POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

State Estimation for Legged Robots: Fusion of Inertial Sensing and Model-Based Control Information

Laurea Magistrale in Automation and Control Engineering - Ingegneria dell'Automazione

Author: Stefano Taborelli

Advisor: Prof. Marco Faroni

Co-advisor: Chris Mcgreavy

Academic year: 2024-2025

---

1. Introduction

Robotics plays an increasingly important role in industrial environments where human intervention is unsafe, impractical, or inefficient. In such contexts, legged robots represent a particularly promising solution when operation over uneven terrain, stairs or obstacles is required.

At CERN, Legged robots are increasingly employed, for inspections and interventions in inaccessible or unsafe areas, characterized by high levels of ionizing radiation, uneven terrain, confined spaces, absence of GNSS, and potential electromagnetic disturbances. In this context, the considered quadrupeds robots are controlled through a model-based architecture built in-house at CERN around a Nonlinear Model Predictive Controller (NMPC) combined with a Whole-Body Controller (WBC), following the structure proposed in [4] and outlined in Figure 1. It is important to underline that this controller relies at each control cycle on an estimate of the current robot state to compute dynamically consistent motion commands. For this reason accurate and reliable pose estimation is not only a navigation requirement, but a fundamental prerequisite for stable locomotion control.

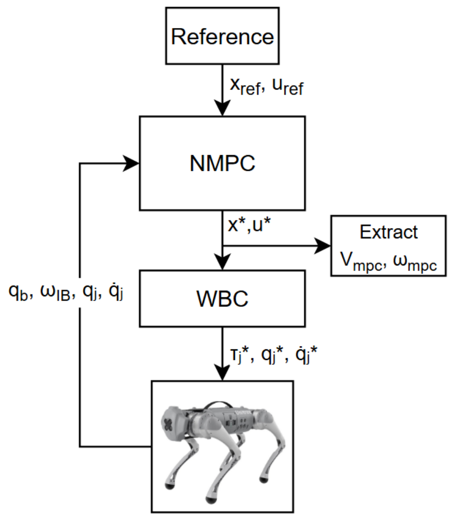

Figure 1: Diagram of the robot control architecture.

Inertial Measurement Units (IMUs) represent the primary building block for motion estimation, but other sources of information are needed to constraint the inevitable drift caused by integration error accumulation in purely inertial navigation.

State-of-the-art solutions for legged robots typically mitigate this issue by fusing IMU data with

leg kinematics, contact constraints, or other external sensors, at the cost of increased modeling and implementation complexity.
The objective of this thesis is to investigates an alternative and exploratory approach that aims to exploit quantities derived from the robot controller itself as an additional information source, and fuse them with inertial data, within a probabilistic state estimation framework.
The work progresses by first developing an IMU-only Extended Kalman Filter (EKF) framework, establishing a baseline architecture, then designing and evaluating in simulation two different dual fusion architectures combining IMU measurements and NMPC outputs. In the first strategy, inertial measurements drive the propagation model while NMPC predictions are used for correction; in the second, the NMPC output is embedded in the process model and inertial quantities are used for correction.
Overall, the study highlights the structural limitations of incorporating controller-derived quantities into the estimation loop underlining the importance of independent measurements to achieve a more robust state estimation.

2. Related work

Inertial navigation has long been a fundamental approach for estimating the motion of autonomous systems in environments where external references are unavailable. The basic principle relies on integrating linear acceleration and angular velocity measurements provided by an Inertial Measurement Unit (IMU) in order to reconstruct velocity, position, and orientation over time. While conceptually straightforward, this process inevitably suffers from error accumulation due to the integration of sensor noise and bias, leading to unbounded drift in the estimated states. For this reason, inertial navigation is rarely used in a purely dead-reckoning configuration. Instead, probabilistic frameworks such as the Extended Kalman Filter (EKF) are commonly employed to fuse inertial data with additional information sources and explicitly model uncertainty and bias dynamics [5].
IMU-only strategies have been extensively investigated, particularly in pedestrian tracking applications as for example in [2]. A widely adopted technique is the use of Zero-Velocity Updates (ZUPT), where stationary phases are detected and exploited as pseudo-measurements to reset velocity offset and indirectly position drift [3]. While effective in scenarios characterized by frequent motion pauses, ZUPT-based approaches are inherently limited in more general robotic applications, including the considered one. These limitations motivate the integration of other constraints to ensure long-term consistency.
In mobile robotics, state estimation typically combines inertial sensing with motion constraints derived from wheel odometry, which provides a continuous and directly observable source of relative displacement. In contrast, legged robots lack continuous rolling contact and operate with intermittent and changing footholds, making the estimation problem more challenging. State-of-the-art approaches for quadruped robots therefore fuse IMU data with leg kinematics and contact information within EKF or error-state formulations as in [1]. When a foot is assumed to be stationary during stance, zero-foot-velocity constraints can be introduced in the world frame, significantly improving observability and limiting drift . Contact-aided inertial estimation has thus become the dominant paradigm in legged robot state estimation.

3. Methodology

The main goal of this thesis is the base pose estimation of a quadruped robot, namely position and orientation of the base frame over time. To this end, after some preliminary analysis on IMU-only navigation and basic correction methods, two dual fusion strategies are investigated. Both strategies are formulated within an Extended Kalman Filter (EKF) framework, where a nonlinear process model is used to propagate the state estimate and an observation model is used to incorporate available information as corrective updates.
The robot NMPC internally exploits a dynamic model of the robot, contact constraints, encoders information and a pose estimate of the base frame to predict the future evolution and therefore the control input over a finite horizon. From the NMPC solution, it is possible to extract predicted base linear and angular velocities, which are computed as part of the optimal motion plan over the prediction horizon.

These quantities are dynamically consistent with the robot model and therefore embed structured information about the expected motion of the base frame. The possibility of directly exploiting such controller-generated velocities within a probabilistic state estimation framework represents the central idea investigated in this thesis. Therefore, the information sources that will be exploited in both cases are: inertial measurements from an IMU (linear acceleration and angular velocity), and NMPC derived quantities (predicted base linear and angular velocities).
In both strategies, the **EKF state** is defined as

$$\mathbf{x}(k) = \begin{bmatrix} \mathbf{p}(k) \\ \mathbf{v}(k) \\ \mathbf{q}(k) \end{bmatrix} \in R^{10}, \tag{1}$$

where $\mathbf{p} \in R^3$ and $\mathbf{v} \in R^3$ denote base position and velocity in the world frame, and $\mathbf{q} \in S^3$ is a unit quaternion describing base orientation.

3.1. IMU-based model EKF with MPC corrections

In the first strategy, inertial measurements are used to propagate the state at high rate, while NMPC outputs are injected as corrective information. The propagation relies on the standard inertial mechanization equations driven by accelerometer and gyroscope measurements expressed in the body frame, $(\mathbf{a}_b(k), \boldsymbol{\omega}_b(k))$.

The discrete-time **process model** is:

$$\begin{aligned} \mathbf{q}(k+1) &= \mathbf{q}(k) \otimes \delta\mathbf{q}(\boldsymbol{\omega}_b(k)\,\Delta t)\,, \\ \mathbf{a}_w(k) &= \mathbf{R}(\mathbf{q}(k))\;\mathbf{a}_b(k), \\ \mathbf{v}(k+1) &= \mathbf{v}(k) + (\mathbf{a}_w(k) + \mathbf{g})\,\Delta t, \\ \mathbf{p}(k+1) &= \mathbf{p}(k) + \mathbf{v}(k)\Delta t + \tfrac{1}{2}\,(\mathbf{a}_w(k) + \mathbf{g})\,\Delta t^2, \end{aligned} \tag{2}$$

where $\mathbf{R}(\mathbf{q})$ is the rotation matrix associated with $\mathbf{q}$, $\mathbf{g} = [0, 0, -9.81]^\top$ is gravity in the world frame, and $\delta\mathbf{q}(\cdot)$ is the incremental quaternion obtained from the Euler-Rodrigues formula coming from the variation of the angular velocity.

The **observation model** exploits the NMPC predicted base linear velocity as a measurement:

$$\mathbf{z}(k) = \mathbf{v}_{\text{mpc}}(k), \qquad \mathbf{h}(\mathbf{x}(k)) = \mathbf{v}(k), \tag{3}$$

which directly corrects the velocity state. Position and orientation are corrected indirectly through the EKF cross-covariance terms, since they are coupled with velocity in the process model through $\mathbf{R}(\mathbf{q}(k))$.
When a measurement is available, the EKF **corrective update** is performed through the standard innovation form:

$$\begin{aligned} \mathbf{y}(k) &= \mathbf{z}(k) - \mathbf{h}\big(\hat{\mathbf{x}}^-(k)\big)\,, && (4) \\ \mathbf{S}(k) &= \mathbf{H}(k)\,\mathbf{P}^-(k)\,\mathbf{H}^\top(k) + \mathbf{R}(k), && (5) \\ \mathbf{K}(k) &= \mathbf{P}^-(k)\,\mathbf{H}^\top(k)\,\mathbf{S}^{-1}(k), && (6) \\ \hat{\mathbf{x}}(k) &= \hat{\mathbf{x}}^-(k) + \mathbf{K}(k)\,\mathbf{y}(k), && (7) \\ \mathbf{P}(k) &= (\mathbf{I} - \mathbf{K}(k)\mathbf{H}(k))\,\mathbf{P}^-(k), && (8) \end{aligned}$$

where $(\hat{\mathbf{x}}^-, \mathbf{P}^-)$ denote the predicted state and covariance, $\mathbf{H}$ is the Jacobian of $\mathbf{h}(\cdot)$, and $\mathbf{R}$ is the measurement noise covariance associated with $\mathbf{v}_{\text{mpc}}$.
In this formulation the quaternion is included directly in the state. As a consequence, the EKF correction applies an additive update to $\mathbf{q}$ in (7), followed by normalization to enforce the unit-norm constraint. This is a first-order approximation (locally replacing S^3 with its tangent space) and is acceptable under the assumption of small orientation corrections; a more rigorous alternative would be an error-state representation, which is not pursued here to keep the formulation compact.

3.2. MPC-model-based EKF with IMU corrections

The second strategy explores the dual configuration: NMPC outputs drive the propagation, while inertial information provides the correction. The state definition remains unchanged, but the **process model** is built from the NMPC angular velocity $\boldsymbol{\omega}_{\text{mpc}}(k)$ and from a linear acceleration $\mathbf{a}_{\text{mpc}}$, obtained through numerical differentiation of the NMPC predicted base velocity.

$$\mathbf{a}_{\text{mpc}}(k) \approx \frac{\mathbf{v}_{\text{mpc}}(k) - \mathbf{v}_{\text{mpc}}(k-1)}{\Delta t}. \tag{9}$$

This trick yields a propagation model formally analogous to the IMU-based case (2) is obtained, at the cost of introducing noise due to numerical differentiation. One difference is that the gravity term does not explicitly appear since $\mathbf{a}_{\text{mpc}}$ is derived from the NMPC velocity that was originally in body frame.
The key difference with respect to the IMU-driven propagation is that the translational dynamics are no longer explicitly expressed as

$\mathbf{R}(\mathbf{q})\mathbf{a}_b$, which remove the coupling between translational and rotational states in the linearized model. As a consequence, in order to effectively constrain both translation and orientation, two corrective quantities are used in the observation model. In this implementation, they are chosen as IMU-derived references: a base linear velocity estimate $\mathbf{v}_{\text{imu}}(k)$ (obtained by integrating gravity-compensated accelerations) and an orientation reference $\mathbf{q}_{\text{imu}}(k)$ (obtained by integrating gyroscope measurements). The **measurement model** is therefore:

$$\mathbf{z}(k) = \begin{bmatrix} \mathbf{v}_{\text{imu}}(k) \\ \mathbf{q}_{\text{imu}}(k) \end{bmatrix}, \qquad \mathbf{h}(\mathbf{x}(k)) = \begin{bmatrix} \mathbf{v}(k) \\ \mathbf{q}(k) \end{bmatrix} \quad (10)$$

The EKF correction step follows the same update equations (7)-(8), with an observation Jacobian consistent with the stacked measurement.

Despite their conceptual simplicity, both strategies exhibit structural limitations. First, the NMPC-derived quantities are not independent from the estimate, since the state used by the EKF is fed back to the controller and influences $(\mathbf{v}_{\text{mpc}}, \boldsymbol{\omega}_{\text{mpc}})$. This estimator-controller coupling may lead to overconfidence or even estimate divergence, and its impact strongly depends on the relative weighting imposed by the tuning of the noise covariance matrices. Second, the MPC-driven propagation relies on numerically differentiating $\mathbf{v}_{\text{mpc}}$ to obtain $\mathbf{a}_{\text{mpc}}$, which amplifies high-frequency noise and may inject artificial dynamics if not properly accounted for. Finally, IMU bias states are neglected here due to the preliminary and simulation-based nature of the study; while bias modeling can be naturally integrated in the IMU-driven propagation, it becomes significantly less straightforward in the MPC-driven formulation, where the process model is not driven by inertial measurements, reducing estimator flexibility for real-world deployment.

4. Experimental Results

The objective of the experimental tests is to evaluate whether the proposed fusion strategies can effectively improve quadruped base pose estimation and to assess the practical impact of the structural limitations discussed in the previous section. In particular, the experiments aim at understanding the standalone behavior of each propagation model, the effect of introducing the corrective term, and the interaction between the estimator and the controller when the estimated state is fed back into the NMPC. All experiments are performed in simulation using MuJoCo, where both the quadruped model and the NMPC controller are implemented (simulation window in Figure 2).

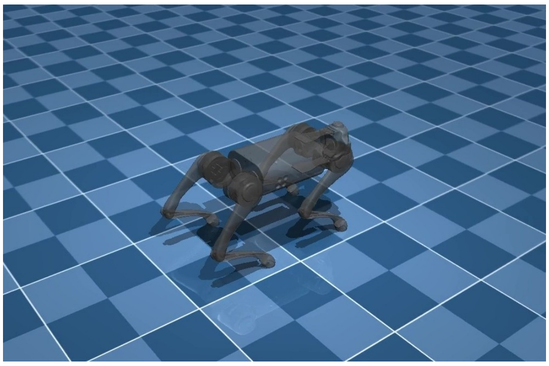

Figure 2: MuJoCo simulation framework.

The robot is commanded to follow a straight-line trajectory at constant velocity. Orientation estimates (quaternion components) are not reported below, as they remain highly accurate in simulation due to the ideal IMU assumption.

4.1. IMU-model-based EKF with MPC corrections

As a first step, the inertial propagation model alone is evaluated without applying any correction and without feeding the estimate back to the controller.

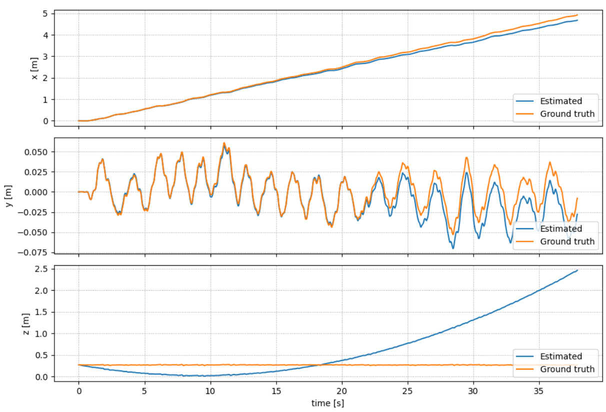

Figure 3: Estimated position and ground truth. Just model propagation

In this configuration, the estimated position closely follows the ground truth along the hor-

izontal axes, confirming the correct implementation of the inertial mechanization (Figure 3). A noticeable drift appears on the vertical (Z) axis, likely related to contact modeling and initialization effects. This behavior is interpreted as equivalent to a bias term and is not considered critical at this stage, since bias modeling is not included in the current formulation.
When the NMPC velocity correction is introduced and the estimated state is fed back into the controller, the overall estimation performance degrades. Figure 4 shows that the position estimate progressively deviates from the ground truth when significant weight is assigned to the NMPC-based measurement.

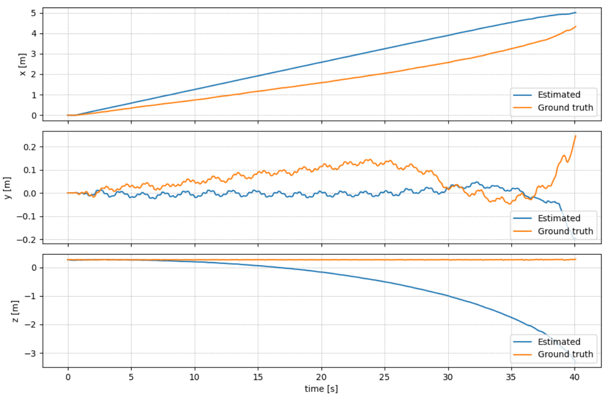

Figure 4: Estimated position and ground truth using IMU-based propagation with NMPC corrections and full state feedback.

This behaviour can be explained by analysing the NMPC predicted velocity. As shown in Figure 5, the NMPC velocity significantly deviates from the ground-truth velocity during the trajectory.

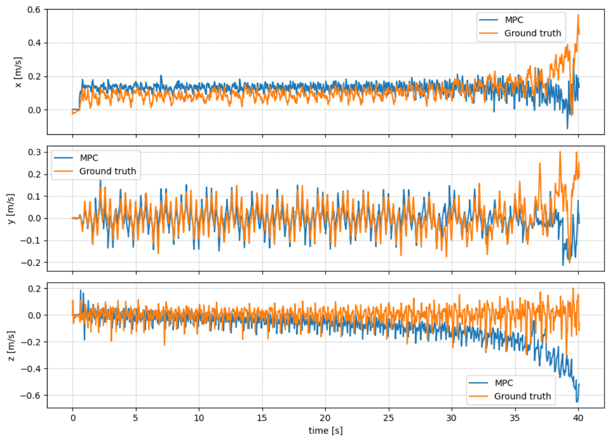

Figure 5: Comparison between NMPC predicted velocity and ground truth.

Since this quantity is used as a corrective measurement, any systematic mismatch directly biases the EKF update and shifts the estimate away from the true state.
When excessive confidence is assigned to the NMPC correction (small R), the estimator may even diverge. This behaviour is consistent with the previously discussed coupling effect: the NMPC relies on the estimated state, and its output is then re-injected into the EKF as a measurement, potentially leading to overconfident and self-reinforcing errors.

4.2. NMPC-model-based EKF with IMU corrections

In this second strategy the NMPC outputs drive the propagation, while IMU-derived quantities provide independent corrections. Here, the estimator–controller coupling affects the filter through the process model, manifesting as a model mismatch rather than influencing the correction term.
The standalone behavior of the NMPC-based propagation already reveals limitations. Even without feeding the estimated state back into the controller, the propagated solution exhibits significant instability, diverging from the ground truth. This indicates that the NMPC-derived system model does not provide a robust basis for state propagation.
When the full estimator is tested and the estimated state is fed back into the controller, the performance does not substantially improve.

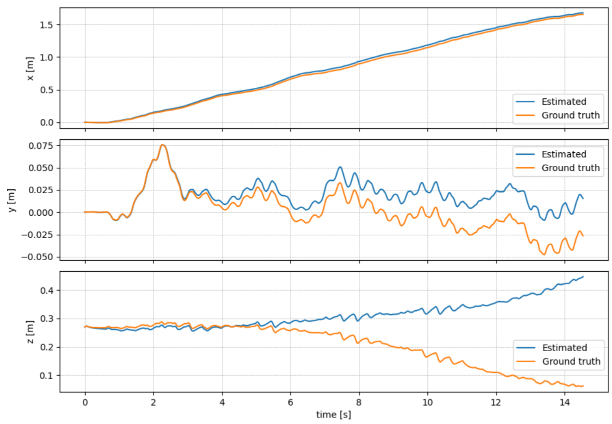

Figure 6: Estimated position and ground truth using NMPC-based propagation with IMU corrections and full state feedback into the controller.

Even if assigning greater weight to the IMU corrections, as observed in Figure 6, stabilizes the estimate and prevents divergence, increasing the influence of the NMPC-derived quantities again leads to performance degradation and instability.

An additional test is conducted by introducing Gaussian noise to IMU and encoder measurements, based on datasheet specifications. As expected, the estimation accuracy degrades, but no total divergence is observed. This suggests that measurement noise alone is not the dominant limiting factor of the proposed architectures. Rather, the main issues appear to be structural, namely the coupling between estimation and control in the first strategy and the weak propagation model (including numerical differentiation effects) in the second strategy.

Overall, the results indicate that the performance of both approaches is strongly constrained by structural limitations that cannot be resolved through tuning alone.

5. Conclusions and Future Developments

This work investigated whether controller-generated quantities could be effectively exploited as additional information within an EKF state estimator, with the objective of supporting inertial pose estimation and mitigating drift.

The results highlight that, although conceptually simple, the use of NMPC-derived quantities as either propagation inputs or corrective measurements introduces structural limitations. In both explored strategies, inertial information remained the dominant contributor to estimation performance. When significant weight was assigned to controller-derived quantities, instability or performance degradation emerged. This behavior is primarily attributed to the intrinsic coupling between estimator and controller: the state estimate is fed back into the NMPC, whose output is then reintroduced into the estimator, weakening the independence assumptions underlying the Kalman filter framework.

Moreover, the use of derived quantities proved to be structurally fragile. State estimators benefit from direct, physically grounded and independent measurements, whereas controller outputs inherently embed correlations that limit their robustness as estimation constraints. The comparative analysis also revealed that performance differences between the two dual strategies were often dominated by covariance tuning rather than by structural advantages, further indicating the limited robustness of both formulations.

The natural step toward a structurally robust estimator is therefore to follow established quadruped literature and integrate contact based constraints within an IMU-driven EKF framework. In practice, this means retaining inertial propagation with explicit bias modeling, while exploiting reliable foot contacts and leg kinematics to introduce zero-foot-velocity constraints in the world frame, significantly improving observability and long-term consistency. Additional improvements could focus on combining multiple sources of correction together(e.g., Zero Velocity Updates events for long pauses and future SLAM-based pose estimates) adapting the correcting strategy based on the availability of the different measurements.

References

[1] Michael Bloesch, Marco Hutter, Mark A. Hoepflinger, Stefan Leutenegger, Christian Gehring, C. David Remy, and Roland Siegwart. State estimation for legged robots – consistent fusion of leg kinematics and imu. In *Robotics: Science and Systems (RSS)*, Sydney, Australia, 2012.

[2] Eric Foxlin. Pedestrian tracking with shoe-mounted inertial sensors. *IEEE Computer Graphics and Applications*, 25(6):38–46, 2005.

[3] Isaac Skog, John-Olof Nilsson, and Peter Händel. Evaluation of zero-velocity detectors for foot-mounted inertial navigation systems. In *2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pages 1–6. IEEE, 2010.

[4] Jean-Pierre Sleiman, Farbod Farshidian, Maria Vittoria Minniti, and Marco Hutter. A unified mpc framework for whole-body dynamic locomotion and manipulation. *IEEE Robotics and Automation Letters*, 6(3):4688–4695, 2021.

[5] David H. Titterton and John L. Weston. *Strapdown Inertial Navigation Technology*. The Institution of Engineering and Technology, London, UK, 2nd edition, 2004.