

Towards Robotic Exploration of the Solar System: A Control Perspective on DLR's Autonomous Lightweight Rover Units

Kristin Lakatos

Vollständiger Abdruck der von der TUM School of Engineering and Design der Technischen Universität München zur Erlangung einer

Doktorin der Ingenieurwissenschaften (Dr.-Ing.)

genehmigten Dissertation.

Vorsitz: Prof. dr.ir. Daniel J. Rixen

Prüfende der Dissertation: 1. Prof. Dr.-Ing. habil. Paul Kotyczka
2. Prof. Dr.-Ing. Alin Olimpiu Albu-Schäffer

Die Dissertation wurde am 15.04.2025 bei der Technischen Universität München eingereicht und durch die TUM School of Engineering and Design am 06.10.2025 angenommen.

Preface

I would like to express my deepest gratitude to the many individuals who have supported me throughout my PhD journey. First and foremost, I would like to thank my supervisor PD Dr.-Ing. habil. Paul Kotyczka for his outstanding supervision, constructive feedback, and personal support during the last years, and Prof. Dr.-Ing. Alin Albu-Schäffer for serving as a second reviewer.

During my time at the Institute of Robotics and Mechatronics of the German Aerospace Center (DLR), I was fortunate to work with an exceptional team. I would like to thank everyone who has helped me during all this time from 2014 to 2025. Your contributions have made a significant impact on my work, and I am grateful for the opportunity to work with such a talented group of individuals.

In particular, I would like to thank my mentor Alexander Dietrich, for all the discussions, professional and fun conversations, and constructive feedback since my Master's thesis back in 2014. Moreover, I would like to thank Armin Wedler, who is the heart and the soul of the mobile robotics team, and who convinced me to stay at DLR when I wanted to leave after the first years. I also want to acknowledge the contributions of my colleagues in the mobile robotics team, whose expertise and camaraderie made my work a pleasure: Martin Schuster, Yunis Scheeler, Ryo Sakagami, Riccardo Giubilato, Moritz Kuhne, Bernhard Rebele, and all the others. Special thanks go to the rover hardware and electronics team, especially Annika Maier, Bertram Willberg, Bernhard Vodermayr and Josef Reill, for their invaluable assistance. My deepest gratitude goes to Florian Schmidt, who always found the time for helping me and sat next to me in the lab so many hours that I can't count them all. Very special thanks are due to Lukas Burkhard for the great work during his Master's thesis, which is an integral part of this thesis, and also for his patience with me during stressful times. Last but not least I want to thank my office colleague Mallikarjuna Vayugundla for his help and encouragement, the good conversations, and all the coffees.

In addition, I would like to thank our project partners from e:fs for the great cooperation, which I am very glad to continue in the future. In particular, I would like to thank Niklas Baldauf for the outstanding collaboration during the DeLeMIS project. The results of the test campaign contributed significantly to this thesis.

Finally, I also extend my gratitude to the many people who have supported me personally, including particularly my family: My parents, Marion and Hans, and my sisters Karen and Lena,

for their unwavering love, encouragement, and support; my mother-in-law Akiko, for her kindness, understanding and help. And finally my husband Dominic, for being an amazing partner and supporting me all along the way.

München, April 2025

Kristin Lakatos

Abstract

The exploration of the solar system poses significant challenges to robotic systems, particularly those operating on planetary surfaces with complex terrain and uncertain conditions. This PhD thesis addresses the control of two prototypes of planetary exploration rovers, by presenting novel control approaches tailored to the unique demands of wheeled mobile manipulators in extraterrestrial environments.

Autonomous exploration of the solar system involves a wide range of tasks, including sampling and analyzing geological formations, deploying scientific instruments, setting up infrastructure such as communication arrays and life support systems, and navigating through uncharted terrain. These tasks require precise control over the rover's motion, localization, and interaction with the environment to ensure successful data collection, instrument deployment, and safe operation.

Whole-body control approaches offer a promising solution to these challenges by integrating multiple degrees of freedom (DOF) of mobile platform and robotic manipulator into a single control framework. By extending the capabilities of the rover's manipulator arm, whole-body control enables the execution of higher-level autonomous tasks. Moreover, increasing the workspace of the manipulator through the DOF of the platform increases the redundancy of the system, allowing for more efficient task execution.

The thesis is divided into two main chapters. The first part focuses on whole-body control concepts for the LRU2 prototype, which is equipped with a robotic arm to perform manipulation tasks. The first topic of this chapter is the development of a platform torque interface, which distributes the actuation torques of the wheels in order to generate a desired wrench acting at the platform center. Building upon this prerequisite, a whole-body impedance controller is presented, which extends classical Cartesian impedance control to a robot arm mounted on a mobile platform. The kinematic redundancy is resolved via quadratic optimization, taking actuator limits into account. The second part further extends this approach to a passivity-based tracking controller for the complete system. This includes the automatic generation of steering angles along with the torque control law for arm and platform propulsion. In this way, both locomotion and manipulation tasks can be solved using the same control framework. A switching approach is proposed which enables the regulation of a static equilibrium at the end of a given time-dependent trajectory. The underlying passivity-based control concept is then evaluated in detail on a holonomic robot arm, particularly addressing convergence properties and the tuning of the controller gains. A comprehensive evaluation of the proposed control methods is conducted by performing experiments with the LRU2.

The second chapter addresses the wheel-ground interaction, a critical aspect of rover operation. According to the focus of this chapter, the proposed approaches are validated using the second rover prototype LRU1, which is equipped with a scientific camera instead of a robotic arm. Two distinct approaches are presented. The first is a parameter-based slip estimation method, which aims at the improvement of the wheel odometry measurements in the context of the rover's self-localization. The method is implemented and validated through data obtained during a Moon-analog mission on Mt. Etna. Second, a modular platform tracking controller is proposed, which combines a nonlinear controller with a linear model predictive controller (MPC) and a deep neural network. The latter is used to learn the unmodeled wheel-ground interaction. The performance of this novel approach is evaluated with data from a test campaign in DLR's planetary exploration lab, driving the rover on a lunar soil simulant.

Through these contributions, this thesis advances the state of the art in control theory for robotic exploration systems, enabling more efficient, stable, and robust navigation and manipulation in complex planetary terrain. By bridging the gap between the development of a theoretical framework and practical implementation, this research supports the future success of robotic planetary exploration of the solar system.

List of Symbols and Abbreviations

List of symbols

Throughout this thesis, all scalar quantities are described by plain letters, whereas matrices and vectors are printed in bold. The total derivative w. r. t. time t of a variable is abbreviated by dots, for example $\dot{\mathbf{q}} = \frac{d}{dt}\mathbf{q}$, $\ddot{\mathbf{q}} = \frac{d^2}{dt^2}\mathbf{q}$. Error variables are indicated by a tilde, and are defined according to $\tilde{(\cdot)} = (\cdot) - (\cdot)_{\text{des}}$.

Various of the following mathematical symbols are introduced in the text with different subscripts, superscripts, and symbols. They are listed without further specification here, their meaning and dimension is specified when the respective variable is introduced in the text. Please note that this list is not exhaustive, but it covers the most significant and recurring quantities throughout the text.

General

\mathcal{A}	set
λ	eigenvalue
t, τ	time
i, j, k	indices for numbering
n, m, l, N	numbers, dimensions
ϵ, c	constants
Δ, δ	difference operator
F	force
M	torque
P	power
E	energy
m	mass
v	translational velocity
ω	angular velocity
V	potential function, Lyapunov function
I	unit matrix
R	rotation matrix
N	null space projector

Robot model

x, y, z	translational degrees of freedom
θ	yaw angle
l	distance
ξ	position of the platform center in a fixed coordinate frame (X, Y)
v_{plf}	rover velocity in local (x, y) frame
φ	steering angles
r	degree of steerability
ϕ	kinematic parameters representing the steering configuration
v_w	translational wheel velocities
f_w	wheel propulsion forces
\mathbf{A}	nonholonomic constraint matrix of the platform
ν	instantaneously feasible platform velocity
Σ	matrix representing the null space of \mathbf{A}
β	lateral driving angle
κ	curvature
ν	translational velocity tangential to the platform path
\mathbf{B}	platform input matrix
η	joint angles of a robot arm
\mathbf{M}	inertia matrix
\mathbf{C}	Coriolis/centrifugal matrix
\mathbf{g}	gravity torques
\mathbf{q}	coordinates of the mobile manipulator
τ	joint torques
\mathbf{A}	nonholonomic constraint matrix
λ	Lagrange multipliers
\mathbf{u}	actuator torques
\mathbf{B}	input matrix
\mathbf{v}	instantaneously feasible velocities
\mathbf{S}	matrix in the null space of \mathbf{A}
\mathbf{B}	input matrix
\mathbf{x}	task space coordinates
\mathbf{J}	task Jacobian matrix
\mathbf{F}	Cartesian wrench

Whole-body control

\mathbf{Q}	weighting matrices
\mathbf{e}	platform wrench error
\mathbf{D}	damping matrix
ζ	damping factors
\mathbf{K}	stiffness matrix
J	cost function
\mathbf{s}	velocity error (sliding variable)
ψ, θ, ϕ	yaw, pitch, roll Euler angles
$\Omega, \mathbf{H}, \mathbf{G}, \mathbf{L}$	gain matrices
σ	sliding variable (arm velocities)

Wheel-ground interaction

p	slip parameters
Ξ	measurement noise
ρ	bogie angle
ζ	6-DOF rover pose
h	output function
H	parameter Jacobian matrix
W	weighting matrix
α	update rate
F, G	linearized system matrices
Ψ	matrix integral
Φ	state transition matrix
C	parameter influence matrix
s	total body slip
u	control input
A, B	linearized state and input matrix
Q, P, R	weighting matrices
Γ	input feature matrix
γ	input feature vector
NN	neural network output
π	platform telemetry data

List of Acronyms

ARCHES	Autonomous Robotic Networks to Help Modern Societies
DLR	Deutsches Zentrum für Luft- und Raumfahrt (German Aerospace Center)
DNN	Deep Neural Network
DGPS	Differential GPS
DOF	degree(s) of freedom
EKF	Extended Kalman Filter
EOM	Equations of Motion
EtherCAT	Ethernet for Control Automation Technology
ESA	European Space Agency
FPGA	Field Programmable Gate Array
GPS	Global Positioning System
GPU	Graphics Processing Unit
GT	Ground Truth
ICR	Instantaneous Center of Rotation
IMU	Inertial Measurement Unit
IPEM	Integrated Prediction Error Minimization
JAXA	Japan Aerospace Exploration Agency
LB-MPC	Learning-Based Model Predictive Control
LQR	Linear Quadratic Regulator
LRU	Lightweight Rover Unit
MER	Mars Exploration Rover
MMX	Martian Moons eXploration
MSL	Mars Science Laboratory
MPC	Model Predictive Control

NASA	National Aeronautics and Space Administration
NLC	Non-Linear Controller
OBC	On-Board Computer
PCI-E	Peripheral Component Interconnect Express
QP	Quadratic Programming
RAFCON	RMC advanced flow control
RMC	Robotik und Mechatronik Zentrum (Robotics and Mechatronics Center)
ROBEX	Robotische Exploration unter Extrembedingungen (Helmholtz alliance for Robotic Exploration of Extreme Environments)
ROS	Robot Operating System
RMS	Root Mean Square
RMSE	Root Mean Square Error
RTC	Realtime Computer
SEA	Serial Elastic Actuator
SLAM	Simultaneous Localization and Mapping
USB	Universal Serial Bus
VO	Visual Odometry
WMM	Wheeled Mobile Manipulator
WMR	Wheeled Mobile Robot
WO	Wheel Odometry

List of Abbreviations

cmd	commanded
con	contact
des	desired
eff	effective
max	maximum
min	minimum
msr	measured
opt.	optimal
p.d.	positive definite
plf	platform
pos.	position
pred	predicted
proj.	projected
ref	reference
rot	rotational
s.p.d.	symmetric positive definite
trans	translational
vel.	velocity
virt	virtual
wgc	wheel-ground contact

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Problem statement	3
1.3	Related works	4
1.3.1	Compliant control for holonomic robots	4
1.3.2	Whole-body control for wheeled mobile manipulators	5
1.3.3	Wheel-ground interaction	8
1.4	Contributions and outline	9
2	System Model	13
2.1	The Lightweight Rover Units: two prototypes of planetary exploration rovers . .	13
2.1.1	Mechanical design	13
2.1.2	Software architecture	16
2.1.3	Modeling assumptions	17
2.1.4	Control architecture	18
2.1.5	Measurement methods	20
2.2	Mathematical modeling	21
2.2.1	Mathematical modeling of a wheeled platform	21
2.2.2	Dynamic equations of a serial robotic arm	26
2.2.3	Combined dynamic equations of a wheeled mobile manipulator	27
2.2.4	Definition of the task space	29
3	Whole-Body Control	31
3.1	Platform torque control	31
3.1.1	Optimization-based distribution of the wheel forces	32
3.1.2	Experiments	33
3.2	Whole-body impedance control	36
3.2.1	Cartesian impedance control	36
3.2.2	Redundancy resolution via quadratic optimization	38
3.2.3	Experiments	39
3.2.4	Résumé	42
3.3	Whole-body passivity-based tracking control	44
3.3.1	Joint-level tracking controller design	45
3.3.2	Task-space tracking control	50

3.3.3	End of trajectory: regulation of a static equilibrium	53
3.3.4	Experiments	54
3.3.5	Résumé	63
3.4	Passivity-based tracking control for holonomic robots	66
3.4.1	Brief introduction of PD+ and SLC	67
3.4.2	Generalized robot tracking controller	68
3.4.3	Discussion of some properties	69
3.4.4	Experiments	71
3.4.5	Résumé	74
4	Wheel-Ground Interaction	77
4.1	Slip modeling and estimation: improving the wheel odometry	78
4.1.1	Parameterized slip model	79
4.1.2	Calibration of the slip model	80
4.1.3	Experimental validation of the slip-compensated wheel odometry	84
4.1.4	Towards slip-compensated control	90
4.1.5	Résumé	91
4.2	Learning-based trajectory tracking control	92
4.2.1	Controller architecture	92
4.2.2	Experiments	97
4.2.3	Résumé	102
5	Discussion	105
5.1	Evaluation of the presented approaches in the context of planetary exploration	105
5.2	Justification of the modeling assumptions	106
5.3	Feasibility in planetary exploration scenarios	108
6	Conclusion	111
Appendix		115
A.1	Theorems and Fundamentals	115
A.1.1	Nonautonomous system	115
A.1.2	Exponential stability	115
A.1.3	Brockett's theorem	116
A.1.4	Min-max theorem	116
A.1.5	Null space and range space	116
A.1.6	Singular value decomposition	117
A.1.7	Computation of the Moore-Penrose pseudoinverse	117
A.2	The ROBEX and ARCHES space-analog missions	119
List of Figures		121
List of Tables		123
Bibliography		125

CHAPTER 1

Introduction

...to boldly go where no one has gone before!

From the opening monologue of the TV series *Star Trek: The Next Generation*,
created by Gene Roddenberry.

These famous opening lines from the iconic television series *Star Trek* outline one of the central themes of the show: exploration, or travel into the unknown in search of new knowledge. It is not for nothing that *Star Trek* is one of the most famous and lasting TV formats ever produced, comprising twelve TV series and thirteen movies as of today. As pointed out in [Ger07], the worldwide fascination for the franchise has its roots in the utopia of humanity willing to surpass itself and change for the better. Thereby, the immersive atmosphere of a positive future is created by using the great unknown of space as a template.

The idea of exploring new worlds has fascinated humanity ever since the first humans gazed into the star-studded night sky. Nowadays, our civilization can look back on a number of manned and unmanned space missions to our nearest celestial bodies. Even if these missions devour a tremendous amount of resources, and there is some inevitable criticism arising from this fact, the exploration of the unknown has inspired generations of humans by some means or other. Only this deep and profoundly human fascination for the miracles waiting beyond our horizon can explain the efforts of so many individuals and organizations dedicated to space exploration, considering the lasting social and environmental problems that humanity is facing here on earth. Apart from the few individuals who were selected to undertake the challenges and risks of manned spaceflight, machines have been an indispensable aspect of planetary exploration from the very beginning. Looking into the future, robots are expected to play an increasingly important role in space exploration due to the rapid technical evolution in the last decades.

1.1 Motivation

Recently, the European Space Agency (ESA) presented a new roadmap for space exploration until and beyond the year 2040. It states that “At the core of the strategy is a bold vision to

establish continuous, sustainable, and responsible human and robotic exploration of the Solar System by providing unique contributions and benefiting society. [...] More specifically, surface autonomous mobility with sample acquisition and analysis will significantly enhance the in-situ science, prospecting for in-situ resource utilization and sample return selection. It will also free-up a predicted scarce crew time by fostering collaborative robotics, supported by artificial intelligence.” [Eur24, pp. 2, 14]. From this statement, the importance of robots for future space exploration can be deduced. Thereby, different tasks like sample acquisition and infrastructure setup involve advanced robotic skills, both in locomotion and manipulation, in a potentially harsh and mostly unknown environment.

In [Ell16] and [Sid18], an overview of the past robotic space exploration missions is given. Recapitulating the history, it becomes clear that rovers have been utilized for different types of planetary exploration missions, ranging from small asteroids and comets to the moons of the gas giants and terrestrial-type planets such as Mars, not to forget our Moon. Thereby, planetary rovers are particularly useful due to their capability to operate on the surface and thus to cover a wider area of exploration compared to static landers. From an operational point of view, rovers offer a great amount of versatility, allowing for the considerate selection of in-situ targets for scientific analysis. Apart from geological questions, astrobiological investigation has recently become a topic of interest, imposing even greater demands on planetary rover platforms. However, the low amount of robotic autonomy still limits the scientific outcome of robotic exploration, as the excessive need for human interaction limits the productive time periods of the robotic explorers. Already in 2008, Bajracharya *et al.* summarized their review on the Mars missions of the National Aeronautics and Space Administration (NASA) from *Sojourner* to the twin Mars Exploration Rovers (MER) *Spirit* and *Opportunity*: “The vehicles used to explore the Martian surface require a high degree of autonomy to navigate challenging and unknown terrain, investigate targets, and detect scientific events. Increased autonomy will be critical to the success of future missions” [BMH08, p. 44]. Besides highlighting the need for more mature algorithms which enable autonomy on a high level, the statement mentions the challenges which concern the control of rovers: The *challenging and unknown terrain* needs to be navigated, while unknown wheel-ground contact effects can be a critical issue for mobile robots. In general, wheel slippage affects both traction performance and energy consumption, and leads to a gradual deviation of the vehicle from the intended path. This can result in large drift and poor results of the localization and control systems [Yos09]. In the worst case, traction issues in a loose soil environment can lead to lasting immobility of the system, as happened with *Spirit* on Mars in May 2009 [Sid18, p. 225].

Moreover, the rover has to be able to *investigate targets, and detect scientific events*, which involves the operation and manipulation of scientific instruments in an unknown environment. These tasks have to be fulfilled as autonomously as possible, which is *critical to the success of future missions*. In [Tre+09], the operations of the robotic manipulators on the Mars Phoenix Lander and the MERs *Spirit* and *Opportunity* are compared and analyzed. Thereby, the main technology gaps are identified as “single-command (single Sol) autonomous approach from extended distances (10-20 m away) and instrument placement on a designated science target(s) within the manipulator work volume” [Tre+09, p. 9], and “autonomous manipulation [...] for successful sample acquisition and in situ science investigations in planetary exploration missions” [Tre+09, p. 9].

Looking at these requirements from a control perspective, one can directly derive the need for algorithms which provide well-defined environment interaction properties, both for manipulation and for locomotion, and a combination of these basic skills. Therefore, the key hypothesis of this thesis is that combining the degrees of freedom (DOF) of the mobile platform and the robotic manipulator with unifying control approaches can bring benefits beyond the separate control of the respective subsystems, and enable improvements in both task performance and autonomy. This approach is referred to as *whole-body control*, and has been a subject of interest in terrestrial

robotics for a while. However, due to the limitations of robotic space exploration, it has not been adopted for planetary rovers so far. Thus, this thesis aims for the development of whole-body control, combined with an outlook on the handling of wheel-ground interaction effects, in the context of (future) planetary exploration missions.

1.2 Problem statement

As pointed out before, this work deals with control strategies which are developed for wheeled planetary exploration rovers. The main feature of this class of robots is a locomotion system for traction and maneuverability in rough terrain, which requires appropriate methods for the control of the chassis. Moreover, a lot of the aforementioned tasks require dedicated manipulation capabilities, so that a robot arm is mounted on the wheeled platform.

Several constraints arise from the fact that the complete systems have to be transported to and work autonomously and robustly on foreign celestial bodies; some of the most important are for sure the severe limits on the system's overall mass, the available power, and the computational resources. A limited communication bandwidth and line-of-sight windows add on top, see e. g. [WGM89; Spi89]. Moreover, the space transportation implies exposure to a high amount of ionizing radiation during the cruise, and potentially also on the extraterrestrial body to explore (e. g. the Moon) [Ell16]. All these factors and limitations have to be taken into account in the development of control algorithms for locomotion and manipulation, which distinguishes the suitable approaches considerably from terrestrial applications.

In particular, the scope of this work is the development of control algorithms for a class of planetary exploration rovers, which comprise a mobile undercarriage equipped with steering wheels in combination with a robotic manipulator. Thereby, the rover is required to perform dedicated scientific tasks in the context of a planetary exploration mission, which involve among others

- Driving on unknown and potentially rough terrain,
- Building a map of the environment and localize itself on it,
- Identifying regions and objects of interest by providing measurements and scans of the environment,
- Reaching an object of interest, interacting with it physically (for example performing in-situ measurements, collecting geologic samples, deploying infrastructure devices, et cetera).

This list is by far not exhaustive, but it gives a perspective on the most important challenges of robotic planetary exploration from a control perspective. These challenges include in particular the interaction with an unknown environment. With respect to manipulation, this causes the need for dedicated contact behavior which prevents stability issues. Concerning locomotion, the identification and compensation of wheel-ground interaction effects (both for the self localization of the rover and for locomotion control) are crucial factors for the success of an exploration mission.

Therefore, the goal of this work is to outline a control architecture, which allows for the use of all available DOF in order to tackle the challenges of a planetary exploration scenario, while respecting the limitations that come with it. This concept of whole-body control has the potential to improve the capabilities of planetary rovers beyond the prevailing independent asynchronous operation of platform and robot arm as it is standard in current space missions. The coordinated use of the platform together with the manipulator (instead of a sequential use of the platform

and the manipulator) increases the workspace of the rover while performing manipulation tasks, and it allows to exploit the increased redundancy of the system for an optimized distribution of the actuation torques. Moreover, a greater amount of interaction forces can be applied to the environment by combining the actuation forces of arm and platform. Consequently, whole-body control strategies offer a viable solution for the deployment of smaller robot arms with potentially fewer DOF, which can subsequently yield cost savings in development and transportation. However, the combined use of platform and arm is a source of novel challenges, as the interaction of the wheels on the ground additionally affects the motion of the end effector. It is therefore crucial to properly address this issue, both for autonomous locomotion and manipulation tasks utilizing whole-body control.

1.3 Related works

The related works can be grouped into the two main topics of this work, namely whole-body control for wheeled mobile manipulators and the estimation and compensation of wheel-ground interaction effects for mobile platforms. Prior to the review of these two topics, a general introduction to the field of compliant control for (holonomic) robots is given, as the whole-body control approaches in this thesis rely on the basic ideas of this field.

1.3.1 Compliant control for holonomic robots

The whole-body control algorithms which are the core contribution of this thesis are based on the concept of *passivity-based control*. This control paradigm is based on the idea of shaping the physical energy of a fully actuated mechanical system while preserving the inherent passivity of the robot as a mechanical system. The derived controllers build on low-level joint torque controllers and feature active compliance of the closed-loop system.

A seminal control approach based on this concept is called *impedance control*, and is introduced in [Hog85]. The basic idea is to assign a prescribed dynamic behavior to the robot while it is interacting with the environment. Thereby, the generalized impedance (i.e., the correlation between generalized velocities or displacements and generalized forces) is specified by a mass-spring-damper system, predestining the control concept for tasks involving physical contact with an operator or the environment. Utilizing the *Operational Space Formulation* introduced by Khatib [Kha87], it is possible to implement an impedance-based controller in Cartesian coordinates of the end effector, which is then denoted as *Cartesian impedance control*. A large amount of literature deals with different varieties and applications of impedance control for robot arms, including [FB97; Alb+03; AOH07; Ott08; DWA11], to name just a few. An overview and comparison of basic concepts and principles, implementation strategies, and practical applications is given by [SYZ19]. Thereby, the control objective is mostly to drive the robot to a desired static equilibrium pose or position, which is referred to as *regulation control*.

In general, regulation controllers are useful for stabilizing the robot in the vicinity of the desired equilibrium. If, however, the task includes larger or repetitive motions of the robot, it is oftentimes desirable to plan the complete motion in advance. This results in the definition of trajectories (comprising the desired velocities and accelerations of the robot together with the time-varying desired positions). In this case, *tracking controllers* can be used to stabilize the robot dynamics to the time-varying trajectory. Classical approaches for robotic trajectory tracking include among others algorithms based on feedback linearization, such as *inverse dynamics* or *computed torque control* schemes, see e.g. [Fre82; SV85]. These techniques shape the inertia of the system in order to create a linear closed-loop error dynamics, at the cost of robustness in the presence of modeling uncertainties and external disturbances. To overcome these drawbacks, the first passivity-based

tracking control algorithms have been developed in the 1980s. The shaping of the potential energy for fully actuated mechanical systems can be traced back to [TA81]. In [OS88], the term of passivity-based control was coined, proving that passivity is the key property underlying the closed-loop stability of these control designs. Important representatives of passivity-based tracking controllers are for example the augmented PD controller (also known as PD+ controller) presented in [PP88] and the Slotine-Li controller introduced in [SL87].

In [Die+21], the performance of passivity-based tracking control (precisely an augmented PD controller) for a holonomic robot arm is compared to control approaches based on feedback-linearization. It is shown that the passivity-based controller is more robust w.r.t. modeling errors and external disturbances by means of a theoretical and experimental analysis. This high robustness is an important aspect for robots in unpredictable and unknown environments. Due to the mentioned considerations, the passivity-based control paradigm suits the requirements of planetary exploration, particularly concerning contact with an unstructured environment. Using the theoretical foundations, approaches for regulation and tracking control of a wheeled mobile robot (WMR) with robotic manipulator are developed in Chapter 3. Thereby, the well-known concepts have to be adjusted to a modular system comprising of a wheeled platform and a robotic arm, with fundamentally different mechanical properties and actuator characteristics. In other words, the arm joints feature a low-level torque control loop with joint torque sensing. In contrast, the platform actuators are not equipped with torque sensors, which inherently limits the performance for motor torque control. Moreover, certain platform elements do require an accurate positioning despite the permanent existence of unknown contact forces, which applies in particular to the steering actuators. Thus, adopting passivity-based control strategies for a planetary exploration rover is not straightforward and requires various modifications. An overview of the existing literature on the topic of whole-body control for wheeled mobile manipulators (WMMs) is given in the following.

1.3.2 Whole-body control for wheeled mobile manipulators

In fact, literature on whole-body control for planetary rovers is sparse, as the common practice in space robotics is to separate locomotion and manipulation tasks. An optimal whole-body trajectory planning algorithm for mobile manipulators in planetary exploration is proposed in [Sto+24], taking into account the nonholonomic platform constraints, but without explicitly proposing a corresponding whole-body control algorithm. However, there is a broad interest in whole-body control methods for WMMs in terrestrial applications. In terms of the requirements, the algorithms from field robotics are most likely to match the requirements of planetary exploration (except the computational limitations). For example, [Hyo+19] lists *terrain adaptability*, *base stability* and *mobility*, and *impact resistance* as key features for “a general purpose mobile platform that can traverse irregular terrains such as agricultural or forest landscapes while carrying manipulators and other tools for heavy-duty operations” [Hyo+19, p. 1]. However, the strict separation of locomotion and manipulation is the prevailing strategy also in this domain (see e.g. [Raj+24] for a recent review of state-of-the-art robotic harvesting procedures).

On the other hand, recent publications of whole-body control algorithms for WMMs oftentimes deal with robotic arms on platforms in a domestic or industrial context. The first class of algorithms for whole-body control of WMMs is purely kinematic, i.e., both platform and arm are controlled on velocity level. The authors of [YY92; Ser98; Bay+02; DOG06] derive kinematic control laws for mobile manipulators featuring differential-drive kinematics. In [DOG10], a kinematic control law for a mobile manipulator with steering wheels is proposed based on a framework of input-output feedback linearization. Thereby, the velocity of the steering wheels becomes an input for task execution, together with the manipulator joint accelerations and the driving accelerations of the

base. Among the kinematic control approaches, the work of H. Seraji [Ser93] stands out, as it extends the concept of redundancy to the complete system (not only the manipulator) for the first time. The result is an *implicit coordination* control scheme, which allows to define the control problem without explicitly defining the trajectory of the wheeled platform. This property makes implicit coordination control schemes particularly suited for reactive approaches [PFC07].

The advantage of the kinematic control laws is that all joints are addressed on velocity level, which allows a unified view on platform and arm despite their different dynamic properties. This method simplifies the implementations of algorithms which ensure a decent tracking performance, however, at the cost of a well-defined and passive interaction behavior. Thus, particularly in regard to scenarios involving the presence of human beings in the workspace of the robot, there is an increased endeavor to provide compliant whole-body control concepts for WMMs. Compliant whole-body controllers are frequently implemented to ensure a well-defined contact behavior, as discussed before for holonomic robots. However, most methods use the joint torque interface for the robot arm, while the platform is controlled on velocity level. Thereby, implicit coordination control schemes have been widely adopted for the force control of WMMs, see e.g. [Ume+99; Kan+01; TX01; ONŽ03; OŽN04]. In [WBK07; Whi+09], a Cartesian impedance controller is proposed for the whole-body coordination of a robotic arm and a two-wheeled nonholonomic, mobile platform, with the focus on the redundancy resolution via decoupled null space control. In [Die15], a multi-task whole-body impedance control framework for the humanoid robot Rollin' Justin is presented, which features a torque-controlled upper body mounted on a velocity-controlled mobile platform. The coupling between the upper-body torque control law and the platform velocity interface uses an admittance simulation, which assigns a virtual mass-damper behavior to the platform. Assuming a high-gain kinematic platform controller which is able to follow the velocity commands sufficiently fast, a desired dynamic behavior can be assigned to the platform, overriding its physical properties. Anyway, a proper treatment of the inertia couplings between platform and upper body is required to guarantee stability [Die+16]. Later, the framework is extended to tracking control [DO20]. In [Isk+19], a whole-body impedance control concept for the assistive robotic system EDAN is introduced, coupling the velocity-controlled wheelchair via an admittance interface to the torque-controlled arm. In combination with shared control, the approach enables the realization of complex tasks like opening a door, passing through it and closing it afterwards. A whole-body control approach based on hierarchical quadratic programming is presented in [Kim+19]. The optimization problem is solved on acceleration level, providing continuous task transitions. Again, the mobile base is coupled via an admittance interface. Wu et al. introduce a control framework for a mobile manipulator, based on a weighted whole-body Cartesian impedance controller, which implements three motion modes, i. e., locomotion, manipulation, and modified loco-manipulation [Wu+21]. Once more, the velocity-controlled mobile platform is coupled via an admittance interface, a decoupled dynamics is assumed for the complete robot. Recently, Aguilera et al. proposed a Linear Quadratic Regulator (LQR) controller which enables a mobile manipulator to autonomously control and maneuver a passive nonholonomic shopping cart. It is assumed that the platform velocity controller has high gains and does not slip or skid [AH23].

The advantage of the aforementioned approaches is that the compliant behavior of the robotic arm is combined with the decent tracking performance of a velocity-controlled platform (at least on the prevalent floor coatings). On the other hand, a velocity-controlled platform can only react to external forces and torques if it is equipped with appropriate sensor capabilities. If this is not the case, the underlying velocity controller compensates for any disturbances, rendering the platform insensitive w. r. t. external forces and torques. One possible way to enable a reactive platform behavior despite the kinematic control is to equip the platform with force/torque sensors. Exemplarily, the work of Leboutet et al. [Leb+19] can be mentioned, where an artificial skin senses external forces and localized them on its kinematic tree. Using this information, a control

method is proposed which provides mobile robots with whole-body compliance capabilities in response to multicontact physical interactions with their environment. Thereby, the controller is formulated as a set of quadratic optimization problems, solved in parallel for each limb involved in the interaction process. Experimental results demonstrate the reliability of the proposed method and its applicability to human-robot interaction scenarios. However, this approach requires not only a costly and highly complex artificial skin, but the signal processing also features high computational complexity.

Intuitively, the most natural way to obtain compliant whole-body behavior is to control the wheel actuators on torque level. Indeed, there exist some works which embed platform torque control in a compliant whole-body control framework. However, these approaches are mainly limited to platforms featuring omnidirectional wheels. The works of Djebbari et al. [DBA12] and Holmberg et al. [HK00] can be mentioned as examples, as they apply mobile impedance techniques to robots on holonomic platforms. Another compliant whole-body controller for a wheeled humanoid robot in sloped terrains is proposed in [SPP13]. It is experimentally shown that the controller performs decent in urban-like environments. In [Don+07], an adaptive trajectory tracking controller for a redundantly-actuated omnidirectional mobile manipulator with uncertainties and disturbances is proposed. Thereby, the manipulator joints as well as all platform actuators (steering and propulsion) are controlled on torque level. While the proposed control law guarantees robustness w. r. t. external disturbances in theory, the performance is validated only in simulation. Ren et al. apply a passivity-based tracking control law to a holonomic, omnidirectional mobile platform, adding an additional disturbance rejection term [RDM18]. The aforementioned approaches have in common that the platform features special wheels (cf. omniwheels, powered castor wheels,...), which are not suitable for planetary rovers, as their applicability is limited to even floors. Another class of systems that can be applied in outdoor environments are differential-drive platforms or skid-steer vehicles, where the steering of the platform is produced by a reverse rotation of opposite wheel pairs. An example for a whole-body torque control concept for a differential-drive WMM is [Xia+18]. Therein, a tracking control law combining fuzzy neural networks (FNN) and an extended Kalman filter (EKF) is presented. The FNN is trained to generate a feed-forward torque, while the EKF is used to sequentially update both the output weights and the centers of the FNN.

From the review of the aforementioned publications, it can be noted that most robotic systems have more actuated DOF than the dimension of the task space. This observation holds to a special degree for mobile manipulators, which combine the DOF of the mobile platform and the robot arm. Thus, adequate strategies for redundancy resolution are a necessary part of the development of whole-body control concepts for these kind of robots. A detailed literature review on redundancy resolution strategies goes beyond the scope of this thesis. However, the most relevant approaches will be mentioned for the sake of completeness.

First, it is possible to use numerical optimization strategies in order to resolve the kinematic redundancy of the manipulator. Many examples have already been mentioned in the preceding review of whole-body control approaches for mobile manipulators. On the other hand, the concept of null-space projection resolves the redundancy by performing additional tasks in the null space of the main task, such that the main task is not disturbed by the null space task. Dietrich et al. give an overview of the hierarchical implementation of several tasks via null space projectors in [DOA15]. A formal proof of stability of the derived projected hierarchical impedance control is given in [DOP18], using conditional stability theory. The concept is extended to tracking control in [DO20; Wu+23]. In general, both resolving the redundancy via null-space projections and relying on numerical optimization strategies results in non-integrability of the solution. Thus, these

approaches do not guarantee *cyclicity* of the resulting trajectories.¹ To overcome this potential drawback, a third class of redundancy resolution strategies is based on augmenting the task in order to create an invertible task Jacobian matrix, as proposed in [Ser89].

1.3.3 Wheel-ground interaction

First, it has to be mentioned that a vast amount of literature is concerned with the wheel torque control of vehicles in rough terrain. The works [Lam+04; ID04b; Yua+06; WA07; Kre+08; SSK17; Hyo+19; BK20] can be mentioned as a small selection of examples. The main objective of these works is the distribution of the platform actuation torques w.r.t. terrain-related criteria such as the minimization of slip, the distribution of vertical contact forces, or the power consumption. In general, the control domain that deals with the distribution of actuation inputs in over-actuated system is called *control allocation*, see [JF13] for an overview. A discussion of topics like the modeling of the wheel-terrain interaction (including the estimation of terrain parameters), and traction control with particular attention to planetary exploration rovers can be found in the monograph [ID04a]. However, none of these works applies the developed algorithms in the context of whole-body control.

On the other hand, the topic of wheel-ground interaction on velocity level is of great interest for actual space missions [Hev+13]. An overview of the state of the art, focusing on current and past space missions, can be found in [GI18] and the references therein. It can be concluded that the problem of slip estimation and compensation in planetary exploration is not yet solved, especially using conventional, model-based methods.

The first question addressed in this thesis is how to estimate the slip on body level and incorporate this information in the computation of the wheel odometry. The solution proposed in Section 4.1 is based on a parameterized slip model as proposed by Seegmiller and Kelly in [SK14], combined with a calibration method that relies only on measurements of the start and end position of a trajectory. The theoretical foundation of the method is laid in [See+13] and [Kel04]. Of course, there exists other literature that deals with the improvement of the wheel odometry. An augmentation of the planar 2D model to full three dimensions is treated e.g. by [Tar+99; TM07] or [SK14]. Information on the calibration of the wheel odometry is presented for example in [OB04]. Helmick et al. describe a concept for sensor fusion and slip compensation in [Hel+05]. A Kalman filter is used to merge data from an Inertial Measurement Unit (IMU) and visual odometry (VO) to obtain a pose estimation, which is then compared to a kinematic rover model to determine the amount of slip. Another work that deals with slip detection and estimation is [Oje+06], where the wheel slippage is estimated from motor current measurements. However, the proposed slip correction method works only in the direction of motion, not laterally, and requires information on the terrain. Note that slip estimation and correction is a topic which is also important in other fields besides planetary robotics. An example from field robotics is [Lin+02], where an observation model for odometric sensor data is used in the vehicle's navigation system. The proposed method relies on measurements obtained from torque sensors between the wheels and their hubs to estimate the slip. Another important field for slip estimation is obviously automotive research – although most methods from this area are not relevant in the context of planetary exploration due to the different system characteristics (in contrast to the deformable soils relevant in planetary exploration, the main source of slip in automotive applications are highly dynamic maneuvers on paved roads).

¹When performing repetitive tasks, it is oftentimes desirable that the robot returns to the same configuration after finishing the operation cycle. This property is called “cyclicity” of motion [Ser89].

The final question is how to compensate slip in order to reduce tracking errors of the vehicle. Therefore, a platform tracking controller which compensates slip on body level by means of a deep neural network (DNN) is proposed in Section 4.2. A concise overview of related works in the area of slip-compensated velocity control of planetary rovers is given in the following. Ishigami et al. investigate path following control for a rover with independently steerable wheels in [INY06] and [INY09]. The proposed feedback controller is independent of the terrain features and soil parameters, and compensates for three types of slip, namely, the vehicle sideslip and longitudinal/lateral wheel slip. Longitudinal and lateral slip on body level is estimated online in [Li+14] based on the rover’s state information. However, the method is only validated in simulation. Adding machine learning techniques to enhance the results of model-based slip estimation has recently become a topic of increased interest. Adding a machine learning component to a model predictive control (MPC) law has been mainly proposed in the context of autonomous racing cars, as for example in [Kab+19] or [RB20]. While these approaches are conceptually iteration-based, a concept which is more suited for exploration scenarios was recently presented in [WFM23]. However, the algorithm presented in the present work excludes Gaussian processes and requires a minimal amount of training data. A general overview of the variants of learning-based MPC algorithms can be found in [Hew+20] and [Mes+22].

1.4 Contributions and outline

While the sections above give an overview of the literature in the field of compliant whole-body control and slip consideration on body level, other relevant works are mentioned throughout the work at several places. The review reveals the research gaps targeted in this work:

First, there is a lack of whole-body control algorithms which are tailored for planetary exploration rovers. This particularly includes the use of wheel torque control in combination with a whole-body control concept. Moreover, the problem of coordinating the steering wheels in a reactive way together with a platform torque controller is not covered yet in the target context.

The compliant whole-body control approaches proposed in Chapter 3 provide a unique framework for compliant whole-body control for planetary exploration rovers in contact with an unstructured environment. Thereby, the platform can be controlled by distributing the wheel forces according to a desired platform wrench. This concept is well-suited for local manipulation tasks, where the inherent system redundancy can be used for compensating tracking errors due to unmodeled wheel-ground interaction. The automatic generation of a desired steering configuration, which is incorporated in the whole-body tracking controller in Section 3.3, enables the platform to react to disturbances at the end effector while maintaining the active compliance properties which can protect the rover hardware in case of obstacles. On the other hand, the platform tracking performance can be enhanced by relying on the low-level platform velocity controller for tasks which involve the traverse of longer areas in difficult terrain. In this way, the presented approach can be used both for locomotion tasks as well as for local manipulation tasks, by switching to the instantaneously convenient platform controller. Moreover, a two-phase approach for the regulation of a static equilibrium at the end of the trajectory is presented, which exploits the model structure in order to create an artificial holonomic subsystem. Finally, the passivity-based tracking control law is formally analyzed by applying it to a holonomic robot arm, discussing relevant properties regarding the convergence of the error states and the physical interpretation of the gains.

The second research objective of this work concerns the treatment of slip with regard to the platform velocity controller, as addressed in Chapter 4. Thereby, the research-driving questions can be formulated as

How can slip be modeled and compensated without prior knowledge of the soil properties, using only available on-board sensors and computationally lightweight algorithms? In which way can the obtained slip estimate be used for crucial tasks like self localization and traction control?

The first contribution in this area is the combination of a parameterized slip model with a calibration procedure from the literature. The resulting strategy for slip estimation is used within the computation of the wheel odometry, as presented in Section 4.1. Thereby, the approaches are adapted such that they meet the requirements of an autonomous robotic mission on a foreign planet. An in-depth experimental evaluation of the performance of the selected method is performed with data obtained at a Moon-analog test site, namely on Mt. Etna. The results motivate the development of a slip compensation based on machine learning, which is presented in Section 4.2. The developed modular approach is based on model-predictive control together with a deep neural network. Thereby, the control architecture is designed to reduce the computational complexity, which distinguishes the approach from current learning-based tracking controllers which are developed e. g. in the context of automotive racing scenarios.

The overall contribution of this thesis is the coalescence of the areas of whole-body control and wheel-ground interaction into a control framework, implemented and tested at two prototypical planetary exploration rovers.

The structure of the thesis including the corresponding publications can be inspected in Figure 1.1. **Chapter 2** introduces the hardware design and software architecture of the LRU prototypes in Section 2.1, followed by the derivation of a mathematical model of the system in Section 2.2. In **Chapter 3**, two whole-body control approaches for wheeled mobile manipulators are presented. In order to enable whole-body control on torque level, a platform torque interface is created by distributing the actuation forces of the single wheels in Section 3.1. Building upon this prerequisite, a whole-body impedance controller is developed in Section 3.2, which realizes a virtual spring-damper behavior at the end effector while resolving the kinematic redundancy via numerical optimization. In Section 3.3, the Cartesian impedance controller is extended to a passivity-based whole-body trajectory tracking controller. The resulting control law implements a compliant tracking behavior in joint space or task space, while the desired steering configuration is computed along with a torque control law for the remaining actuators. A thorough theoretical and practical analysis of the holonomic version of the passivity-based control law is performed in Section 3.4, with special consideration of human-robot collaboration. **Chapter 4** deals with the problems which arise from an unmodeled contact between the wheels and the terrain. Therefore, a parameterized slip model, which can be calibrated without relying on continuous ground-truth position data, is presented in Section 4.1. Based on the findings from a Moon-analog test campaign on Mt. Etna, a modular platform tracking controller is described in Section 4.2, which combines a MPC with a DNN for the estimation of the unknown ground contact dynamics. Finally, **Chapter 5** discusses the results of the previous chapters. In Section 5.1, the control approaches from Chapter 3 and 4 are evaluated in the context of a planetary exploration scenario. Afterwards, Section 5.2 discusses the justification of the modeling assumptions, and Section 5.3 addresses the feasibility of the approaches in a real planetary exploration mission. **Chapter 6** concludes the thesis and gives an outlook on future research topics.

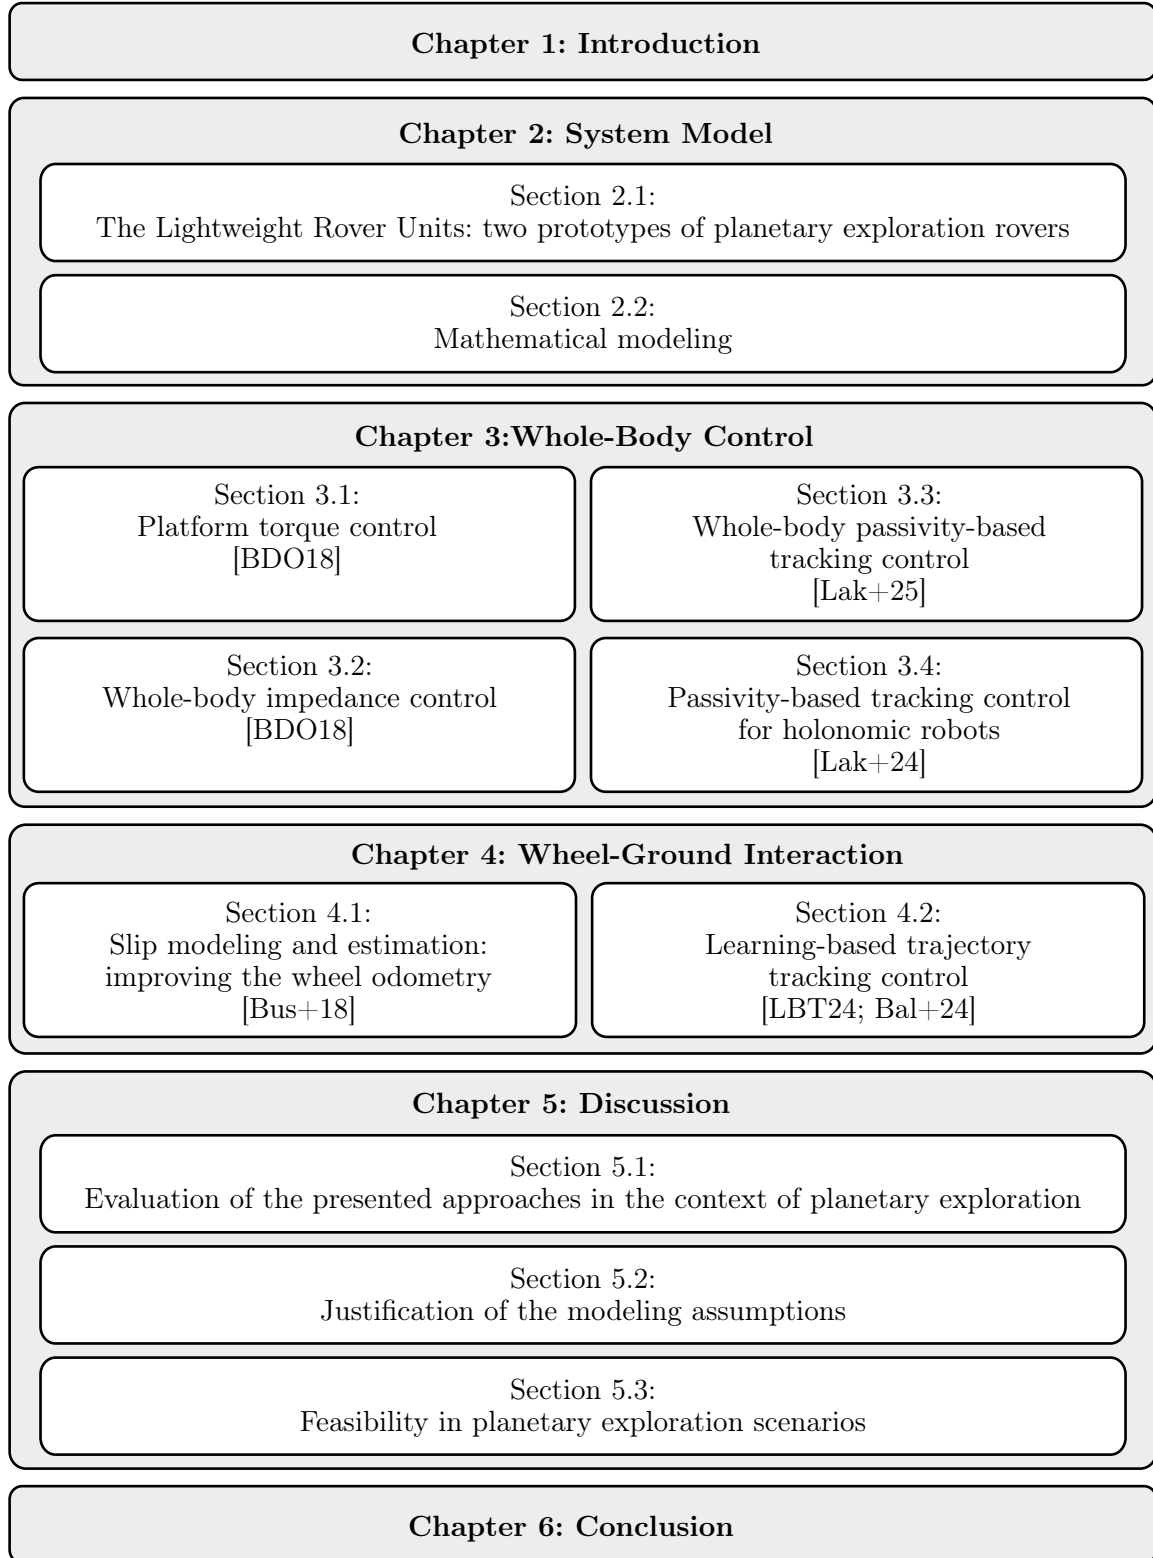


Figure 1.1 Overview of the thesis, including the relation to the publications.

Several parts of this work have evolved from work that has been published before in international journals and conferences. Figure 1.1 shows the structure of the thesis and the relations to the respective publications. The main publications on which this thesis is based are²

- K. Bussmann, A. Dietrich, and C. Ott. “Whole-Body Impedance Control for a Planetary Rover with Robotic Arm: Theory, Control Design, and Experimental Validation”. In: *Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA)*. Brisbane, QLD, May 2018, pp. 910–917
- K. Lakatos, D. Lakatos, P. Kotyczka, and A. Dietrich. “Passivity-Based Tracking Control of a Planetary Rover with Robotic Arm”. In: *IEEE Transactions on Field Robotics* (2025). (Accepted for publication.)
- K. Lakatos, D. Lakatos, X. Wu, P. Kotyczka, and A. Dietrich. “On Passivity-Based Trajectory Tracking for Robotic Manipulators Combining PD+ and Slotine-Li Control”. In: *at - Automatisierungstechnik* 72.12 (Dec. 2024), pp. 1195–1206,

which deal with the topic of whole-body control (see Chapter 3). Apart from these works, the second group of publications deal with the interaction between the wheels and the terrain (cf. Chapter 4), which includes

- K. Bussmann, L. Meyer, F. Steidle, and A. Wedler. “Slip Modeling and Estimation for a Planetary Exploration Rover: Experimental Results from Mt. Etna”. In: *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Madrid, Spain, Oct. 2018, pp. 2449–2456
- K. Lakatos, N. Baldauf, and A. Turnwald. “Towards Learning-Based Trajectory Tracking Control for a Planetary Exploration Rover: Development and Testing”. In: *Proceedings of the 2024 International Conference on Space Robotics (iSpaRo)*. Luxembourg, June 2024, pp. 299–306,

together with the follow-up publication

N. Baldauf, K. Lakatos, A. Meinert, and A. Turnwald. “Towards Learning-Based Trajectory Tracking Control for a Planetary Exploration Rover: Adaptive Model Predictive Control”. In: *Proceedings of the 28th International Conference on System Theory, Control and Computing (ICSTCC)*. Sinaia, Romania, Oct. 2024, pp. 1–6.

²Note that my last name changed from Bussmann to Lakatos in 2022.

CHAPTER 2

System Model


In the following sections, the mechanical design of the rover prototypes is described. Afterwards, a mathematical model of the system is deduced.

2.1 The Lightweight Rover Units: two prototypes of planetary exploration rovers

The main goal of this thesis is to provide control concepts for wheeled mobile robots, suitable for the use in planetary exploration scenarios. In particular, the work deals with the control of two specimen of the DLR's *Lightweight Rover Unit* (LRU). The two prototypical rovers LRU1 and LRU2 are shown in Figure 2.1, while performing typical mission tasks in a Moon-analog environment.

2.1.1 Mechanical design

The first LRU was built at the DLR in 2014. The design was inspired by space concept studies such as the *Mobile Payload Element* from [HMR+13]. A detailed overview of the mechanical design is given in [Wed+15b]. The premise was to build a fast and versatile rover system, which could drive in rough terrain such as Moon and Mars analog sites. A second prototype was build a short time later, such that it was possible to differentiate the hardware and thus the skills in the context of a planetary exploration mission. Thereby, the mechanical design of the chassis was left almost identical. Figure 2.2 shows an excerpt of the platform hardware design. The platform features four wheels with individual propulsion actuators located in the wheel hubs. The axles of the respective steering actuators are located above the wheels, such that the offset is only in z -direction and the wheel can be modeled as centered steering wheel. The two front wheels and the two rear wheels are connected via stiff bogie axles. The rover body containing the inertial measurement unit (IMU) is connected to the bogie axles with serial elastic actuators (SEAs), which provide both passive and (if desired) active elastic suspension. A set of grayscale narrow-angle stereo cameras (Guppy PRO F-125B) is mounted on a pole connected to a pan and a tilt actuator and used for navigation purposes. The third, central camera can be used for adding e. g. a color sensor or a zoom lens, depending on the mission task.



a) LRU1 driving on volcanic terrain.
©DLR/Oprean



b) LRU2 shoveling soil.
©DLR/Oprean

Figure 2.1 The rover prototypes during the ARCHES analog mission on Mt. Etna in summer 2022, performing their typical tasks during a planetary exploration scenario. For details on the ARCHES mission, refer to Appendix A.2.

The wheel design combines the advantages of both wide and narrow wheel types by combining a central rigid ring with a flexible metal surface. Thus, the contact geometry and the rolling resistance of the wheel is adjusted depending on the type of surface. The wheels can be equipped with different kind of elastic spokes, see Figure 2.3. While the wheels of LRU2 are equipped with 12 simple bended blade springs, the wheels of LRU1 feature three laser-sintered titan spokes, which have been numerically optimized with regard to the force distribution while maintaining radial elasticity.

All wheel actuators and the bogie actuators utilize RoboDrive® ILM38 BLDC drive trains with Harmonic Drive® HFUC11-100 gears, with 5 Nm nominal output torque. The pan and tilt actuators are the smaller ILM25 drive units with only 2.4 Nm nominal output torque after the gear. The low-level motor control functionality is provided by commercial Elmo® motor control units, connected to the platform controller via EtherCAT® with 1 kHz cycle time. The steering actuators as well as the pan and tilt actuators are operated using the low-level position controllers. The wheel propulsion actuators as well as the bogie actuators can be used both in velocity control mode and in torque control mode. However, as no torque sensors are integrated in the drive units, scaled measurements of the motor current are used instead of torque measurements, which limits the performance of the low-level torque controllers.

In order to differentiate the skills which the two rover prototypes can contribute to a space mission, different mechanical features and instruments were added to the platform prototypes. Table 2.1 gives an overview of the differences. LRU1 is equipped with a scientific camera head called *ScienceCam*, see Figure 2.1 a). Therefore, it can be used for the exploration and mapping of previously unknown terrain. The ScienceCam was first introduced in [Wed+17], technical details are given in [Sch+20]. It contains an infrared camera, a narrow-angle camera and a pair of

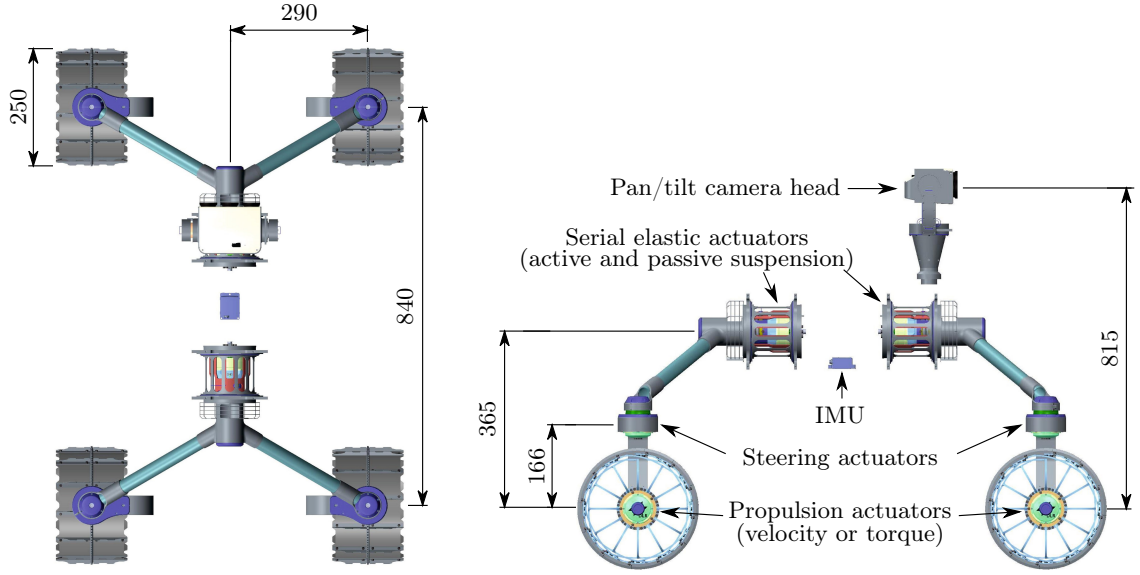


Figure 2.2 Top and side view on the LRU platform. All dimensions are given in mm. Adapted from [Wed+15a].

Feature	LRU1	LRU2
Mechanical design	ScienceCam (spectral and thermal cameras)	Jaco2 robot arm
Wheel design	3 × laser sinter titan spokes	12 × blade springs spokes
Actuated DOF	12	12 + 6

Table 2.1 Distinctive features of the two rover prototypes.

stereo cameras with spectral filter wheels. Thus, regions or objects of interest such as geologically interesting stones can be identified and visually analyzed.

In contrast, LRU2 (see Figure 2.1 b)) carries a 6-DOF robotic manipulator in order to perform manipulation tasks as picking and returning geological samples like stones and sand, or operating additional scientific instruments. The robotic arm is a commercially available Jaco2 by Kinova™. It was built as a robotic assistant, and is mainly supposed to be mounted on an electric wheelchair. According to the data sheet, the total weight of the Jaco2 is 5.3 kg (in the original configuration with the three-fingered hand) and has a maximum reach of 90 cm. The maximum payload is given as 2.5 kg mid-range and 1.5 kg in full extension. In the course of the ROBEX space analog mission (see Appendix A.2), the hand with the three fingers was removed and replaced with the custom *ENVICON docking interface*. The docking interface complies with corresponding adapters at the scientific instruments and payload boxes, and is a key element in the manipulation of different payload systems during a range of planetary exploration mission scenarios. A detailed description of the design and the docking process can be found in [Leh+18]. Later, a revised version of the docking interface featuring also a fully developed electrical interface was used in the ARCHES space analog mission, see [Leh+23] for details.

A detailed description of the LRU2 can be found in [Sch+16] and [Sch+19]. The works highlight the skills of the rover in the context of the national robotics contest *SpaceBotCamp Challenge*, which took place in 2015 in Cologne and focused on autonomous planetary exploration and manipulation.




Figure 2.3 Possible spoke configurations. Left image: Example spoke configurations. Right upper image: Spoke configuration at LRU1 (three laser sintered titan spokes). Right lower image: Spoke configuration at LRU2 (12 blade springs). Images adapted from [Wed+15b].

2.1.2 Software architecture

Due to the need for autonomy in planetary exploration, the computational setup of the LRUs allows to run all relevant software components on board. Therefore, both LRUs share a common computer and software infrastructure, which is depicted in Figure 2.4. Most software components are not required to run in real time. These components are build upon the well-known *Robot Operating System* (ROS) and run on the powerful main on-board computer (OBC). At the moment, the OBC is a 9th generation Core i7 mobile processor (featuring efficient mechanisms for throttling CPU power). The software running on the main computer can be summarized in the categories perception, navigation, search and exploration, and manipulation. A more detailed description of the high-level software architecture can be found in [Sch+19]. Additionally, an Spartan 6 LX75 FPGA extension board is used to perform extensive on-board image processing (dense stereo matching). The extension board as well as the Xsens MTi-10 IMU and cameras are physically connected to the OBC. A separate real-time computer (RTC) is used to run the control algorithms for platform and manipulator. Currently, the RTC is an UP Squared with an Intel® Atom™ processor and a real-time Linux operating system. It is used to run the EtherCAT master software, which sends motor commands to the platform actuators. At LRU2, the real-time computer is additionally used to compute and send motor commands for the Jaco2 arm to a separate BeagleBone Black computer, which is used to interface the motors of the Jaco2 arm via RS485.

The inter-host communication between OBC and RTC is performed using the middleware Links and Nodes (LN) [Sch20], which was developed at DLR to create and manage flexible distributed real-time systems, in particular embedded robotic systems. Moreover, the LN process manager is used to handle all processes and their dependencies which are running together in the course of a so-called *mission*. The sequence of actions performed by (multiple) robotic systems with the goal of fulfilling dedicated tasks is controlled by hierarchical state machines. The sequential control system is implemented in the open source framework RAFCON (“RMC advanced flow control”) [Ger25], which was developed at DLR and introduced in [Bru+16].

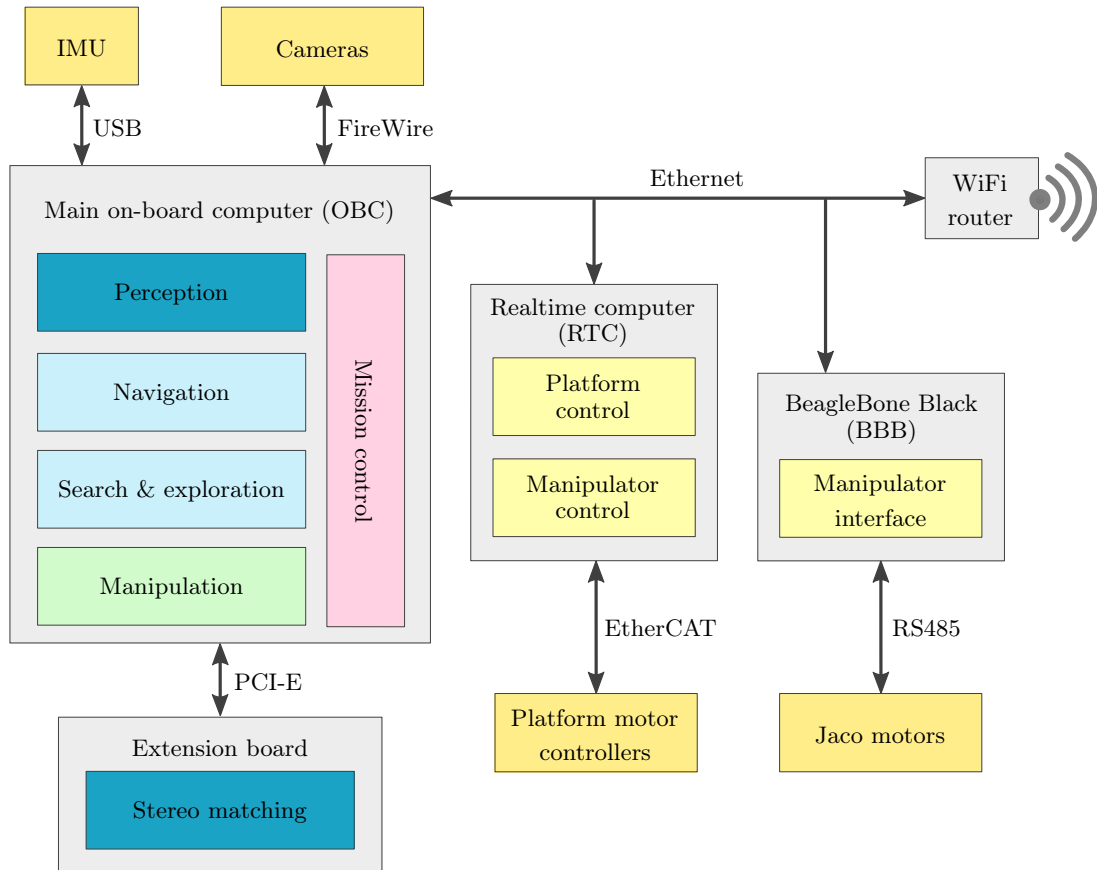


Figure 2.4 Overview of the computer and software structure of the LRU rovers. Adapted from [Sch+19].

2.1.3 Modeling assumptions

Developing control concepts for the LRUs requires mathematical modeling of the rover as a dynamic system. Different kinematic and dynamic models are utilized for the different control approaches. The relatively simple standard controller running at the platform during pure driving operations uses a simple kinematic model of the platform, while the basic arm controllers rely on a standalone dynamic model for arm. Beyond these smaller models, the development of whole-body control approaches requires a dynamic model of the full WMM. The mathematical modeling is described in Section 2.2. Thereby, all equations of the platform and the arm are derived using a set of assumptions. The most important assumptions include

Assumption 1. *All motors are ideal torque sources.*

Assumption 2. *The links and the joints of the robot can be considered as rigid.*

Assumption 3. *The dynamics of the position-controlled steering actuators is sufficiently fast, such that the steering dynamics can be neglected.*

Assumption 4 (for the nominal platform model). *The contact between wheels and ground can be modeled as perfect point contact.*

On the other hand, the following limitations apply in practice (the list is not exhaustive):

- The actuator inputs are limited. The values of the actuator limits are given in Table 2.2.

Limit	Wheel Propulsion	Steering	Arm Joints (1,...,6)
Position in ° (lower/upper)	–	± 150	(–, 50/310, 19/341, –, –, –)
Velocity in °/s	± 360	± 80	$\pm(54, 84, 54, 108, 108, 108)$
Torque in Nm (max. peak)	± 25	–	$\pm(29, 55, 29, 15.7, 15.7, 15.7)$

Table 2.2 LRU1 and LRU2 actuator limits

- Both platform and arm are constructed to be lightweight, thus using flexible structures and materials.
- There are no torque sensors mounted at the wheel hubs, the low-level wheel torque controllers use scaled measurements of the motor current instead.
- During driving, energy is lost, among others due to deformation of the wheel structure (spokes, wheel surface) and deformation of the soil (digging). Wheel-ground contact forces are neither known nor compensated in practice.

The justification of the modeling assumptions w. r. t. the presented control approaches is discussed in Chapter 5, with regard to the experimental results that are presented throughout Chapter 3 and Chapter 4.

2.1.4 Control architecture

A simplified sketch of the control architecture is shown in Figure 2.5, neglecting the feedback of telemetry data to the components on the upper layers. The platform and arm controllers running on the RTC receive commands from the planning layer located at the OBC. From this input, commands for the low-level controllers are computed.

In the basic setup (which is commonly used in advanced mission scenarios), the planning, control and hardware components are divided into platform control and arm control. The platform control pipeline includes the high-level path planning component, which sends platform velocity commands to the platform controller. From this input, actuator commands (i. e., steering angles and wheel velocities) are generated and passed to the low-level motor controllers.¹ The computation of the wheel actuator commands is based on a forward kinematics mapping from the desired body velocity to the respective wheel velocity and steering position commands, which is sometimes referred to as *Ackermann steering* and described in Section 2.2.1.

Remark 1. As mentioned before, the velocity of the steering actuators is limited to 80°/s on the LRUs (cf. Table 2.2). This means that the desired steering values obtained from the platform controller cannot always be achieved immediately. In this case, the corresponding wheel torque (or velocity) commands are set to zero by the low-level controllers until the steering configuration is consistent again. In this way, undesired skidding and slipping is avoided, at the cost of violating the assumption of a perfect steering controller.

In the nominal control setup, the arm joints are kept at a constant position while the platform approaches its goal. Then, a manipulation sequence is performed using one or more of the control modes provided by the Manipulator control module. The available control modes are *Joint Position Control*, *Joint Impedance Control*, *Cartesian Position Control*, *Cartesian Impedance*

¹Additionally, position commands for the pan and tilt actuators are sent in order to adjust the camera position. The bogie actuators can either be used to keep the body aligned to the gravity, or keep a desired constant angle between body and bogie axes. Another possibility is to use the bogie actuators to distribute the vertical wheel forces, as explored in the Master's thesis of Tobias Kindsmüller [Kin17]. However, the control of the pan-tilt and bogie actuators is not in the scope of this work.

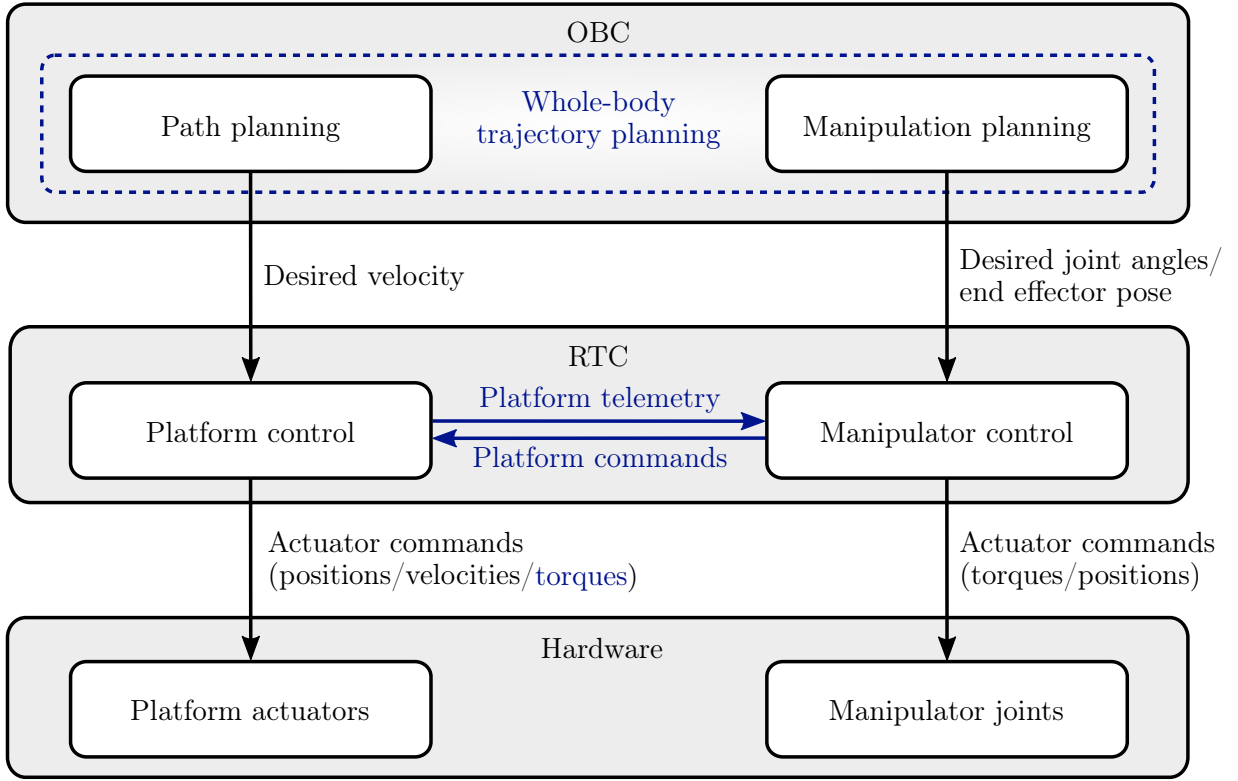


Figure 2.5 A simplified top-down view on the controller structure at the LRU. The planning components are running on the OBC. The RTC receives desired values for the platform velocity and the manipulator joint angles, or the pose of the end effector. Then, the control modules for platform and manipulator compute command values for the respective low-level controllers. Blue lines and variables indicate the additional components required for whole-body control.

Control. Usually, the position control modes are used for free motions of the manipulator, as they provide decent positioning accuracy. On the contrary, the impedance control modes provide a well-defined contact behavior and are thus utilized whenever the manipulator interacts with the environment. The resulting low-level actuator commands are finally passed to the hardware.

This basic setup is used most of the times when performing long and complex missions. Thereby, the separated platform and arm motions are coordinated via state machines. An example of such a complex mission task is the performance of seismic measurements on Mt. Etna, which is detailed from the mission control perspective in [Bru+18].

However, the use of whole-body control algorithms is a promising perspective for planetary exploration missions, especially in regard of more complex manipulation tasks. Thus, the basic control setup can be extended to a whole-body controller by adding components to and communication between the platform controller and the arm controller. First, a torque interface is added to the platform controller, which distributes a desired wrench at the platform center to the wheel torque actuators, using a force equilibrium at the platform center as described in Section 2.2.1. The distribution of the wheel forces is the topic of Section 3.1. Then, using the platform torque interface, torque-based whole-body control algorithms can be integrated into the manipulator control module as shown in Section 3.2 and Section 3.3. Thereby, the platform controller provides the interface to the wheel actuators, and exchanges telemetry and commands with the whole-body controller.

It must be mentioned that the whole-body control approaches developed in the course of this work have been validated experimentally using custom-generated trajectories. In particular, this means that the whole-body controllers together with a whole-body trajectory planner have not been integrated in the high-level mission framework yet, which is indicated by the dashed line surrounding the whole-body trajectory planning component in Figure 2.5. This is a topic of future work for the LRUs and their successor systems.

2.1.5 Measurement methods

Both the implementation and execution of control algorithms and their experimental evaluation require measurements of suitable internal or external sensors. The most important are explained in the following.

Ground-truth position measurements

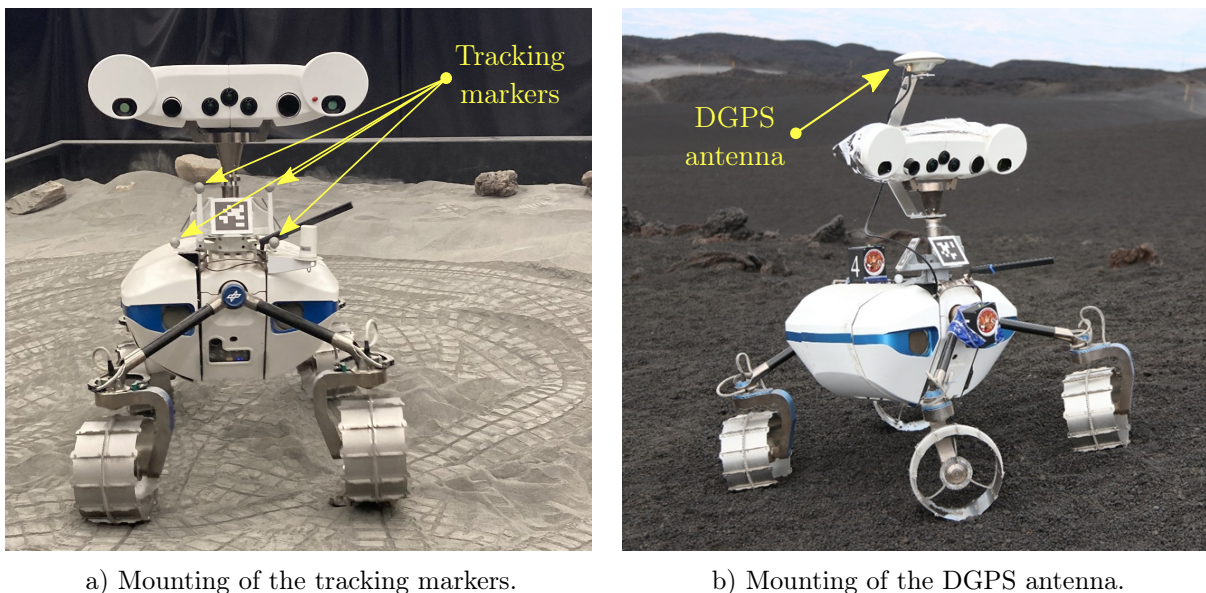


Figure 2.6 Mounting of the tracking markers and the DGPS antenna for ground-truth position measurements at the LRU1.

When performing experiments in the indoor laboratories at the DLR, ground truth position data can be obtained by using a tracking system. In particular, the *Planetary Exploration Lab* (PEL) features an ART-Track 2 tracking system, while in the *mobile lab*, a VICON tracking system is available. Both tracking systems feature a set of cameras mounted on the ceiling, which track artificial markers in the infrared spectrum. The markers consist of a number small retro-reflective balls in a unique configuration, which can be identified and tracked in 6 dimensions (full 3D position and orientation). The LRUs are tracked using rigidly mounted markers at the rover's neck as indicated in Figure 2.6 a). The tracking frequency in the PEL is 60 Hz and covers an area of approx. $10\text{ m} \times 5\text{ m}$, while the VICON system in the mobile lab is capable of up to 250 Hz and covers around $6.5\text{ m} \times 11\text{ m}$.

For outdoor experiments as during the test campaigns on Mt. Etna, a Differential GPS (DGPS) receiver was mounted on each of the LRU robots. Figure 2.6 b) shows the DGPS antenna mounted on LRU1 during the ROBEX mission. As the accuracy of the obtained position measurements depends on several factors such as e.g. the number of satellites in direct line of sight and the

placement of the antenna w.r.t. electromagnetic interference, it varies up to approx. 10 cm. However, as the DGPS position data is only used for evaluating the performance of the wheel odometry computations in Chapter 4 and the self localization of the rover, the accuracy is considered sufficient.² In contrast to the tracking system, DGPS measurements only comprise 3D position measurements (no orientation).

On-board measurements

The joints of the robot arm deliver measurements of their position, velocity and torque at every millisecond. From the position measurements of the arm joints, it is possible to compute the position of the end effector using the forward kinematics (assuming rigid links and joints). The platform actuators send measurements of their position, velocity and motor current in each cycle. It is possible to compute an estimation of the velocity of the platform center, the *body velocity*, from measurements of the single wheel velocities and the steering angles in terms of their kinematic relationship. This measurement of the body velocity by an inverse kinematics mapping is called *wheel odometry*. The kinematic relationships of the platform are elaborated in the following section.

Transforming the wheel odometry to a fixed frame, it is possible to integrate the (inertial) velocity of the platform w.r.t. time, which yields an estimation of the platform position (obtained from measurements of the wheel velocities). For motions in sloped terrain, it is also possible to incorporate measurements of the body angle (IMU and bogie positions) into the computation, and thus obtain a 3D wheel odometry. The implementation of the 3D wheel odometry was part of the Master's Thesis of Lukas Meyer [Mey17], which was performed in the context of the ROBEX space-analog mission and pursued the goal of improving the wheel odometry measurements on Mt. Etna. For details on the 3D wheel odometry and other steps that were undertaken to elaborate the wheel odometry, the reader is referred to the mentioned thesis and the resulting publication [Bus+18], which corresponds substantially to Section 4.1 of this thesis.

The wheel velocity measurements can be fused with IMU measurements and visual odometry using an extended Kalman filter (EKF). The algorithm is introduced in [Sch+12]. The resulting pose estimation w.r.t. a map of the environment is referred to as *self localization* of the platform. It is computed on the OBC.

Finally, the position of the end effector of the mobile manipulator can be obtained by starting at the absolute position of the platform (either ground truth position measurement, self localization or integrated wheel odometry) and extending the kinematic chain with the forward kinematics computation of the robot arm.

2.2 Mathematical modeling

In the following, the mathematical modeling of both the nonholonomic mobile platform and the serial robotic arm are introduced. Then, a combined model of the complete system is deduced. Basic definitions of the rank of a matrix as well as the null space and range space can be found in Appendix A.1.5.

2.2.1 Mathematical modeling of a wheeled platform

A simple model of the mobile platform on a planar surface is described in the following. It is used in the basic controllers running at the LRUs, and it is incorporated in the derivation of the

²A detailed analysis of the pose estimation is performed in [Vay+18], using data from long-range runs of several hundred meters.

equations of motion of the complete WMM. Thereby, perfect wheel-ground contact is assumed, which means that one single contact point per wheel is considered, neglecting the relative motion between wheel and underground (wheel slip) and soil deformation (perfect force transmission). Figure 2.7 shows the used variables and their kinematic relationships.

Kinematic relations of the mobile platform

First, the forward kinematics equations of the platform are derived (for reference see Figure 2.7 a)). The position of the platform center in the inertial frame (X, Y) is given as

$$\boldsymbol{\xi} = \begin{pmatrix} {}^0x \\ {}^0y \\ \theta \end{pmatrix}. \quad (2.1)$$

The platform features $n_w = 4$ independent, centered steering wheels, and $\boldsymbol{\varphi} \in \mathbb{R}^{n_w}$ combines the individual steering angles φ_i for $i = 1, \dots, n_w$. The distances of the i -th steering axis from the platform center in the local (x, y) -frame (also called body frame) are denoted by $l_{x,i}, l_{y,i}$, $i = 1, \dots, n_w$. The rover velocity in the body frame is

$$\mathbf{v}_{\text{plf}} = \begin{pmatrix} v_x \\ v_y \\ \omega \end{pmatrix} \quad (2.2)$$

and can be derived as

$$\mathbf{v}_{\text{plf}} = \mathbf{R}_z(\theta)^T \dot{\boldsymbol{\xi}}, \quad (2.3)$$

with the rotation matrix

$$\mathbf{R}_z(\theta) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (2.4)$$

The translational velocity of the i -th wheel center is given by

$$v_{w,i} = r_w \omega_{w,i}, \quad (2.5)$$


with r_w denoting the radius of the wheel and $\omega_{w,i}$ its rotational velocity. Splitting the wheel velocities $v_{w,i}$ into their local x - and y -components, the respective component can be derived w. r. t. the body velocity as

$$\begin{aligned} v_{x,i} &= v_{w,i} \cos(\varphi_i) = v_x - l_{y,i} \omega, \\ v_{y,i} &= v_{w,i} \sin(\varphi_i) = v_y + l_{x,i} \omega. \end{aligned} \quad (2.6)$$


Using this relation, the steering angles and the wheel velocities can be computed from the body velocity as

$$\varphi_i = \arctan \left(\frac{v_{y,i}}{v_{x,i}} \right), \quad (2.7)$$

$$v_{w,i} = \pm \sqrt{v_{x,i}^2 + v_{y,i}^2}. \quad (2.8)$$



a) Kinematic relations of the mobile platform.



b) Force equilibrium of the mobile platform assuming perfect wheel-ground contact.

Figure 2.7 Platform velocity and torque mappings. The left superscript 0 indicates that a variable is expressed in the inertial frame (X, Y) , while a left superscript b indicates the local (body) frame (x, y) .

Finally, the inverse kinematic relation between the measured wheel velocities v_i and the steering angles φ_i is stacked to

$$\begin{pmatrix} v_{x,1} \\ v_{y,1} \\ v_{x,2} \\ v_{y,2} \\ v_{x,3} \\ v_{y,3} \\ v_{x,4} \\ v_{y,4} \end{pmatrix} = \begin{pmatrix} 1 & 0 & -l_{y,1} \\ 0 & 1 & l_{x,1} \\ 1 & 0 & -l_{y,2} \\ 0 & 1 & l_{x,2} \\ 1 & 0 & -l_{y,3} \\ 0 & 1 & l_{x,3} \\ 1 & 0 & -l_{y,4} \\ 0 & 1 & l_{x,4} \end{pmatrix} \begin{pmatrix} v_x \\ v_y \\ \omega \end{pmatrix}. \quad (2.9)$$

The result (2.9) can then be pseudo-inverted in order to obtain an estimation of the platform velocity \mathbf{v}_{plf} , which is called the *wheel odometry*.

Modeling of a constraint-consistent steering configuration

Assuming perfect wheel ground contact, the motion of each wheel on the planar surface is restricted to rolling, while lateral slipping is not possible. This fact can be modeled by the *Pfaffian constraint*

$$\underbrace{\begin{pmatrix} -\sin(\varphi_i) & \cos(\varphi_i) & l_{x,i} \cos(\varphi_i) + l_{y,i} \sin(\varphi_i) \end{pmatrix}}_{\boldsymbol{\alpha}_i(\varphi_i)} \begin{pmatrix} v_x \\ v_y \\ \omega \end{pmatrix} = \boldsymbol{\alpha}_i(\varphi_i) \mathbf{R}_z(\theta)^T \dot{\boldsymbol{\xi}} = \mathbf{0}. \quad (2.10)$$

The constraints are non-integrable, as there does not exist a function $h : \mathbb{R}^3 \rightarrow \mathbb{R}$ such that

$$\boldsymbol{\alpha}_i(\varphi_i) \mathbf{R}_z(\theta)^T \dot{\boldsymbol{\xi}} = \mathbf{0} \quad \Leftrightarrow \quad \frac{\partial h}{\partial \boldsymbol{\xi}} \dot{\boldsymbol{\xi}} = \mathbf{0}. \quad (2.11)$$

Thus, the constraints are called nonholonomic. Stacking all the n_w row vectors $\boldsymbol{\alpha}_i(\varphi_i)$ to a matrix $\mathbf{A}(\varphi) \in \mathbb{R}^{n_w \times 3}$ yields the nonholonomic constraint equation

$$\mathbf{A}(\varphi) \mathbf{v}_{\text{plf}} = \mathbf{0}. \quad (2.12)$$

For arbitrary values of the steering angles, $\text{rank}(\mathbf{A}) \leq 3$. For $\text{rank}(\mathbf{A}) = 3$, from (2.12) directly follows that $\mathbf{v}_{\text{plf}} = \mathbf{0}$, which means that no planar motion is possible. Thus, the wheels have to be oriented properly, such that $\text{rank}(\mathbf{A}) \leq 2$. Geometrically, this condition can be interpreted as the existence of an *instantaneous center of rotation* (ICR) at each time instant (the position of the ICR w. r. t. the rover frame may vary with time). This fact is illustrated in Figure 2.7 a). It means that at every instant, the velocity of any point on the rigid platform must be orthogonal to the straight line joining this point and the ICR. This particularly implies that the horizontal rotation axes of all the wheels must be concurrent at the ICR in order to allow a platform motion which is consistent with the constraints (2.12). In [CBD96], the *degree of steerability* r of a mobile robot is defined as the number of centered steering wheels that can be oriented independently in order to steer the robot. In the case of a platform with $n_w \geq 2$ independent centered steering wheels, $r = 2$ holds. This means that the angles of two arbitrary steering wheels define the position of the ICR, and the other wheels are oriented accordingly such that $\text{rank}(\mathbf{A}) = r = 2$. In this case, an instantaneously feasible velocity $\boldsymbol{\nu} \in \mathbb{R}^{3-r}$ exists, which lies in the null space of \mathbf{A} .

If all wheels are oriented properly such that $\text{rank}(\mathbf{A}) = r$, it is possible to represent the steering configuration $\boldsymbol{\varphi} \in \mathbb{R}^{n_w}$ by a minimal set of kinematic parameters $\boldsymbol{\phi}(\boldsymbol{\varphi}) \in \mathbb{R}^r$, and perform a velocity transformation

$$\mathbf{v}_{\text{plf}} = \boldsymbol{\Sigma}(\boldsymbol{\phi}) \boldsymbol{\nu} \quad (2.13)$$

with a full-rank matrix $\mathbf{\Sigma} \in \mathbb{R}^{3 \times (3-r)}$ satisfying $\mathbf{A}\mathbf{\Sigma} = \mathbf{0}$. Then, $\boldsymbol{\nu}$ is the instantaneously feasible velocity w. r. t. the steering configuration $\boldsymbol{\phi}$.

Note that the mapping $\Psi : \mathbf{v}_{\text{plf}} \mapsto (\boldsymbol{\phi}, \boldsymbol{\nu})$ is not unique. Several possibilities have been considered in literature, see e. g. [CBD96], [Con+08]. However, the following choice of $(\boldsymbol{\phi}, \boldsymbol{\nu})$ is considered for the whole-body control approaches developed in Chapter 3:

$$\begin{pmatrix} v_x \\ v_y \\ \omega \end{pmatrix} = \begin{pmatrix} \cos(\beta) \\ \sin(\beta) \\ \kappa \end{pmatrix} \nu, \quad (2.14)$$

where κ denotes the curvature and β is the lateral driving angle, see Figure 2.7 a). In this case, the kinematic configuration is parameterized by $\boldsymbol{\phi} = (\beta, \kappa)$, and ν equals the translational velocity tangential to the path of the platform. Note that the mapping Ψ is 1 : 1, and can be locally inverted according to

$$\begin{aligned} \nu &= \pm \sqrt{v_x^2 + v_y^2} \\ \beta &= \arctan\left(\frac{v_y}{v_x}\right) \\ \kappa &= \frac{\omega}{\nu}. \end{aligned} \quad (2.15)$$

Obviously, singularities occur at $v_x = v_y = 0$.

From (2.6) and (2.14), the steering angles and wheel velocities can be computed w. r. t. given values of κ , β and ν as

$$\varphi_i(\boldsymbol{\phi}) = \arctan\left(\frac{(\sin(\beta) + l_{x,i}\kappa)\nu}{(\cos(\beta) - l_{y,i}\kappa)\nu}\right), \quad (2.16)$$

$$v_{w,i}(\boldsymbol{\phi}, \boldsymbol{\nu}) = \pm \nu \sqrt{(\sin(\beta) + l_{x,i}\kappa)^2 + (\cos(\beta) - l_{y,i}\kappa)^2} \quad (2.17)$$

for $i = 1, \dots, 4$.

Mapping from the body wrench to the wheel torques

The torque equilibrium at the center of the platform in the local rover frame (x, y) (indicated by the left-superscript b for “body frame”) can be established as

$$\begin{aligned} \underbrace{\begin{pmatrix} F_x \\ F_y \\ M_z \end{pmatrix}}_{\mathbf{b}\boldsymbol{\tau}_{\text{plf}}} &= \begin{pmatrix} \sum_i^4 F_{x,i} \\ \sum_i^4 F_{y,i} \\ \sum_i^4 (-F_{x,i}l_{y,i} + F_{y,i}l_{x,i}) \end{pmatrix} = \\ &= \underbrace{\begin{pmatrix} \cos(\varphi_1) & \cos(\varphi_2) & \cos(\varphi_3) & \cos(\varphi_4) \\ \sin(\varphi_1) & \sin(\varphi_2) & \sin(\varphi_3) & \sin(\varphi_4) \\ \Delta_1 & \Delta_2 & \Delta_3 & \Delta_4 \end{pmatrix}}_{\mathbf{B}(\boldsymbol{\varphi})} \underbrace{\begin{pmatrix} f_{w,1} \\ f_{w,2} \\ f_{w,3} \\ f_{w,4} \end{pmatrix}}_{\mathbf{f}_w}, \end{aligned} \quad (2.18)$$

with

$$\Delta_i = -l_{y,i} \cos(\varphi_i) + l_{x,i} \sin(\varphi_i), \quad i = 1, \dots, 4, \quad (2.19)$$

see Figure 2.7 b). Note that the platform wrench generated by the four wheel forces is arbitrary in general. However, only the part of the wrench which lies in the null space of the nonholonomic constraints leads to a platform motion. The remaining part of the platform wrench lies in the range space of the constraint matrix and leads to internal tensions.³ One possibility to distinguish between the components of the platform wrench is via orthogonal projection in the null space of the constraint Jacobian matrix [Agh05]:

$$\boldsymbol{\tau}_{\text{plf}} = \boldsymbol{\tau}_{\text{plf}}^{\parallel} + \boldsymbol{\tau}_{\text{plf}}^{\perp} \quad (2.20)$$

$$\boldsymbol{\tau}_{\text{plf}}^{\parallel} = \mathbf{N}_A \mathbf{B} \mathbf{f}_w \quad (2.21)$$

$$\boldsymbol{\tau}_{\text{plf}}^{\perp} = (\mathbf{I} - \mathbf{N}_A) \mathbf{B} \mathbf{f}_w. \quad (2.22)$$

with $\boldsymbol{\tau}_{\text{plf}}^{\parallel}$ denoting the “acting” part of the platform wrench, and $\boldsymbol{\tau}_{\text{plf}}^{\perp}$ being the “passive” part leading to internal tensions. The null space projector \mathbf{N}_A can be computed for example using the relation

$$\mathbf{N}_A = \mathbf{I} - \mathbf{A}^+ \mathbf{A}, \quad (2.23)$$

with \mathbf{A}^+ denoting the Moore-Penrose pseudoinverse of the (rank-deficient) matrix \mathbf{A} [Agh05].⁴ From (2.21)–(2.22) and (2.12) then follows that

$$\mathbf{v}_{\text{plf}}^T \boldsymbol{\tau}_{\text{plf}}^{\parallel} = \mathbf{v}_{\text{plf}}^T \boldsymbol{\tau} \quad \text{and} \quad \mathbf{v}_{\text{plf}}^T \boldsymbol{\tau}_{\text{plf}}^{\perp} = 0. \quad (2.24)$$

Another possibility to eliminate the passive part of the platform wrench from the dynamic equations is to write the dynamic equations in terms of the constraint-consistent velocities from (2.13), which is shown in Section 2.2.3.

2.2.2 Dynamic equations of a serial robotic arm

A serial robotic arm with n_m joint angles $\boldsymbol{\eta} \in \mathcal{Q} \subset \mathbb{R}^{n_m}$ can be modeled by a set of nonlinear, second-order, ordinary differential equations, called the *equations of motion* (EOM) or the *rigid-body dynamics equations* of the robot. Thereby, \mathcal{Q} describes the admissible joint angles in a subset of $[-\pi, \pi)$ each, given by the operation limits of the robot. The EOM can be derived in different ways, for example by using Lagrange’s equations or the Newton-Euler formalism (a detailed derivation of the EOM of an open-chain manipulator can be found e. g. in [MLS94, pp. 169ff.]). They can be written as

$$\mathbf{M}_m(\boldsymbol{\eta}) \ddot{\boldsymbol{\eta}} + \mathbf{C}_m(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}}) \dot{\boldsymbol{\eta}} + \mathbf{g}_m(\boldsymbol{\eta}) = \boldsymbol{\tau}_m + \boldsymbol{\tau}_{m,\text{ext}}. \quad (2.25)$$

The following well-known properties hold:

1. The inertia matrix $\mathbf{M}_m(\boldsymbol{\eta}) \in \mathbb{R}^{n_m \times n_m}$ is symmetric positive definite (also written as s.p.d. or $\mathbf{M}_m(\boldsymbol{\eta}) \succ \mathbf{0}$) [MLS94, p. 171] and uniformly bounded [GSS98].
2. The effect of Coriolis and centrifugal forces can be written as $\mathbf{C}_m(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}}) \dot{\boldsymbol{\eta}}$. The choice of $\mathbf{C}_m(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}}) \in \mathbb{R}^{n_m \times n_m}$ is not unique. However, it can be chosen such that $\dot{\mathbf{M}}_m = \mathbf{C}_m + \mathbf{C}_m^T$, or $\dot{\mathbf{M}}_m - 2\mathbf{C}_m$ is skew-symmetric, respectively [MLS94, p. 171]. This property of the Coriolis matrix can be utilized to show passivity of (2.25) w. r. t. the power port $(\dot{\boldsymbol{\eta}}, \boldsymbol{\tau}_m)$, using the kinetic and potential energy $\frac{1}{2} \dot{\boldsymbol{\eta}}^T \mathbf{M}_m(\boldsymbol{\eta}) \dot{\boldsymbol{\eta}} + V_g(\boldsymbol{\eta})$ as storage function ($V_g(\boldsymbol{\eta})$ denotes the gravity potential).

³For a mathematical definition of the null space and the range space of a matrix, refer to Appendix A.1.5.

⁴A definition of the Moore-Penrose pseudoinverse and its computation in terms of the singular value decomposition of \mathbf{A} is shown in Appendix A.1.7. For alternative methods for the computation of \mathbf{N}_A based on the singular value decomposition of \mathbf{A} , see Appendix A.1.6.

Gravitational forces can be derived as $\mathbf{g}(\boldsymbol{\eta}) = (\partial V_g(\boldsymbol{\eta})/\partial \boldsymbol{\eta})^T \in \mathbb{R}^{n_m}$. The considered robot features a torque interface in each joint and is assumed to be fully actuated through $\boldsymbol{\tau}_m \in \mathbb{R}^{n_m}$. External torques are denoted by $\boldsymbol{\tau}_{m,\text{ext}} \in \mathbb{R}^{n_m}$.

2.2.3 Combined dynamic equations of a wheeled mobile manipulator

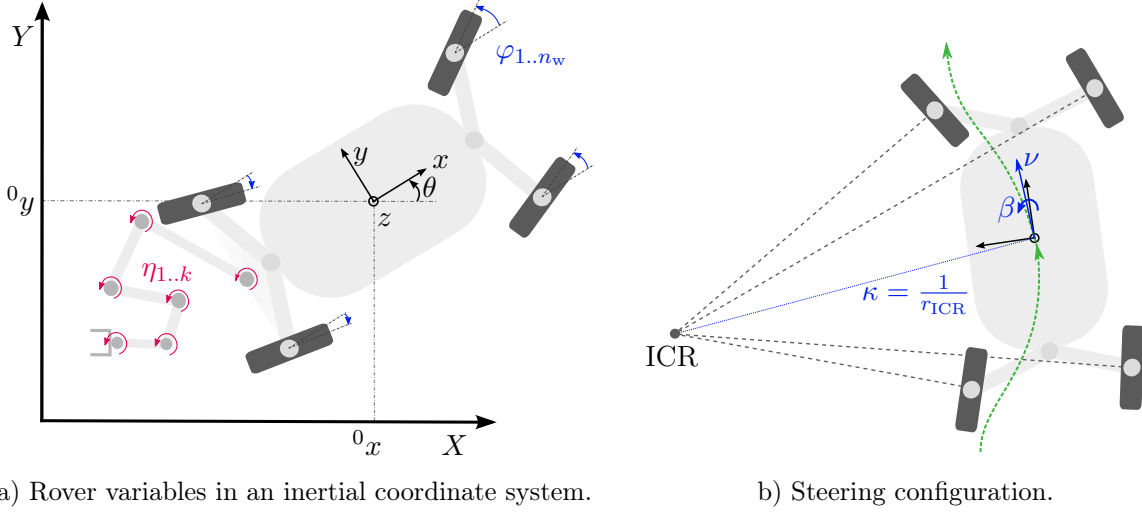


Figure 2.8 Sketch of the rover. The steering parameterization can be modeled using the lateral steering angle β and the curvature κ , alternatively the Cartesian coordinates of the ICR.

Figure 2.8 a) shows the kinematic structure of the LRU2, which comprises a robotic arm mounted on a mobile platform. The combined system coordinates are chosen as $\mathbf{q} = (\boldsymbol{\xi}^T \boldsymbol{\eta}^T)^T \in \mathbb{R}^n$ with $n = 3 + n_m$. The steering actuators are controlled on position level, and it is assumed that the dynamics of the low-level steering controllers can be neglected compared to the dynamics of the torque-controlled actuators. The steering parameters are then incorporated in the system equations as kinematic parameters instead of state variables. Using this simplification, the EOM of the system can be written as

$$\bar{\mathbf{M}}(\mathbf{q})\ddot{\mathbf{q}} + \bar{\mathbf{C}}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \bar{\mathbf{g}}(\mathbf{q}) + \mathbf{A}(\mathbf{q}, \boldsymbol{\phi})^T \boldsymbol{\lambda} = \bar{\boldsymbol{\tau}} + \bar{\boldsymbol{\tau}}_{\text{ext}} \quad (2.26)$$

$$\mathbf{A}(\mathbf{q}, \boldsymbol{\phi})\dot{\mathbf{q}} = \mathbf{0}, \quad (2.27)$$

where again $\bar{\mathbf{M}}(\mathbf{q}) \in \mathbb{R}^{n \times n}$ is the symmetric, positive definite inertia matrix, $\bar{\mathbf{C}}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} \in \mathbb{R}^n$ models centrifugal/Coriolis torques, and $\bar{\mathbf{g}}(\mathbf{q}) \in \mathbb{R}^n$ accounts for gravitational effects. The generalized forces acting on the system due to actuation are denoted by $\bar{\boldsymbol{\tau}} \in \mathbb{R}^n$. They can be written as

$$\bar{\boldsymbol{\tau}} = \begin{pmatrix} \boldsymbol{\tau}_{\text{plf}} \\ \boldsymbol{\tau}_m \end{pmatrix} = \underbrace{\begin{pmatrix} \mathbf{R}_z(\theta)^T \mathbf{B}(\boldsymbol{\phi}) & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}}_{\bar{\mathbf{B}}(\mathbf{q}, \boldsymbol{\phi})} \underbrace{\begin{pmatrix} \mathbf{f}_w \\ \boldsymbol{\tau}_m \end{pmatrix}}_{\mathbf{u}}, \quad (2.28)$$

where $\mathbf{u} \in \mathbb{R}^{n_w + n_m}$ are the actuator forces and $\bar{\mathbf{B}}(\mathbf{q}, \boldsymbol{\phi}) \in \mathbb{R}^{n \times (n_w + n_m)}$ maps the actuator forces to generalized forces acting on \mathbf{q} . Additionally, generalized external forces $\bar{\boldsymbol{\tau}}_{\text{ext}} \in \mathbb{R}^n$ may be present.

As the robot arm is holonomic, the constraint equation (2.27) can be written out as

$$\mathbf{A}(\mathbf{q}, \boldsymbol{\phi})\dot{\mathbf{q}} = (\mathbf{A}(\boldsymbol{\phi})\mathbf{R}_z(\theta)^T \quad \mathbf{0}) \begin{pmatrix} \dot{\boldsymbol{\xi}} \\ \dot{\boldsymbol{\eta}} \end{pmatrix} = \mathbf{0}. \quad (2.29)$$

The constraint forces corresponding to (2.27) have to be added to (2.26) as $\mathbf{A}(\mathbf{q}, \phi)^T \boldsymbol{\lambda}$. Thereby, the Lagrangian multipliers $\boldsymbol{\lambda} \in \mathbb{R}^{n_w}$ can be interpreted as reaction forces orthogonal to the feasible motion directions, which prevent a motion of the system in directions which would violate the constraints. The constraint forces can be eliminated by transforming the velocities to an instantaneously feasible subset $\mathbf{v} \in \mathbb{R}^{n-r}$ as proposed before:

$$\dot{\mathbf{q}} = \mathbf{S}(\mathbf{q}, \phi) \mathbf{v}, \quad (2.30)$$

such that $\mathbf{A}(\mathbf{q}, \phi) \mathbf{S}(\mathbf{q}, \phi) = \mathbf{0}$, introducing the full-rank matrix

$$\mathbf{S}(\mathbf{q}, \phi) = \begin{pmatrix} \mathbf{R}_z(\theta) \boldsymbol{\Sigma}(\phi) & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \in \mathbb{R}^{n \times (n-r)}. \quad (2.31)$$

Then, (2.26) can be transformed to

$$\mathbf{M}(\mathbf{q}, \phi) \dot{\mathbf{v}} + \mathbf{C}(\mathbf{q}, \mathbf{v}, \phi, \dot{\phi}) \mathbf{v} + \mathbf{g}(\mathbf{q}, \phi) = \mathbf{B}(\mathbf{q}, \phi) \mathbf{u} + \boldsymbol{\tau}_{\text{ext}} \quad (2.32)$$

with

$$\begin{aligned} \mathbf{M}(\mathbf{q}, \phi) &= \mathbf{S}(\mathbf{q}, \phi)^T \bar{\mathbf{M}}(\mathbf{q}) \mathbf{S}(\mathbf{q}, \phi), \\ \mathbf{C}(\mathbf{q}, \mathbf{v}, \phi, \dot{\phi}) &= \mathbf{S}(\mathbf{q}, \phi)^T \bar{\mathbf{M}}(\mathbf{q}) \dot{\mathbf{S}}(\mathbf{q}, \dot{\mathbf{q}}, \phi, \dot{\phi}) + \mathbf{S}(\mathbf{q}, \phi)^T \bar{\mathbf{C}}(\mathbf{q}, \dot{\mathbf{q}}) \mathbf{S}(\mathbf{q}, \phi), \\ \mathbf{g}(\mathbf{q}, \phi) &= \mathbf{S}(\mathbf{q}, \phi)^T \bar{\mathbf{g}}(\mathbf{q}), \\ \mathbf{B}(\mathbf{q}, \phi) &= \mathbf{S}(\mathbf{q}, \phi)^T \bar{\mathbf{B}}(\mathbf{q}, \phi), \\ \boldsymbol{\tau}_{\text{ext}}(\mathbf{q}, \phi) &= \mathbf{S}(\mathbf{q}, \phi)^T \bar{\boldsymbol{\tau}}_{\text{ext}}. \end{aligned} \quad (2.33)$$

It can easily be verified that the properties $\mathbf{M} \succ \mathbf{0}$ and the skew-symmetry of $\dot{\mathbf{M}} - 2\mathbf{C}$ hold for (2.33). Further, in [GSS98] it is shown that $\mathbf{M}(\mathbf{q}, \phi)$ is uniformly bounded in \mathbf{q} . The boundedness of $\mathbf{M}(\mathbf{q}, \phi)$ w.r.t. ϕ depends on the particular choice of ϕ .

Summarized, (2.30) and (2.32) form a state-space model which can be used for the design of whole-body control laws:

$$\dot{\mathbf{q}} = \mathbf{S}(\mathbf{q}, \phi) \mathbf{v} \quad (2.34)$$

$$\mathbf{M}(\mathbf{q}, \phi) \dot{\mathbf{v}} + \mathbf{C}(\mathbf{q}, \mathbf{v}, \phi, \dot{\phi}) \mathbf{v} + \mathbf{g}(\mathbf{q}, \phi) = \boldsymbol{\tau} + \boldsymbol{\tau}_{\text{ext}}(\mathbf{q}, \phi), \quad (2.35)$$

with $\boldsymbol{\tau} = \mathbf{B}(\mathbf{q}, \phi) \mathbf{u}$. Thereby, (2.34) models the kinematic properties of the platform, and is therefore referred to as *kinematic subsystem*. On the other hand, (2.35) models the dynamics of the combined mobile robot, and is thus called *dynamic subsystem*.

Remark 2. Writing (2.34)–(2.35) in the generic form of a state-space model $\dot{\mathbf{z}} = \mathbf{f}(\mathbf{z}) + \mathbf{g}(\mathbf{z}) \mathbf{u}$ with the state $\mathbf{z} \in \mathbb{R}^{(2n-r)} = (\mathbf{q}^T \ \mathbf{v}^T)^T$, one can immediately see that

$$D = \text{span}\{\mathbf{f}, \mathbf{g}_1, \dots, \mathbf{g}_{n-r}\} \quad (2.36)$$

does not map onto \mathbb{R}^{2n-r} (i.e. $\dim D < 2n-r$), as the column dimension of $\mathbf{S}(\mathbf{q}, \phi)$ is only $n-r$. Thus, the system does not fulfill Brockett's third necessary condition [Bro+83] (cf. Theorem 2, Appendix A.1.3). This means that it is not possible to stabilize an equilibrium point $\mathbf{z}_0^* = (\mathbf{q}_0^*, \mathbf{0})$ asymptotically by smooth time-invariant state feedback. This well-known result is e.g. stated in [CBD96].

2.2.4 Definition of the task space

Most of the times, the control objective involves a specified motion of the end effector, given e. g. in Cartesian coordinates. A convenient way to introduce the coordinates of the end effector into the respective control law is by defining the *task space* coordinates $\mathbf{x}(\mathbf{q}) \in \mathbb{R}^m$. In particular, $m = 6$ if the task space equals the Cartesian pose of the end effector. The transformation between n joint-space velocities $\dot{\mathbf{q}}$ and m task velocities $\dot{\mathbf{x}}$ can be written as

$$\dot{\mathbf{x}} = \mathbf{J}(\mathbf{q})\dot{\mathbf{q}} \quad (2.37)$$

with the task Jacobian matrix

$$\mathbf{J}(\mathbf{q}) = \frac{\partial \mathbf{x}(\mathbf{q})}{\partial \mathbf{q}}. \quad (2.38)$$

Conversely, a Cartesian wrench $\mathbf{F} \in \mathbb{R}^6$ at the end effector can be derived from the actuator torques as

$$\boldsymbol{\tau} = \mathbf{J}(\mathbf{q})^T \mathbf{F}, \quad (2.39)$$

which is a direct result of the conservation of power between joint space and task space:

$$\dot{\mathbf{q}}^T \boldsymbol{\tau} = \dot{\mathbf{x}}^T \mathbf{F}. \quad (2.40)$$

CHAPTER 3

Whole-Body Control

The main research question addressed in this thesis is how to operate a WMR for manipulation tasks, utilizing the complete workspace of holonomic arm and mobile platform. Thereby, the robot must be able to interact with previously unknown environments, including physical contact situations. This applies both for industrial human-robot collaboration and for planetary exploration scenarios, where robots are used for tasks which need a well-defined contact behavior. One example is the laser-induced breakdown spectroscopy, which was performed on Mt. Etna using the LRU2 [Leh+23]. Impedance-based control approaches are widely used in this context, as they are particularly suited for physical contact situations. In general, the concept of *passivity-based control* shapes the energy of the closed-loop system such that a desired equilibrium is stabilized, maintaining the inherent passivity of the robot dynamics.

While an exhaustive amount of literature exists concerning the passivity-based control of holonomic robot arms, the concept is not so common in the context of mobile robots, as mobile platforms are oftentimes controlled on velocity level. To overcome these limitations, a platform torque interface is developed in Section 3.1 by an optimization-based distribution of the wheel propulsion torques. Based on the platform torque interface, a whole-body impedance controller is then presented in Section 3.2. Thereby, quadratic optimization is used to cope with kinematic redundancy and actuator limitations. However, the approach assumes a constant steering configuration, which predestines the whole-body impedance controller for local manipulation tasks. Consequently, the next step is to incorporate a steering control law in the impedance-based control concept, which leads to the whole-body trajectory tracking control law presented in Section 3.3. Some interesting properties of the resulting control law are then analyzed (considering the simplified case of a holonomic robot arm) in Section 3.4.

3.1 Platform torque control

In order to create a physically compliant control law for the complete system (2.34)–(2.35), it is inevitable to actuate the wheels on torque level. This means creating a resulting wrench acting on the platform center by coordinating the steering angles and distributing the wheel propulsion torques. Thus, the first step towards whole-body control is the implementation of a torque interface for the platform, such that high-level control approaches can interface the platform

subsystem via a wrench acting on the platform coordinates ξ .¹ The question how to distribute forces and torques in over-actuated mechanical systems with a redundant set of actuators is a well-established research field in platform motion control, called *control allocation*. An overview of existing methods and approaches can be found in [JF13].

3.1.1 Optimization-based distribution of the wheel forces

One possibility to create a generalized platform torque interface for an arbitrary commanded wrench is to utilize the orthogonal projection approach from Section 2.2.1. Recall that the acting platform wrench is

$$\boldsymbol{\tau}_{\text{plf}}^{\parallel} = \mathbf{N}_A \mathbf{B} \mathbf{f}_w, \quad (3.1)$$

while

$$\boldsymbol{\tau}_{\text{plf}}^{\perp} = (\mathbf{I} - \mathbf{N}_A) \mathbf{B} \mathbf{f}_w \quad (3.2)$$

is the passive part of the platform wrench which does not contribute to the platform motion. Consequently, the complete wrench acting on the platform is the sum of the acting and the passive wrench:

$$\boldsymbol{\tau}_{\text{plf}} = \boldsymbol{\tau}_{\text{plf}}^{\parallel} + \boldsymbol{\tau}_{\text{plf}}^{\perp}. \quad (3.3)$$

Obviously, the platform is redundantly actuated, $n_w = 4$ actuation torques map to a three-dimensional wrench. A possible solution for solving (3.1) for the wheel forces \mathbf{f}_w is to use a generalized inverse of $\mathbf{N}_A \mathbf{B}$. However, as the achievable actuator torques are limited in practice, it is advantageous to formulate a quadratic optimization problem

$$\mathbf{f}_w^* = \arg \min \left(\frac{1}{2} \mathbf{f}_w^T \mathbf{Q}_f \mathbf{f}_w + \frac{1}{2} \tilde{\boldsymbol{\tau}}_{\text{plf}}^{\parallel T} \mathbf{Q}_{\parallel} \tilde{\boldsymbol{\tau}}_{\text{plf}}^{\parallel} + \frac{1}{2} \tilde{\boldsymbol{\tau}}_{\text{plf}}^{\perp T} \mathbf{Q}_{\perp} \tilde{\boldsymbol{\tau}}_{\text{plf}}^{\perp} \right) \quad (3.4)$$

$$\text{s. t. } \mathbf{N}_A \mathbf{B} \mathbf{f}_w - \boldsymbol{\tau}_{\text{plf,des}}^{\parallel} = \tilde{\boldsymbol{\tau}}_{\text{plf}}^{\parallel} \quad (3.5)$$

$$(\mathbf{I} - \mathbf{N}_A) \mathbf{B} \mathbf{f}_w - \boldsymbol{\tau}_{\text{plf,des}}^{\perp} = \tilde{\boldsymbol{\tau}}_{\text{plf}}^{\perp} \quad (3.6)$$

$$\mathbf{f}_{w,\min} < \mathbf{f}_w < \mathbf{f}_{w,\max}, \quad (3.7)$$

with $\mathbf{Q}_f \in \mathbb{R}^{4 \times 4}$ is the input weighting matrix, $\mathbf{Q}_{\parallel} \in \mathbb{R}^{3 \times 3}$ weights the error between the commanded and the resulting active platform wrench, and $\mathbf{Q}_{\perp} \in \mathbb{R}^{3 \times 3}$ is a weighting matrix for the inactive actuation torques. In this way, it is possible to weight the acting input wrench separately to the passive input wrench. It is also possible to chose $\boldsymbol{\tau}_{\text{plf,des}}^{\perp} = \mathbf{0}$ instead of $\boldsymbol{\tau}_{\text{plf,des}}^{\perp} = (\mathbf{I} - \mathbf{N}_A) \boldsymbol{\tau}_{\text{plf,des}}$ in order to minimize the amount of internal forces acting on the platform structure. Indeed, adding more constraints to (3.5)–(3.7) is straightforward, for example limits on the permitted wheel force related to the terrain friction properties as proposed in [ID04b]. The optimization problem can be solved by using a realtime-capable solver for quadratic problems. One example is the open-source software package qpOASES [Fer+14], which implements a parametric active set method for solving convex quadratic programming (QP) problems and for computing critical points of non-convex QP problems in C++, providing customizable interfaces to Matlab, Simulink and Python (among others). The algorithm is described in detail in [FBD08].

¹Parts of the theory as well as the experimental results of Section 3.1 and Section 3.2 have been published in [BDO18].

Remark 3. Instead of separately weighting the active and passive platform wrench, a simplified version of (3.4)–(3.7) can be implemented as

$$\mathbf{f}_w^* = \arg \min \left(\frac{1}{2} \mathbf{f}_w^T \mathbf{Q}_f \mathbf{f}_w + \frac{1}{2} \mathbf{e}^T \mathbf{Q}_e \mathbf{e} \right) \quad (3.8)$$

$$\text{s. t. } \mathbf{B} \mathbf{f}_w - \boldsymbol{\tau}_{\text{plf,des}} = \mathbf{e} \quad (3.9)$$

$$\mathbf{f}_{w,\min} < \mathbf{f}_w < \mathbf{f}_{w,\max} . \quad (3.10)$$

In this way, the sum of the acting platform wrench and the passive platform wrench is weighted with a single matrix $\mathbf{Q}_e \in \mathbb{R}^{3 \times 3}$. By choosing the desired platform wrench such that $\boldsymbol{\tau}_{\text{plf,des}}^\perp = \mathbf{0}$, the passive platform wrench is automatically minimized in (3.8)–(3.10).

Remark 4. Instead of dividing the platform wrench in its acting and passive part, another possibility is to directly use the reduced platform wrench $\boldsymbol{\tau}_{\text{plf}}^r \in \mathbb{R}^{3-r}$ for optimizing the wheel propulsion torques (as $r = 2$, $\boldsymbol{\tau}_{\text{plf}}^r$ is scalar for the platform model described in Section 2.2.1). Recalling Section 2.2.3, from (2.33) with (2.28) and (2.31) it can be deduced that the constraint-consistent platform input torque is

$$\boldsymbol{\tau}_{\text{plf}}^r = \boldsymbol{\Sigma}^T \mathbf{B} \mathbf{f}_w . \quad (3.11)$$

In this case, the optimization problem can be established as

$$\mathbf{f}_w^* = \arg \min \left(\frac{1}{2} \mathbf{f}_w^T \mathbf{Q}_f^r \mathbf{f}_w + \frac{1}{2} \mathbf{e}_r^T \mathbf{Q}_e^r \mathbf{e}_r \right) \quad (3.12)$$

$$\text{s. t. } \boldsymbol{\Sigma}^T \mathbf{B} \mathbf{f}_w - \boldsymbol{\tau}_{\text{plf,des}}^r = \mathbf{e}_r \quad (3.13)$$

$$\mathbf{f}_{w,\min} < \mathbf{f}_w < \mathbf{f}_{w,\max} . \quad (3.14)$$

This approach was adopted for the implementation of the passivity-based tracking controller, which is described in Section 3.3.

3.1.2 Experiments

The remaining question is how to choose the weighting matrices. This question was addressed by performing an experimental study with the LRU2, which is shown in [BDO18]. Thereby, an impedance control law is applied to the mobile platform torque interface, in order to regulate the platform position to a desired value in x -direction.² The resulting impedance wrench is generated by the wheels. Note that during the experiments, the optimization of the wheel forces was implemented as (3.8)–(3.10), with the desired platform wrench satisfying $\boldsymbol{\tau}_{\text{plf,des}} = \mathbf{N}_A \boldsymbol{\tau}_{\text{plf,des}}$.

Experiment 3.1 – Comparison of different choices of the weighting matrix \mathbf{Q}_f : In the first experiment, a cost function minimizing the wheel forces is compared to a cost function which minimizes the computed wheel power consumption $P_w = \mathbf{v}_w^T \mathbf{f}_w$. For this purpose, the weighting matrix \mathbf{Q}_f is chosen as $\mathbf{Q}_f^f = \mathbf{I}$ and $\mathbf{Q}_f^v = \text{diag}(v_{w,1}^2, v_{w,2}^2, v_{w,3}^2, v_{w,4}^2)$ (comprising the measured wheel velocities). While performing this experiment, the front wheels are located on a surface with different friction properties than the rear wheels, which mimics the locally different soil properties that can occur during driving in outdoor environments.

Figure 3.1 a) depicts the reference trajectory and the resulting platform motion in x -direction for the two cases of force and power optimization. The following plots b) and c) show the wheel

²In practice, a trapezoidal trajectory in x -direction is commanded, to avoid jumps in the input signal. However, the experiments are performed using an impedance law without velocity and acceleration feed-forward terms as introduced in (3.15).

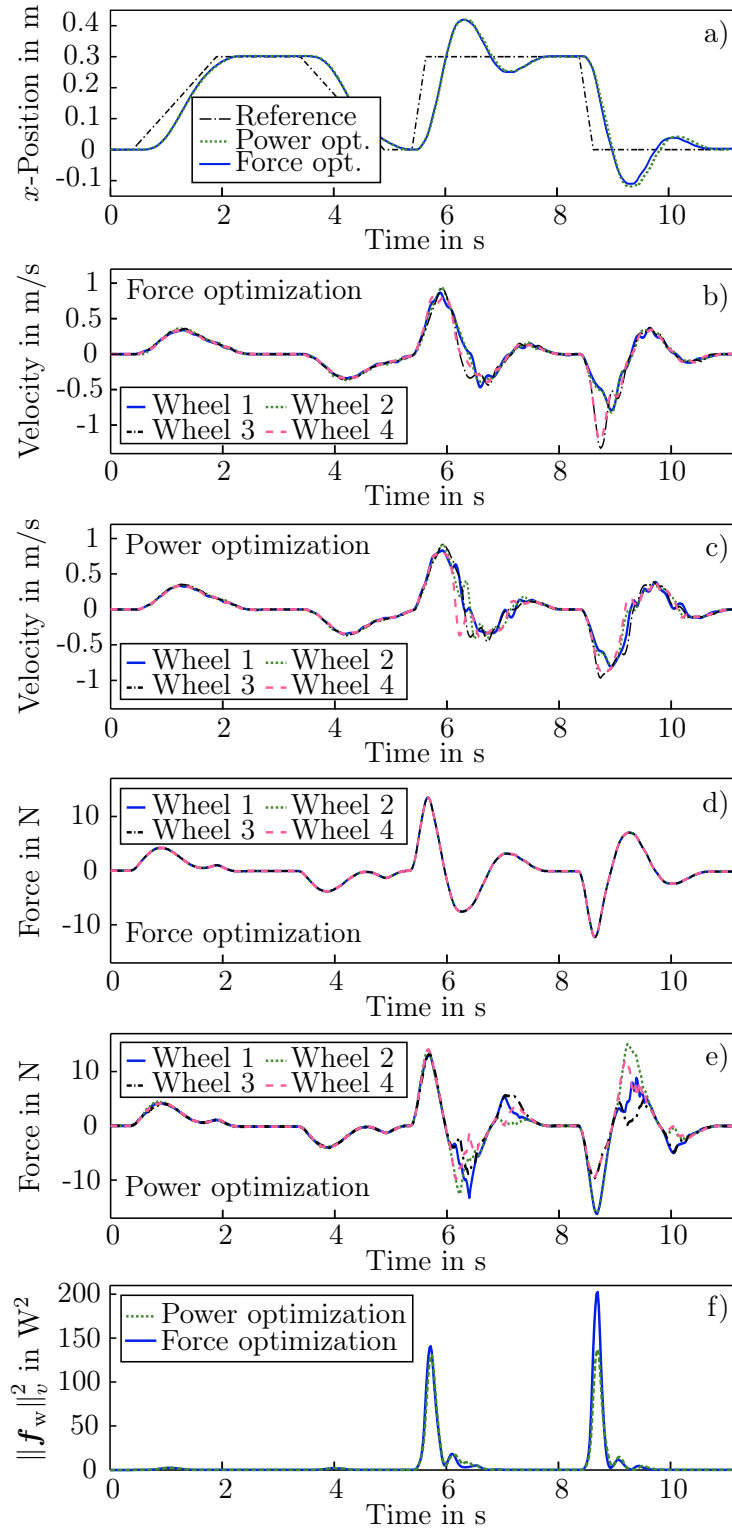


Figure 3.1 Experiment 3.1: Comparison of different choices of the weighting matrix \mathbf{Q}_f . a) Reference trajectory and system response. b)–e) Wheel velocity measurements and corresponding optimized wheel force commands. f) Comparison of the power-based norm for the two optimization approaches. From [BDO18].

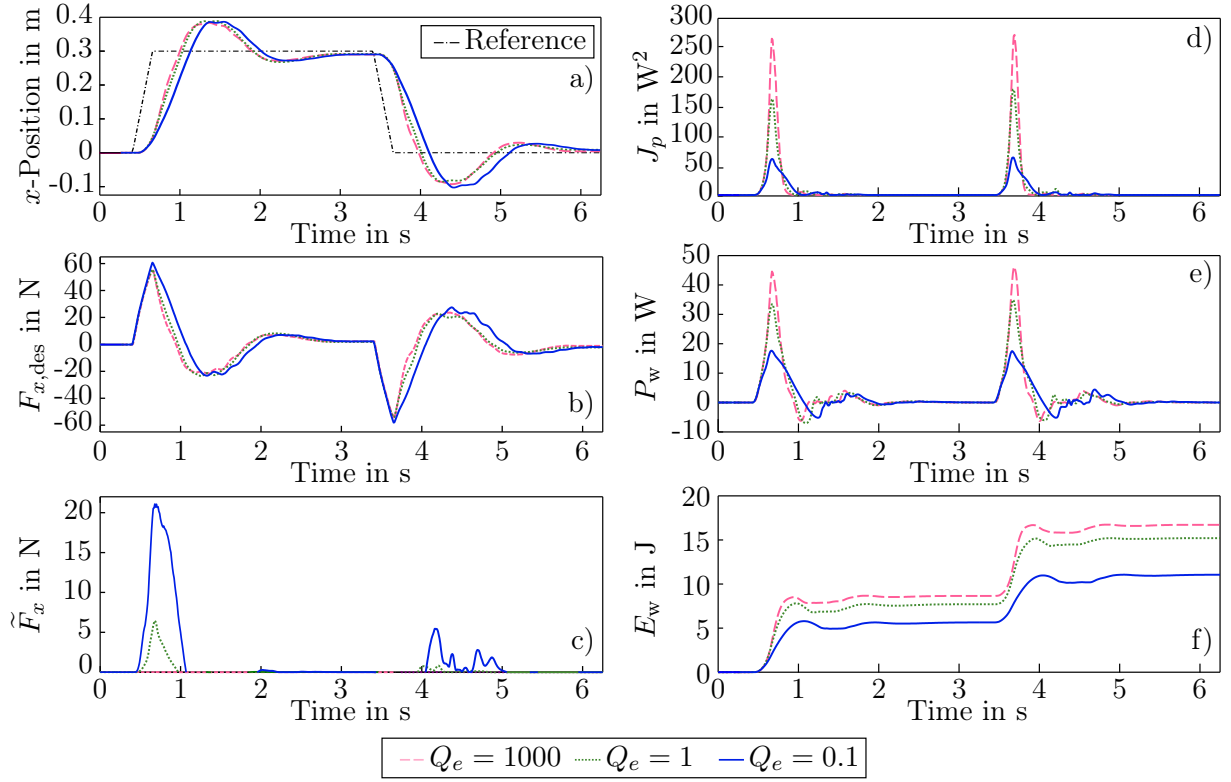


Figure 3.2 Experiment 3.2: Comparison of different values of the weighting factor Q_e . a) Reference trajectory and resulting platform motion. b) Force $F_{x,\text{des}}$, computed by the impedance controller. c) Force error $\tilde{F}_x = F_{x,\text{des}} - F_x$. The resulting force F_x is the sum of all wheel forces in x -direction. d) Power-based cost function $J_p = \frac{1}{2} \mathbf{f}_w^T \mathbf{Q}_f^v \mathbf{f}_w + \frac{1}{2} \mathbf{e}^T \mathbf{Q}_e \mathbf{e}$. e) Power consumption of the wheels, computed as $P_w = \mathbf{v}_w^T \mathbf{f}_w$. f) Energy consumption of the wheels (computed as $E_w = \int_0^t P_w d\tau$). From [BDO18].

velocity measurements for the two optimization approaches, while the corresponding wheel force commands are plotted in d) and e). The first part of the trajectory is from $t = 0$ s until $t = 5$ s. It features a trapezoidal reference signal with a small slope, which produces moderate commanded forces. Thus, the two optimization procedures generate almost identical wheel force commands, which indicates that no significant wheel slippage occurs. From $t = 5$ s to $t = 10$ s, the second part of the trajectory features a larger slope, which leads to a higher amplitude in the commanded force. Plots b) and c) show that the rear wheels 3 and 4 start to slip especially during the fast backward motion at $8 \text{ s} < t < 10 \text{ s}$. As the velocity measurements are not used for the command generation in case of force optimization, the respective wheel velocities in plot b) are high, especially at $t \approx 9$ s. On the other hand, the commanded forces depend on the measured velocities in the case of power optimization. From Figure 3.1 c) and e), it can be verified that the commanded wheel forces of the rear wheels 3 and 4 are reduced after an initial increase of the corresponding velocities. As a result, the commanded force is redistributed to the front wheels 1 and 2. Figure 3.1 f) shows the effect of the different optimization policies on the power-based norm³ $\|\mathbf{f}_w\|_v^2 = \frac{1}{2} \mathbf{f}_w^T \mathbf{Q}_f^v \mathbf{f}_w$. It can be observed that the value of the power norm is lower for the power optimization approach, particularly for the case of slippage of the two rear wheels.

³This norm is assumed to be directly related to the overall power consumption of the rover's driving unit, even though the (electrical) power consumption of the rover is not measured directly.

Experiment 3.2 – Comparison of different choices of the weighting matrix \mathbf{Q}_e : The goal of the second experiment is to compare different choices of the weighting matrix, here chosen as $\mathbf{Q}_e = Q_e \mathbf{I}$ with the weighting factor Q_e . Therefore, another trapezoidal trajectory is commanded, which features steep changes of the desired position. The choice of the trajectory is supposed to induce a high command force and thus slip in all wheels. Figure 3.2 a) shows the commanded trajectory together with the resulting platform motion in x -direction. The commanded force $F_{x,\text{des}}$ in Figure 3.2 b) is produced by the high-level impedance controller. The force error $\tilde{F}_x = F_{x,\text{des}} - F_x$ is depicted in Figure 3.2 c), comparing different values of Q_e . Thereby, the resulting force $F_x = \sum_{i=1}^4 f_{w,i} \cos(\varphi_i)$ is the sum of all wheel forces in x -direction, obtained by a power optimization according to (3.4). Figure 3.2 d)–f) show the power-based cost function $J_p = \frac{1}{2} \mathbf{f}_w^T \mathbf{Q}_f^v \mathbf{f}_w + \frac{1}{2} \mathbf{e}^T \mathbf{Q}_e \mathbf{e}$, the computed power consumption of the wheels $P_w = \mathbf{v}_w^T \mathbf{f}_w$, and the resulting energy $E_w = \int_0^t P_w d\tau$, respectively. It can be seen that the resulting force is very close to the commanded force for the high weight $Q_e = 1000 \text{ m}^2/\text{s}^2$. For $Q_e = 1 \text{ m}^2/\text{s}^2$, the force error in Figure 3.2 c) (green dashed line) is still small compared to the total value of the commanded force. Also the resulting platform motion is similar to the case when $Q_e = 1000 \text{ m}^2/\text{s}^2$, cf. Figure 3.2 a). However, Figure 3.2 e) reveals a significantly lower power consumption for $Q_e = 1 \text{ m}^2/\text{s}^2$, which is possible due to a reduced wheel slippage. Comparing the (computed) energy consumption in plot f), it becomes clear that approximately 1 J is saved each time a rapid movement occurs. With regard to planetary exploration scenarios with an instantaneously limited power supply, the importance of the energy savings must be highlighted when compared to the rather marginal losses in tracking performance. Decreasing the weighting factor further to $Q_e = 0.1 \text{ m}^2/\text{s}^2$, the deviation between desired and generated force grows up to 20 N, which leads to a noticeably slower system response, cf. Figure 3.2 c) and a) (blue solid line). On the other hand, Figure 3.2 f) shows energy savings up to 4 J per acceleration process compared to $Q_e = 1000 \text{ m}^2/\text{s}^2$. Noticeably, the goal position is still reached for all choices of Q_e , as the high-level impedance controller compensates for errors in the platform torque interface.

3.2 Whole-body impedance control

Building upon the platform torque interface, a whole-body controller utilizing the platform as an additional, virtual joint can be straightforwardly implemented.⁴ A proper choice for the realization of a defined interaction behavior of the robot is an impedance controller, which has been introduced in the seminal work of [Hog85]. The idea of impedance control is to influence the mechanical impedance of the robot in contact with an environment, that is, the mapping from (generalized) velocities to (generalized) forces [Ott08]. This leads to a well-defined interaction behavior of the robot in contact with an environment, which can be physically described as an admittance [Hog85] (i. e., the environment maps forces to velocities).

3.2.1 Cartesian impedance control

One possibility to implement physical compliance at the end effector is Cartesian impedance control [AOH07]. Thereby, a desired wrench $\mathbf{F}_{\text{imp}} \in \mathbb{R}^6$ is computed w. r. t. the Cartesian coordinates of the end effector $\mathbf{x} \in \mathbb{R}^6$. It is derived as the gradient of a virtual spring potential V_{imp} , a damping term, and possibly dynamic feed-forward terms:

$$\mathbf{F}_{\text{imp}} = - \left(\frac{\partial V_{\text{imp}}}{\partial \mathbf{x}} \right)^T - \mathbf{D}_x(\mathbf{q})\dot{\mathbf{x}} + \mathbf{M}_c\ddot{\mathbf{x}}_{\text{des}} + \mathbf{C}_c\dot{\mathbf{x}}_{\text{des}}. \quad (3.15)$$

⁴Recall that Section 3.2 is mainly based on the results presented in [BDO18].

Thereby, the spring potential can be defined as $V_{\text{imp}} = \frac{1}{2} \tilde{\mathbf{x}}^T \mathbf{K}_x \tilde{\mathbf{x}}$, with the Cartesian error $\tilde{\mathbf{x}} = \mathbf{x} - \mathbf{x}_{\text{des}}$ and the s.p.d. stiffness matrix is $\mathbf{K}_x \in \mathbb{R}^{6 \times 6}$. The damping term contains the damping matrix $\mathbf{D}_x(\mathbf{q}) \in \mathbb{R}^{6 \times 6}$. In general, the choice of the damping matrix is arbitrary as long as it is positive definite. Thus, defining a constant damping matrix could appear as a simple and functional choice. However, in practice it is convenient to specify desired damping ratios $\boldsymbol{\zeta} \in \mathbb{R}^6$ in every Cartesian direction. In this case, the damping matrix becomes configuration-dependent. Knowledge of the inertia and stiffness matrices enables the computation of the desired damping ratios, for example by using the *Double Diagonalization* approach from [Alb+03]. In this case, the damping matrix can be formally written as

$$\mathbf{D}(\mathbf{q}) = \mathbf{D}(\mathbf{M}(\mathbf{q}), \mathbf{K}_x, \boldsymbol{\zeta}). \quad (3.16)$$

The additional feed-forward terms $\mathbf{M}_c \ddot{\mathbf{x}}_{\text{des}}$ and $\mathbf{C}_c \dot{\mathbf{x}}_{\text{des}}$ may be added to improve trajectory tracking (alternatively, they can be omitted without affecting the contact properties of the closed-loop system). Thereby, the matrices

$$\mathbf{M}_c = (\mathbf{J}^{M+})^T \mathbf{M} \mathbf{J}^{M+}, \quad (3.17)$$

$$\mathbf{C}_c = (\mathbf{J}^{M+})^T \mathbf{C} \mathbf{J}^{M+} + (\mathbf{J}^{M+})^T \mathbf{M} \dot{\mathbf{J}} \quad (3.18)$$

are the task-space inertia and Coriolis/centrifugal matrices, which contain the dynamically consistent pseudoinverse of \mathbf{J} [Kha87]:

$$\mathbf{J}^{M+} = \mathbf{M}^{-1} \mathbf{J}^T (\mathbf{J} \mathbf{M}^{-1} \mathbf{J}^T)^{-1}. \quad (3.19)$$

Finally, the desired wrench \mathbf{F}_{imp} can be projected to the actuator torques utilizing the relationship between the joint torques and the corresponding task-space wrench (2.39), and adding a gravity compensation on joint level:

$$\boldsymbol{\tau} = \mathbf{g}(\mathbf{q}) + \mathbf{J}(\mathbf{q})^T \mathbf{F}_{\text{imp}}. \quad (3.20)$$

The joint torques which are generated by the impedance controller, that is $\boldsymbol{\tau} = \mathbf{J}(\mathbf{q})^T \mathbf{F}_{\text{imp}}$, are called *range space torques*, as they consist of a linear combination of the columns of the transposed task Jacobian matrix and thus define the range space of $\mathbf{J}(\mathbf{q})^T$ (for a mathematical definition of the range space and the null space of a matrix, see Definition 2 in Appendix A.1.5). On the other hand, as $\mathbf{J} \in \mathbb{R}^{m \times n}$ with $m < n$, there exist redundant DOF which can be used to fulfill secondary tasks. Thereby, it is oftentimes desirable that the secondary task is executed as well as possible, but without disturbing the task on the first level. This can be achieved by projecting the control input of the secondary task into the null space of the primary task via

$$\boldsymbol{\tau}_{\mathcal{N}} = \mathbf{N} \boldsymbol{\tau}_0, \quad (3.21)$$

where $\boldsymbol{\tau}_0 \in \mathbb{R}^n$ is an arbitrary torque command originating from a secondary task, $\mathbf{N} \in \mathbb{R}^{n \times n}$ is the *null space projector* and $\boldsymbol{\tau}_{\mathcal{N}} \in \mathbb{R}^n$ denotes the corresponding *null space torque*.⁵ The null space projector is said to be *statically consistent* if it fulfills $\mathbf{N}(\mathbf{q}) \mathbf{J}(\mathbf{q})^T = \mathbf{0}$ at any steady state $\dot{\mathbf{q}} = \ddot{\mathbf{q}} = \mathbf{0}$. This means that the secondary task does not generate interfering forces in the task space of the primary task at any static equilibrium. If, additionally, the secondary task does not induce accelerations in the task space during transient motions, it is called *dynamically consistent* [Kha87]. In this case, the condition $\mathbf{J}(\mathbf{q}) \mathbf{M}(\mathbf{q})^{-1} \mathbf{N}(\mathbf{q}) = \mathbf{0}$ must always hold [Die15, pp. 66–67]. A detailed explanation of null space projectors and the corresponding properties of static and

⁵In this work, only two task levels are considered, due to the small degree of redundancy of the LRU2. However, extensions to multiple task levels are possible, see e.g. [Die15, Chapter 4] for details.

dynamic consistency can be found in [Die15, Chapter 4, pp. 63ff.] and the references therein. Assuming a non-singular task Jacobian matrix, the null space projector can be computed as

$$\mathbf{N} = \mathbf{I} - \mathbf{J}^T(\mathbf{J}^T\mathbf{W}^+)^{-1}\mathbf{J}, \quad (3.22)$$

with the s.p.d. weighting matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ and the weighted pseudoinverse

$$\mathbf{J}^{\mathbf{W}^+} = \mathbf{W}^{-1}\mathbf{J}^T(\mathbf{J}\mathbf{W}^{-1}\mathbf{J}^T)^{-1}. \quad (3.23)$$

One possible choice is $\mathbf{W} = \mathbf{I}$, which results in $\mathbf{J}^{\mathbf{W}^+} = \mathbf{J}^+$ (the Moore-Penrose pseudoinverse). The resulting projector is computationally cheap and fulfills the static consistency condition. Another popular choice is $\mathbf{W} = \mathbf{M}$, which results in a dynamically consistent null space projector [Kha87]. The computation of different null space projectors in terms of the singular value decomposition of the task Jacobian matrix \mathbf{J} is derived in [Die15, Appendix B, pp. 165ff.].

3.2.2 Redundancy resolution via quadratic optimization

For a non-redundant and holonomic robot under Cartesian impedance control, passivity and asymptotic stability can be concluded using the invariance principle of LaSalle [LaS68]. However, most robots possess more DOF than the $m = 6$ DOF of a Cartesian task at the end effector. Thus, the choice of actuator torques is not unambiguous. As mentioned before, there are different possibilities for redundancy resolution, for example the task augmentation as introduced in [Ser89], or the implementation of a strict task hierarchy via null space projectors as e.g. shown in [SS91] or [DOA13]. The advantage of the aforementioned approaches is the analytical proof of stability. On the other hand, constraints such as joint torque limits are not considered in general.

An elegant solution to the problem of kinematic redundancy is again to apply numerical optimization strategies, which provides the opportunity to include additional constraints such as unilateral torque limits or permissible joint angle ranges. Following this approach, the global optimization problem can be established as

$$\boldsymbol{\tau}^* = \arg \min (J_u + J_{\mathcal{R}} + J_{\mathcal{N}}) \quad (3.24)$$

$$\text{s. t. } (\mathbf{I} - \mathbf{N})\boldsymbol{\tau} = \mathbf{J}^T \mathbf{F}, \quad (3.25)$$

$$\mathbf{N}\boldsymbol{\tau} = \boldsymbol{\tau}_{\mathcal{N}}, \quad (3.26)$$

$$\boldsymbol{\tau}_{\min} < \boldsymbol{\tau} < \boldsymbol{\tau}_{\max}. \quad (3.27)$$

with the permissible lower and upper bounds $\boldsymbol{\tau}_{\min}$ and $\boldsymbol{\tau}_{\max}$, and with

$$J_u = \frac{1}{2}\boldsymbol{\tau}^T \mathbf{Q}_u \boldsymbol{\tau}, \quad (3.28)$$

$$J_{\mathcal{R}} = \frac{1}{2}(\mathbf{F}_{\text{imp}} - \mathbf{F})^T \mathbf{Q}_{\mathcal{R}}(\mathbf{F}_{\text{imp}} - \mathbf{F}), \quad (3.29)$$

$$J_{\mathcal{N}} = \frac{1}{2}(\mathbf{N}\boldsymbol{\tau}_{\mathcal{N},\text{des}} - \boldsymbol{\tau}_{\mathcal{N}})^T \mathbf{Q}_{\mathcal{N}}(\mathbf{N}\boldsymbol{\tau}_{\mathcal{N},\text{des}} - \boldsymbol{\tau}_{\mathcal{N}}). \quad (3.30)$$

Thereby, J_u is the component responsible for the distribution of the input torques $\boldsymbol{\tau}$, which can be weighted by the selection of $\mathbf{Q}_u \in \mathbb{R}^{n \times n}$. A task-related component $J_{\mathcal{R}}$ is used to realize the desired impedance wrench at the end effector. The associated weighting matrix is $\mathbf{Q}_{\mathcal{R}} \in \mathbb{R}^{m \times m}$. A secondary task can be added by the torque-related term $J_{\mathcal{N}}$, which comprises the weighting matrix $\mathbf{Q}_{\mathcal{N}} \in \mathbb{R}^{n \times n}$. By (3.26), it is ensured that the related torque $\boldsymbol{\tau}_{\mathcal{N}}$ is in the null space of the task Jacobian matrix and thus does not disturb the task execution. Therefore, the dynamically consistent null space projector

$$\mathbf{N} = \mathbf{I} - \mathbf{J}^T(\mathbf{J}^T\mathbf{M}^+)^{-1}\mathbf{J} \quad (3.31)$$

is chosen.

3.2.3 Experiments

Finally, the approach is validated with experiments at the LRU2 [BDO18]. The performance of the undisturbed closed-loop system is investigated with a regulation task at the end effector. Then, the behavior in a physical contact situation (an operator disturbing the end effector) is analyzed while hitting artificially lowered joint torque limits. The following experiments are also shown in the supplementary video of [BDO18]. During the experiments, the desired Cartesian stiffness is set to 450 N/m for translations and 12.5 Nm/rad for rotations, all damping ratios are chosen as 0.5 s^{-1} . The steering angles are constantly set to $\varphi_i = 0$, $i = 1, \dots, 4$, such that only a platform motion in x -direction is possible. According to the results of the preceding section, the wheel torque distribution problem is parameterized with $\mathbf{Q}_f = \mathbf{Q}_f^v$, $\mathbf{Q}_e = \text{diag}(10, 10, 10)$. A joint impedance controller for all six joints with a stiffness of 20 Nm/rad is implemented as null space task, with the desired joint angles close to the initial configuration. An additional damping in the null space is not implemented due to the large intrinsic joint damping of the Jaco arm. The following weights are chosen for the optimization (the units of the weighting matrices are chosen consistently such that J is unit-free): $\mathbf{Q}_u = \text{diag}(v_{w,1}^2, v_{w,2}^2, v_{w,3}^2, v_{w,4}^2, 0, 0, 0, 0, 0)$, $\mathbf{Q}_R = \text{diag}(1, 1, 1, 1, 1, 1)$, $\mathbf{Q}_N = \text{diag}(0, 1, 1, 1, 1, 1, 1)$.

Experiment 3.3 – Regulation of the x -coordinate of the end effector: The experiment shows the undisturbed regulation of a desired pose of the end effector. The desired and measured trajectory in x -direction is shown in Figure 3.3 a) (the desired values of all other DOF are kept constant). It is visible that the low stiffness of the impedance controller (3.15) results in a steady-state error at the end effector. This error can be explained by unmodeled friction, especially between wheels and ground. In Figure 3.3 b), the range space torques $\boldsymbol{\tau}_R = (\mathbf{I} - \mathbf{N})\boldsymbol{\tau}$ are plotted. It can be verified that between 11 s and 16 s, the commanded forces are constant while the robot does not move. In particular, the solid blue line shows a commanded platform force in x -direction of approx. 10 N, which however does not induce a platform motion.

The null space torques $\boldsymbol{\tau}_N = \mathbf{N}\boldsymbol{\tau}$ are a result of the projected joint impedance law, which keeps the joints close to their initial configuration without affecting the wrench at the end effector. They are shown in Figure 3.3 c). Note that the remaining error in the position of the end effector could be compensated by the arm and thus be significantly smaller without a dedicated null space reference torque (i.e., choosing $\mathbf{Q}_N = \mathbf{0}$). In this experiment, no actuator torque limits are met, meaning that both range space force and null space torque can be accomplished without limitations. Thus, both the range space and the null space error norms ($\sqrt{J_R}$ and $\sqrt{J_N}$, respectively) are small, as Figure 3.3 d) shows.

Experiment 3.4 – Physical interaction with the platform: Afterwards, the differences between the presented whole-body controller and approaches such as [Die+16] are experimentally investigated. In particular, the difference between the presented platform torque interface to a platform velocity controller combined with an admittance interface is highlighted. The main difference between a whole-body controller with a torque-controlled platform (like the whole-body impedance controller at hand) and a velocity-controlled platform is that platform torque control enables active compliance, whereas the admittance simulation used together with a platform velocity controller in [Die+16] only mimics impedance behavior w.r.t. interactions with the end effector. That means that a platform velocity controller does not feature active compliance w.r.t. interactions with the platform. While this can be seen as a disadvantage (thinking of human-robot interaction or unknown environments in general), an advantage of the admittance interface is that a desired platform inertia and damping can be assigned, while in the platform torque control approach, the physical inertia and damping properties of the platform appear unchanged. The actively compliant

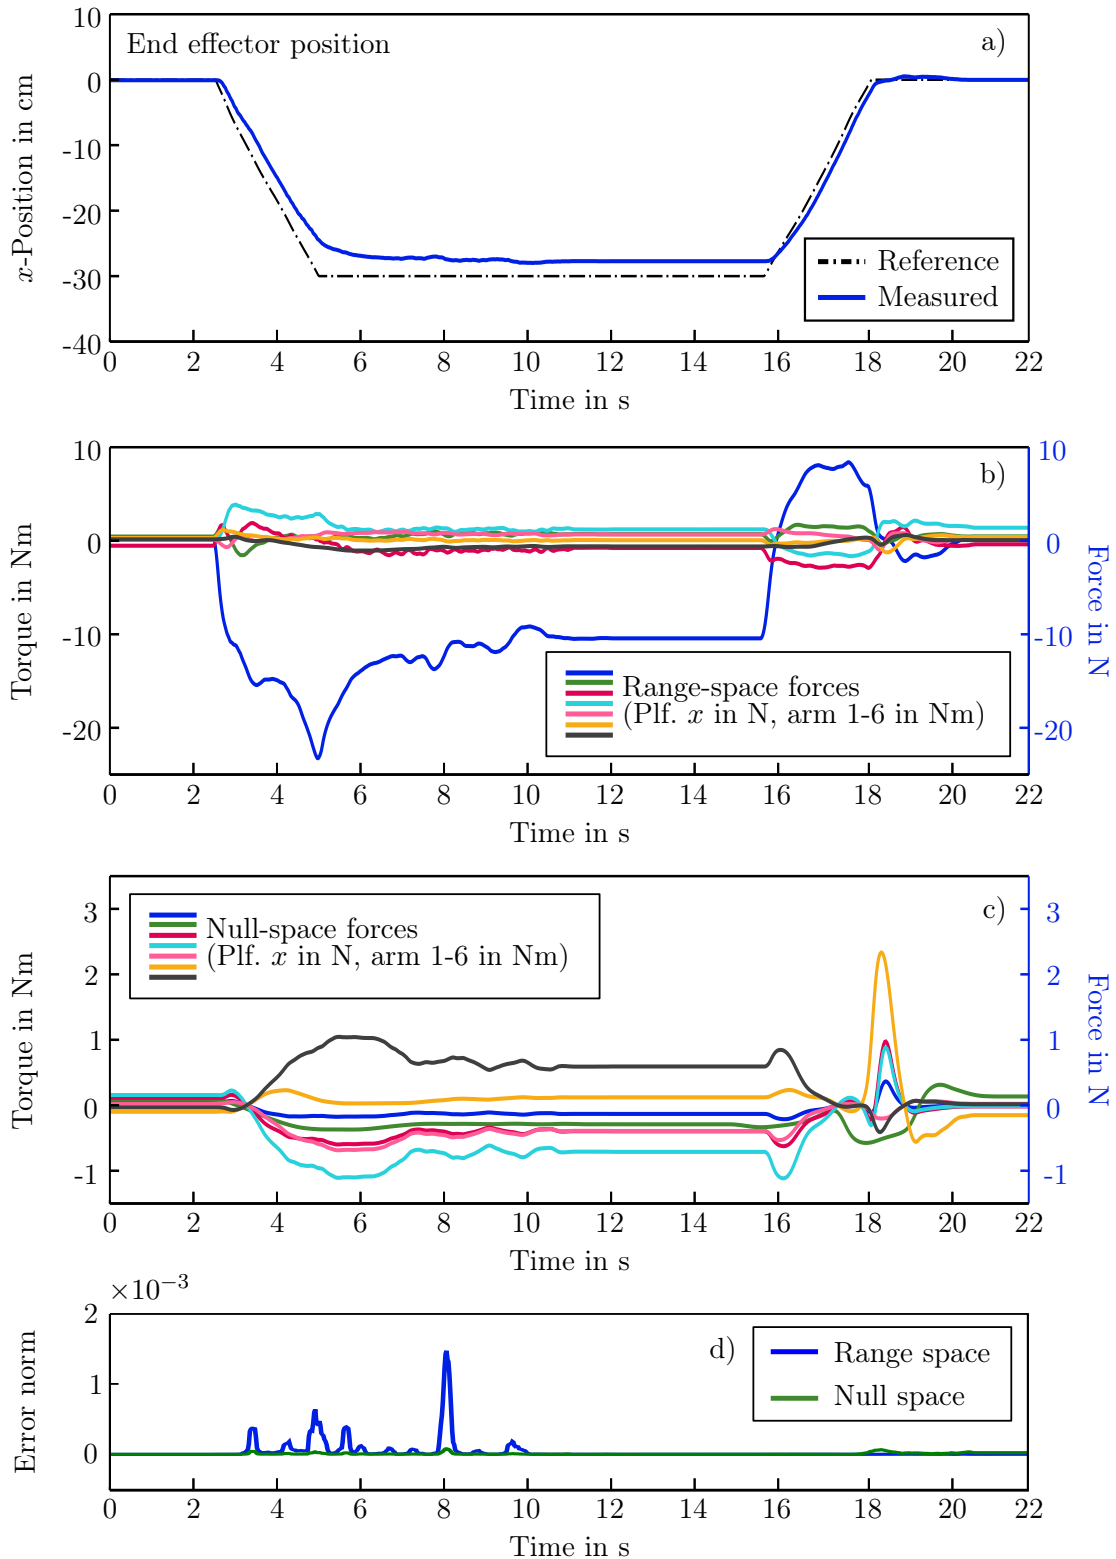


Figure 3.3 Experiment 3.3: Regulation of the x -coordinate of the end effector. a) Commanded and measured x -coordinate of the end effector. b) Commanded joint torques in the range space of the transposed Jacobian matrix, created by optimization of the end effector wrench. Blue line: Platform force in x -direction created by all wheels. Other lines: Joint torques 1, \dots , 6 of the robotic arm. c) Commanded joint torques in the null space of the task Jacobian matrix. d) Norm of the force/torque errors in the range space $\sqrt{J_{\mathcal{R}}}$ and the null space $\sqrt{J_{\mathcal{N}}}$ of the task Jacobian matrix. From [BDO18].




Figure 3.4 Experiment 3.4: Regulation of the end effector during physical interaction with the platform. a) Commanded and measured x -coordinate of the end effector. b) Position of the platform in x -direction. c) Measured joint velocities in the robotic arm. From [BDO18].

behavior of the platform is shown in Figure 3.4. In this experiment, a user actively moves the wheeled subsystem back and forth in the null space of the end effector. Figure 3.4 a) shows the x -coordinate of the end effector, which deviates only a couple of centimeters from the desired position, while the platform is pushed and pulled around ± 30 cm, as shown in Figure 3.4 b). During this physical interaction with the platform, the arm joints compensate for the moving platform, as the measured arm velocities show in Figure 3.4 c). This active compliance property of the locomotion subsystem can be beneficial for a WMM, both for local autonomous manipulation tasks in the presence of obstacles, and also for human-robot interaction scenarios. In the latter scenario, active compliance of the platform can increase the safety of the interaction.

Experiment 3.5 – Physical interaction at the end effector in the presence of joint torque limits: The third whole-body impedance control experiment examines the behavior of the closed-loop system in the case of active constraints. Therefore, the maximum admissible torque of joint 2 (elbow joint) is artificially lowered to 5 Nm (nominally, the data sheet states 18 Nm peak torque for joint 2). An operator performs physical interaction with the end effector. For this experiment, the main task was weighted ten times higher than the null space task by setting all diagonal entries of \mathbf{Q}_N to 0.1. The resulting behavior can be inspected in Figure 3.5 (the interaction period from approx. $1.5\text{ s} < t < 8\text{ s}$ is highlighted in yellow). Figure 3.5 a) shows the deviation of the x -coordinate of the end effector due to the physical interaction. From $1.5\text{ s} < t < 4\text{ s}$, the end effector is pushed, which leads to a positive error. Afterwards, the end effector is pulled by the operator, resulting in a negative deviation. The resulting torques in the constrained joint 2 are shown in Figure 3.5 b). The blue, dot-dashed line depicts the projected impedance torque, which is the torque that would be generated by a classical impedance controller according to $\boldsymbol{\tau}_{\text{imp}} = \mathbf{J}(\mathbf{q})^T \mathbf{F}_{\text{imp}}$. The null-space reference torque is plotted in green dot-dashed. The solid lines show the range space (blue) and null space (green) torques, which result from the optimization problem (3.24)–(3.27). The pink line is the sum of range space and null space torque, which equals the final actuator command torque. First, it can be verified that the optimized actuator torque stays within the admissible bounds of ± 5 Nm. Therefore, the null space torque is allowed to deviate from the reference in order to compensate for the range space torque, which is weighted ten times higher in the optimization. An example for this effect can be observed around $t = 2\text{ s}$, where the range space torque stays close to its reference of up to 12 Nm, at the cost of a large error in the null space torque. A similar effect happens between $t = 4\text{ s}$ and $t = 6\text{ s}$, where the torque hits the limit twice. However, the errors in the range space cannot fully be compensated by deviations of the null space torques in this case, even if effects similar to the aforementioned are observed around $t = 5\text{ s}$. Therefore, the error norm in the range space rises in this time interval, which is shown in Figure 3.5 d). Interestingly, a similar behavior can be observed in the other joints without active joint limits. The torques of joint 3 are shown exemplarily in Figure 3.5 c). Even if this joint does not meet a constraint itself, it shows also a deviation of both null space and range space torque around $t = 2\text{ s}$ and $4\text{ s} < t < 6\text{ s}$. By inspecting the error norms in Figure 3.5 d), it can be verified that the optimization-based impedance controller is able to implement the desired wrench at the end effector to a large extent, at the cost of errors in the null space. Finally, the experiment confirms that the closed-loop system remains stable and converges to the setpoint after the end of the physical interaction, in spite of the joint torque limitations.

3.2.4 Résumé

The implementation of the well-known Cartesian impedance control law as optimization problem is beneficial in different senses. First, different types of redundancy such as kinematic redundancy and actuation redundancy (e.g. the control allocation problem of the platform) can be solved in

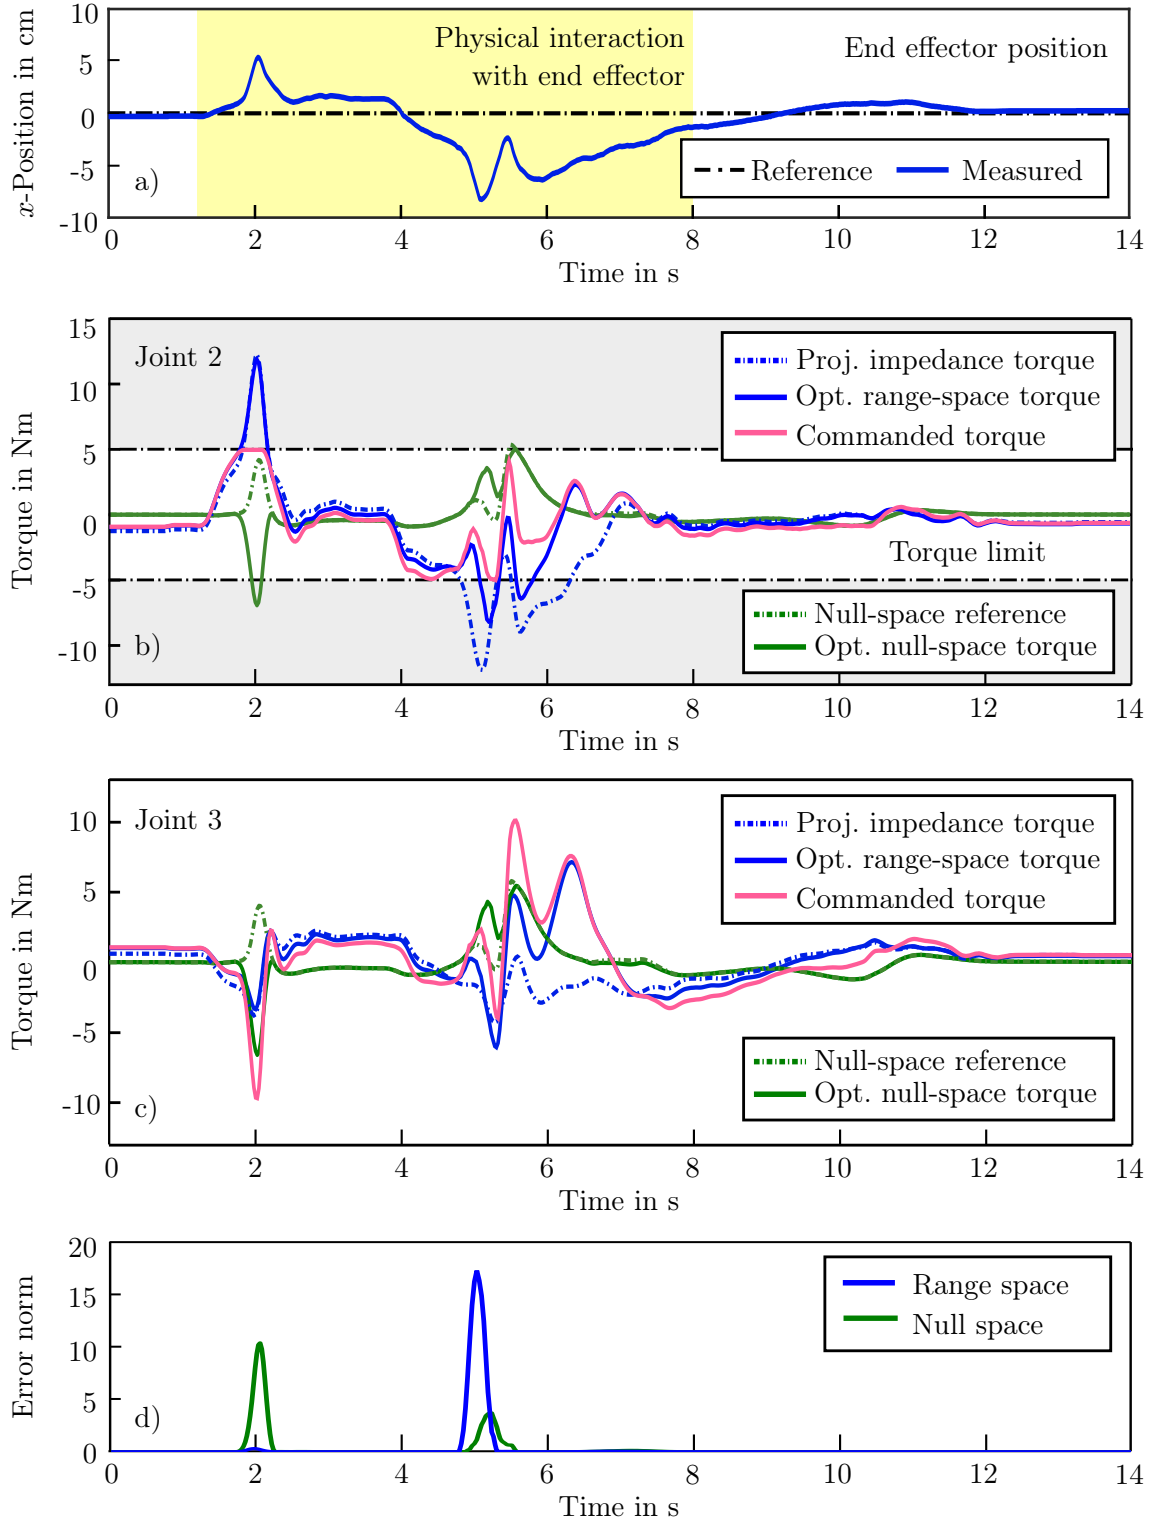


Figure 3.5 Experiment 3.5: Physical interaction at the end effector in the presence of joint torque limits. a) Commanded and measured x -coordinate of the end effector. b) Torques generated at the second joint (constrained to 5Nm). c) Torques generated at the third joint (unconstrained). d) Norm of the force/torque errors in range space and null space. From [BDO18].

a unified way. The set of constraints can be extended arbitrarily, according to the specifications of the system and the scenario. The specific formulation as quadratic optimization problem with linear constraints facilitates fast and online computation of the solution. Last, in the case of active constraints, the sum of null space and range space torques is optimized according to the weight of the task. Due to this fact, it is possible to still fulfill certain tasks despite active limitations, at the cost of the performance of the null space task.

The presented approach enlarges the configuration space of the robotic arm by one DOF, which is actuated by the torque-controlled mobile platform. Clearly, this fact is beneficial for local manipulation tasks, as the degree of redundancy is increased and thus additional tasks can be fulfilled besides the main manipulation task. However, the steering configuration is assumed to be static. The underlying idea is that steering the wheels in rough terrain can cause problems concerning the wheel-ground interaction, such as twisting or side-slipping of the whole rover. Thus, as explained in Section 2.1.4, the platform motion is performed before the manipulation sequence starts in the basic controller setup. In this case, a static steering configuration is selected in advance, which is able to support the respective manipulation task. Particularly, this implies that (2.32) becomes a holonomic equation, with the instantaneously appointed platform motion direction as holonomic coordinate.

One remaining question is how to select the steering configuration, which corresponds to the directions in which the compliance of the wheeled subsystem should be implemented. Possible approaches are strongly application-dependent and range from reactive admittance-based steering parameterizations as it is used in the whole-body impedance controller in [Die15] to workspace-dependent steering, to name but a few. Desirably, the steering configuration is computed automatically using the provided task. This automatic generation of the steering angles is referred to as *implicit coordination*. In the following, an extension of the presented whole-body impedance controller is presented, which incorporates the implicit coordination of the wheels while tracking a desired trajectory.

3.3 Whole-body passivity-based tracking control

As mentioned before, the missing feature for a successful mobile manipulation task is the automatic generation of the steering angles w. r. t. a given task. It is a well-known consequence of Brockett's Theorem that the dynamic model of a robot subject to nonholonomic constraints cannot be stabilized asymptotically to a static equilibrium using time-invariant state feedback control. One possibility is to use time-variant state feedback instead. For example, it is possible to stabilize the dynamic model of a WMM (2.34)–(2.35) to a time-variant trajectory. In the following, a whole-body trajectory tracking control law is presented, which achieves compliant tracking of a joint-level or task-space trajectory for a mobile robot equipped with a robotic arm.⁶ Thereby, the joints of the robot arm as well as the wheel propulsion actuators feature a torque interface, while the steering actuators are controlled on position level. A passivity-based control concept is proposed for the torque-controlled joints, while the position commands for the steering actuators are generated along with the torque control law such that the closed-loop dynamics of the complete system is stabilized to the desired trajectory. This feature qualifies the control law for tasks which cover a larger workspace and thus need platform locomotion (in contrast to the whole-body impedance controller from the previous section, which assumes a static steering configuration).

⁶Up to some changes, the following section corresponds to [Lak+25].

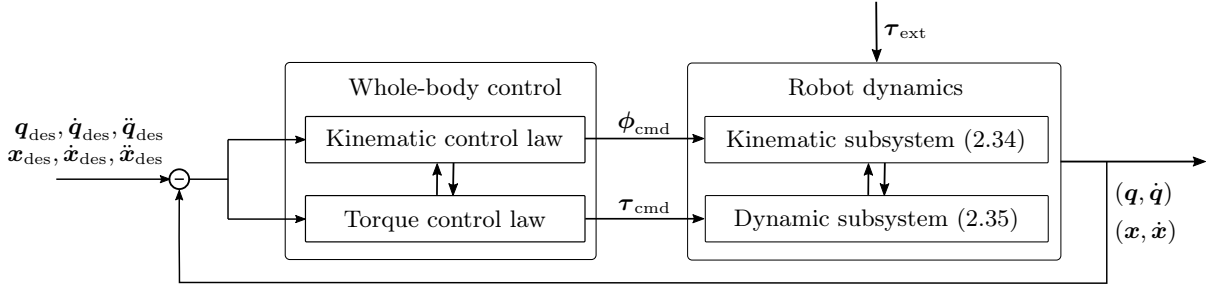


Figure 3.6 Whole-body control concept

A significant advantage of passivity-based control approaches is that the closed-loop system cannot be destabilized by certain unmodeled effects like friction. Thus, these control laws are called *robust* w.r.t. disturbances and modeling uncertainties. This fact predestines passivity-based control for scenarios, where the robot is in contact with an unknown and unstructured environment (indeed, this is always the case for a planetary rover traversing rough terrain). In comparison, control approaches laws using feedback linearization techniques (as e.g. *inverse dynamics* or *computed torque control* schemes) can indeed be destabilized, which is a result of the pre-multiplication of the closed-loop gains with the inverse of the inertia matrix. This fact is theoretically discussed and practically shown e.g. in [Die+21].

The concept of the whole-body tracking controller is sketched in Figure 3.6. A torque control law computes the input for the dynamic subsystem (2.35), while the wheels are oriented according to a kinematic control law acting on the kinematic subsystem (2.34), such that exponential convergence to a given time-varying trajectory is ensured. In the following, the controller is introduced and discussed in joint space, followed by the extension to the task space. Thereby, a switching approach is developed which makes it possible to achieve asymptotic stability of the time-dependent trajectory tracking together with the regulation of a static endpoint at the end of the trajectory. Finally, a detailed discussion of the underlying passivity-based control algorithm is presented, comparing gains and tracking performance at a holonomic robot arm to the well-known PD+ and Slotine-Li controllers.

3.3.1 Joint-level tracking controller design

Let $\mathbf{q}_{\text{des}}(t)$, $\dot{\mathbf{q}}_{\text{des}}(t)$, $\ddot{\mathbf{q}}_{\text{des}}(t)$ be a given (uniformly continuous and bounded) trajectory in the joint space. The tracking error is defined as $\tilde{\mathbf{q}}(t) = \mathbf{q} - \mathbf{q}_{\text{des}}(t)$, with the time derivatives $\dot{\tilde{\mathbf{q}}}(t)$, $\ddot{\tilde{\mathbf{q}}}(t)$. As of now, dependencies on time will be omitted in the notation and only written if strictly necessary for the understanding. The following assumptions are made:

Assumption 5. *In the following considerations, the low-level steering controllers are assumed to be perfect, such that $\boldsymbol{\phi} \equiv \boldsymbol{\phi}_{\text{ref}}$ holds.⁷ Moreover, the desired trajectory is chosen such that the kinematic parameters $\boldsymbol{\phi}_{\text{ref}}$ and their first derivatives $\dot{\boldsymbol{\phi}}_{\text{ref}}$ are bounded.*

Further, the external torque acting on the platform $\boldsymbol{\tau}_{\text{plf,ext}}$ is split into a part which models the nominal wheel-ground interaction forces (in particular all kinds of friction between wheels and

⁷Clearly, the validity of this assumption depends strongly of the hardware capabilities in combination with the choice of the desired trajectory. A discussion of this question part of Chapter 5, based on the experimental results shown in this chapter.

ground), denoted by $\tau_{\text{plf,wgc}}$, and the other external forces acting on the platform (e.g. forces which occur in contact with obstacles or a human operator), $\tau_{\text{plf,con}}$, such that

$$\tau_{\text{plf,ext}} = \tau_{\text{plf,wgc}} + \tau_{\text{plf,con}}. \quad (3.32)$$

Assumption 6. *The nominal wheel-ground contact forces $\tau_{\text{plf,wgi}}$ are known and can be compensated by the platform controller.*

Thus, a compensation term

$$\tau_{\text{wgc}} = \begin{pmatrix} \tau_{\text{plf,wgc}} \\ \mathbf{0} \end{pmatrix} \quad (3.33)$$

can be computed and applied to the combined equations of platform and arm. The remaining external forces $\tau_{\text{con}} = \tau_{\text{ext}} - \tau_{\text{wgc}}$ are

$$\tau_{\text{con}} = \begin{pmatrix} \tau_{\text{plf,con}} \\ \tau_{\text{m,ext}} \end{pmatrix}. \quad (3.34)$$

Control law

In the following, a control law based on the concept of passivity-based control is proposed, which basically means that the kinetic and the potential energy of the open-loop robot dynamics (2.35) are modified by the control law, such that the equilibrium is shifted to the desired trajectory. Therefore, the following torque control law is applied to the dynamic subsystem (2.35):

$$\tau = M\dot{v}_{\text{ref}} + C v_{\text{ref}} + g - \tau_{\text{wgc}} - H(q, \phi_{\text{ref}})s - S(q, \phi_{\text{ref}})^T K \tilde{q}, \quad (3.35)$$

with a uniformly bounded s.p.d. damping matrix $H(q, \phi_{\text{ref}}) \in \mathbb{R}^{(n-r) \times (n-r)}$ and the constant diagonal p.d. stiffness matrix $K \in \mathbb{R}^{n \times n}$. The reference feasible velocity v_{ref} and the kinematic parameters ϕ_{ref} are chosen as the solution of

$$S(q, \phi_{\text{ref}})v_{\text{ref}} = \dot{q}_{\text{des}} - \Omega \tilde{q} \quad (3.36)$$

with the s.p.d. diagonal gain matrix $\Omega \in \mathbb{R}^{n \times n}$, and

$$s = v - v_{\text{ref}} \quad (3.37)$$

denotes the velocity error.

From (2.34) and (3.37), and inserting (3.35) into (2.35), the closed-loop equations in error coordinates are derived as

$$\dot{\tilde{q}} = S(q, \phi_{\text{ref}})(s + v_{\text{ref}}) - \dot{q}_{\text{des}} \quad (3.38)$$

$$M\dot{s} + (C + H)s + S^T K \tilde{q} = \tau_{\text{con}}. \quad (3.39)$$

As a consequence, the closed-loop dynamics is stabilized to the desired trajectory, as the following proposition shows. For this reason, (3.36) is referred to as the *kinematic stability condition* in the following.

Proposition 1.

In the case of free motion, the closed-loop error dynamics (3.38)–(3.39) has a uniformly exponentially stable equilibrium at $(\tilde{q}^, s^*) = (\mathbf{0}, \mathbf{0})$, if the reference values $\phi_{\text{ref}}, v_{\text{ref}}$ satisfy (3.36).*

Proof. The storage function $V_v : (\mathbb{R}^3 \times \mathcal{Q}) \times \mathbb{R}^{n-r} \times \mathbb{R}^+ \rightarrow \mathbb{R}$ is chosen as

$$V_v(\tilde{\mathbf{q}}, \mathbf{s}, t) = \frac{1}{2} \mathbf{s}^T \mathbf{M}(\tilde{\mathbf{q}}, \phi_{\text{ref}}, t) \mathbf{s} + \frac{1}{2} \tilde{\mathbf{q}}^T \mathbf{K} \tilde{\mathbf{q}}, \quad (3.40)$$

which is p.d. in $(\tilde{\mathbf{q}}, \mathbf{s})$ for all t . Note that V_v is time-dependent, which necessitates the use of theorems for non-autonomous systems in the stability analysis. The total time derivative of (3.40) is

$$\dot{V}_v(\tilde{\mathbf{q}}, \mathbf{s}, t) = \mathbf{s}^T \boldsymbol{\tau}_{\text{con}} - \mathbf{s}^T \mathbf{H}(\tilde{\mathbf{q}}, \phi_{\text{ref}}, t) \mathbf{s} + \tilde{\mathbf{q}}^T \mathbf{K} (\mathbf{S}(\tilde{\mathbf{q}}, \phi_{\text{ref}}, t) \mathbf{v}_{\text{ref}} - \dot{\mathbf{q}}_{\text{des}}(t)). \quad (3.41)$$

Inserting (3.36) into (3.41) yields

$$\dot{V}_v(\tilde{\mathbf{q}}, \mathbf{s}, t) = \mathbf{s}^T \boldsymbol{\tau}_{\text{con}} - \mathbf{s}^T \mathbf{H}(\tilde{\mathbf{q}}, \phi_{\text{ref}}, t) \mathbf{s} - \tilde{\mathbf{q}}^T \mathbf{K} \boldsymbol{\Omega} \tilde{\mathbf{q}}. \quad (3.42)$$

For the case of undisturbed motion, i.e., $\boldsymbol{\tau}_{\text{con}} = \mathbf{0}$, the storage function $V_v(\tilde{\mathbf{q}}, \mathbf{s}, t)$ is a strict Lyapunov function for (3.38)–(3.39) with the time derivative \dot{V}_v being negative definite in $(\tilde{\mathbf{q}}, \mathbf{s})$. In particular, this implies boundedness of the state $(\tilde{\mathbf{q}}, \mathbf{s})$, and thus boundedness of \mathbf{q} .

Next, it can be shown that V_v is decrescent. Let $\lambda_i(\cdot)$ denote the i -th eigenvalue of a matrix. From the uniform boundedness of the eigenvalues of the inertia matrix and applying the min-max theorem (see Appendix A.1.4) it follows that there exist time-independent lower and upper bounds for $V_v(\tilde{\mathbf{q}}, \mathbf{s}, t)$:

$$\underline{V}_v(\tilde{\mathbf{q}}, \mathbf{s}) = \frac{1}{2} \inf_{\mathbf{q}, \phi_{\text{ref}}} \min_i \lambda_i(\mathbf{M}(\mathbf{q}, \phi_{\text{ref}})) \|\mathbf{s}\|^2 + \frac{1}{2} \min_i \lambda_i(\mathbf{K}) \|\tilde{\mathbf{q}}\|^2 \quad (3.43)$$

$$\overline{V}_v(\tilde{\mathbf{q}}, \mathbf{s}) = \frac{1}{2} \sup_{\mathbf{q}, \phi_{\text{ref}}} \max_i \lambda_i(\mathbf{M}(\mathbf{q}, \phi_{\text{ref}})) \|\mathbf{s}\|^2 + \frac{1}{2} \max_i \lambda_i(\mathbf{K}) \|\tilde{\mathbf{q}}\|^2. \quad (3.44)$$

Further, the time derivative (3.42) is bounded by

$$\dot{\overline{V}}_v(\tilde{\mathbf{q}}, \mathbf{s}) = - \inf_{\mathbf{q}, \phi_{\text{ref}}} \min_i \lambda_i(\mathbf{H}(\mathbf{q}, \phi_{\text{ref}})) \|\mathbf{s}\|^2 - \min_i \lambda_i(\mathbf{K} \boldsymbol{\Omega}) \|\tilde{\mathbf{q}}\|^2. \quad (3.45)$$

Thus, the inequalities

$$\underline{V}_v(\tilde{\mathbf{q}}, \mathbf{s}) \leq V_v(\tilde{\mathbf{q}}, \mathbf{s}, t) \leq \overline{V}_v(\tilde{\mathbf{q}}, \mathbf{s}) \quad (3.46)$$

and

$$\dot{V}_v(\tilde{\mathbf{q}}, \mathbf{s}, t) \leq \dot{\overline{V}}_v(\tilde{\mathbf{q}}, \mathbf{s}) < 0 \quad (3.47)$$

hold for all $(\tilde{\mathbf{q}}, \mathbf{s}) \neq (\mathbf{0}, \mathbf{0})$. According to the exponential stability theorem (cf. Appendix A.1.2), this proves uniform exponential stability of $(\tilde{\mathbf{q}}^*, \mathbf{s}^*) = (\mathbf{0}, \mathbf{0})$. \square

Remark 5. From (3.42), passivity of the closed-loop system (3.38)–(3.39) w.r.t. the power port $(\boldsymbol{\tau}_{\text{con}}, \mathbf{s})$ can be concluded immediately.

Remark 6. From (3.38) and (3.36), it follows that

$$\dot{\tilde{\mathbf{q}}} = -\boldsymbol{\Omega} \tilde{\mathbf{q}} + \mathbf{S}(\mathbf{q}, \phi_{\text{ref}}) \mathbf{s}, \quad (3.48)$$

which can be interpreted as a stable dynamics of \mathbf{q} with the input \mathbf{s} . It can be concluded that the convergence of $\mathbf{s} \rightarrow \mathbf{0}$ leads to the convergence $\tilde{\mathbf{q}} \rightarrow \mathbf{0}$ as $t \rightarrow \infty$.

Remark 7. Note that the Lyapunov function (3.40) is radially unbounded for $(\tilde{\mathbf{q}}, \mathbf{s}) \in (\mathbb{R}^3 \times \mathcal{Q}) \times \mathbb{R}^{n-r}$. In this sense, the proof of stability holds “globally”.

Discussion of the controller gains

For practical use, a question of interest is how to select the controller gains. Clearly, the values of \mathbf{H} , \mathbf{K} , and $\mathbf{\Omega}$ influence both the convergence properties of the closed-loop error dynamics and the stiffness and damping characteristics of the robot in interaction with the environment. The influence of the gains considering the interaction properties of a holonomic arm discussed in detail in Section 3.4. For a holonomic system, it can be shown that the gains can be used both to parameterize the convergence rate of the tracking errors and to specify a desired contact stiffness and damping. In principle, these considerations can also be used as guidelines for the nonholonomic system at hand. However, another important aspect arises for the considered system in practice, and that is the performance of the underlying steering position controller.

From (3.36), it follows that the value of $\mathbf{\Omega}$ directly influences the effect of the tracking error on the computation of the steering configuration. Thus, selecting higher values for $\mathbf{\Omega}$ results in a more “sensitive” steering behavior. In order to analyze the influence of \mathbf{K} and \mathbf{H} , the control law (3.35) is written in terms of the unconstrained system matrices (the transformations are given in (2.33)). Moreover, the derivative of (3.36) w. r. t. time

$$\dot{\mathbf{S}}\mathbf{v}_{\text{ref}} + \mathbf{S}\dot{\mathbf{v}}_{\text{ref}} = \ddot{\mathbf{q}}_{\text{des}} - \mathbf{\Omega}\ddot{\mathbf{q}} \quad (3.49)$$

is used, and \mathbf{H} is assumed to be chosen such that it can be written as

$$\mathbf{H} = \mathbf{S}^T \bar{\mathbf{H}} \mathbf{S}. \quad (3.50)$$

Sorting the result yields

$$\boldsymbol{\tau} = \mathbf{S}^T \bar{\mathbf{M}} \ddot{\mathbf{q}}_{\text{des}} + \mathbf{S}^T \bar{\mathbf{C}} \dot{\mathbf{q}}_{\text{des}} + \mathbf{S}^T \bar{\mathbf{g}} - \boldsymbol{\tau}_{\text{wgc}} - \mathbf{S}^T (\bar{\mathbf{M}} \mathbf{\Omega} + \bar{\mathbf{H}} \mathbf{\Omega}) \dot{\tilde{\mathbf{q}}} - \mathbf{S}^T (\bar{\mathbf{C}} \mathbf{\Omega} + \bar{\mathbf{H}} \mathbf{\Omega} + \mathbf{K}) \tilde{\mathbf{q}}. \quad (3.51)$$

The first four terms model feed-forward action, gravity compensation, and friction compensation. The fifth term is the effective feedback of the velocity error $\dot{\tilde{\mathbf{q}}}$; thus, the corresponding gain $-\mathbf{S}^T (\bar{\mathbf{M}} \mathbf{\Omega} + \bar{\mathbf{H}} \mathbf{\Omega})$ is called the *effective D-gain*. Accordingly, the feedback gain multiplied by the position error $\tilde{\mathbf{q}}$ is referred to as the *effective P-gain*. It can be observed that both the effective P-gain and D-gain are combinations of matrices of the state space model with the gain matrices $\bar{\mathbf{H}}$, \mathbf{K} , and $\mathbf{\Omega}$, left-multiplied by \mathbf{S}^T . Thus, for a given $\mathbf{\Omega}$, the values of \mathbf{K} and $\bar{\mathbf{H}}$ parameterize the convergence of the system in the direction of the instantaneously feasible velocity. In particular, the effective P-gain $\mathbf{K}_{\text{eff}} = -\mathbf{S}^T (\bar{\mathbf{C}} \mathbf{\Omega} + \bar{\mathbf{H}} \mathbf{\Omega} + \mathbf{K})$ models the stiffness of the closed-loop system in the direction of the instantaneously feasible motion, while the effective D-gain influences the associated damping. In this context, it can be observed that the convergence rate of the Lyapunov function (3.40) depends on \mathbf{H} and the product $\mathbf{K} \mathbf{\Omega}$, see (3.45) and (3.47). This means that weighting \mathbf{K} versus $\mathbf{\Omega}$ while keeping their product constant yields a constant decrease rate of the error state $(\tilde{\mathbf{q}}, \mathbf{s})$. Thereby, the behavior of the closed-loop system can be “shifted” between the steering sensitivity (parametrized by $\mathbf{\Omega}$) and the parameterization of the closed-loop stiffness tangential to the trajectory, influenced by \mathbf{K} .

Summarized, a reasonable tuning strategy for the controller gains is to first choose the value for $\mathbf{\Omega}$, depending on the performance of the low-level steering controllers. Then, select \mathbf{K} and \mathbf{H} such that a desired tracking behavior along the instantaneously feasible motion direction results. A detailed understanding of the influence of \mathbf{K} and \mathbf{H} in the feasible motion direction can be derived from the insights gained in Section 3.4, where the passivity-based tracking control law is applied to a holonomic system.

Extension to kinematic platform control

In practice, most mobile platforms rather feature a kinematic control interface (based on the wheel velocity controllers) than a torque interface which implements torques at the single wheels. One of the reasons for this fact is the need for complicated modeling and compensation of wheel-ground contact forces. Thus, it is worth mentioning that the control algorithm at hand can also be extended to systems without a platform torque interface. In this case, the kinematic stability conditions are solved w. r. t. the kinematic parameters ϕ_{ref} and the platform reference velocity v_{ref} , which are then implemented by the low-level platform velocity controller directly (without computing an actuation torque for the platform).

Assumption 7. *The closed-loop error dynamics of the velocity-controlled platform is sufficiently fast, such that $v \simeq v_{\text{ref}}$ can be assumed.*

For proving the stability under platform velocity control, the dynamic equations of the mobile manipulator (2.34)–(2.35), as well as the kinematic stability condition (3.36), are written w. r. t. both platform and arm variables:

1. Kinematic subsystem (2.34):

$$\begin{pmatrix} \dot{\xi} \\ \dot{\eta} \end{pmatrix} = \begin{pmatrix} \mathbf{R}_z(\theta)\Sigma(\phi_{\text{ref}}) & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \nu_{\text{ref}} \\ \dot{\eta} \end{pmatrix}, \quad (3.52)$$

2. Dynamic subsystem (2.35):

$$\begin{pmatrix} \mathbf{M}_p & \mathbf{M}_{pm} \\ \mathbf{M}_{pm}^T & \mathbf{M}_m \end{pmatrix} \begin{pmatrix} \dot{\nu}_{\text{ref}} \\ \ddot{\eta} \end{pmatrix} + \begin{pmatrix} \mathbf{C}_p & \mathbf{C}_{pm} \\ \mathbf{C}_{mp} & \mathbf{C}_m \end{pmatrix} \begin{pmatrix} \nu_{\text{ref}} \\ \dot{\eta} \end{pmatrix} + \begin{pmatrix} \mathbf{g}_p \\ \mathbf{g}_m \end{pmatrix} = \begin{pmatrix} \boldsymbol{\tau}_{\text{plf}} \\ \boldsymbol{\tau}_m \end{pmatrix} + \begin{pmatrix} \boldsymbol{\tau}_{\text{plf,ext}} \\ \boldsymbol{\tau}_{m,\text{ext}} \end{pmatrix}, \quad (3.53)$$

3. Kinematic stability condition (3.36):

$$\mathbf{R}_z(\theta)\Sigma(\phi_{\text{ref}})\nu_{\text{ref}} = \dot{\xi}_{\text{des}} - \boldsymbol{\Omega}_p \tilde{\xi} \quad (3.54)$$

$$\dot{\eta}_{\text{ref}} = \dot{\eta}_{\text{des}} - \boldsymbol{\Omega}_m \tilde{\eta}. \quad (3.55)$$

First, the platform subsystem is considered. Defining the platform position error $\tilde{\xi} = \xi - \xi_{\text{des}}$, from the first line of (3.52) with (3.54) follows directly the exponentially stable platform dynamics

$$\dot{\tilde{\xi}} = -\boldsymbol{\Omega}_p \tilde{\xi}. \quad (3.56)$$

Note that the platform velocity controller compensates for external forces acting on the platform, for the dynamic couplings of the arm, as well as for the remaining platform dynamics in (3.53). Second, the arm subsystem is analyzed using the second line of (3.53), assuming free motion of the arm ($\boldsymbol{\tau}_{m,\text{ext}} = \mathbf{0}$):

$$\mathbf{M}_{pm}^T \dot{\nu}_{\text{ref}} + \mathbf{M}_m \ddot{\eta} + \mathbf{C}_{mp} \nu_{\text{ref}} + \mathbf{C}_m \dot{\eta} + \mathbf{g}_m = \boldsymbol{\tau}_m. \quad (3.57)$$

Splitting the torque control law (3.35) in a similar manner results in

$$\boldsymbol{\tau}_m = \mathbf{M}_{pm}^T \dot{\nu}_{\text{ref}} + \mathbf{M}_m \ddot{\eta}_{\text{ref}} + \mathbf{C}_{mp} \nu_{\text{ref}} + \mathbf{C}_m \dot{\eta}_{\text{ref}} + \mathbf{g}_m - \mathbf{K}_m \tilde{\eta} - \mathbf{H}_m \boldsymbol{\sigma}, \quad (3.58)$$

with $\tilde{\eta} = \eta - \eta_{\text{des}}$ and $\boldsymbol{\sigma} = \dot{\eta} - \dot{\eta}_{\text{ref}}$. Note that (3.58) contains a compensation of the inertia couplings between platform and arm. Applying (3.58) to (3.57) yields the closed-loop error dynamics of the arm

$$\mathbf{M}_m \dot{\boldsymbol{\sigma}} + (\mathbf{C}_m + \mathbf{H}_m) \boldsymbol{\sigma} + \mathbf{K}_m \tilde{\eta} = \mathbf{0}. \quad (3.59)$$

Then, uniform exponential stability of the equilibrium $(\tilde{\eta}^*, \dot{\tilde{\eta}}^*) = (\mathbf{0}, \mathbf{0})$ of (3.59) can be straightforwardly shown, following the same line of argumentation as in the case of platform torque control, using the positive definite Lyapunov function

$$V_\eta(\tilde{\eta}, \sigma, t) = \frac{1}{2} \sigma^T \mathbf{M}_m \sigma + \frac{1}{2} \tilde{\eta}^T \mathbf{K}_m \tilde{\eta}. \quad (3.60)$$

With (3.55), the total time derivative is computed as

$$\dot{V}_\eta = -\sigma^T \mathbf{H}_m \sigma - \tilde{\eta}^T \mathbf{K}_m \Omega_m \tilde{\eta}, \quad (3.61)$$

which is negative definite in $(\tilde{\eta}, \sigma)$. As

$$\sigma = \dot{\tilde{\eta}} + \Omega_m \tilde{\eta}, \quad (3.62)$$

the derivative is clearly also negative definite in $(\tilde{\eta}, \dot{\tilde{\eta}})$. This shows uniform global asymptotic stability of the equilibrium $(\tilde{\eta}^*, \dot{\tilde{\eta}}^*) = (\mathbf{0}, \mathbf{0})$. Exponential stability can be shown analogously to the torque-controlled case.

Remark 8. Since the platform velocity controller automatically compensates for external forces acting on the platform, only free motion of the robot arm has to be assumed for the proof of stability, whereas the wheel-ground contact forces do not need to be known or compensated. On the one hand, this eliminates a source of model uncertainties (recall that a perfect compensation of wheel-ground contact forces is assumed in the torque-controlled case, which is hard to implement in practice). On the other hand, it means that the platform does not react to other interaction forces, i.e., it does not feature active compliance.

To demonstrate the effect of using platform velocity control compared to platform torque control, the results of an experimental comparison are shown in Section 3.3.4.

3.3.2 Task-space tracking control

In the following, the control approach is extended to the task-space tracking case. The task velocities $\dot{\mathbf{x}}$ can be expressed w.r.t. the constraint-consistent velocities \mathbf{v} using

$$\dot{\mathbf{x}} = \mathbf{J}(\mathbf{q})\dot{\mathbf{q}} = \mathbf{J}(\mathbf{q})\mathbf{S}(\mathbf{q}, \phi)\mathbf{v}. \quad (3.63)$$

Control law

The control law is chosen as

$$\boldsymbol{\tau} = \mathbf{M}\dot{\mathbf{v}}_{\text{ref}} + \mathbf{C}\mathbf{v}_{\text{ref}} + \mathbf{g} - \boldsymbol{\tau}_{\text{wgc}} - \mathbf{H}\mathbf{s} - \mathbf{S}^T \mathbf{J}^T \mathbf{K}_x \tilde{\mathbf{x}} \quad (3.64)$$

with the constant, diagonal, p.d. matrix $\mathbf{K}_x \in \mathbb{R}^{m \times m}$. Using (2.35), (3.37), (3.63), and (3.64), the closed-loop error dynamics in task space are computed as

$$\dot{\tilde{\mathbf{x}}} = \mathbf{J}\mathbf{S}(\phi_{\text{ref}})(\mathbf{s} + \mathbf{v}_{\text{ref}}) - \dot{\mathbf{x}}_{\text{des}}, \quad (3.65)$$

$$\mathbf{M}\dot{\mathbf{s}} + (\mathbf{C} + \mathbf{H})\mathbf{s} + \mathbf{S}^T \mathbf{J}^T \mathbf{K}_x \tilde{\mathbf{x}} = \boldsymbol{\tau}_{\text{con}}. \quad (3.66)$$

The reference trajectory $(\phi_{\text{ref}}, \mathbf{v}_{\text{ref}})$ is chosen according to the task-space kinematic stability condition

$$\mathbf{J}(\mathbf{q})\mathbf{S}(\mathbf{q}, \phi_{\text{ref}})\mathbf{v}_{\text{ref}} = \dot{\mathbf{x}}_{\text{des}} - \Omega_x \tilde{\mathbf{x}} \quad (3.67)$$

with the constant, diagonal, p.d. gain matrix $\Omega_x \in \mathbb{R}^{m \times m}$.

Proposition 2.

For a non-redundant robotic system in free motion, i. e., $\boldsymbol{\tau}_{\text{con}} = \mathbf{0}$, the equilibrium $(\tilde{\mathbf{x}}^*, \mathbf{s}^*) = (\mathbf{0}, \mathbf{0})$ of the closed-loop equations (3.65)–(3.66) is uniformly locally exponentially stable if the reference trajectory is chosen according to (3.67).

Proof. Choosing the Lyapunov function candidate

$$V_x(\tilde{\mathbf{x}}, \mathbf{s}, t) = \frac{1}{2} \mathbf{s}^T \mathbf{M} \mathbf{s} + \frac{1}{2} \tilde{\mathbf{x}}^T \mathbf{K}_x \tilde{\mathbf{x}} \quad (3.68)$$

and assuming the case of free motion, i. e., $\boldsymbol{\tau}_{\text{con}} = \mathbf{0}$, yields the time derivative

$$\dot{V}_x(\tilde{\mathbf{x}}, \mathbf{s}, t) = -\mathbf{s}^T \mathbf{H} \mathbf{s} + \tilde{\mathbf{x}}^T \mathbf{K}_x (\mathbf{J} \mathbf{S}(\phi_{\text{ref}}) \mathbf{v}_{\text{ref}} - \dot{\mathbf{x}}_{\text{des}}). \quad (3.69)$$

Inserting (3.67) into (3.69) results in

$$\dot{V}_x(\tilde{\mathbf{x}}, \mathbf{s}, t) = -\mathbf{s}^T \mathbf{H} \mathbf{s} - \tilde{\mathbf{x}}^T \mathbf{K}_x \boldsymbol{\Omega}_x \tilde{\mathbf{x}}, \quad (3.70)$$

which is negative definite in $(\tilde{\mathbf{x}}, \mathbf{s})$, similar to (3.42). If the robot is non-redundant, i. e., $m = n$, the mapping $\mathbf{x} = \mathbf{h}(\mathbf{q})$ is locally 1:1 and invertible. Then, (3.68) is positive definite for the complete state of the robot, thus it is indeed a Lyapunov function for the system (3.65)–(3.66). Proceeding analogously to Section 3.3.1, it can be shown that V_x is decrescent by computing the respective time-invariant bounds for (3.68) and (3.69). Consequently, the desired task-space trajectory $(\tilde{\mathbf{x}}, \mathbf{s}) = (\mathbf{0}, \mathbf{0})$ can be proven to be uniformly exponentially stable. \square

Moreover, from (3.65) and (3.67) follows that

$$\dot{\tilde{\mathbf{x}}} = \mathbf{J} \mathbf{S} \mathbf{s} - \boldsymbol{\Omega}_x \tilde{\mathbf{x}}, \quad (3.71)$$

and thus $(\tilde{\mathbf{x}}^*, \mathbf{s}^*) = (\mathbf{0}, \mathbf{0}) \rightarrow (\tilde{\mathbf{x}}^*, \dot{\tilde{\mathbf{x}}}^*) = (\mathbf{0}, \mathbf{0})$. Note that for task-space coordinates, stability statements can in general only be made locally (e. g. due to task space singularities or orientation tasks defined in $\text{SO}(3)$).

Redundancy resolution

If the robot is redundant, i. e., $m < n$, (3.68) is only positive semi-definite and thus, no statements about stability can be made from (3.70). In this case, the null space of the robot has to be properly defined in order to guarantee convergence of the redundant DOF. One possibility is to utilize a quadratic optimization approach, comparable to the optimization-based impedance controller presented in Section 3.2.2. However, as these approaches can be computationally expensive and thus might not be applicable under the hardware restrictions of a space mission, an alternative and simpler solution how to resolve the redundancy shall be presented, which also allows an analytical proof of asymptotic stability.⁸

Therefore, the robot coordinates \mathbf{q} are split into $\mathbf{q}_x \in \mathbb{R}^m$ (m coordinates that are used to fulfill the task such that $\mathbf{x} = \mathbf{h}(\mathbf{q}_x)$ is locally 1:1 and invertible) and the remaining $\mathbf{q}_y \in \mathbb{R}^{(n-m)}$ coordinates, which are used to define the null space. A desired joint-space trajectory for the null space coordinates is given as $\mathbf{q}_{y,\text{des}}(t), \dot{\mathbf{q}}_{y,\text{des}}(t), \ddot{\mathbf{q}}_{y,\text{des}}(t)$. Then, also the reference velocities are split:

$$\mathbf{v}_{\text{ref}} = \begin{pmatrix} \mathbf{v}_{x,\text{ref}} \\ \dot{\mathbf{q}}_{y,\text{ref}} \end{pmatrix}, \quad (3.72)$$

⁸Note that the presented approach is closely related to the concept of *task augmentation* presented in [Ser89], where the task space is augmented by kinematic functions in order to generate a quadratic and invertible task Jacobian matrix.

and thus

$$\mathbf{s} = \begin{pmatrix} \mathbf{v} - \mathbf{v}_{x,\text{ref}} \\ \dot{\mathbf{q}}_y - \dot{\mathbf{q}}_{y,\text{ref}} \end{pmatrix}. \quad (3.73)$$

An additional torque input is added to the control law, which acts only on the null space coordinates:

$$\boldsymbol{\tau} = \mathbf{M}\dot{\mathbf{v}}_{\text{ref}} + \mathbf{C}\mathbf{v}_{\text{ref}} + \mathbf{g} - \boldsymbol{\tau}_{\text{wgc}} - \mathbf{H}\mathbf{s} - \mathbf{S}^T \mathbf{J}^T \mathbf{K}_x \tilde{\mathbf{x}} - \boldsymbol{\tau}_y \quad (3.74)$$

with

$$\boldsymbol{\tau}_y = \begin{pmatrix} \mathbf{0} \\ \mathbf{G}\tilde{\mathbf{q}}_y \end{pmatrix} \quad (3.75)$$

with $\mathbf{G} \in \mathbb{R}^{(n-m) \times (n-m)}$ being s.p.d. and constant, and $\tilde{\mathbf{q}}_y = \mathbf{q}_y - \mathbf{q}_{y,\text{des}}(t)$. Now, a quadratic potential for the null space coordinates is added to the Lyapunov function candidate in order to make it positive definite:

$$V_{x,y}(\mathbf{s}, \mathbf{q}, t) = \frac{1}{2} \mathbf{s}^T \mathbf{M} \mathbf{s} + \frac{1}{2} \tilde{\mathbf{x}}^T \mathbf{K}_x \tilde{\mathbf{x}} + \frac{1}{2} \tilde{\mathbf{q}}_y^T \mathbf{G} \tilde{\mathbf{q}}_y. \quad (3.76)$$

With $\boldsymbol{\tau}_{\text{con}} = \mathbf{0}$, the time derivative becomes

$$\begin{aligned} \dot{V}_{x,y} = & -\mathbf{s}^T \mathbf{H} \mathbf{s} + \tilde{\mathbf{x}}^T \mathbf{K}_x (\mathbf{J} \mathbf{S} \mathbf{v}_{\text{ref}} - \dot{\mathbf{x}}_{\text{des}}) \\ & - \mathbf{s}^T \boldsymbol{\tau}_y + \tilde{\mathbf{q}}_y^T \mathbf{G} \dot{\tilde{\mathbf{q}}}_y. \end{aligned} \quad (3.77)$$

with

$$\mathbf{s}^T \boldsymbol{\tau}_y = (\dot{\mathbf{q}}_y - \dot{\mathbf{q}}_{y,\text{ref}})^T \mathbf{G} \tilde{\mathbf{q}}_y. \quad (3.78)$$

In order to make the time derivative negative definite, the kinematic stability condition (3.67) is used again. It yields $m < n$ equations, which need to be solved for n variables. Thus, $n - m$ additional equations have to be specified to create a unique solution. These side conditions are now defined via the null-space coordinates, such that the whole $n \times n$ system of equations becomes

$$\mathbf{J}(\mathbf{q}_x, \mathbf{q}_y) \mathbf{S}(\phi_{\text{ref}}) \mathbf{v}_{\text{ref}} = \dot{\mathbf{x}}_{\text{des}} - \boldsymbol{\Omega}_x \tilde{\mathbf{x}} \quad (3.79)$$

$$\dot{\mathbf{q}}_{y,\text{ref}} = \dot{\mathbf{q}}_{y,\text{des}} - \mathbf{L} \tilde{\mathbf{q}}_y \quad (3.80)$$

with $\mathbf{L} \in \mathbb{R}^{(n-m) \times (n-m)}$ p.d. and diagonal. This renders (3.77) negative definite:

$$\dot{V}_{x,y} = -\mathbf{s}^T \mathbf{H} \mathbf{s} - \tilde{\mathbf{x}}^T \mathbf{K}_x \boldsymbol{\Omega}_x \tilde{\mathbf{x}} - \tilde{\mathbf{q}}_y^T \mathbf{G} \mathbf{L} \tilde{\mathbf{q}}_y. \quad (3.81)$$

One can directly conclude exponential stability of a well-defined configuration trajectory

$$\mathbf{q}^*(t) = \begin{pmatrix} \mathbf{q}_x(t) \\ \mathbf{q}_y(t) \end{pmatrix} = \begin{pmatrix} \mathbf{h}^{-1}(\mathbf{x}_{\text{des}}(t)) \\ \mathbf{q}_{y,\text{des}}(t) \end{pmatrix}. \quad (3.82)$$




Figure 3.7 Concept of the two-phase proof of stability. From [Lak+25].

3.3.3 End of trajectory: regulation of a static equilibrium

For nonholonomic systems, proofs of stability for smooth trajectory-following controllers are only valid as long as the trajectory does not stop, such that it would become a regulation problem instead [Bro+83]. Therefore, time-variant (see e.g. [TdM96]) or switching (see e.g. [ODV02]) control laws are commonly utilized in order to regulate nonholonomic systems to a constant desired endpoint. Another approach can be found in [GH14], which is based on stabilizing the dynamics to a neighborhood of a given posture by modifying the system output function.

In contrast to the aforementioned approaches, the idea here is to switch the system characteristics artificially when the trajectory stops. This is achieved by “locking” the kinematic parameters in the close vicinity of the desired goal, and thus creating an artificial holonomic system at the cost of losing certain DOF temporarily. For that reason, the following approach is only applicable in the case of task-space control, where the convergence to the desired endpoint can still be realized due to the inherent kinematic redundancy. In the following, $m = n - r$ is assumed, to ensure that the robot kinematics is non-redundant at the end of the trajectory. Note that it is also possible to extend the following considerations $m < n - r$, e.g. by using redundancy resolution techniques as described before. The stabilization of the endpoint of a desired task-space trajectory can be achieved by the two-phase approach as depicted in Figure 3.7. The first phase consists of the trajectory tracking already addressed in Section 3.3.2. The system is driven from the initial state $x_{1,start}$ at time $t = t_{1,start}$ to a state $x(t_{1,end}) = x_{1,end}$. At the time when the trajectory stops, $\dot{x}_{des} \equiv 0$. At that time, the second phase (“regulation”) begins, and the terminal state of the first phase becomes the initial condition for phase two. Note that the proof of uniform exponential stability presented in Section 3.3.2 guarantees bounded states $x_{1,end}$, so that the proper initialization $x_{2,start} = x_{1,end}$ with $t_{1,end} = t_{2,start}$ for the proof of asymptotic stability in phase 2 has already been ensured. Thereby, ϕ_{ref} is set to zero such that a holonomic system is created (with an arbitrary but static $\phi_{ref} \equiv \bar{\phi}_{ref}$). The idea is sketched in Figure 3.8.

For the regulation phase, two reasonable possibilities to choose the control law are examined in the following. The first possibility is to keep the kinematic stability condition and solve $JS(q, \bar{\phi}_{ref})v_{ref} = -\Omega_x \tilde{x}$ for v_{ref} only, such that $\dot{V}_x = -sHs - \tilde{x}K_x\Omega_x\tilde{x}$. This shows asymptotic stability of the origin $(\tilde{x}^*, s^*) = (0, 0)$. As shown before, this implies stability of $(\tilde{x}^*, \dot{\tilde{x}}^*) = (0, 0)$. Note that switching to a static configuration $\dot{\phi}_{ref} = 0$ reduces the number of active DOF of the artificial holonomic systems from n to $n - r$. Thus, using the redundancy resolution strategy described in Section 3.3.2, the number of null space coordinates q_y has to be reduced such that the resulting number of equations equals $n - r$, i.e., r former null space coordinates are added to fulfill the task, the remaining $n - r - m$ coordinates form the new null space. In particular, for $m = n - r$, no redundancy resolution has to be performed after the end of trajectory.

The second possibility is to set $v_{ref} = 0$, yielding a classical impedance control law for the regulation phase. In this case, asymptotic convergence to the set $\mathcal{M} := \{s = v = 0, \tilde{x} = 0\}$ can be concluded using the invariance principle of LaSalle. For non-redundant systems, this




Figure 3.8 Sketch of trajectory tracking in the task space. At the end of the trajectory, the steering configuration is locked ($\dot{\phi}_{ref} = 0$). In this way, a *virtual holonomic joint* is created. This turns the system into an artificially holonomic one, at the cost of losing two DOF. From [Lak+25].

implies convergence of all system states to the set $\mathcal{M}_r = \{\dot{q} = 0, q = q^* \mid x(q^*) = x_{des}\}$. Of course, convergence of the regulation task depends on the workspace of the remaining “artificial” holonomic robot. It needs to be guaranteed that the initial state of the robot is sufficiently close to the desired state at the end of the trajectory. This can be supported by a) choosing a suitable trajectory, b) choosing the controller gains sufficiently high, and c) ensuring that phase 1 is long enough. In this case, the aforementioned proof of exponential stability guarantees the fulfillment of the requirement.

3.3.4 Experiments

First, the joint-level trajectory tracking experiments are performed only with the platform subsystem, using the platform torque interface. Tracking experiments in task space using the complete rover follow, including the regulation of a static equilibrium at the end of the trajectory. During all experiments, the platform velocity and position are computed from the wheel odometry. Wheel-ground contact forces are not considered, such that $\tau_{wgc} = 0$ in the control law. While the experiments in this section analyze the tracking performance mainly w. r. t. different platform control modes, the impact of different gains choices is experimentally investigated at the example of a purely holonomic robot arm in Section 3.4.




Figure 3.9 LRU2 driving on volcanic soil during Experiment 3.6.

Tracking a joint-space trajectory with the platform subsystem

Considering only the platform subsystem, the analytical solution of the kinematic stability condition (3.36) is

$$v_{\text{ref}} = \pm \sqrt{(v_{y,\text{des}} - \Omega_y {}^b\tilde{y})^2 + (v_{x,\text{des}} - \Omega_x {}^b\tilde{x})^2} \quad (3.83)$$

$$\kappa_{\text{ref}} = \frac{\omega_{\text{des}} - \Omega_{\theta} \tilde{\theta}}{v_{\text{ref}}} \quad (3.84)$$

$$\beta_{\text{ref}} = \arctan \left(\frac{v_{y,\text{des}} - \Omega_y {}^b\tilde{y}}{v_{x,\text{des}} - \Omega_x {}^b\tilde{x}} \right) \quad (3.85)$$

with $({}^b\tilde{x}, {}^b\tilde{y}, \tilde{\theta})$ denoting the coordinate errors in rover frame, and $\Omega_{x/y/\theta} > 0$ the scalar gains in the respective coordinate. When computing β_{ref} using the four-quadrant inverse tangent (e. g. the Matlab `atan2` function), the positive sign of the reference velocity v_{ref} must be selected. The wheel propulsion torques are distributed according to (3.12)–(3.14), such that the reduced platform wrench $\tau_{\text{plf}}^r = \mathbf{S}(\mathbf{q}, \phi_{\text{ref}})^T \tau_{\text{plf}}$ is generated by the platform torque interface. The corresponding weights are chosen as $\mathbf{Q}_f^r = \text{diag}(v_{w,i}) + 0.1\mathbf{I}$, $\mathbf{Q}_e^r = 100$.

Experiment 3.6 – Tracking a sinusoidal trajectory with the platform on volcanic soil: A sinusoidal example trajectory in (x, y, θ) is commanded to the platform with a constant reference velocity in x -direction of 0.05 m/s. The trajectory and the platform path in the x - y -plane are depicted in Figure 3.10. In Figure 3.11, reference and measured position are plotted, followed by the error values for the position coordinates. The controller gains are chosen as $\mathbf{\Omega} = \text{diag}(0.1, 0.1, 0.1)$ 1/s and $\mathbf{K} = \text{diag}(1500 \text{ N/m}, 1500 \text{ N/m}, 150 \text{ Nm/rad})$. The (scalar) damping gain H is computed w. r. t. \mathbf{K} using the system mass of $m = 42 \text{ kg}$ and the damping ratio $\zeta = 0.5$ according to $H = 2\zeta\sqrt{\mathbf{S}^T \mathbf{K} m}$. By inspecting Figure 3.11 it can be verified that the trajectory is followed properly, and the errors are in the magnitude of 1–2 cm. The errors result mainly from the




Figure 3.10 Experiment 3.6: Tracking a sinusoidal trajectory with the platform on volcanic soil. Desired (black dot-dashed) and measured (blue) trajectory for the mobile platform in (x, y) -coordinates. From [Lak+25].




Figure 3.11 Experiment 3.6: Tracking a sinusoidal trajectory with the platform on volcanic soil. a) Desired and measured trajectory in (x, y, θ) . b) Errors in (x, y, θ) directions. From [Lak+25].

Gain	Unit	Experiment 3.7 (Tracking)	Experiment 3.8 (Contact)
Ω_x	1/s	diag(0.7, 0.5, 0.5, 0.1, 0.1, 0.1)	diag(0.7, 0.5, 0.5, 0.3, 0.1, 0.1)
$\mathbf{K}_{\text{trans}}$	N/m	diag(750, 750, 750)	diag(525, 525, 525)
\mathbf{K}_{rot}	Nm/rad	diag(75, 75, 75)	diag(52.5, 17.5, 17.5)
ζ	–	(0.7, 0.7, 0.7, 0.7, 0.7, 0.7)	(0.7, 0.7, 0.7, 0.7, 0.7, 0.7)
\mathbf{L}	1/s	diag(0.1, 0.1, 0.1)	diag(0.1, 0.1, 0.1)
\mathbf{G}	Nm/rad	diag(50, 50, 50)	diag(3, 3, 3)

Table 3.1 Gains used for the experiments 3.7–3.8. The stiffness gain \mathbf{K} is given as a combination of the translational stiffness $\mathbf{K}_{\text{trans}}$ and the rotational stiffness \mathbf{K}_{rot} . From [Lak+25].

unmodeled wheel-ground interaction forces, as the experiment is performed on volcanic soil, see Figure 3.9. Thus, parts of the force transmitted through the wheels are used for soil deformation instead of platform locomotion.

Note that the small tracking errors which are observed in Experiment 3.6 occur at comparatively high stiffness values of $\mathbf{K} = \text{diag}(1500 \text{ N/m}, 1500 \text{ N/m}, 150 \text{ Nm/rad})$. However, the stiffness values are reduced for the following task-space tracking experiments. In this case, the platform torque controller features high errors and significant delays, among others due to the limitations of the hardware that are listed in Section 2.1.3.

Task-space trajectory tracking with the complete mobile manipulator

The following experiments show the performance of the whole-body tracking controller in task space, particularly comparing the use of the platform torque interface with a velocity-controlled platform. Additionally, the 6-DOF arm mounted on the rover platform is now included in the experiments. Therefore, a 6-DOF Cartesian trajectory is assigned to the end effector, which can be inspected in Figure 3.12. The desired Cartesian DOF which do not appear in the figure are kept constant at their initial values. Note that the command to stop the trajectory is manually triggered by the operator. Thus, the final time and value of the trajectories vary between the datasets. The end of the trajectory is discussed in Section 3.3.3.

Experiment 3.7 – Cartesian tracking control of the end effector: The goal of the experiment is to compare the performance of the whole-body controller using the platform torque interface and the platform velocity interface. Therefore, the same trajectory is commanded twice, the results are compared pairwise. The controller gains used for Experiment 3.7 can be inspected in Table 3.1. Again, the damping matrix \mathbf{H} is computed depending on the configuration, using the *Double Diagonalization* approach from [Alb+03]. The corresponding damping factors are given by ζ .

During the tracking phase, the kinematic redundancy is resolved following the simple task augmentation approach described in Section 3.3.2. Thereby, the positions of the last three joints of the arm (joints number 4, 5, 6) are used as additional task coordinates \mathbf{q}_y . After the trajectory stops, only the position of joint 6 remains as null-space condition, while joints 4 and 5 are additionally used to fulfill the Cartesian task.

In Figure 3.13, a 3D plot of both the platform trajectory and the trajectory of the end effector is shown. The initial frame is given by the position of the platform at $t = 0 \text{ s}$. Figure 3.14 shows a comparison of the resulting tracking errors for both platform torque control and platform velocity control. It can be observed that at the commanded high velocities (translational velocity of the

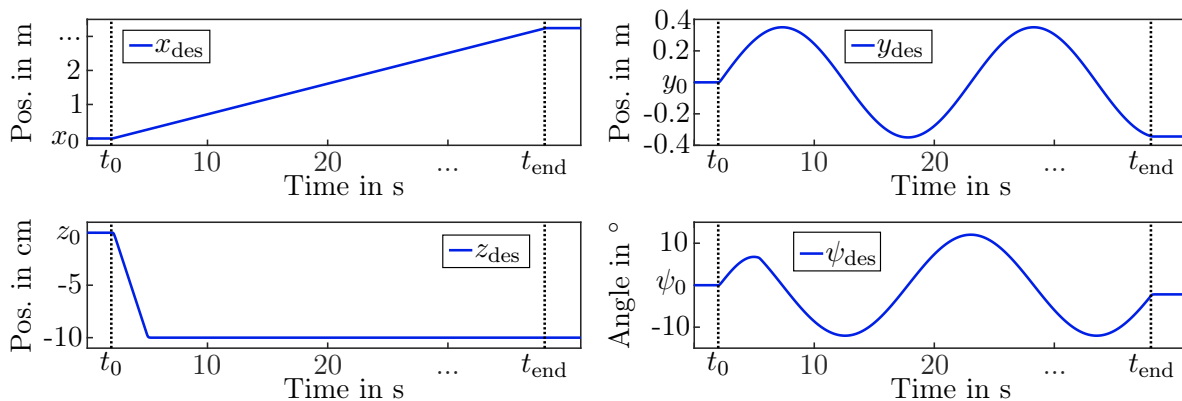


Figure 3.12 Experiment 3.7: Commanded task-space trajectory for the tracking experiment. (x, y, z) denote Cartesian coordinates of the end effector in a fixed initial frame, ψ is the yaw angle corresponding to a yaw-pitch-roll (ψ, θ, ϕ) Euler angle parameterization. The trajectory continues periodically until the end of trajectory is commanded at t_{end} . At that time, all desired values remain constant. From [Lak+25].

end effector up to 0.14 m/s), tracking errors are significant for platform torque control. Notably, the tracking performance is significantly better with platform velocity control. To quantify the tracking performance, the root mean square error (RMSE) for the platform velocity control is computed as 2.2 cm for translational deviations and 4.6° for rotational deviations (parameterized in Euler angles). For the platform torque controller, the translational RMSE is 4.8 cm, the rotational RMSE is 9.9°. The values are visualized in the bar diagram Figure 3.15. Thus, the total tracking error at the end effector is approximately 2.2 times larger for platform torque control than for platform velocity control. This can be explained with the superior performance of the wheel velocity controllers compared to the wheel torque controllers. The performance of wheel velocity control versus wheel torque control can also be verified by inspecting Figure 3.16, which shows the tracking of the reference values for the platform. The position-controlled steering angles are depicted in Figure 3.16 b) and c), which show the tracking of β_{ref} and κ_{ref} . It can be verified that the underlying steering controllers can follow the steering configuration properly, apart from initial re-orientation effects in the first approx. 2 s after the start of the trajectory. Figure 3.16 a) shows the tracking of the reference velocity v_{ref} . It can be observed that the reference velocity is limited to 0.5 m/s. This is a safety feature implemented to protect robot and user. The limit is active in the first 3 s, when the steering configuration is not yet reached and thus, no lateral velocity is created by the wheel torque or velocity controllers. Afterwards, the platform velocity controller is able to follow the reference velocity, while the wheel torque controller does produce significant tracking errors, combined with time delays of up to 2.5 s (cf. Figure 3.16). Again, this can be explained by the missing compensation of wheel-ground interaction effects, and the lack of performance of the low-level wheel torque controllers combined with lower stiffness gains \mathbf{K}_x compared to Experiment 3.6. The influence of the errors produced by the wheel torque controllers can also be inspected in Figure 3.17. The upper two plots a) and b) show the commanded torques and the resulting joint positions for the first three arm joints. While in the kinematic control case the arm positions remain more or less constant after the initial phase (which also contains the 10 cm ramp in z -direction), it can be verified that the commanded joint torques for the arm partially compensate for the platform errors due in the case of platform torque control. The effect is also illustrated in the video that accompanies [Lak+25].

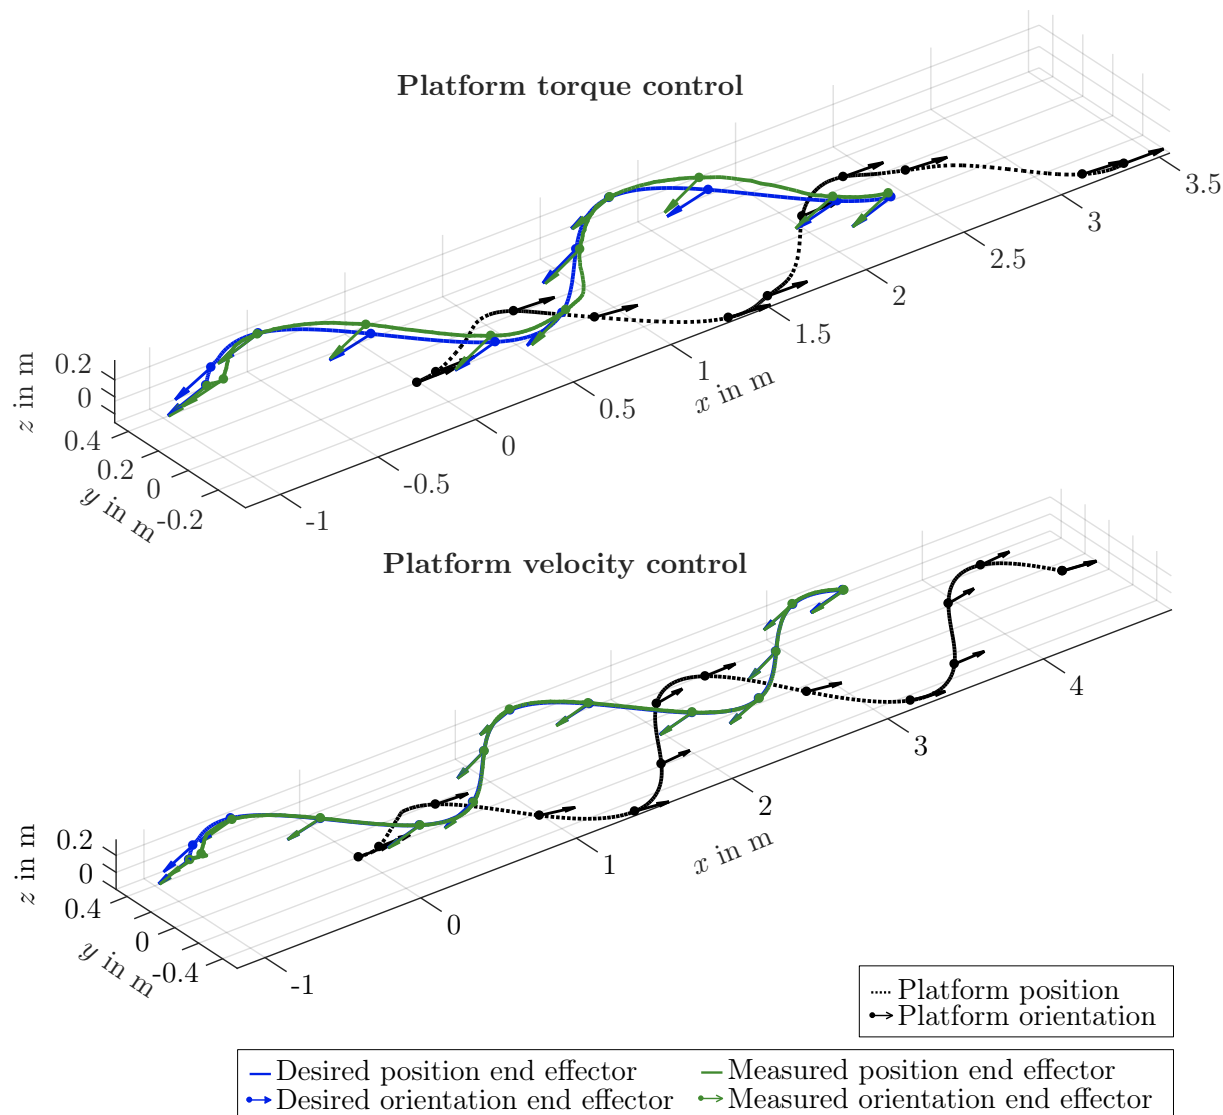


Figure 3.13 Experiment 3.7: Trajectory tracking. 3D plot of the platform and the position and orientation of the end effector. The position of the platform center at $t = 0$ s is $(x, y, z) = (0, 0, 0)$. From [Lak+25].

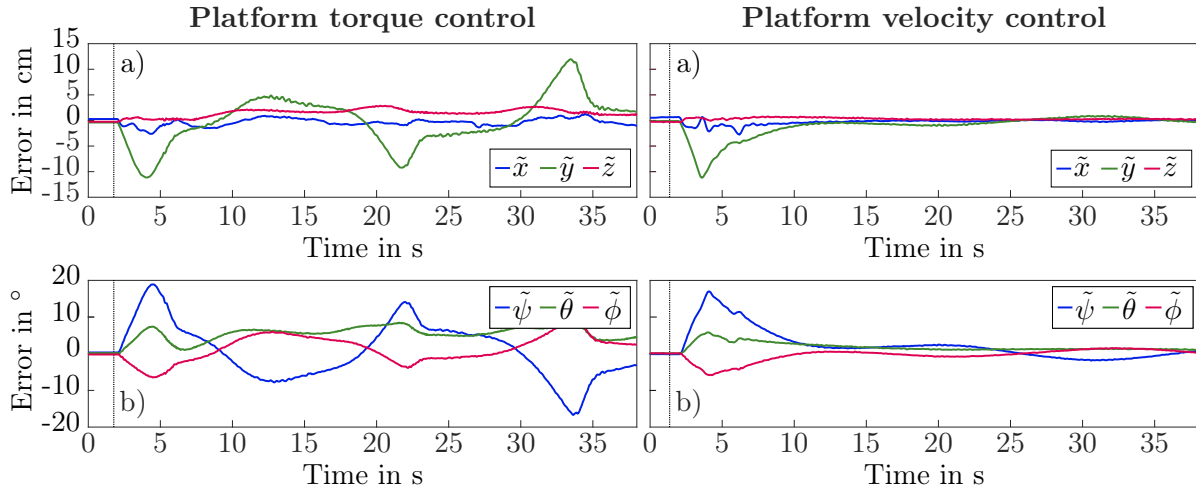


Figure 3.14 Experiment 3.7: Comparison of the tracking errors for platform torque and velocity control. a) Translational errors ($\tilde{x}, \tilde{y}, \tilde{z}$). b) Rotational errors parameterized in YPR-Euler angles ($\tilde{\psi}, \tilde{\theta}, \tilde{\phi}$). From [Lak+25].




Figure 3.15 Experiment 3.7: Comparison of the RMSE of the coordinates of the end effector for platform velocity control and platform torque control. From [Lak+25].

The joint positions for joints 4 to 6 (additional task coordinates) can be inspected in Figure 3.17 c).

Experiment 3.8 – Trajectory tracking with physical contact: In the following experiment, the behavior of the controller in the presence of external disturbances is investigated. The experiment is performed with the wheel torque controller, demonstrating active compliance of the platform. In order to highlight the compliance of the controlled system, the stiffness gains are chosen softer than in Experiment 3.7, see Table 3.1. The end effector tracks a line trajectory with $\dot{x}_{\text{des}} = -5 \text{ cm/s}$, while an operator disturbs both platform and end effector manually. A series of snapshots from the experiment can be inspected in Figure 3.18.

It can be verified that both platform and end effector are disturbed in different directions while the end effector follows the desired trajectory. The resulting positions of platform and end effector can be inspected in Figure 3.19. Figure 3.20 shows the Euclidean norm of the translational and rotational DOF of the end effector. The snapshot pictures from Figure 3.18 are linked to the shown data by printing the respective letter from the caption of the subfigures upon the plot. It can be observed that the rotational errors due to the disturbance are not converging after releasing the system at $t = 26 \text{ s}$, probably due to the small gains. However, the translational errors

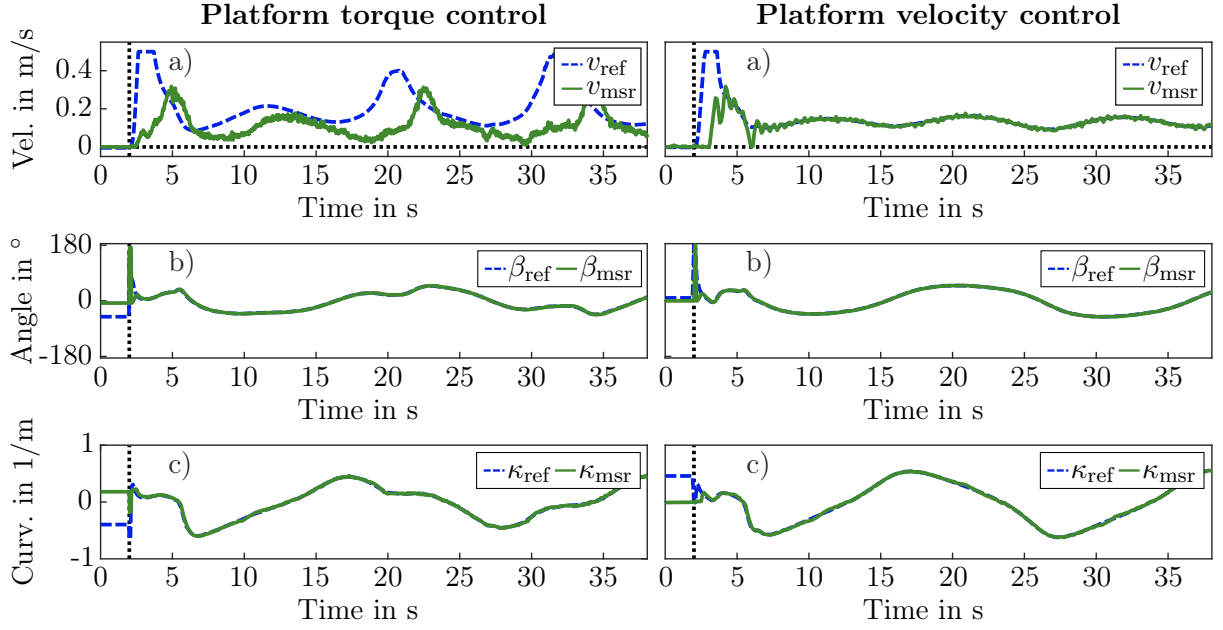


Figure 3.16 Experiment 3.7: Comparison of the performance of platform torque and velocity controller. a) Tracking of the reference velocity v . b) Tracking of the lateral driving angle β . c) Tracking of the curvature κ . From [Lak+25].

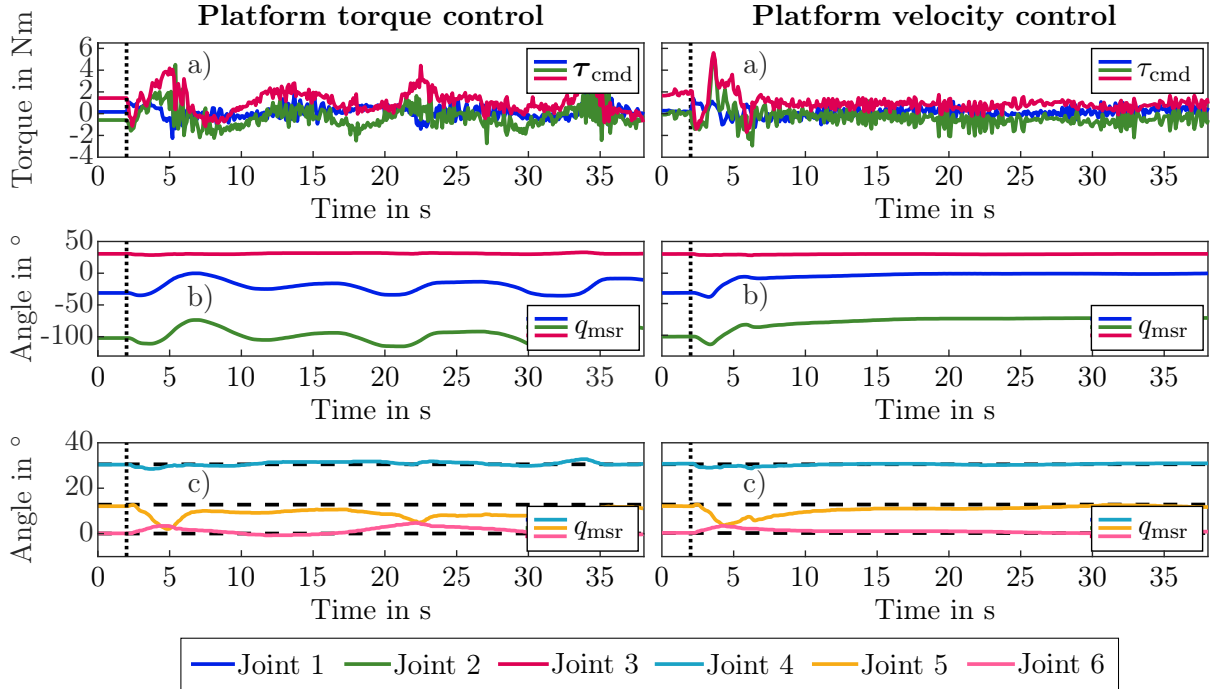


Figure 3.17 Experiment 3.7: Comparison of commanded torques and measured angles for the robot arm. a) Commanded torques for joints 1, 2, 3. b) Measured angles for joints 1, 2, 3. c) Reference values and measured angles for joints 4, 5, 6 (additional joint-space potential for redundancy resolution). From [Lak+25].

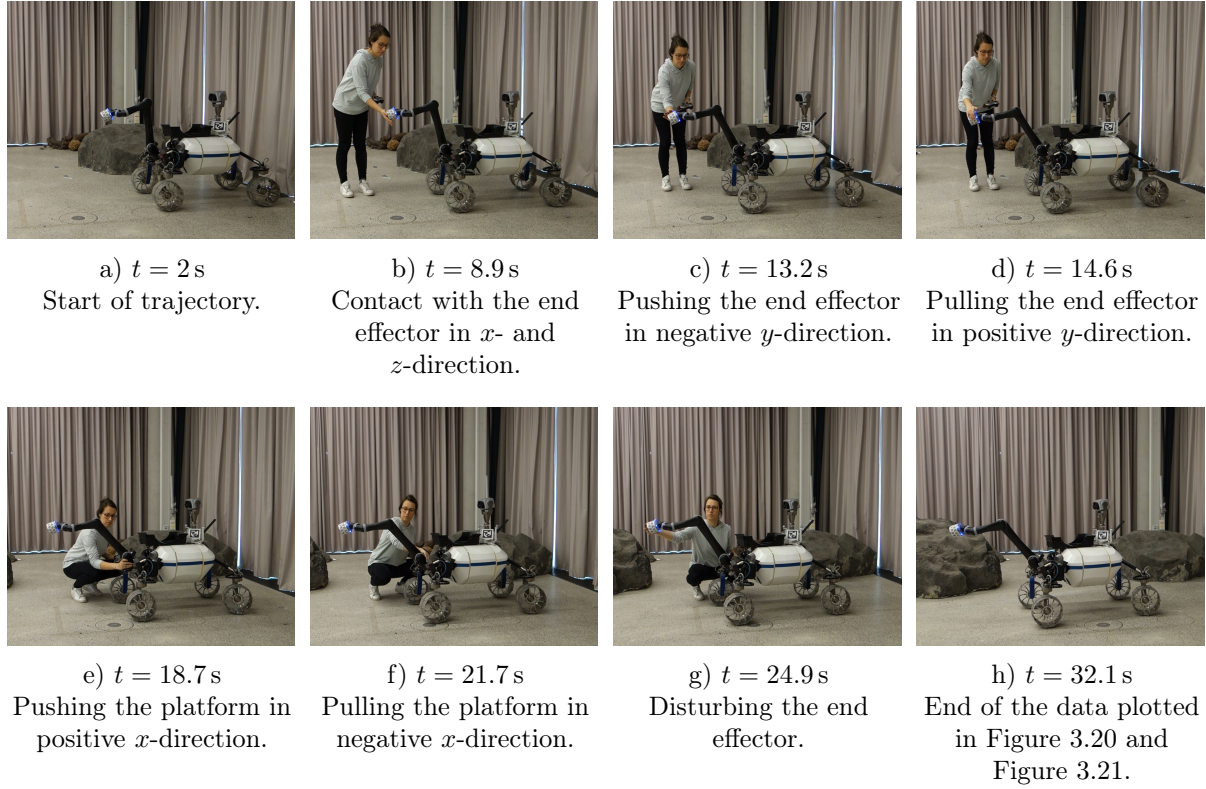


Figure 3.18 A series of snapshots from Experiment 3.8. The manual disturbance of both end effector and platform are depicted with their respective timestamps. The snapshots are linked to the corresponding data (Figure 3.20 and Figure 3.21) by the subcaptions a)–h). From [Lak+25].

are in the magnitude of centimeters and diminish when no external force is present. The effect of the external disturbance on the platform is visible in Figure 3.21, where the kinematic parameters of the platform are depicted. It can be seen that the reference values for the lateral driving angle β_{ref} and the curvature κ_{ref} change depending on the deviation of the end effector, e.g. at a), b) and c), when the end effector is disturbed in different directions. The active compliance of the platform becomes visible especially in the upper plot, where the tangential velocity is depicted. At e), the operator stops the platform manually. Therefore, the measured velocity v_{msr} becomes zero. However, the arm is (partly) capable to compensate for the external disturbance by stretching out, as can be observed in Figure 3.20 and Figure 3.18 d)–f).




Figure 3.19 Experiment 3.8: Measured position of the platform and the end effector. Both the end effector and the platform are disturbed manually during the tracking phase. From [Lak+25].




Figure 3.20 Experiment 3.8: Euclidean norm of translational (blue) and rotational (green) errors at the end effector. Phases with disturbance of the end effector are highlighted in yellow, disturbance of the platform is highlighted in green. The letters a)–h) at the upper line of the plot correspond to the snapshots in Figure 3.18. From [Lak+25].

Regulation of a desired pose of the end effector

To verify the two-phase concept for a static end point regulation, the end of trajectory of Experiment 3.7 is inspected in detail in Figure 3.22 and Figure 3.23. Here, only the last two seconds before the end of trajectory are depicted to show the convergence in detail. In Figure 3.22, one can see that the errors converge continuously after stopping the trajectory. The energy function $V_{x,y}$ from (3.68) is plotted in Figure 3.23. It is the sum of the potential energy-like part $V_{x,y,\text{pot}} = \frac{1}{2} \tilde{\mathbf{x}}^T \mathbf{K}_x \tilde{\mathbf{x}}$ (plotted in blue dashed), the kinetic-energy-like part $V_{x,y,\text{kin}} = \frac{1}{2} \dot{\mathbf{s}}^T \mathbf{M} \dot{\mathbf{s}}$ (green dash-dotted), and the potential energy of the null space function, $V_{x,y,\text{ns}} = \frac{1}{2} \tilde{\mathbf{q}}_y^T \mathbf{G} \tilde{\mathbf{q}}_y$ (yellow dash-dotted), which are also visible in the plot. Note that the null space energy jumps due to the drop of null space coordinates at $t = t_{\text{end}}$. It can be verified that the sum of energies increases slightly after the end of the trajectory, and then continues converging to zero.

3.3.5 Résumé

In comparison to the whole-body impedance controller from Section 3.2, the presented passivity-based tracking controller offers the possibility to change the direction of the possible platform motion by steering the wheels automatically, depending on the task. This makes the approach available not only for local manipulation tasks, but also for driving the platform to a desired destination. In the context of a planetary exploration mission, the controller can be used globally, both for locomotion (in this case, a joint-space tracking controller using the platform velocity interface is suggested in order to exploit the superior platform tracking performance), and for manipulation (switching to the platform torque interface to support active whole-body compliance in the vicinity of the object of manipulation). A particular advantage in this context is the possibility to utilize the same control law for a regulation task, covering also the end of the trajectory.

The derivation of the control law is based on a number of specific assumptions. First, the kinematic parameters ϕ and their derivatives $\dot{\phi}$ are assumed to be bounded. It has to be mentioned that this assumption depends on the particular choice of the kinematic parameters ϕ , and thus may limit the physical capabilities of a mobile platform. For example, the kinematic structure of the LRU2 allows the system to turn on the spot, but the choice of a bounded κ_{ref} does not include this kind of motion. If instead the position of the ICR is chosen as kinematic parameter, the

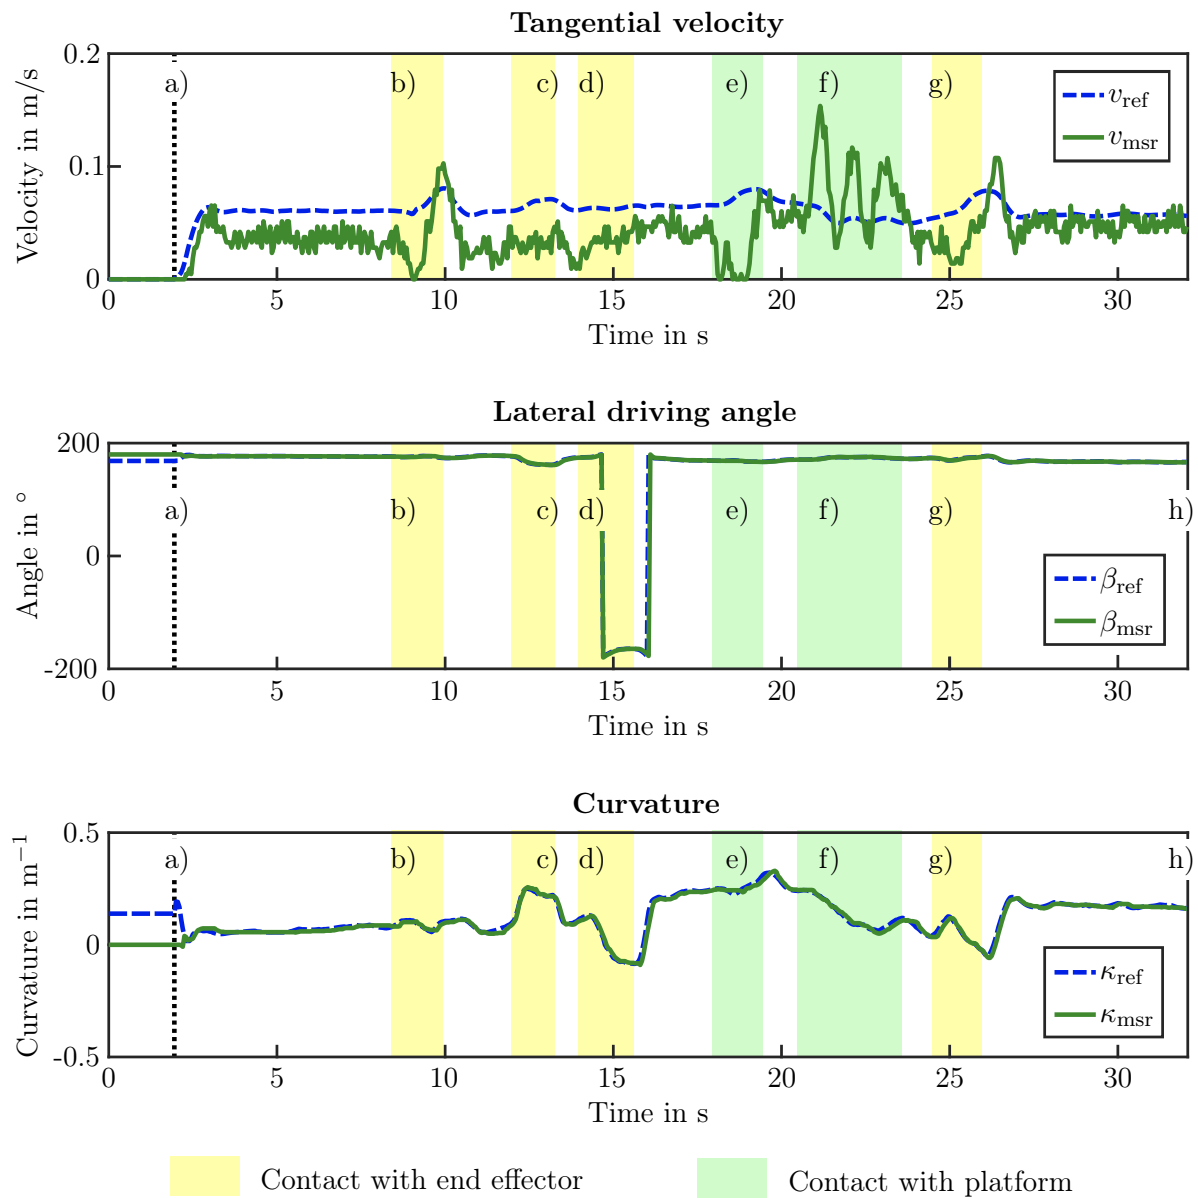


Figure 3.21 Experiment 3.8: Kinematic platform parameters during the contact experiment. Physical contact with the end effector is highlighted in yellow, disturbance of the platform is highlighted in green. The letters a)–h) correspond to the snapshots in Figure 3.18. From [Lak+25].




Figure 3.22 Experiment 3.7 (platform torque control): Euclidean norm of the translational (blue) and rotational (green) errors at the end effector. The plot shows the time span two seconds before until four seconds after the end of the trajectory. From [Lak+25].




Figure 3.23 Experiment 3.7 (platform torque control): Energy function at the end of the trajectory. The storage function $V_{x,y}$ (magenta) is the sum of the kinetic part $V_{x,y,kin}$ (green dash-dotted), the task space potential $V_{x,y,pot}$ (blue dashed), and the additional joint space potential, $V_{x,y,ns}$ (yellow dash-dotted). From [Lak+25].

platform cannot move straight, as the ICR coordinates would become infinite. Still, a reasonable subset of trajectories can be implemented assuming a smart choice of kinematic parameters. However, there might always be instantaneous violations of this assumption (e.g. β can jump due to the value range of the arctan function). Furthermore, a detailed discussion of the underlying modeling assumptions, in particular the performance of the low-level steering and wheel propulsion controllers, follows in Chapter 5. These assumptions mainly concern the hardware capabilities, and are thus no limitation of the algorithms in general. Indeed, the experiments verify that the control approach is applicable even if the hardware does not perfectly comply with the assumptions (the limited steering velocity can be mentioned as an example).

Summarized, the presented algorithm offers a broad field of application in different senses. First, it is based on classical passivity-based control concepts, which are fairly well-known and understood in theory and practice. The implementation of the kinematic solutions for the parts of the system equations subject to nonholonomic constraints can be solved and implemented straightforwardly. Second, it is applicable to a wide variety of robotic systems, particularly nonholonomic WMRs with different platform characteristics. The automatic generation of steering angles from a given task-space trajectory renders a dedicated platform trajectory planning redundant. Hence, the algorithm is particularly suited for scenarios with unknown or unforeseen external disturbances. This includes, besides the autonomous performance of different tasks in an unknown environment, the broad field of applications involving human-robot collaboration.

For all these reasons, the mentioned control law seems promising not only for the use in the context of planetary exploration, but also in industrial applications involving human-robot collaboration. Consequently, the following section can be understood as an excursion, exploring the basic properties of the passivity-based control law by means of a holonomic robot arm and in regard of human-robot interaction.

3.4 Passivity-based tracking control for holonomic robots

The whole-body control approach presented in Section 3.3 has shown great potential for the control of nonholonomic planetary rovers in a space exploration scenario. On the other hand, advantages of the control law in regard of human-robot collaboration were identified, particularly the active compliance of manipulator and platform. The properties of the passivity-based torque control law, particularly the convergence behavior and the meaning of the gains w.r.t. contact properties, are investigated in this section. In order to facilitate the theoretical and experimental analysis, and leaving aside the peculiarities of the wheeled subsystem, the aforementioned tracking control law is applied to a purely holonomic robot arm. The resulting control law can be understood as a combination of the well-known PD+ control law [PP88] and the Slotine-Li controller (SLC) [SL87]. Therefore, the control law is referred to as *Generalized Robot Tracking Controller* (GTC). While mentioned a couple of times in the literature before (in fact, [SH90] first mentioned the control law, but without performing a detailed theoretical or experimental analysis of the control law itself), a detailed investigation of the properties of the GTC in the context human-robot cooperation was only recently performed in [Lak+24].⁹ For the sake of comparability and clearness, the following considerations assume perfect model knowledge, which makes a parameter adaption law as originally added in [SL87] and [SH90] obsolete.

⁹Consequently, the following section partly corresponds to [Lak+24].

3.4.1 Brief introduction of PD+ and SLC

In the following, the PD+ as well as the Slotine-Li controller are briefly recapitulated, before proceeding to a detailed discussion of the GTC.

PD+ controller

The *augmented PD* or *PD+* controller is first introduced in [PP88]. The control law is

$$\tau_{\text{PD}+} = \mathbf{M}_m(\boldsymbol{\eta})\ddot{\boldsymbol{\eta}}_{\text{des}} + \mathbf{C}_m(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}})\dot{\boldsymbol{\eta}}_{\text{des}} + \mathbf{g}_m(\boldsymbol{\eta}) - \mathbf{K}_m\tilde{\boldsymbol{\eta}} - \mathbf{D}_m(\boldsymbol{\eta})\dot{\tilde{\boldsymbol{\eta}}} \quad (3.86)$$

with the s.p.d. gain matrices $\mathbf{K}_m \in \mathbb{R}^{n_m \times n_m}$ and $\mathbf{D}_m(\boldsymbol{\eta}) \in \mathbb{R}^{n_m \times n_m}$. Applying (3.86) to (2.25) yields the closed-loop dynamics

$$\mathbf{M}_m(\tilde{\boldsymbol{\eta}}, t)\ddot{\tilde{\boldsymbol{\eta}}} + (\mathbf{C}_m(\tilde{\boldsymbol{\eta}}, \dot{\tilde{\boldsymbol{\eta}}}, t) + \mathbf{D}_m(\tilde{\boldsymbol{\eta}}, t))\dot{\tilde{\boldsymbol{\eta}}} + \mathbf{K}_m\tilde{\boldsymbol{\eta}} = \boldsymbol{\tau}_{m,\text{ext}}. \quad (3.87)$$

For $\boldsymbol{\tau}_{m,\text{ext}} = \mathbf{0}$, it is possible to show uniform exponential stability of $(\tilde{\boldsymbol{\eta}}^*, \dot{\tilde{\boldsymbol{\eta}}}^*) = (\mathbf{0}, \mathbf{0})$ with the Lyapunov function

$$V_{\text{PD}+}(\tilde{\boldsymbol{\eta}}, \dot{\tilde{\boldsymbol{\eta}}}, t) = \frac{1}{2}\dot{\tilde{\boldsymbol{\eta}}}^T \mathbf{M}_m(\tilde{\boldsymbol{\eta}}, t)\dot{\tilde{\boldsymbol{\eta}}} + \frac{1}{2}\tilde{\boldsymbol{\eta}}^T \mathbf{K}_m\tilde{\boldsymbol{\eta}} + \epsilon \dot{\tilde{\boldsymbol{\eta}}}^T \mathbf{M}_m(\tilde{\boldsymbol{\eta}}, t)\tilde{\boldsymbol{\eta}} \quad (3.88)$$

(and some intermediate considerations). Thereby, $\epsilon > 0$ is a sufficiently small constant. A detailed proof of stability is given e. g. [MLS94, pp. 194–195]. One big advantage of the use of a PD+ controller is the direct realization of a desired contact stiffness and damping through \mathbf{K}_m and $\mathbf{D}_m(\boldsymbol{\eta})$, respectively. However, the convergence rate of the tracking errors cannot be computed exactly, as the value of ϵ is not specified.

Slotine-Li controller

The control law, omitting the parameter adaption due to the assumption of a perfect model knowledge, is derived as [SL87]

$$\tau_{\text{SLC}} = \mathbf{M}_m(\boldsymbol{\eta})\ddot{\boldsymbol{\eta}}_{\text{ref}} + \mathbf{C}_m(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}})\dot{\boldsymbol{\eta}}_{\text{ref}} + \mathbf{g}_m(\boldsymbol{\eta}) - \mathbf{H}_m(\boldsymbol{\eta})\boldsymbol{\sigma}, \quad (3.89)$$

introducing the s.p.d. gain matrix $\mathbf{H}_m(\boldsymbol{\eta}) \in \mathbb{R}^{n_m \times n_m}$. The *reference trajectory* is constructed as a combination of the desired velocities and the weighted position errors as

$$\dot{\boldsymbol{\eta}}_{\text{ref}} = \dot{\boldsymbol{\eta}}_{\text{des}} - \boldsymbol{\Omega}_m\tilde{\boldsymbol{\eta}}, \quad (3.90)$$

with the diagonal p.d. matrix $\boldsymbol{\Omega}_m \in \mathbb{R}^{n_m \times n_m}$. The so-called *sliding variable* $\boldsymbol{\sigma}$ is defined as

$$\boldsymbol{\sigma} = \dot{\boldsymbol{\eta}} - \dot{\boldsymbol{\eta}}_{\text{ref}} = \dot{\tilde{\boldsymbol{\eta}}} + \boldsymbol{\Omega}_m\tilde{\boldsymbol{\eta}}. \quad (3.91)$$

The resulting closed-loop dynamic equations are

$$\mathbf{M}_m(\tilde{\boldsymbol{\eta}}, t)\dot{\boldsymbol{\sigma}} + (\mathbf{C}_m(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma}, t) + \mathbf{H}_m(\tilde{\boldsymbol{\eta}}, t))\boldsymbol{\sigma} = \boldsymbol{\tau}_{m,\text{ext}} \quad (3.92)$$

$$\dot{\tilde{\boldsymbol{\eta}}} = \boldsymbol{\sigma} - \boldsymbol{\Omega}_m\tilde{\boldsymbol{\eta}}. \quad (3.93)$$

In [SOK90], uniform exponential stability of the origin $(\tilde{\boldsymbol{\eta}}^*, \boldsymbol{\sigma}^*) = (\mathbf{0}, \mathbf{0})$ for the undisturbed system is shown using the Lyapunov function

$$V_{\text{SLC}}(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma}, t) = \frac{1}{2}\boldsymbol{\sigma}^T \mathbf{M}_m(\tilde{\boldsymbol{\eta}}, t)\boldsymbol{\sigma} + \tilde{\boldsymbol{\eta}}^T \boldsymbol{\Omega}_m \mathbf{H}_m\tilde{\boldsymbol{\eta}}, \quad (3.94)$$

assuming that \mathbf{H}_m is constant and diagonal. Moreover, it is possible to directly show passivity of (3.92)–(3.93) with input $\boldsymbol{\tau}_{m,\text{ext}}$ and output $\boldsymbol{\sigma}$ using (3.94) as storage function. One advantage of the SLC compared to the PD+ controller is that the convergence of the error state $(\tilde{\boldsymbol{\eta}}, \dot{\tilde{\boldsymbol{\eta}}})$ can directly be specified through \mathbf{H}_m and $\boldsymbol{\Omega}_m$. On the other hand, assigning a decoupled interaction stiffness and damping is not possible by choice of the gains.

3.4.2 Generalized robot tracking controller

In fact, some theoretical and practical drawbacks arise from the use of PD+ controller and SLC. The most obvious is that for both control laws, the stability analysis uses Lyapunov functions which are not directly related to the physical energy of the system. Thus, computing the convergence rate of the tracking errors does not necessarily give information about the decay of potential and kinetic energy of the robot. Also, tuning the gains is restricted to either influencing contact properties directly (PD+) or tuning the pure tracking performance (SLC).

To address the aforementioned disadvantages, the GTC will be reviewed in the following from [Lak+24]. Summarized, the GTC features the following beneficial properties:

- a) Uniform exponential stability with a more physically motivated Lyapunov function than for PD+ (3.88) and SLC (3.94), which allows to examine potential as well as kinetic "pseudo-energy" (originating from the deviations from the desired trajectory).
- b) Independent parameterization of contact stiffness and damping gains (in contrast to SLC).
- c) The possibility to impose an exponential convergence rate directly by choice of the gains (in contrast to PD+).

The control law was first mentioned in the literature in [SH90]. A slightly modified version¹⁰ can be derived by applying the passivity-based tracking control law from the previous section to the manipulator subsystem only. The resulting control law is

$$\tau_{\text{GTC}} = \mathbf{M}_m(\boldsymbol{\eta})\ddot{\boldsymbol{\eta}}_{\text{ref}} + \mathbf{C}_m(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}})\dot{\boldsymbol{\eta}}_{\text{ref}} + \mathbf{g}_m(\boldsymbol{\eta}) - \mathbf{K}_m\tilde{\boldsymbol{\eta}} - \mathbf{H}_m(\boldsymbol{\eta})\boldsymbol{\sigma}, \quad (3.95)$$

with the uniformly bounded s.p.d. gain matrix $\mathbf{H}_m(\boldsymbol{\eta}) \in \mathbb{R}^{n_m \times n_m}$ and the diagonal p.d. gain matrix $\mathbf{K}_m \in \mathbb{R}^{n_m \times n_m}$. Applying (3.95) together with (3.90) and (3.91) to the dynamic equations of the robot arm (2.25) results in the closed-loop equations

$$\mathbf{M}_m(\tilde{\boldsymbol{\eta}}, t)\dot{\boldsymbol{\sigma}} + (\mathbf{C}_m(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma}, t) + \mathbf{H}_m(\tilde{\boldsymbol{\eta}}, t))\boldsymbol{\sigma} + \mathbf{K}_m\tilde{\boldsymbol{\eta}} = \boldsymbol{\tau}_{m,\text{ext}} \quad (3.96)$$

$$\dot{\tilde{\boldsymbol{\eta}}} = \boldsymbol{\sigma} - \boldsymbol{\Omega}_m\tilde{\boldsymbol{\eta}}. \quad (3.97)$$

A corresponding Lyapunov function candidate $V_{\text{GTC}} : \mathcal{Q} \times \mathbb{R}^n \times \mathbb{R}^+ \rightarrow \mathbb{R}$ is proposed as

$$V_{\text{GTC}}(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma}, t) = \frac{1}{2}\boldsymbol{\sigma}^T \mathbf{M}_m(\tilde{\boldsymbol{\eta}}, t)\boldsymbol{\sigma} + \frac{1}{2}\tilde{\boldsymbol{\eta}}^T \mathbf{K}_m\tilde{\boldsymbol{\eta}}, \quad (3.98)$$

which is positive definite in $(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma})$ for all t . The time derivative of (3.98) along the trajectories of (3.96)–(3.97) is readily computed as

$$\dot{V}_{\text{GTC}}(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma}, t) = \boldsymbol{\sigma}^T \boldsymbol{\tau}_{m,\text{ext}} - \boldsymbol{\sigma}^T \mathbf{H}_m(\tilde{\boldsymbol{\eta}}, t)\boldsymbol{\sigma} - \tilde{\boldsymbol{\eta}}^T \mathbf{K}_m \boldsymbol{\Omega}_m \tilde{\boldsymbol{\eta}}. \quad (3.99)$$

For $\boldsymbol{\tau}_{m,\text{ext}} = \mathbf{0}$, (3.99) is negative definite in $(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma})$ for all t due to the definition of \mathbf{K}_m and $\boldsymbol{\Omega}_m$ as diagonal p.d. matrices. Following the same line of argumentation as in the proof of Proposition 1, it is then possible to determine time-independent upper and lower bounds for $V_{\text{GTC}}(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma}, t)$:

$$\underline{V}_{\text{GTC}}(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma}) = \frac{1}{2} \inf_{\boldsymbol{\eta} \in \mathcal{Q}} \min_i \lambda_i(\mathbf{M}_m(\boldsymbol{\eta})) \|\boldsymbol{\sigma}\|^2 + \frac{1}{2} \min_i \lambda_i(\mathbf{K}_m) \|\tilde{\boldsymbol{\eta}}\|^2 \quad (3.100)$$

$$\overline{V}_{\text{GTC}}(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma}) = \frac{1}{2} \sup_{\boldsymbol{\eta} \in \mathcal{Q}} \max_i \lambda_i(\mathbf{M}_m(\boldsymbol{\eta})) \|\boldsymbol{\sigma}\|^2 + \frac{1}{2} \max_i \lambda_i(\mathbf{K}_m) \|\tilde{\boldsymbol{\eta}}\|^2 \quad (3.101)$$

¹⁰Modification mainly in terms of notation. Also, [SH90] introduced scalar gains instead of gain matrices.

Control Law	$\tilde{\eta}$: Effective P-Gain	$\dot{\tilde{\eta}}$: Effective D-Gain
PD+	$-\mathbf{K}_m$	$-\mathbf{D}_m(\eta)$
SLC	$-\mathbf{C}_m(\eta, \dot{\eta})\Omega_m - \mathbf{H}_m(\eta)\Omega_m$	$-\mathbf{M}_m(\eta)\Omega_m - \mathbf{H}_m(\eta)$
GTC	$-\mathbf{C}_m(\eta, \dot{\eta})\Omega_m - \mathbf{H}_m(\eta)\Omega_m - \mathbf{K}_m$	$-\mathbf{M}_m(\eta)\Omega_m - \mathbf{H}_m(\eta)$

Table 3.2 Comparison of the effective control gains of PD+, SLC and GTC. From [Lak+24].

such that $\underline{V}_{\text{GTC}}(\tilde{\eta}, \sigma) \leq V_{\text{GTC}}(\tilde{\eta}, \sigma, t) \leq \bar{V}_{\text{GTC}}(\tilde{\eta}, \sigma)$ holds. Moreover, (3.99) is bounded by

$$\dot{\bar{V}}_{\text{GTC}}(\tilde{\eta}, \sigma) = -\inf_{\eta \in \mathcal{Q}} \min_i \lambda_i(\mathbf{H}_m(\eta)) \|\sigma\|^2 - \min_i \lambda_i(\mathbf{K}_m \Omega_m) \|\tilde{\eta}\|^2, \quad (3.102)$$

such that

$$\dot{V}_{\text{GTC}}(\tilde{\eta}, \sigma, t) \leq \dot{\bar{V}}_{\text{GTC}}(\tilde{\eta}, \sigma) < 0 \quad \forall (\tilde{\eta}, \sigma) \neq (\mathbf{0}, \mathbf{0}). \quad (3.103)$$

Uniform exponential stability of the equilibrium $(\tilde{\eta}^*, \sigma^*) = (\mathbf{0}, \mathbf{0})$ follows directly from the exponential stability theorem (see Theorem 1 in Appendix A.1.2).

Remark 9. Again, the Lyapunov function (3.98) is radially unbounded for $\eta \in \mathcal{Q}$. In this sense, the proof of stability holds “globally” for the relevant working space of the robot.

Remark 10. For the sake of completeness, it should be mentioned that the stability properties derived for $(\tilde{\eta}, \sigma)$ do also apply for the full system error state $(\tilde{\eta}, \dot{\tilde{\eta}})$, as σ can be considered as input of the asymptotically stable $\tilde{\eta}$ -dynamics (3.97). Thus, convergence of σ directly implies convergence of $\dot{\tilde{\eta}}$, cf. [SL91, p. 399].

Remark 11. With (3.98) and (3.99), passivity of (3.96)–(3.97) w.r.t. the power port $(\tau_{m,\text{ext}}, \sigma)$ can be concluded.

3.4.3 Discussion of some properties

As mentioned above, the GTC combines components from the PD+ and the SLC. In particular, both of these classical control algorithms can be derived as corner cases of the GTC:

- For $\Omega_m \rightarrow \mathbf{0}$, (3.95) becomes a PD+ controller (3.86).
- For $\mathbf{K}_m \rightarrow \mathbf{0}$, (3.95) becomes a Slotine-Li controller (3.89).

Comparison of the control gains: In order to compare the control gains, the SLC and GTC control laws (3.89) and (3.95) are written in the form of the PD+ control law (3.86), i.e., the “+”-part equals the first three terms of (3.86), and is contained in all three control laws. The other parts are sorted by $\tilde{\eta}$ and $\dot{\tilde{\eta}}$. The resulting proportional and derivative gains are denoted as *effective* control gains. They are listed in Table 3.2.

It can be seen that the SLC as well as the GTC introduce dynamic couplings in the feedback terms. For the SLC, stiffness and damping are coupled via the choice of $\mathbf{H}_m(\eta)$ and Ω_m . The additional proportional gain \mathbf{K}_m in the GTC allows a “decoupling” of effective stiffness and damping (i.e., you can change the effective stiffness by choosing a suitable \mathbf{K}_m without changing the effective damping). This might be beneficial under certain circumstances, e.g. when intrinsic elastic properties of a system can be used in order to perform a motion more efficiently. In practice, it is possible to implement and tune an SLC, and then add the additional stiffness term $\mathbf{K}_m \tilde{\eta}$ in order to adjust or improve convergence and/or compliance behavior.

Parameterization of stiffness and damping: An interesting possibility for gain tuning is to look at \mathbf{K}_m as influencing the behavior in physical contact situations. From (3.96) and (3.97), it can be verified that the perceived contact stiffness behavior is $(\mathbf{H}_m \boldsymbol{\Omega}_m + \mathbf{K}_m) \tilde{\boldsymbol{\eta}} = \boldsymbol{\tau}_{m,\text{ext}}$ by regarding the static case $\dot{\boldsymbol{\eta}} = \dot{\boldsymbol{\eta}}_{\text{des}} = \mathbf{0}$ and $\ddot{\boldsymbol{\eta}} = \ddot{\boldsymbol{\eta}}_{\text{des}} = \mathbf{0}$. This means that with a given choice of \mathbf{H}_m and $\boldsymbol{\Omega}_m$, it is possible to enforce a desired contact behavior $\mathbf{K}_{m,\text{eff}} \tilde{\boldsymbol{\eta}} = \boldsymbol{\tau}_{m,\text{ext}}$ by setting

$$\mathbf{K}_m = \mathbf{K}_{m,\text{eff}} - \mathbf{H}_m \boldsymbol{\Omega}_m. \quad (3.104)$$

However, \mathbf{K}_m must be positive definite to ensure stability with the given Lyapunov function (3.98). In other words, the perceived contact stiffness for the GTC should be chosen always equal to or larger than $\mathbf{H}_m \boldsymbol{\Omega}_m$ (which is the perceived contact stiffness of the SLC). In order to overcome this theoretical restriction, another Lyapunov function candidate is proposed:

$$V_{GTC}^*(\tilde{\boldsymbol{\eta}}, \boldsymbol{\sigma}, t) = \frac{1}{2} \boldsymbol{\sigma}^T \mathbf{M}_m \boldsymbol{\sigma} + \frac{1}{2} \tilde{\boldsymbol{\eta}}^T \mathbf{K}_m \tilde{\boldsymbol{\eta}} + \tilde{\boldsymbol{\eta}}^T \bar{\mathbf{H}}_m \boldsymbol{\Omega}_m \tilde{\boldsymbol{\eta}}, \quad (3.105)$$

where $\bar{\mathbf{H}}_m$ equals the gain matrix \mathbf{H}_m from (3.95), which is additionally required to be constant and diagonal. Then, the time derivative w. r. t. (3.96)-(3.97) is

$$\dot{V}_{GTC}^* = - \begin{pmatrix} \tilde{\boldsymbol{\eta}} \\ \boldsymbol{\sigma} \end{pmatrix}^T \begin{pmatrix} \mathbf{K}_m \boldsymbol{\Omega}_m + 2\bar{\mathbf{H}}_m \boldsymbol{\Omega}_m^2 & -\bar{\mathbf{H}}_m \boldsymbol{\Omega}_m \\ -(\bar{\mathbf{H}}_m \boldsymbol{\Omega}_m)^T & \bar{\mathbf{H}}_m \end{pmatrix} \begin{pmatrix} \tilde{\boldsymbol{\eta}} \\ \boldsymbol{\sigma} \end{pmatrix} \quad (3.106)$$

which is negative definite if $\bar{\mathbf{H}}_m \succ \mathbf{0}$ and

$$\mathbf{K}_m \boldsymbol{\Omega}_m + 2\bar{\mathbf{H}}_m \boldsymbol{\Omega}_m^2 - (\bar{\mathbf{H}}_m \boldsymbol{\Omega}_m)^T \bar{\mathbf{H}}_m^{-1} \bar{\mathbf{H}}_m \boldsymbol{\Omega}_m \succ \mathbf{0} \quad (3.107)$$

according to the Schur complement condition [Zha05]. Choosing $\mathbf{K}_m = \gamma \bar{\mathbf{H}}_m \boldsymbol{\Omega}_m$, it can be verified that

$$(\gamma + 1) \bar{\mathbf{H}}_m \boldsymbol{\Omega}_m^2 \succ \mathbf{0}, \quad (3.108)$$

which holds for $\gamma > -1$. This shows that the effective stiffness of the GTC can be changed from arbitrarily close to zero up to higher stiffness values compared to the SLC. The effect is also shown in the experiments at the end of this section.

Convergence properties: Similar to the considerations in Section 3.3.1, it can be observed that the decrease rate of the energy function depends on the gains \mathbf{H}_m and $\mathbf{K}_m \boldsymbol{\Omega}_m$ (cf. (3.102) and (3.103)). In this case, weighting \mathbf{K}_m versus $\boldsymbol{\Omega}_m$ “shifts” the behavior of the closed-loop system between convergence of the position error within the sliding variable in (3.97) and the parameterization of a desired stiffness (choice of \mathbf{K}_m).

A remarkable observation can be made by inspecting the Lyapunov function and its derivative at $\boldsymbol{\tau}_{m,\text{ext}} = \mathbf{0}$: Making the particular choice of gains

$$\mathbf{H}_m(\boldsymbol{\eta}) := \alpha \mathbf{M}_m(\boldsymbol{\eta}) \quad (3.109)$$

$$\boldsymbol{\Omega}_m := \alpha \mathbf{I} \quad (3.110)$$

with some constant $\alpha > 0$, it can be verified that

$$\dot{V}_{GTC} = -2\alpha V_{GTC}. \quad (3.111)$$

and thus

$$V_{GTC}(t) = V_{GTC}(t_0) e^{-2\alpha(t-t_0)}. \quad (3.112)$$

Gain	Unit	Experiments 3.9 and 3.10 (Tracking)	Experiment 3.11 (Contact)
$\mathbf{\Omega}_m$	1/s	diag(30, 30, 30, 30, 30, 30, 30)	diag(30, 30, 30, 30, 30, 30, 30)
\mathbf{H}_m	Nm s/rad	diag(50, 50, 50, 20, 10, 5, 5)	diag(15, 15, 15, 6, 3, 1.5, 1.5)

Table 3.3 Gains used for the experiments 3.9–3.11.

Note that similar considerations exist for the SLC, cf. [SL91, pp. 409ff.]. However, the storage function $V_{\text{GTC}}(t)$ can be interpreted as a “pseudo-energy”, originating from the deviations from the desired trajectory. In that sense, the convergence rate is physically more meaningful as the convergence rate of the SLC.

As the considerations above hold independent of \mathbf{K}_m , tuning the GTC boils down to the choice of α_m , specifying the decay rate of the storage function. Then, \mathbf{K}_m can be chosen independently, weighting the “potential energy”-like term vs. the “kinetic energy”-like term in (3.98) and thus, again, as shifting a constant convergence rate between $\tilde{\boldsymbol{\eta}}$ and $\boldsymbol{\sigma}$.

Remark 12. For the particular choice (3.109)–(3.110), a configuration-dependent contact stiffness results:

$$\mathbf{K}_{m,\text{eff}} = \alpha^2 \mathbf{M}_m(\boldsymbol{\eta}) + \mathbf{K}_m. \quad (3.113)$$

This means that it is only possible to specify a contact stiffness at a given operating point. Moreover, the position and velocity error feedback terms in the control law (3.95) become inertia-dependent. In general, this can be problematic in terms of robustness, as unmodeled disturbances as e.g. friction can destabilize the system (see e.g. [Die+21] for a detailed discussion).

Remark 13. It is possible to extend (3.109) by a second (positive, possibly time-dependent) parameter $\beta(t)$, such that

$$\mathbf{H}_m(\boldsymbol{\eta}) := (\alpha + \beta(t)) \mathbf{M}(\boldsymbol{\eta}) \quad (3.114)$$

$$\mathbf{\Omega}_m := \alpha \mathbf{I}. \quad (3.115)$$

In this case, the equality (3.112) becomes a “less or equal”, which means that the given exponential function becomes an upper bound for the real convergence.

Remark 14. For the sake of completeness, it shall be mentioned that the GTC algorithm can straightforwardly be extended to the task space, analog to the passivity-based tracking controller for nonholonomic systems. Details are given in [Lak+24].

3.4.4 Experiments

Finally, an experimental study of the tracking and interaction performance of the GTC is performed [Lak+24], using a commercially available robot arm, the KUKA LWR IV+ [Bis+10]. The manipulator features seven revolute joints, see Figure 3.24 a). The controllers for the experiment are running in 1 kHz. The resulting joint torque commands are processed by low-level joint torque controllers operating at 3 kHz, cf. [AOH07], [Ott+08]. A selection of experimental results is shown in a video available at <https://www.youtube.com/watch?v=rQcEfocGjVY>.

Experiment 3.9 – Tracking performance: The first experiment is performed to examine the tracking performance of the GTC in comparison to the SLC. Note that the effective P- and D-gains are configuration-dependent for SLC and GTC, while they are constant for PD+ (cf. Table 3.2), thus a comparison of the tracking performance of GTC and PD+ is not considered

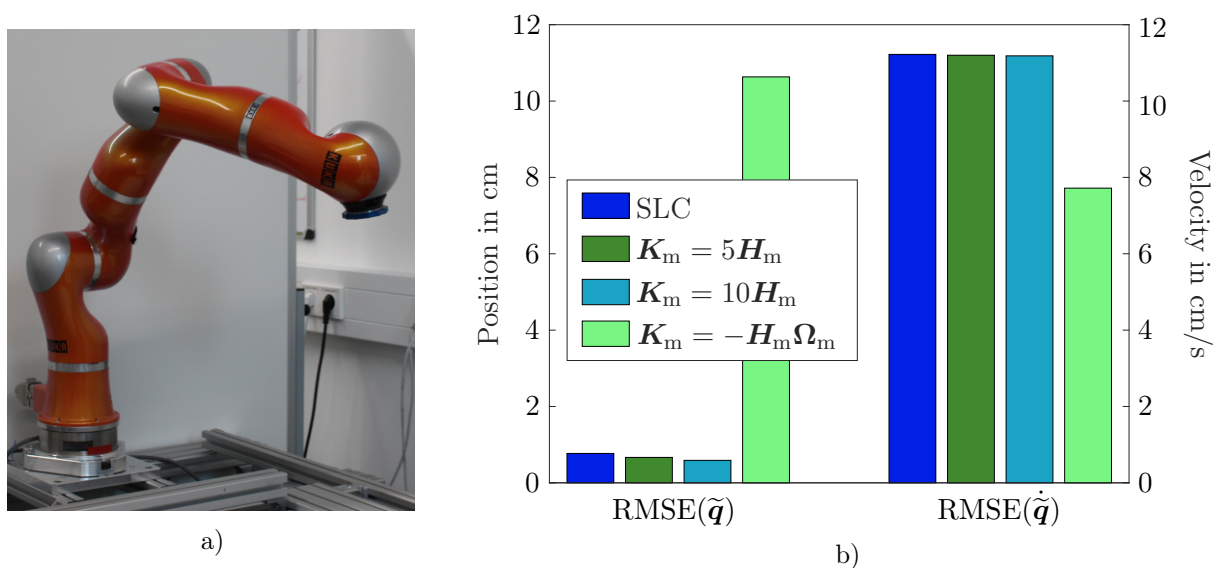


Figure 3.24 a) The KUKA LWR IV+ in its initial pose for the tracking experiments 3.9 and 3.10. b) Comparison of joint position and velocity errors for different stiffness parameterizations (cf. Experiment 3.9, undisturbed trajectory tracking). From [Lak+24].

meaningful. Thereby, the gains of the baseline SLC implementation are manually tuned to reach decent tracking performance. The resulting gains Ω_m and H_m are given in Table 3.3, while the value of K_m is set to different values in order to compare the result on the tracking performance of the GTC. In order to incorporate the different load and motor characteristics of the single joints, K_m is chosen proportional to H_m .

Figure 3.25 a) shows the desired trajectory for all seven joints of the LWR, while Figure 3.25 b) and c) depict the Euclidean norm of the joint position errors and the joint velocity errors, respectively. It can be verified that the position error norm decreases with higher values of K_m , whereas the velocity error does not change significantly. This finding can be explained with the effect of K_m on the effective gains (cf. Table 3.2). In fact, adding a positive $K_m\tilde{\eta}$ to the GTC control law increases the effective P-gain, but does not influence the effective D-gain.

The second question is how a negative value of K_m (and thus a decreased effective stiffness) influences the tracking performance of the GTC. Therefore, a value of $K_m = -H_m\Omega_m$ is added to the analysis. Figure 3.24 b) shows the resulting RMSE of the corresponding positions and velocities, in comparison to the SLC and two parameterization with increased effective P-gain. The results validate that choosing K_m negative up to $K_m = -H_m\Omega_m$ yields a stable behavior, while the perceived stiffness $K_{m,\text{eff}}$ of the SLC controller is reduced until $K_{m,\text{eff}} = 0$. In that case, the position error is significantly higher than for the positive effective stiffness values, as the position tracking errors do not converge anymore but sum up over the course of the trajectory. However, the velocity error does still converge for $K_m = -H_m\Omega_m$, as the $\text{RMSE}(\dot{\tilde{q}})$ confirms. In the video, it is visible that the overall tracking performance still looks comparable to the other stiffness parameterizations. This shows that it is indeed possible to choose the effective P-gain and effective D-gain completely independent. Remarkably, it is possible to assign an arbitrary contact stiffness (up to zero effective stiffness) to a given SLC implementation by using the GTC.

Experiment 3.10 – Tracking in the presence of external torques: For the second experiment, the same trajectory as in Experiment 3.9 is commanded. In order to investigate the effect of




Figure 3.25 Experiment 3.9: Undisturbed trajectory tracking. The gains are selected according to Table 3.3. a) Desired trajectory $\boldsymbol{\eta}_{\text{des}}$ in joint space. b) Euclidean norm of joint position errors $\|\tilde{\boldsymbol{\eta}}\|_2$ for different values of \mathbf{K}_m . c) Euclidean norm of joint velocity errors $\|\dot{\tilde{\boldsymbol{\eta}}}\|_2$ for different values of \mathbf{K}_m . From [Lak+24].

disturbances on the tracking performance, virtual external torques of ± 20 Nm are applied at the first three joints between $t = 1$ s and $t = 3$ s, see Figure 3.26 a). The tracking results for two parameterizations are shown in Figure 3.26 b) and c).

Another question investigated with this experiment is how good the theoretically computed contact stiffness would be perceived in a contact situation with the robot moving on a trajectory. This can be estimated comparing the measured joint deflections caused by the virtual external torques to the expected joint deflections which are computed for the static case as $\tilde{\boldsymbol{\eta}}_{\text{comp.}} = \mathbf{K}_{m,\text{eff}}^{-1} \boldsymbol{\tau}_{\text{ext,virt.}}$. Figure 3.26 d) shows the achieved joint stiffness for joint 2 exemplarily. The dashed lines denote the computed joint deflections, the other lines are the measured joint position errors. Deviations originate among other things from the Coriolis terms in the effective P-gain (cf. Table 3.2), which is present during the execution of a trajectory. Still, it is visible that the measured joint deflections fit the stiffness parameterization.

Experiment 3.11 – Physical interaction: Experiment 3.11 shows a static regulation case involving physical interaction with a human operator. Thereby, the robot is commanded to maintain the initial configuration $\boldsymbol{\eta}_{\text{des}} = (-\pi/2, \pi/2, 0, 0, 0, 0, 0)$ rad. The external torques applied by the operator are estimated using the momentum-based observer from [De +06; HDA17], with the observer gain chosen as 25 s^{-1} . In order to investigate the effect of the perceived stiffness during interaction, the parameter \mathbf{H}_m is reduced by a factor of $2/3$, cf. see Table 3.3. \mathbf{K}_m was chosen as $\mathbf{K}_m = -\boldsymbol{\Omega}_m \mathbf{H}_m$, $\mathbf{K}_m = 0 \text{ s}^{-1} \mathbf{H}_m$ (which is equal to the SLC), and $\mathbf{K}_m = 100 \text{ s}^{-1} \mathbf{H}_m$, respectively. The stability and the contact behavior can be inspected in Figure 3.27 and Figure 3.28, which show the estimated external joint torques (plot f)) together the measured joint deflections (plot g)). Snapshots of the corresponding operator actions complement the measurements, indicated by a)–e).




Figure 3.26 Experiment 3.10: Trajectory with virtual external torques. The gains were chosen as in Experiment 3.9. a) Virtual external torques applied to the first three joints. b) Euclidean norm of joint position errors $\|\tilde{\eta}\|_2$ for the GTC compared to the SLC. c) Euclidean norm of joint velocity errors $\|\dot{\tilde{\eta}}\|_2$ for the GTC compared to the SLC. d) Measured joint position deviation of joint 2 vs. computed joint position error. The error was computed for the static case as $\tilde{\eta}_{comp.} = K_{m,eff}^{-1} \tau_{m,ext,virt}$. From [Lak+24].

The first observation is that the GTC is stable with all chosen gains, even with very high perceived stiffness $K_{m,eff} = K_m + H_m \Omega_m = \text{diag}(1950, 1950, 1950, 780, 390, 195, 195)$ Nm/rad and with zero perceived stiffness at $K_m = -H_m \Omega_m$. Second, it can be verified that the external torques are in the same order of magnitude for all parameterizations. However, the corresponding joint errors are smaller and converge faster for the higher stiffness $K_m = 100 \text{ s}^{-1} H_m$, whereas the application of external torques leads to remaining deviations in the joint angles for zero effective stiffness. This result validates both the stability and the theoretically resulting perceived stiffness property of the GTC.

3.4.5 Résumé

The GTC algorithm can be used for trajectory tracking of robotic arms in the joint space, an extension to the task space is possible. Regarding the implementation, the GTC is not expensive (adding a position feedback term to an SLC-implementation is sufficient). As pointed out in Section 3.4.3, it is then possible to create both a PD+ controller and an SLC as corner cases by proper selection of the respective gains. From this insight, different possible gain tuning

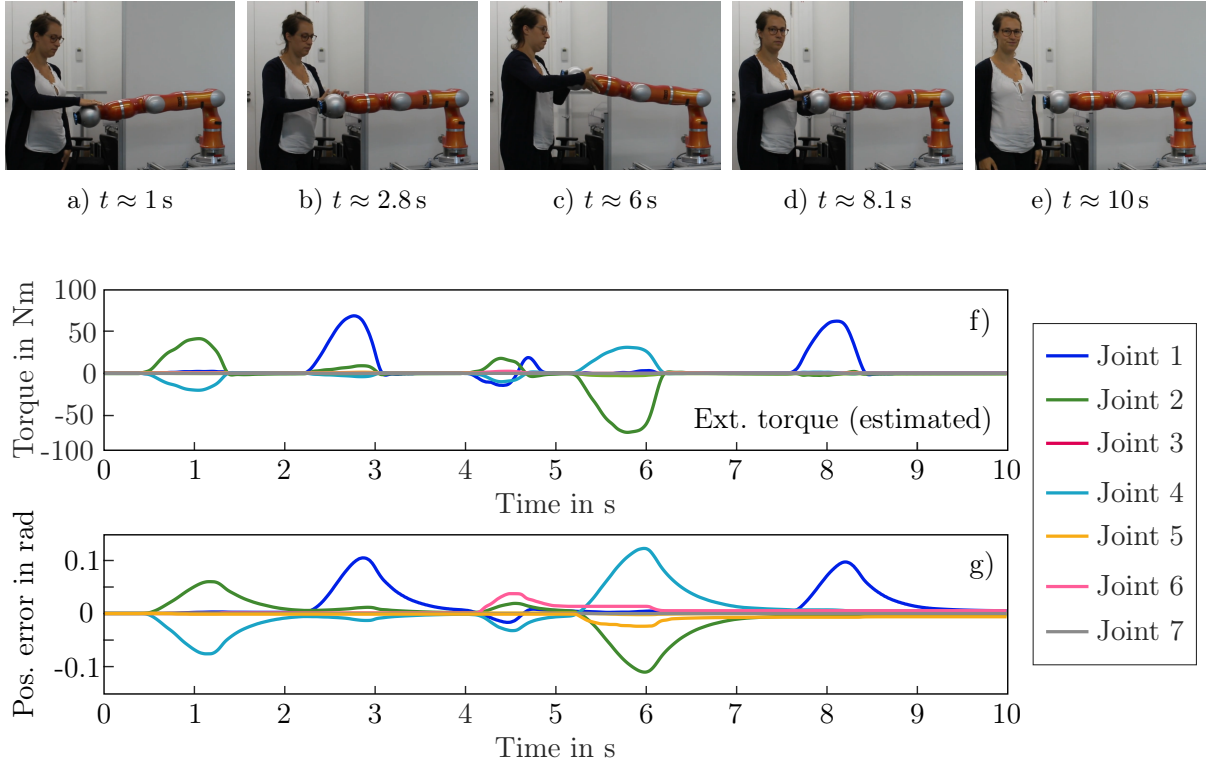


Figure 3.27 Experiment 3.11: $\mathbf{K}_m = \mathbf{0}$ (SLC). Physical interaction with a human operator. $\mathbf{\Omega}_m$ and \mathbf{H}_m are given in Table 3.3. a)–e) Snapshots from the video. f) Joint torque disturbances, estimated using a disturbance observer. Note that the estimations are only qualitative and are plotted to give an impression of the actions of the human. g) Measured joint position errors. From [Lak+24].

strategies arise. One of several possibilities is experimentally examined, where an existing SLC implementation is used as a baseline w.r.t. the trajectory tracking behavior. The additional P-gain is then used to adjust the stiffness according to the task (either tracking or interaction).

Remarkably, it could be shown that the effective stiffness can not only be increased but also be reduced compared to the SLC, while maintaining the full velocity tracking performance. If the model is sufficiently good (or the controller is used together with a parameter adaption part as introduced in [SH90]), this creates the possibility to realize a high tracking performance, while the contact stiffness is still low. This property predestines the GTC among others to human-robot interaction scenarios.

Regarding the use in planetary exploration scenarios, the presented analyses of the GTC can contribute to a deeper understanding of the whole-body tracking control law proposed in Section 3.3, and thus facilitate a task-consistent choice of the controller gains.

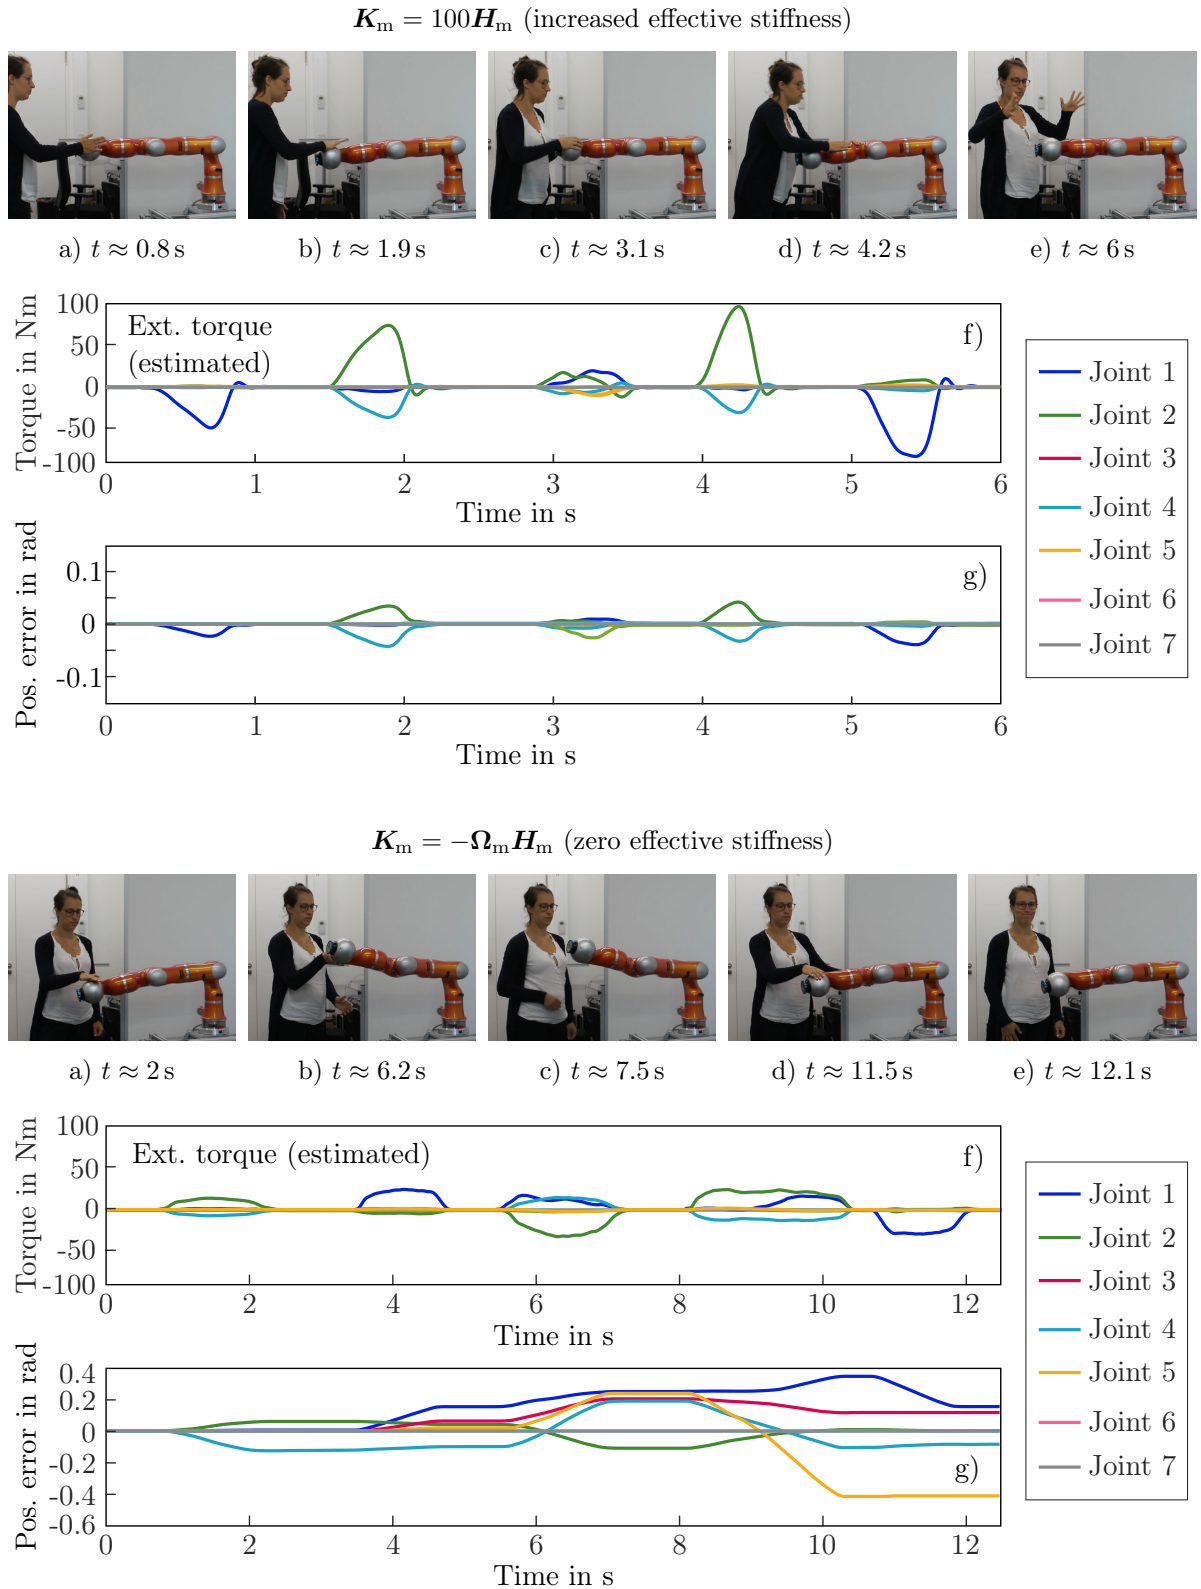


Figure 3.28 Experiment 3.11: Physical interaction with a human operator. Upper figures: $K_m = 100H_m$ (increased effective stiffness). Lower figures: $K_m = -\Omega_m H_m$ (zero effective stiffness). a)–e) Snapshots from the video. f) Joint torque disturbances, estimated using a disturbance observer. Note that the estimations are only qualitative and are plotted to give an impression of the actions of the human. g) Measured joint position errors. From [Lak+24].

CHAPTER 4

Wheel-Ground Interaction

From the preceding chapter, it becomes clear that one of the limiting factors for the tracking accuracy at the end effector is the presence of unmodeled wheel-ground interaction effects. Therefore, it is clearly desirable to estimate and compensate effects like slip in order to improve the whole-body tracking performance.

But, even more crucial when it comes to autonomous exploration, wheel slip effects negatively affect essential autonomy skills such as the self localization of the rover, if it uses wheel odometry measurements [Yos09]. This is particularly the case when traversing longer distances in rough terrain. An impressive example for the issues arising from an erroneous self localization are the Mars Exploration Rovers *Spirit* and *Opportunity*. Even though the design average traverse speed of the rovers was limited to 100 m/day (among others due to the risks inherent in target designation beyond 100 m), Spirit's wheels sank into a sand trap while climbing a slope in 2009, after traversing 7.73 km in 6 years and 3 months [Ell16, p. 68].

Summarized, slip estimation and compensation are substantial factors for the success of ground-based planetary exploration missions. Not only the exploration part is concerned, where longer distances have to be traversed without the possibility to obtain external measurements of the rover position. Also local manipulation tasks utilizing the compliant whole-body controller rely on a proper treatment of the wheel-ground interaction. In view of the scenario, the chosen approaches for slip estimation and compensation have to consider both limitations of the hardware and the on-board computers, and external conditions (for example the lack of ground truth position data).

Let us assume that the rover is deployed on a foreign celestial body, where the terrain properties cannot always be assumed to be known a priori. The first question which arises in the context of wheel-ground interaction is how to model the observed effects without a priori knowledge of the soil. In this case, the behavior of the robot in combination with the environment can only be determined by the available on-board sensors of the rover, without relying on external measurements. The obtained information can then be used in order to improve the on-board self localization of the robot. Thus, the first part of this chapter explores a possibility to estimate the slip and compensate it in the wheel odometry measurements, using a parameterized slip model. The resulting corrected velocity measurement is called *slip-aware wheel odometry*. The performance of the slip-aware wheel odometry is evaluated with respect to the rover's pose estimation algorithm, based on experiments in the context of the ROBEX Moon-analog mission. The experiments imply that the applied parameterized slip model does indeed improve the

correctness of the wheel odometry measurements, but the performance is still stretched to its limits at the Moon-analog test site.

Recapitulating these results, machine learning is identified as a promising approach. Consequently, the second part of the chapter presents the setup and the results of a test campaign which was performed in order to evaluate the use of machine-learning techniques in the context of planetary exploration. Thereby, a deep neural network (DNN) component is combined with a modular tracking control concept. The resulting tracking performance is analyzed in comparison to a conventional tracking controller.

4.1 Slip modeling and estimation: improving the wheel odometry

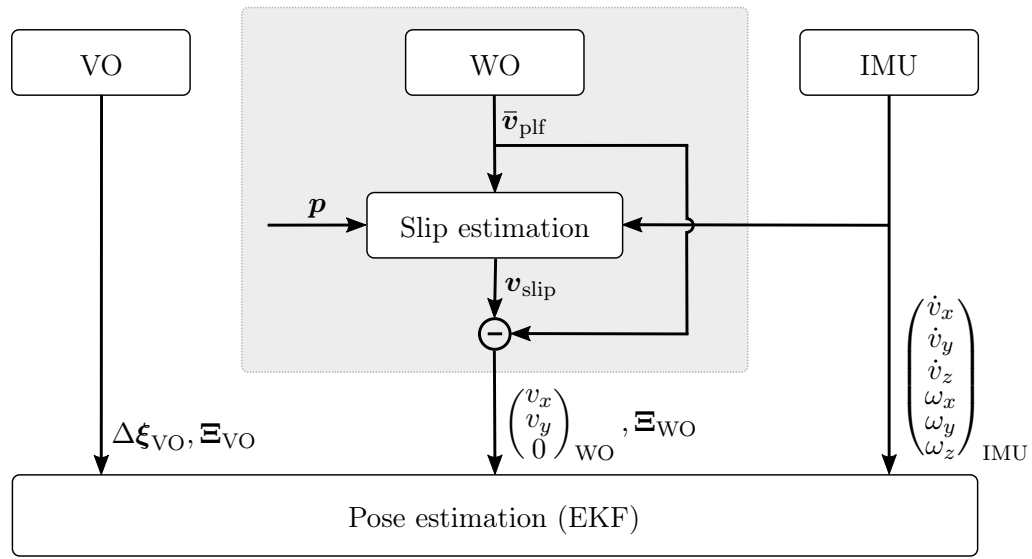


Figure 4.1 Structure of the Extended Kalman Filter (EKF) for pose estimation. The EKF fuses input measurements from the visual odometry (VO), $\Delta\xi_{VO}$ (pose difference between two samples), with measurements from the IMU (the IMU biases are estimated in the filter), and the estimated translational velocity from the wheel odometry (WO), considering the measurement noise Ξ_{VO} and Ξ_{WO} . The WO computation is extended by a slip estimation module, which models the slip on body level using the parameters \mathbf{p} . Note that only translational velocities from the WO are used in the EKF. The WO measurements are computed in the instantaneous plane of motion, thus $v_z = 0$.

In a planetary exploration scenario, the rover is used to explore previously unknown terrain without ground-truth position measurements such as GPS. In this case, the rover gradually builds a map of its environment and localizes itself in this map, using simultaneous localization and mapping (SLAM) techniques. An elaboration of the collaborative localization and mapping algorithms used on the LRUs is given in [Sch+19].

In general, the rover pose w. r. t. a fixed frame in the global map is estimated by an on-board fusion of different sensor signals using an Extended Kalman Filter (EKF). The filter concept is based on the concepts introduced in [Sch+12] and [SRB14], and is used for self-localization in unknown environments. It is clear that the quality of the self-localization is a crucial factor for the success of autonomous robotic exploration tasks.

The structure of the EKF can be inspected in Figure 4.1. The Pose Estimation module (the actual EKF) receives input values from three main sensor units: The first input are measurements

from the visual odometry (VO), which consist of 3D-pose differences $\Delta\boldsymbol{\xi}_{\text{VO}}$ in a comparatively low frequency of 14 Hz. The second input are measurements of the translational accelerations and the rotational velocities of the rover body, provided by an IMU at 100 Hz. Last, the translational velocity computed from the wheel odometry is fed into the EKF. The latter two inputs feature high rates of 100 Hz, but also higher measurement noise $\boldsymbol{\Xi}$. In order to incorporate the different sensor characteristics in the filter output, the covariances of each input are provided and processed by the filter algorithm as detailed in [Sch+12].

In order to improve the self-localization in the context of the ROBEX Moon-analog mission on Mt. Etna, the wheel odometry was extended by a model-based slip estimation. The slip estimation module relies on the raw wheel odometry measurements and IMU data, and uses a parameterized model of the occurring slip. The parameters are given by the vector \boldsymbol{p} . The slip model, the calibration procedure and experimental data from Mt. Etna are shown in the following.¹

4.1.1 Parameterized slip model

As mentioned before, the EKF uses measurements of the platform velocity as input. Therefore, instead of modeling the slip of each wheel, slip is modeled on body level, i. e., in the platform body frame. The corrected platform velocity $\boldsymbol{v}_{\text{plf}} \in \mathbb{R}^3$ is modeled as

$$\boldsymbol{v}_{\text{plf}} = \bar{\boldsymbol{v}}_{\text{plf}} - \boldsymbol{v}_{\text{slip}}(\bar{\boldsymbol{v}}_{\text{plf}}, \boldsymbol{p}, \boldsymbol{F}). \quad (4.1)$$

Thereby, the velocity measurement $\bar{\boldsymbol{v}}_{\text{plf}}$ is generated by an inverse kinematics mapping of the wheel velocity measurements to the body velocity, which is referred to as *wheel odometry*, cf. Section 2.1.5. The unknown term $\boldsymbol{v}_{\text{slip}}$ models the slip on body level, which equals the deviation between the actual (ground-truth) and the measured body velocity. It depends on the measured velocity $\bar{\boldsymbol{v}}_{\text{plf}}$, combined with the linear forces acting on the platform $\boldsymbol{F} \in \mathbb{R}^3$, and n_p slip parameters $\boldsymbol{p} = (p_1, \dots, p_{n_p})$.

For the application on the volcanic soil of Mt. Etna, different parameterizations of the velocity error $\boldsymbol{v}_{\text{slip}}$ were compared. In the end, it turned out that the best results were achieved by the six-parameter slip-velocity model of [SK14], where the linear and angular velocity error components are calculated as

$$\begin{aligned} v_{\vec{x}, \text{slip}} &= (p_1 \frac{F_{\vec{x}}}{F_{\vec{z}}} + p_2) \nu \\ v_{\vec{y}, \text{slip}} &= p_3 \frac{F_{\vec{y}}}{F_{\vec{z}}} \nu \\ \vec{\omega}_{\text{slip}} &= (p_4 \frac{F_{\vec{y}}}{F_{\vec{z}}} + p_5) \nu + p_6 \omega \end{aligned} \quad (4.2)$$

Thereby, ν denotes the velocity in the instantaneous driving direction, cf. Figure 2.7 a). The translational velocities and the linear forces acting on the rover body are rotated into the instantaneous driving direction, indicated by the subscripts $(\vec{x}, \vec{y}, \vec{z})$. The relation between the body frame (x, y, z) and the instantaneous driving frame $(\vec{x}, \vec{y}, \vec{z})$ comprises two rotations. First, the body frame (x, y, z) is rotated around the x -axis to the plane which is spanned by the (averaged) position of the four wheels. The rotation between the body frame (which is defined statically at the center the rover body) and the wheel plane is defined by the angle of the bogie actuators, together with the deflection of the passive spring elements of the SEAs, cf. Figure 2.2. The chosen bogie control approach adjusts the bogie positions to $\rho_1 = \rho_2 = \bar{\rho}$ and keeps the body

¹The following section is mainly based on the Master's thesis of Lukas Burkhard (formerly Meyer) [Mey17] and the resulting publication [Bus+18].

aligned to the direction of the gravity. The different passive spring deflections of the SEAs are assumed to be negligible for the modeling of the instantaneous drive frame, although they are used to maintain ground contact with all wheels while traversing difficult terrain. Afterwards, a rotation around the z -axis by the lateral driving angle β is performed, such that

$$\begin{pmatrix} v_{\vec{x}} \\ v_{\vec{y}} \\ \vec{\omega} \end{pmatrix} = \begin{pmatrix} \cos(\beta) & -\sin(\beta) & 0 \\ \sin(\beta) & \cos(\beta) & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\bar{\rho}) & -\sin(\bar{\rho}) \\ 0 & \sin(\bar{\rho}) & \cos(\bar{\rho}) \end{pmatrix} \begin{pmatrix} v_x \\ v_y \\ \omega \end{pmatrix}, \quad (4.3)$$

or, in short,

$$\vec{v}_{\text{plf}} = \mathbf{R}_z(\beta) \mathbf{R}_x(\bar{\rho}) \mathbf{v}_{\text{plf}}. \quad (4.4)$$

Note that in the nominal platform model (assuming no slip), the velocity component perpendicular to the instantaneous driving direction is zero:

$$v_{\vec{x}} \equiv \nu \quad (4.5)$$

$$v_{\vec{y}} \equiv 0. \quad (4.6)$$

However, side-slip effects may be present and are thus modeled by $v_{\vec{y},\text{slip}}$.

Every parameter p_1 to p_6 in (4.2) captures a different physical source of slip. In particular, p_1 and p_3 model slope-dependent longitudinal slip and slope-dependent lateral slip, respectively. The slope dependency arises from the decreased normal contact forces between the wheels and the ground at inclined terrain. Another relevant effect during driving is the plastic deformation of the loose soil. This causes constant (i. e., not slope-dependent) rolling resistance, which is modeled by p_2 . Thus, the term $p_2\nu$ captures both the *compaction resistance* and the *bulldozing resistance* of the soil, see [Bek69, p. 452]. The last three parameters affect the rotational slip $\vec{\omega}_{\text{slip}}$. Thereby, over- and understeering effects are represented by p_4 , p_5 models wheel asymmetry, and p_6 accounts for skidding [SK14].

Remark 15. Of course there exist a great variety of alternative possibilities to parameterize slip models. In order to find a tradeoff between model complexity and performance, a comparison using experimental data obtained from the DLR test facilities was performed [Mey17]. The analysis revealed that the considered alternatives (including the consideration of acceleration dependency, the distinction between uphill and downhill slip, or the 15-parameter approach from [See+13]) did not improve the results significantly. However, they require a higher number of parameters, which implies a greater number of required calibration trajectories and a higher computation complexity. Thus, the presented six-parameter model is considered in the following, as it features a reasonably high modeling accuracy in combination with low computational and experimental effort.

If the rover drives in an area with homogeneous soil, the parameters can be assumed to be constant. However, a constant parameterization cannot be used in strongly varying terrain. Alternatively, online-parameter adaption strategies can be considered. In each of these cases, the question of a proper calibration of the slip parameters arises, which has to be answered whilst taking into account the constraints of the respective scenario. A static calibration algorithm which does not rely on continuous ground truth position measurements is presented in the following.

4.1.2 Calibration of the slip model

The open question from the preceding section is how to define a proper set of parameters \mathbf{p} for the current soil. Most methods to find \mathbf{p} require measurements of the rover pose for all

times $t \in [t_0, t_{\text{end}}]$ of a trajectory. However, in an exploration scenario, such continuous tracking information might not be available. On the other hand, it is relatively simple to measure certain static poses of a vehicle, e. g. by referencing the system to landmarks or static fiducial markers. Thus, for a suitable parameter identification approach, measurements of start and end point of the respective trajectories must be sufficient. One existing approach which needs no continuous information on the real trajectory is the *Integrated Prediction Error Minimization* (IPEM) approach introduced in [See+13]. The theoretical background of IPEM is recapitulated from [See+13] in the following.

The main idea is that the set of parameters \mathbf{p} can be identified by comparing the predicted output of the nominal system model with the measured output of the real system at the end of a driven trajectory with known start point. Thereby, it is only required to measure the start and end point of a trajectory w. r. t. a fixed frame, and to record the system inputs \mathbf{u} during the drive.

In general, the pose $\boldsymbol{\zeta} \in \mathbb{R}^6$ of the rover driving in some arbitrary terrain evolves according to the model

$$\dot{\boldsymbol{\zeta}}(t) = \mathbf{f}(\boldsymbol{\zeta}(t), \mathbf{v}_{\text{plf}}), \quad (4.7)$$

$$\boldsymbol{\xi}(t) = \mathbf{h}(\boldsymbol{\zeta}(t)). \quad (4.8)$$

The solution of (4.7) can be obtained by integration w. r. t. time as

$$\boldsymbol{\zeta}(t) = \boldsymbol{\zeta}(t_0) + \int_{t_0}^t \mathbf{f}(\boldsymbol{\zeta}(\tau), \mathbf{v}_{\text{plf}}(\mathbf{p}, \tau)) d\tau. \quad (4.9)$$

The output measurement $\boldsymbol{\xi}$ is the pose $\boldsymbol{\zeta}(t_{\text{end}})$ at the end of a driven trajectory (in general, the dimension of $\boldsymbol{\xi}$ is $n_{\xi} \leq 6$, depending on the measurement method). The platform velocity \mathbf{v}_{plf} serves as system input in this model. Using system identification techniques, the parameters \mathbf{p} can be determined such that modeling errors are minimized. As only the start point and the end point can be measured, the dynamic equations (4.7)–(4.8) have to be integrated until the end of the trajectory at t_{end} :

$$\begin{aligned} \boldsymbol{\zeta}(t_{\text{end}}) &= \boldsymbol{\zeta}(t_0) + \int_{t_0}^{t_{\text{end}}} \mathbf{f}(\boldsymbol{\zeta}(\tau), \mathbf{v}_{\text{plf}}(\mathbf{p}, \tau)) d\tau \\ &= \mathbf{g}(\boldsymbol{\zeta}(t_0), \mathbf{v}_{\text{plf}}(\mathbf{p}, \cdot)), \end{aligned} \quad (4.10)$$

where $\mathbf{v}_{\text{plf}}(\mathbf{p}, \cdot)$ denotes the system input for all $\tau \in [t_0, t_{\text{end}}]$. The system output is defined as the (ground-truth) measurement of the end point of a trajectory:

$$\boldsymbol{\xi}_{\text{msr}} = \boldsymbol{\zeta}_{\text{msr}}(t_{\text{end}}). \quad (4.11)$$

The predicted system output depends on the parameters \mathbf{p} and is written as

$$\boldsymbol{\xi}_{\text{est}} = \boldsymbol{\zeta}_{\text{pred}}(t_{\text{end}}). \quad (4.12)$$

Then, the error $\boldsymbol{\xi}_{\text{err}}$ between the estimated and the real output is derived as

$$\boldsymbol{\xi}_{\text{err}} = \boldsymbol{\xi}_{\text{msr}} - \boldsymbol{\xi}_{\text{est}} = \boldsymbol{\xi}_{\text{msr}} - \mathbf{g}(\boldsymbol{\zeta}(t_0), \mathbf{v}_{\text{plf}}(\mathbf{p}, \cdot)). \quad (4.13)$$

Now let $\mathbf{p}^* \in \mathbb{R}^{n_p}$ be an optimal set of parameters which minimizes the output error, such that $\boldsymbol{\xi}_{\text{err}}(\mathbf{p}^*) \approx 0$. Linearizing (4.13) at $\mathbf{p} = \mathbf{p}^*$ yields

$$\boldsymbol{\xi}_{\text{err}} \approx \left. \frac{\partial \mathbf{g}(\boldsymbol{\zeta}(t_0), \mathbf{v}_{\text{plf}}(\mathbf{p}, \cdot))}{\partial \mathbf{p}} \right|_{\mathbf{p}=\mathbf{p}^*} (\mathbf{p}^* - \mathbf{p}) = \mathbf{H} \Delta \mathbf{p}. \quad (4.14)$$

The matrix

$$\mathbf{H} = \left(\frac{\partial \mathbf{g}}{\partial p_1} \dots \frac{\partial \mathbf{g}}{\partial p_{n_p}} \right) \in \mathbb{R}^{n_\xi \times n_p} \quad (4.15)$$

is called the *parameter Jacobian matrix* of a trajectory and models the influence of each parameter on the estimated end point of the corresponding trajectory. It is then possible to determine \mathbf{p} by iteration according to

$$\mathbf{p}_{k+1} = \mathbf{p}_k + \alpha \Delta \mathbf{p}_k. \quad (4.16)$$

Several update steps k can be performed with the update rate $\alpha \in]0, 1]$ in order to obtain a static calibration of \mathbf{p} . Alternatively, a recursive adaption of \mathbf{p} during driving is discussed in [See+13]. This approach uses an EKF to estimate \mathbf{p} by considering piecewise sections of the drive, comparing it to directly available ground truth measurements.

For a successful parameter calibration, the influence of measurement noise should be kept as small as possible, and overfitting should be avoided. Therefore, (4.14) should be rendered overdetermined by a number of n_t appropriate trajectories such that $n_\xi \cdot n_t \gg n_p$. Stacking the respective parameter Jacobian matrices yields

$$\Delta \mathbf{p} = \begin{pmatrix} \mathbf{H}_1 \\ \vdots \\ \mathbf{H}_{n_t} \end{pmatrix}^{\mathbf{W}^+} \begin{pmatrix} \boldsymbol{\xi}_{\text{msr},1} - \mathbf{h}_1(\mathbf{p}) \\ \vdots \\ \boldsymbol{\xi}_{\text{msr},n_t} - \mathbf{h}_{n_t}(\mathbf{p}) \end{pmatrix}. \quad (4.17)$$

The s.p.d. weighting matrix $\mathbf{W} \in \mathbb{R}^{(n_\xi \cdot n_t) \times (n_\xi \cdot n_t)}$ offers the possibility to weight the trajectories against each other, for example if the measurements accuracy differs (e.g. due to different measurement methods).

In order to find an analytical expression for the parameter Jacobian matrix \mathbf{H} , the following simplifications can be considered:

1. Measurement of the rover start and end pose $\boldsymbol{\xi}$ in a fixed 2D-frame ${}^0(x, y, \theta)$.
2. Linearization of (4.7) around the predicted trajectory, evaluated at the present parameter estimates [See+13].

Applying these modifications, the parameter Jacobian matrix can be derived analytically as follows. The first simplification implies that only measurements in an (inclined) 2D plane are required, which is a direct consequence of the endpoint-only dependency of IPeM, and simplifies the measurement procedure. However, the driving itself must be done in arbitrary terrain in order to capture slope-dependent slip effects. This principle is sketched in Figure 4.2 a). In this case, only the planar part of $\boldsymbol{\zeta}$ is considered in (4.7) ff., i.e., $n_\xi = 3$ in the following.

For the second simplification, (4.7) is linearized according to

$$\delta \dot{\boldsymbol{\zeta}}(t) = \mathbf{F}(t) \delta \boldsymbol{\zeta}(t) + \mathbf{G}(t) \delta \mathbf{v}_{\text{plf}}(t) \quad (4.18)$$

with

$$\mathbf{F}(t) = \frac{\partial \mathbf{f}}{\partial \boldsymbol{\zeta}}, \quad \mathbf{G}(t) = \frac{\partial \mathbf{f}}{\partial \mathbf{v}_{\text{plf}}}. \quad (4.19)$$

Then, the general solution for linear systems (see e.g. [Sch08, p. 36]) is used to compute the solution of (4.18) to

$$\delta \boldsymbol{\zeta}(t) = \boldsymbol{\Phi}(t, t_0) \delta \boldsymbol{\zeta}(t_0) + \int_{t_0}^t \boldsymbol{\Phi}(t, \tau) \mathbf{G}(\tau) \delta \mathbf{v}_{\text{plf}}(\mathbf{p}, \tau) d\tau, \quad (4.20)$$

introducing the *state transition matrix* $\Phi \in \mathbb{R}^{3 \times 3}$. Finally, (4.10) can be re-written as [Kel04]

$$\begin{aligned}\zeta(t_{\text{end}}) &= \mathbf{g}(\zeta(t_0), \mathbf{v}_{\text{plf}}(\mathbf{p}, \cdot)) = \\ &= \Phi(t_{\text{end}}, t_0)\zeta(t_0) + \int_{t_0}^{t_{\text{end}}} \Phi(t_{\text{end}}, \tau) \mathbf{G}(\tau) \mathbf{v}_{\text{plf}}(\mathbf{p}, \tau) d\tau.\end{aligned}\quad (4.21)$$

Inserting (4.21) into (4.15), and moving the derivative inside the integral yields the following analytic expression for the parameter Jacobian matrix:

$$\mathbf{H} \approx \int_{t_0}^{t_{\text{end}}} \Phi(t_{\text{end}}, \tau) \mathbf{G}(\tau) \frac{\partial \mathbf{v}_{\text{plf}}(\mathbf{p}, \tau)}{\partial \mathbf{p}} d\tau. \quad (4.22)$$

Note that for this derivation, a linearized system model is considered and thus the resulting parameter Jacobian matrix is an approximation. In general, (4.22) is only valid for small changes in the parameters. The computation of the components of \mathbf{H} in (4.22) is explained in the following.

First, the equations of motion (4.7) for a planar rover pose are

$$\dot{\xi} = \mathbf{R}_z(\theta) \mathbf{v}_{\text{plf}}, \quad (4.23)$$

which means that the linearized system matrices \mathbf{F} and \mathbf{G} can be derived from (4.19) as

$$\mathbf{F} = \begin{pmatrix} 0 & 0 & -\sin(\theta)v_x - \cos(\theta)v_y \\ 0 & 0 & \cos(\theta)v_x - \sin(\theta)v_y \\ 0 & 0 & 0 \end{pmatrix} \quad (4.24)$$

and

$$\mathbf{G} = \mathbf{R}_z(\theta). \quad (4.25)$$

Following the derivations in [Kel04], the state transition matrix Φ can be computed using the matrix integral

$$\Psi(t, \tau) = \int_{\tau}^t \mathbf{F}(\tilde{\tau}) d\tilde{\tau} = \int_{\tau}^t \begin{pmatrix} 0 & 0 & -{}^0\dot{y}(\tilde{\tau}) \\ 0 & 0 & {}^0\dot{x}(\tilde{\tau}) \\ 0 & 0 & 0 \end{pmatrix} d\tilde{\tau} = \begin{pmatrix} 0 & 0 & -({}^0y(t) - {}^0y(\tau)) \\ 0 & 0 & {}^0x(t) - {}^0x(\tau) \\ 0 & 0 & 0 \end{pmatrix}. \quad (4.26)$$

Since Ψ is strictly upper rectangular,² the transition matrix can then be computed as

$$\Phi(t, \tau) = e^{\Psi(t, \tau)} = \mathbf{I} + \Psi(t, \tau) = \begin{pmatrix} 1 & 0 & -({}^0y(t) - {}^0y(\tau)) \\ 0 & 1 & {}^0x(t) - {}^0x(\tau) \\ 0 & 0 & 1 \end{pmatrix}. \quad (4.27)$$

It models the propagation of the wheel odometry error over the time and its impact on the predicted pose. In fact, the difference terms in the last row of the transition matrix account for the fact that errors in the rotational odometry measurement at the beginning of the trajectory have an higher impact on the predicted pose than measurement errors occurring close to the end of the trajectory. Second, due to the linear slip model (4.2), (4.1) can be written as

$$\mathbf{v}_{\text{plf}}(\mathbf{p}) = \bar{\mathbf{v}}_{\text{plf}} - \mathbf{C}(\bar{\mathbf{v}}_{\text{plf}}, \mathbf{F})\mathbf{p}, \quad (4.28)$$

²In general, the matrix exponential of a matrix \mathbf{A} can be computed by the infinite power series $\exp(\mathbf{A}) = \mathbf{I} + \mathbf{A} + \mathbf{A}^2/2! + \mathbf{A}^3/3! + \dots$. The strict upper rectangularity means that \mathbf{A}^2 and all higher-level powers of \mathbf{A} vanish, such that $\exp(\mathbf{A}) = \mathbf{I} + \mathbf{A}$.

with the *parameter influence matrix*

$$\mathbf{C} = \begin{pmatrix} \frac{F_x}{F_z}\nu & \nu & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{F_y}{F_z}\nu & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{F_y}{F_z}\nu & \nu & \omega \end{pmatrix}. \quad (4.29)$$

Using these prerequisites, (4.22) can be written as

$$\mathbf{H} \approx - \int_{t_0}^{t_{\text{end}}} \boldsymbol{\Phi}(t_{\text{end}}, \tau) \mathbf{R}_z(\theta(\tau)) \mathbf{C}(\tau) d\tau. \quad (4.30)$$

Remark 16. If 6 DOF ground truth measurements are available, it is possible to apply a 2.5D approach, where the pose is projected on a 2D plane. The advantage is that the test site for trajectory measurements is not restricted to planar, inclined surfaces anymore. This approach is described in detail in [See+13].

Remark 17. An alternative to the estimated analytical parameter Jacobian matrix from (4.30) is the use of a numerical derivative in (4.15), by computing

$$\frac{\partial \mathbf{g}}{\partial p_i} \approx \frac{\mathbf{g}(\boldsymbol{\zeta}(t_0), \mathbf{v}_{\text{plf}}(\mathbf{p} + \delta \mathbf{p}_i, \cdot)) - \mathbf{g}(\boldsymbol{\zeta}(t_0), \mathbf{v}_{\text{plf}}(\mathbf{p}, \cdot))}{\epsilon}, \quad i = 1, \dots, n_p. \quad (4.31)$$

Thereby, the vector $\delta \mathbf{p}_i \in \mathbb{R}^{n_p}$ introduces a small change ϵ on the i th parameter, all other entries are zero. The advantage of the numerical approach is that it allows for a use of the full 6-DOF rover pose, if a respective measurement is available. On the one hand, this can improve the precision of the results. On the other hand, it increases the computational requirements, as each trajectory has to be evaluated n_p times. Moreover, the numerical differentiation can lead to numerical errors. Thus, using the analytical solution (4.30) is advantageous in applications where fast computation is required, for example online parameter estimation.

In the following, an experimental validation of the proposed approach for slip estimation is shown, which was performed the ROBEX mission on Mt. Etna.

4.1.3 Experimental validation of the slip-compensated wheel odometry

The test site on Mt. Etna is located west of the Cisternazza Crater and features volcanic soil from an eruption in 2001 [Lan+03]. The area is considered to be Moon-analog [PGB12]. The soil mainly consists of sharp-edged gravel with a median particle diameter of approx. 0.5-1 cm.

Before performing the actual calibration of the slip parameters, an additional hardware-specific effect was observed at the LRU1 and incorporated into the mathematical system model. Due to the elastic design of the spokes (cf. Section 2.1.1), torsional compliance is introduced between the steering actuators and the surface of the wheels. In loose soil, this mechanical property causes a systematic error between the commanded and the measured steering angles. The correlation between the steering actuator torque and the corresponding offset of the steering angle $\Delta \phi_i, i = 1, \dots, 4$ was determined experimentally. The measurement data and the interpolated stiffness curve derived from the measurements can be inspected in Figure 4.5 b).

The next step was the calibration of the slip model as described in the preceding section. The calibration area is shown in Figure 4.3. The trajectories cover approximately 20 m in both directions from the initial frame, and are selected such that terrain with different slopes is passed. However, an almost planar area with $\approx 5^\circ$ inclination is chosen for the measurement of start

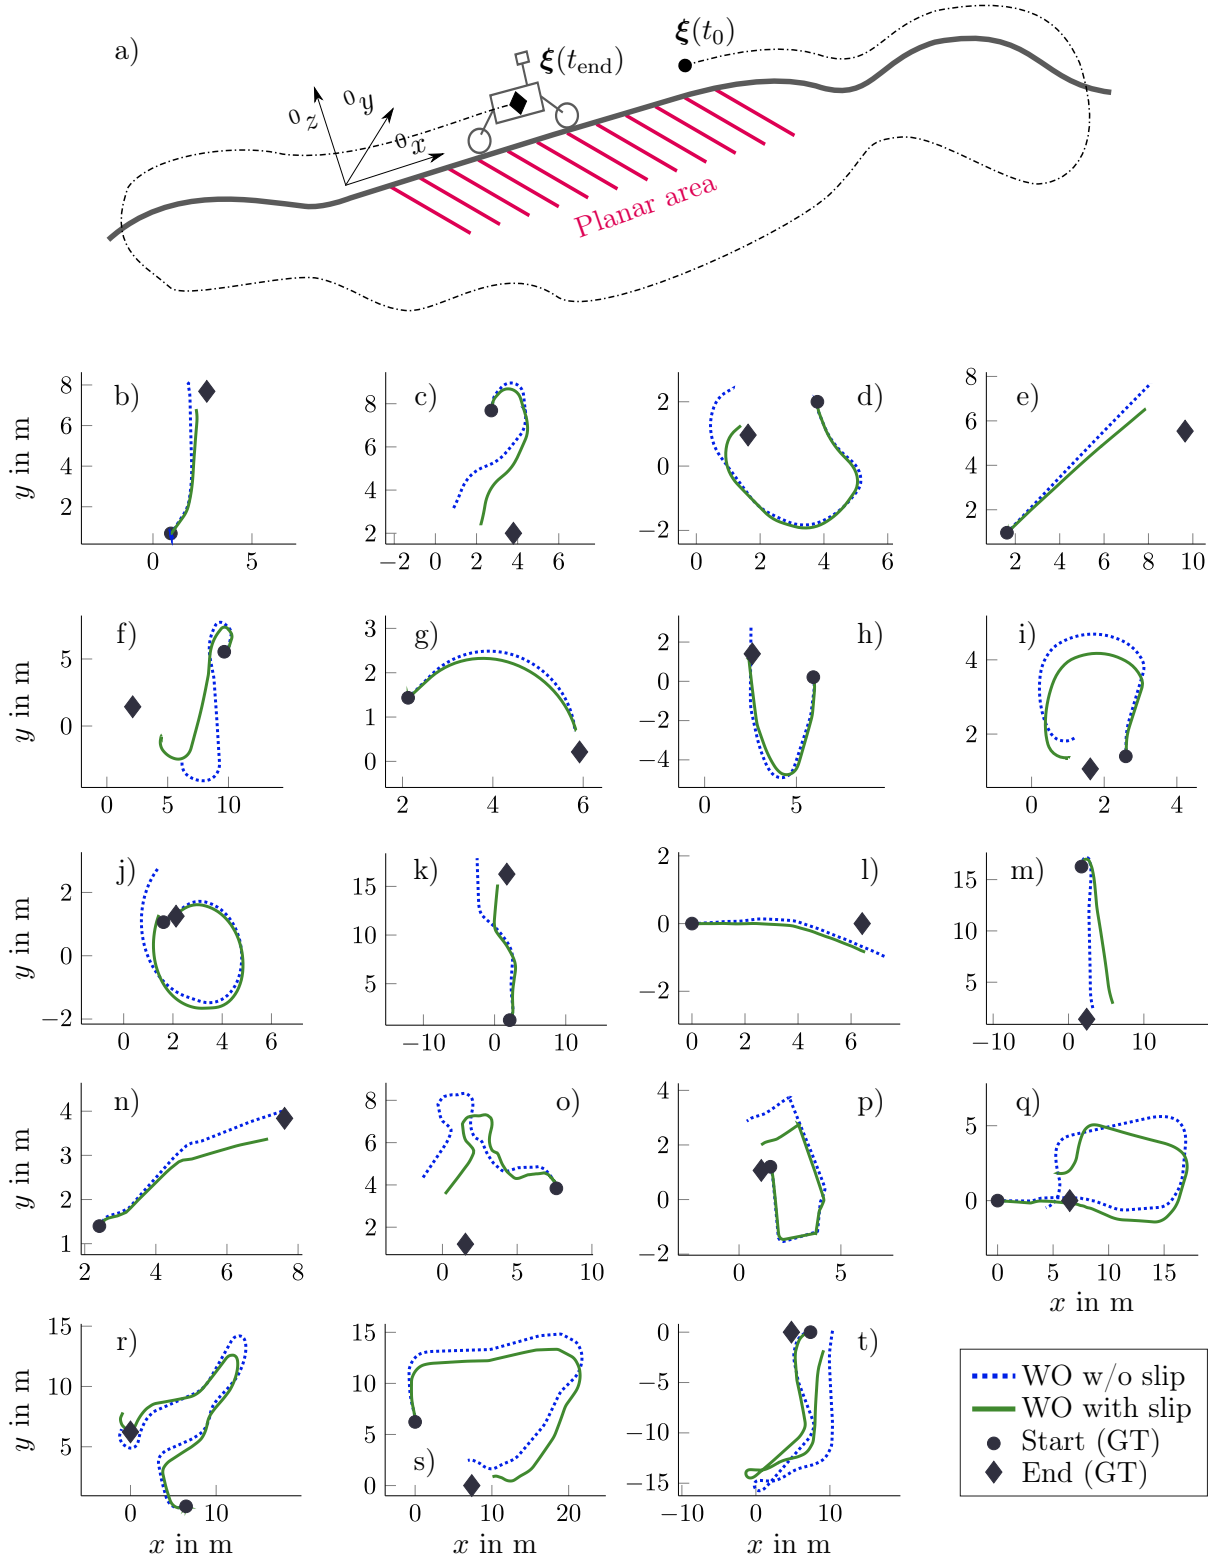


Figure 4.2 a) Setup of the calibration: A fixed coordinate system is constructed on a planar area next to the origin. b)–t) Set of 19 calibration trajectories, used for calibrating the slip model on Mt. Etna. Start and end points are measured w. r. t. the fixed frame and serve as ground-truth (GT) measurements for the calibration. The wheel odometry measurements are integrated to position measurements, and plotted before (without slip, blue dashed) and after (with slip, green) the calibration of the slip parameters. From [Bus+18].

	Value	Unit
p_1	0.6367	—
p_2	0.1089	—
p_3	0.2465	—
p_4	-0.0076	rad/m
p_5	0.0061	rad/m
p_6	0.0733	—



Table 4.1 Calibrated slip parameters.
From [Bus+18].

Figure 4.3 The LRU1 at the calibration site on Mt. Etna. A static reference coordinate system is created manually at a sloped, approximately planar test area.

and end points of the trajectories, as sketched in Figure 4.2 a). Figure 4.2 b)–t) shows the set of 19 calibration trajectories that were used to obtain the six slip parameters shown in Table 4.1. The length of the calibration trajectories varies between 5 m and 40 m. The trajectories were selected to cover a variety of the most typical driving situations of the rover at this particular Moon-analog test site. Some of the most important features of a suitable calibration dataset were identified as

- All modeled types of slip should be present in the data set
- Different command modes should be represented, in particular “car-like” driving ($v_{y,\text{des}} = 0$) and “crab-like” driving (sideways translation with $\omega_{\text{des}} = 0$), as these modes are employed during the autonomous rover operations in the context of the ROBEX mission.

Moreover, a subset of the trajectories does explicitly not contain loop closures, as they can distort the calibration result.³ The shorter trajectories b)–p) were chosen such that one of the modeled slip effects is dominant in each trajectory. The longer trajectories q)–t) contain a combination of all slip effects. The slip parameters were calibrated w. r. t. (4.16), with an update rate of $\alpha = 0.7$, a step size of 0.01 s, and 5 iterations. The trajectories o) and s) were weighted with a factor of 0.5 in the calibration with regard to their total length. The result of the calibration can be inspected in Table 4.1. It can be observed that the parameters p_1 and p_3 describing the slope-dependent slip have the highest values, while the rotational slip parameters are rather small. However, a comparison of the parameter values does only provide limited information, as they are multiplied with different terms in (4.2).

In a second step, the calibrated slip-aware wheel odometry was validated by performing a teleoperated long-range drive at the ROBEX test site on Mt. Etna. The driven path distance is approx. 900 m. A plot of the DGPS position data embedded in the terrain of the test site is shown in Figure 4.4. Starting next to the base camp at ①, the rover drives nearly lateral to the slope of the terrain, thereby crossing the calibration area. At ②, the path turns to follow the slope downhill for

³If the starting point of a trajectory equals its end point, the scaling of the trajectory is undefined.




Figure 4.4 Overview of the test site with DGPS data of the long-range run (yellow line). The numbers indicate selected regions along the trajectory. The length of the trajectory is approximately 900 m. From [Bus+18].

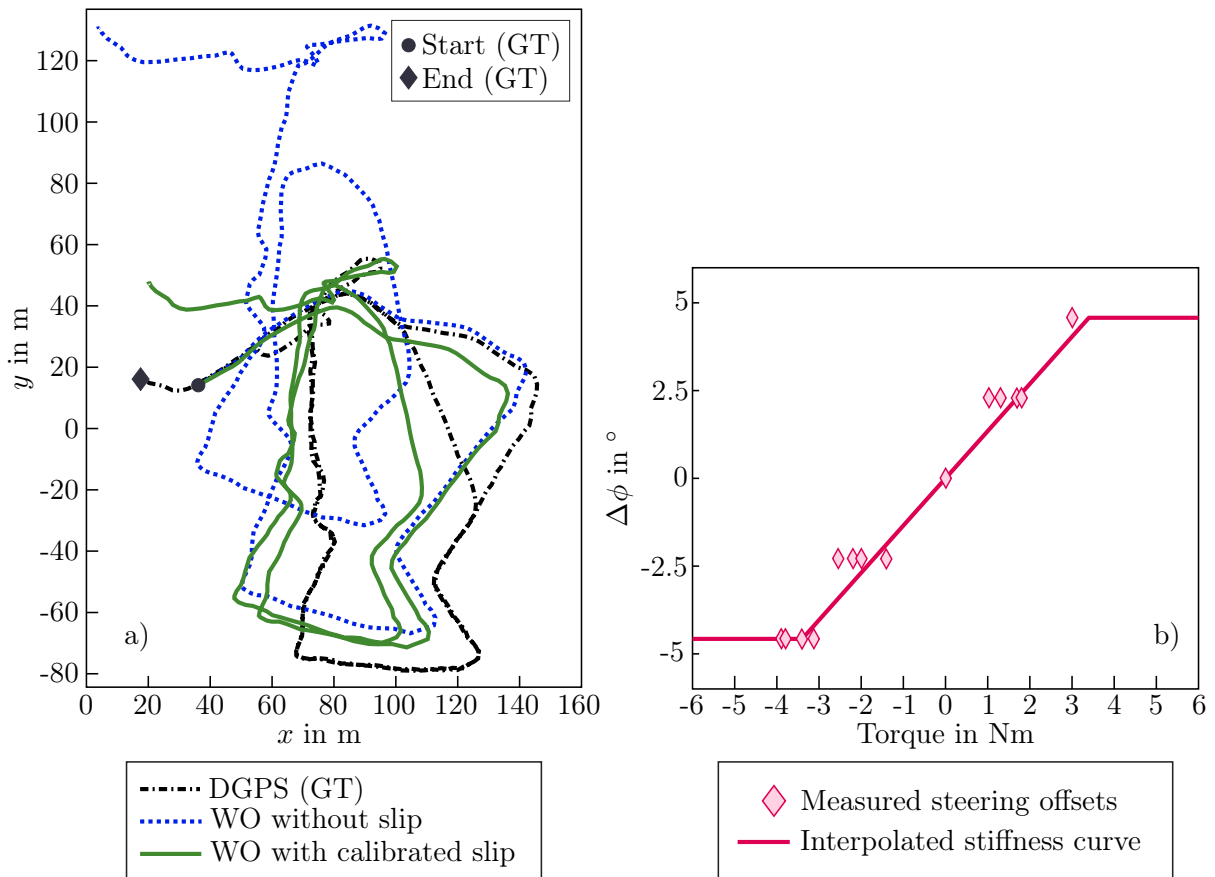


Figure 4.5 a) Wheel odometry position measurements (integrated velocity measurements) of the long-range drive, compared to DGPS data (ground truth). The plot shows the integral of the translational velocities from the WO, combined with the angular velocity from the EKF output. b) Saturated linear spring characteristics, derived from the measurements of the steering angle offset of the four wheels $\Delta\phi_i$, $i = 1, \dots, 4$ for different applied steering torques. From [Bus+18].




Figure 4.6 Data from the long-range run on Mt. Etna. a) The distance error is plotted against the traversed distance, for different WO computations. Uphill passages are highlighted in yellow. b) The value of the six slip parameters, obtained by a simulated online calibration. At approx. 400-425 m, the rover passes loose soil in the vicinity of the calibration area (highlighted in green). From [Bus+18].

approx. 100 m ②, passing close to the rim of the Cisternazza Crater on the left side of the vehicle. At the lowest point, a relatively flat area is traversed ③, until the lap is completed by following the uphill part ④. Afterwards, a second lap (with small deviations at ②) is completed following the same path.⁴ Finally, a third small circle is driven close to ① before returning to the base camp.

Figure 4.5 a) shows the integrated wheel velocity measurements. Note that for the sake of clarity, only the integrals of the translational velocity components from the WO are plotted, combined with the rotation output from the EKF. This representation was chosen due to the structure of the EKF (cf. Figure 4.1), where only the translational velocities from the WO are used for the sensor fusion. Thus, only the translational WO measurements are relevant for the pose estimation. The plot shows that the driven distance is significantly overestimated without the slip estimation (blue dotted line). This is particularly evident when the rover drives uphill, e.g. around ④. The slip-compensated wheel odometry (green solid line) looks closer to the ground truth DGPS measurement (black dot-dashed line). However, the driven distance is slightly underestimated at all parts of the trajectory. As this effect appears independent of the slope, the hypothesis is that the constant rolling resistance is overestimated, i. e., p_2 was estimated too high during the calibration, whereas the values of p_1 and p_3 are sufficiently exact. Figure 4.6 a) verifies this assumption. The plot shows the accumulated error between the estimated and the real driven distance, plotted over the absolute driven distance. Uphill passages are highlighted in yellow. Again, it can be seen that the WO without slip (blue dotted) is strongly slope-dependent, and overestimates the driven distance in general. The slip-aware WO (solid green) underestimates the distance nearly independent of the terrain inclination. To verify the hypothesis of the overestimated rolling resistance, another wheel odometry measurement was computed offline, with a manually tuned parameter $p_2 = 0.036$. The resulting distance errors is plotted as solid magenta line. Clearly, adapting the parameter p_2 as mentioned before removes the main part of the accumulated error, which confirms the hypothesis.

Remark 18. One explanation for the overestimation of the constant rolling resistance could be that the calibration area was close to the base camp and thus frequently passed by persons. This loosened the soil and increased the rolling resistance, causing increased wheel slippage in this particular area. Obviously, this problem would be eliminated in space applications.⁵

As a consequence of the observed calibration errors, an online adaption of the parameters was considered as a possibility to overcome this problem. As a validation of this concept, the online parameter estimation method proposed by [See+13] was applied to the data, using DGPS data as ground truth. Thereby, another EKF is used to estimate the parameters, using piece-wise sections of the trajectory of approximately 2 m length. The calibrated parameters from Table 4.1 were used as starting values. The resulting distance error is plotted as yellow dot-dashed line, the corresponding parameter values can be inspected in Figure 4.6 b). It can be seen that all parameters except p_2 do not change their value significantly over the driven trajectory, which confirms the result of the offline calibration procedure. Noticeably, the value of p_2 decreases over the driven distance, supporting again the hypothesis of the loosened soil at the calibration site. It is visible that around 400-425 m (close to ①) the rover crossed the calibration site, causing a local rise of the parameter p_2 which is highlighted in green in Figure 4.6 b). Note that during the last 150 m of the trajectory, a loss in DGPS precision can be observed, which negatively affects

⁴Note that the deviations visible at ④ in the second lap are a result of a loss in DGPS precision at the last 150 m of the trajectory.

⁵While repetitive passing may in theory change the soil compaction and cause a change in the soil parameters, Ding et al. show in an experimental study (different rover wheels were tested in a single-wheel testbed filled with a lunar soil simulant), that the influence of repetitive passing on wheel–soil interaction models is negligible [Din+11].

the parameter estimation. Due to the slow convergence of the parameter p_2 over the course of the trajectory, the driven distance is still underestimated, as Figure 4.6 a) shows. Nevertheless, it can be concluded that the online calibration can correct erroneous calibration results over the course of longer driving distances, and cope with (slowly) varying terrain properties. However, it is important to highlight that the online parameter adaption can only be performed in the case that reference measurements of the rover position are available. In the described software architecture, the position data from the self-localization cannot simply be used as a reference, as this would induce cross-correlations and disturb the estimation results. If on the other hand the slip estimation is only used in the context of a control-related slip compensation, a reference signal fused from IMU and VO can be used, as proposed e. g. in [Hel+05] and [Hel+06]. In the following, a concept for a controller architecture including a slip compensation is briefly introduced and discussed.

4.1.4 Towards slip-compensated control

The objective of the slip estimation which was presented in the preceding section was to compensate the slip effects in the wheel odometry measurements, in order to improve the rover's self localization. The results indicate that slip can be estimated, at least to a certain extent, by the described parameterized approach (assuming that the soil properties are sufficiently constant). Consequently, the question arises if and how the obtained slip estimation can be used in order to improve the tracking performance of the platform controller. Although first steps in this direction were developed and examined in [Mey17], the concepts were not elaborated and extensively tested in the space-analog context. However, for the reason of completeness, Figure 4.7 shows the proposition of a concept for an extension of the platform controller by a slip compensation module. The upper layer denotes the high-level software, including the path planner that generates a desired body velocity and the pose estimation (cf. Figure 4.1). The lower layer contains the rover hardware with the motors and the sensors. In between, the control layer is located, which comprises of the extended slip estimation as presented before, and the proposed slip-aware platform velocity controller. The architecture makes sure that there is no cross-correlation between the slip-compensated input of the EKF and the slip estimation which is used for control. Therefore, the computed velocities from the extended wheel odometry are compared to the velocity output of the pose estimation $\mathbf{v}_{\text{plf,EKF}}$, resulting in an enhanced estimation of the body slip, $\hat{\mathbf{v}}_{\text{slip}}$. The latter can then be used for control purposes, which is denoted as *slip compensation* in the block diagram. In practice, this slip compensation can be understood in different ways, depending on the control objective. One possibility is to reduce the total body slip $|s|$, which can be defined as

$$|s| = \begin{cases} \frac{\|\mathbf{v}_{\text{slip}}\|_2}{\|\bar{\mathbf{v}}_{\text{plf}}\|_2} & \text{for } \|\bar{\mathbf{v}}_{\text{plf}}\|_2 > c, \\ 0 & \text{otherwise} \end{cases} \quad (4.32)$$

with $c > 0$ a constant, small threshold velocity which avoids the singularity for $\|\bar{\mathbf{v}}_{\text{plf}}\| \rightarrow 0$. A valid approach in this case is to adjust the platform velocity command based on the estimated slip velocity $\hat{\mathbf{v}}_{\text{slip}}$, such that the amount of slip is reduced. As this approach targets to an improved traction, it is referred to as *traction control* in [Mey17].

Another possibility is to actually compensate for the estimated slip velocity by adding it to the velocity command. In this case, the result is an improvement of the tracking performance, at the cost of an increased amount of slip and thus an increased energy consumption. Clearly, the choice of a suitable control objective and thus algorithm depends on the scenario and the constraints that come with it. As energy efficiency was considered more important than precise velocity tracking, the traction control approach was employed and verified during the ROBEX

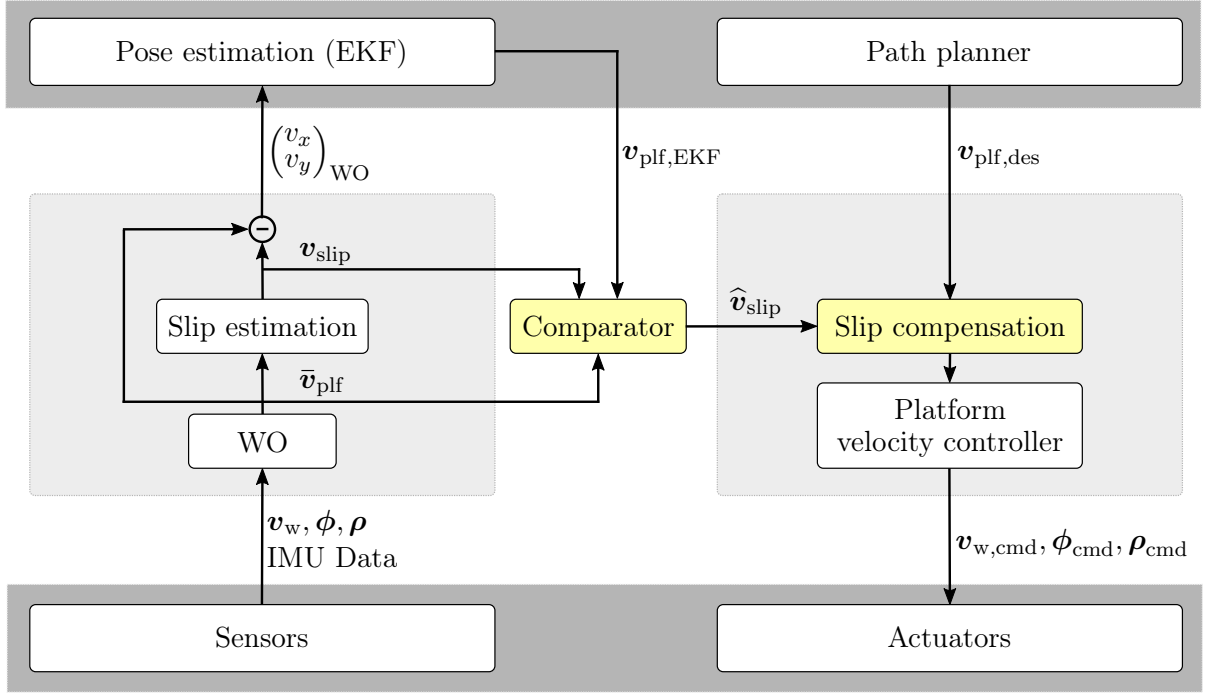


Figure 4.7 Concept of a slip-aware controller module within the LRU control and self-localization framework. The control layer is located between the high-level software (pose estimation, path planner) and the hardware layer (sensors and actuators). The idea is to compare the velocities computed by a static slip model with the EKF output, in order to obtain an enhanced slip estimation \hat{v}_{slip} . This value can then be used for slip compensation within the platform velocity controller.

mission. Some results are presented in [Mey17]. However, in different scenarios (e. g. when driving in confined environments, or heading for a dedicated object of interest), the tracking performance can be more important than a reduction of the total slip velocity. Thus, a slip-compensated tracking controller is presented in the further course of this chapter.

4.1.5 Résumé

Summarized, the presented approach has shown promising results for estimating slip on body level. Experiments in the course of a space-analog mission show that the approach is suitable in a planetary exploration scenario, where no continuous ground truth position measurements are available. It is possible to perform the calibration of the slip parameters in a fully autonomous way, for example after landing on a foreign planet close to the landing site by reference to visual markers on the lander, or expressive landmarks close to the landing site. In that case, a parameterized slip model is provided, which can be used for improving the wheel-odometry measurements, facilitating the exploration of distant areas. However, modeling the slip with a constant parameter model requires homogeneous soil and thus constant soil properties. Data collected at the Moon-analog test site on Mt. Etna shows that this assumption is not always justified, especially when larger distances are traversed. Thus, an online adaption of the parameters has been identified as promising approach. In this case, the wheel odometry measurement cannot be fused into the self localization. Instead, methods like those proposed in [Hel+05] or [Hel+06] could be used, where the wheel odometry is compared and corrected w. r. t. a reference signal from fusing only IMU and VO measurements. Moreover, WO can be used as backup sensor in

the self localization if one of the reference sensors fails. An alternative architecture is briefly introduced in Section 4.1.4, where an enhanced slip measurement is derived by comparing the WO output to the EKF output, which can then be used in the controller.

Nevertheless, the experiments show that there is room for improvement. Comparing the results obtained with the LRU1 on Mt. Etna to the results from [See+13] and [SK14], it can be noticed that the former show less improvements in the wheel odometry measurements than the results presented in the aforementioned publications. A possible explanation is that the proposed 6-parameter slip model is tailored to skid-steered vehicles, which create the vehicle yaw rate by counter rotation of opposing wheel pairs. For these kind of vehicles, rotational slip effects like over-and understeering and wheel asymmetry are obviously more significant. Also, the volcanic soil on Mt. Etna features different properties than the muddy arable lands, grasslands and paved roads which were used in the mentioned publications. In general, it can be assumed that there might exist also non-linear correlations between body velocity and slip, which are not captured by the linear slip model (4.2). Therefore, a learning-based approach for the compensation of unmodeled wheel-ground contact effects is presented in the following section.

4.2 Learning-based trajectory tracking control

In the previous section, it has been shown that a parameterized static model-based slip estimation still leads to significant deviations between wheel odometry measurements and ground truth. In [Vay+18], a detailed analysis of the performance of the rover's on-board pose estimation is performed on a data set of another long-range drive on Mt. Etna. It is shown that the slip-aware wheel odometry measurements improve the self localization. Still, the best results are obtained by leaving the wheel odometry out of consideration in nominal driving situations. This finally led to the conclusion that there is a need for adaptive algorithms for slip estimation, which can capture a broader range of wheel-ground interaction effects. Machine-learning techniques were identified as a promising approach in this context [Bus+18].

Therefore, a modular platform tracking controller containing a learning-based component for the estimation and compensation of unmodeled wheel-ground interaction effects is described in the following. The control architecture builds upon a nonlinear trajectory tracking controller, which is then extended by a (learning-based) model predictive controller (MPC). Model predictive control is a well-established concept in the control of autonomous vehicles, which allows to compute a control input that optimizes a given cost function, considering the computed nominal system dynamics over a given prediction horizon. For an overview of the principles of MPC, the reader is referred to the extensive literature. Es examples, [May14] and [GP11] can be mentioned. Existing variants of MPC combined with machine learning are given in the overview publications [Mes+22] and [Hew+20]. In this work, the errors of the state estimation of the MPC are learned by a deep neural network (DNN), which is embedded in the MPC formulation. The concept is first outlined in [Bal22] and [BT23], a comparative simulation study using a Gazebo simulation of the LRU1 is shown in [Bal+23]. Finally, a test campaign with the LRU1 is performed at the DLR's Planetary Exploration Lab, comparing the tracking performance of the learning-based MPC with the conventional controllers. The control approach, the test campaign, and a summary of the results is described in the following.⁶

4.2.1 Controller architecture

The structure of the learning-based MPC is sketched in Figure 4.8. The rover hardware together with the low-level steering controller are interfaced by a nonlinear controller (NLC) which stabilizes

⁶The theoretical part of this section is mainly based on the two consecutive publications [LBT24] and [Bal+24].

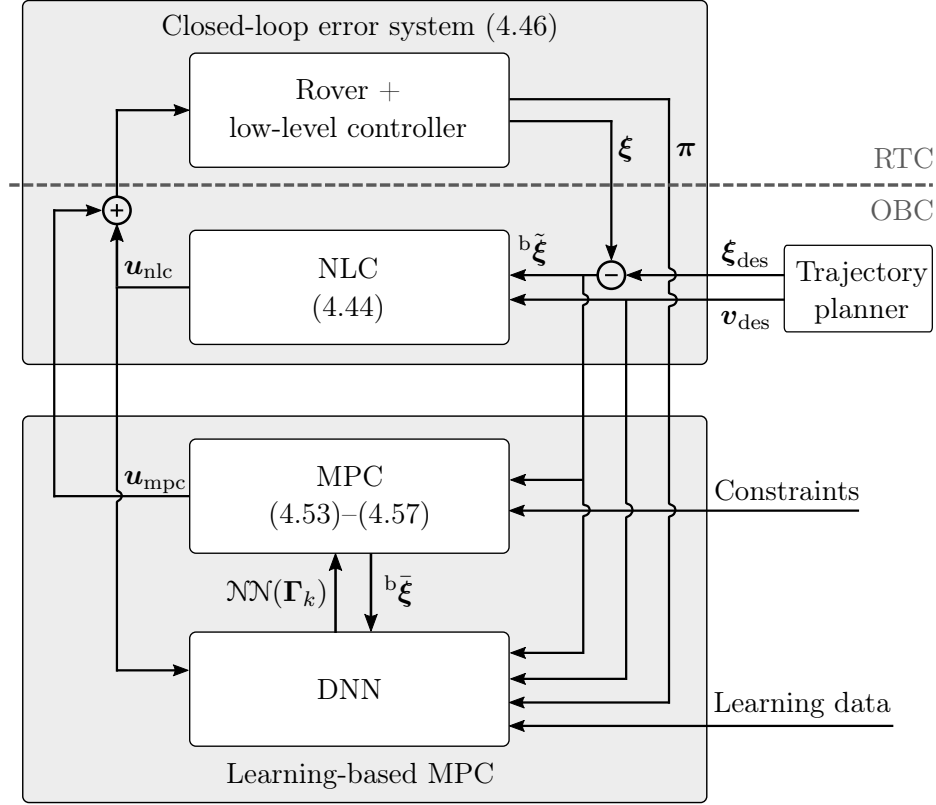


Figure 4.8 Control architecture of the platform with learning-based MPC. The rover system with the low-level steering controller together with the NLC forms the closed-loop error system (4.46). The learning-based MPC can be added in order to improve the tracking performance, using a linearized model of the closed-loop error dynamics (4.54), together with telemetry data from the rover and the output of the NLC. The gray dashed line denotes the physical border between the RTC and the OBC.

the dynamics to a desired trajectory from a trajectory planner component. A learning-based MPC is implemented as outer control loop. The MPC problem is formulated using a linearized error dynamics model, in order to reduce the computational effort. Additionally, a deep neural network (DNN) can be trained, which is then used to estimate the residual error of the MPC error state estimation. The proposed architecture is modular, this means that each component (NLC, MPC, DNN) is an optional add-on to the underlying components. Moreover, the DNN is directly incorporated in the MPC problem and is thus subject to the associated consideration of the constraints.

In order to enable a high level of autonomy, all relevant software components are designed to run on the on-board computers.⁷ In particular, this includes learning-based controller modules. Thereby, it is important to note that no online learning is performed, the neural network is instead trained offline and used as a frozen graph during driving. Still, the neural network can be re-trained and updated after the rover has performed a driving operation, as the training of the neural network is possible in reasonable times on the OBC. An experimental validation of this fact is given at the end of this section, after a brief presentation of the components of the modular learning-based MPC.

⁷The physical implementation of the algorithms on the real-time computer (RTC) and the main on-board computer (OBC) are indicated in Figure 4.8 by the gray dashed line.

Nonlinear tracking control

A nonlinear tracking controller (NLC) is implemented, based on the control law proposed in [Kan+90]. It builds upon the low-level platform velocity controller described in Section 2.1.4, using the platform velocity command $\mathbf{v}_{\text{plf,cmd}}$ as an interface. The low-level platform velocity controller then maps the velocity command to commanded steering angles and wheel velocities, as described in Section 2.2.1.

In order to create a velocity control law which is generic for different kinds of mobile platforms, including but not limited to cars, bicycles, and tricycles, the system input is reduced to the tangential velocity v_x and the yaw rate ω , assuming $v_y \equiv 0$. The reduced mapping of the inputs to the velocity of the rover is

$$\dot{\xi} = \begin{pmatrix} {}^0\dot{x} \\ {}^0\dot{y} \\ \dot{\theta} \end{pmatrix} = \begin{pmatrix} \cos(\theta) & 0 \\ \sin(\theta) & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} v_x \\ \omega \end{pmatrix}. \quad (4.33)$$

Recall that error coordinates are written as $\tilde{(\cdot)} = (\cdot) - (\cdot)_{\text{des}}$. The error coordinates in the local (body) frame can be written as

$${}^b\tilde{\xi} = \begin{pmatrix} {}^b\tilde{x} \\ {}^b\tilde{y} \\ \tilde{\theta} \end{pmatrix} = \begin{pmatrix} \cos(\theta) & \sin(\theta) & 0 \\ -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{pmatrix} (\xi - \xi_{\text{des}}). \quad (4.34)$$

From these relations, the closed-loop error dynamics can be derived following [Kan+90]. The derivative of (4.34) w.r. t. time is

$${}^b\dot{\tilde{x}} = [-s\theta(\dot{x} - \dot{x}_{\text{des}}) + c\theta(\dot{y} - \dot{y}_{\text{des}})]\dot{\theta} + c\theta(\dot{x} - \dot{x}_{\text{des}}) + s\theta(\dot{y} - \dot{y}_{\text{des}}) \quad (4.35)$$

$${}^b\dot{\tilde{y}} = [-c\theta(\dot{x} - \dot{x}_{\text{des}}) - s\theta(\dot{y} - \dot{y}_{\text{des}})]\dot{\theta} - s\theta(\dot{x} - \dot{x}_{\text{des}}) + c\theta(\dot{y} - \dot{y}_{\text{des}}) \quad (4.36)$$

$$\dot{\tilde{\theta}} = \dot{\theta} - \dot{\theta}_{\text{des}}, \quad (4.37)$$

with the abbreviated notation $s(\cdot) = \sin(\cdot)$ and $c(\cdot) = \cos(\cdot)$, and omitting the left superscript 0 for variables expressed in the fixed frame (X, Y) . Using (4.34) and (4.33), the above terms can be simplified to

$${}^b\dot{\tilde{x}} = {}^b\tilde{y}\omega + c\theta\dot{x} + s\theta\dot{y} - c\theta\dot{x}_{\text{des}} - s\theta\dot{y}_{\text{des}} \quad (4.38)$$

$${}^b\dot{\tilde{y}} = -{}^b\tilde{x}\omega - s\theta\dot{x} + c\theta\dot{y} + s\theta\dot{x}_{\text{des}} - c\theta\dot{y}_{\text{des}} \quad (4.39)$$

$$\dot{\tilde{\theta}} = \omega - \omega_{\text{des}}. \quad (4.40)$$

Using $v_x = c\theta\dot{x} + s\theta\dot{y}$ and $v_y = -s\theta\dot{x} + c\theta\dot{y} = 0$ yields

$${}^b\tilde{x} = {}^b\tilde{y}\omega + v_x - c\theta\dot{x}_{\text{des}} - s\theta\dot{y}_{\text{des}} \quad (4.41)$$

$${}^b\tilde{y} = -{}^b\tilde{x}\omega + s\theta\dot{x}_{\text{des}} - c\theta\dot{y}_{\text{des}}. \quad (4.42)$$

Then, $s\theta = s(\tilde{\theta} + \theta_{\text{des}})$ and $c\theta = c(\tilde{\theta} + \theta_{\text{des}})$ are inserted and expanded using trigonometric addition theorems. Using the equality $s\theta\dot{x}_{\text{des}} = c\theta\dot{y}_{\text{des}}$ from (4.33), and $v_{x,\text{des}} = c\theta_{\text{des}}\dot{x}_{\text{des}} + s\theta_{\text{des}}\dot{y}_{\text{des}}$ and $v_{y,\text{des}} = -s\theta_{\text{des}}\dot{x}_{\text{des}} + c\theta_{\text{des}}\dot{y}_{\text{des}} = 0$ finally leads to the error dynamics

$${}^b\dot{\tilde{\xi}} = \begin{pmatrix} {}^b\tilde{y}\omega + v_x - v_{x,\text{des}}\cos(\tilde{\theta}) \\ -{}^b\tilde{x}\omega + v_{x,\text{des}}\sin(\tilde{\theta}) \\ \omega - \omega_{\text{des}} \end{pmatrix}. \quad (4.43)$$

Assuming perfect platform velocity control, i.e., $v_x \equiv v_{x,\text{cmd}}$ and $\omega \equiv \omega_{\text{cmd}}$, a nonlinear control law which stabilizes the platform to a given trajectory $v_{x,\text{des}}(t) > 0, \omega_{\text{des}}(t)$ is

$$\mathbf{u}_{\text{nlc}} = \begin{pmatrix} v_{x,\text{nlc}} \\ \omega_{\text{nlc}} \end{pmatrix} = \begin{pmatrix} v_{x,\text{des}} \cos(\tilde{\theta}) - k_x {}^b\tilde{x} \\ \omega_{\text{des}} - k_y v_{x,\text{des}} {}^b\tilde{y} - k_\theta \sin(\tilde{\theta}) \end{pmatrix}. \quad (4.44)$$

The final control input is

$$\mathbf{u}_{\text{cmd}} = \begin{pmatrix} v_{x,\text{cmd}} \\ \omega_{\text{cmd}} \end{pmatrix} = \begin{pmatrix} v_{x,\text{nlc}} \\ \omega_{\text{nlc}} \end{pmatrix} + \mathbf{u}_{\text{mpc}} \quad (4.45)$$

with $\mathbf{u}_{\text{mpc}} = (v_{x,\text{mpc}} \ \omega_{\text{mpc}})^T$ is the additional control input of the MPC. Inserting (4.45) and (4.44) into (4.43) yields the closed-loop error dynamics

$${}^b\dot{\tilde{\boldsymbol{\xi}}} = \begin{pmatrix} (\omega_{\text{des}} - k_\theta \sin(\tilde{\theta}) - k_y v_{x,\text{des}} {}^b\tilde{y}) {}^b\tilde{y} - k_x {}^b\tilde{x} + {}^b\tilde{y}\omega_{\text{mpc}} + v_{x,\text{mpc}} \\ -(\omega_{\text{des}} - k_\theta \sin(\tilde{\theta}) - k_y v_{x,\text{des}} {}^b\tilde{y}) {}^b\tilde{x} + v_{x,\text{des}} \sin(\tilde{\theta}) - {}^b\tilde{x}\omega_{\text{mpc}} \\ -k_\theta \sin(\tilde{\theta}) - k_y v_{x,\text{des}} {}^b\tilde{y} + \omega_{\text{mpc}} \end{pmatrix}. \quad (4.46)$$

Then, (4.46) is linearized around the equilibrium ${}^b\tilde{\boldsymbol{\xi}}^* = \mathbf{0}$ in order to obtain a linear parameter-varying (LPV) system

$${}^b\dot{\tilde{\boldsymbol{\xi}}} = \mathbf{A} {}^b\tilde{\boldsymbol{\xi}} + \mathbf{B} \mathbf{u}_{\text{mpc}} \quad (4.47)$$

with

$$\mathbf{A}(k_x, k_y, k_\theta, v_{x,\text{des}}, \omega_{\text{des}}) = \begin{pmatrix} -k_x & \omega_{\text{des}} & 0 \\ -\omega_{\text{des}} & 0 & v_{x,\text{des}} \\ 0 & -k_y v_{x,\text{des}} & -k_\theta \end{pmatrix} \quad (4.48)$$

and

$$\mathbf{B} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}. \quad (4.49)$$

The linearized system can be written as discrete-time system as

$${}^b\tilde{\boldsymbol{\xi}}_{k+1} = \mathbf{A}_d {}^b\tilde{\boldsymbol{\xi}}_k + \mathbf{B}_d \mathbf{u}_{\text{mpc},k}. \quad (4.50)$$

Thereby, the input \mathbf{u}_{mpc} is assumed constant over the interval $[t_k, t_{k+1}]$ (zero-order hold). For a constant sampling rate $\Delta t = t_{k+1} - t_k$, the discretized system matrices \mathbf{A}_d and \mathbf{B}_d can be computed as

$$\mathbf{A}_d = \Phi(\Delta t) = \Phi(t_{k+1}, t_k) = e^{\mathbf{A}\Delta t}, \quad (4.51)$$

$$\mathbf{B}_d = \int_{t_k}^{t_{k+1}} \Phi(t_{k+1}, \tau) \mathbf{B} d\tau = -\mathbf{A}^{-1}(\mathbf{I} - \Phi(\Delta t))\mathbf{B}, \quad (4.52)$$

see e.g. [Sch08, p. 38]. The difference equation (4.50) is then used as a basis for the design of the learning-based MPC, which is described in the following.

Learning-based MPC

Based on the linearized system dynamics, the additional control inputs \mathbf{u}_{mpc} can be obtained together with the corresponding error state prediction ${}^b\tilde{\boldsymbol{\xi}}$ by solving the MPC problem as described in [Bal+24]:

$$\{{}^b\tilde{\boldsymbol{\xi}}^*, \mathbf{u}_{\text{mpc}}^*\} = \arg \min_{{}^b\tilde{\boldsymbol{\xi}}, \mathbf{u}_{\text{mpc}}} \left({}^b\tilde{\boldsymbol{\xi}}_N^T \mathbf{P} {}^b\tilde{\boldsymbol{\xi}}_N + \sum_{k=0}^{N-1} \left({}^b\tilde{\boldsymbol{\xi}}_k^T \mathbf{Q} {}^b\tilde{\boldsymbol{\xi}}_k + \mathbf{u}_{\text{mpc},k}^T \mathbf{R} \mathbf{u}_{\text{mpc},k} \right) \right) \quad (4.53)$$

$$\text{s.t. } {}^b\tilde{\boldsymbol{\xi}}_{k+1} = \mathbf{A}_d {}^b\tilde{\boldsymbol{\xi}}_k + \mathbf{B}_d \mathbf{u}_{\text{mpc},k} + \mathcal{NN}(\mathbf{\Gamma}_k)_{k+1} \quad (4.54)$$

$${}^b\tilde{\boldsymbol{\xi}}_0 = {}^b\tilde{\boldsymbol{\xi}}(t_0) \quad (4.55)$$

$${}^b\tilde{\boldsymbol{\xi}}_k \in \mathcal{X}, \mathbf{u}_{\text{mpc},k} \in \mathcal{U}, k = 0, \dots, N-1 \quad (4.56)$$

$${}^b\tilde{\boldsymbol{\xi}}_N \in \mathcal{X}_f \quad (4.57)$$

The quadratic cost function contains the symmetric positive semi-definite weighting matrices $\mathbf{Q} \in \mathbb{R}^{3 \times 3}$, $\mathbf{P} \in \mathbb{R}^{3 \times 3}$, and the s.p.d. matrix $\mathbf{R} \in \mathbb{R}^{2 \times 2}$. N denotes the prediction horizon of the discretized nominal closed-loop LPV system dynamics (4.50). The state and input variables are constrained to their respective constraint sets \mathcal{X} and \mathcal{U} , the set of terminal constraints is \mathcal{X}_f . The term $\mathcal{NN}(\mathbf{\Gamma}_k)$ in (4.54) denotes an additional input term from a deep neural network, which is used to compensate for the remaining prediction error of the linear MPC. If it is set to zero, (4.53)–(4.57) is a “classical” linear MPC. The feature matrix $\mathbf{\Gamma}_k$ is the input of the DNN. It consists of l sequences of the input feature vectors $\boldsymbol{\gamma}_s \in \mathbb{R}^{n_\gamma}$ with $s = k-l, \dots, k$:

$$\mathbf{\Gamma}_k = (\boldsymbol{\gamma}_{k-l} \dots \boldsymbol{\gamma}_k) \in \mathbb{R}^{n_\gamma \times l}. \quad (4.58)$$

Thereby, the feature vector at a single timestamp s is defined as

$$\boldsymbol{\gamma}_s = \left({}^b\tilde{\boldsymbol{\xi}}_s^T \quad {}^b\tilde{\boldsymbol{\xi}}_i^T \quad \mathbf{u}_{\text{nlc},s}^T \quad \mathbf{v}_{\text{des},j}^T \quad \boldsymbol{\pi}_s^T \right)^T. \quad (4.59)$$

It contains the current state ${}^b\tilde{\boldsymbol{\xi}}_s \in \mathbb{R}^3$, the N stacked last predicted states ${}^b\tilde{\boldsymbol{\xi}}_i \in \mathbb{R}^{3N}$ of the time steps $i = s-1-N, \dots, s-1$, the input of the NLC at the current timestamp $\mathbf{u}_{\text{nlc},s} \in \mathbb{R}^2$, the reference velocities $\mathbf{v}_{\text{des},j} \in \mathbb{R}^{2N}$ with $j = s, \dots, s+N$ and the current telemetry data $\boldsymbol{\pi}_s = (\theta \ \phi \ \rho_1 \ \rho_2)^T \in \mathbb{R}^4$. The telemetry data provided by the rover hardware consists of IMU measurements of the body pitch and roll angle (θ and ϕ , respectively), and the angular measurements of the front and rear bogie axles ρ_1 and ρ_2 .

The function $\mathcal{NN}(\mathbf{\Gamma}_k)$ then returns a sequence of predicted errors of the MPC, $\Delta {}^b\hat{x}$, $\Delta {}^b\hat{y}$ and $\Delta \hat{\theta}$, starting from $k+1$ over the prediction horizon. Therefore, the DNN output dimension results in $\mathbb{R}^{3 \times N}$:

$$\mathcal{NN}(\mathbf{\Gamma}_k) = \begin{pmatrix} \Delta {}^b\hat{x}_{k+1} & \dots & \Delta {}^b\hat{x}_{k+N} \\ \Delta {}^b\hat{y}_{k+1} & \dots & \Delta {}^b\hat{y}_{k+N} \\ \Delta \hat{\theta}_{k+1} & \dots & \Delta \hat{\theta}_{k+N} \end{pmatrix} \quad (4.60)$$

In the MPC condition (4.54), the respective $(k+1)$ -th column of $\mathcal{NN}(\mathbf{\Gamma}_k)$ is selected, such that $\mathcal{NN}(\mathbf{\Gamma}_k)_{k+1}$ with

$$\mathcal{NN}(\mathbf{\Gamma}_k)_1 = \begin{pmatrix} \Delta {}^b\hat{x}_{k+1} \\ \Delta {}^b\hat{y}_{k+1} \\ \Delta \hat{\theta}_{k+1} \end{pmatrix} \text{ for } k = 0, \dots, \mathcal{NN}(\mathbf{\Gamma}_k)_N = \begin{pmatrix} \Delta {}^b\hat{x}_{k+N} \\ \Delta {}^b\hat{y}_{k+N} \\ \Delta \hat{\theta}_{k+N} \end{pmatrix} \text{ for } k = N-1 \quad (4.61)$$

is added at the k -th iteration of (4.54), iterating $k = 0, \dots, N-1$ in each optimization time step. Note that the MPC state prediction, including the error prediction by the DNN, is applied and

optimized over the entire prediction horizon. This means that the DNN is directly integrated in the MPC and the associated consideration of the constraints. A detailed description of the used DNN can be found in [Bal+24]. The major components of the network architecture include a *Long Short-Term Memory* layer, two Dense layers with different activation functions and two Dropout layers to avoid overfitting. The average number of neurons per layer is 25. The network was trained with data from the real rover, which was obtained during the test campaign as described in the next section. A data preprocessing pipeline was implemented, using standardization techniques in order to ensure that all features have a similar scale. This step prevents any particular feature from dominating the model training process.

4.2.2 Experiments

The presented approach is tested using the LRU1 at the *Planetary Exploration Lab* (PEL), which is an indoor Moon/Mars-analog testbed for rough-terrain rover systems at DLR. This test site contains a sand box of $10\text{ m} \times 5\text{ m}$, which can be filled with different types of soil. Additionally, the PEL features a mechanism to create variable slopes of 0 to 30° in the rear part of the testbed. The sloped area covers approximately $3\text{ m} \times 5\text{ m}$. The tests are performed on the EAC-1A soil, a common Moon-analog soil which is characterized in [Eng+20]. Small hills and bumps are modeled in the testbed, see Figure 4.10 a).

The idea of the test setup is to compare the performance of different tracking control approaches in a meaningful way. Thus, the position measurements used for the respective feedback control laws are provided as ground-truth measurements from the tracking system. In this way, it is ensured that the tracking performance of the controllers can be evaluated without being biased by errors of on-board pose estimation of the rover. A set of trajectories is provided by a trajectory planner component. During all test drives, the slope of the rear part of the testbed is adjusted to 15° , a minimum of 3 laps are completed with a constant desired velocity of $v_{x,\text{des}} = 0.3\text{ m/s}$. The gains of the NLC are selected as $K_x = 0.045\text{ s}^{-1}$, $K_y = 3.529\text{ m}^{-2}$, and $K_\theta = 1.258\text{ m}^{-1}$. The weights of the MPC are chosen as $\mathbf{Q} = \mathbf{I}$ and $\mathbf{R} = 10\mathbf{I}$. The prediction horizon is set to $N = 10$, this yields a dimension of the feature vectors of $n_\gamma = 59$. The number of feature sequences is chosen as $l = 10$.

Experiment 4.1 – Comparison of NLC and MPC: First, the performance of the MPC without the learning component is compared to the basic NLC. The results of two test trajectories can be inspected in Figure 4.9. The upper plots a) show the projected view on the (x, y) -plane of the testbed with the position data obtained from the tracking system. The position of the rover is plotted in blue (NLC) and green (MPC). It can be verified that both the NLC and the linear MPC are able to track the reference trajectory. However, significant errors occur, which can be explained with the uneven terrain, cf. Figure 4.10 a). In particular, large deviations occur using the NLC only in the sloped area, which can be explained by downhill lateral and side-slip. This again verifies the findings from the evaluation of the parameterized slip model (see Section 4.1.3), where it is shown that slope-dependent slip and side slip are the dominant components of the parameterized slip model. However, it is evident by looking at the projected position data that the MPC is able to significantly reduce the errors due to downhill slip (noticeably without an embedded slip model).

Figure 4.9 b) shows the Euclidean norms of the tracking errors in local x and y -direction. Note that the high initial errors depend on the starting position and time of the rover w. r. t. the reference trajectory, which is not exactly the same for the different drives. Thus, comparing the data is only reasonable after the first approx. 20 s. From this time on, the error norms are in the same order of magnitude for both trajectories. However, the Peanut trajectory shows larger deviations at certain times, e. g. at $t \approx 30\text{ s}$ and $t \approx 100\text{ s}$ for the NLC and at $t \approx 50\text{ s}$

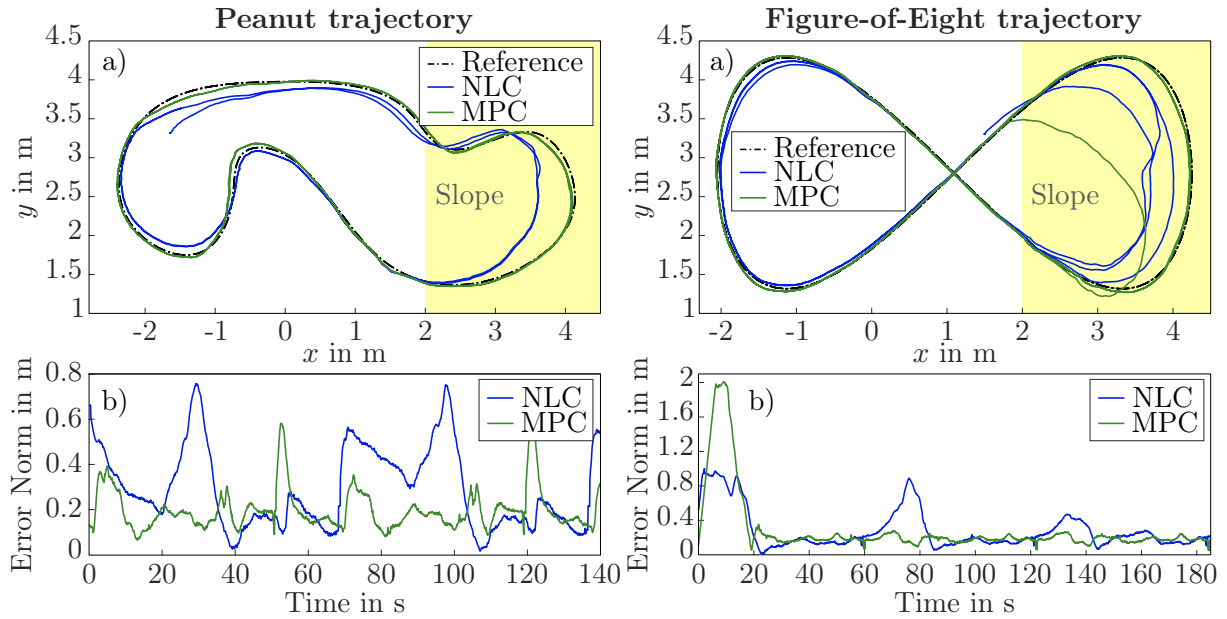


Figure 4.9 Experiment 4.1: Comparison of the tracking performance of the NLC and the MPC with two different example trajectories. a) Reference and measured path for the NLC and the MPC. The yellow area indicates the slope of 15° . b) Comparison of the Euclidean norm of the tracking errors $\sqrt{{}^b\tilde{x}^2 + {}^b\tilde{y}^2}$ for the NLC and the MPC.




Figure 4.10 a) The LRU1 driving on the sloped area of the PEL during the Experiments 4.1 and 4.2. The terrain was modeled with small surface irregularities, disturbing the performance of the tracking controller. b) Experiment 4.3: Comparison of the RMS of the translational tracking errors between the NLC and the MPC driving two example trajectories. Note that the total RMSE is only computed for $t \geq 20$ s, such that the difference in the initial position of the rover does not distort the results.




Figure 4.11 Experiment 4.2, Drive 2: Comparison of the MPC error state with the prediction of the DNN (only the next step in the prediction horizon ($\Delta {}^b\hat{x}_{k+1}, \Delta {}^b\hat{y}_{k+1}$) is plotted as a solid line). The actual MPC error state (${}^b\hat{x}_{k+1}, {}^b\hat{y}_{k+1}$) is plotted as dot-dashed line. a) Comparison in local x -direction. b) Comparison in local y -direction. From [LBT24].

and $t \approx 120$ s for the MPC. This can be explained with the sharp curves that the rover has to pass during the trajectory. At these instants, the wheels have to be reoriented in order to adjust to the new direction. In this case, the low-level platform velocity controller sets the wheel velocity commands to zero until the steering positions correspond to their desired values. This behavior is implemented in order to protect the hardware from undesired skidding and slipping, cf. Section 2.1.4. However, this behavior is not covered by the mathematical system model used within the MPC.

Figure 4.10 b) shows the RMSE of the NLC and the MPC for the respective times $t \geq 20$ s. Again, it can be verified that the Peanut trajectory is more demanding, such that the tracking errors are larger for both controllers. However, the MPC reduces the total translational RMSE by about 40 % compared to the NLC for both trajectories. This verifies the modular structure described in the preceding section and serves as a baseline for the following experiment, which evaluates the learning-based addition to the MPC.

Experiment 4.2 – Evaluation of the learning-based MPC: The training of the neural network was conducted with a training data set from the initial run of a Figure-of-Eight trajectory (Drive 1). Thereby, the training was performed on the OBC and lasted around 3 minutes (the size of the data set was 4 MB). As an initial verification of the effect of the training, Drive 2 was performed on the same trajectory with the basic MPC. During the drive, the DNN was predicting the error states in *open loop*, i. e., the predictions were *not* fed back into the MPC. In Figure 4.11, the predicted error states from the DNN ($\Delta {}^b\hat{x}_{k+1}, \Delta {}^b\hat{y}_{k+1}$) are compared to the internal error states (${}^b\hat{x}_{k+1}, {}^b\hat{y}_{k+1}$) of the MPC. It can be verified that the DNN is able to predict a great amount of the errors, even if the driven trajectory differs from the training data. This confirms the success of the learning-based MPC approach. Consequently, as a next step, the Figure-of-Eight trajectory was driven with the learning-based MPC in *closed loop*, meaning that the error predictions from the DNN were this time included in the MPC (Drive 3). The left two plots in Figure 4.12 show a comparison of the tracking performance between the first three Drives 1–3 (training – open-loop – closed-loop) on the Figure-of-Eight trajectory. The data is only plotted for $t \geq 20$ s, such that the initial convergence to the trajectory from different starting positions is not displayed. In Figure 4.12 b) it can be seen that the Euclidean norm of the translational errors is similar for Drives 1 (basic MPC, blue dotted, generation of training data) and 2 (basic MPC, solid green,




Figure 4.12 Experiment 4.2: Comparison of the tacking performance for two different trajectories. a) Path in the x - y -plane. b) Euclidean error norm $\sqrt{{}^b\tilde{x}^2 + {}^b\tilde{y}^2}$.

For the Figure-of-Eight trajectory, the ground-truth MPC Drive 1 (blue) is compared to the LB-MPC open-loop Drive 2 (green) and the LB-MPC closed-loop Drive 3 (magenta). The right plots show the Strange Eight trajectory, with the MPC ground truth Drive 4 again in blue, the LB-MPC open-loop Drive 5 in magenta and the LB-MPC closed-loop Drive 6 after a re-training of the network in orange. Note that the data is only shown for $t \geq 20$ s, when the effects due to the difference in the initial position have vanished.

open-loop drive), which validates the repeatability of results in the test setup. The solid magenta line shows Drive 3, that is the closed-loop drive with the LB-MPC. It can be verified that the tracking performance is improved and the error norm is reduced during the complete drive.

Afterwards, the behavior of the LB-MPC is tested using a second validation trajectory, the *Strange Eight*, which can be inspected in the right plots of Figure 4.12. First, another ground-truth drive with the basic MPC is performed (Drive 4, plotted as blue dotted line). Figure 4.13 shows a comparison of the commanded and the measured velocities v_x and ω for the two ground-truth drives of the two trajectories (Drive 1 and Drive 4). It is visible that the Strange Eight validation trajectory induces larger variations in both velocity commands. One can even see that the yaw rate is hitting the limit of -1 rad/s at the Strange Eight trajectory, and is thus limited by the low-level controller. Moreover, the low-level controller sets the wheel velocity commands to zero when the steering configuration changes too fast, which causes instantaneous drops to zero of the measured rover velocities. This effect can be observed for both trajectories in Figure 4.13 a), although it happens more often and for longer periods at the Strange Eight trajectory. These two limitations induced by the low-level controller are the main cause of the larger error norm during




Figure 4.13 Experiment 4.2, Drives 1 and 4: Ground-truth drives with the basic MPC. The plots show a comparison of commanded and measured velocities v_x and ω for the two trajectories. The measured values (from the WO measurements) are depicted in cyan, the output of the MPC is plotted in blue, and the processed velocity command of the low-level controller is plotted in magenta dot-dashed.




Figure 4.14 Experiment 4.2, Drives 1–3 (Figure-of-Eight) and Drives 4–6 (Strange Eight): The total translational RMSE is compared for two different trajectories, computed for $t \geq 20$ s in order to neglect initialization effects.

the open-loop drive of the Strange Eight compared to the Figure-of-Eight trajectory, which is visible in Figure 4.12 b).

Next, a closed-loop drive with the LB-MPC is performed for both trajectories. The runs are plotted in magenta in Figure 4.12. It can be seen that the DNN component is able to reduce the Euclidean norm of the translational tracking errors at most times, especially at the sloped parts of the trajectory. However, significant deviations from the trajectory can be observed while traversing the flat area, which indicates a deficient learning success, probably due to the small learning data set. Here, it is again important to emphasize that the training was performed using data only from a single drive of approx. 3 laps, executed on a different trajectory (Drive 1, Figure-of-Eight).

As a last step, the training of the DNN is repeated, adding the data of Drive 4 and Drive 5 to the initial training data from Drive 1. One last drive is performed with the LB-MPC comprising the re-trained network, Drive 6. The results are plotted in orange, and verify the success of the consecutive training. Figure 4.14 sums up the results by showing the total RMSE of the translational tracking for all Drives 1–6. The left three bars compare the performance while driving the Figure-of-Eight trajectory, the right three bars show the drives on the Strange Eight trajectory. The formerly mentioned results can be verified again by this depiction. Moreover, the plot shows that the second training cycle of the DNN (performed while driving the Strange Eight trajectory) reduces the total RMSE even compared to the closed-loop drive on the Figure-of-Eight trajectory. This implies that the adding more learning data indeed improves the result of the learning.

4.2.3 Résumé

Summarized, a modular approach for the platform tracking control of a planetary rover was presented and experimentally validated. The presented algorithms are designed to be suitable in planetary exploration scenarios. Therefore, the components are designed to be computationally lightweight compared to state-of-the-art tracking controllers developed for terrestrial applications (as e.g. autonomous racing). In particular, the linearization of the closed-loop error system allows to employ a linear MPC instead of a nonlinear MPC. Moreover, the high-level controller components do not require to be implemented in real time. The integration of a deep neural network allows to compensate for unmodeled and possibly nonlinear wheel-ground interaction effects. The implementation of the DNNs as frozen graphs thereby further enhances computational efficiency. Indeed, the computation time of the LB-MPC including the optimization problem is measured as 31 ms per cycle in average during the test drives, the DNN inference takes 90 μ s in average on the OBC (specifications of the OBC can be found in Section 2.1.2). Last, the selected network architecture requires only small training data sizes for effective learning, as the experiments show. However, it must be mentioned that no dedicated analysis is performed which validates that the control approach is able to run on today's radiation-hardened space-qualified processors. In fact, the computational resources are one of the greatest challenges in space robotics up to now. However, current developments from recent activities like the Mars 2020 mission give cause to the hope that in near future, more powerful processors will be used in planetary exploration. This will be discussed in detail in the following chapter.

Summarized, the aforementioned points together with the experimental results highlight the great potential of the approach for the application in autonomous space exploration. However, future work is planned in order to further improve the results. First and foremost, the implementation of online learning is an important feature which can enhance the performance in varying terrains. Further, the next big step in the development of a learning-based slip estimation and compensation is the closer conjunction of the perception pipeline to the learning-based controller.

In particular, it is planned to add information from the terrain segmentation to the input features of the DNN, such that different terrain types can be distinguished properly.

CHAPTER 5

Discussion

The preceding chapters propose control-related methods which implement crucial skills for a wheeled robot during an autonomous planetary exploration mission. Therefore, the first part of the discussion addresses an evaluation of the presented approaches in the context of planetary exploration. All presented algorithms were tested at prototypical hardware of two planetary exploration rovers, namely LRU1 and LRU2. Thereby, the theoretical development of the approaches relies on several assumptions, which are listed in Section 2.1.3. The validity and justification of the assumptions w.r.t. the hardware capabilities is discussed in the second part of this chapter. However, one of the most important questions is whether the developed algorithms could be used in a real space mission, that is, could they be implemented on a present or (near-) future space robot? This question is addressed in the third part of the discussion.

5.1 Evaluation of the presented approaches in the context of planetary exploration

In Chapter 3, different whole-body control algorithms were presented. All proposed approaches are based on impedance control and passivity-based control concepts, which were selected due to their beneficial properties like the well-defined contact behavior and the robustness w.r.t. unmodeled dynamics and external disturbances. In Section 3.2, a whole-body Cartesian impedance controller was introduced, which can be used for the regulation of a static pose of the end effector. Thereby, the platform is used as an additional DOF while maintaining a static steering configuration. Indeed, the idea of this approach is to provide a classical Cartesian impedance controller for local manipulation tasks, while the formulation as a quadratic optimization problem allows to distribute the joint torques according to a task-related cost function and considering constraints such as actuator limits. In this way, tasks with intense environment contact forces can be conducted, like some invasive methods for geologic sample acquisition as e.g. shoveling and drilling. The static steering configuration is beneficial in this case, as it prevents the platform from undesired slipping and skidding due to wheel steering on deformable soil.

However, some manipulation tasks can comprise the need of an extended manipulator workspace, for example in the presence of obstacles. In these cases, it can be necessary to steer the wheels in order to change the workspace of the platform. The whole-body tracking control algorithm

introduced in Section 3.3 implements dynamic control of the steering angles in a reactive way, complying with the passivity-based whole-body torque control law. Thereby, using the platform torque interface ensures active compliance of both platform and arm, which predestines the control law for interaction with an unstructured environment. If, however, the manipulation task requires a larger and more precise platform motion, it is possible to use the same control law but switch to the platform velocity interface. In this way, the higher tracking precision of the platform velocity controller on unknown ground is exploited. Finally, the proposed regulation procedure for the end of the trajectory can be utilized for combining locomotion with delicate manipulation tasks, using the same control framework. A potential use case for this kind of flexible whole-body control could be for example the collection of left-behind samples, comparable to the planned Mars Sample Return (MSR) mission.¹

From the aforementioned evaluation it is evident that whole-body compliance (which is achieved by the use of the platform torque controller) is beneficial for sensitive, local manipulation tasks, whereas the platform velocity controller is preferred for traversing longer distances in difficult terrain. For this reason, Chapter 4 addresses the problems of slip estimation and compensation w. r. t. platform velocity control. Thereby, Section 4.1 presents a parameter-based slip model, which can be calibrated without the need of continuous ground-truth position measurements. It is therefore well suited for the slip estimation as part of a slip-aware self-localization algorithm in autonomous planetary exploration. However, due to the structural coupling to the pose estimation of the rover, the slip estimation cannot readily be used for slip compensation within the platform velocity controller. Therefore, a modular platform velocity tracking controller is proposed in Section 4.2, which compensates for tracking errors originating in unmodeled wheel-ground interaction by utilizing a machine learning approach. Summarized, the slip estimation and compensation methods presented in Chapter 4 are valuable tools for the locomotion and self-localization performance of the rover on planetary surfaces. The reasonable next step is combining these approaches with the whole-body tracking controller from Section 3.3 to improve the overall system tracking performance. The result is a whole-body control framework which covers the main scenarios and tasks expected during a planetary exploration mission, and enables a higher level of autonomy by promoting control on task level.

5.2 Justification of the modeling assumptions

Throughout the work, control approaches have been derived based on the mathematical modeling of the rover system. However, as mentioned in Section 2.1.3, the mathematical modeling relies on a number of assumptions. In particular, these assumptions concern the hardware characteristics and the properties of the wheel-ground contact. Although the assumptions do not match the real hardware in each detail (a list of practical limitations can also be found in Section 2.1.3), each of the presented control approaches has been tested on the real rover prototype, oftentimes even in space-analog environments. Thus, it can be argued that all presented methods have proven their applicability by performing the respective experiments.

The first class of assumption concerns the structure of the mathematical system model: *All motors are ideal torque sources. The links and joints of the robot are rigid.* These are common assumptions in the control of robots, as they result in the simple second-order dynamics (2.35), where the actuation torques are directly acting on the respective joint accelerations. This mathematical structure is particularly beneficial for the development of passivity-based control concepts. Apparently, these assumptions do not correspond to the real behavior of the robot in

¹As of today, the status of the MSR mission is uncertain due to funding issues.

any case. The modeling of the motors as ideal torque sources does hold approximately, as long as the commanded torques do not hit the actuator torque limits (cf. Table 2.2). This requirement can be supported by parameterizing both task and controller in a feasible way – however, unforeseen external disturbances can cause higher demanded actuator torques. To overcome this problem, approaches like the optimization-based whole-body impedance controller presented in Section 3.2 can be utilized, which exploit the robot’s redundancy to distribute the actuation torques such that the limits are not exceeded.

The next question is if the rigid model of the robot is justified with regard to the lightweight structure of the mechanical components. In fact, the joint elasticities of the Jaco arm (originating mainly from the Harmonic Drive gears) are not considered in the control algorithm, as motor-side joint position measurements are used in the feedback control loop. Moreover, the links of both LRU and Jaco are elastic due to the lightweight structure of the rover. This elasticity can lead to different kinds of disturbances, from oscillations to errors in the forward kinematics computation. The treatment of these kind of errors is well established in literature (see e.g. the overview papers [BV04; KSY15]), and is thus not within the scope of this work. Concerning the erroneous forward kinematics, visual servoing approaches can be adopted in order to improve the manipulation accuracy [CH08; HS19]. Summarized, the mentioned modeling errors and limitations strongly depend on the available hardware and are no limitations of the proposed algorithms in theory.

The second class of assumptions concerns the behavior of the mobile platform in interaction with the environment: *The dynamics of the position-controlled steering actuators is sufficiently fast, such that the steering dynamics can be neglected. The wheel-ground interaction can be modeled as perfect point contact. Wheel velocities and wheel torques are translated to body velocity and body wrench without losses.*

First, the velocity of the steering actuators is limited to $80^\circ/\text{s}$, as pointed out in Section 2.1.3. Thus, wheel propulsion velocity and torque commands cannot always be applied immediately. In the context of the whole-body tracking controller, this means that the assumption of perfect steering control, i.e., $\phi \equiv \phi_{\text{ref}}$, clearly depends on the choice of the kinematic parameters and the hardware design. It can be supported by a careful selection of the desired trajectories w.r.t. the steering capabilities of the system at hand. However, the kinematic redundancy of a wheeled mobile manipulator can be used to compensate for errors originating from the steering controller on task level.

Second, the assumption of a perfect wheel-ground contact does not fully represent the reality on most relevant terrain types, as shown in Chapter 4. However, the presented methods for the estimation and compensation of unknown wheel-ground interaction effects provide the possibility to compensate for disturbances originating from the wheel-ground interaction. Again, the kinematic redundancy of the complete system can be used to compensate for the remaining errors in task space.

Last, controlling single wheels on torque level is demanding at the hardware at hand, due to missing torque sensors in the wheel propulsion actuators and the inherently unknown wheel-ground interaction forces. In fact, many wheeled platforms are rather controlled on velocity level. A great advantage of the whole-body tracking control algorithm presented in Section 3.3 is that it allows to utilize the platform velocity interface instead of the platform torque interface. A benefit is a potentially better tracking performance of the mobile platform, while a potential drawback is that the active compliance of the platform is lost. Depending on the actual task and scenario, the possibility to switch between the platform control modes can make the best of the compliant whole-body control law, by offering high tracking accuracy (possibly in combination with the slip compensation approaches presented in Chapter 4) for tasks involving larger platform locomotion, while providing whole-body compliance for local manipulation tasks.

Summarized, it has been experimentally verified that the proposed whole-body control approaches are suitable for the regulation and tracking of a pose or trajectory of the end effector, considering the limitations of the prototypical rover hardware.

5.3 Feasibility in planetary exploration scenarios

The algorithms in this thesis were designed in the context of present and future robotic planetary exploration missions. In particular, the considered circumstances are based on the robotic exploration of Earth's moon and Mars. The question that arises naturally concerns the feasibility of the proposed approaches in such a scenario. The main aspect of this question is the following: *Are the algorithms applicable w. r. t. the circumstances, i. e., a) with the available sensor data and b) considering the computational capabilities of a space-qualified robotic vehicle?*

The first part of the question can be answered in a nutshell with *yes*. All algorithms are only relying on on-board sensor measurements, without the need for external information such as ground-truth position data or a priori knowledge about the terrain.

Concerning the second part of the question, it has to be admitted that it is not proven that the proposed algorithms are suited for today's radiation-hardened processors, such as the RAD 750 processor which is used for NASA's Curiosity and Perseverance rovers (featuring a maximum clock cycle time of 200 MHz). The upcoming Mars Sample Return mission can be named as an example: Bowkett et al. identify the online inversion of the Jacobian matrix of a 7-DOF robotic arm on such a processor as one of the main computational challenges concerning the implementation of a closed-loop control law in Cartesian coordinates of the end effector [Bow+24].

In fact, there was (deliberately) no runtime analysis of the proposed algorithms performed on real space-qualified computers, looking out for future developments. It is expected that "Future computational resources such as multicore architectures and FPGAs will offer enhanced complexity navigation algorithms" [Ell16, p. 152]. As of today, FPGAs have been implemented in a limited capacity in several rovers including the MERs. Moreover, already the last robotic Mars exploration mission included the flying system Ingenuity as a technology demonstrator, equipped with a 2.26 GHz Quad-core SnapdragonTM801 processor with 2 GB Random Access Memory (RAM) and 32 GB Flash memory. The processor is among others used to perform the expensive realtime sensor fusion and filter propagation for the use in flight control [Bal+18]. The inner flight-control loops are implemented using a pair of hot-swappable dualredundant automotive-grade processors [BAG21]. Compared to conventional space hardware, these terrestrial processors feature an increased danger of data or instruction corruption caused by cosmic radiation. The solution applied to Ingenuity was as simple as providing redundancy in all relevant components, i. e., holding two copies of memory and double check operations as much as possible. For example, if a fault was detected in software or memory of one of the flight computer units, that would trigger the swap of the active microcontroller [Gil24]. In fact, the full mission success criteria for Ingenuity were not only accomplished but by far exceeded. Therefore, two identical Ingenuity-sized helicopters are developed as part of the MSR mission. This second generation of Mars rotorcrafts is heavily based on the heritage designs of Ingenuity, but will use an even newer model of the Snapdragon as primary flight computer [Wit+23]. Thus, it can be anticipated that the computation capabilities of wheeled exploration robots will exceed the present ones by far in the near future.

But even with the computational limits of the present space-qualified processors, advanced algorithms can be implemented by an intelligent use of resources combined with robust algorithms. For example, the offline pre-computation of the inverse Jacobian matrix around a baseline trajectory as a solution for the computational bottlenecks is proposed in [Bow+24]. These kind of solutions have been explored since the early days of model-based control, when computational

resources were comparably limited to today's space-qualified processors. Of course, these techniques are also conceivable for the algorithms proposed in this work. Another approach is to consider a reduction of the computational cycle times. However, it has to be noted that reduced cycle times can limit the bandwidth and thus the achievable stiffness values of the passivity-based control concepts from Chapter 3. Finally, building modular control architectures can reduce the computational requirements, as the example of the learning-based MPC from Section 4.2 shows. Thereby, the computationally expensive high-level control components are not required to run in realtime, as they build upon a simple low-level controller which already provides satisfactory tracking performance.

Summarized, what qualifies the proposed compliant whole-body control approaches for a space exploration scenario is the well-defined contact behavior, which is beneficial in unknown and unstructured environments. Moreover, the inherent robustness of passivity-based control w.r.t. disturbances and modeling errors is an undeniable advantage, thinking for example of the possibly higher signal-to-noise ratio of certain sensors in space applications (for example, this concerns force and torque measurements, as some of the strain gauge materials used in terrestrial applications cannot be used on Mars [Bow+24]). In fact, compliant closed-loop control approaches in task coordinates will make their way to space in the near future, among others due to the advanced manipulation requirements of the MSR mission [Bow+24]. Thus, compliant whole-body control can be imagined to be the next step in future planetary exploration.

CHAPTER 6

Conclusion

In the current work, a comprehensive examination of robotic planetary exploration is conducted from the viewpoint of control theory and practice. The most important aspects for future ground-based exploration scenarios have been identified in the field of locomotion and manipulation. Therefore, the research objective is to provide algorithms which enable a higher level of autonomy in task execution. Looking on the existing unmanned missions (particularly the current and near-future NASA Mars missions), a considerable potential for improvement is foreseen by compliant manipulation approaches, including particularly whole-body control strategies. Thereby, unlike in terrestrial applications like industrial and domestic robotics, the wheel-ground interaction is a substantial factor for the successful task execution.


The first part of the thesis deals with the coordinated whole-body control of a rover prototype equipped with a robotic arm. A platform torque interface is created, proposing an optimization-based solution for the control allocation of the wheel torques. Based on this prerequisite, a whole-body Cartesian impedance controller is presented, which implements a desired spring-damper behavior at the end effector while resolving the redundancy via quadratic optimization. The mentioned control approach is well-suited for local manipulation tasks, assuming a predefined steering configuration. In order to provide a more general control concept for tasks involving significant platform locomotion, a tracking controller which provides an implicit coordination scheme (that is, the steering configuration is computed together with the torque control law without explicitly specifying a platform trajectory) is proposed based on the idea of compliant whole-body behavior. The tracking control law allows both the use of the platform torque interface and the use of the platform velocity controller. Moreover, a concept how to stabilize a static equilibrium at the end of the trajectory is provided. The underlying passivity-based torque control law is eventually analyzed in detail at the example of a holonomic robot arm, with special attention to interaction and convergence properties.

The second part of this work is concerned with the interaction between wheels and underlying terrain. Thereby, the focus is laid on approaches which do neither require ground-truth position measurements nor existing knowledge of the terrain properties. First, a parameter-based slip model is adopted together with a suitable calibration procedure. The method is validated extensively with experiments on Moon-analog terrain. It is shown that the slip estimation improves the wheel odometry measurements significantly. However, it can be observed that the model-based slip estimation is not able to capture all relevant (possibly non-linear) effects. Proceeding from these

insights, a modular platform tracking controller which combines MPC with a machine learning component is proposed and evaluated by a series of test drives on Moon-analog soil. Thereby, the control architecture is designed such that the overall computational load is significantly reduced w. r. t. similar existing methods.

In summary, the control methods which were proposed and evaluated in this work enable future planetary rovers to use the combined degrees of freedom of platform and arm in autonomous locomotion and manipulation scenarios. Thereby, both established and new methods of compliant whole-body control are introduced to the field of planetary exploration robotics. As a next step, the whole-body control approaches can be combined with the proposed approaches for slip estimation and compensation to further improve the resulting regulation and tracking performance.

Looking out for future development in the area of space-qualified processors, the algorithms proposed in this work are expected to be applicable on a (near-)future generation of space robots. Noticeably, the insights from the parameterized slip estimation from [Bus+18] influence an upcoming space mission even today. The Martian Moons eXploration (MMX) mission by the Japan Aerospace Exploration Agency (JAXA) features the major scientific goal of clarifying the origin of the two Martian moons Phobos and Deimos, the currently expected launch date is 2026. As part of the mission, the rover *Idefix* (a joint contribution of the DLR and the Centre National d'Etudes Spatiales CNES) will land on Phobos and act as a scientific and technology demonstration platform. In the context of the Autonomous Navigation Experiment performed by DLR, a velocity correction module is implemented based on the insights from the slip estimation on Mt. Etna [Vay+25]. This example shows that research performed on terrestrial prototypes still has formative influence on actual space exploration, despite the current differences in computational capacities. However, it is expected that more powerful processors will gain increasing importance in near-future space exploration, enabling the use of complex control algorithms as the ones proposed in this work. In the end, whole-body control is a way to increase the autonomy of planetary exploration rovers, and thus eventually saves costs and increases the scientific outcome of robotic space exploration.



©Martin Görner

A.1 Theorems and Fundamentals

The algorithms presented in Chapter 3 and Chapter 4 rely on some basic theorems, which will be introduced in the following chapter.

A.1.1 Nonautonomous system

Definition 1. Let a nonautonomous dynamic system be described by

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t) \quad (\text{A.1})$$

where $\mathbf{f} : \mathcal{D} \times \mathbb{R}^+ \rightarrow \mathbb{R}^n$ is piecewise continuous in t and locally Lipschitz in \mathbf{x} on $\mathcal{D} \times \mathbb{R}^+$, and $\mathcal{D} \subset \mathbb{R}^n$ is a domain that contains the origin $\mathbf{x} = \mathbf{0}$. The origin is an equilibrium point for (A.1) at $t = 0$ if

$$\mathbf{f}(\mathbf{0}, t) = \mathbf{0}, \quad \forall t \geq 0. \quad (\text{A.2})$$

An equilibrium point at the origin could be a translation of a nonzero equilibrium point or, more generally, a translation of a nonzero solution of the system.

Note that the uppermentioned definition is adopted from [Kha02, p. 147]. For further common definitions and theorems used in Lyapunov stability theory, the reader is referred to [Kha02, Chapter 4] or similar works.

A.1.2 Exponential stability

The exponential stability theorem is recalled from [Kha02, p. 154, Theorem 4.10], where also the proof can be found.

Theorem 1. Let $\mathbf{x} = \mathbf{0}$ be an equilibrium point for (A.1) and $\mathcal{D} \subset \mathbb{R}^n$ be a domain containing $\mathbf{x} = \mathbf{0}$. Let $V : \mathcal{D} \times \mathbb{R}^+ \rightarrow \mathbb{R}$ be a continuously differentiable function such that

$$k_1 \|\mathbf{x}\|^a \leq V(\mathbf{x}, t) \leq k_2 \|\mathbf{x}\|^a \quad (\text{A.3})$$

$$\frac{\partial V}{\partial t} + \frac{\partial V}{\partial \mathbf{x}} \mathbf{f}(\mathbf{x}, t) \leq -k_3 \|\mathbf{x}\|^a, \quad (\text{A.4})$$

$\forall t \geq 0$ and $\forall \mathbf{x} \in \mathcal{D}$, where k_1, k_2, k_3 and a are positive constants. Then, $\mathbf{x} = \mathbf{0}$ is exponentially stable. If the assumptions hold globally, then $\mathbf{x} = \mathbf{0}$ is globally exponentially stable.

A.1.3 Brockett's theorem

The following theorem is also known as Brockett's theorem [Bro+83].

Theorem 2. *Let $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$ be given with $\mathbf{f}(\mathbf{x}_0, \mathbf{0}) = \mathbf{0}$ and $\mathbf{f}(\cdot, \cdot)$ continuously differentiable in a neighborhood of $(\mathbf{x}_0, \mathbf{0})$. A necessary condition for the existence of a continuously differentiable control law which makes $(\mathbf{x}_0, \mathbf{0})$ asymptotically stable is*

1. *The linearized system has no uncontrollable modes associated with eigenvalues with positive real part.*
2. *There exists a neighborhood N of $(\mathbf{x}_0, 0)$ such that for each $\boldsymbol{\xi} \in N$ there exists a control $\mathbf{u}_\xi(t)$ defined for all $t > 0$ that drives the solution of $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}_\xi)$ from the point $\mathbf{x} = \boldsymbol{\xi}$ at $t = 0$ to $\mathbf{x} = \mathbf{x}_0$ at $t = \infty$.*
3. *The mapping $\gamma : N \times \mathbb{R}^m \rightarrow \mathbb{R}^n$, N a neighborhood of the origin, defined by $\gamma : (\mathbf{x}, \mathbf{u}) \rightarrow \mathbf{f}(\mathbf{x}, \mathbf{u})$ should be onto an open set of the origin.*

A.1.4 Min-max theorem

The min-max-theorem is a direct result of the from the Courant-Fischer theorem [Cou20; Fis05] and states that

$$\min_i \lambda_i(\mathbf{B}) \|\mathbf{x}\|^2 \leq \mathbf{x}^T \mathbf{B} \mathbf{x} \leq \max_i \lambda_i(\mathbf{B}) \|\mathbf{x}\|^2 \quad \forall \mathbf{x} \quad (\text{A.5})$$

for a s.p.d. matrix \mathbf{B} .

A.1.5 Null space and range space

Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ be an arbitrary matrix. The following definition of the range space, the null space, and the rank \mathbf{A} are recalled from [GL96, p. 49].

Definition 2. The *range space* of \mathbf{A} is defined by

$$\mathcal{R}(\mathbf{A}) = \{\mathbf{y} \in \mathbb{R}^m : \mathbf{y} = \mathbf{A}\mathbf{x} \text{ for some } \mathbf{x} \in \mathbb{R}^n\}. \quad (\text{A.6})$$

The range space is spanned by the columns of $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_n]$, i.e.,

$$\mathcal{R}(\mathbf{A}) = \text{span}\{\mathbf{a}_1, \dots, \mathbf{a}_n\}. \quad (\text{A.7})$$

The null space of \mathbf{A} is defined as

$$\mathcal{N}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{A}\mathbf{x} = \mathbf{0}\}. \quad (\text{A.8})$$

The *rank* of a matrix \mathbf{A} is defined by

$$\text{rank}(\mathbf{A}) = \dim(\mathcal{R}(\mathbf{A})). \quad (\text{A.9})$$

It can be shown that $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}^T)$. The matrix \mathbf{A} is *rank deficient* if

$$\text{rank}(\mathbf{A}) < \min\{m, n\}. \quad (\text{A.10})$$

For $\mathbf{A} \in \mathbb{R}^{m \times n}$,

$$\dim(\mathcal{N}(\mathbf{A})) + \text{rank}(\mathbf{A}) = n. \quad (\text{A.11})$$

A.1.6 Singular value decomposition

Theorem 3. [GL96, p. 70] For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, there exist orthogonal matrices

$$\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_m] \in \mathbb{R}^{m \times m} \text{ and } \mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_n] \in \mathbb{R}^{n \times n} \quad (\text{A.12})$$

such that

$$\mathbf{U}^T \mathbf{A} \mathbf{V} = \mathbf{\Sigma}, \quad (\text{A.13})$$

with

$$\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_p) \in \mathbb{R}^{m \times n}, \quad p = \min\{m, n\}, \quad (\text{A.14})$$

and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$.

The proof can be found in [GL96, p. 70]. The values σ_i are the *singular values* of \mathbf{A} and the vectors \mathbf{u}_i and \mathbf{v}_i are the *ith left singular vector* and the *ith right singular vector*, respectively, satisfying

$$\mathbf{A} \mathbf{v}_i = \sigma_i \mathbf{u}_i \quad (\text{A.15})$$

$$\mathbf{A}^T \mathbf{u}_i = \sigma_i \mathbf{v}_i \quad (\text{A.16})$$

for $i = 1, \dots, \min\{m, n\}$. If the rank of \mathbf{A} is given as $\text{rank}(\mathbf{A}) = r$, the singular values of \mathbf{A} are

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > \sigma_{r+1} = \dots = 0, \quad (\text{A.17})$$

and

$$\mathcal{N}(\mathbf{A}) = \text{span}\{\mathbf{v}_{r+1}, \dots, \mathbf{v}_n\} \quad (\text{A.18})$$

$$\mathcal{R}(\mathbf{A}) = \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_r\}. \quad (\text{A.19})$$

Splitting the matrices \mathbf{U} and \mathbf{V} accordingly, i. e.,

$$\mathbf{U} = [\mathbf{U}_1 \quad \mathbf{U}_2], \quad \mathbf{V} = [\mathbf{V}_1 \quad \mathbf{V}_2] \quad (\text{A.20})$$

with $\mathbf{U}_1 \in \mathbb{R}^{m \times r}$, $\mathbf{U}_2 \in \mathbb{R}^{m \times (m-r)}$, $\mathbf{V}_1 \in \mathbb{R}^{n \times r}$, $\mathbf{V}_2 \in \mathbb{R}^{n \times (n-r)}$, one can obtain orthogonal projections onto the subspaces of \mathbf{A} as follows [GL96, p. 75]:

$$\mathbf{V}_1 \mathbf{V}_1^T = \text{projection on to } \mathcal{R}(\mathbf{A}^T) \quad (\text{A.21})$$

$$\mathbf{V}_2 \mathbf{V}_2^T = \text{projection on to } \mathcal{N}(\mathbf{A}) \quad (\text{A.22})$$

$$\mathbf{U}_1 \mathbf{U}_1^T = \text{projection on to } \mathcal{R}(\mathbf{A}) \quad (\text{A.23})$$

$$\mathbf{U}_2 \mathbf{U}_2^T = \text{projection on to } \mathcal{N}(\mathbf{A}^T). \quad (\text{A.24})$$

A.1.7 Computation of the Moore-Penrose pseudoinverse

The following statements are adapted from [GL96, pp. 257–258]. The *Moore-Penrose pseudoinverse* $\mathbf{A}^+ \in \mathbb{R}^{n \times m}$ of a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is defined as the unique matrix satisfying the *Moore-Penrose conditions*

$$1. \quad \mathbf{A} \mathbf{A}^+ \mathbf{A} = \mathbf{A}, \quad (\text{A.25})$$

$$2. \quad \mathbf{A}^+ \mathbf{A} \mathbf{A}^+ = \mathbf{A}^+, \quad (\text{A.26})$$

$$3. \quad (\mathbf{A} \mathbf{A}^+)^T = \mathbf{A} \mathbf{A}^+, \quad (\text{A.27})$$

$$4. \quad (\mathbf{A}^+ \mathbf{A})^T = \mathbf{A}^+ \mathbf{A}. \quad (\text{A.28})$$

The matrix \mathbf{A}^+ exists for every matrix \mathbf{A} . If \mathbf{A} has full rank, \mathbf{A}^+ can be computed as

1. If \mathbf{A} has linearly independent columns,

$$\mathbf{A}^+ = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T. \quad (\text{A.29})$$

In this case, \mathbf{A}^+ is a left-pseudoinverse of \mathbf{A} , i. e., $\mathbf{A}^+ \mathbf{A} = \mathbf{I}$.

2. If \mathbf{A} has linearly independent rows,

$$\mathbf{A}^+ = \mathbf{A}^T (\mathbf{A} \mathbf{A}^T)^{-1}. \quad (\text{A.30})$$

In this case, \mathbf{A}^+ is a right-pseudoinverse of \mathbf{A} , i. e., $\mathbf{A} \mathbf{A}^+ = \mathbf{I}$.

If $\text{rank}(\mathbf{A}) = n = m$, then $\mathbf{A}^+ = \mathbf{A}^{-1}$. For rank-deficient matrices \mathbf{A} with $\text{rank}(\mathbf{A}) = r$, \mathbf{A}^+ can be computed in terms of the singular value decomposition as

$$\mathbf{A}^+ = \mathbf{V} \mathbf{\Sigma}^+ \mathbf{U}^T \quad (\text{A.31})$$

with

$$\mathbf{\Sigma}^+ = \text{diag} \left(\frac{1}{\sigma_1}, \dots, \frac{1}{\sigma_r}, \dots, 0 \right) \in \mathbb{R}^{n \times m}. \quad (\text{A.32})$$

A.2 The ROBEX and ARCHES space-analog missions

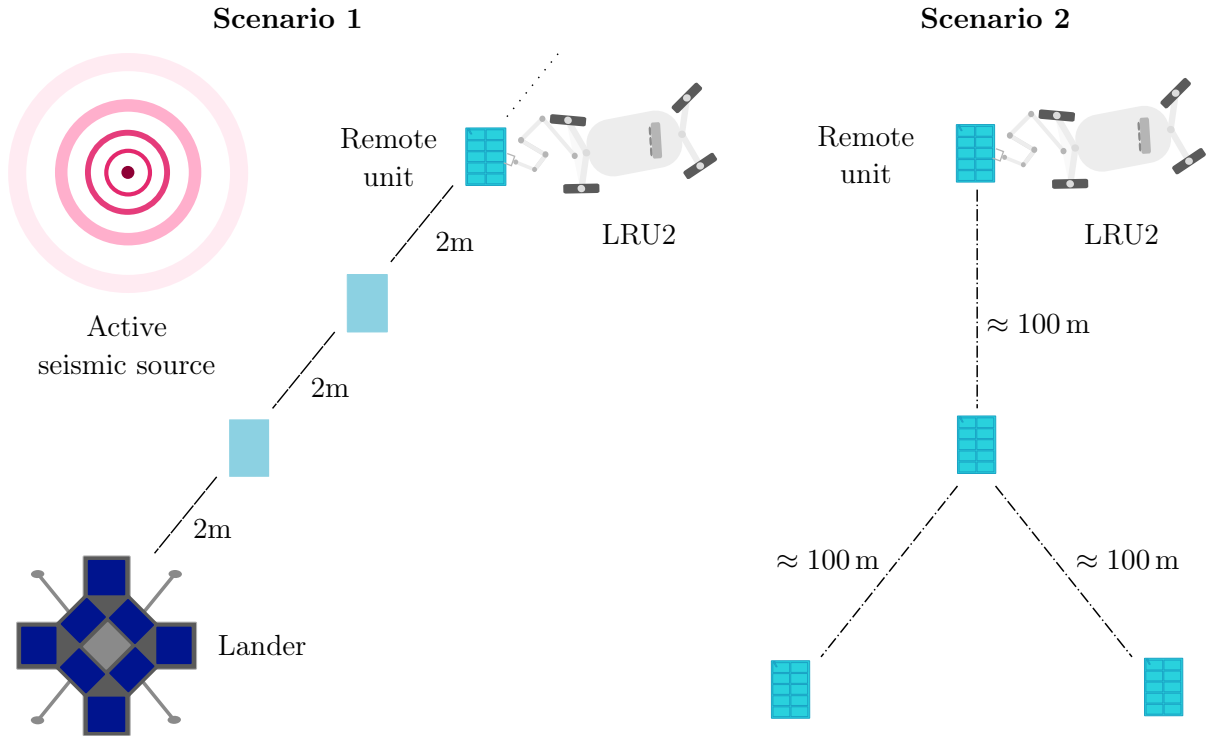


Figure A.1 The two scientific scenarios which were performed autonomously by the LRU2 during the ROBEX space analog mission. Scenario 1: Active seismic measurements at equidistant measure points. Scenario 2: Deployment of a seismic network for long-term geological measurements. Adapted from [Sch+19].

Apart from the rather theoretical considerations derived from past and future space missions, extensive tests with robotic prototypes in a space-analog scenario and environment form the spine of the development of algorithms throughout this work. Thereby, the wheeled terrestrial rover prototypes LRU1 and LRU2 are used to perform several locomotion and manipulation tasks in a harsh and unknown outdoor environment, which is an invaluable source of experience and gives insights into the challenges of planetary exploration.

Two of the most extensive space-analog missions involving the LRUs were located at Mt. Etna, Sicily, a certified Moon-analog test site [PGB12]. The robotic demonstrations were performed at the volcano’s flank at an altitude of approximately 2600 m with the two prototypical planetary rovers LRU1 and LRU2.

The first of the two demonstration missions took place in June/July 2017 in the context of a project named *Robotic Exploration of Extreme Environments* (ROBEX). Details about the mission can be found e. g. in [Wed+15a; Wed+17], and [Wed+18]. The scenario comprised active and passive seismic measurements as shown in Figure A.1, which were performed autonomously by LRU2. This rover prototype is equipped with a robotic arm, which enables it to perform both locomotion and manipulation tasks. The manipulation of the payload boxes for the seismic measurements and a sample-return experiment is described in [Leh+18], [Bru+18] performs a post-mortem analysis of the seismometer deployment using DLR’s task programming framework RMC advanced flow control (RAFCON). Beyond the main scenarios, additional scientific questions were addressed in the course of the demo mission. Among others, a slip-aware wheel odometry

was developed and tested, which is detailed in Section 4.1. The performance of the locomotion capabilities and the pose estimation of the LRU1 was validated through a number of long-range drives. Two of them are available with the data set publication [Vay+18], together with an analysis of the on-board self localization of the rover.



Figure A.2 Cooperative exploration of the Cisternazza crater in the context of the ARCHES space analog mission in 2022 on Mt. Etna. LRU1 has the winch mechanism mounted on its body. The end of the tether is connected to a hook, which is clutched to the robot arm on LRU2. From [Bur+24].

The second space-analog mission took place in June/July 2022 at the same test site on Mt. Etna, as final technology demonstration of the ARCHES project (*Autonomous Robotic Networks to Help Modern Societies*). The scenario involved several scientific experiments, exploration, and infrastructure tasks in the Moon analog environment, among others the setup of a Low-Frequency Radio Array (LOFAR) array. In contrast to the ROBEX project, the challenges of the scenario were tackled by multi-robot cooperation in a heterogeneous team (i.e., LRU1 and LRU2 and DLR's hexacopter ARDEA). A detailed description of the scenario and the results can e.g. be found in [Sch+20; Wed+22], while the mission control software which is essential for the complex interaction of the robotic team is detailed in [Sak+23]. A robotic experiment involving measurements with a Laser-induced Breakdown Spectrometer is described in [Leh+23].

As an additional experiment, the Cisternazza crater was explored by the two rover prototypes LRU1 and LRU2, performing a panoramic scan of the interior of the Crater. Figure A.2 shows a snapshot from the experiment, highlighting the importance of the collaboration: LRU1 (equipped with a scientific camera head) is abseiling into the crater in order to take a panoramic scan, while LRU2 is securing LRU1 from the crater rim. Thereby, LRU1 overcomes slopes with a maximum inclination of 28° thanks to the tether-winch support, which exceeds the rover's stand-alone capability by far (the rover is usually able to traverse slopes of $10\text{--}18^\circ$, depending on the soil). A detailed description of the collaborative crater exploration experiment is given in [Bur+24].

List of Figures

1.1	Overview of the thesis, including the relation to the publications.	11
2.1	The rover prototypes during the ARCHES analog mission on Mt. Etna in summer 2022, performing their typical tasks during a planetary exploration scenario. . . .	14
2.2	Top and side view on the LRU platform.	15
2.3	Overview of possible spoke configurations for the LRUs.	16
2.4	Overview of the computer and software structure of the LRU rovers.	17
2.5	Simplified top-down view on the controller structure at the LRU.	19
2.6	Mounting of the tracking markers and the DGPS antenna for ground-truth position measurements at the LRU1.	20
2.7	Platform velocity and torque mappings.	23
2.8	Possible steering parameterizations.	27
3.1	Experiment 3.1: Comparison of different choices of the weighting matrix \mathbf{Q}_f	34
3.2	Experiment 3.2: Comparison of different values of the weighting factor Q_e	35
3.3	Experiment 3.3: Regulation of the x -coordinate of the end effector.	40
3.4	Experiment 3.4: Regulation at the end effector during physical interaction with the platform.	41
3.5	Experiment 3.5: Physical interaction at the end effector in the presence of joint torque limits.	43
3.6	Whole-body control concept	45
3.7	Concept of the two-phase proof of stability.	53
3.8	Sketch of trajectory tracking in the task space.	54
3.9	LRU2 driving on volcanic soil during Experiment 3.6.	55
3.10	Experiment 3.6: Tracking a sinusoidal trajectory with the platform on volcanic soil.	56
3.11	Experiment 3.6: Trajectory and tracking errors.	56
3.12	Experiment 3.7: Commanded task-space trajectory for the tracking experiment. .	58
3.13	Experiment 3.7: 3D plot of the platform and the position and orientation of the end effector.	59
3.14	Experiment 3.7: Comparison of the tracking errors for platform torque and velocity control.	60
3.15	Experiment 3.7: Comparison of the RMSE of the coordinates of the end effector for platform velocity control and platform torque control.	60

3.16	Experiment 3.7: Comparison of the performance of platform torque and velocity controller.	61
3.17	Experiment 3.7: Comparison of commanded torques and measured angles for the robot arm.	61
3.18	A series of snapshots from Experiment 3.8.	62
3.19	Experiment 3.8: Measured position of the platform and the end effector.	62
3.20	Experiment 3.8: Euclidean norm of translational and rotational errors at the end effector.	63
3.21	Experiment 3.8: Kinematic platform parameters during the contact experiment. .	64
3.22	Experiment 3.7 (platform torque control): Euclidean norm of the translational and rotational errors at the end effector.	65
3.23	Experiment 3.7 (platform torque control): Energy function at the end of the trajectory.	65
3.24	a) The KUKA LWR IV+. b) Joint position and velocity errors from Experiment 3.9	72
3.25	Experiment 3.9: Undisturbed trajectory tracking.	73
3.26	Experiment 3.10: Trajectory with virtual external torques.	74
3.27	Experiment 3.11: Physical interaction with a human operator, $\mathbf{K}_m = \mathbf{0}$ (SLC). . .	75
3.28	Experiment 3.11: Physical interaction with a human operator. $\mathbf{K}_m = 100\mathbf{H}_m$ and $\mathbf{K}_m = -\mathbf{\Omega}_m\mathbf{H}_m$	76
4.1	Structure of the Extended Kalman Filter (EKF) for pose estimation.	78
4.2	Overview of calibration setup and the calibration trajectories for the parameterized slip model.	85
4.3	The LRU1 at the calibration site on Mt. Etna.	86
4.4	Overview of the test site with DGPS data of the long-range run.	87
4.5	Results of the long-range drive. a) Wheel odometry position measurements. b) Saturated linear spring characteristics, derived from the measurements of the steering angle offset.	87
4.6	Data from the long-range run on Mt. Etna. a) Distance error. b) Value of the slip parameters, obtained by a simulated online calibration.	88
4.7	Concept of a slip-aware controller module within the LRU control and self-localization framework.	91
4.8	Control architecture of the platform with learning-based MPC.	93
4.9	Experiment 4.1: Comparison of the tracking performance of the NLC and the MPC with two different example trajectories.	98
4.10	a) LRU1 driving on the sloped area of the PEL. b) Experiment 4.3: Comparison of the RMS of the translational tracking errors between the NLC and the MPC driving two example trajectories.	98
4.11	Experiment 4.2, Drive 2: Comparison of the MPC error state with the prediction of the DNN.	99
4.12	Experiment 4.2: Comparison of the tracking performance for two different trajectories.	100
4.13	Experiment 4.2, Drives 1 and 4: Ground-truth drives with the basic MPC.	101
4.14	Experiment 4.2, Drives 1–3 (Figure-of-Eight) and Drives 4–6 (Strange Eight). Comparison of the total RMSE for two trajectories.	101
A.1	The two scientific scenarios of the ROBEX space analog mission.	119
A.2	Cooperative exploration of the Cisternazza crater in the context of the ARCHES space analog mission in 2022 on Mt. Etna.	120

List of Tables

2.1	Distinctive features of the two rover prototypes.	15
2.2	LRU1 and LRU2 actuator limits	18
3.1	Gains used for the experiments 3.7–3.8	57
3.2	Comparison of the effective control gains of PD+, SLC and GTC.	69
3.3	Gains used for the experiments 3.9–3.11.	71
4.1	Calibrated slip parameters.	86

Bibliography

Note that my last name changed from Bussmann to Lakatos in July 2022.

- [Agh05] F. Aghili. “A unified approach for inverse and direct dynamics of constrained multibody systems based on linear projection operator: applications to control and simulation”. en. In: *IEEE Transactions on Robotics* 21.5 (Oct. 2005), pp. 834–849 (cited on page 26).
- [AH23] S. Aguilera and S. Hutchinson. “Control of Cart-Like Nonholonomic Systems Using a Mobile Manipulator”. In: *Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Oct. 2023, pp. 6801–6808 (cited on page 6).
- [Alb+03] A. Albu-Schäffer, C. Ott, U. Frese, and G. Hirzinger. “Cartesian impedance control of redundant robots: recent results with the DLR-light-weight-arms”. In: *Proceedings of the 2003 IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 3. Taipei, Taiwan: IEEE, 2003, pp. 3704–3709 (cited on pages 4, 37, 57).
- [AOH07] A. Albu-Schäffer, C. Ott, and G. Hirzinger. “A Unified Passivity-based Control Framework for Position, Torque and Impedance Control of Flexible Joint Robots”. In: *International Journal of Robotics Research* 27.1 (Jan. 2007), pp. 23–39 (cited on pages 4, 36, 71).
- [BAG21] J. Balaram, M. Aung, and M. P. Golombek. “The Ingenuity Helicopter on the Perseverance Rover”. In: *Space Science Reviews* 217.4 (May 2021), p. 56 (cited on page 108).
- [Bal+18] B. Balaram, T. Canham, C. Duncan, H. F. Grip, W. Johnson, J. Maki, A. Quon, R. Stern, and D. Zhu. “Mars Helicopter Technology Demonstrator”. In: *Proceedings of the 2018 AIAA Atmospheric Flight Mechanics Conference*. Kissimmee, Florida: American Institute of Aeronautics and Astronautics, Jan. 2018 (cited on page 108).
- [Bal+23] N. Baldauf, A. Turnwald, T. Lubiniecki, K. Lakatos, and N. Panagiotopoulos. “Learning-Based Motion Control of a Rover on Unknown Ground”. In: *Papers of ESA GNC-ICATT 2023*. ESA, July 2023 (cited on page 92).

- [Bal+24] N. Baldauf, K. Lakatos, A. Meinert, and A. Turnwald. “Towards Learning-Based Trajectory Tracking Control for a Planetary Exploration Rover: Adaptive Model Predictive Control”. In: *Proceedings of the 28th International Conference on System Theory, Control and Computing (ICSTCC)*. Sinaia, Romania, Oct. 2024, pp. 1–6 (cited on pages 11, 12, 92, 96, 97).
- [Bal22] N. Baldauf. “Untersuchung lernbasierter Methoden zur Regelung eines mobilen Roboters mit Modellunsicherheiten”. MA thesis. Technische Hochschule Ingolstadt, 2022 (cited on page 92).
- [Bay+02] B. Bayle, J.-Y. Fourquet, F. Lamiraux, and M. Renaud. “Kinematic Control of Wheeled Mobile Manipulators”. In: *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Vol. 2. Sept. 2002, pp. 1572–1577 (cited on page 5).
- [BDO18] K. Bussmann, A. Dietrich, and C. Ott. “Whole-Body Impedance Control for a Planetary Rover with Robotic Arm: Theory, Control Design, and Experimental Validation”. In: *Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA)*. Brisbane, QLD, May 2018, pp. 910–917 (cited on pages 11, 12, 32–36, 39–41, 43).
- [Bek69] M. Bekker. *Introduction to Terrain-Vehicle Systems*. Univ. of Michigan Press, 1969 (cited on page 80).
- [Bis+10] R. Bischoff, J. Kurth, G. Schreiber, R. Koeppe, A. Albu-Schäffer, A. Beyer, O. Eiberger, S. Haddadin, A. Stemmer, G. Grunwald, and G. Hirzinger. “The KUKA-DLR Lightweight Robot arm - a new reference platform for robotics research and manufacturing”. In: *Proceedings of the 41st International Symposium on Robotics (ISR) and the 6th German Conference on Robotics (ROBOTIK)*. 2010, pp. 1–8 (cited on page 71).
- [BK20] S. Barthelmes and U. Konigorski. “Model-Based Chassis Control System for an over-Actuated Planetary Exploration Rover”. In: *at - Automatisierungstechnik* 68.1 (Jan. 2020), pp. 58–71 (cited on page 8).
- [BMH08] M. Bajracharya, M. W. Maimone, and D. Helmick. “Autonomy for Mars Rovers: Past, Present, and Future”. In: *Computer* 41.12 (Dec. 2008), pp. 44–50 (cited on page 2).
- [Bow+24] J. Bowkett, M. Dolci, J. Aldrich, D. P. Moreno, A. Boettcher, J. Kim, A. Okon, J. Townsend, and C. Collins. “Challenges in Closed-loop Compliant Motion Control for Planetary Robotics”. In: *Proceedings of the 2024 IEEE Aerospace Conference*. Mar. 2024, pp. 1–11 (cited on pages 108, 109).
- [Bro+83] R. W. Brockett et al. “Asymptotic stability and feedback stabilization”. In: *Differential geometric control theory* 27.1 (1983), pp. 181–191 (cited on pages 28, 53, 116).
- [Bru+16] S. G. Brunner, F. Steinmetz, R. Belder, and A. Dömel. “RAFCON: A Graphical Tool for Engineering Complex, Robotic Tasks”. In: *Proceedings of the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Oct. 2016, pp. 3283–3290 (cited on page 16).
- [Bru+18] S. G. Brunner, P. Lehner, M. J. Schuster, S. Riedel, R. Belder, D. Leidner, A. Wedler, M. Beetz, and F. Stulp. “Design, Execution, and Postmortem Analysis of Prolonged Autonomous Robot Operations”. In: *IEEE Robotics and Automation Letters* 3.2 (Apr. 2018), pp. 1056–1063 (cited on pages 19, 119).

-
- [BT23] N. Baldauf and A. Turnwald. “Iterative learning-based model predictive control for mobile robots in space applications”. In: *Proceedings of the 27th International Conference on Methods and Models in Automation and Robotics (MMAR)*. 2023, pp. 434–439 (cited on page 92).
 - [Bur+24] L. Burkhard, R. Sakagami, K. Lakatos, H. Gmeiner, P. Lehner, J. Reill, M. G. Müller, M. Durner, and A. Wedler. “Collaborative Multi-Rover Crater Exploration: Concept and Results from the ARCHES Analog Mission”. In: *Proceedings of the 2024 IEEE Aerospace Conference*. Mar. 2024, pp. 1–14 (cited on page 120).
 - [Bus+18] K. Bussmann, L. Meyer, F. Steidle, and A. Wedler. “Slip Modeling and Estimation for a Planetary Exploration Rover: Experimental Results from Mt. Etna”. In: *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Madrid, Spain, Oct. 2018, pp. 2449–2456 (cited on pages 11, 12, 21, 79, 85–88, 92, 112).
 - [BV04] M. Benosman and G. L. Vey. “Control of Flexible Manipulators: A Survey”. In: *Robotica* 22.5 (Sept. 2004), pp. 533–545 (cited on page 107).
 - [CBD96] G. Campion, G. Bastin, and B. D’Andréa-Novel. “Structural Properties and Classification of Kinematic and Dynamic Models of Wheeled Mobile Robots”. In: *IEEE Transactions on Robotics and Automation* 12.1 (1996), pp. 47–62 (cited on pages 24, 25, 28).
 - [CH08] F. Chaumette and S. Hutchinson. “Visual servoing and visual tracking”. In: *Handbook of Robotics* (2008), pp. 563–583 (cited on page 107).
 - [Con+08] C. P. Connette, A. Pott, M. Hägele, and A. Verl. “Control of an Pseudo-omnidirectional, Non-holonomic, Mobile Robot based on an ICM Representation in Spherical Coordinates”. In: *Proceedings of the 47th IEEE Conference on Decision and Control (CDC)*. Dec. 2008, pp. 4976–4983 (cited on page 25).
 - [Cou20] R. Courant. “Über die Eigenwerte bei den Differentialgleichungen der mathematischen Physik”. In: *Mathematische Zeitschrift* 7.1 (Mar. 1920), pp. 1–57 (cited on page 116).
 - [DBA12] S. Djebrani, A. Benali, and F. Abdessemed. “Modelling and control of an omnidirectional mobile manipulator”. en. In: *International Journal of Applied Mathematics and Computer Science* 22.3 (Sept. 2012), pp. 601–616 (cited on page 7).
 - [De +06] A. De Luca, A. Albu-Schäffer, S. Haddadin, and G. Hirzinger. “Collision Detection and Safe Reaction with the DLR-III Lightweight Manipulator Arm”. In: *Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Oct. 2006, pp. 1623–1630 (cited on page 73).
 - [Die+16] A. Dietrich, K. Bussmann, F. Petit, P. Kotyczka, C. Ott, B. Lohmann, and A. Albu-Schäffer. “Whole-body impedance control of wheeled mobile manipulators: Stability analysis and experiments on the humanoid robot Rollin’ Justin”. In: *Autonomous Robots (AURO): Special Issue on Whole-Body Control of Contacts and Dynamics for Humanoid Robots* 40.3 (Mar. 2016), pp. 505–517 (cited on pages 6, 39).
 - [Die+21] A. Dietrich, X. Wu, K. Bussmann, M. Harder, M. Iskandar, J. Engelsberger, C. Ott, and A. Albu-Schäffer. “Practical Consequences of Inertia Shaping for Interaction and Tracking in Robot Control”. In: *Control Engineering Practice* 114 (Sept. 2021), p. 104875 (cited on pages 5, 45, 71).

- [Die15] A. Dietrich. *Whole-Body Impedance Control of Wheeled Humanoid Robots*. Vol. 116. Springer Tracts in Advanced Robotics. Springer International Publishing, 2015 (cited on pages 6, 37, 38, 44).
- [Din+11] L. Ding, H. Gao, Z. Deng, K. Nagatani, and K. Yoshida. “Experimental Study and Analysis on Driving Wheels’ Performance for Planetary Exploration Rovers Moving in Deformable Soil”. In: *Journal of Terramechanics* 48 (Jan. 2011), pp. 27–45 (cited on page 89).
- [DO20] A. Dietrich and C. Ott. “Hierarchical Impedance-Based Tracking Control of Kinetically Redundant Robots”. In: *IEEE Transactions on Robotics* 36.1 (2020), pp. 204–221 (cited on pages 6, 7).
- [DOA13] A. Dietrich, C. Ott, and A. Albu-Schäffer. “Multi-Objective Compliance Control of Redundant Manipulators: Hierarchy, Control, and Stability”. In: *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Nov. 2013, pp. 3043–3050 (cited on page 38).
- [DOA15] A. Dietrich, C. Ott, and A. Albu-Schäffer. “An overview of null space projections for redundant, torque-controlled robots”. In: *The International Journal of Robotics Research* 34.11 (Sept. 2015), pp. 1385–1400 (cited on page 7).
- [DOG06] A. De Luca, G. Oriolo, and P. Giordano. “Kinematic modeling and redundancy resolution for nonholonomic mobile manipulators”. en. In: *Proceedings of the 2006 IEEE International Conference on Robotics and Automation (ICRA)*. Orlando, FL, USA: IEEE, 2006, pp. 1867–1873 (cited on page 5).
- [DOG10] A. De Luca, G. Oriolo, and P. R. Giordano. “Kinematic control of nonholonomic mobile manipulators in the presence of steering wheels”. en. In: *Proceedings of the 2010 IEEE International Conference on Robotics and Automation (ICRA)*. Anchorage, AK: IEEE, May 2010, pp. 1792–1798 (cited on page 5).
- [Don+07] Dong Xu, Dongbin Zhao, Jianqiang Yi, and Xiangmin Tan. “Robust Adaptive Tracking Control of Omnidirectional Wheeled Mobile Manipulators”. In: *Proceedings of the 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*. San Diego, CA, USA: IEEE, Oct. 2007, pp. 3598–3603 (cited on page 7).
- [DOP18] A. Dietrich, C. Ott, and J. Park. “The Hierarchical Operational Space Formulation: Stability Analysis for the Regulation Case”. In: *IEEE Robotics and Automation Letters* 3.2 (Apr. 2018), pp. 1120–1127 (cited on page 7).
- [DWA11] A. Dietrich, T. Wimböck, and A. Albu-Schäffer. “Dynamic Whole-Body Mobile Manipulation with a Torque Controlled Humanoid Robot via Impedance Control Laws”. In: *Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Sept. 2011, pp. 3199–3206 (cited on page 4).
- [Ell16] A. Ellery. *Planetary Rovers: Robotic Exploration of the Solar System*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016 (cited on pages 2, 3, 77, 108).
- [Eng+20] V. S. Engelschön, S. R. Eriksson, A. Cowley, M. Fateri, A. Meurisse, U. Kueppers, and M. Sperl. “EAC-1A: A Novel Large-Volume Lunar Regolith Simulant”. In: *Scientific Reports* 10.1 (Mar. 2020), p. 5473 (cited on page 97).
- [Eur24] European Space Agency (ESA). *EXPLORE 2024. The European Exploration Strategy*. 2024. URL: <https://esamultimedia.esa.int/docs/HRE/Explore%5C%5F2040.pdf> (visited on 10/30/2024) (cited on page 2).

-
- [FB97] E. D. Fasse and J. F. Broenink. “A Spatial Impedance Controller for Robotic Manipulation”. In: *IEEE Transactions on Robotics and Automation* 13.4 (Aug. 1997), pp. 546–556 (cited on page 4).
 - [FBD08] H. J. Ferreau, H. G. Bock, and M. Diehl. “An online active set strategy to overcome the limitations of explicit MPC”. In: *International Journal of Robust and Nonlinear Control* 18.8 (2008), pp. 816–830 (cited on page 32).
 - [Fer+14] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl. “qpOASES: a parametric active-set algorithm for quadratic programming”. In: *Mathematical Programming Computation* 6.4 (Dec. 2014), pp. 327–363 (cited on page 32).
 - [Fis05] E. Fischer. “Über quadratische Formen mit reellen Koeffizienten”. In: *Monatshefte für Mathematik und Physik* 16.1 (Dec. 1905), pp. 234–249 (cited on page 116).
 - [Fre82] E. Freund. “Fast Nonlinear Control with Arbitrary Pole-Placement for Industrial Robots and Manipulators”. In: *The International Journal of Robotics Research* 1.1 (Mar. 1982), pp. 65–78 (cited on page 4).
 - [Ger07] L. Geraghty. *Living with Star Trek: American Culture and the Star Trek Universe*. Annotated edition. London England: PAPERBACKSHOP UK IMPORT, June 2007 (cited on page 1).
 - [Ger25] German Aerospace Center, Institute of Robotics and Mechatronics. *RAFCON*. 2025. URL: <https://dlr-rm.github.io/RAFCON/> (visited on 04/13/2025) (cited on page 16).
 - [GH14] C. Gruber and M. Hofbaur. “Practically stabilizing motion control of mobile robots with steering wheels”. In: *Proceedin2014 IEEE Conference on Control Applications (CCA)*. Oct. 2014, pp. 1312–1317 (cited on page 53).
 - [GI18] R. Gonzalez and K. Iagnemma. “Slippage estimation and compensation for planetary exploration rovers. State of the art and future challenges”. In: *Journal of Field Robotics* 35.4 (June 2018), pp. 564–577 (cited on page 8).
 - [Gil24] J. W. Gilbert. “Assessing Computing Reliability Design Practices for Low-SWAP Spacecraft Avionics by Comparison of the Ingenuity Helicopter and Dragonfly Missions”. Available at <https://jscholarship.library.jhu.edu/items/ebd2b798-dcdc-4575-96c2-3b864cfdd252>. [Accessed Dec. 19, 2024]. MA thesis. Johns Hopkins University, May 2024 (cited on page 108).
 - [GL96] G. H. Golub and C. F. V. Loan. *Matrix Computations*. 3rd ed. Baltimore, Maryland: John Hopkins University Press, 1996 (cited on pages 116, 117).
 - [GP11] L. Grüne and J. Pannek. “Nonlinear Model Predictive Control”. In: *Nonlinear Model Predictive Control*. London: Springer London, 2011, pp. 43–66 (cited on page 92).
 - [GSS98] F. Ghorbel, B. Srinivasan, and M. W. Spong. “On the uniform boundedness of the inertia matrix of serial robot manipulators”. In: *Journal of Robotic Systems* 15.1 (1998), pp. 17–28 (cited on pages 26, 28).
 - [HDA17] S. Haddadin, A. De Luca, and A. Albu-Schäffer. “Robot collisions: A survey on detection, isolation, and identification”. In: *IEEE Transactions on Robotics* 33.6 (Dec. 2017), pp. 1292–1312 (cited on page 73).
 - [Hel+05] D. Helmick, Yang Cheng, D. Clouse, M. Bajracharya, L. Matthies, and S. Roumeliotis. “Slip compensation for a Mars rover”. In: *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Edmonton, Alta., Canada: IEEE, 2005, pp. 2806–2813 (cited on pages 8, 90, 91).

- [Hel+06] D. M. Helmick, S. I. Roumeliotis, Y. Cheng, D. S. Clouse, M. Bajracharya, and L. H. Matthies. “Slip-compensated path following for planetary exploration rovers”. In: *Advanced Robotics* 20.11 (2006), pp. 1257–1280 (cited on pages 90, 91).
- [Hev+13] M. Heverly, J. Matthews, J. Lin, D. Fuller, M. Maimone, J. Biesiadecki, and J. Leichty. “Traverse Performance Characterization for the Mars Science Laboratory Rover”. In: *Journal of Field Robotics* 30.6 (2013), pp. 835–846 (cited on page 8).
- [Hew+20] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger. “Learning-Based Model Predictive Control: Toward Safe Learning in Control”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 3.1 (May 2020), pp. 269–296 (cited on pages 9, 92).
- [HK00] R. Holmberg and O. Khatib. “Development and control of a holonomic mobile robot for mobile manipulation tasks”. In: *International Journal of Robotics Research* 19.11 (2000), pp. 1066–1074 (cited on page 7).
- [HMR+13] R. Haarmann, Q. Mühlbauer, L. Richter, et al. “Mobile Payload Element (MPE): concept study of a small, outonomous and innovative sample fetching rover”. In: *Proceedings of the 12th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*. 2013 (cited on page 13).
- [Hog85] N. Hogan. “Impedance Control: An Approach to Manipulation: Part I - Theory, Part II - Implementation, Part III - Applications”. In: *Journal of Dynamic Systems, Measurement, and Control* 107 (Mar. 1985), pp. 1–24 (cited on pages 4, 36).
- [HS19] Y. Huang and J. Su. “Visual Servoing of Nonholonomic Mobile Robots: A Review and a Novel Perspective”. In: *IEEE Access* 7 (2019), pp. 134968–134977 (cited on page 107).
- [Hyo+19] S.-H. Hyon, Y. Ida, J. Ishikawa, and M. Hiraoka. “Whole-Body Locomotion and Posture Control on a Torque-Controlled Hydraulic Rover”. In: *IEEE Robotics and Automation Letters* 4.4 (Oct. 2019), pp. 4587–4594 (cited on pages 5, 8).
- [ID04a] K. Iagnemma and S. Dubowsky. *Mobile Robots in Rough Terrain: Estimation, Motion Planning, and Control with Application to Planetary Rovers*. Springer Science & Business Media, July 2004 (cited on page 8).
- [ID04b] K. Iagnemma and S. Dubowsky. “Traction Control of Wheeled Robotic Vehicles in Rough Terrain with Application to Planetary Rovers”. In: *The International Journal of Robotics Research* 23.10-11 (Oct. 2004), pp. 1029–1040 (cited on pages 8, 32).
- [INY06] G. Ishigami, K. Nagatani, and K. Yoshida. “Path Following Control with Slip Compensation on Loose Soil for Exploration Rover”. In: *Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Oct. 2006, pp. 5552–5557 (cited on page 9).
- [INY09] G. Ishigami, K. Nagatani, and K. Yoshida. “Slope Traversal Controls for Planetary Exploration Rover on Sandy Terrain”. In: *Journal of Field Robotics* 26.3 (2009), pp. 264–286 (cited on page 9).
- [Isk+19] M. Iskandar, G. Quere, A. Hagengruber, A. Dietrich, and J. Vogel. “Employing Whole-Body Control in Assistive Robotics”. In: *Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Macau, China: IEEE, Nov. 2019, pp. 5643–5650 (cited on page 6).
- [JF13] T. A. Johansen and T. I. Fossen. “Control allocation—A survey”. In: *Automatica* 49.5 (May 2013), pp. 1087–1103 (cited on pages 8, 32).

-
- [Kab+19] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger. “Learning-Based Model Predictive Control for Autonomous Racing”. In: *IEEE Robotics and Automation Letters* 4.4 (2019), pp. 3363–3370 (cited on page 9).
 - [Kan+01] S. Kang, K. Komoriya, K. Yokoi, T. Koutoku, and K. Tanie. “Utilization of Inertial Effect in Damping-Based Posture Control of Mobile Manipulator”. In: *Proceedings of the 2001 IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 2. May 2001, pp. 1277–1282 (cited on page 6).
 - [Kan+90] Y. Kanayama, Y. Kimura, F. Miyazaki, and T. Noguchi. “A Stable Tracking Control Method for an Autonomous Mobile Robot”. In: *Proceedings of the 1990 IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 1. May 1990, pp. 384–389 (cited on page 94).
 - [Kel04] A. Kelly. “Linearized Error Propagation in Odometry”. In: *International Journal of Robotics Research* 23 (2004), pp. 179–218 (cited on pages 8, 83).
 - [Kha02] H. K. Khalil. *Nonlinear Systems (Third Edition)*. Prentice Hall, 2002 (cited on page 115).
 - [Kha87] O. Khatib. “A Unified Approach for Motion and Force Control of Robot Manipulators: The Operational Space Formulation”. In: *IEEE Journal of Robotics and Automation* RA-3.1 (Feb. 1987), pp. 43–53 (cited on pages 4, 37, 38).
 - [Kim+19] S. Kim, K. Jang, S. Park, Y. Lee, S. Y. Lee, and J. Park. “Whole-Body Control of Non-holonomic Mobile Manipulator Based on Hierarchical Quadratic Programming and Continuous Task Transition”. In: *Proceedings of the 2019 IEEE 4th International Conference on Advanced Robotics and Mechatronics (ICARM)*. July 2019, pp. 414–419 (cited on page 6).
 - [Kin17] T. Kindsmüller. “Vertical Wheel Force Control for a Planetary Rover”. Available at <https://elib.dlr.de/112394/>. [Accessed 22 Oct. 2024]. MA thesis. Technische Universität München, May 2017 (cited on page 18).
 - [Kre+08] A. Krebs, T. Thueer, E. Carrasco, and R. Siegwart. “Towards torque control of the CRAB rover”. In: *Proceedings of the 9th International Symposium on Artificial Intelligence, Robotics and Automation in Space*. 2008, p. 8 (cited on page 8).
 - [KSY15] C. T. Kiang, A. Spowage, and C. K. Yoong. “Review of Control and Sensor System of Flexible Manipulator”. In: *Journal of Intelligent & Robotic Systems* 77.1 (Jan. 2015), pp. 187–213 (cited on page 107).
 - [Lak+24] K. Lakatos, D. Lakatos, X. Wu, P. Kotyczka, and A. Dietrich. “On Passivity-Based Trajectory Tracking for Robotic Manipulators Combining PD+ and Slotine-Li Control”. In: *at - Automatisierungstechnik* 72.12 (Dec. 2024), pp. 1195–1206 (cited on pages 11, 12, 66, 68, 69, 71–76).
 - [Lak+25] K. Lakatos, D. Lakatos, P. Kotyczka, and A. Dietrich. “Passivity-Based Tracking Control of a Planetary Rover with Robotic Arm”. In: *IEEE Transactions on Field Robotics* (2025). (Accepted for publication.) (Cited on pages 11, 12, 44, 53, 54, 56–65).
 - [Lam+04] P. Lamon, A. Krebs, M. Lauria, R. Siegwart, and S. Shooter. “Wheel torque control for a rough terrain rover”. In: *Proceedings of the 2004 IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 5. New Orleans, LA, USA: IEEE, 2004, pp. 4682–4687 (cited on page 8).

- [Lan+03] G. Lanzafame, M. Neri, V. Acocella, A. Billi, R. Funiciello, and G. Giordano. “Structural features of the July–August 2001 Mount Etna eruption: evidence for a complex magma supply system”. In: *Journal of the Geological Society* 160.4 (July 2003), pp. 531–544 (cited on page 84).
- [LaS68] J. P. LaSalle. “Stability Theory for Ordinary Differential Equations”. In: *Journal of Differential Equations* 4.1 (Jan. 1968), pp. 57–65 (cited on page 38).
- [LBT24] K. Lakatos, N. Baldauf, and A. Turnwald. “Towards Learning-Based Trajectory Tracking Control for a Planetary Exploration Rover: Development and Testing”. In: *Proceedings of the 2024 International Conference on Space Robotics (iSpaRo)*. Luxembourg, June 2024, pp. 299–306 (cited on pages 11, 12, 92, 99).
- [Leb+19] Q. Leboutet, E. Dean-Leon, F. Bergner, and G. Cheng. “Tactile-Based Whole-Body Compliance With Force Propagation for Mobile Manipulators”. In: *IEEE Transactions on Robotics* 35.2 (Apr. 2019), pp. 330–342 (cited on page 6).
- [Leh+18] P. Lehner, S. Brunner, A. Dömel, H. Gmeiner, S. Riedel, B. Vodermayr, and A. Wedler. “Mobile Manipulation for Planetary Exploration”. In: *Proceedings of the 2018 IEEE Aerospace Conference*. Mar. 2018, pp. 1–11 (cited on pages 15, 119).
- [Leh+23] P. Lehner, R. Sakagami, W. Boerdijk, A. Dömel, M. Durner, G. Franchini, A. Prince, K. Lakatos, D. L. Risch, L. Meyer, B. Vodermayr, E. Dietz, S. Frohmann, F. Seel, S. Schröder, H.-W. Hübers, A. Albu-Schäffer, and A. Wedler. “Mobile Manipulation of a Laser-induced Breakdown Spectrometer for Planetary Exploration”. In: *Proceedings of the 2023 IEEE Aerospace Conference*. Mar. 2023, pp. 1–19 (cited on pages 15, 31, 120).
- [Li+14] W. Li, H. Gao, H. Yang, N. Li, L. Ding, and Z. Deng. “A Method to Online Estimate Wheel’s Slippage for Planetary Rover”. In: *Proceeding of the 11th World Congress on Intelligent Control and Automation*. June 2014, pp. 2469–2474 (cited on page 9).
- [Lin+02] D. Lindgren, T. Hague, P. P. Smith, and J. Marchant. “Relating Torque and Slip in an Odometric Model for an Autonomous Agricultural Vehicle”. In: *Autonomous Robots* 13.1 (July 2002), pp. 73–86 (cited on page 8).
- [May14] D. Q. Mayne. “Model Predictive Control: Recent Developments and Future Promise”. In: *Automatica* 50.12 (Dec. 2014), pp. 2967–2986 (cited on page 92).
- [Mes+22] A. Mesbah, K. P. Wabersich, A. P. Schoellig, M. N. Zeilinger, S. Lucia, T. A. Badgwell, and J. A. Paulson. “Fusion of Machine Learning and MPC under Uncertainty: What Advances are on the Horizon?” In: *Proceedings of the 2022 American Control Conference (ACC)*. June 2022, pp. 342–357 (cited on pages 9, 92).
- [Mey17] L. Meyer. “Pose Estimation and Traction Control for a Planetary Exploration Rover”. MA thesis. Technische Universität München, 2017 (cited on pages 21, 79, 80, 90, 91).
- [MLS94] R. M. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994 (cited on pages 26, 67).
- [OB04] L. Ojeda and J. Borenstein. “Methods for the Reduction of Odometry Errors in Over-Constrained Mobile Robots”. In: *Autonomous Robots* 16.3 (May 2004), pp. 273–286 (cited on page 8).

-
- [ODV02] G. Oriolo, A. De Luca, and M. Vendittelli. “WMR Control via Dynamic Feedback Linearization: Design, Implementation, and Experimental Validation”. In: *IEEE Transactions on Control Systems Technology* 10.6 (Nov. 2002), pp. 835–852 (cited on page 53).
 - [Oje+06] L. Ojeda, D. Cruz, G. Reina, and J. Borenstein. “Current-Based Slippage Detection and Odometry Correction for Mobile Robots and Planetary Rovers”. In: *IEEE Transactions on Robotics* 22.2 (Apr. 2006), pp. 366–378 (cited on page 8).
 - [ONŽ03] D. Omrčen, B. Nemec, and L. Žlajpah. “Torque-Velocity Control Algorithm for on-Line Obstacle Avoidance for Mobile Manipulators”. In: *Proceedings of the 2003 IEEE International Conference on Industrial Technology*. Vol. 2. Dec. 2003, pp. 784–789 (cited on page 6).
 - [OS88] R. Ortega and M. Spong. “Adaptive Motion Control of Rigid Robots: A Tutorial”. In: *Proceedings of the 27th IEEE Conference on Decision and Control (CDC)*. Vol. 2. Dec. 1988, pp. 1575–1584 (cited on page 5).
 - [Ott+08] C. Ott, A. Albu-Schäffer, A. Kugi, and G. Hirzinger. “On the Passivity-Based Impedance Control of Flexible Joint Robots”. In: *IEEE Transactions on Robotics* 24.2 (Apr. 2008), pp. 416–429 (cited on page 71).
 - [Ott08] C. Ott. *Cartesian Impedance Control of Redundant and Flexible-Joint Robots*. Vol. 49. Springer Tracts in Advanced Robotics. Springer Publishing Company, Berlin Heidelberg, 2008 (cited on pages 4, 36).
 - [OŽN04] D. Omrčen, L. Žlajpah, and B. Nemec. “Autonomous motion of a mobile manipulator using a combined torque and velocity control”. In: *Robotica* 22.6 (Nov. 2004), pp. 623–632 (cited on page 6).
 - [PFC07] V. Padois, J.-Y. Fourquet, and P. Chiron. “Kinematic and dynamic model-based control of wheeled mobile manipulators: a unified framework for reactive approaches”. In: *Robotica* 25.2 (Mar. 2007), pp. 157–173 (cited on page 6).
 - [PGB12] L. Preston, M. Grady, and S. Barber. *TN2: The Catalogue of Planetary Analogues*. Published: The Planetary and Space Sciences Research Institute, The Open University, UK. Dec. 2012 (cited on pages 84, 119).
 - [PP88] B. Paden and R. Panja. “Globally asymptotically stable ‘PD+’ controller for robot manipulators”. In: *International Journal of Control* 47.6 (June 1988). Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/00207178808906130>, pp. 1697–1712 (cited on pages 5, 66, 67).
 - [Raj+24] V. Rajendran, B. Debnath, S. Mghames, W. Mandil, S. Parsa, S. Parsons, and A. Ghalamzan-E. “Towards Autonomous Selective Harvesting: A Review of Robot Perception, Robot Design, Motion Planning and Control”. In: *Journal of Field Robotics* 41.7 (2024), pp. 2247–2279 (cited on page 5).
 - [RB20] U. Rosolia and F. Borrelli. “Learning How to Autonomously Race a Car: A Predictive Control Approach”. In: *IEEE Transactions on Control Systems Technology* 28.6 (Nov. 2020), pp. 2713–2719 (cited on page 9).
 - [RDM18] C. Ren, Y. Ding, and S. Ma. “Passivity-Based Active Disturbance Rejection Control of an Omnidirectional Mobile Robot”. In: *2018 IEEE 8th Annual International Conference on CYBER Technology in Automation, Control, and Intelligent Systems (CYBER)*. July 2018, pp. 1513–1518 (cited on page 7).

- [Sak+23] R. Sakagami, S. G. Brunner, A. Dömel, A. Wedler, and F. Stulp. “ROSMC: A High-Level Mission Operation Framework for Heterogeneous Robotic Teams”. In: *Proceedings of the 2023 IEEE International Conference on Robotics and Automation (ICRA)*. May 2023, pp. 5473–5479 (cited on page 120).
- [Sch+12] K. Schmid, F. Ruess, M. Suppa, and D. Burschka. “State Estimation for highly dynamic flying Systems using Key Frame Odometry with varying Time Delays”. In: *Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Oct. 2012, pp. 2997–3004 (cited on pages 21, 78, 79).
- [Sch+16] M. J. Schuster, C. Brand, S. G. Brunner, P. Lehner, J. Reill, S. Riedel, T. Bodenmüller, K. Bussmann, S. Buttner, A. Domel, W. Friedl, I. Grix, M. Hellerer, H. Hirschmüller, M. Kassecker, Z.-C. Marton, C. Nissler, F. Ruess, M. Suppa, and A. Wedler. “The LRU Rover for Autonomous Planetary Exploration and Its Success in the SpaceBotCamp Challenge”. In: *Proceedings of the International Conference on Autonomous Robot Systems and Competitions (ICARSC)*. Bragança, Portugal: IEEE, May 2016, pp. 7–14 (cited on page 15).
- [Sch+19] M. J. Schuster, S. G. Brunner, K. Bussmann, S. Büttner, A. Dömel, M. Hellerer, H. Lehner, P. Lehner, O. Porges, J. Reill, S. Riedel, M. Vayugundla, B. Vodermayr, T. Bodenmüller, C. Brand, W. Friedl, I. Grix, H. Hirschmüller, M. Kassecker, Z.-C. Márton, C. Nissler, F. Ruess, M. Suppa, and A. Wedler. “Towards Autonomous Planetary Exploration: The Lightweight Rover Unit (LRU), its Success in the SpaceBotCamp Challenge, and Beyond”. In: *Journal of Intelligent & Robotic Systems (JINT)* 93.3-4 (Nov. 2019), pp. 461–494 (cited on pages 15–17, 78, 119).
- [Sch+20] M. J. Schuster, M. G. Müller, S. G. Brunner, H. Lehner, P. Lehner, R. Sakagami, A. Dömel, L. Meyer, B. Vodermayr, R. Giubilato, M. Vayugundla, J. Reill, F. Steidle, I. von Barga, K. Bussmann, R. Belder, P. Lutz, W. Stürzl, M. Smíšek, M. Maier, S. Stoneman, A. F. Prince, B. Rebele, M. Durner, E. Staudinger, S. Zhang, R. Pöhlmann, E. Bischoff, C. Braun, S. Schröder, E. Dietz, S. Frohmann, A. Börner, H.-W. Hübers, B. Foing, R. Triebel, A. O. Albu-Schäffer, and A. Wedler. “The ARCHES Space-Analogue Demonstration Mission: Towards Heterogeneous Teams of Autonomous Robots for Collaborative Scientific Sampling in Planetary Exploration”. In: *IEEE Robotics and Automation Letters* 5.4 (2020), pp. 5315–5322 (cited on pages 14, 120).
- [Sch08] G. Schulz. *Regelungstechnik 2: Mehrgrößenregelung, Digitale Regelungstechnik, Fuzzy-Regelung*. 2nd ed. Munich, Germany: Oldenbourg Wissenschaftsverlag GmbH, 2008 (cited on pages 82, 95).
- [Sch20] F. Schmidt. *Links and Nodes source code*. 2020. URL: <https://gitlab.com/links%5C%5Fand%5C%5Fnodes/links%5C%5Fand%5C%5Fnodes> (visited on 10/17/2024) (cited on page 16).
- [See+13] N. Seegmiller, F. Rogers-Marcovitz, G. A. Miller, and A. Kelly. “Vehicle Model Identification by Integrated Prediction Error Minimization”. In: *International Journal of Robotics Research* 32.8 (July 2013), pp. 912–931 (cited on pages 8, 80–82, 84, 89, 92).
- [Ser89] H. Seraji. “Configuration control of redundant manipulators: theory and implementation”. In: *IEEE Transactions on Robotics and Automation* 5.4 (1989), pp. 472–490 (cited on pages 8, 38, 51).

-
- [Ser93] H. Seraji. “An On-Line Approach to Coordinated Mobility and Manipulation”. In: *Proceedings of the 1993 IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 1. May 1993, pp. 28–35 (cited on page 6).
- [Ser98] H. Seraji. “A Unified Approach to Motion Control of Mobile Manipulators”. In: *The International Journal of Robotics Research* 17.2 (Feb. 1998), pp. 107–118 (cited on page 5).
- [SH90] N. Sadegh and R. Horowitz. “Stability and Robustness Analysis of a Class of Adaptive Controllers for Robotic Manipulators”. In: *The International Journal of Robotics Research* 9.3 (June 1990), pp. 74–92 (cited on pages 66, 68, 75).
- [Sid18] A. A. Siddiqi. *Beyond Earth: A Chronicle of Deep Space Exploration, 1958-2016*. Second edition. NASA SP 2018-4041. Available as a free download at <https://www.nasa.gov/history/history-publications-and-resources/nasa-history-series/beyond-earth/> [Accessed Dec. 05, 2024]. Washington, DC: National Aeronautics and Space Administration, Office of Communications, NASA History Division, 2018 (cited on page 2).
- [SK14] N. Seegmiller and A. Kelly. “Enhanced 3D Kinematic Modeling of Wheeled Mobile Robots”. In: *Robotics: Science and Systems X*. Robotics: Science and Systems Foundation, July 2014 (cited on pages 8, 79, 80, 92).
- [SL87] J.-J. E. Slotine and W. Li. “On the Adaptive Control of Robot Manipulators”. In: *The International Journal of Robotics Research* 6.3 (Sept. 1987), pp. 49–59 (cited on pages 5, 66, 67).
- [SL91] J.-J. E. Slotine and W. Li. *Applied Nonlinear Control*. Englewood Cliffs, NJ: Prentice Hall, 1991 (cited on pages 69, 71).
- [SOK90] M. Spong, R. Ortega, and R. Kelly. “Comments on ‘Adaptive manipulator control: a case study’ by J. Slotine and W. Li”. In: *IEEE Transactions on Automatic Control* 35.6 (1990), pp. 761–762 (cited on page 67).
- [Spi89] A. J. Spiessbach. “Hazard Avoidance For A Mars Rover”. In: *Mobile Robots III*. Vol. 1007. SPIE, Mar. 1989, pp. 77–84 (cited on page 3).
- [SPP13] L. Sentis, J. Petersen, and R. Philippsen. “Implementation and stability analysis of prioritized whole-body compliant controllers on a wheeled humanoid robot in uneven terrains”. In: *Autonomous Robots* 35.4 (2013), pp. 301–319 (cited on page 7).
- [SRB14] K. Schmid, F. Ruess, and D. Burschka. “Local reference filter for life-long vision aided inertial navigation”. In: *Proceedings of the 17th International Conference on Information Fusion (FUSION)*. IEEE. July 2014, pp. 1–8 (cited on page 78).
- [SS91] B. Siciliano and J.-J. Slotine. “A General Framework for Managing Multiple Tasks in Highly Redundant Robotic Systems”. In: *Proceedings of the 5th IEEE International Conference on Advanced Robotics (ICAR)*. Vol. 2. Pisa, Italy, June 1991, pp. 1211–1216 (cited on page 38).
- [SSK17] A. Siravuru, S. V. Shah, and K. M. Krishna. “An Optimal Wheel-Torque Control on a Compliant Modular Robot for Wheel-Slip Minimization”. In: *Robotica* 35.2 (Feb. 2017), pp. 463–482 (cited on page 8).
- [Sto+24] F. Storiale, E. Ferrentino, F. Salvioi, K. Kapellos, and P. Chiacchio. *Optimal Whole Body Trajectory Planning for Mobile Manipulators in Planetary Exploration and Construction*. May 2024. arXiv: 2405.14363 [cs] (cited on page 5).

- [SV85] M. Spong and M. Vidyasagar. “Robust Nonlinear Control of Robot Manipulators”. In: *Proceedings of the 24th IEEE Conference on Decision and Control (CDC)*. Fort Lauderdale, FL, USA: IEEE, Dec. 1985, pp. 1767–1772 (cited on page 4).
- [SYZ19] P. Song, Y. Yu, and X. Zhang. “A Tutorial Survey and Comparison of Impedance Control on Robotic Manipulation”. In: *Robotica* 37.5 (May 2019), pp. 801–836 (cited on page 4).
- [TA81] M. Takegaki and S. Arimoto. “A New Feedback Method for Dynamic Control of Manipulators”. In: *Journal of Dynamic Systems, Measurement, and Control* 103.2 (June 1981), pp. 119–125 (cited on page 5).
- [Tar+99] M. Tarokh, G. McDermott, S. Hayati, and J. Hung. “Kinematic Modeling of a High Mobility Mars Rover”. In: *Proceedings of the 1999 IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 2. May 1999, pp. 992–998 (cited on page 8).
- [TdM96] B. Thuilot, B. d’Andréa-Noël, and A. Micaelli. “Modeling and feedback control of mobile robots equipped with several steering wheels”. In: *IEEE Transactions on Robotics and Automation* 12.3 (June 1996), pp. 375–390 (cited on page 53).
- [TM07] M. Tarokh and G. McDermott. “A Systematic Approach to Kinematics Modeling of High Mobility Wheeled Rovers”. In: *Proceedings of the 2007 IEEE International Conference on Robotics and Automation (ICRA)*. Apr. 2007, pp. 4905–4910 (cited on page 8).
- [Tre+09] A. Trebi-Ollennu, R. Volpe, R. Bonitz, M. Robinson, and J. Carsten. “In Situ Robotic Arm Operations”. In: *IEEE Robotics & Automation Magazine* 16.4 (Dec. 2009), pp. 34–43 (cited on page 2).
- [TX01] J. Tan and N. Xi. “Unified Model Approach for Planning and Control of Mobile Manipulators”. In: *Proceedings of the 2001 IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 3. May 2001, pp. 3145–3152 (cited on page 6).
- [Ume+99] Y. Umeda, D. Nakamura, T. Murakami, and K. Ohnishi. “Hybrid Position/Force Control of a Mobile Manipulator Based on Cooperative Task Sharing”. In: *Proceedings of the IEEE International Symposium on Industrial Electronics (ISIE)*. Vol. 1. July 1999, pp. 139–144 (cited on page 6).
- [Vay+18] M. Vayugundla, F. Steidle, M. Smisek, M. J. Schuster, K. Bussmann, and A. Wedler. “Datasets of Long Range Navigation Experiments in a Moon Analogue Environment on Mount Etna”. In: *Proceedings of the 50th International Symposium on Robotics (ISR)*. June 2018, pp. 1–7 (cited on pages 21, 92, 120).
- [Vay+25] M. Vayugundla, T. Bodenmüller, L. Burkhard, M. J. Schuster, B.-M. Steinmetz, M. Sewtz, N. Borgsmüller, F. Buse, W. Stürzl, R. Giubilato, F. Schuler, M. G. Müller, M. Kuhne, L. Jörg, A. Lund, A. Wedler, R. Triebel, M. Smisek, and M. Grebenstein. “The DLR Autonomous Navigation Experiment with the IDEFIX Rover: Software Architecture, Autonomous Navigation Features and Preliminary Operations Concept”. In: *Proceedings of the 2025 IEEE Aerospace Conference*. Big Sky, Montana, Mar. 2025 (cited on page 112).
- [WA07] K. J. Waldron and M. E. Abdallah. “An Optimal Traction Control Scheme for Off-Road Operation of Robotic Vehicles”. In: *IEEE/ASME Transactions on Mechatronics* 12.2 (Apr. 2007), pp. 126–133 (cited on page 8).

-
- [WBK07] G. D. White, R. M. Bhatt, and V. N. Krovi. “Dynamic redundancy resolution in a nonholonomic wheeled mobile manipulator”. In: *Robotica* 25.2 (Mar. 2007), pp. 147–156 (cited on page 6).
- [Wed+15a] A. Wedler, M. Hellerer, B. Rebele, H. Gmeiner, B. Vodermayr, T. Bellmann, S. Barthelmes, R. Rosta, C. Lange, L. Witte, N. Schmitz, M. Knapmeyer, A. Czeluske, L. Thomsen, C. Waldmann, S. Flögel, M. Wilde, and Y. Takei. “ROBEX – Components and methods for the planetary exploration demonstration mission”. In: *Proceedings of the 13th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*. ESA Website, 2015 (cited on pages 15, 119).
- [Wed+15b] A. Wedler, B. Rebele, J. Reill, M. Suppa, H. Hirschmüller, C. Brand, M. Schuster, B. Vodermayr, H. Gmeiner, A. Maier, B. Willberg, K. Bussmann, F. Wappler, M. Hellerer, and R. Lichtenheldt. “LRU-lightweight rover unit”. In: *Proceedings of the 13th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*. 2015 (cited on pages 13, 16).
- [Wed+17] A. Wedler, M. Vayugundla, H. Lehner, P. Lehner, M. Schuster, S. Brunner, W. Stürzl, A. Dömel, H. Gmeiner, B. Vodermayr, B. Rebele, I. Grix, K. Bussmann, J. Reill, B. Willberg, A. Maier, P. Meusel, F. Steidle, M. Smisek, M. Hellerer, M. Knapmeyer, F. Sohl, A. Heffels, L. Witte, C. Lange, R. Rosta, N. Toth, and S. Völk. “First Results of the ROBEX Analogue Mission Campaign: Robotic Deployment of Seismic Networks for Future Lunar Missions”. In: *Proceedings of the 68th International Astronautical Congress (IAC)*. Adelaide, Australia: International Astronautical Federation (IAF), Sept. 2017 (cited on pages 14, 119).
- [Wed+18] A. Wedler, K. Bussmann, A. Dömel, M. Drauschke, H. Gmeiner, I. L. Grix, H. Lehner, M. Vayugundla, M. G. Müller, J. Reill, M. Schuster, W. Stürzl, B. Vodermayr, P. Lehner, S. Brunner, and M. Görner. “From the ROBEX experiment toward the Robotic Deployment and Maintenance of Scientific Infrastructure for future Planetary Exploration Missions”. In: *42nd COSPAR Scientific Assembly*. Vol. 42. Pasadena, California, July 2018 (cited on page 119).
- [Wed+22] A. Wedler, M. G. Müller, M. Schuster, M. Durner, P. Lehner, A. Dömel, F. Steidle, M. Vayugundla, R. Sakagami, L. Meyer, M. Smisek, W. Stürzl, N. Schmitz, B. Vodermayr, A. Fonseca Prince, A. Ehreiser, E. Staudinger, R. Pöhlmann, S. Zhang, M. Hellerer, R. Lichtenheldt, D. Franke, A. F. X. Pignede, W. Schindler, M. Schütt, B. Rebele, W. Boerdijk, R. Giubilato, J. Reill, M. Kuhne, J. Lee, S. Schröder, S. Frohmann, F. Seel, E. Dietz, R. Triebel, N. Y.-S. Lii, E. Bischoff, C. Braun, S. Kille, K. Wormnes, A. Pereira, W. Carey, A. P. Rossi, L. Thomsen, T. Graber, T. Krüger, A. Börner, P. Irmisch, K. Bussmann, G. Paar, A. Bauer, S. Völk, H. Rauer, H.-W. Hübers, J. Bals, S. Hohmann, T. Asfour, B. Foing, and A. O. Albu-Schäffer. “Finally! Insights into the ARCHES Lunar Planetary Exploration Analogue Campaign on Etna in Summer 2022”. In: *Proceedings of the 73rd International Astronautical Congress (IAC)*. Paris, France: International Astronautical Federation (IAF), Sept. 2022 (cited on page 120).
- [WFM23] J. Wang, M. T. H. Fader, and J. A. Marshall. “Learning-Based Model Predictive Control for Improved Mobile Robot Path Following Using Gaussian Processes and Feedback Linearization”. In: *Journal of Field Robotics* 40.5 (2023), pp. 1014–1033 (cited on page 9).

- [WGM89] B. H. Wilcox, D. B. Gennery, and A. H. Mishkin. “Mars Rover Local Navigation And Hazard Avoidance”. In: *Mobile Robots III*. Ed. by W. J. Wolfe. Vol. 1007. International Society for Optics and Photonics. Boston, MA: SPIE, Mar. 1989, pp. 72–76 (cited on page 3).
- [Whi+09] G. D. White, R. M. Bhatt, C. P. Tang, and V. N. Krovi. “Experimental Evaluation of Dynamic Redundancy Resolution in a Nonholonomic Wheeled Mobile Manipulator”. In: *IEEE/ASME Transactions on Mechatronics* 14.3 (June 2009), pp. 349–357 (cited on page 6).
- [Wit+23] S. Withrow-Maser, W. Johnson, N. Schatzman, L. Young, H. Cummings, C. Malpica, L. Meyn, NASA Langley Research Center, B. Allan, Jet Propulsion Laboratory, T. Tzanetos, H. Grip, Science and Technology Corporation, W. Koning, A. Chan, A. Ruan, AeroVironment Inc., B. Pipenberg, and M. Keennon. “Mars Sample Recovery Helicopter: Rotorcraft to Retrieve the First Samples from the Martian Surface”. In: *Proceedings of the Vertical Flight Society 79th Annual Forum*. West Palm Beach, Florida USA: The Vertical Flight Society, May 2023, pp. 1–8 (cited on page 108).
- [Wu+21] Y. Wu, E. Lamon, F. Zhao, W. Kim, and A. Ajoudani. “Unified Approach for Hybrid Motion Control of MOCA Based on Weighted Whole-Body Cartesian Impedance Formulation”. In: *IEEE Robotics and Automation Letters* 6.2 (Apr. 2021), pp. 3505–3512 (cited on page 6).
- [Wu+23] X. Wu, C. Ott, A. Albu-Schäffer, and A. Dietrich. “Passive Decoupled Multitask Controller for Redundant Robots”. In: *IEEE Transactions on Control Systems Technology* 31.1 (Jan. 2023), pp. 1–16 (cited on page 7).
- [Xia+18] K. Xia, H. Gao, L. Ding, G. Liu, Z. Deng, Z. Liu, and C. Ma. “Trajectory Tracking Control of Wheeled Mobile Manipulator Based on Fuzzy Neural Network and Extended Kalman Filtering”. In: *Neural Computing and Applications* 30.2 (July 2018), pp. 447–462 (cited on page 7).
- [Yos09] K. Yoshida. “Achievements in Space Robotics”. In: *IEEE Robotics & Automation Magazine* 16.4 (Dec. 2009), pp. 20–28 (cited on pages 2, 77).
- [Yua+06] Yuan Ping Li, T. Zielinska, M. Ang, and Wei Lin. “Wheel-ground interaction modelling and torque distribution for a redundant mobile robot”. In: *Proceedings of the 2006 IEEE International Conference on Robotics and Automation (ICRA)*. Orlando, FL, USA: IEEE, 2006, pp. 3362–3367 (cited on page 8).
- [YY92] Y. Yamamoto and X. Yun. “Coordinating Locomotion and Manipulation of a Mobile Manipulator”. In: *Proceedings of the 31st IEEE Conference on Decision and Control (CDC)*. Vol. 3. Dec. 1992, pp. 2643–2648 (cited on page 5).
- [Zha05] F. Zhang, ed. *The Schur Complement and Its Applications*. Numerical Methods and Algorithms 4. New York: Springer Science and Business Media, 2005 (cited on page 70).