

Right Model, Right Time: Real-Time Cascaded-Fidelity MPC for Bipedal Walking

Franek Stark^{1,2}, Felix Wiebe^{1,2}, Shubham Vyas^{1,2}, Dennis Mronga¹, Frank Kirchner^{1,2}

Abstract—This paper presents a multi-phase whole-body model predictive control (MPC) approach for bipedal walking, combining a detailed whole-body model in the near horizon with a simplified single-rigid-body model in the later prediction steps. This reduces computational complexity while retaining prediction capabilities. The resulting nonlinear optimal control problem is solved entirely within the general-purpose, off-the-shelf nonlinear MPC framework *acados*, using sequential quadratic programming (SQP). Given a contact schedule and a target walking speed, the controller optimizes joint torques without depending on pre-selected footstep locations. The controller is validated in MuJoCo simulation on the 18-DoF bipedal robot *HyPer-2*.

I. INTRODUCTION

Legged robotics has advanced rapidly in recent years, with dynamic locomotion successfully demonstrated on quadrupeds and, later, on bipedal robots. Thereby, two different control paradigms are prevalent: reinforcement learning (RL) and model predictive control (MPC). While RL has shown impressive results and benefits strongly from improved simulations and computing power, it still faces well-known challenges such as time-consuming reward shaping, sample inefficiency, lack of generalization and open questions regarding safety. For this reason, MPC remains attractive, because it allows the controller to exploit system dynamics online and reason about constraints directly. However, for high-dimensional problems, such as bipedal walking in humanoids, the computational complexity of the optimal control problem remains too high, especially when aiming for an online solution. Hence, many works rely on reduced-order models in MPC, in combination with instantaneous Whole-Body Control (WBC), which uses the full dynamics [1]. A common limitation of these approaches is that contact forces and body motions are planned independently. Here, additional heuristics make systematic tuning of the different control levels difficult and the use of simplified models in MPC may lead to dynamically infeasible motion plans. Recently, whole-body MPC for bipedal robots has become feasible due to advances in fast numerical optimal control, demonstrated with differential dynamic programming (DDP)-based solvers [2] and sequential quadratic programming (SQP)-based approaches [3], [4]. These approaches use a single full-fidelity model throughout the prediction horizon, which limits horizon length under real-time constraints. Multi-phase (cascaded-fidelity) MPC formulations address this by using a detailed model in the early steps of the prediction




Fig. 1. The multi-phase model predictive control (MPC) predicts the first N_{wb} steps using the whole-body (WB) dynamics, while the following N_{srb} steps are predicted using single-rigid-body (SRB) dynamics. This way prediction accuracy and solve time can be balanced effectively.

horizon and simpler models in later steps [5]–[7], reducing problem complexity while preserving predictive capability. These works rely on problem-specific DDP solvers [5], [6] or a custom SQP solver [7], and the schedule-design framework of [8] is restricted to offline settings. In parallel, recent results suggest that tailored SQP solvers are highly competitive for constrained nonlinear MPC [9] and are available in mature general-purpose frameworks such as *acados* [10], which has direct support for multi-phase problems [11].

Motivated by these developments, this paper presents a whole-body multi-phase MPC for bipedal walking. The proposed controller employs a whole-body (WB) model in the early steps of the horizon and a single-rigid-body (SRB) model in the later steps (see Fig. 1) to reduce computational cost, while solving the resulting nonlinear optimal control problem in a real-time SQP setting with only a small number of iterations per control cycle. In contrast to prior multi-phase MPC works [5]–[7], which rely on custom DDP solvers or hand-built SQP layers, the entire formulation is implemented in *acados* [10], a general-purpose nonlinear MPC framework. We additionally provide an empirical characterization of how the WB/SRB split, discretization steps, and number of SQP iterations jointly affect real-time feasibility and tracking accuracy. Our controller was successfully tested on the 18 DoF bipedal robot *HyPer-2* in a MuJoCo simulation using purely torque control and a control frequency of 100 Hz.

II. MODEL PREDICTIVE CONTROL FORMULATION

The proposed MPC scheme is formulated as the following two-phase optimal control problem, in which a short-horizon

¹All authors are with the Robotics Innovation Center at the German Research Center for Artificial Intelligence (DFKI), Bremen, Germany

²FS, FW, SV, and FK are also with the University Bremen, Germany

Corresponding author's email: franek.stark@dfki.de

WB model is followed by a reduced SRB model.

$$\begin{aligned}
\min_{\mathcal{X}, \mathcal{U}} \quad & \sum_{k=0}^{N_{wb}-1} L_{wb}(\mathbf{x}_{wb,k}, \mathbf{u}_{wb,k}) + L_t(\mathbf{x}_{wb, N_{wb}}) \\
& + \sum_{k=0}^{N_{srb}-1} L_{srb}(\mathbf{x}_{srb,k}, \mathbf{u}_{srb,k}) + L_f(\mathbf{x}_{srb, N_{srb}}) \\
\text{s.t.} \quad & \mathbf{x}_{wb,0} = \bar{\mathbf{x}}, \\
& \mathbf{x}_{wb,k+1} = f_{wb}(\mathbf{x}_{wb,k}, \mathbf{u}_{wb,k}) \quad k = 0 \dots N_{wb} - 1, \\
& 0 \geq h_{wb}(\mathbf{x}_{wb,k}, \mathbf{u}_{wb,k}) \quad k = 0 \dots N_{wb} - 1, \\
& 0 \geq h_t(\mathbf{x}_{wb, N_{wb}}), \\
& \mathbf{x}_{srb,0} = f_t(\mathbf{x}_{wb, N_{wb}}), \\
& \mathbf{x}_{srb,k+1} = f_{srb}(\mathbf{x}_{srb,k}, \mathbf{u}_{srb,k}) \quad k = 0 \dots N_{srb} - 1, \\
& 0 \geq h_{srb}(\mathbf{x}_{srb,k}, \mathbf{u}_{srb,k}) \quad k = 0 \dots N_{srb} - 1, \\
& 0 \geq h_f(\mathbf{x}_{srb, N_{srb}}).
\end{aligned} \tag{1}$$

Here, $\mathcal{X}$ and $\mathcal{U}$ denote the set of all states and inputs over both phases. The cost function is composed of the stage costs of the WB and SRB phases $L_{wb}(\cdot)$ and $L_{srb}(\cdot)$, a transition cost at the phase boundary $L_t(\cdot)$, and a terminal cost at the end of the prediction horizon $L_f(\cdot)$. The first phase starts from the current state estimate $\bar{\mathbf{x}}$ and is subject to the WB dynamics $f_{wb}(\cdot)$ and constraints $h_{wb}(\cdot)$. At the switching point, the transition map $f_t(\cdot)$ initializes the SRB state from the terminal WB state, while additional transition constraints are imposed through $h_t(\cdot)$. The second phase then evolves according to the SRB dynamics $f_{srb}(\cdot)$, constraints $h_{srb}(\cdot)$ and the terminal constraint set $h_f(\cdot)$. All terms are detailed in the following.

A. Whole Body phase

The WB-phase considers the full bipedal dynamics. The state is parameterized as $\mathbf{x}_{wb,k} = [\delta \mathbf{q}_k^\top, \mathbf{v}_k^\top]^\top$, where $\delta \mathbf{q}_k \in \mathbb{R}^{n_j+6}$ denotes a local coordinate representation of the generalized positions around the initial generalized position $\mathbf{q}_0 \in \mathbb{R}^{n_j+7}$, with n_j the number of actuated joints, the floating-base orientation expressed in the tangent space of $SO(3)$, and $\mathbf{v}_k \in \mathbb{R}^{n_j+6}$ the generalized joint velocities. The full configuration for prediction step k can be recovered as $\mathbf{q}_k = \mathbf{q}_0 \oplus \delta \mathbf{q}_k$, where $\oplus$ denotes the manifold update. The control input is defined as $\mathbf{u}_{wb,k} = [\boldsymbol{\tau}_k^\top, \mathbf{w}_{wb,l,k}^\top, \mathbf{w}_{wb,r,k}^\top]^\top \in \mathbb{R}^{n_j+12}$, where $\boldsymbol{\tau} \in \mathbb{R}^{n_j}$ denotes the vector of actuated joint torques and $\mathbf{w}_{wb,i,k} = [\mathbf{f}_{wb,i,k}^\top, \mathbf{m}_{wb,i,k}^\top]^\top \in \mathbb{R}^6$ the contact wrenches for the left and right foot ($i \in \{l, r\}$), composed of the contact force $\mathbf{f}_{wb,i,k} \in \mathbb{R}^3$ and contact moment $\mathbf{m}_{wb,i,k} \in \mathbb{R}^3$. The system dynamics are discretized using forward Euler, yielding

$$f_{wb}(\mathbf{x}_{wb,k}, \mathbf{u}_{wb,k}) = \begin{bmatrix} \delta \mathbf{q}_{k+1} \\ \mathbf{v}_{k+1} \end{bmatrix} = \begin{bmatrix} \delta \mathbf{q}_k + \Delta t_{wb} \mathbf{v}_k \\ \mathbf{v}_k + \Delta t_{wb} \mathbf{a}_k \end{bmatrix} \tag{2}$$

where $\mathbf{a}_k$ denotes the generalized acceleration and Δt_{wb} the discretization step. To reduce computational complexity,

dynamics terms are fixed around the initial $(\mathbf{q}_0, \mathbf{v}_0)$ such that

$$\mathbf{a}_k = \mathbf{M}(\mathbf{q}_0)^{-1} \left(\mathbf{B} \boldsymbol{\tau}_k - \mathbf{n}(\mathbf{q}_0, \mathbf{v}_0) + \sum_{i \in \{l, r\}} \mathbf{J}_{c_i}(\mathbf{q}_0)^\top \mathbf{w}_{wb,i,k} \right), \tag{3}$$

where $\mathbf{M}(\mathbf{q}_0) \in \mathbb{R}^{(n_j+6) \times (n_j+6)}$ is the generalized inertia matrix, $\mathbf{B} \in \mathbb{R}^{(n_j+6) \times n_j}$ is the actuation selection matrix, $\mathbf{J}_{c_i}(\mathbf{q}_0) \in \mathbb{R}^{6 \times (n_j+6)}$ is the Jacobian of contact i , and $\mathbf{n}(\mathbf{q}_0, \mathbf{v}_0) \in \mathbb{R}^{n_j+6}$ is the combined vector of Coriolis-Centrifugal and gravity forces. The WB phase is augmented with the following constraints, which together define $h_{wb}(\cdot)$:

1) *Wrench cone*: Each $\mathbf{w}_{wb,i,k}$ is constrained to a rectangular wrench cone as derived in [12].

2) *Contact schedule*: The force-contact constraints couple the normal contact force to the prescribed contact mode by enforcing $0 \leq [\mathbf{f}_{wb,i,k}]_z \leq c_{wb,i,k} f_z^{\max}$, where $c_{wb,i,k} \in \{0, 1\}$ is prescribed a priori and indicates whether foot i is planned to be in contact at step k , $[\cdot]_z$ extracts the z -component of the respective force and $f_z^{\max}$ is the maximal z -force.

3) *No slip*: The no-slip constraints enforce zero foot velocity for active contacts, i.e., $c_{wb,i,k} \mathbf{J}_{c_i}(\mathbf{q}_0) \mathbf{v}_k = 0$.

4) *Contact kinematics*: For each foot i , a ground-contact consistency constraint is imposed as

$$0 \leq h_i(\mathbf{q}_k) \leq (1 - c_{wb,i,k}) h_{\max}, \tag{4}$$

$$c_{wb,i,k} \phi_i(\mathbf{q}_k) = 0, \quad c_{wb,i,k} \theta_i(\mathbf{q}_k) = 0, \tag{5}$$

where the functions $h_i(\cdot)$, $\phi_i(\cdot)$, and $\theta_i(\cdot)$ map the generalized joint positions to the foot height, foot roll and foot pitch angle of foot $i \in \{l, r\}$ using forward kinematics. This enforces zero height and planar contact under the assumption of planar terrain, while during swing allowing a maximum height of $h_{\max}$.

5) *Joint limits*: Joint position, velocity and torque limits are enforced as box constraints.

The cost function of the WB-phase is

$$\begin{aligned}
L_{wb}(\cdot) = & \|\delta \mathbf{q}_k - \delta \mathbf{q}_{\text{ref}}\|_{\mathbf{W}_q}^2 + \|\mathbf{v}_k - \mathbf{v}_{\text{ref}}\|_{\mathbf{W}_v}^2 \\
& + \sum_{i \in \{l, r\}} \|h_i(\mathbf{q}_k) - h_{i,k,\text{ref}}\|_{w_h}^2 + \|\phi_i(\mathbf{q}_k)\|_{w_\phi}^2 \\
& + \|\theta_i(\mathbf{q}_k)\|_{w_\theta}^2 + \|\psi_i(\mathbf{q}_k)\|_{w_\psi}^2 \\
& + \|\psi_{i,\text{dis}}(\mathbf{q}_k)\|_{w_{y_b}}^2 + \|\dot{\mathbf{p}}_i(\mathbf{q}_k, \mathbf{v}_k) - \dot{\mathbf{p}}_{i,\text{ref}}\|_{\mathbf{W}_p}^2 \\
& + \|\mathbf{u}_{wb,k}\|_{\mathbf{W}_u}^2,
\end{aligned} \tag{6}$$

where $\delta \mathbf{q}_{\text{ref}}$, $\mathbf{v}_{\text{ref}}$ denote the local configuration increment and generalized velocity references, respectively. Furthermore $h_{i,\text{ref}}$ denotes the height reference of foot i at prediction step k , $\psi_i(\mathbf{q}_k)$ denotes the foot yaw in body coordinates obtained via forward kinematics. Further, $\psi_{i,\text{dis}}(\mathbf{q}_k)$ is the foot-hip distance along the body y -axis of foot i , respectively. Finally, $\dot{\mathbf{p}}_i(\mathbf{q}_k, \mathbf{v}_k)$ denotes the Cartesian foot velocity with $\dot{\mathbf{p}}_{i,\text{ref}}$ being its reference. The diagonal weighting matrices and scalars $\mathbf{W}_q$, $\mathbf{W}_v$, w_h , w_ϕ , w_θ , w_ψ , w_{y_b} , $\mathbf{W}_p$, and $\mathbf{W}_u$ determine the relative importance of the individual objectives. While the dynamic constraint is linear in the decision variables,

the overall WB-phase remains nonlinear due to the kinematic constraints and the forward-kinematics-dependent cost terms, which are evaluated at the predicted $\mathbf{q}_k$.

B. Single Rigid Body phase

In the SRB phase, the robot is approximated as a single rigid body with lumped mass and inertia, while the contact forces and contact locations are optimized. The state is parameterized as $\mathbf{x}_{\text{srb},k} = [\delta\mathbf{q}_{\text{b},k}, \mathbf{v}_{\text{b},k}, \mathbf{p}_{\text{l},k}, \mathbf{p}_{\text{r},k}]^\top$, where $\delta\mathbf{q}_{\text{b},k} \in \mathbb{R}^6$ and $\mathbf{v}_{\text{b},k} \in \mathbb{R}^6$ denote local coordinate representation of the floating base position and the velocity. Hence they match the parameterization of the WB-phase state such that the floating base position $\mathbf{q}_{\text{b},k}$ can be recovered from $\mathbf{q}_0$ as in the WB-phase. Instead of capturing the further joints in the state, the foot positions relative to the body frame are directly included as $\mathbf{p}_{i,k} \in \mathbb{R}^3$. The control input is given by $\mathbf{u}_{\text{srb},k} = [\mathbf{w}_{\text{srb},\text{l},k}, \mathbf{w}_{\text{srb},\text{r},k}, \dot{\mathbf{p}}_{\text{l},k}, \dot{\mathbf{p}}_{\text{r},k}]^\top$, with $\dot{\mathbf{p}}_{i,k}$ as the foot velocities in body and the wrenches analogue to WB-phase as $\mathbf{w}_{\text{srb},i,k} = [\mathbf{f}_{\text{srb},i,k}, \mathbf{m}_{\text{srb},i,k}]^\top \in \mathbb{R}^6$. The state update is

$$f_{\text{srb}}(\cdot) = \begin{bmatrix} \delta\mathbf{q}_{\text{b},k+1} \\ \mathbf{v}_{\text{b},k+1} \\ \mathbf{p}_{\text{l},k+1} \\ \mathbf{p}_{\text{r},k+1} \end{bmatrix} = \begin{bmatrix} \delta\mathbf{q}_{\text{b},k} + \Delta t_{\text{srb}} \mathbf{v}_{\text{b},k} \\ \mathbf{v}_{\text{b},k} + \Delta t_{\text{srb}} \mathbf{a}_{\text{b},k} \\ \mathbf{p}_{\text{l},k} + \Delta t_{\text{srb}} \dot{\mathbf{p}}_{\text{l},k} \\ \mathbf{p}_{\text{r},k} + \Delta t_{\text{srb}} \dot{\mathbf{p}}_{\text{r},k} \end{bmatrix}, \quad (7)$$

with the acceleration being expressed in local robot frame

$$\mathbf{a}_{\text{b},k} = \begin{bmatrix} \frac{1}{m} \left(\sum_{i \in \{1,\text{r}\}} \mathbf{f}_{\text{srb},i,k} \right) - \mathbf{R}_{\text{b}}^{\text{w}}(\mathbf{q}_0) \mathbf{g} \\ \mathbf{I}(\mathbf{q}_0)^{-1} \left(\sum_{i \in \{1,\text{r}\}} [\mathbf{p}_{i,k} - \mathbf{p}_{\text{c}}(\mathbf{q}_0)] \times \mathbf{f}_{\text{srb},i,k} + \mathbf{m}_{\text{srb},i,k} \right) \end{bmatrix} \quad (8)$$

m is the total robot mass, $\mathbf{g}$ is the gravity vector transformed into body frame by $\mathbf{R}_{\text{b}}^{\text{w}}(\mathbf{q}_0)$, $\mathbf{I}(\mathbf{q}_0)$ is the rigid-body inertia used in the SRB phase, and $\mathbf{p}_{\text{c}}(\mathbf{q}_0)$ denotes the center-of-mass (COM) position relative to the body frame, both linearized around $\mathbf{q}_0$. The SRB-phase is subject to the following constraints, which together define $h_{\text{srb}}(\cdot)$ and $h_{\text{f}}(\cdot)$:

1) *Kinematic reachability*: Each foot is restricted to a kinematically admissible region. Since the relative foot positions are states of the SRB model, this is imposed directly as

$$\underline{\mathbf{p}}_i \leq \mathbf{p}_{i,k} \leq \bar{\mathbf{p}}_i, \quad \text{for } i \in \{1,\text{r}\}, \quad (9)$$

where the lateral bounds are chosen side-dependently to preserve the left/right foot assignment.

2) *Wrench cones*: As in the WB phase, all contact wrenches $\mathbf{w}_{\text{srb},i,k}$ are constrained into rectangular wrench cones.

3) *Contact schedule*: The normal contact forces are coupled to the contact schedule as in the WB-phase.

4) *Contact kinematics*: Ground consistency is enforced as

$$0 \leq {}^{\text{w}}\mathbf{p}_{i,k} \leq (1 - c_{\text{srb},i,k}) h_{\text{max}}, \quad i \in \{1,\text{r}\}, \quad (10)$$

where $c_{\text{srb},i,k}$ specifies the contact schedule as in WB-phase and $[\cdot]_{\text{z}}$ extracts the z -component of the respective foot position approximated into world coordinates by

$${}^{\text{w}}\mathbf{p}_{i,k} = [\mathbf{q}_{\text{b},k}]_{\text{trans}} + \mathbf{R}_{\text{b}}^{\text{w}}(\mathbf{q}_0) \mathbf{p}_{i,k}. \quad (11)$$

Here, $[\cdot]_{\text{trans}}$ extracts the translational part of the floating-base generalized positions or velocities.

5) *No slip*: The no-slip condition in the SRB phase can be directly imposed through foot velocities,

$$c_{\text{srb},i,k} ([\mathbf{v}_{\text{b},k}]_{\text{trans}} + \dot{\mathbf{p}}_{i,k}) = \mathbf{0}, \quad i \in \{1,\text{r}\}. \quad (12)$$

The cost function of the SRB-phase is

$$L_{\text{srb}}(\cdot) = \|\delta\mathbf{q}_{\text{b},k} - \delta\mathbf{q}_{\text{b},\text{ref}}\|_{\mathbf{W}_{\text{qb}}}^2 + \|\mathbf{v}_{\text{b},k} - \mathbf{v}_{\text{b},\text{ref}}\|_{\mathbf{W}_{\text{vb}}}^2 + \sum_{i \in \{1,\text{r}\}} \|\mathbf{p}_{i,k} - \mathbf{p}_{i,\text{ref}}\|_{\mathbf{W}_{\text{p}}}^2 + \|[\mathbf{p}_{i,k}]_{\text{z}} - h_{i,k,\text{ref}}\|_{\mathbf{w}_{\text{h}}}^2 + \|\mathbf{u}_{\text{srb},k}\|_{\mathbf{W}_{\text{usrb}}}^2, \quad (13)$$

where $\delta\mathbf{q}_{\text{b},\text{ref}}$, $\mathbf{v}_{\text{b},\text{ref}}$ denote the local floating-base configuration increment and base velocity references, respectively. Furthermore, $\mathbf{p}_{i,\text{ref}}$ is the relative position reference of foot $i \in \{1,\text{r}\}$ in the body frame. The weighting matrices and scalars $\mathbf{W}_{\text{qb}}$, $\mathbf{W}_{\text{vb}}$, $\mathbf{W}_{\text{p}}$, $\mathbf{w}_{\text{h}}$, and $\mathbf{W}_{\text{usrb}}$ determine the relative importance of the individual objectives. The final cost function $L_{\text{f}}(\cdot)$ uses the same terms, except the input dependent.

C. Transition phase

The transition-phase maps the WB-state to the SRB-state as

$$f_{\text{t}}(\cdot) = \begin{bmatrix} \delta\mathbf{q}_{\text{b},0} \\ \mathbf{v}_{\text{b},0} \\ \mathbf{p}_{\text{l},0} \\ \mathbf{p}_{\text{r},0} \end{bmatrix} = \begin{bmatrix} [\delta\mathbf{q}_{N_{\text{wb}}}]_{\text{fb}} \\ [\mathbf{v}_{N_{\text{wb}}}]_{\text{fb}} \\ {}^{\text{b}}p_{\text{l}}(\mathbf{q}_{N_{\text{wb}}}) \\ {}^{\text{b}}p_{\text{r}}(\mathbf{q}_{N_{\text{wb}}}) \end{bmatrix}, \quad (14)$$

where $[\cdot]_{\text{fb}}$ extracts the floating-base components (first 6 elements) of the corresponding whole-body state, and ${}^{\text{b}}p_i(\mathbf{q}_{N_{\text{wb}}})$, $i \in \{1,\text{r}\}$, denotes the foot position in body frame obtained from the forward kinematics at the terminal whole-body configuration. The final state of the WB-phase $\mathbf{x}_{\text{wb},N_{\text{wb}}}$ is constrained and penalized in the transition phase by $h_{\text{t}}(\cdot)$ and $L_{\text{t}}(\cdot)$, which both assemble from the same components as $h_{\text{wb}}(\cdot)$ and $L_{\text{wb}}(\cdot)$, omitting the input dependent terms.

III. RESULTS

The controller is evaluated in a MuJoCo simulation on the torque controlled bipedal robot *HyPer-2*, developed at DFKI RIC. The robot has a height of 110 cm, weighs 22.00 kg and has 6 torque-controlled joints per leg. In this work it is considered without arms and hence has in total 18 DoFs including the floating base. The robot dynamics and kinematics are generated symbolically with Pinocchio [13] and CasADi [14], and used to formulate the MPC in acados [10]. Initially, the MPC is manually tuned to achieve stable walking. Therefore, the ratio $\alpha = \frac{N_{\text{wb}}}{N_{\text{wb}} + N_{\text{srb}}}$ is set to $\alpha = 0.5$ and the total number of prediction steps to $N = N_{\text{wb}} + N_{\text{srb}} = 10$. It was observed that increasing the discretization step size along the prediction horizon is beneficial. Hence the step sizes were set to $\Delta t_{\text{wb}} = 0.02\text{ s}$, $\Delta t_{\text{srb}} = 0.1\text{ s}$. The weights in (6) and (13) are tuned so that the robot maintains a fixed target height and follows a fixed target speed (see Appendix A). A simple walking gait with a per-foot stance time of 0.5 s and a double support time of 0.1 s is used. The swing-foot height is generated as a smooth arc with 3 cm clearance, whose derivative defines vertical velocity targets that are zero




Fig. 2. Average NLP cost, solve time, height and velocity error for bipedal walking over no. of SQP iterations. RT capable settings are marked in green.

at lift-off and touchdown. Stable walking under pure torque control required a 50 Hz to 100 Hz control rate, leaving a maximum solve time of 10 ms. The solver parameters, shown in Appendix B, were manually tuned to ensure stability while also finding solutions quickly. However, as shown in [8], it is not strictly necessary to solve the optimization problem until convergence, since a “good-enough” solution is usually adequate. Consequently, the number of SQP iterations remains one of the main factors influencing the solve time. To investigate this further, we repeat the following experiment several times: The bipedal robot walks forward for 5 s at a target speed of 0.3 meters per second. We increase the number of SQP iterations in each run, starting at 1. The control frequency was set to 100 Hz, regardless of the solve time. Fig. 2 shows the average solve time, non-linear program (NLP) cost, height and velocity error over the maximum number of SQP-iterations, respectively. It shows that solve time is smaller than 10 ms




Fig. 3. Average solve time, height and velocity error for bipedal walking. Here, the ratio $\alpha = N_{wb}/(N_{wb} + N_{srb})$ is varied, while N is kept fixed ($N = 10$). RT capable settings are marked in green.

when using up to 3 SQP iterations and the solve time increases linearly with the number of iterations. The height and velocity errors are largely minimized within the first 3 iterations, while the solver cost does not change much after the first 3 iterations. Hence, 3 was chosen as the iteration limit for the controller and all further experiments. A video demonstrating the simulated *HyPer-2* robot walking under real-time torque control is available online¹. To investigate the relationship

¹Video of the robot walking in simulation: <https://youtu.be/1VuRR2P5hxM>




Fig. 4. Average solve time, height and velocity error for bipedal walking. Here, Δt_{wb} is varied and $\Delta t_{srb} = 5 \Delta t_{wb}$.

between the prediction horizon of whole-body and single rigid body model, the ratio α is varied, while keeping $N = 10$ fixed. The results are shown in Fig. 3. When using $N_{wb} \leq 2$, that is $\alpha \in \{0.1, 0.2\}$, the robot quickly became unstable. When increasing N_{wb} , both, the height and velocity error reduce, where the former has its minimum at $\alpha = 0.6$. At the same time, the solve time increases linearly, such that for $\alpha = 0.6$ the target solve time of 10 ms is not met. For $\alpha > 0.7$, the robot fell. Finally, we vary the discretization step sizes to study the influence of the resulting prediction horizon length, $T = N_{wb} \Delta t_{wb} + N_{srb} \Delta t_{srb}$. Here, Δt_{wb} was varied directly, while Δt_{srb} was computed according to $\Delta t_{srb} = 5 \Delta t_{wb}$. The results are shown in Fig. 4. Unlike in the two previous experiments, the discretization step sizes only have a minor influence on the solve time. The robot is unstable for $\Delta t_{wb} < 15$ ms ($T < 90$ ms). While the velocity error is largely independent of the discretization step size, the height error is smallest at approximately $\Delta t_{wb} = 21$ ms ($T = 126$ ms). Beyond a step size of 21 ms, the robot exhibits increasing knee buckling and approaches a singular configuration, eventually falling for step sizes above $\Delta t_{wb} > 37$ ms ($T > 222$ ms).

IV. DISCUSSION AND FUTURE WORK

The results demonstrate that real-time whole-body MPC with a multi-phase horizon can be achieved within a general-purpose, off-the-shelf nonlinear MPC framework, without requiring a custom solver tailored to legged robotics. This lowers the implementation barrier and allows direct benefit from ongoing advances in numerical optimization. Stable walking is obtained with only a small number of SQP iterations, and the reduced model complexity in the further part of the prediction horizon decreases solve time effectively. However, the analysis is not yet conclusive. Varying the solve time in reality degrades control frequency. The horizon composition via α also changes the effective prediction horizon T due to the different discretization step sizes in each phase. As a result, the instabilities we observe for larger values of α (e.g., $\alpha > 0.7$) may also be attributed to a shortened horizon.

Future work will therefore focus on resolving this through more controlled experiments that independently vary horizon length and model composition. Extending the framework to further model abstractions, and relaxing the prescribed contact schedule are other promising directions. Finally, validation on the real *HyPer-2* system is required to assess stability under realistic conditions.

ACKNOWLEDGEMENTS

This work was done in the projects CoEx (grant number 01IW24008) and ActGPT (grant number 01IW25002) both funded by the Federal Ministry of Research, Technology and Space (BMFTR) and is supported with funds from the federal state of Bremen for setting up the Underactuated Robotics Lab (265/004-08-02-02-30365/2024-102966/2024-740847/2024). Further, this work has been partially supported by the German Federal Ministry of Research, Technology and Space (BMFTR) under the Robotics Institute Germany (RIG). Many thanks to the acados team, for enabling such results.

APPENDIX

A. Cost function weights

Table I summarizes the diagonal entries of all weighting matrices and scalars used in the WB and SRB stage and terminal costs. In the WB phase, the foot z -velocity costs are contact-stage dependent. During swing, moderate costs are used, while in approach (one timestep before contact), larger weights promote a controlled touchdown. The foot-hip distance regularization enforces more natural walking and also works against foot self-collisions. In the SRB phase, foot position weights enforce natural below-hip foot placement.

TABLE I
DIAGONAL COST WEIGHTS USED IN THE WB AND SRB PHASES.

Symbol	Description	Value
<i>WB phase</i>		
$\mathbf{W}_q$	base pos./ori. $(x, y, z, \phi, \theta, \psi)$	$(0, 0, 50, 10, 10, 1)$
	joint position regularization	1×10^{-8}
$\mathbf{W}_v$	base vel. $(v_x, v_y, v_z, \omega_x, \omega_y, \omega_z)$	$(40, 40, 10, 10, 10, 1)$
	joint velocity	0.05
w_h	foot height	10
w_ϕ, w_θ	foot roll/pitch	30
w_ψ	foot yaw	60
w_{yb}	foot-hip distance in body y	10
$\mathbf{W}_p$	foot velocity (swing)	$(1, 1, 1)$
	foot velocity (approach)	$(10, 10, 30)$
$\mathbf{W}_u$	joint torque regularization	1×10^{-13}
	contact wrench regularization	1×10^{-8}
<i>SRB phase</i>		
$\mathbf{W}_{qb}$	base pos./ori. $(x, y, z, \phi, \theta, \psi)$	$(0, 0, 50, 10, 10, 1)$
$\mathbf{W}_{vb}$	base vel. $(v_x, v_y, v_z, \omega_x, \omega_y, \omega_z)$	$(40, 40, 10, 10, 10, 1)$
$\mathbf{W}_p$	foot position in body xy	100
w_h	foot height	10
$\mathbf{W}_p$	foot velocity	2
$\mathbf{W}_{u, \text{srb}}$	contact wrench regularization	1×10^{-8}

B. Solver setup

The acados SQP solver is configured to use a Gauss-Newton Hessian approximation and the sparsity exploiting interior-point method (IPM) solver HPIPM [15] in BALANCE for the quadratic program (QP)-subproblems. Warm-starting is enabled at both the NLP and QP levels, including initialization of the first QP from the previous NLP solution. An adaptive Levenberg-Marquardt regularization is employed to improve numerical robustness, together with projection-based regularization of the QP. A fixed-step globalization strategy is used, avoiding line-search overhead and favoring predictable runtimes. The adaptive QP tolerance strategy

ADAPTIVE_CURRENT_RES_JOINT balances accuracy and speed during each SQP iteration.

REFERENCES

- [1] D. Kim, J. Di Carlo, B. Katz, G. Bledt, and S. Kim, "Highly Dynamic Quadruped Locomotion via Whole-Body Impulse Control and Model Predictive Control," Sep. 2019. [Online]. Available: <http://arxiv.org/abs/1909.06586>
- [2] E. Dantec *et al.*, "Whole-Body Model Predictive Control for Biped Locomotion on a Torque-Controlled Humanoid Robot," in *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*, Nov. 2022, pp. 638–644. [Online]. Available: <https://ieeexplore.ieee.org/document/10000129>
- [3] M. Y. Galliker *et al.*, "Planar Bipedal Locomotion with Nonlinear Model Predictive Control: Online Gait Generation using Whole-Body Dynamics," in *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*, Nov. 2022, pp. 622–629. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/10000132>
- [4] C. Khazoom, S. Hong, M. Chignoli, E. Stanger-Jones, and S. Kim, "Tailoring Solution Accuracy for Fast Whole-Body Model Predictive Control of Legged Robots," *IEEE Robotics and Automation Letters*, vol. 9, no. 12, pp. 11074–11081, Dec. 2024. [Online]. Available: <https://ieeexplore.ieee.org/document/10669220/>
- [5] H. Li, R. J. Frei, and P. M. Wensing, "Model Hierarchy Predictive Control of Robotic Systems," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3373–3380, Apr. 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/9361258>
- [6] H. Li and P. M. Wensing, "Cafe-Mpc: A Cascaded-Fidelity Model Predictive Control Framework With Tuning-Free Whole-Body Control," *IEEE Transactions on Robotics*, vol. 41, pp. 837–856, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10759808/>
- [7] C. Khazoom, "Tailoring Complexity of Model-Based Controllers for Legged Robots," Ph.D. dissertation, Massachusetts Institute of Technology, May 2025. [Online]. Available: <https://hdl.handle.net/1721.1/163429>
- [8] C. Khazoom, S. Heim, D. Gonzalez-Diaz, and S. Kim, "Optimal Scheduling of Models and Horizons for Model Hierarchy Predictive Control," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 9952–9958. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/10160528>
- [9] A. Jordana, S. Kleff, A. Meduri, J. Carpentier, N. Mansard, and L. Righetti, "Structure-Exploiting Sequential Quadratic Programming for Model-Predictive Control," *IEEE Transactions on Robotics*, vol. 41, pp. 4960–4974, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/11112619>
- [10] R. Verschueren *et al.*, "acados—a modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, vol. 14, no. 1, pp. 147–183, Mar. 2022. [Online]. Available: <https://doi.org/10.1007/s12532-021-00208-8>
- [11] J. Frey, K. Baumgärtner, G. Frison, and M. Diehl, "Multi-Phase Optimal Control Problems for Efficient Nonlinear Model Predictive Control With acados," *Optimal Control Applications and Methods*, vol. 46, no. 2, pp. 827–845, 2025. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/oca.3234>
- [12] S. Caron, Q.-C. Pham, and Y. Nakamura, "Stability of surface contacts for humanoid robots: Closed-form formulae of the Contact Wrench Cone for rectangular support areas," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. Seattle, WA, USA: IEEE, May 2015, pp. 5107–5112. [Online]. Available: <http://ieeexplore.ieee.org/document/7139910/>
- [13] J. Carpentier *et al.*, "The pinocchio c++ library – a fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives," in *IEEE International Symposium on System Integrations (SII)*, 2019.
- [14] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.
- [15] G. Frison and M. Diehl, "HPIPM: a high-performance quadratic programming framework for model predictive control," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 6563–6569, 2020. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S2405896320303293>