S}^i \mathbf{p} \leq \mathbf{s}^i + M(1 - a_i)$ will be satisfied for any value of $\mathbf{p}$, in other words, the constraint is inactive. $\sum_{i=0}^m a_i = 1$ implies that $\exists i, a_i = 1$. The obtained behaviour is thus the desired one: if (5) is true then $\mathbf{p}$ is in contact with a surface.

2) Problem formulation

The final MIP problem is obtained by combining contact and surface constraints as follows. For simplicity, and without loss of generality, we assume that the candidate contact surfaces are the same for each step.

$$\begin{aligned} \text{find } & \mathbf{P}, \mathbf{A} = [\mathbf{a}_0, \dots, \mathbf{a}_m] \in \{0, 1\}^{(m+1) \times n} \\ \text{min } & l(\mathbf{P}) \\ \text{s.t. } & \mathbf{K}\mathbf{P} \leq \mathbf{k}, \\ & \mathbf{P} \in \mathcal{C}, \\ & \forall j \in \{0, \dots, n-1\}: \\ & \forall i, \mathbf{S}^i \mathbf{p}_j \leq \mathbf{s}^i + M(1 - a_i^j); \sum_{i=0}^m a_i^j = 1, \end{aligned} \quad (6)$$

where $\mathbf{P} = [\mathbf{p}_0 \dots \mathbf{p}_n] \in \mathbb{R}^{3 \times n}$ is the vector comprising the variable n next foot positions (6 or 8 in our experiments); $\mathbf{a}^j = [a_0^j, \dots, a_m^j]$ is the vector of binary variables associated with the j -th optimised foot position; $l(\mathbf{P})$ is an objective function; $\mathcal{C}$ is a set of user-defined convex constraints (in our case, constraints on initial contact positions); $\mathbf{K}$ and $\mathbf{k}$ are constant matrix and vector representing the linearised kinematic constraints on the position of each effector with respect to the others [12], [15], [46].

3) Cost function details

As mentioned in [24], two quadratic costs are used in the problem. The first attempts at regularising the footstep locations using Raibert's heuristic. Since the optimisation is triggered at the beginning of each step (between 3 and 5 Hz), and not at each time step (50 Hz), an approximation of the Raibert heuristic is applied. The idea is to interpolate the base position given the desired velocity and place the foot accordingly to the estimated hip position. The second term penalises the distance between the hip and the foot

location. This aims at penalising solutions close to reaching the robot's kinematic limits.

$$l(\mathbf{P}) = \sum_j^n w_1 \|\mathbf{p}_j - \mathbf{p}_j^*\|_{x,y}^2 + w_2 \|\mathbf{p}_j - \mathbf{p}_{\text{hip},j}^*\|^2 \quad (7)$$

where $\mathbf{p}_j^*$ is the extrapolated foot position taking into account the linear and angular reference velocity at the corresponding contact phase. The weights w_1 and w_2 are set to 1 and 0.5. This penalisation only considers the ℓ^2 norm on the x- and y-axis. Similarly, $\mathbf{p}_{\text{hip},j}^*$ is the extrapolated hip position from the desired velocity. The computation of the reference foot location is detailed in Sec. VI.

VI. FOOT TRAJECTORY GENERATION

Our *Footstep and Collision free-trajectory* (Fig. 2 - red) algorithms compute the end-effector trajectory as a Bezier curve given the current robot state (COM position/orientation), the next moving foot within the horizon of the MPC, the contact plan obtained by the Surface Planner, the timings of contact depending on the gait chosen, as well as a desired target velocity for the robot. At the MPC frequency, set to 50 Hz, a first quadratic program (QP) computes the next foot location and a second QP calculates its trajectory during the flight. Their formulation differs from our previous work [24] as the base velocity is not optimised here. The total computation time for the foot trajectory is negligible in the pipeline, as evidenced by Table 2.

A. FOOT POSITION OPTIMISATION

We use the Raibert's heuristic [48] as it is commonly done on quadruped robots [49]–[51], which provide coherent foot locations according to the base position, gait period, current and reference velocity on flat ground. The objective is to extend it in 3D and satisfy the surface selection constraint. The target 2D position of the foot $\mathbf{p}^*$ is computed as follows:

$$\mathbf{p}^* = \mathbf{p}_{\text{hip}} + \frac{T_s}{2} \mathbf{v}_{\text{ref}} + k(\mathbf{v}_{\text{base}} - \mathbf{v}_{\text{ref}}) + \sqrt{\frac{h}{g}} \mathbf{v}_{\text{base}} \times \omega_{\text{ref}} \quad (8)$$

where T_s is the stance phase time extracted from the phase-based gait pattern, $\mathbf{v}_{\text{base}}$, and $\mathbf{v}_{\text{ref}}$ are respectively the current base velocity and the reference velocity commanded by the user. Finally, h , g , k and ω_{ref} are respectively the nominal height of the robot, the gravity constant, a user-defined feedback coefficient, and the reference angular velocity around the z-axis. The estimated hip position at contact $\mathbf{p}_{\text{hip}}$ accounts for the accumulated delay in the estimation of the base position, since the moving average is used to reject oscillations of the same period as gait, as discussed in section VIII-G. The final position $\mathbf{p}$ is the closest to $\mathbf{p}^*$ that lies on the selected surface $\mathcal{S}$ under the locally-

approximated kinematic constraints $\mathcal{K}$:

$$\min_{\mathbf{p}} \quad \|\mathbf{p} - \mathbf{p}^*\|^2 \quad (9a)$$

$$\text{s.t.} \quad S\mathbf{p} \leq s, \quad (9b)$$

$$\mathbf{p} \in \mathcal{K}. \quad (9c)$$

Unlike our previous work [24], the reference base velocity is not optimized in this QP, as it is unconstrained in the WB-MPC. Additionally, in the specific case of a walking gait, we have removed the feedback term $k(\mathbf{v}_{base} - \mathbf{v}_{ref})$ for two reasons. First, the slower dynamics of walking gait reduce the effectiveness of this term. Second, the inherent oscillatory motion introduces an additional delay of approximately half a period due to the filtering applied to reject these oscillations (as discussed in Sec. VIII-G).

B. COLLISION FREE TRAJECTORY




FIGURE 7. Adaptation of the end-effector trajectory to climb a step. The active constraint corresponds to the half-plane crossed by the reference trajectory -red curve-. Both curves are discretised to perform the optimisation and the orange points highlight the ones under constraints.

A collision-free trajectory $\mathbf{p}(t) : \mathbb{R} \rightarrow \mathbb{R}^3$ connects the current foot location to the optimised location. $\mathbf{p}(t)$ aims at following a reference trajectory while avoiding obstacles and keeping computational time low as it is parameterized using a Bezier curve of degree d on the Bernstein basis $(B_i)_{i \leq d}$:

$$\mathbf{p}(t) = \sum_{i=0}^d B_i^d\left(\frac{t}{T}\right) \mathbf{P}_i, \quad (10)$$

where $\mathbf{P} = [\mathbf{P}_0 \ \dots \ \mathbf{P}_d]^T$ are the $d + 1$ control points and T is the total time of the trajectory, empirically set to 600 ms for a walking motion.

1) Reference trajectory

The trajectory $\mathbf{p}_{ref}(t) : \mathbb{R} \rightarrow \mathbb{R}^3$ is composed of a degree 6 polynomial curve of degree 6 on the z -axis and a degree 5 one on the x and y axes. All of which have with the first 3 control points fixed to ensure the continuity in position, velocity and acceleration of the curve from the current state.

End velocity and acceleration are set to 0 to avoid slippage at the end position. Additionally, the height at $\frac{T}{2}$ is fixed on the z -axis, constraining all degrees of freedom.

2) Collision avoidance

The trajectory $\mathbf{p}(t_k)$ is a 3D curve of degree $d = 7$ that follows $\mathbf{p}_{ref}(t)$ while avoiding collisions. To do so, we enforce collision avoidance along the foot trajectory. This is achievable iteratively [52] by solving a sequence of QPs and adding constraints where collisions occur. For computational efficiency, we empirically identify the points likely to be in a collision and add collision constraints to them (Fig. 7 - Yellow dots). These $n_c + 1$ points $\mathbf{p}(t_k) = \mathbf{A}_k \mathbf{P}$, $\forall k \in [0, \dots, n_c]$ are linearly defined by Bezier's control points, with $\mathbf{A}_k \in \mathbb{R}^{3 \times (d+1)}$. We choose the active constraint to be the half-space traversed by the reference curve (Fig. 7 - pink half-space). We thus write $\forall k, \mathbf{S}_m^i \mathbf{p}(t_k) \geq s_m^i$, where $\mathbf{S}_m^i \in \mathbb{R}^3$ and $s_m^i \in \mathbb{R}$ define what constitutes the current half-space. All the collected constraints are then stacked into a single matrix and vector $\mathbf{G}$ and $\mathbf{h}$, leading to the QP:

$$\min_{\mathbf{p}} \quad \sum_{k=0}^{n_c} \|\mathbf{p}(t_k) - \mathbf{p}_{ref}(t_k)\|^2 \quad (11a)$$

$$\text{s.t.} \quad \mathbf{p}(0) = \mathbf{p}_{ref}(0), \quad \mathbf{p}(T) = \mathbf{p}_{ref}(T), \quad (11b)$$

$$\dot{\mathbf{p}}(0) = \dot{\mathbf{p}}_{ref}(0), \quad \dot{\mathbf{p}}(T) = \mathbf{0}_{\mathbb{R}^3}, \quad (11c)$$

$$\ddot{\mathbf{p}}(0) = \ddot{\mathbf{p}}_{ref}(0), \quad \ddot{\mathbf{p}}(T) = \mathbf{0}_{\mathbb{R}^3}, \quad (11d)$$

$$\mathbf{G}\mathbf{P} \leq \mathbf{h}. \quad (11e)$$

This approach can be expanded to avoid local obstacles along the trajectory by imposing constraints that depend on the environment's heightmap. Currently, our trajectory avoidance is based only on obstacles perceived as walkable surfaces, which has been found to work effectively in many scenarios. Nevertheless, if an obstacle is along the trajectory but not detected as such, it will not be avoided. This highlights the importance of including height-map information when available.

VII. MOTION GENERATION

Given the current state of the robot (6D position and velocity), the next contact sequence and the end-effector trajectories, our *Whole-Body MPC* (WB-MPC) and *Riccati gain controller* elements (Fig. 2 - red and violet) generates the motion of the legged robot. The WB-MPC optimises a trajectory at 50 Hz and the Riccati gain controller applies it at 400 Hz.

A. OPTIMAL CONTROL FORMULATION

The legged robot generates motions through a model predictive controller that relies on the robot's full-body dynamics. The formulation is based on two previous papers [36], [37]. We solve the optimal control (OC) problems using Crocoddyl's advanced solvers [34]. Our pipeline relies on two different OC formulations based on forward and inverse dynamics. Both yield similar performances, demonstrating our method's generality. Both formulations can be written

as [36], [37]:

$$\begin{aligned}
 \min_{\{\mathbf{q}, \mathbf{v}\}, \{\boldsymbol{\tau}\} \text{ or } \{\dot{\mathbf{v}}, \boldsymbol{\lambda}\}} \quad & \ell_N(\mathbf{x}_N) + \sum_{k=0}^{N-1} \int_{t_k}^{t_{k+1}} \ell_k(\mathbf{q}_k, \mathbf{v}_k, \boldsymbol{\lambda}_k, \boldsymbol{\tau}_k) dt \\
 \text{s.t.} \quad & \mathbf{q}_{k+1} = \mathbf{q}_k \oplus \int_{t_k}^{t_{k+1}} \mathbf{v}_{k+1} dt, \\
 & \mathbf{v}_{k+1} = \mathbf{v}_k + \int_{t_k}^{t_{k+1}} \dot{\mathbf{v}}_k dt, \\
 & \begin{bmatrix} \dot{\mathbf{v}}_k \\ -\boldsymbol{\lambda}_k \end{bmatrix} = \begin{bmatrix} \mathbf{M}_k & \mathbf{J}_c^\top \\ \mathbf{J}_c & \end{bmatrix}^{-1} \begin{bmatrix} \boldsymbol{\tau}_b \\ -\mathbf{a}_c \end{bmatrix} \quad (\text{forward dyn.}) \\
 \text{or} \quad & \text{ID}(\mathbf{q}_k, \mathbf{v}_k, \dot{\mathbf{v}}_k, \boldsymbol{\lambda}_k) = \mathbf{0}, \quad (\text{inverse dyn.})
 \end{aligned} \tag{12a}$$

where the robot state $\mathbf{x} = (\mathbf{q}, \mathbf{v})$ contains the generalized position and velocity vectors. More precisely, $\mathbf{q} \in \mathbb{SE}(3) \times \mathbb{R}^{n_j}$ with the generalized velocity lying in the tangent space $\mathbf{v} \in \mathfrak{se}(3) \times \mathbb{R}^{n_j}$ where n_j is the number of articulated joints and $\oplus$ denotes the integration operator in $\mathbb{SE}(3)$ inspired by [53] and used in [34].

For the forward-dynamics formulation, the input command is $\mathbf{u} = \{\boldsymbol{\tau}\}$ corresponds to the joint torques. k , $\mathbf{M}_k$ and $\mathbf{J}_c$ corresponds to the discrete-time t_k , the generalised mass matrix and the contact Jacobian, respectively. $\boldsymbol{\tau}_b$ includes the joint torque command, Coriolis and gravitation terms and $\mathbf{a}_c$ is the desired acceleration which results from the rigid contact constraint (contact point velocity is null) and includes the Baumgarte stabilisation term [54]. The impulse dynamics described in more detail in [36] which allow velocity changes at impact have also been omitted from the formulation. This contact dynamics formulation comes from the application of the Gauss principle of least constraint [46], as described in detail in [55]. By handling this constraint in the backward pass, the decision variables can be condensed with forward dynamics [34], [36], [55] and contact forces can be removed. This drastically reduces the number of decision variables.

Regarding the inverse-dynamics formulation, where $\mathbf{u} = \{\dot{\mathbf{v}}, \boldsymbol{\lambda}\}$, the contact forces and the acceleration become decision variables of the problem [37], [56]. We can perform efficient factorizations by applying nullspace parametrization as proposed in [37]. Furthermore, an alternative inverse dynamic equation that allows us to remove joint torques from the control vector is possible. Both strategies reduce the size of the problem needed to enable MPC applications.

Regularisation terms included in the cost function ℓ are summarized in Table 1. These different formulations are transcribed into quadratic terms, where $\mathbf{q}_a, \dot{\mathbf{q}}_a \in \mathbb{R}^{12}$ are respectively the actuated joint positions and velocities. $\mathbf{p}_{\text{base}} \in \mathbb{SE}(3)$ describes the base pose in the world frame. $\mathbf{p}_{\mathcal{G}_k} \ominus \mathbf{p}_{\text{ref}, \mathcal{G}_k} = \mathbf{p}_{\mathcal{G}_k}^{-1} \mathbf{p}_{\text{ref}, \mathcal{G}_k}$ represents the pose error in $\mathbb{SE}(3)$ of the feet relative to the reference position. It is

TABLE 1. Details of the costs used in the OCP

Name	Formulation	Cost
State bound	$\mathbf{x}_{\min} \leq \mathbf{x} \leq \mathbf{x}_{\max}$	10^3
Base orientation reg.	$\log(\mathbf{p}_{\text{base}}^{-1})$	$[0_{\mathbb{R}^3} \ 10^2 \ 10^2 \ 0]$
Base velocity reg.	$\ \dot{\mathbf{p}}_{\text{base}}\ $	10
Joint position reg.	$\ \mathbf{q}_a - \mathbf{q}_{a, \text{ref}}\ $	0.01
Joint velocity reg.	$\ \dot{\mathbf{q}}_a\ $	1
Torque reg.	$\ \mathbf{u}_k\ $	0.1
Force reg.	$\ \boldsymbol{\lambda}_{\mathcal{C}_k}\ $	1
Friction cone	$\mathbf{C} \boldsymbol{\lambda}_{\mathcal{C}_k} \leq \mathbf{c}$	10
Feet position track.	$\log(\mathbf{p}_{\mathcal{G}_k}^{-1} \mathbf{p}_{\text{ref}, \mathcal{G}_k})$	10^6
Feet velocity track.	$\dot{\mathbf{p}}_{\text{ref}, \mathcal{G}_k} - \dot{\mathbf{p}}_{\mathcal{G}_k}$	10^4

penalised by considering it in its tangent space [53] with the log function that maps $\mathbb{SE}(3)$ to $\mathfrak{se}(3)$. Feet position and velocity tracking error are expressed in the robot's inertial frame. The base position is not penalised whereas its rotation is penalised around the x- and y-axis only, corresponding to the roll and pitch angles. $\mathbf{C}$ and $\mathbf{c}$ are the matrix and vector of the linearised friction cone. $\mathcal{C}_k$ represents the set of feet in contact with the ground and inversely $\mathcal{G}_k$ is the set of feet in the swing phase for the node k . Costs are set to zero if the corresponding feet are in the contact phase.

B. RICCATI-GAIN CONTROLLER

OC formulations with reduced-order dynamics require an instantaneous whole-body controller for tracking computed forces and maintaining robot balance [49], [51], [57]. Such instantaneous controllers may compete with the MPC policy and not necessarily generate the motion predicted by them [58]. At a higher cost in computing time, which has a non-negligible influence, many benefits can appear when including whole-body dynamics in the OC problem, e.g., imposing the joint effort limits. Another benefit is to derive local feedback controllers from optimisation principles. This control policy looks like this:

$$\boldsymbol{\tau}_d = \boldsymbol{\tau}_{ff} + \mathbf{K}(\mathbf{x} \ominus \mathbf{x}^*), \tag{13a}$$

$$= -\mathbf{Q}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{Q}_{\mathbf{u}} - \mathbf{Q}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{Q}_{\mathbf{u}\mathbf{x}}(\mathbf{x} \ominus \mathbf{x}^*), \tag{13b}$$

where $\ominus$ denotes the $\mathbb{SE}(3)$ error. $\mathbf{Q}_{\mathbf{u}\mathbf{u}}$, $\mathbf{Q}_{\mathbf{u}}$ and $\mathbf{Q}_{\mathbf{u}\mathbf{x}}$ are respectively the partial derivatives of the value function for the joint-effort command and the state [36]. While the MPC runs every 0.02 s (50 Hz), the discretisation is set to 0.01 s. Instead, the low-level Riccati-gain controller runs at 400 Hz, and an interpolation step is necessary. For this, the feedback term is computed by using an interpolation of the reference state $\mathbf{x}^*$ by integrated the dynamic with the contact model dynamic 12. The optimal effort command is then provided at 400 Hz in addition to a joint impedance controller based on the joint command and velocity obtained from $\mathbf{x}^*$.

VIII. RESULTS

A. IMPLEMENTATION

Communication between each module is achieved via a low-latency ROS communication layer (TCP, no-delay) [59]. Three onboard computers (Intel(R) Core(TM) i7-5600U CPU @ 2.6GHz) share the main tasks of locomotion, perception and estimation. Point-cloud framework, heightmap generator, state estimation and Riccati-gain controller are performed onboard. We used an additional computer (Intel(R) Core(TM) i5-8365U CPU @ 1.60GHz) to extract the segmented planes from the heightmap, post-process the extracted surfaces and run the mixed-integer program in a different thread. A final computer (Intel(R) Core(TM) i9-9900KF CPU @ 3.60GHz) is used to run the MPC which sends the plan to the Riccati-gain lower controller at 50 Hz.

B. EXPERIMENTS

The pipeline has been tested in various scenarios, first with onboard perception and then with a model of the environment. This is to emphasize the motion generation part and break away from the perception constraints. During all experiments, the user commands the robot's velocity with a joystick. Some of the experiments (1.3, 1.4 and 2.1) were conducted using the inverse dynamic formulation and presented in the related paper [37] as evidence of the formulation's effectiveness. We present these experiments again in this paper to specifically highlight the planning aspect. We release the environments tested for this paper in a public repository (<https://github.com/thomascbrs/walkgen-environments>). The source code for our planner is available as an open source package (<https://github.com/loco-3d/slim.git>) and we will release the rest of our code upon acceptance of the paper.

1) Using onboard perception

We evaluated our complete pipeline on three major scenarios, listed below.

- *Experiment 1.1* (Fig. 1): The first experiment is a 5-minute experiment, representative of the type of environment our perceptive locomotion pipeline enables. The robot starts at the bottom of an industrial staircase with 7 steps, each 17 cm high and 29 cm deep. The average terrain slope is 30 degrees. Once at the top, the robot makes a U-turn on the industrial platform and performs its descent.
- *Experiment 1.2* (Fig. 8-a): The second scenario corresponds to two platforms of 1 by 1 meters connected by a piece of wood placed diagonally to the right of the assembly. Once the robot is on the first platform, we manually remove the ground from the contact surface list to prevent the robot from moving forward. A 20x30 cm block was then added to the left of the assembly, and after a few seconds, the platform is detected and the robot moved forward. This demonstrates the pipeline's reactive capabilities.

- *Experiment 1.3* (Fig. 8-b): It is the same configuration as the previous experiment with 2 pieces of wood connecting the two platforms. This experiment highlights the accurate execution of footstep plans.
- *Experiment 1.4* (Fig. 8-c): Using the robot onboard perception, the final experiment aims to mix various terrains types and heights. Since the metallic plate was not fixed, it slid at the end of the experiment, showing the stability of the controller.

2) Using a model of the environment

To further evaluate the locomotion capabilities of our pipeline, we conducted additional experiments in which the contact surfaces are known a priori, and not given by Plane-Seg. The robot's pose with respect to the world frame is still estimated using the onboard sensors (LIDAR and proprioceptive sensors) and in general, the rest of the framework remains the same.

- *Experiment 2.1* (Fig. 8-d) : The first experiment was to climb stairs with 2 missing steps. We removed steps 2 and 6 of the stairs. It corresponds to climbing a slope of 30 degrees with two gaps of 34 cm. This pushes the robot to its kinematic and the actuation limits. This, therefore, highlights the benefit of taking into account the entire robot model to plan the motion and adapt the posture. This justifies our approach, as discussed in Sec. VIII-E.
- *Experiment 2.2* (Fig. 8-e): The second experiment is similar and corresponds to descending stairs with 1 missing step (step 6). It corresponds to crossing a slope of 30 degrees with a gap of 34 cm.
- *Experiment 2.3* (Fig. 8-f) : The last experiment was conducted on a platform of size 1 m × 1 m and 38 cm in height. It is higher than the robot's height in its nominal position. It is the only experiment where we had to increase the safety margin around the obstacle (outer margin in Sec. IV) to 12 cm since the robot's shoulders are at the same height as the obstacle and collide with it.

3) Simulation experiments

A final set of experiments were run in the PyBullet simulator [60] to show more dynamic gaits such as trotting. This was done to illustrate the diversity of terrains the framework can target. The setup and communication protocol are identical to the one used on the hardware and described previously in Sec. VIII-A.

- *Experiment 3.1* (Fig. 8-g): The first experiment corresponds to a set of stepping stones of different sizes and heights crossed with a trotting gait.
- *Experiment 3.2* (Fig. 8-h): The second experiment involves the ICRA 2023 Quadruped Challenge, which comprises a parkour task featuring a range of obstacles, including pallets, inclined ramps, rounded rubber ramps, and a stack of plastic crates. This challenge



FIGURE 8. Screenshots of different experiments highlighting our architecture's capabilities. The first three rows of results were obtained using the onboard camera and a complete perception pipeline for active obstacle detection. Rows (d), (e), and (f) were performed without perception setup but instead used a pre-computed model of the environment to overcome the perception system limitations and test the controller's limits. The last two rows were obtained through simulation to further test our pipeline in challenging scenarios. Video accessible at <https://youtu.be/bNVxTh0eIXI>.

serves as a valuable benchmark for evaluating our pipeline capabilities. For most parts, our pipeline navigates parkour using a dynamic trot. For some of them, we had to marginally modify some parameters. The wooden pallets with gaps between the wooden slats imply many potential surfaces for each contact, up to 12, even in the pre-selection section. Hence, we had to reduce the surface planner's horizon to optimise the next 4 contact sequences (and not 6 as used previously). Step height need to be increased by up to 25 cm on flat terrain with wood panels of 20 cm. It would be ideal to include height map directly in trajectory optimization. This is rather than solely relying on surfaces for obstacle avoidance. as discussed previously in Sec. VI. Furthermore, when dealing with a stack of plastic crates, adapting the gait to a walking gait (in which only one foot moves at a time) is more stable and enables safe passage across the terrain, compared to a more dynamic trotting gait.

C. EVALUATION OF THE PERCEPTION PIPELINE




FIGURE 9. Perception output during experiment 1-(ii) before ground removal. Black surfaces: initial convex surfaces; coloured surfaces: filtered.

In this section, we evaluate the perception pipeline, from the heightmap to the resulting surfaces filtered and reshaped for security margins. Fig. 9 shows the scene during experiment 1.1, before the manual removal of the ground. The necessity of the filtering routine, described in Sec. IV is highlighted in this example since the convex planes




FIGURE 10. Height map and resulting surfaces during experiment 1-(i) (first row) and experiment 1-(ii) (second row). (<https://youtu.be/bNVxTh0eIXI>).

extracted from the height map overlap. Additionally, we note that the security margin is applied inside the obstacles to avoid walking on the edges. This is due in part to estimation errors, which are around evaluated to 3-4 cm. More height maps and resulting potential surfaces are shown in Fig 10, taken from experiments 1.1 and 1.2. It is interesting to note the evolution of the terrain estimation around the robot during these experiments. When climbing or descending stairs, the robot's camera only identifies the next 1 or 2 stairs ahead. During the second experiment 1.2, where an obstacle was added in front of the robot, it took several iterations of the probabilistic algorithm to update the heightmap with the detected object and consequently 2 to 3 seconds to obtain a feasible surface to walk on.

D. COMPUTATION PERFORMANCE

Table 2 presents the computation time statistics for each module of the pipeline during experiment 1.1, i.e., climbing up and down stairs with active perception.

1) Surface Processing

It takes 90 ms on average to post-process the incoming surfaces from perception with Algorithm IV-C3 presented in Sec. IV. In comparison of the plane's update frequency, which is between 0.5 Hz and 1 Hz with a non-optimised code, it represents roughly an increase of 5-10%.

2) Surface Selection

Surface selection, as described in Sec. V encompasses the pre-selection step to reduce the number of potential surfaces

TABLE 2. Computing time

Name	Mean	Min	Max
<i>Surface Processing</i>			
Number of surfaces processed	23.71	8.	45.
Surface processing [ms]	88.89	19.49	164.22
<i>Surface Selection</i>			
Number of potential surfaces	3.08	1.	9.
Pre-selection [ms]	17.16	5.47	47.46
MIP [ms]	98.62	32.95	253.15
<i>Foot Trajectory</i>			
Foot location [ms]	0.14	0.08	0.74
Foot trajectory [ms]	0.082	0.007	0.33
<i>Motion Generation (MPC)</i>			
Solve time [ms]	9.29	8.02	13.52
Total time [ms]	13.27	8.76	18.50

considered by the mixed-integer program. It takes around 120 ms to find the next surface of contact in this experiment. Pre-selection reduces the potential surfaces from 23 to 3 for each foot. During this experiment, we optimise over 8 contact phases, which correspond to 8-foot positions optimised with a walking gait (1 contact is created/broken for each contact phase). The mixed-integer optimisation starts at the beginning of each upcoming foot trajectory and the next surface information needs to be received before the beginning of the upcoming cycle. In this experiment, the foot trajectory is 600 ms, which is enough for the maximum time taken. However, for a trotting gait, the foot trajectory is set to 300 ms. We must reduce the number of foot locations to 6, thus allowing 3 contact phases (2-foot locations optimised for each contact phase) to ensure a safe margin regarding the computing time. We note that the computing time of mixed-integer optimisation depends heavily on the number of potential surfaces for each contact and the number of contact locations optimised. For a detailed analysis of the MIP computation times, we refer the reader to [16].

3) Foot trajectory

The foot trajectory encompasses the optimisation of the footstep inside the attributed surface and the end-effector curve optimisation as described in Sec. VI. It represents only 1% of every MPC step.

4) Motion generation (MPC)

On average the MPC step takes around 13 ms, including the update of the optimal control problem and solving this latter with 1 iteration which takes around 9 ms. This is below the maximum 20 ms time required for an MPC running at 50 Hz.

E. ANALYSIS OF THE WHOLE-BODY MPC

In scenarios involving climbing steps, the torque limit is reached on at least one actuator in each of these experiments. It is in this case that the whole-body MPC becomes crucial as it adjusts body posture to reduce joint torques. Experiment 2.1 is a representative case. First, we observe that the COM motion leans forward and close to the feet, almost in contact with the stairs, at the moment of giving the last motion to cross the gap (Fig. 11). The torque limit is reached on both joints of the hind right leg during this motion, as shown in Fig. 12. We analyse the following quantities: torques, angular position and angular velocities at the HFE joint (hip flexion/extension) and the KFE joint (knee flexion/extension). The HAA joints (hip abduction/adduction) are less prone to reach torque limits. While crossing the gap, we can observe two torque peaks that correspond exactly to the robot configuration shown in Fig. 12. In this instance, the overshoot of the torque command above the joint limits (a few Nm) can be attributed to two main reasons. First, the constraints on torque limits can be violated (Sec. VII) as the Riccati controller guarantees joint limits within a neighbourhood. Second, the robot is not directly torque-controlled by the Riccati controller (feedforward + state feedback). A PD term is also used to achieve more robust and compliant behavior, which, in addition to the Riccati term, can lead to small overshoots.

F. EVALUATION OF THE COLLISION-FREE FOOT TRAJECTORY

The end-effector trajectory is re-computed at 50 Hz in order to get robust tracking. Some minor modifications were necessary when transitioning from simulation to hardware. An offset of -1 cm has been added on the z-axis to accommodate for perception errors, reaching 2/3 cm at the end of our longest experiments (1.1), and ensuring contact creation occurs. A finer control of the feet's trajectory according to the contact detection was not necessary in our case, as the whole-body MPC is robust to state estimation errors (Sec. VIII-H). Additionally, the swing-foot trajectory and footstep are no longer updated once 70% of the flight phase has passed in order to avoid a sliding contact after a sudden change in the target position. Finally, the end-effector velocity feedback is not taken into account in the optimisation due to large uncertainties about the end-effector velocities. The foot velocity is therefore assumed to be tracked properly and the initial velocity while re-computing the curve is taken from the previous control cycle.

Fig. 13 shows a swing-foot trajectory while the robot crosses an obstacle of 20 cm and while climbs the stairs during experiment 1.1. We observe the state estimation errors on the foot position when landing on the ground, resulting in the foot slightly bouncing.




FIGURE 11. Posture adjustment (experiment 2.1). The hind right leg reaches the torque limits on both HFE (hip flexion/extension) and KFE (knee flexion/extension) when crossing the gap. This corresponds to the peak torque at 52s. The whole-body MPC adjusts the body posture to compensate. The body leans forward and lowers as much as possible, reaching the kinematic limit of the hind leg to reduce torque on it. (<https://youtu.be/bNVxTh0eixI>).




FIGURE 12. Torque, angular position and velocities of the joints HAA (hip abduction/abduction), HFE (hip flexion/extension) and KFE (knee flexion/extension) of the hind right leg during experiment 2.1. Blue quantities are computed by the state feedback controller and orange ones are measured on the robot. The maximum torque limit is approximately 40 N.m and is represented by red dashed curves. Each actuator can reach 12 m.s^{-1} . Two torque peaks reaching the boundaries can be observed around the 40 s and 52 s which correspond to the robot configurations shown above.

G. EVALUATION OF THE VELOCITY TRACKING AND DESIGN CHOICES

Not constraining the robot's pose and velocity to follow a reference trajectory has a significant impact on the COM




FIGURE 13. Side view of the front-left foot's trajectory while crossing a 20 cm obstacle (left) and the staircase (right) using a walking gait pattern at a reference velocity of 0.1 m.s^{-1} . The grey curves on the left graph are reference Bezier curves computed during each control cycle. Their transparency indicates the prediction horizon. More transparent curves represent further predictions for the future. Grey areas on the right correspond to the cumulative position of surfaces detected by the camera.




FIGURE 14. Screenshots of the motion resulting from our OCP regularization setup for a walking gait pattern. The blue area between represents the support polygon and the yellow circle is the Centre Of Pressure (COP) applied.

trajectory, especially in a walking gait pattern when only one foot is lifted at a time (Fig. 14). Indeed, Fig. 15-(a) shows that the base trajectory during a forward walk oscillates in a range of 10 cm around the middle of the feet positions.




FIGURE 15. Base trajectory during a forward walking gait of period 2.4s. The joystick command is in 2 separate steps of 0.05m.s^{-1} and 0.1m.s^{-1} along the x-axis. To reject disturbances, the filter is a moving average on the walk period. (a) base position on the ground floor plane. (b) and (c) base velocity along respectively the x and y axes. Small disturbances at first in the filtered quantities are due to the filter initialisation phase.

However, constraining the base position and velocity would result in a fixed base orientation. This can be explained by the motion stability found in the optimization process. The centre of pressure is well located in the middle of the support polygon. The COM position is dragged inside this region resulting in this wave motion. Since it is a periodic motion, the state position has been filtered with a moving average for the walk period (Fig. 15), rejecting all frequencies in synchronisation with the gait. This behaviour does not appear with a trotting gait since two opposite legs are left at the same time and the resulting oscillation is much smaller.

H. ROBUSTNESS OF THE PIPELINE

A crucial point is to understand the repeatability of our experiments and how the locomotion controller adapts to unforeseen events. We observed a significant difference in the reliability whether on-board perception was used or perfect knowledge of the environment was given. The climbing stairs scenario was reliably repeated ten times with active perception, with a 80% success rate, with errors mainly due to state estimation errors and drift. For climbing down the stairs the experiment was successfully repeated twice, but perception issues made the experiment difficult, as the camera position did not allow to clearly perceive all

the stairs on the way down, resulting in unfeasible contact planning problems in some instances. In the absence of the perceptive part (experiments 2.1, 2.2 and 2.3), when the environment is perfectly known, the locomotion pipeline is stable and the experiments were successfully carried out on the first attempt and repeated twice.

An interesting point attributed to the whole-body MPC is the ability to recover to unplanned situations, as in experiment 2.2 when ANYmal misses a step during the descent. Similarly, at the end of experiment 1.4, a metal plate slips when walking on it. In both cases, the robot recovered properly (Fig. 16).

1) Importance of the preview window

In our previous work, we have conducted an ablation study to analyse the effect of the different elements of the pipeline on the success rate [24]. In particular it was established that for the solo robot, planning several steps ahead was an essential contribution to the success rates. In the experiments considered in the present work, the motions are not dynamic enough to establish empirically that a long preview window as a strong benefit to the robustness of the pipeline.




FIGURE 16. Screenshots of the robot recovering from an unplanned event. Top: the metal plate slips while the robot walks on it. Bottom: the two front feet slip and miss the intended step.

I. COMPARISON WITH THE STATE OF THE ART

The end-to-end implementation of frameworks similar to ours [30], [32] are not publicly available at the time of submission, neither the associated benchmarks and the quantity of work and resources to re-implement them is not reasonable, which prevents a quantitative comparison of the approaches.

IX. DISCUSSION

In this paper, we propose a complete pipeline for perceptive quadrupedal locomotion; from onboard perception to locomotion generation and control. Our experimental results highlighted capabilities such as crossing challenging

scenarios like climbing and descending industrial stairs or obstacle parkour with moving obstacles. Therefore, we have experimentally validated our architecture. This is based on the strong design assumption that a high-level planner, such as mixed-integer optimization, only selects the stepping contact surfaces. The utility of mixed-integer optimisation for design integration was partially validated in our previous work [24]. We used a different robot and integrated perception into our approach along with more complex experiments which reinforced interest in our decomposition.

A. PERCEPTION PIPELINE

We have demonstrated the efficacy of our perceptive pipeline based on the probabilistic terrain mapping method [28] and the Plane-Seg algorithm [35] which allow us to obtain satisfactory behaviour, in particular with our filtering work. Additionally, we demonstrated that our architecture can successfully navigate complex parkour terrains with moving objects or industrial staircases. Still, while our architecture is capable of navigating challenging scenarios when perception is removed, improving perception capabilities could enhance the overall robustness and reliability of the system. This would be particularly beneficial in scenarios such as the staircase where vision detection is restricted to only a few steps ahead due to our single depth camera with a 20% downward angle. Additionally, the algorithm used to extract convex planes is not incremental and does not take into account previously computed planes, making it vulnerable to faulty scans in some cases. The implementation of a short term memory for the plane decomposition system would contribute to the robustness of the approach and will be considered in future work.

B. COLLISION AVOIDANCE OF THE BODY

The trade-off proposed in our approach is to strongly regularise the OCP around a reference end-effector trajectory while avoiding obstacles, which proved to work in a wide range of environments. This occurs even if a collision with the environment is not specifically considered in our approach. This could be considered within the MPC but would require a dedicated study since it is a challenging nonlinear optimisation problem. The safety margin around obstacles was sufficient in almost all scenarios to avoid collisions, except for the 40 cm step in 2.3, where we had to increase these margins to 10 cm. As mentioned before, our approach ensured that there was no collision in the end-effector trajectory. It could have been done inside the MPC with additional terms in the objective. However, this would increase complexity and possibly affect convergence rates. Knee collision is implicitly considered by the margin around obstacles. However, this could be addressed explicitly by planning the foot trajectory including the whole leg, as is the case in [25], although re-planning at 50 Hz might become unfeasible.

C. COMPARISON WITH THE STATE OF THE ART

Two recent works [30], [32] present interesting similarities to our architecture. We can first observe the decomposed approach used in both cases with the foot location optimised outside the MPC. We have underlined the technical difficulties that would allow a objective comparison of the frameworks, although we would like to integrate the publicly available perception module proposed in [29]. However, this approach does not return a selection of potential candidate surfaces and only one candidate, such that further integration will be required to make the approach compatible with our framework.

An important question when comparing the approaches relate to the fact that we put an emphasis on having a horizon windows as large as possible when planning. Other work suggest that planning only for the next step might be sufficient [28]. The experiments in this paper do not allow to infirm this hypothesis. However, in our previous work it was established that for the Solo robot, the preview was indeed required [24]. We hypothesise that in the current state of things the preview is not required because the considered motions are not dynamic enough, but future work will be needed to reach stronger conclusions.

D. IMPROVEMENT POINTS

As often occurs in model-based approaches to quadruped robots, the main limitation is the gait fixed beforehand. It would be interesting to replace the high-level planner with an acyclic planner to optimise the timing of contacts and the type of gait. This could lead to computing-time issues especially due to the increase in problem complexity. To take this further, it would be interesting to integrate this into a global planner to achieve autonomous behaviour. Although we consider a stable walking gait necessary for the most challenging scenarios, we also acknowledge the potential benefits of implementing a trotting gait on the hardware. However, transitioning to a trotting gait, which works well in simulation, proved more challenging to implement on the robot. We believe that this limitation could be overcome with a faster controller and more integrative efforts. Finally, our foot placement is based on Raibert's heuristic, which is optimal on flat ground when considering an inverse linear pendulum model for the robot. We extend it in 3D which produces satisfying foot placement as we demonstrated in our experiment. Nevertheless, this remains a heuristic and does not ensure the position of feet is feasible in terms of torque power limit.

X. CONCLUSION

In this paper, we present a complete methodology for crossing challenging terrains; from terrain perception to locomotion generation and control. We have demonstrated our pipeline on various terrains, such as an industrial staircase or a parkour-type environment with a moving object in the scene. These experiments have allowed us to further validate our approach based on a sub-division of the global problem.

First, a high-level planner formulated as a mixed-integer optimisation selects only the next surfaces of contact with a horizon of a few contacts (6 to 8 in our experiments). To achieve this, one must adapt the perception to extract relevant potential surfaces as convex planes. For motion generation, we rely on an efficient whole-body MPC and a linear state feedback controller. Collision-free trajectory and footstep adaptation on the high-level planner are optimised separately. The OCP problem is then strongly regularized around this reference end-effector trajectory while robot's posture is adapted by the MPC. This represents a wise choice of design to leverage whole-body optimisation capabilities. To move further, we would like to use a more complex high-level planner to optimise contact timing to cross even more challenging terrains, for example, dynamic motions that include jumping over gaps or obstacles.

REFERENCES

- [1] Marco Hutter, Christian Gehring, Dominic Jud, Andreas Lauber, C. Dario Bellicoso, Vassilios Tsounis, Jemin Hwangbo, Karen Bodie, Peter Fankhauser, Michael Bloesch, Remo Diethelm, Samuel Bachmann, Amir Melzer, and Mark Hoepfner. Anymal - a highly mobile and dynamic quadrupedal robot. In *IEEE/RSJ Int. Conf. Intell. Rob. Sys. (IROS)*, pages 38–44, 2016.
- [2] Unitree. B1. <https://www.unitree.com/b1/>, 2021. Accessed: December 2022.
- [3] Boston Dynamics. Spot. <https://www.bostondynamics.com/>, 2016.
- [4] Dario Bellicoso, Marko Bjelonic, Lorenz Wellhausen, Kai Holtmann, Fabian Günther, Marco Tranzatto, Péter Fankhauser, and Marco Hutter. Advances in real-world applications for legged robots. *J. Field Robots.*, 2018.
- [5] T. Bretl. Motion planning of multi-limbed robots subject to equilibrium constraints: The free-climbing robot problem. *The Int. J. of Rob. Res. (IJRR)*, 25, 2006.
- [6] Kris Hauser, Timothy Bretl, Jean-Claude Latombe, Kensuke Harada, and Brian Wilcox. Motion planning for legged robots on varied terrain. *The Int. J. of Rob. Res. (IJRR)*, 27(11-12), 2008.
- [7] Adrien Escande, Abderrahmane Kheddar, Sylvain Miossec, and Sylvain Garsault. Planning support contact-points for acyclic motions and experiments on hrp-2. In *Experimental Robotics*, pages 293–302. Springer, 2009.
- [8] Karim Bouyarmane, Adrien Escande, Florent Lamiraux, and Abderrahmane Kheddar. Potential field guide for humanoid multicontacts acyclic motion planning. In *IEEE Int. Conf. Rob. Autom. (ICRA)*, pages 1165–1170, 2009.
- [9] Steve Tonneau, Andrea Del Prete, Julien Pettré, Chonhyon Park, Dinesh Manocha, and Nicolas Mansard. An efficient acyclic contact planner for multipled robots. *IEEE Trans. Robot. (T-RO)*, 34(3), 2018.
- [10] Masaki Murooka, Iori Kumagai, Mitsuharu Morisawa, Fumio Kanehiro, and Abderrahmane Kheddar. Humanoid Loco-Manipulation Planning Based on Graph Search and Reachability Maps. *IEEE Robot. Automat. Lett. (RA-L)*, 6(2), 2021.
- [11] Andrew Short and Tirthankar Bandyopadhyay. Legged motion planning in complex three-dimensional environments. *IEEE Robot. Automat. Lett. (RA-L)*, 3(1):29–36, 2018.
- [12] Steve Tonneau, Pierre Fernbach, Andrea del Prete, Julien Pettré, and Nicolas Mansard. 2PAC: Two Point Attractors for Center of Mass Trajectories in Multi Contact Scenarios. *ACM Trans. Graphics*, 37(5), 2018.
- [13] Pierre Fernbach, Steve Tonneau, Olivier Stasse, Justin Carpentier, and Michel Taix. C-CROC: Continuous and Convex Resolution of Centroidal Dynamic Trajectories for Legged Robots in Multicontact Scenarios. *IEEE Trans. Robot. (T-RO)*, 36, 2020.
- [14] Igor Mordatch, Emanuel Todorov, and Zoran Popovic. Discovery of complex behaviors through contact-invariant optimization. *ACM Trans. Graphics*, 31, 2012.
- [15] Alexander W Winkler, Dario C Bellicoso, Marco Hutter, and Jonas Buchli. Gait and trajectory optimization for legged systems through phase-based end-effector parameterization. *IEEE Robot. Automat. Lett. (RA-L)*, 3, 2018.
- [16] Daeun Song, Pierre Fernbach, Thomas Flayols, Andrea Del Prete, Nicolas Mansard, Steve Tonneau, and Young J. Kim. Solving Footstep Planning as a Feasibility Problem Using L1-Norm Minimization. *IEEE Robot. Automat. Lett. (RA-L)*, 6, 2021.
- [17] Michael Posa, Cecilia Cantu, and Russ Tedrake. A direct method for trajectory optimization of rigid bodies through contact. *The Int. J. of Rob. Res. (IJRR)*, 33, 2014.
- [18] Olivier Melon, Mathieu Geisert, David Surovik, Ioannis Havoutis, and Maurice Fallon. Reliable trajectories for dynamic quadrupeds using analytical costs and learned initializations. In *IEEE Int. Conf. Rob. Autom. (ICRA)*, pages 1410–1416, 2020.
- [19] Olivier Melon, Romeo Orsolino, David Surovik, Mathieu Geisert, Ioannis Havoutis, and Maurice Fallon. Receding-horizon perceptive trajectory optimization for dynamic legged locomotion with learned initialization. In *IEEE Int. Conf. Rob. Autom. (ICRA)*, pages 9805–9811, 2021.
- [20] Robin Deits and Russ Tedrake. Footstep planning on uneven terrain with mixed-integer convex optimization. *IEEE Int. Conf. Hum. Rob. (ICHR)*, 2014.
- [21] B. Aceituno-Cabezas, C. Mastalli, H. Dai, M. Focchi, A. Radulescu, D. G. Caldwell, J. Cappelletto, J. C. Grieco, G. Fernandez-Lopez, and C. Semini. Simultaneous contact, gait, and motion planning for robust multilegged locomotion via mixed-integer convex optimization. *IEEE Robot. Automat. Lett. (RA-L)*, 2018.
- [22] Brahayam Ponton, Majid Khadiv, Avadesh Meduri, and Ludovic Righetti. Efficient Multicontact Pattern Generation With Sequential Convex Approximations of the Centroidal Dynamics. *IEEE Trans. Robot. (T-RO)*, 37, 2021.
- [23] Lorenzo Amatiucci, Joon-Ha Kim, Jemin Hwangbo, and Hae-Won Park. Monte carlo tree search gait planner for non-gaited legged system control. In *IEEE Int. Conf. Rob. Autom. (ICRA)*, pages 4701–4707, 2022.
- [24] Fanny Risbourg, Thomas Corbères, Pierre-Alexandre Léziart, Thomas Flayols, Nicolas Mansard, and Steve Tonneau. Real-time footstep planning and control of the solo quadruped robot in 3d environments. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12950–12956, 2022.
- [25] Matt Zucker, Nathan Ratliff, Martin Stolle, Joel Chestnutt, J Andrew Bagnell, Christopher G Atkeson, and James Kuffner. Optimization and learning for rough terrain legged locomotion. *The Int. J. of Rob. Res. (IJRR)*, 30, 2011.
- [26] J. Zico Kolter, Mike P. Rodgers, and Andrew Y. Ng. A control architecture for quadruped locomotion over rough terrain. In *IEEE Int. Conf. Rob. Autom. (ICRA)*, pages 811–818, 2008.
- [27] Mrinal Kalakrishnan, Jonas Buchli, Peter Pastor, Michael Mistry, and Stefan Schaal. Learning, planning, and control for quadruped locomotion over challenging terrain. *The Int. J. of Rob. Res. (IJRR)*, 30, 2011.
- [28] Péter Fankhauser, Michael Bloesch, and Marco Hutter. Probabilistic terrain mapping for mobile robots with uncertain localization. *IEEE Robot. Automat. Lett. (RA-L)*, 3, 2018.
- [29] Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science Robotics*, 7, 2022.
- [30] Fabian Jenelten, Ruben Grandia, Farbod Farshidian, and Marco Hutter. TAMOLS: Terrain-aware motion optimization for legged systems. *IEEE Trans. Robot. (T-RO)*, 38, 2022.
- [31] Ruben Grandia, Andrew J. Taylor, Aaron D. Ames, and Marco Hutter. Multi-layered safety for legged robots via control barrier functions and model predictive control. *CoRR*, abs/2011.00032, 2020.
- [32] Ruben Grandia, Fabian Jenelten, Shaohui Yang, Farbod Farshidian, and Marco Hutter. Perceptive locomotion through nonlinear model-predictive control. *IEEE Trans. Robot. (T-RO)*, 39(5):3402–3421, 2023.
- [33] Justin Carpentier and Nicolas Mansard. Multicontact Locomotion of Legged Robots. *IEEE Trans. Robot. (T-RO)*, 34, 2018.
- [34] Carlos Mastalli, Rohan Budhiraja, Wolfgang Merkt, Guilhem Saurel, Bilal Hammoud, Maximilien Naveau, Justin Carpentier, Ludovic Righetti, Sethu Vijayakumar, and Nicolas Mansard. Crocodyl: An efficient and versatile framework for multi-contact optimal control. In *IEEE Int. Conf. Rob. Autom. (ICRA)*, pages 2536–2542, 2020.
- [35] Maurice Fallon and Matt Antone. Plane Seg – Robustly and Efficiently Extracting Contact Regions from Depth Data, 2019.
- [36] Carlos Mastalli, Wolfgang Merkt, Guiyang Xin, Jaehyun Shim, Michael Mistry, Ioannis Havoutis, and Sethu Vijayakumar. Agile maneuvers in legged robots: a predictive control approach, 2022.

[37] Carlos Mastalli, Saroj Prasad Chhatoi, Thomas Corbères, Steve Tonneau, and Sethu Vijayakumar. Inverse-Dynamics MPC via Nullspace Resolution. *IEEE Trans. Robot. (T-RO)*, 2023.

[38] Michael Bloesch, Marco Hutter, Mark Hoepfner, Stefan Leutenegger, Christian Gehring, C. David Remy, and Roland Siegwart. State estimation for legged robots - consistent fusion of leg kinematics and IMU. In *Proceedings of Robotics: Science and Systems*, Sydney, Australia, July 2012.

[39] M. Hutter, C. Gehring, A. Lauber, F. Gunther, C. D. Bellicoso, V. Tsounis, P. Fankhauser, R. Diethelm, S. Bachmann, M. Bloesch, H. Kolvenbach, M. Bjelonic, L. Isler, and K. Meyer. Anymal - toward legged robots for harsh environments. *Advanced Robotics*, 31(17), 2017.

[40] Mahes Visvalingam and J. Duncan Whyatt. Line generalisation by repeated elimination of points. *Cartographic Journal*, 30:46–51, 1993.

[41] azrafe7. *hxGeomAlgo*, 2013.

[42] Robin Deits and Russ Tedrake. Efficient mixed-integer planning for uavs in cluttered environments. *IEEE Int. Conf. Rob. Autom. (ICRA)*, 2015.

[43] F. Grimminger, A. Meduri, M. Khadiv, J. Viereck, M. Wüthrich, M. Naveau, V. Berenz, S. Heim, F. Widmaier, T. Flayols, J. Fiene, A. Badri-Spröwitz, and L. Righetti. An open torque-controlled modular robot architecture for legged locomotion research. *IEEE Robot. Automat. Lett. (RA-L)*, 5(2), 2020.

[44] E.G. Gilbert, D.W. Johnson, and S.S. Keerthi. A fast procedure for computing the distance between complex objects in three-dimensional space. *IEEE Journal on Robotics and Automation*, 4(2):193–203, 1988.

[45] Justin Carpentier, Guilhem Saurel, Gabriele Buondonno, Joseph Mirabel, Florent Lamiraux, Olivier Stasse, and Nicolas Mansard. The pinocchio c++ library : A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives. In *IEEE Int. Sym. System Integration (SII)*, pages 614–619, 2019.

[46] Wieber P.-B. *Holonomy and Nonholonomy in the Dynamics of Articulated Motion*, pages 411–425. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

[47] Johan Lofberg. Big-m and convex hulls. <https://yalmip.github.io/tutorial/bigmandconvexhulls>, 2022.

[48] Marc H Raibert. *Legged robots that balance*. MIT press, 1986.

[49] Kim, Jared Di Carlo, Benjamin Katz, Gerardo Bledt, and Sangbae Kim. Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control. *arXiv: 1909.06586*, 2019.

[50] Jared Di Carlo, Patrick M. Wensing, Benjamin Katz, Gerardo Bledt, and Sangbae Kim. Dynamic locomotion in the mit cheetah 3 through convex model-predictive control. In *IEEE/RSJ Int. Conf. Intell. Rob. Sys. (IROS)*, pages 1–9, 2018.

[51] Pierre-Alexandre Léziart, Thomas Flayols, Felix Grimminger, Nicolas Mansard, and Philippe Souères. Implementation of a reactive walking controller for the new open-hardware quadruped solo-12. In *IEEE Int. Conf. Rob. Autom. (ICRA)*, pages 5007–5013, 2021.

[52] Mylène Campana, Florent Lamiraux, and Jean-Paul Laumond. A gradient-based path optimization method for motion planning. *Advanced Robotics*, 30(17-18), 2016.

[53] José-Luis Blanco. A tutorial on se(3) transformation parameterizations and on-manifold optimization. Technical report, University of Malaga, September 2010.

[54] J. Baumgarte. Stabilization of constraints and integrals of motion in dynamical systems. *Computer Methods in Applied Mechanics and Engineering*, 1:1–16, 1972.

[55] Rohan Budhiraja, Justin Carpentier, Carlos Mastalli, and Nicolas Mansard. Differential dynamic programming for multi-phase rigid contact dynamics. In *IEEE Int. Conf. Hum. Rob. (ICHR)*, pages 1–9, 2018.

[56] Tom Erez and Emanuel Todorov. Trajectory optimization for domains with contacts using inverse dynamics. *IEEE/RSJ Int. Conf. Intell. Rob. Sys. (IROS)*, 2012.

[57] Carlos Mastalli, Ioannis Havoutis, Michele Focchi, Darwin G. Caldwell, and Claudio Semini. Motion planning for quadrupedal locomotion: Coupled planning, terrain mapping, and whole-body control. *IEEE Trans. Robot. (T-RO)*, 36(6):1635–1648, 2020.

[58] Pierre-Alexandre Léziart, Thomas Corbères, Thomas Flayols, Steve Tonneau, Nicolas Mansard, and Philippe Souères. Improved Control Scheme for the Solo Quadruped and Experimental Comparison of Model Predictive Controllers. *IEEE Robot. Automat. Lett. (RA-L)*, 7, 2022.

[59] Morgan Quigley. Ros: an open-source robot operating system. In *IEEE Int. Conf. on Rob. and Aut. (ICRA)*, 2009.

[60] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2021.



THOMAS CORBÈRES received his M.Sc. degree in aerospace Engineering engineering from ISAE-SUPAERO Toulouse, France in 2020.

He is currently pursuing a Ph.D. degree in robotics at the University of Edinburgh under the University of S. Tonneau. His research interests include legged locomotion, model predictive control, and contact planning.



CARLOS MASTALLI received an M.Sc. degree in mechatronics engineering from the Simón Bolívar University, Caracas, Venezuela, in 2013 and a Ph.D. degree in bio-engineering and robotics from the Istituto Italiano di Tecnologia, Genoa, Italy, in 2017.

He is currently an Assistant Professor at Heriot-Watt University, Edinburgh, U.K. He is the Head of the Robot Motor Intelligence (RoMI) Lab affiliated with the National Robotarium and Edinburgh Centre for Robotics. He is also appointed as Research Scientist at IHMC, USA. Previously, he conducted cutting-edge research in several world-leading labs: Istituto Italiano di Tecnologia (Italy), LAAS-CNRS (France), ETH Zürich (Switzerland), and the University of Edinburgh (UK). His research focuses on building athletic intelligence for robots with legs and arms. Carlos' research work is at the intersection of model predictive control, numerical optimization, machine learning, and robot co-design.



WOLFGANG MERKT received the B.Eng.(Hns) degree in mechanical engineering with management and the M.Sc.(R) and Ph.D. degrees in robotics and autonomous systems from the University of Edinburgh, Edinburgh, U.K., in 2014, 2015 and 2019, respectively.

He is currently a Postdoctoral Researcher at the Oxford Robotics Institute, University of Oxford. During his Ph.D., he worked on trajectory optimization and warm starting optimal control for high-dimensional systems and humanoid robots under the supervision of S. Vijayakumar. His research interests include fast optimization-based methods for planning and control, loco-manipulation, and legged robots.



IOANNIS HAVOUTIS received the M.Sc. in artificial intelligence and Ph.D. in Informatics degrees from the University of Edinburgh, Edinburgh, U.K. He is a Lecturer in Robotics at the University of Oxford.

He is part of the Oxford Robotics Institute and a co-lead of the Dynamic Robot Systems group. His focus is on approaches for dynamic whole-body motion planning and control for legged robots in challenging domains. From 2015 to 2017, he was a postdoc at the Robot Learning and Interaction Group, at the Idiap Research Institute. Previously, from 2011 to 2015, he was a senior postdoc at the Dynamic Legged System lab the Istituto Italiano di Tecnologia.



JAEHYUN SHIM received a B.Eng. degree in mechanical engineering from the University of Tokyo, Tokyo, Japan in 2015 and an M.A.Sc. degree in mechanical engineering from the University of British Columbia, Vancouver, BC, Canada, in 2018. He is currently a Software Engineer at the University of Edinburgh. He worked on developing autonomous delivery robots as a Software Engineer at ROBOTIS, Seoul, South Korea. His research interests include optimization-based

planning and legged robots.



MAURICE FALLON studied Electronic Engineering at University College Dublin. His Ph.D. research in the field of acoustic source tracking was carried out in the Engineering Department of the University of Cambridge.

After his Ph.D. he moved to MIT as a post doctoral researcher and later research scientist in the Marine Robotics Group (2008-2012). From 2012-2015 he was the perception lead of MIT's team in the DARPA Robotics Challenge. From

2015, he was a Lecturer at University of Edinburgh where he led research in collaboration with NASA's humanoid robotics program. He moved to Oxford in April 2017 and took up the Royal Society University Research Fellowship in October 2017.



NICOLAS MANSARD is CNRS researcher ("chargé de recherche CNRS") since 2009. He received the M.Sc. in computer science of Univ. Grenoble in 2003 and the PhD in robotics of Univ. Rennes in 2006.

He was then post-doctoral researcher in Stanford Univ. with O. Khatib in 2007 and in JRL-Japan with A. Kheddar in 2008. He was invited researcher in Univ. Washington with E.Todorov in 2014. He received the CNRS Bronze Medal in

2015 (one medal is award in France in automatic/robotic/signal-processing every year). His main research interests include the motion generation, planning and control of complex robots, with a special regard in humanoid robotics. His expertise covers sensor-based (vision and force) control, numerical mathematics for control, bipedal locomotion and locomotion planning. He published more than 70 papers in international journals and conferences and supervised 10 PhD thesis. He is currently associate editor of IEEE Transactions on Robotics.



THOMAS FLAYOLS is a contract researcher at LAAS-CNRS in the Gepetto group, where he graduated from his PhD in 2018. Before that he graduated from Ecole Normale Supérieure de Cachan (ENS) with a Master Degree in Embedded Systems and Industrial Computing.

His research interest focuses on force control of legged robots for locomotion and stabilization. He has a particular interest in experimental validation of estimation and control techniques. He is also a

mechatronic engineer attached to open hardware design applied to actuation and perception for robotic. Currently, he is actively contributing to the design of the open source quadruped robot ODRI initiated by the MPI of Tübingen.



SETHU VIJAYAKUMAR received the Ph.D. degree in computer science and engineering from the Tokyo Institute of Technology, Tokyo, Japan, in 1998.

He is Professor of Robotics and a Director with the Institute of Perception, Action, and Behaviour, School of Informatics, University of Edinburgh, U.K., where he holds the Royal Academy of Engineering Microsoft Research Chair in Learning Robotics. He also has additional appointments

as an Adjunct Faculty with the University of Southern California, Los Angeles, CA, USA and a Visiting Research Scientist with the RIKEN Brain Science Institute, Tokyo. His research interests include statistical machine learning, robotics, planning and optimization in autonomous systems to human motor control, and optimality. Dr. Vijayakumar is a Fellow of the Royal Society of Edinburgh.



STEVE TONNEAU received the Ph.D. degree in humanoid robotics from the INRIA/IRISA, France, in 2015.

He is a Lecturer at the University of Edinburgh, Edinburgh, U.K.. Previously, he was a Postdoctoral Researcher at LAAS-CNRS in Toulouse, France. His research focuses on motion planning based on the biomechanical analysis of motion invariants. Applications include computer graphics animation as well as robotics.

...