Online Refinement of Uncertainty Sets for Robust MPC of Quadrupedal Robots Using Convex Cone Programming

Ananya Trivedi[1], Sarvesh Prajapati[1], Mark Zolotas[1], Taşkın Padır[1]

I. Introduction

Utilizing the linearized Single Rigid Body Dynamics (SRBD) model for Model Predictive Control (MPC) of quadrupedal robots has shown great promise [1], [2]. The performance of such simplified models often degrades under significant modeling uncertainties. Domain randomization techniques in reinforcement learning offer a viable solution to the aforementioned challenge [3]. However, these data-driven approaches frequently fall short of providing interpretability in real-world contexts. Additionally, the training process can be both time-intensive and require significant computational resources.

Alternatively, parametric and additive uncertainty in robot dynamics can be addressed using Robust MPC (RMPC) and Stochastic MPC (SMPC). In practice, solutions from RMPC can be conservative if an overapproximation of the unknown uncertainty sets is used [4], [5]. SMPC helps mitigate this challenge by taking into account a distribution of uncertainties and providing a customizable threshold for the maximum probability of constraint violation. Nevertheless, existing SMPC methods are computationally intensive making them unsuitable for real-time MPC applications [6].

In this paper, we separate the uncertainties linked to the SRBD model from those involving surface friction estimation. Our approach is termed Adaptive Robust MPC (AR-MPC), akin to Adaptive MPC [7], as we begin with an initial estimation of the uncertainty sets and aim for convergence to their accurate values using data gathered during robot operation. These evolving uncertainty sets are integrated into an RMPC framework to maintain robustness against worst-case modeling errors. Ensuring robustness is achieved with only a minimal increase in computational demand. In the following section, we detail the steps involved in implementing our method.

II. Method

We begin by defining the following variables: $\theta \in \mathbb{R}^3$ as Euler angles, $p \in \mathbb{R}^3$ as the center of mass (CoM) position, $\omega \in \mathbb{R}^3$ for angular velocity, $\dot{p} \in \mathbb{R}^3$ as CoM velocity, $g \in \mathbb{R}$ for the scalar of gravity acceleration, and $f_i = [f_i^x, f_i^y, f_i^z]^\top \in \mathbb{R}^3$ to denote contact forces for the i^{th} foot. State-space x and control action u are:

$$x \triangleq \left[\theta^\top, p^\top, \omega^\top, \dot{p}^\top, -g\right]^\top, \quad u \triangleq \left[f_1^\top, \ldots, f_4^\top\right]^\top$$

[1]Institute for Experiential Robotics, Northeastern University, Boston, Massachusetts, USA. {trivedi.ana, prajapati.s, m.zolotas,t.padir}@northeastern.edu

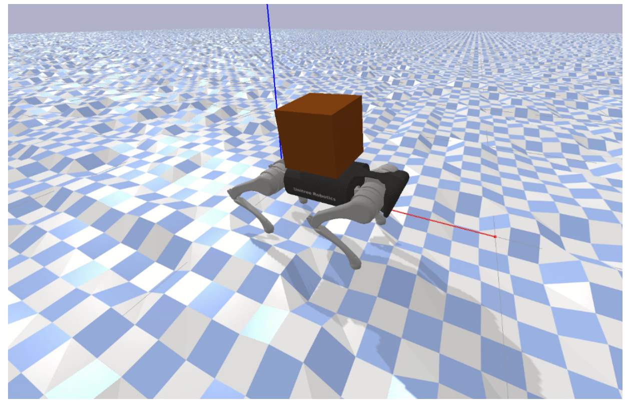

Fig. 1: Adaptive Robust MPC aims to generate optimal ground reaction forces despite unknown payloads, uneven terrains and slippery surfaces.

The linearized SRBD model from [2] can be expressed as:

$$x_{i+1} = A_i x_i + B_i u_i \tag{1}$$

Here, at time instant i, A_i is the state transition matrix and B_i is the control selection matrix. The matrix B_i is a function of the robot's inertia $I_{3\times3}$, mass m, and stance-foot positions r. We now consider bounded yet unknown uncertainties in these parameters, namely $\Delta I_{3\times3}$, Δm, and Δr. Additionally, to account for terrain compliance, we consider unmodeled foot-contact forces Δf. Under these assumptions, the discretized form of the linearized SRBD model derived in [2] can be extended as:

$$\begin{aligned} x_{i+1} &= A_i x_i + (B_i + \Delta B_i)(u_i + \Delta u_i) \\ &= A_i x_i + B_i u_i + E\theta_i^a \end{aligned} \tag{2}$$

$E\theta_i^a$ represents the time-varying cumulative effects of modeling uncertainties. The least conservative domain of $E\theta_i^a$ is referred to as the Feasible Parameter Set (FPS) [8]. While the FPS is not known in practice, it can be approximated based on a history of past robot observations. This approximation is more accurate for a larger look-back horizon p. We approximate the FPS as a polytope, Ω:

$$\Omega = \{E\theta_i : H(E\theta_i) \le h\} \tag{3}$$

Here, H and h are functions of the vector of variables:

$$\{A_0, \ldots, A_{-p}, B_0, \ldots, B_{-p}, x_0, \ldots, x_{-p}, u_0, \ldots, u_{-p}\} \tag{4}$$

The online update expressions for Ω are recursive in nature and can be computed efficiently. Further, the size of Ω reduces with time. This helps alleviate the conservatism associated with standard Tube-Based RMPC algorithms [9].

A. Control Law and Friction Cone Constraints

Control laws that work well in practice for robust and stochastic optimal control formulations consist of a feedback term to account for disturbances and an additive feedforward term [6]. At iteration i of the RMPC, the control action is expressed as:

$$u_i \in \mathbb{R}^{12} = \underbrace{K_i E \theta_i}_{\text{feedback term}} + \underbrace{v_i}_{\text{feedforward term}} \tag{5}$$

Here, K_i is a stabilizing feedback gain obtained by solving the corresponding Discrete Algebraic Ricatti Equation (DARE).

Similar to [4], we consider the bounded uncertainty friction cone-constraints at iteration i of the RMPC as:

$$(C_{\text{nom}} + C_{\text{dist}} z) u_i = (C_{\text{nom}} + C_{\text{dist}} z)(K_i E \theta_i + v_i) \leq 0 \tag{6}$$

$$z \in \mathcal{Z} \triangleq \{z \in \mathbb{R} \mid -\rho \leq z \leq \rho\} \tag{7}$$

The matrix C_{nom} depends on the nominal friction coefficient μ_{nom}, while C_{dist} depends on the expected error in the friction coefficient μ_{dist}. The scalar ρ can modulate the effects of μ_{dist} such that a higher value implies more frictional uncertainty. By introducing additional decision variables λ, [4] converted bounded friction uncertainties into additional linear constraints in their optimization formulation. Applying a similar approach to our control law shown in Eqn. 5, Eqn. 6 can then be reformulated as:

$$C_{\text{nom}}(K_i E \theta_i + v_i) + \lambda_i^\top \alpha \rho \leq 0, \tag{8}$$

$$\lambda_i^\top \beta - C_{\text{dist}}(K_i E \theta_i + v_i) = 0. \tag{9}$$

$$\lambda_i \geq 0 \tag{10}$$

Here, α and β are constant matrices. As is common in RMPC approaches, the effects of θ_i on the constraints in Eqn. 8 can be replaced by their worst-case contributions. This results in the following maximization linear program:

$$\gamma_i = \max_{\theta_i \in \Omega} \; C_{\text{nom}} K_i E \theta_i \tag{11}$$

$$\theta_i^* = \arg\max_{\theta_i \in \Omega} \; C_{\text{nom}} K_i E \theta_i \tag{12}$$

Finally, the constraints in Eqns. 8 and 9 take the form:

$$\gamma_i + C_{\text{nom}} v_i + \lambda_i^\top \alpha \rho \leq 0 \tag{13}$$

$$\lambda_i^\top \beta - C_{\text{dist}}(K_i E \theta_i^* + v_i) = 0 \tag{14}$$

B. Robust MPC and Friction Cone Uncertainty Adaptation

In summary, the RMPC problem involves solving the following Quadratic Programming problem:

$$\begin{aligned} \min_{x[.], v[.], \lambda[.]} \; & \sum_{i=0}^{N-1} \left(||x_i - x_{i,\text{desired}}||_Q^2 + ||v_i||_R^2 \right) \\ \text{s.t.} \quad & \text{Eqn. 1, Eqn. 10, Eqn. 13, Eqn. 14} \end{aligned} \tag{15}$$

Eqn. 15 is in the form of a standard convex-MPC problem over a prediction horizon N. Q and R are tuning weights to penalize trajectory tracking errors and large control actions, respectively. Eqn. 1 is a linear equality constraint on decision variables x and v, while Eqn. 13 and Eqn. 14 are linear inequality and equality constraints on decision variables λ and v. Eqn. 15 can be efficiently solved to global optima using QP solvers, such as [10], [11].

Similar in principle to the online refinement of the FPS, we also propose adjusting the domain of z in Eqn. 7. Any optimal solution of Eqn. 15, ($x^*[.]$, $v^*[.]$, $\lambda^*[.]$), must also satisfy the constraint Eqn. 13. Using this fact, we can formulate the following minimization Linear Program in ρ:

$$\rho^* = \min_{\rho} \rho \tag{16}$$

$$\text{s.t.} \quad \gamma_i + C_{\text{nom}} v_i^* + {\lambda_i^*}^\top \alpha \rho \leq 0, \; \forall i = 0 \ldots N-1$$

Eqn. 16 finds the least conservative value ρ^* that still satisfies the MPC solution. Thus, we slowly converge to the true friction constraint factor $\mu_{nom} + \mu_{dist} \rho^*$.

Algorithm 1 summarizes the proposed AR-MPC.

Algorithm 1 Adaptive Robust MPC

```
Given:
    p: Lookback horizon for Ω estimation
    Q: State tracking weights
    R: Control weights
    Ω: Current estimate of FPS
    N: MPC Horizon
    μ_nom, μ_dist: Nominal, disturbance friction coefficients
    ρ: Current estimate of scalar for friction uncertainty

Step 1:
for i = 0 to N − 1 do
    K_i ← DARE Solution
    γ_i, θ_i ← Eqns. 11 and 12
end for

Step 2:
Ω ← Update using expressions from Eqns. 3 and 4
x*[.], v*[.], λ*[.] ← Solve Eqn. 15
ρ ← Solve Eqn. 16

Repeat Step 1
```

III. Conclusion

In this paper, we proposed an algorithm for robust ground reaction force-based control during quadrupedal locomotion. We analyzed robustness against set-bounded modeling and friction estimation uncertainties. By adapting these sets online, we aim to achieve reduced conservatism in comparison to standard tube-based RMPC algorithms. Our hardware experiments will be designed to assess robustness to these aspects. We intend to load the robot with heavy payloads and evaluate navigation across a grassy terrain while tracking a desired CoM trajectory. This terrain will also feature wooden blocks of arbitrary height and orientation. Currently, we are conducting simulation tests using PyBullet [12] as in Figure 1 and hardware tests on a Unitree Go1-Edu quadruped robot.

References

[1] Y. Yang, T. Zhang, E. Coumans, J. Tan, and B. Boots, "Fast and efficient locomotion via learned gait transitions," in *Conference on robot learning*. PMLR, 2022, pp. 773–783.

[2] J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim, "Dynamic locomotion in the mit cheetah 3 through convex model-predictive control," in *2018 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2018, pp. 1–9.

[3] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.

[4] S. Xu, L. Zhu, H.-T. Zhang, and C. P. Ho, "Robust convex model predictive control for quadruped locomotion under uncertainties," *IEEE Transactions on Robotics*, 2023.

[5] A. Pandala, R. T. Fawcett, U. Rosolia, A. D. Ames, and K. A. Hamed, "Robust predictive control for quadrupedal locomotion: Learning to close the gap between reduced-and full-order models," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6622–6629, 2022.

[6] A. Gazar, "Stochastic predictive control for legged robots," Ph.D. dissertation, Universität Tübingen, 2024.

[7] M. Sombolestan and Q. Nguyen, "Adaptive force-based control of dynamic legged locomotion over uneven terrain," *IEEE Transactions on Robotics*, 2024.

[8] M. Bujarbaruah, X. Zhang, M. Tanaskovic, and F. Borrelli, "Adaptive stochastic mpc under time-varying uncertainty," *IEEE Transactions on Automatic Control*, vol. 66, no. 6, pp. 2840–2845, 2020.

[9] C. Brasseur, A. Sherikov, C. Collette, D. Dimitrov, and P.-B. Wieber, "A robust linear mpc approach to online generation of 3d biped walking motion," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 595–601.

[10] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, "qpoases: A parametric active-set algorithm for quadratic programming," *Mathematical Programming Computation*, vol. 6, pp. 327–363, 2014.

[11] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "Osqp: An operator splitting solver for quadratic programs," *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020.

[12] E. Coumans and Y. Bai, "Pybullet quickstart guide," 2021.