



- Hey mate! Do you think they can see us?



UNIVERSIDAD DE MÁLAGA



AUTONOMOUS NAVIGATION OF PLANETARY ROVERS

- Martin Azkarate Vecilla -
PhD Thesis

- Look son, that's the origins of our species...

- Woh! And how did they arrive here in the first place?

- God built them...




YEAR 3.000 ON MARS...





UNIVERSIDAD
DE MÁLAGA

AUTOR: Martin Azkarate Vecilla

 <https://orcid.org/0000-0003-3284-5422>

EDITA: Publicaciones y Divulgación Científica. Universidad de Málaga



Esta obra está bajo una licencia de Creative Commons Reconocimiento-NoComercial-SinObraDerivada 4.0 Internacional:

<http://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>

Cualquier parte de esta obra se puede reproducir sin autorización pero con el reconocimiento y atribución de los autores.

No se puede hacer uso comercial de la obra y no se puede alterar, transformar o hacer obras derivadas.

Esta Tesis Doctoral está depositada en el Repositorio Institucional de la Universidad de Málaga (RIUMA): riuma.uma.es

AUTONOMOUS NAVIGATION
OF
PLANETARY ROVERS

MARTIN AZKARATE VECILLA
PHD THESIS

Doctor of Philosophy in Mechatronics Engineering
Department of Systems Engineering and Automation
Faculty of Industrial Engineering
University of Malaga



UNIVERSIDAD DE MÁLAGA

in collaboration with




Supervisor: Dr. Carlos J. Pérez Del Pulgar
Advisor: Prof. Alfonso J. García Cerezo
ESA Advisor: Gianfranco Visentin

Malaga, 2021



UNIVERSIDAD
DE MÁLAGA

Martin Azkarate Vecilla: *Autonomous Navigation of Planetary Rovers*
PhD Thesis, 2021

 <https://orcid.org/0000-0003-3284-5422>

Some rights reserved.





DECLARACIÓN DE AUTORÍA Y ORIGINALIDAD DE LA TESIS PRESENTADA PARA OBTENER EL TÍTULO DE DOCTOR

D./Dña MARTIN AZKARATE VECILLA

Estudiante del programa de doctorado INGENIERÍA MECATRÓNICA de la Universidad de Málaga, autor/a de la tesis, presentada para la obtención del título de doctor por la Universidad de Málaga, titulada: AUTONOMOUS NAVIGATION OF PLANETARY ROVERS

Realizada bajo la tutorización de ALFONSO J. GARCÍA CEREZO y dirección de CARLOS J. PÉREZ DEL PULGAR MANCEBO (si tuviera varios directores deberá hacer constar el nombre de todos)

DECLARO QUE:

La tesis presentada es una obra original que no infringe los derechos de propiedad intelectual ni los derechos de propiedad industrial u otros, conforme al ordenamiento jurídico vigente (Real Decreto Legislativo 1/1996, de 12 de abril, por el que se aprueba el texto refundido de la Ley de Propiedad Intelectual, regularizando, aclarando y armonizando las disposiciones legales vigentes sobre la materia), modificado por la Ley 2/2019, de 1 de marzo.

Igualmente asumo, ante a la Universidad de Málaga y ante cualquier otra instancia, la responsabilidad que pudiera derivarse en caso de plagio de contenidos en la tesis presentada, conforme al ordenamiento jurídico vigente.

En Málaga, a 12 de JULIO de 2021

Fdo.: MARTIN AZKARATE VECILLA Doctorando/a	Fdo.: ALFONSO JOSE GARCIA CEREZO Tutor/a
Fdo.: CARLOS JESÚS PÉREZ DEL PULGAR MANCEBO Director/es de tesis	

*Hay palabras que valen más que mil imágenes, y obras que valen más que mil palabras,
por ello;*

“Buenos pensamientos, buenas palabras, y buenas obras.”

ZARATHUSTRA

A mi Hermano, por hacerme sentir orgulloso.
A mi Padre, por nunca esperar menos de mí.
A mi Madre, por no parar hasta convencerme.
A María, por ser mi fiel compañera de viaje en la vida.
A todos aquellos que creen en la ciencia y en la investigación como motores de vida.

ABSTRACT

From the experience and lessons learnt out of planetary rover missions such as the Mars Exploration Rovers or Curiosity, and looking forward to the upcoming missions of ExoMars and the Mars Sample Return programme, this thesis aims to study and enhance the autonomous navigation capabilities onboard the rovers, in order to increase the overall mission scientific return. As we see more and more autonomy being embarked in space missions, planetary rovers are already relying on self-navigating functionalities to fulfil their mission objectives. The constraints given by the space environment and the limitations in communications found in Mars exploration missions render the implementation of autonomous navigation capabilities as the most efficient solution, or the only, to extend the traversed distance per sol (a Martian day) beyond the few tens of meters.

Autonomous Navigation is a complex capability that relies on the implementation of several functionalities and their orchestrated execution in order to perform. Those functionalities are, as a minimum, Localisation, Perception or Mapping, Path Planning, and Path or Trajectory Control. In this thesis, these core robotics functions are investigated and newly developed methods that focus on the planetary rovers scenario are proposed. These methods take into consideration the constraints and conditions found in Mars missions, making them particularly fit and targeted for rover systems.

Different ways to combine and execute these functionalities are studied in order to compose the Guidance, Navigation and Control (GNC) subsystem of a rover that is capable of navigating autonomously. The potential variation in the terrain conditions found across a rover mission is taken into consideration, adapting the behaviour and functionalities run by the control architecture, with the objective of maximising the length of traverse per sol. A navigation mode is proposed targeting the relatively benign terrain cases. This mode relies solely on a stereo camera that mimics the ExoMars Localisation Cameras (LocCam) sensor to both localise the rover and avoid the hazards along the path. This is done without requiring periodic stops of the rover, increasing the effective traverse speed. The navigation mode is experimentally demonstrated in an exhaustive field test campaign totalling over 500 m of autonomous traverse. For rough terrain cases, a navigation mode that builds upon the first mode is proposed. It integrates a newly developed Simultaneous Localisation and Mapping (SLAM) component for planetary rovers that allows for improved accuracy of rover

localisation and smoother detection and avoidance of terrain hazards. The map produced by the SLAM function is used in a Global Localisation component to correct the accumulated drift in long-range traverses. These two components are experimentally validated making use of a representative dataset gathered in a field test campaign run in a planetary analogue terrain near the Teide Volcano in the Canary Islands. Finally, the thesis proposes the design of a complete GNC architecture, that integrates the two navigation modes and could see a potential exploitation in future Mars missions such as the Sample Fetch Rover.

RESUMEN

A partir de la experiencia obtenida de las misiones de rovers planetarios, como los Mars Exploration Rovers o Curiosity, y en vista de las futuras misiones de ExoMars y el programa Mars Sample Return, esta tesis tiene como objetivo estudiar y mejorar las capacidades a bordo de los rovers para incrementar el retorno científico general de la misión. A medida que vemos misiones espaciales en las que cada vez se embarca una mayor capacidad de autonomía, los rovers planetarios dependerán de las funcionalidades de navegación autónoma para cumplir con los objetivos de su misión. Las restricciones dadas por el entorno espacial y las limitaciones en las comunicaciones, que se dan en las misiones de exploración a Marte, hacen que la implementación de capacidades de navegación autónoma sea la solución más eficiente, o tal vez la única, para extender la distancia recorrida por sol (día marciano) más allá de las pocas decenas de metros.

La navegación autónoma es una capacidad compleja que se basa en la implementación de varias funcionalidades y la ejecución orquestada de éstas para su desempeño. Esas funcionalidades son, como mínimo, Localización, Percepción o Mapeo, Planificación de Ruta y Control de Trayectoria. En esta tesis se investigan todas estas funciones robóticas esenciales y se proponen nuevos métodos que se centran en el caso particular de los rovers planetarios. Éstos tienen en consideración las limitaciones y condiciones que se encuentran en las misiones a Marte, lo que las hace particularmente adecuadas y específicas para estos sistemas rover.

Se han estudiado diferentes formas de combinar y ejecutar estas funcionalidades, que en conjunto componen el subsistema de Guiado, Navegación y Control (GNC) de un rover que es capaz de navegar de forma autónoma. Se toma en consideración la potencial variación en las condiciones del terreno encontradas en una misión rover, adaptando el comportamiento y las funcionalidades ejecutadas por la arquitectura de control, con el objetivo de maximizar la distancia recorrida por sol. Se propone un modo de navegación enfocado a los casos de terreno de dificultad media o baja. Este modo se basa únicamente en una cámara estéreo, que imita al sensor LocCam de ExoMars, para localizar el rover y evitar los obstáculos a lo largo del camino. Esto se consigue sin requerir que el rover se pare periódicamente, lo que aumenta la velocidad efectiva. El modo de navegación se valida de forma experimental en una exhaustiva campaña de pruebas de campo que acumula más de 500 m de recorrido autónomo. Para los casos

de terreno más dificultoso, se propone un modo de navegación que, basándose en el primer modo, añade nuevas capacidades al rover. Éste integra un nuevo componente de Localización y Mapeo Simultaneos (SLAM) desarrollado para rovers planetarios que permite mejorar la precisión de la localización del rover y una detección más lejana de los peligros del terreno para poder eviatarlos de forma más fluida. El mapa producido por la función SLAM se utiliza posteriormente en un componente de Localización Global para corregir la deriva acumulada en travesías de largo alcance. Estos dos componentes son validados experimentalmente utilizando un conjunto de datos representativos adquiridos en una campaña de prueba de campo llevada a cabo en un terreno planetario análogo cerca del volcán Teide en Tenerife (Islas Canarias). Finalmente, la tesis propone el diseño de la arquitectura GNC completa, que integra ambos modos de navegación y podría ver una potencial explotación en futuras misiones a Marte como el Sample Fetch Rover del programa Mars Sample Return.

ACKNOWLEDGMENTS

First of all, I owe my deepest gratitude to my main PhD supervisor Dr. Carlos J. Pérez Del Pulgar for the assistance provided and for the advice, patience and comments he has made in the elaboration of the papers and this thesis. I am also grateful to him for giving me the opportunity to perform this thesis and for pushing me to believe this was possible.

I would also like to express my gratitude to Gianfranco Visentin, for allowing me to conduct this thesis in parallel to my full time job and for being such an inspiration for the last decade. There are intelligent people that impress you when you meet them, and there are few brilliant minds that inspire you every day.

To all members and alumni of the Automation and Robotics Section of the European Space Agency, the brightest group of colleagues one could ever have. And, among them, to Levin Gerdes in particular.

CONTENTS

I THESIS DISSERTATION

1	INTRODUCTION	3
1.1	Planetary Exploration: Mission to Mars	3
1.1.1	Roving Mars: The Challenges of the Red Planet	4
1.1.2	Increasing the scientific return with Autonomy	5
1.1.3	Autonomous Navigation in Planetary Missions	6
1.2	Thesis Contributions	7
1.3	Context and Motivation	9
1.4	Thesis Outline	10
2	A REVIEW OF THE STATE OF THE ART	11
2.1	Localisation	12
2.1.1	Relative Localisation	12
2.1.2	Global Localisation	19
2.2	Perception	24
2.3	Simultaneous Localisation and Mapping (SLAM)	28
2.4	Path Planning	29
2.5	Trajectory Control	36
2.6	Autonomous Navigation	39
3	OVERVIEW OF THESIS CONTRIBUTIONS	43
3.1	Previous Relevant Work	44
3.1.1	Development of Planetary Rover Prototypes for R&D activities	44
3.1.2	Improving Rover Locomotion Capabilities with Wheel Walking	48
3.1.3	Remote Rover Operations Field Test Campaigns	51
3.2	Dataset Acquisition Field Tests	54
3.3	Efficient Autonomous Navigation	56
3.3.1	Trajectory Control	57
3.3.2	Hazard Detection	58
3.3.3	Path Planner	60
3.3.4	Experimental Results	63
3.4	SLAM for Autonomous Planetary Rovers with Global Localisation	65
3.4.1	SLAM	66
3.4.2	Global Localisation	67

3.4.3	Experimental Results	68
3.5	Design of a complete GNC Architecture	72
3.6	Closing remarks	73
4	CONCLUSIONS	75
	BIBLIOGRAPHY	79
II MAIN AUTHORED PUBLICATIONS		
A	EFFICIENT AUTONOMOUS NAVIGATION FOR PLANETARY ROVERS WITH LIMITED RESOURCES	91
B	SLAM FOR AUTONOMOUS PLANETARY ROVERS WITH GLOBAL LOCALIZATION	93
C	A GNC ARCHITECTURE FOR PLANETARY ROVERS WITH AUTONOMOUS NAVIGATION	95
D	DESIGN, TESTING AND EVOLUTION OF MARS ROVER TESTBEDS FOR ESA PLANETARY EXPLORATION	97
III RESUMEN		
E	RESUMEN	101

LIST OF FIGURES

- Figure 1.1 [MSR](#) programme conceptual representation (credits: European Space Agency ([ESA](#))) 4
- Figure 1.2 Martian landscape (credits: National Aeronautics and Space Administration ([NASA](#))) 5
- Figure 2.1 Basic conceptual representation of autonomous navigation functionalities 11
- Figure 2.2 Left: [ExoMars](#) wheel being tested in a Single Wheel Testbed. Right: [ExoMars](#) wheel under Locomotion tests (Oettershagen et al., 2019) 13
- Figure 2.3 Illustration of the rover kinematics model using contact points for improved 3D delta pose estimation (Hidalgo-Carrió et al., 2017). 14
- Figure 2.5 Illustration of a rover traverse and the network with PanCam and NavCam images for bundle adjustment (Li et al., 2004). 18
- Figure 2.6 Representation of the global localisation using two-way Doppler with Mars Odyssey (Guinn, 2001). 20
- Figure 2.7 NavCam mosaic from Opportunity in Meridiani Planum. Opportunity landed in the 20 m diameter Eagle Crater in the middle of this mosaic (Matthies et al., 2007). 20
- Figure 2.8 Skyline matching method representation used in [VIPER](#) algorithm (Chiodini et al., 2017). 21
- Figure 2.9 Reconstructed Opportunity wide baseline NavCam [DEM](#) and [HiRISE DEM](#) at Victoria Crater in 3D showing possibilities of registration (Tao, Muller, and Poole, 2016). 22
- Figure 2.10 Example of co-registered wide baseline NavCam [ORI](#) and [HiRISE ORI](#) (left) and corrected traverse (right) for Curiosity (Tao, Muller, and Poole, 2016). 23
- Figure 2.11 Points on the mesh built from [M2020](#) images are rendered from multiple virtual viewpoints (left). Synthetic templates are matched to real [SFR](#) images to compute points in camera frame (middle). Both point sets are aligned to estimate the [SFR](#) pose (right, conceptual rover depicted) (Pham et al., 2021). 24
- Figure 2.12 [M2020](#) flight NavCams (Maki et al., 2020). 26

- Figure 2.13 [ExoMars](#) camera sensors on top of the mast (Miles, Gunn, and Coates, 2020). 26
- Figure 2.14 Virtual representation of the [GESTALT](#) system for terrain assessment (Maimone, Leger, and Biesiadecki, 2007). 27
- Figure 2.15 Representation of keyframe selection approach in (Hidalgo-Carrió, Poulakis, and Kirchner, 2018) for its Adaptive Graph [SLAM](#) solution. 30
- Figure 2.16 Grid-based path planning from s to g . *Any-angle* algorithm solution (in blue) finds smooth paths without being restricted in the direction of connecting nodes (Ferguson and Stentz, 2005a) compared to the more classic angle-constrained path planner solution (in black). 32
- Figure 2.17 Typical suboptimal path obtained from randomised planning (red path) and the corresponding shortcut path (green) (Elbanhawi and Simic, 2014). 33
- Figure 2.18 Mosaic looking backward on the tracks of a Field D* checkout on [MER](#) (Carsten et al., 2009). 34
- Figure 2.19 Demonstration of how wheel placement planning can obtain a path solution in the presence of rocks with specific paths for each wheel (Ono et al., 2015). 34
- Figure 2.20 [ExoMars](#) path planning with kino-dynamic constraints principle (Bora et al., 2017). 35
- Figure 2.21 Illustration of trajectory control concept (Bora et al., 2017). 37
- Figure 2.22 Conceptual representation *pure pursuit* working principle as implemented by [NASA](#) on [MER](#) (Helmick et al., 2006). The carrot coordinates represent the instantaneous waypoint along the path towards which the rover is steered. 38
- Figure 2.23 Perseverance rover and Ingenuity helicopter on Mars soon after landing, ready to begin their missions. 40
- Figure 2.24 [ExoMars](#) [GTM](#) during final validation tests in Turin. 41
- Figure 2.25 Early conceptual design of [SFR](#). 42
- Figure 3.1 Computer rendered image of the [ExoMars](#) mission rover, Rosalind Franklin 45
- Figure 3.2 The [BEMA](#) Triple Bogie configuration and right bogie beam with actuator locations (Poulakis et al., 2015). 45
- Figure 3.3 [ExoTeR](#) at the [PRL](#) of [ESA](#) 46
- Figure 3.4 [MaRTA](#) at the [PRL](#) of [ESA](#) 47
- Figure 3.5 [HDPR](#) during a field test near [ESA-ESTEC](#) 48
- Figure 3.6 Wheel Walking test in loose sand trap. 50

- Figure 3.7 [ExoTeR](#) at DLR performing a Wheel Walking test. 51
- Figure 3.8 [ExoTeR](#) ready for Remote Operations tests at [SEROM](#) in CNES. 52
- Figure 3.9 Example [DEM](#) data gathered from the rover's telemetry. 53
- Figure 3.10 [ExoTeR](#) at post-egress challenging traverse. Left to the rover is the rock that it eventually hit. 53
- Figure 3.11 [HDPR](#) during the field test campaign in Katwijk. 55
- Figure 3.12 [HDPR](#) during the field test campaign in Tenerife. 55
- Figure 3.13 *c-pursuit* behaviour representation: Vector fields of direction to *look-ahead point* from all possible inside the safety corridor. 57
- Figure 3.14 Hazard Detector Initial Test 59
- Figure 3.15 The underlying geometry on which the hazard detection is based. The tolerated distances from the camera are computed in such a way that hazards would be detected at certain heights. 60
- Figure 3.16 Graphical example case where the rover finds a hazard on its way (a) and plans an alternate path to avoid it (b). 62
- Figure 3.17 Panoramic view of the field test setup. 63
- Figure 3.18 Efficient Autonomous Navigation test traverse. The red arrows represent the pre-planned waypoints while the thin blue line shows the actual traverse. The total planned traverse length was 171 m and the actual executed trajectory length 191 m. 64
- Figure 3.19 The high-level design diagram of the [SLAM](#) system. 67
- Figure 3.20 The particle cloud (red dots) of the particle filter distributed around the rover pose. The trajectory of the rover with respect to the initial state is represented by the orange line. The current local elevation map colors represent the range of height values from high (red) to low (purple). 68
- Figure 3.21 Graphical illustration of how template matching is performed between the local gradient image (template) and the global gradient image (reference). 69

Figure 3.22	Perspective view of the SLAM process visualization. The trajectory of the rover is represented by the red line. The current local elevation map and global map (background) colors represent the relative separate range of height values of each map from high (red) to low (purple). 70
Figure 3.23	Relative Localization experiments traverses 70
Figure 3.24	Orthographic views of the environment and the rover's traverses in it. The black and red colors correspond to the ground truth and SLAM paths accordingly. The pose correction is visible in the end of the traverse. 71
Figure 3.25	Schematic concept of the two-level GNC architecture for Autonomous Navigation of Planetary Rovers. Orange coloured boxes refer to sensors, green to developed components and greyed components to re-used from open-source community. 73

LIST OF TABLES

Table 2.1	Summary table with localisation accuracy and frequency for each method 16
Table 3.1	Summary table of results of Traverse 2 71

ACRONYMS

UMA	University of Malaga
ESA	European Space Agency
ESTEC	European Space Research and Technology Centre
PRL	Planetary Robotics Lab
NASA	National Aeronautics and Space Administration
JPL	Jet Propulsion Laboratory
CNES	French Space Agency
ADS	Airbus Defense and Space
BEMA	Bogie Electro-Mechanical Actuator
CFA	Cumulative Fractional Area
CV	Computer Vision
DARPA	Defense Advanced Research Projects Agency
DEM	Digital Elevation Model
DLR	German Space Agency
DTE	Direct To Earth
EKF	Extended Kalman Filters
ExoMars	Exobiology on Mars
ExoTeR	ExoMars Testing Rover
FDIR	Fault Detection Isolation and Recovery
FMM	Fast Marching Method
FM*	Heuristic Fast Marching
FoV	Field of View
GCS	Ground Control Station
GESTALT	Grid-based Estimation of Surface Traversability Applied to Local Terrain
GNC	Guidance Navigation and Control
GNSS	Global Navigation Satellite System

GTM	Ground Test Model
HDPR	Heavy Duty Planetary Rover
HiRISE	High Resolution Imaging Science Experiment
IBA	Incremental Bundle Adjustment
ICP	Iterative Closest Point
ICR	Instantaneous Center of Rotation
IMU	Inertial Measurement Unit
ISS	International Space Station
LocCam	Localisation Camera
M2020	Mars2020
MaRTA	Martian Rover Testbed for Autonomy
MER	Mars Exploration Rovers
MOGA	Multiple-frame Odometry-compensated Global Alignment
MRO	Mars Reconnaissance Orbiter
MSL	Mars Science Laboratory
MSR	Mars Sample Return
NavCam	Navigation Camera
LiDAR	Light Detection And Ranging
OBC	On Board Computer
ORI	Ortho-Rectified Image
PanCam	Panoramic Cameras
PTU	Pan-Tilt Unit
RAMS	Reliability Availability Maintainability and Safety
RMC	Robotics and Mechatronics Center
ROCC	Rover Operations Control Centre
RoI	Region of Interest
RTG	Radioisotope Thermoelectric Generators
SEROM	<i>Site d'Essai pour les Rovers Mobiles</i>
SFR	Sample Fetch Rover
SLAM	Simultaneous Localisation and Mapping
SPARTAN	SPARing Robotics Technologies for Autonomous Navigation

ToF	Time of Flight
TRL	Technology Readiness Level
UHF	Ultra High Frequency
VIPER	Visual Position Estimator for Rovers
VO	Visual Odometry

Part I

THESIS DISSERTATION

INTRODUCTION

1.1 PLANETARY EXPLORATION: MISSION TO MARS

Thanks to robotic planetary exploration, nowadays humans reach remote places that they could not imagine to get a glance of otherwise. While rich and audacious entrepreneurs race and bet on the year in which we will become a multi-planetary species, one could argue we have actually long crossed that line. Before mankind ever sets foot on Mars, our knowledge of the red planet will be such that we might even know that life ever existed on Mars. Moreover, in one decade from now we could be on the verge of bringing a Martian soil sample back to Earth. The Exobiology on Mars ([ExoMars](#)) mission of [ESA](#) on one side, and the [NASA-ESA](#) collaboration on the Mars Sample Return ([MSR](#)) program on the other side, are paving the way to make these objectives reality. Similar to how previous missions like Mars Science Laboratory ([MSL](#)), Mars Exploration Rovers ([MER](#)), Phoenix or the early Mars Pathfinder have made possible for us to even think of such a complex mission like [MSR](#) in the first place (see [Figure 1.1](#)).

There is no doubt of the invaluable science that planetary missions bring to society and the knowledge that we have gained about Mars since the Viking mission landers relayed their first pictures from Mars. Mission after mission, the learning of the Martian geology and potential habitability is enabled, and continuously supported, by a range of spacecraft spanning from orbiting satellites over landing probes to exploratory rovers. Eventually, robotics exploration usually acts as a precursor to human exploration, rather than replacing it. The instruments and mechanisms carried by planetary rovers allow performing in-situ science before any potential human arrival, and let prepare for the challenges that a future manned exploration could entail. However, let us focus on how these first in-situ experiments are made possible and how planetary rovers are engineered to succeed in their mission objectives and execute their tasks.




Figure 1.1: MSR programme conceptual representation (credits: ESA)

1.1.1 *Roving Mars: The Challenges of the Red Planet*

As difficult as exciting that the Entry, Descent and Landing phase of a mission to Mars can be, the challenges that come after a safe landing are no less overwhelming. As the book by Steve W. Squyres with the same title explains, *Roving Mars* is the great adventure. Its harsh environment with a thin atmosphere (1% of Earth's atmospheric pressure) lets surface temperature cycle in a range of up to 100° within a sol. At the same time, the lack of a magnetic field exposes the planet surface directly to all cosmic radiation coming from Space, making it even more inhospitable. Besides, there is weather on Mars, and surprisingly, wind can pick up dust devils or even create dust storms that increase the atmospheric optical depth to so high values that the sun flux is barely noticed at the Martian surface. Considering that the sun is the main source of power of most planetary rovers, except for those running on nuclear Radioisotope Thermoelectric Generators (RTG)s, dust storms are long dark nights for rovers that can bring their mission to an end. On top of all that, remoteness is arguably the most constraining challenge, since the distance of 2.66 au between Mars and Earth (at their farthest orbital coordinates) means that an electromagnetic signal can take up to 22 min to propagate between the two planets. Additionally, such a signal would require a massive, powerful antenna that rovers generally lack due to system constraints in mass and power budget. Instead, rovers rely on one or several Mars orbiting satellites that can relay their communications with Earth. However, rovers and their orbiting companions do not have visibility at all times and therefore, communication windows only happen few times a sol, if any.




Figure 1.2: Martian landscape (credits: [NASA](#))

This means that rovers need to spend most of the time of the day performing their tasks “on their own”. Finally, the Martian terrain is the last challenge to overcome. As seen in the landscape image of [Figure 1.2](#), in order to navigate rovers need to find their way across obstacles, slopes and a mix of rocky and sandy soil terrain characteristics. These can lead the rover to stuck situations like sand traps made of very loose soil that might only be noticed once it is too late.

All in all, rigorous engineering processes that come up with designs that fulfil all those hard constraints have brought up systems capable of withstanding the Martian environment and roving on it. The system engineering challenge is widely recognised in planetary robotics and fitting all subsystems and functionalities within the tight mass and power budgets requires an engineering masterpiece to be solved. But moreover, it is thanks to the autonomous features and capabilities onboard the rovers that Mars exploration missions have so far been successful in providing the scientific data they were designed for.

1.1.2 *Increasing the scientific return with Autonomy*

A planetary rover is like a spacecraft on the surface of a planet. As such, it comprises several subsystems typically found on any spacecraft like the power generation, management and energy storage, thermal control, structures and mechanisms, radiation protection, communications and data handling and storage system. And, of course, we should add to all this the more robotics-related subsystems of mobility and navigation. Housekeeping telemetry allows monitoring the health status of the overall system on Ground. Yet, given the

limitations in communications, it is mandatory to have an onboard Mission Management System that monitors the available resources and can react to events in order to guarantee the safety of the rover and prevents it from taking any action that could put it at risk. Despite the implementation of these self-monitoring and other Fault Detection Isolation and Recovery (FDIR) features, one would not define rovers as autonomous systems just yet, since we are still far from seeing a planetary rover with embarked deliberative capabilities that could change its course of action on its daily tasks. However, one can notice an increasing reliance on autonomous capabilities of the rover to accomplish mission objectives, which is a direct conclusion of the ambitious scientific targets they are built for.

Rovers can often be seen as the vessels of the scientific payload they carry, because in contrast to landers, rovers can traverse the planetary surface to access multiple scientifically interesting targets that would hardly be reachable from a static landing site. Therefore, the more efficient the traverse of a rover is, the greater the potential scientific return of the overall mission can be. For all these reasons, autonomous navigation is meant to play a key role in future planetary rover missions.

1.1.3 *Autonomous Navigation in Planetary Missions*

On Earth, many examples of autonomous navigation system applications exist. The Defense Advanced Research Projects Agency (DARPA) Grand and Urban Challenges demonstrated already more than a decade ago the feasibility of autonomously driven vehicles. Nowadays, many high-tech companies in the automotive industry are developing full self-driving vehicles capable of safely transporting humans in urban traffic areas. However, and despite these advances in technology, we seldom find examples of autonomously navigating systems in space. It was not until recently that the MER and Curiosity rovers have exploited some autonomous navigation capabilities and features, yet within limited scenarios. Upcoming missions are defining their system requirements to include autonomous navigation capabilities onboard, but still the advances are slow. The aforementioned challenges and constraints, imposed by the space environment and space systems, typically leave the solutions adopted in Earth applications unfeasible. Rovers on Mars do not have a connection to the data servers or any Global Navigation Satellite System (GNSS) that are available on Earth. Furthermore, rovers are traversing difficult and unstructured surfaces such as the one shown in Figure 1.2, with just some knowledge of the terrain deduced from orbiter imagery. Although these images are useful for the evaluation of potential landing sites and hazardous zones, they cannot

offer the same level of detail as numerous services provide on our home planet. Consequently, there will be obstacles which are too small to be captured by the orbiter's maps and yet too large for the rover to overcome, and therefore need to be detected and avoided safely by the rover. Another important consequence of the harsh environmental conditions such as vacuum, extreme temperatures, and radiation is that these have created a permanent technology gap for the available electronic components and avionics for space systems. The latest space-qualified computers are only as capable as a 10-year-old computer on Earth. Additionally, sensors such as scanning LiDARs, embarked in many self-driving applications on Earth, have not yet been qualified for space, as their spinning mirrors remain incapable of sustaining the structural stresses during launch. Limitations in system mass and power budgets also play a major role in the design of the system where component selection must be highly optimised. Finally, due to the critical nature and unique opportunities that space missions represent, strong Reliability Availability Maintainability and Safety (RAMS) requirements must be adhered to throughout their development. In particular, all flight software is thoroughly tested against RAMS and many of the commonly used C++ libraries that run on Earth applications do not qualify. All these constraints render the implementations of autonomous navigation features and Computer Vision (CV) algorithms for space a very challenging task, where computational power and efficiency are two of the most important design drivers.

1.2 THESIS CONTRIBUTIONS

This thesis has explored the possibilities to increase the performance and efficiency of the different functionalities running onboard the rover that build up for autonomous navigation and has aimed to come up with solutions that increase the overall locomotion and traverse performance of planetary rovers. It has targeted both the design of novel algorithms and its experimental validation. All in all, the main contributions achieved throughout this thesis can be summarised in:

- **Design and integration of planetary exploration robotics systems for R&D purposes:**

In order to facilitate the core studies of this thesis and provide testing platforms where experimental validation can be conducted, a significant amount of work during the thesis focused on the development of rover platform testbeds that could support this research, including their preparation and final conditioning to perform specific field test campaigns. Special focus was given to the study and design of the mobile chassis

kinematic chain and the control of it in order to improve the locomotion performances of rovers and therefore increase its navigation capabilities.

- **Real sensor dataset acquisition and publication:**

The experimental validation outside of simulated environments is considered relevant for this thesis. At the same time, it is acknowledged the effort involved in the execution of tests on real platforms, specially when performed outside of the lab environment. Therefore, the acquisition of real sensor datasets performed during field testing campaigns in representative analogue terrain is considered a valuable contribution of this thesis. These datasets have not only been available for the development of algorithms in this thesis but have also been published and made accessible to the community.

- **Investigate new methods for functionalities that support autonomous navigation for planetary exploration:**

The core of the work in this thesis focused on studying the algorithmic designs and approaches of different functionalities running onboard the rover related to its navigation. Solutions for trajectory control, path planning, localisation and perception were investigated and novel methods for all of these functions have been developed and introduced that particularly target the conditions found in Mars planetary rover exploration missions.

The trajectory control component is an improved version of the original *pure-pursuit* algorithm that guarantees tracking the path within a bounded safety corridor. The path planning algorithm stems from the *Fast Marching Method*, optimising both the smoothness of the path and component execution time. An efficient hazard detector allows the perception pipeline to be running in parallel to the rover motion. A full [SLAM](#) solution that improves the rover localisation and provides a map of the local terrain in order to navigate safely and efficiently. And a global localisation component based on *Template Matching* that enables long range autonomous navigation by correcting the drift in the rover estimated pose.

- **Experimental validation of implemented algorithms in real analogue conditions:**

An important objective of the thesis focused on the performance validation of the developed algorithms in representative field test campaigns. This shall serve to increase the Technology Readiness Level ([TRL](#)) of their implementations and to demonstrate the overall suitability of the thesis contributions for future Mars rover missions. In this thesis, the developed

and implemented algorithms have been validated either using the aforementioned datasets acquired in previous field test campaigns or in by means of explicit field test campaigns dedicated to their demonstration.

- **Design of a GNC architecture for planetary rovers with autonomous navigation capabilities:**

The Guidance Navigation and Control (**GNC**) of a rover is the subsystem implementing the navigation functionalities and commanding the rover to drive it towards the next target exploration location. The ultimate contribution of the thesis has been the definition of a complete **GNC** architecture that orchestrates the execution of the previously designed and experimentally validated algorithms. This contribution built upon all previous ones and encompasses all developed methods leveraging on their use case in the global picture.

It is worth mentioning here that the work performed and described in this thesis has been done in collaboration with other members and visiting researchers at the Planetary Robotics Lab (**PRL**) of **ESA**, and therefore the contributions are a result of a combined effort. In any case, given the personal lab manager role taken for the last seven years, the work here presented has been either personally developed, or led with guidance and close supervision.

1.3 CONTEXT AND MOTIVATION

This thesis has been conducted in the context of the research studies on robotics technology for Mars rover missions done in the **PRL** of **ESA** and in collaboration with the Space Robotics Lab of the University of Malaga (**UMA**). Common interests and fields of research have led to the continuous collaboration between both labs that is running since 2016, when the first joint activity was kicked off.

The results of this thesis are a demonstration of the evolution of the studies of the **PRL** in one of the main areas of space robotics and for which the lab provides support to actual planetary missions such as **ExoMars** or the future Sample Fetch Rover (**SFR**). The thesis also adds to the continuous development and increase of expertise of the lab in this relevant field of research.

This research was primarily motivated by the addition of autonomously navigating capabilities to the **ExoMars** rover mission, which was reconsidered to be included in the mission baseline after being descopeed for some years. In addition, the relevancy of the investigations and results in this study are also highlighted by the **SFR** mission requirements on fast navigation. These are set in

order to traverse long distances in limited time and fulfil the mission objectives, making [SFR](#) the fastest autonomously navigating rover up to date.

1.4 THESIS OUTLINE

The thesis is divided in four chapters, including this introduction and the last one dedicated to main conclusions and description of potential future work.

The second chapter presents a state of the art analysis on the core topics of this thesis. In particular, it deepens on the pillars of autonomous navigation and the different functionalities that compose it. For each of the functions analysed it provides an explanation of the theoretical principles used by the approaches found in the literature.

The third chapter outlines the main contributions of the thesis. It is divided in sections where each of them highlights and explains briefly a singular contribution that forms a significant part that builds up for the thesis. Given the format of the thesis by compendium of publications, sections provide most important principles and conclusions of each contribution without aiming to deepen into implementation details. For that, each section is accompanied by a relevant publication where detailed information about each contribution can be found and theoretical foundations and experimental results are thoroughly explained. Therefore, the chapter serves as a condensed summary of the main thesis contributions.

The thesis is structured and written in a concise way, to allow the reader quickly grasp the concepts behind the main contributions and relies on the given references to the published articles to provide an extension to each of these concepts.

A REVIEW OF THE STATE OF THE ART

This chapter provides a review of the state of the art on the different core building blocks or functionalities that compose the Guidance Navigation and Control (GNC) architecture of a planetary rover with autonomous navigation. In short, Guidance is responsible for the rover path planning, Navigation is related to localization and mapping, and Control is the subsystem able to command and steer the rover, typically to follow a predefined trajectory. As already mentioned in [Chapter 1](#), rovers need more and more autonomous capabilities onboard to complete their missions, and autonomous navigation is seen as a fundamental added value that can significantly increase the science return.

Generally speaking, autonomous rovers must be capable of robustly localising themselves in the Martian surface with minimum drift and of scanning the terrain ahead in order to detect any hazards to plan a safe path to the next navigation target. Once a safe path has been computed they need to follow it, controlling at all times that the rover trajectory stays within a predefined safety corridor along the path. Once the edge of the last scan of terrain is reached the process is repeated by starting the generation of a new model of the current environment in front. As an standard approach (see [Figure 2.1](#)), this process of sensing, planning and acting is iterated until the final navigation target is reached (Correal and Pajares, 2011).

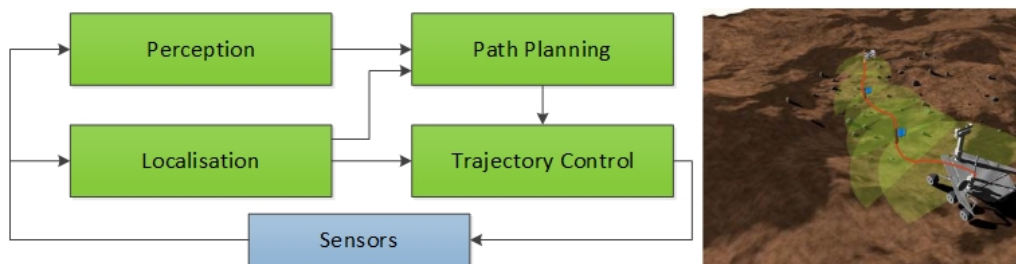


Figure 2.1: Basic conceptual representation of autonomous navigation functionalities

This chapter separates the autonomous navigation in four main functionalities of: localisation, perception, path planning and trajectory control. Each one of these functions is reviewed in dedicated sections providing the most relevant references in the literature and latest advances in the field of planetary missions. Additionally, a section reviewing the application of Simultaneous Localisation and Mapping (SLAM) techniques to planetary missions is included, which relates to the optimised combination of the localisation and perception functions. Finally, the last section addresses the existing designs of complete GNC architectures with autonomous navigation.

2.1 LOCALISATION

Localisation is the function providing an estimate of the rover pose –position and attitude– with respect to a given reference frame. On one side, relative localisation shall estimate the incremental displacement of the rover with regards to a previous point in the traverse such as the starting point of the mission. On the other side, global or absolute localisation aims at positioning the rover within a given global reference fixed to the planetary geodesy. Here, we will address both localisation modes separately, given the different nature of the approaches to solve them.

2.1.1 *Relative Localisation*

There exists different techniques to solve the relative localisation, or in essence, dead reckoning. This is defined as the process of calculating the current position of some moving object by using a previously determined position, or fix, and then incorporating estimations of speed, heading direction, and course over elapsed time. Depending on the sensor measurements used for state estimation, we could name this process as Wheel Odometry, Inertial Odometry or Visual Odometry.

The first one uses the wheel encoders for the position and speed measurements of the different motor joints, that can be used to infer the amount of turns and the direction in which the wheels have turned. Simple kinematic models can then be used to estimate the body motion in the local body frame. Wheel odometry by itself can only provide a rough position estimate and is generally limited to a planar motion estimation. On irregular, uneven terrain with slopes and rocks such as in planetary mission, the simplification for a 2D motion assumption is no longer valid and the process will quickly lose track of the rover attitude, which in turn will produce further errors in rover position. Another challenge that Wheel Odometry faces is the inability to estimate the slip. Wheel-

soil interaction has been deeply studied in the literature of terramechanics, proposing parametric models that include the longitudinal and lateral slip within them. In those studies, it is common to use Single Wheel Testbed facilities for dynamics characterisation and models validation, similar to the one seen in [Figure 2.2](#). However, the dimensionality of the parameter space is yet too large without yielding accurate or valid solutions for generalisation. Recent works in (Gonzalez and Iagnemma, 2018a,b) have studied different methods to estimate and compensate for slip. These however rely on the input from several sensor data other than just wheel encoders. Achievable performance with wheel odometry data only will therefore highly depend on the terrain itself, the dynamics of the rover and their interaction. Overall, even for a simple, planar case solutions using this method will produce a drift equivalent to 10 % of the traversed distance, quickly deviating even further as the orientation track is lost.





Figure 2.2: Left: [ExoMars](#) wheel being tested in a Single Wheel Testbed. Right: [ExoMars](#) wheel under Locomotion tests (Oettershagen et al., 2019)

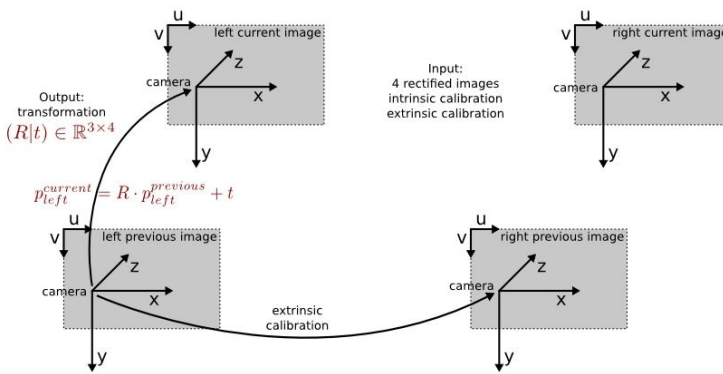
Inertial odometry uses proprioceptive sensors such as accelerometers and gyroscopes, or a full Inertial Measurement Unit (IMU) on its dead reckoning process. Body axis acceleration and angular rate sensor data is basically filtered and integrated over time to estimate the rover speed vector and orientation. However, the slow dynamics of rover systems and noise level in the accelerometer sensor is prone to produce high errors in the estimation of the rover translation. The process is more suitable for estimating the rover attitude with the possibility of using the gravity vector sensed by the accelerometers in static conditions to correct the possible gyro bias drift over time (Trawny and Roumeliotis, 2005). Inertial and Wheel Odometry are commonly combined to provide together a full 3D pose estimation where the IMU is mainly used to estimate the rover attitude and the wheel encoder data is used to translate the body frame according to a defined motion model of the chassis kinematics. Generally speaking, this method is computationally cheap and suitable for terrains with good traction.

using stereo correlation from left-to-right image. The process is repeated in the subsequent image pair and then features from both consecutive stereo pairs are matched together. Images should have sufficient overlap between them, which is dependent on the rover speed and frequency at which the process runs. If sufficient features, or temporal landmarks, are matched and are robustly detected, then the transformation that minimises the error between all matched landmarks is computed (see [Figure 2.4](#)). VO has the advantage of providing a potentially very accurate localisation estimate which basically depends on the accuracy at which the 3D position of the features can be estimated, and their reliability. At the same time, this process is not subject to errors due to slip, since any translation error induced by slip is already detected in the visual tracking process. As such, VO can generally provide performances better than 2% (Núñez, Vázquez-Martín, and Bandera, 2011; Shaw et al., 2013). The main disadvantage of VO is its complexity and the amount of image processing computations that it entails. These have been difficult to embark in space grade processing units, but luckily the slow motion dynamics of the rovers have been helpful to run these algorithms at sufficient frequency.

Finally, approaches to relative localisation often decide to combine the data coming from different odometers through the implementation of Sensor Fusion. The highly accurate but infrequent VO is combined with the less accurate but highly frequent Inertial and Wheel odometry (Bora et al., 2017). They complement each other and serve to detect spurious errors in any of the pipelines by cross checking them, which minimises the drift in dead reckoning in the longer run. The combination of several sources of odometry estimations is also of benefit for the sensor fusion process that generally takes the form of a Bayesian filter. In this case each measurement input serves to reduce the covariance or uncertainty of the process and increase the overall accuracy. Solutions for sensor fusion are typically variations of Kalman filters or Particle filters (Murphy, 1999; Thrun, Burgard, and Fox, 2005). It should be noted that, as initially explained, relative localisation odometry measurements are giving incremental estimations of motion that we could refer to as *delta poses*. It is important to not mistake those measurements as absolute pose estimations of the current time. As such, the measurements are only providing information relative to the latest incremental displacement with respect to the previous pose state. A solution named Stochastic Cloning that copes with this particular conditions is found in the literature (Helmick et al., 2004) and has been more recently applied in (Moreno, Azkarate, and Marchegiani, 2021), which extends the state vector of the filter with a copy of the previous state estimation and expresses the measurement as the difference between the previous and current state variables. Field test



(a) Left-to-right stereo correlated features and up-to-down temporal matches



(b) Frames involved in stereo and temporal images

Figure 2.4: Graphical illustration of visual odometry process using stereo correlated temporal matched features (Núñez, Vázquez-Martín, and Bandera, 2011)

experiments using this method report results of 1.5% for a medium length traverse.

The Table 2.1 summarises the general performances of the different methods for relative localisation described above.

Method	Accuracy	Frequency
Wheel Odometry	10%	10-100Hz
Inertial Odometry	5-10%	10-100Hz
Visual Odometry	1-5%	0.5Hz
Sensor Fusion	1-2%	10Hz

Table 2.1: Summary table with localisation accuracy and frequency for each method

When looking at real planetary rover missions, in the first Mars Pathfinder rover mission of 1997, Sojourner’s localisation was calculated via simple dead

reckoning with wheel encoders and a solid-state gyro (Volpe et al., 2000). However, the approach relied on operators correcting daily its position and orientation by recognizing it on lander imagery (Matthies et al., 1995), and was therefore limited to the exploration on the lander vicinity. Sojourner managed to traverse around 100 m in total during a mission that lasted approximately three months. For the next rover missions, the NASA's MER, the localisation pipeline was significantly improved by the addition of a full IMU sensor and the Sun finding function (Ali et al., 2005), that serves to provide absolute heading corrections for the rover azimuth. But more importantly, it was thanks to the implementation of VO techniques that robust onboard relative localisation was first achieved (Cheng, Maimone, and Matthies, 2005; Maimone, Cheng, and Matthies, 2007; Matthies et al., 2007). Additionally, Jet Propulsion Laboratory (JPL) also developed an Incremental Bundle Adjustment (IBA) method that would be used to process offline on Ground images taken by PanCam, NavCam, and HazCam cameras (Li et al., 2004) (see Figure 2.5). A network of connected images is generated using tie-points detected in overlapping images and an optimisation process generates refined and accurate rover poses at image acquisition locations, improving even the results of VO. While this was significantly useful for the generation of high quality landing site topographic mapping products, it shall be noted that IBA was never run onboard given the heavy computations that it requires in its tie-points generation and optimisation process.

The next generation of NASA rovers, the MSL mission Curiosity rover and the latest Mars2020 (M2020) mission Perseverance rover, have followed very similar approaches for localisation to those just explained, with some improvements in sensor performance and redundancy (Grotzinger et al., 2012). However, due to the extremely time consuming process that the VO algorithm requires, this function has been kept only for particularly challenging terrain where slippage is of significant concern, as it notoriously reduces the traversed distance per sol when compared to other navigation approaches that do not require such precise algorithms to run (Gong, 2015).

In Europe, the first Mars rover mission still to be launched, the Rosalind Franklin rover of the ExoMars 2022 mission, fuses visual and inertial odometry for its relative localisation (Bora et al., 2017; Silva et al., 2013). Following the logic of explanations above, the rover aims to run the VisLoc VO algorithm (Shaw et al., 2013; Townson, Woods, and Carnochan, 2018; Ward et al., 2016) at 0.1 Hz while wheel and inertial odometry are combined to run at 1 Hz. The approach yields accurate enough solutions for the control of a system that roves at an average speed of 1.1 cm s^{-1} . As for the next European rover mission to Mars, the SFR mission in collaboration with NASA for the MSR programme, the design makes full reuse of the ExoMars solution, with almost identical sensors

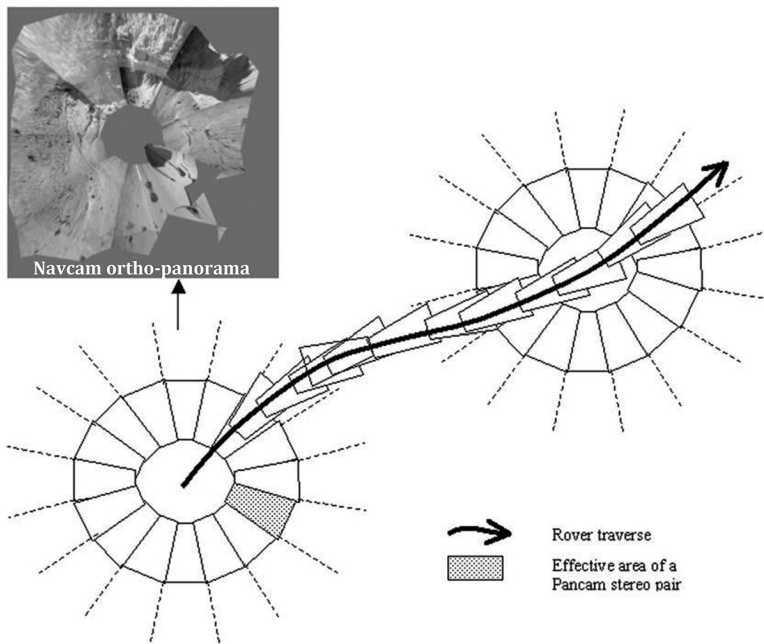


Figure 2.5: Illustration of a rover traverse and the network with PanCam and NavCam images for bundle adjustment (Li et al., 2004).

and software implementation. The only difference lies on the frequency of algorithms execution, that is increased by a factor of five in a rover that will run six times faster. This is possible thanks to the upgrade on the On Board Computer (OBC) capabilities with respect to *ExoMars*, jumping from LEON2 to a LEON4 architecture.

While nowadays we can conclude that VO is a matured technology in planetary rovers, the challenge still remains in reducing the execution time of it. The image processing of stereo pairs and 3D features correspondence matching are computationally heavy processes that can take up to seconds to complete when run on the currently available space-grade processors (Lentaris et al., 2018). An alternative to these constraints is found in HW accelerated solutions (Kostavelis et al., 2011; Malinowski, 2019), that use FPGAs to implement the stereo processing algorithms with parallel execution, allowing many repeated calculations to be performed concurrently and reduce significantly the execution time. Interestingly, the SPARING Robotics Technologies for Autonomous Navigation (SPARTAN) (Avilés et al., 2018; Kostavelis et al., 2014; Lentaris et al., 2015) development has been considered for SFR as a HW accelerated VO solution, developing even a breadboard and performing a test campaign. Eventually, the

VisLoc solution of [ExoMars](#) has been kept also for this mission due to the overall system trade-off and design reuse.

2.1.2 Global Localisation

Global localisation methods are implemented to correct the inherent drift in relative localisation solutions that grows together with the traversed distance. It should be noted that drift corrections by means of loop closure identification and graph optimisation methods are not addressed here. This is a technique hardly usable in planetary rovers, since they do not frequently revisit places throughout their traverse. Additionally, these solutions typically require computational and memory resources that would be difficult to fit in space grade processors.

Global localisation solutions were first investigated to determine the landing point location of planetary missions. Early approaches used a Ultra High Frequency ([UHF](#)) two-way Doppler tracking system. Mars Pathfinder exploited this technique using its Direct To Earth ([DTE](#)) communications antenna. With the arrival of the Odyssey Orbiter to Mars in 2001, [MER](#) could improve the accuracy achieved by this technique up to 50 m thanks to the rover-to-orbit geometry changing (Guinn, 2001) (see [Figure 2.6](#)). Mars orbiting satellites eventually became the main source of reference data with which global localisation techniques have been developed. In general, global localisation methods aim to triangulate or co-register telemetry data with identified landmarks that appear in both rover and orbital imagery. The most relevant orbital reference data to mention here would be that provided by High Resolution Imaging Science Experiment ([HiRISE](#)) (McEwen et al., 2007). Over time, techniques have improved with higher landmark identification accuracy for both orbital and rover data products and the landing location of missions has been refined and corrected, as can be seen in (Tao, Muller, and Poole, 2016).

As [MER](#) missions started to traverse further away from their landing location, global techniques to correct the rover localisation drift along the mission were necessary for the continuation of their scientific missions. Initially, [MER](#) mission teams exploited the [IBA](#) method introduced in the previous section to keep track of the rover localisation as it roved accurately on Ground (see [Figure 2.7](#)). However, even this incremental method, that depends heavily on the tie-points that are used to connect the rover images, will also build up drift in the range of 2% of the traversed distance.

An early global method presented in the '90s proposed to use skyline signatures captured by a rover panoramic picture, and match it to predicted skyline signatures at locations of the global map (see [Figure 2.8](#)). The Visual Position Estimator for Rovers ([VIPER](#)) system (Cozman, Krotkov, and Guestrin, 2000)

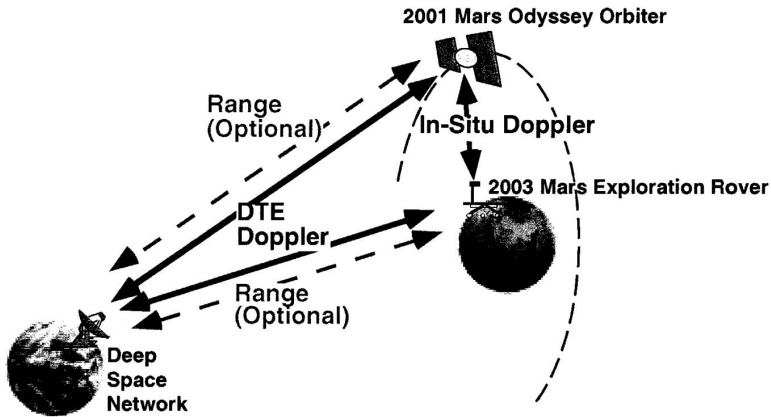


Figure 2.6: Representation of the global localisation using two-way Doppler with Mars Odyssey (Guinn, 2001).




Figure 2.7: NavCam mosaic from Opportunity in Meridiani Planum. Opportunity landed in the 20m diameter Eagle Crater in the middle of this mosaic (Matthies et al., 2007).

is probably the most well-known example of this technique, which yields a localisation solution with an accuracy in the order of 100 m, and more recently exploited in (Chiodini et al., 2017) with MER image data to achieve slightly better results.

It was later with the introduction of feature-based solutions that the accuracy of global techniques was significantly improved up to few meters. Techniques vary depending on the type of features they use to correlate with orbital maps. These can be elevation features, appearance or intensity features, or otherwise a network of identified landmarks or regions of interest.

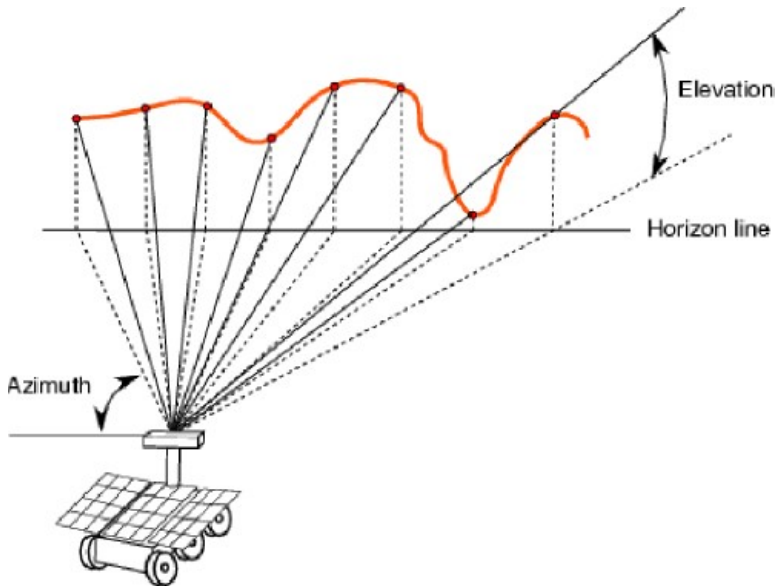


Figure 2.8: Skyline matching method representation used in [VIPER](#) algorithm (Chiodini et al., 2017).

The first one uses a Digital Elevation Model ([DEM](#)) built by the navigation cameras of the rover and compare it to the elevation maps generated from the scientific cameras in orbiting satellites (Carle and Barfoot, 2010; Carle, Furgale, and Barfoot, 2010; Van Pham, Maligo, and Lacroix, 2013). The main drawback of this approach is the relatively low resolution of [DEMs](#) built from orbital data, which requires the generation of a notably large-size 3D map on-board the rover to compensate (see [Figure 2.9](#)). In (Furgale, Carle, and Barfoot, 2010), the Multiple-frame Odometry-compensated Global Alignment ([MOGA](#)) algorithm results demonstrate to outperform those of [VIPER](#). The second one uses an Ortho-Rectified Image ([ORI](#)) instead of a [DEM](#), which is a 2D projection on the ground plane of the intensity data values of a 3D reconstructed map (Sheshadri et al., 2012; Tao, Muller, and Poole, 2016). This method has the advantage that it can potentially work in flat areas without significant elevation features as long as the terrain shows variations in albedo. Additionally, the raw orbital imagery is directly an [ORI](#)-type of image, which means that it does not require stereo processing to produce the reference maps and therefore these are available in the highest possible resolution. In the third category, we find methods that first extract landmarks from the orbital reference maps and build a global network of points of interest with them. Then, similar local network is built as the rover traverses. As soon as enough rocks or landmarks are introduced in the local network their distribution pattern can be matched between both networks

(Boukas, Gasteratos, and Visentin, 2018; Hwangbo, Di, and Li, 2009). The benefits of this technique are that points of interest such as rocks or outcrops can be identified easily from orbital data without having to distinguish them uniquely. Two similarly sized rock enhance the network of landmarks placed at two different locations. The correspondence between local and orbital networks is done by landmark location with some rough characteristics of size and type without having to match specific features of elevation or intensity which could be differently sensed by the two sources of data. The main drawback is the large distances that need to be traversed while registering the local landmarks network before any solution can be computed, with the corresponding errors introduced in the network by the drift of the relative localisation. Finally, approaches using Machine Learning techniques to solve the global localisation problem are also being studied in latest works (Moreno et al., 2021; Naguib et al., 2018; Wu et al., 2019).




Figure 2.9: Reconstructed Opportunity wide baseline NavCam DEM and HiRISE DEM at Victoria Crater in 3D showing possibilities of registration (Tao, Muller, and Poole, 2016).

Nowadays, most promising results in real missions are obtained with the ORI feature matching techniques. The arrival of the Mars Reconnaissance Orbiter (MRO) and the provision HiRISE imagery meant that reference maps of up to 0.25 cm/pixel can be obtained compared to the 1 m/pixel of DEMs (see Figure 2.10). Some works have tried to improve the resolution of reference maps by further processing of orbital data using shape-from-shadowing to generate DEMs of up to 0.25 cm/pixel (Gupta et al., 2014) and super resolution restoration with multi-pass images and subpixel information to resolve a 5 cm resolution of the area (Tao and Muller, 2016).

It is worth mentioning that none of these techniques has been embarked in any rover system so far and they have only been applied on the Ground




Figure 2.10: Example of co-registered wide baseline NavCam ORI and HiRISE ORI (left) and corrected traverse (right) for Curiosity (Tao, Muller, and Poole, 2016).

Control Station (GCS) on Earth. They have served to correct the pose estimation of the rover and update it via telecommands, but no rover had these techniques implemented onboard. No mission so far saw this need because the traverse distance per sol did not generate enough drift that meant a risk for the continuation of the mission and the completion of its objectives. Even nowadays, Curiosity's global position is corrected regularly, even daily, on Ground by operators manually matching to the orbital maps an ORI made out of the few stereo images taken by the rover at its end-of-sol location. This paradigm has changed now with the upcoming SFR mission where the tight mission timeline is imposing considerably higher traverse speeds and distance per sol. In order to avoid any dependencies on Ground pose corrections that could keep the rover temporarily stationary and jeopardise the timeline, a global localisation solution has been included in the rover baseline design. While its development is still in early breadboarding phase, the approach is to be based on a combined DEM and ORI matching technique. Additionally, the global localisation shall perform in two different scenarios, first during the main mission traverse and second at the Sample Depot operations where Perseverance will have dropped its sample tubes. In the former scenario, reference maps will be taken from HiRISE and loaded onboard the rover. In the latter, the reference maps will be generated from the M2020 rover camera images. Thanks to the higher resolution

of these reference maps, it is believed that absolute pose estimations within the Depot maps can be obtained with an accuracy better than 10 cm. Latest works by NASA-JPL in (Pham et al., 2021) show the results in a representative field test campaign where their in-depot global localisation technique based on Virtual Template Synthesis and Matching approach is used (Figure 2.11).




Figure 2.11: Points on the mesh built from M2020 images are rendered from multiple virtual viewpoints (left). Synthetic templates are matched to real SFR images to compute points in camera frame (middle). Both point sets are aligned to estimate the SFR pose (right, conceptual rover depicted) (Pham et al., 2021).

Finally, another source of in-depot localisation that SFR might be able to exploit is to co-register its position based on the relative localisation against detected sample tubes. As sample tubes will be very accurately localised within the depot maps, these can serve as “geo-referenced” landmarks. Besides, the relative position between SFR and the samples will need to be accurately detected in order to have successful autonomous grasping operations. In (Dafttry et al., 2021), NASA-JPL presents the machine vision techniques used for the detection are relative pose estimation of tubes.

2.2 PERCEPTION

Perception is the term commonly used in robotics for the mapping functionality or in other words the modelling of the environment. Maps can be of different types depending on the data they encode, but in general we refer to a 3D representation of the terrain, typically a point cloud or a DEM once the point cloud is structured. The terrain model is then further processed on the perception pipeline in order to detect obstacles or hazards and extract from this the terrain’s traversability information. This is usually expressed in the shape of a 2D grid noting the presence of hazards on each cell that represents a piece of terrain of a certain resolution. The traversability grids are commonly the input data for the next main functionality in the navigation task, i.e. the path planning. This grid is translated into a cost map with values giving the associated risk or penalty for traversing a cell.

In order to model the terrain in front, the rover needs to embark some kind of exteroceptive sensor. One of the most common sensors would be the

cameras, in particular, optical stereo benches. Examples of perception functions using sensors such as Time of Flight (ToF) cameras or Light Detection And Ranging (LiDAR)s are common in terrestrial applications these days, given their capability to directly provide accurate 3D data points. Because of this, spinning mirror LiDARs became very popular in first self-driving car solutions. Few examples of LiDAR sensors onboard satellite missions exist, that were used for relative navigation in docking maneuvers to the International Space Station (ISS) (Christian and Cryan, 2013). However, these are not (yet) suitable for space rovers due to their high mass and power consumption their current technology requires. In that regard, they are active sensors, i.e. they emit light, as opposed to optical cameras which are passive. So far, no planetary rover has been equipped with a sensor of this kind and no rover mission seems to be keen on qualifying such a sensor for higher TRL. A difficult challenge in that qualification process is actually the vibration and stress of the launch phase that such sensor is meant to withstand. Currently emerging technologies for flash and solid state LiDARs could potentially bring a new era of sensors for space, yet it is difficult to foresee when this could happen. On the contrary, optical cameras have been long qualified for space and all planetary missions embark several of these sensors (Maki et al., 2012; Maki et al., 2020; Miles, Gunn, and Coates, 2020) (see Figure 2.12 & Figure 2.13). Stereo cameras permit to reconstruct the environment by correlating pixels from left-to-right images thanks to accurate calibration processes. An stereo rectified image pair is accurately calibrated when a pixel of the left camera can find its corresponding pixel in the right camera at the same image row, corresponding to the epipolar line. This allows to estimate the (disparity and then) depth of each pixel in the camera frame and eventually obtain a point cloud.

Similar to the case of VO, the process of estimating the depth of every pixel for 3D mapping the environment requires heavy computational resources that are hard to embed into the processing modules of space avionics OBCs. Solutions to solve the estimation of traversability of the path ahead without needing to generate a terrain DEM have been studied in the literature (Mandelbaum et al., 1998; Williamson and Thorpe, 1998). These fit potential ground planes to the disparities provided by a stereo camera and hazards are detected by measuring deviations from the average height in the fitted ground plane. While this is a potentially more efficient and faster computing solution, it becomes less suitable when applied to space applications due to the difficulty of fitting a ground plane to uneven and unstructured planetary surfaces.

Looking at actual rover missions, the early Sojourner had limited embarked navigation functionalities. It was with the arrival of MERs that a rover first implemented mapping capabilities (Goldberg, Maimone, and Matthies, 2002).




Figure 2.12: *M2020* flight NavCams (Maki et al., 2020).




Figure 2.13: *ExoMars* camera sensors on top of the mast (Miles, Gunn, and Coates, 2020).

However, the amount of time that the process would take, in the order of minutes, each time the mapping and traversability analysis functions were used, significantly reduced the net traverse speed of the rover. In *MER* these functions were part of the Grid-based Estimation of Surface Traversability Applied to Local Terrain (*GESTALT*) (see Figure 2.14) system (Biesiadecki and Maimone, 2006; Maimone, Leger, and Biesiadecki, 2007) and due to the time it took to run them, *MER* operations always tried to prioritise the use of the *Blind Mode* of navigation. In this mode rover drivers on Ground provided a safe path for

the rover to follow in the form of a sequence of waypoints, without needing to stop every few meters for the mapping cycle. NASA Curiosity and Perseverance rovers have profited from faster computing modules to reduce the execution times of these functions and increase the use of more autonomous navigating modes (Maimone, 2017; Rankin et al., 2020, 2021).




Figure 2.14: Virtual representation of the **GESTALT** system for terrain assessment (Maimone, Leger, and Biesiadecki, 2007).

For the upcoming **SFR** mission a solution based on the HW accelerated **SPARTAN** algorithms was considered and bread boarded (Kostavelis et al., 2016; Lentaris et al., 2019). Despite the improved execution time of this solution, eventually the design approach of **SFR** will be based on the re-use of the **ExoMars** mission heritage algorithms implementation, similarly to the decision taken for the **VO** implementation. Interestingly, the **ExoMars** rover mission has redundant flight implementations of its perception. It contains two separate SW images to choose from, one developed by Airbus Defense and Space (**ADS**) (McManamon, Lancaster, and Silva, 2013) and another one by French Space Agency (**CNES**) (Moreno, 2013). Both implementations take a similar general approach for **DEM** generation through stereo disparity calculation and traversability analysis by running different safety checks over the elevation map grid. Finally, for the **SFR** mission, it has been decided to merge these two solutions into a single flight implementation.

Out of space missions, Machine Learning techniques are more and more employed in order to directly classify the terrain in safely traversable or hazardous areas (Andrakhanov and Stuchkov, 2017). A more interesting case is found in latest works in (Dimastrogiovanni, Cordes, and Reina, 2020) in which the

camera data of the terrain ahead is combined with other proprioceptive sensor data when that terrain is being traversed. This leads to increase the overall knowledge and learn the tractive characteristics of such terrain to eventually be able to predict those before actually traversing it. However, these techniques for the terrain classification and obstacle detection are still far from being taken to space implementations, since the lack of non-determinism of machine learning technology makes it difficult to be adopted in space missions.

2.3 SIMULTANEOUS LOCALISATION AND MAPPING (SLAM)

SLAM has been a very popular research field in the mobile robotics community over the last 30 years. SLAM solutions aim at interlacing and solving together these two crucial functions in robotics. The idea of computing both localisation and mapping simultaneously is inherent to the problem statement where a mobile robot is trying to localise itself and traverse in an a priori unknown terrain. The localisation needs to be computed with respect to a map that is not given but rather built as the platform moves. At the same time, for the map to be built the localisation from where the latest perception of the environment was taken needs to be computed accordingly. Therefore, instead of treating these as independent processes SLAM tries to combine the data process of both functions to come up with an optimised solution that improves the accuracy of both functions. Therefore, the key difference of SLAM is on the combination and optimisation aspects, opposed to the more classic approach of solving each problem separately.

SLAM has been extensively investigated in the research community with many applications for robotics on Earth. Extensive surveys of the different approaches to SLAM exist in the literature (Bailey and Durrant-Whyte, 2006; Cadena et al., 2016; Dissanayake et al., 2011; Durrant-Whyte and Bailey, 2006; Stachniss, Leonard, and Thrun, 2016; Thrun, Burgard, and Fox, 2005), providing a comprehensive description of the SLAM categories and most popular implementations, that go from probabilistic methods using Particle Filters or Extended Kalman Filters (EKF), or scan matching methods, to more recent graph optimisation techniques with bundle adjustment in the presence of loop closures. The aim of this section in this thesis is not to reproduce (even partially) the work on those surveys but instead to elaborate on the applicability of SLAM techniques to the planetary rovers scenario and conditions.

SLAM is known to be a computationally greedy process. This fact has highly hindered its adoption in space applications due to the difficulties to run it in space processing modules that lack the significant computational resources. Instead, autonomously navigating planetary rovers have opted to separately

compute the localisation of the rover and build the map of the environment. In fact, not only these are independent processes, but they even also happen to be not computed in a concurrent or parallel processing approach. These are typically run sequentially, with the relative localisation function running while driving and the mapping only happening once the rover is stopped, similar to the navigation concept explained in (Correal and Pajares, 2011).

Actually, in the history of planetary rover missions there has been no rover that embarked a *SLAM* implementation so far. However, with the advances and qualification of more powerful space processors with multi-core architectures (Furano et al., 2020; Lentaris et al., 2018, 2019), the possibilities to embark a *SLAM* function in space rovers have increased. Yet, not all *SLAM* approaches are valid for an *active* and *online* *SLAM* solution. For example, Visual *SLAM* solutions may track features or landmarks in the scene to update the pose estimate, similar to *VO*. But, the map of sparse landmarks does not necessarily produce a sufficiently dense model of the environment to be able to evaluate its traversability. Therefore, these do not provide a valid approach to rover navigation on their own. In (Shaukat, 2015), three potential *GNC* architectures for future rover missions are proposed which include a *SLAM* filter on them. While the reference scenario used in that article is the *SFR* mission, this mission has finally discarded the implementation of this function (see Section 2.6 for more details on the navigation capabilities of *SFR*). Works in (Yeomans et al., 2017) propose a Teach-and-Repeat framework that creates a graph of localised landmarks at the teaching phase and is then able to autonomously reproduce this traverse with certain robustness to changing conditions. Interestingly, this approach could be of use for *SFR*'s return traverse where the rover could undo the path back to the lander directly avoiding the hazards encountered on the way towards the sample depot. Finally, it is worth mentioning the work in (Hidalgo-Carrió, Poulakis, and Kirchner, 2018), that demonstrates how the computational load of a Graph *SLAM* method can be effectively reduced by adapting the selection of *keyframes*, based on the inertial odometry error, to minimise the amount of these needed for the *SLAM* solution to converge in the back-end (see Figure 2.15). That work particularly targeted the planetary rover case achieving a reduction of up to 75% in the number of processed *keyframes* with a minimum loss in the solution accuracy.

2.4 PATH PLANNING

Path planning is the functionality that finds a safely traversable path for the rover to follow that shall take it from its current position to a given target location. The input to this module is typically the traversability map, or more

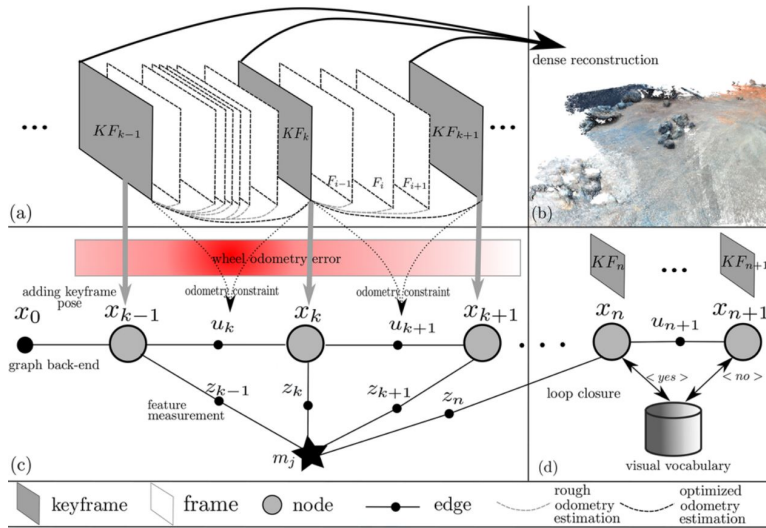


Figure 2.15: Representation of keyframe selection approach in (Hidalgo-Carrió, Poulakis, and Kirchner, 2018) for its Adaptive Graph SLAM solution.

directly its cost map, generated at the end of the perception task, and the output is usually a sequence of waypoints the rover shall follow as accurately as possible. This is an optimisation problem generally solved by analysing the possible paths that connect the starting and end positions and evaluating them against a given metric or criteria that is used to assign a resulting cost to each evaluated path. The path with lowest total cost is the one finally chosen. This path is defined as *optimal* if it actually minimises the cost function among all possible paths that take the rover from its initial state to the final target location. Like in many, if not every, optimisation problem, a good solver needs to balance between finding the global minima, without getting stuck in a local one, and the algorithm’s execution time, making sure it does not keep looking for an even better nonexistent solution. As such, a planning algorithm is defined as *complete* if it will always find a path in finite time when one exists, and will let us know in finite time if no path exists. Similarly, a planning algorithm is *optimal* if it will always find a globally optimal path without getting stuck in a local minima (Ferguson, Likhachev, and Stentz, 2005). Despite the uneven, unstructured nature of planetary terrain, wheeled rover path planning is commonly reduced to a planar problem using a 2D grid to represent the traversability map input, as already mentioned in the previous section. Despite this simplification, the altitude information is not fully obviated, and slopes data can be encoded within the cost associated to traverse a certain map location.

Classic graph search algorithms have been used to implement grid-based path planning solvers (Dijkstra, 1959; Hart, Nilsson, and Raphael, 1968) due to the similar ways of how these problems can be defined and approached. A* can be probably considered as the basic algorithm from which the principles of mobile robot path planning have originally developed (Nilsson, 1980). This is an optimal algorithm that uses an heuristic function to quickly and roughly estimate the cost of a potential path to the goal in order to guide the search towards the most promising paths, which brings major savings in computation time. Based on A*, other algorithms were developed that added interesting capabilities to it. Dynamic A* (D*) (Stentz and Hebert, 1995) and variations of it (Ferguson and Stentz, 2005b; Koenig and Likhachev, 2002b) permit to efficiently repair a previously planned path when new sensory information reveals changes or more details along the path. This way it can avoid re-running the algorithm over the entire graph each time repairing is needed to obtain a new optimal path. An alternative approach for scenarios where a robot might need to react quickly and the execution time may be bound are referred to as *anytime* algorithms (Likhachev et al., 2008). These typically construct an initial, possibly highly suboptimal, solution very quickly, then improve the quality of this solution while time permits. This is achieved by inflating the heuristic values used in A* (weighted A*) which provides substantial speed-ups at the cost of solution optimality. The work in (Ferguson, Likhachev, and Stentz, 2005) is a comprehensive guide where all these algorithms are thoroughly described and their applicable cases explained.

An interesting improvement to the dynamically replanning solutions was brought by the Field D* algorithm (Ferguson and Stentz, 2005a, 2006) that added the capability to produce completely smooth paths without being restricted to waypoints connecting cell-centric positions, a.k.a. *any-angle* algorithms (see Figure 2.16). This capability can be more or less relevant depending on the grid cell size or grid resolution over which the algorithms run. Using high resolution grids smoother paths can be produced even without using an *any-angle* path planner. But high-resolution grids also mean a higher computational cost due to the increased number of states or nodes to evaluate. Another solution with *any-angle* planning capabilities and with improved computational cost is found in the Fast Marching Method (FMM) (Sethian, 1999). It uses the so-called *eikonal* equation to define the propagation of a wave over a potential field that represents the arrival time of the wave to each node in the grid. The applicability of this method to planetary exploration case has been already studied in (Garrido, Álvarez, and Moreno, 2016).

A completely different branch of algorithms worth mentioning are those based on probabilistic methods, opposed to the deterministic algorithms where




Figure 2.16: Grid-based path planning from s to g . *Any-angle* algorithm solution (in blue) finds smooth paths without being restricted in the direction of connecting nodes (Ferguson and Stentz, 2005a) compared to the more classic angle-constrained path planner solution (in black).

all previous ones are categorised. When the dimensionality of the problem is increased these methods provide valid solutions in optimised time at the expense of finding suboptimal solutions (see Figure 2.17). This is relevant for example when planning the motion of kinematic chains with a high number of degrees of freedom. Previous deterministic approaches would take significant amount of time to explore the state space of such high dimensional problems. A comprehensive study of randomised or sampling-based algorithms can be found in (LaValle, 2006) and later on in (Elbanhawi and Simic, 2014). Other relatively more modern methods based on neural network, fuzzy logic and nature-inspired algorithms are extensively explained in the survey (Mac et al., 2016). However, rover platform navigation is a relatively low dimensional problem, which is further simplified by the approach of the path planning function that solves this as a 2D planar motion case. For this reason, probabilistic methods are commonly not used for the rover navigation use case.

Looking at the planetary rovers scenario, the path planning problem is usually divided first in two stages, the global planning and the local planning. The global planning is performed using orbital maps data and processing it to generate traversability grids of relatively low resolution. The global planning happens typically on Ground and produces a rough path with somewhat separated waypoints for the rover to follow. This global path is the one that guides the rover main mission objective. As the rover starts its traverse, and the perception modules provide more detailed, high resolution traversability maps, the rover might need to plan at local scale a path that avoids any obstacles that it finds on the way while still moving forward along the global path. The local planner is

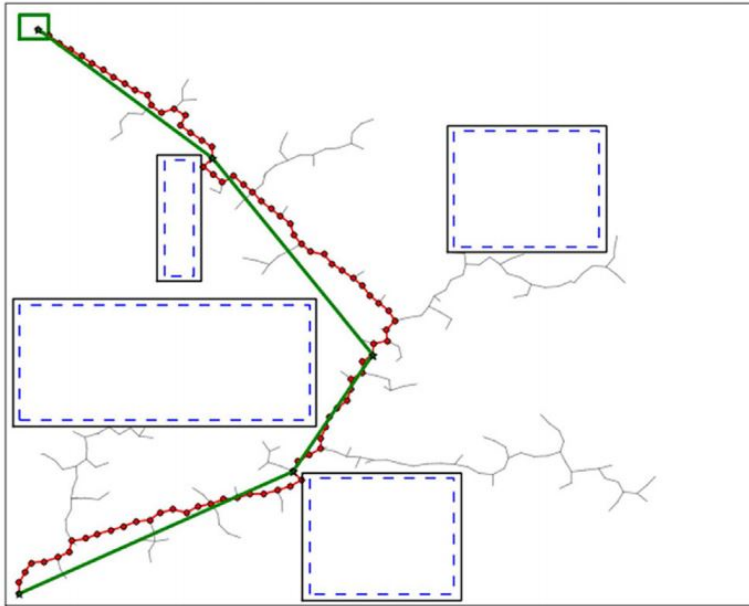


Figure 2.17: Typical suboptimal path obtained from randomised planning (red path) and the corresponding shortcut path (green) (Elbanhawi and Simic, 2014).

actually in charge of this (re)planning task. Of course, the local planner is only necessary for those navigation modes where the rover is not “blindly” following the path provided by Ground.

While Sojourner was said to include autonomous capabilities onboard (Mishkin et al., 1998), its daily operations were highly guided by humans on Ground. In reality, we could say MER were the first rovers to implement onboard path planning. The original local path planner integrated in the GESTALT system (see Section 2.2) used a set of potential arc motions evaluated under a voting system that would rank the arcs according to traversability, steering-time and goal reaching criteria. This system worked well enough to keep the rover safe and reach the goal while avoiding small hazards. However, a relatively big-sized hazard or non-traversable area in the rover’s field of view would lead the planner to get stuck without finding a possible way out. This unfortunate situation meant the rover had to wait for an alternative path surrounding the hazard to be uploaded from Ground with the consequent loss of effective mission time. GESTALT was later on upgraded to include the Field D* algorithm as global path planner (Carsten et al., 2007, 2009) and its potential field was also used at the local level arcs evaluation. That allowed the rover to find possible ways out autonomously even when larger hazards were encountered. Figure 2.18 shows an example of the resulting path obtained with Field D* on MER.




Figure 2.18: Mosaic looking backward on the tracks of a Field D* checkout on MER (Carsten et al., 2009).

As MSL progressed on its exploration, rocky fields created significant wear on its wheels, highly degrading its performance (Graser et al., 2020). This led to an additional enhancement of its onboard path planning, based on a finer terrain obstacles classification and wheel-placement planning (Ono et al., 2015), allowing the rover to traverse rocky terrain minimising the wheel deterioration (see Figure 2.19). All in all, Field D* is still nowadays the core path planning algorithm onboard NASA's Curiosity and Perseverance rovers (Rankin et al., 2021).




Figure 2.19: Demonstration of how wheel placement planning can obtain a path solution in the presence of rocks with specific paths for each wheel (Ono et al., 2015).

The solution adopted for the European ExoMars rover is originally based on the A* algorithm for the local path planner (Bora et al., 2017). This decision is based

on the fact that at each planning cycle the local search graph is significantly different to the previous one, and therefore a new optimal local path from scratch is computed. Dynamically replanning algorithms have sometimes led to increased execution time when the newly acquired data resulted in a significant update of the map (Ferguson, Likhachev, and Stentz, 2005), which is the case for the scale and range at which the local planner reasons. The *ExoMars* planner also considers the kino-dynamic constraints of the non-holonomic platform, and avoids paths with high curvatures that would require point turn commands to be executed (see Figure 2.20). Ideally, the solution is composed of a series of arcs that the rover can fulfil with simple Ackermann motion commands. The series of estimated Ackermann motions are included in the path planning solution to be used later on at the trajectory control module (see Section 2.5).

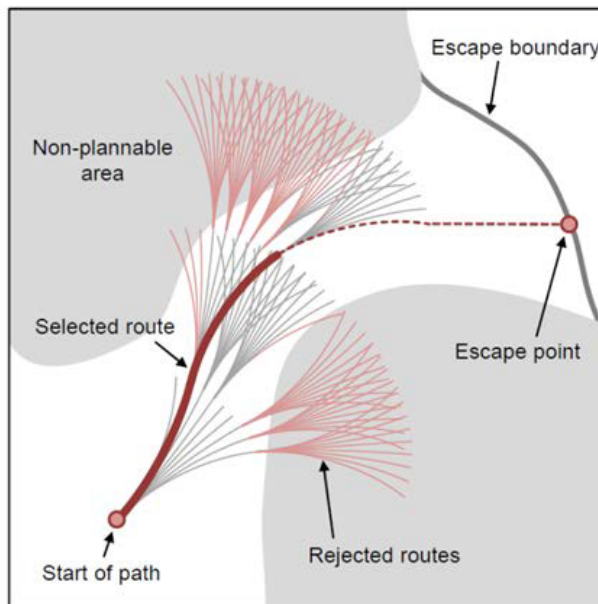


Figure 2.20: *ExoMars* path planning with kino-dynamic constraints principle (Bora et al., 2017).

On the other hand, the rover relies on a global path computed on Ground by operators to guide it towards its overall mission goal, and use the onboard perception and path planning only in cases where the operators judgement of the path safety is hindered either by the limited resolution of the orbital map or by the range and visibility of the end-of-sol rover camera local perceptions. CNES proposed some improvements to A* in (Rusu et al., 2013), based on the algorithm named Fringe Retrieving A* (Sun, Yeoh, and Koenig, 2009), which transforms the previous search tree to the initial search tree for the current A*

iteration. This approach is executing more informed A* searches with initial information instead of running from scratch, and also avoids memory and time consumption for the cells in the search tree which can be reused for the current path planning. Additionally, the same paper proposes the inclusion onboard of global planning capabilities using the D* Lite algorithm (Koenig and Likhachev, 2002a). However, due to limitations on computing performance these improvements have not been finally integrated in the flight software of CNES's navigation solution for ExoMars. For the SFR mission, the path planning task will be implemented by ADS and the current baseline does not foresee major modifications with regards to the ExoMars algorithms.

2.5 TRAJECTORY CONTROL

Trajectory control is the module in charge of commanding the platform in order to follow as close as possible the sequence of waypoints that come out as solution of the path planning problem. It shall steer the rover along the path and apply correcting manoeuvres whenever the rover deviates from the nominal safe path. This function uses the pose estimation updates from the localisation module at each control cycle to calculate the error with respect to the nominal path, which usually comprises a lateral deviation error and a direction or heading error. The subsequent command is computed as a function of the current position in the path and its deviation, in order to reduce it while still progressing towards the next waypoint in the path. Feedback and feedforward control mechanisms are used for this purpose, with a tradeoff between control effort and control error. Even a performing trajectory control function will experience some residual control error. Contributing to the control error are the localisation noise and drift, and any slip that results from the terramechanics that govern the rover motion dynamics at the given terrain of traverse. This is why a *safety corridor* needs to be defined, corresponding to the maximum permissible control error. The width of this corridor, or lateral deviation, is in fact taken into account in the path planning function to ensure that enough margin is taken with respect to the detected hazards. Usually, the module runs in parallel to the rover motion, without needing it to stop. The execution frequency of this control cycle is usually in the range of 1 Hz to 10 Hz, due to the slow dynamics and speeds of planetary rovers traverse, moving only at few centimeters per second.

In the literature, path tracking controllers are known to carry out the trajectory following function. In (Snider, 2009), a comprehensive survey is presented, comparing the most popular solutions and their main characteristics and applicability. Generally, trajectory control algorithms are divided in two major classes,

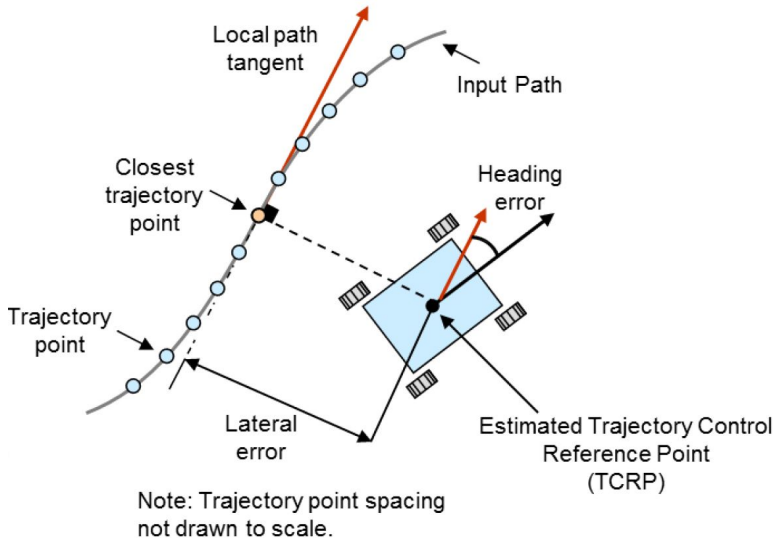


Figure 2.21: Illustration of trajectory control concept (Bora et al., 2017).

the ones that are based on classic control theory and the ones that are based on geometric relations. In the first group, a model of the plant is defined using kinematic or dynamic equations that describe the resulting motion with respect to a given input command. The plant can be then analysed to apply many of the classic control theory approaches such as non-linear control (De Luca, Oriolo, and Samson, 1998), model predictive control (Ollero and Amidi, 1991) or optimal control (Peng and Tomizuka, 1991; Rajamani, 2011). While these can be robust solutions for a wide variety of scenarios, their control laws are usually more complex and gain parameters less intuitively tuned. In recent works in (Pitkänen et al., 2017) a control law using the Instantaneous Center of Rotation (ICR) approach to control a platform with independently reconfigurable, steerable and actuated wheels. In the second group, the control command is computed as the result of geometric relations and constraints. The benefit of these approaches is that they are easily explainable and their parameters tuned according to geometric dimensions. These controllers work well particularly in slow driving scenarios where dynamics do not take much effect. The most well-known algorithm in this category is *Pure Pursuit* (Coulter, 1992). This popular controller and similar ones based on it with some variations have been used in many mobile robotics applications (Amidi and Thorpe, 1991), including successful cars participating in the popular DARPA Grand Challenge (Thrun et al., 2006; Urmson et al., 2006) and later on in the Urban Challenge (Patz et al., 2008; Urmson et al., 2008). Works in (Samuel, Hussein, and Mohamad, 2016) review the most relevant geometry based algorithms that came after the original

Pure Pursuit. More recently in (Chen et al., 2018) a solution that mixes the *Pure Pursuit* path tracker with a proportional-integral controller is presented aiming to get the best of both control theory and geometry based class controllers.




Figure 2.22: Conceptual representation *pure pursuit* working principle as implemented by NASA on MER (Helmick et al., 2006). The carrot coordinates represent the instantaneous waypoint along the path towards which the rover is steered.

For the planetary missions scenario, both types of path trackers could be potentially used. Given the slow speeds of the rovers, geometry based as well as controllers using simple kinematic models could well perform. First NASA rovers opted for the *pure pursuit* approach, that was renamed as *carrot heading* (see Figure 2.22) but implemented equally. With the experience of MER and the frequent occasions in which rovers experienced high slippage, and extension to the controller was implemented that would add corrective measures when the estimated slip vector was non-zero (Helmick et al., 2006). This was done to address one of the main drawbacks of this geometry based trajectory controller, since the original *Pure Pursuit* does not have a proper mechanism to fight the steady state error. Later on, and due to the wheel damage experienced by Curiosity during its mission, an additional speed control layer was added to the trajectory control. This aimed to reduce the torque exerted by each wheel by individually adjusting the speed command using a body kinematic model that estimated the wheels' contact points (Toupet et al., 2020).

For *ExoMars*, the controller is based on the platform's kinematic model with combined feedback and feedforward commands, where the feedforward com-

mand is the one calculated at the path planning algorithm (see [Section 2.4](#)) (Bora et al., 2017; Silva et al., 2013). The estimated slip vector is used for correcting the propagation of the rover position which indirectly influences the eventual control command. For the [SFR](#) rover, despite the increased traverse speeds, the baseline design assumes the same controller will be still valid and perform well for the mission.

2.6 AUTONOMOUS NAVIGATION

All the above studied functions or modules, when implemented together, render a rover capable of navigating autonomously. Autonomous Navigation, often abbreviated as *AutoNav*, is the capability or navigation mode of a rover that relies entirely on the onboard implementation and orchestrated execution of these functionalities in order to guide the traverse with little-to-none human guidance. Ground operators may still provide the global targets that the rover shall follow to reach its next scientific target, based on the analysis of the low resolution orbital maps. Yet, the rover shall be capable of guiding itself at the high resolution local scale.

It is concluded that for a rover to navigate autonomously many software components need to run onboard as part of its [GNC](#) architecture. However, the challenges and costs of bringing these algorithms up to space systems have made missions to adopt little by little the capabilities that build up towards autonomous navigation. As a consequence, for many years, [MER](#) were mainly operated in the so-called *Blind mode*, releasing the rover from the need of running most of these algorithms (Biesiadecki, Leger, and Maimone, 2007). As missions evolved so did the rover onboard capabilities and [MER](#) have eventually been driving significant distances autonomously over the last years (Maimone, 2017). This was mainly thanks to the implementation of the [GESTALT](#) system for terrain assessment and other autonomous functions onboard the rovers (Biesiadecki and Maimone, 2006; Maimone, Leger, and Biesiadecki, 2007). These articles give a detailed overview of the implemented functions in the [MER](#) flight software and how they were combined in order to achieve most efficient autonomous traverses in different environmental conditions. An interesting function added to the extended mission of [MER](#) was the *Visual Target Tracking* which helped dealing with target location uncertainties and rover localisation drift during the last part of approach to a target. The main difference between [MER](#) and [MSL](#) missions was not the addition of more autonomous navigating capabilities, but rather the frequency or usage made of out of them (Rankin et al., 2020, 2021). Curiosity included in its baseline all navigating functions that were added to Spirit and Opportunity throughout their mission. Together with the more

capable computing processor onboard, Curiosity has used, for the most part of its traverse, navigation modes other than the operator-guided blind one. Unfortunately, due to the heavy wear on wheels experienced by Curiosity over the last years (Graser et al., 2020), operations nowadays are again avoiding the use of *AutoNav*, limiting the daily traverses to distances that can be safely and accurately planned by operators on Ground (shorter than 100 m). This is done to avoid even the relatively small rocks that *AutoNav* would simply neglect due to not being bigger than a safe threshold to consider them as potential hazards. As a consequence, *VO* is now almost constantly run in Curiosity to make sure the rover does not deviate significantly from the operators' pre-planned path. As an additional autonomy feature, although not related to navigation, Curiosity has implemented autonomous science capabilities that allows it to detect dust devils or clouds over the images taken during the traverse, automatically selecting images known to be interesting for later downlink (Castano et al., 2008). Finally, looking at the *M2020* Perseverance rover, and despite its carbon copy of *MSL* external appearance, it should be noted the significant increase in computing resources of the latest rover. Not only it embarks a more powerful processor, but it also integrates a co-processor unit in order to run many of the autonomous navigation functions in parallel. So algorithms such as *VO* and those that comprise *AutoNav* become much more cost effective as can run concurrently to other main processor tasks. Therefore, while Perseverance will still rove at 4 cm s^{-1} (same as Curiosity), effectively it will be able to traverse faster, completing longer distances in a single sol.



Figure 2.23: Perseverance rover and Ingenuity helicopter on Mars soon after landing, ready to begin their missions.

As for the European counterpart, the *ExoMars* rover most basic operational baseline assumes the use of the *Follow Path* navigation mode, where an obstacle-

free path pre-computed on Ground is sent to the rover for it to follow without the need of mapping the environment or planning its path (Winter et al., 2015). This is similar to the original *Blind mode* of MER and MSL. Eventually, the ExoMars mission decided to include the *AutoNav* functionality onboard (Winter et al., 2017), leveraging on all the developments that had been happening for more than a decade in different European agencies, research institutes and industrial partners. Due to some project programmatic constraints, CNES is eventually also supplying to the mission another *AutoNav* implementation based on their algorithms developed during the last two decades on the EDRES system (Bousquet, 2011; Moreno, 2013). Thanks to this, Rosalind Franklin might validate on Mars, to the highest TRL, various navigation algorithms. At the moment, with ExoMars mission getting ready to be launched in one year, the Ground Test Model (GTM) of ExoMars is performing the final validation test at the Rover Operations Control Centre (ROCC) in Turin (see Figure 2.24).



Figure 2.24: ExoMars GTM during final validation tests in Turin.

For SFR (see Figure 2.25), as already mentioned in previous sections, a strong approach to the re-use of ExoMars algorithms is proposed, due to the tight project schedule to meet its launch date. However, given the fast navigating requirements of the MSR campaign, SFR needed to improve the autonomous navigating performances of ExoMars. Latest work in (Marc and Weclowski, 2018; Weclowski and Marc, 2019) demonstrates the use of an alternative, effectively faster, navigation mode. Its use is targeted for benign terrain and, despite having many common algorithms with *AutoNav*, it reduces the amount of

safety checks to speed up the traverse stopping for less time and with longer distance intervals. This *Follow Path with Safety Checks (FOPSA)* mode, is definitely increasing the distance traversed per sol with respect to the original *AutoNav*, but still requires the rover to stop periodically for mapping, i.e. *DEM* generation. Further optimisations of the *AutoNav* were originally considered in *SFR* in order to get it closer to a *continuous drive* traverse mode. Parallelisation of tasks computation by exploiting the multi-core co-processor integrated in the *SFR* avionics has been studied in the early phases of the project. This included the possibility to embark an additional *FPGA* co-processing unit to run part of the navigation algorithms (Lentaris et al., 2019). However, this design approach is currently discarded for system level energetic and thermal constraints that make the parallelisation efforts hardly bring any noticeable reduction in the total traverse time. Besides, a *continuous drive* does not appear to be necessary to fulfil the mission timeline, considering that *SFR* already drives six times faster than *ExoMars*.




Figure 2.25: Early conceptual design of *SFR*.

OVERVIEW OF THESIS CONTRIBUTIONS

This chapter gives a concise description of the main contributions of the thesis. These are accompanied by references to one or more publications made on them. As such, the sections hereafter provide a summary of the conducted research for this thesis and invites the reader to pick up on the various references for more detailed information. This is considered to be in line with the actual format of the thesis by compendium of publications.

Contributions are presented in chronological order, and show the evolution followed over the years of research for this thesis. It begins with a section dedicated to describe work performed in advance to the actual thesis development. This work served in many cases as stepping stones for the later work performed during the thesis and therefore was considered relevant to be included in this manuscript. This section starts describing the conception of robotics systems and integration of full rover platforms that have served the Planetary Robotics Lab (PRL) of ESA for its continuous testing and support in R&D activities and mission programs. This is followed by the introduction of two test campaigns in the frame of the ExoMars mission. The first one was instrumental to demonstrate the relevant role of a particular locomotion mode, i.e. Wheel Walking, for the mission, improving significantly the traction capabilities of the rover. The second allowed the ExoMars team to gain knowledge and important experience on Remote Operations, demonstrating the challenges of rover missions with respect to their remote control and operation. The Remote Operations campaigns showed also the importance of embarking autonomous features in rovers in order to accomplish mission objectives. This closes the section describing the introductory work, and coincides with the moment in time in which the lab decided to focus its research on the field of autonomous navigation.

The following four sections are each dedicated to one main contribution of the thesis and, as already mentioned, are supported by the work introduced up to there. This starts in Section 3.2 with the description of Dataset Acquisition field tests carried out to provide representative sensor data taken on Mars analogue terrains on Earth. The preparation for these field tests and the datasets

gathered have allowed and inspired the PRL members to conduct the research activities on the field of autonomous navigation in which this thesis is focused. In Section 3.3, an autonomous navigation mode suitable for relatively benign terrain is presented. Novel components for trajectory control, path planning and hazard detection are introduced and their integrated execution to efficiently navigate is demonstrated in a field test campaign with several traverses of medium and long range. In Section 3.4, the development of a new SLAM algorithm is described and the results of validating its performance with previously acquired datasets in analogue field tests are presented. All contributions culminate in Section 3.5, where they are combined to introduce the design of a complete GNC architecture that could be exploited in future Mars planetary rover missions such as the SFR mission of the MSR programme.

3.1 PREVIOUS RELEVANT WORK

3.1.1 *Development of Planetary Rover Prototypes for R&D activities*

Planetary rovers are complex multidisciplinary systems that require a sound integration of their mechanical, electrical and software designs. A good understanding of these three aspects and the system engineering approach that encompasses them are instrumental in order to perform any research in robotics. So far, all Mars surface mobile missions have opted for a wheel-based solution for their locomotion subsystems, mainly due to the reduced complexity and mass and power required when compared to other solutions such as leg- or track-based locomotion. Therefore, in this work, we have focused on the design and challenges faced by wheeled rovers.

In the PRL of ESA we have developed in the last decade several prototype platforms that have served as testbeds to perform valuable research activities and test campaigns. These rovers have been built with a certain degree of space representativeness mimicking several aspects of real rover missions, and in particular the ExoMars rover system (see Figure 3.1), due to its unique role in ESA's planetary exploration programmes up until recent years. Here we briefly show three lab rover prototypes that have supported the work performed in the elaboration of this thesis: the ExoMars Testing Rover (ExoTeR), the Martian Rover Testbed for Autonomy (MaRTA) and the Heavy Duty Planetary Rover (HDPR).

ExoTeR (see Figure 3.3) is a half-scaled prototype version of the ExoMars rover system. It was first designed and built between 2008 and 2010 and has gone through further development and integration of several robotics systems during the last decade. It mimicks the ExoMars locomotion system kinematic configuration, known as the Triple Bogie (see Figure 3.2). Opposed to the Rocker Bogie




Figure 3.1: Computer rendered image of the [ExoMars](#) mission rover, Rosalind Franklin




Figure 3.2: The [BEMA](#) Triple Bogie configuration and right bogie beam with actuator locations (Poulakis et al., 2015).

system of the [NASA](#) rovers, this one features three points of attachment to the rover body and does not require a differential beam across to provide platform stability. Each wheel module comprises driving (DRV), steering (STR) (except for the two center wheels) and deployment (DEP) actuators. Deployment actuators permit the system to be stowed by putting the wheels upwards to optimise volume and system accommodation, particularly during the launch and cruise mission phases. Deployment actuators can be further exploited during the surface mission, enabling the implementation of a locomotion mode referred to as Wheel Walking (see [Section 3.1.2](#) for more details on this mode). [ExoTeR](#) also comprises nowadays a mast with a Pan-Tilt Unit (PTU) and a robotic manipulator. The first one allows to embark sensors at the top of it with the possibility to point them in any desired direction to perceive the environment. As for the second, despite [ExoMars](#) does not include a manipulator, it has allowed the [PRL](#)

to perform already relevant research activities within the scope of the sample fetching mission scenario of *SFR*. In terms of sensors, *ExoTeR* embarks two stereo cameras for localisation and navigation, providing similar functionality to the *ExoMars* Localisation Camera (*LocCam*) and Navigation Camera (*NavCam*). It also has a full *IMU* sensor and wheel and bogie pivot encoders for rover odometry estimation.




Figure 3.3: *ExoTeR* at the PRL of ESA

To *ExoTeR* followed *MaRTA* (see Figure 3.4), built between 2017 and 2019 with the same half-scaled approach of *ExoMars*' design. Due to the evolution of the design of the Bogie Electro-Mechanical Actuator (*BEMA*) locomotion system of the *ExoMars* rover (see Figure 3.2), *MaRTA* has slightly different dimensions and kinematic configuration compared to *ExoTeR*, while still maintaining the Triple Bogie design approach. Its thinner wheels and additional steering actuator units in the middle wheels make *MaRTA* more representative of the latest *ExoMars* rover design. It has already integrated the same mast and *PTU* as *ExoTeR* and, currently, an upgraded redesigned version of *ExoTeR*'s manipulator is being developed to be soon integrated in *MaRTA*. It is also embarking, or meant to embark soon, a set of sensors similar to those in *ExoTeR* that will allow it to implement all the necessary navigation functionalities. Finally, *MaRTA* also included in its design several significant improvements that came out of the lessons learnt from building and testing *ExoTeR*, making it a more robust and reliable research platform.

Thorough details of these two *ExoMars*-representative platforms are given in (Azkarate et al., 2021), an article recently submitted and currently under review.




Figure 3.4: MaRTA at the PRL of ESA

For the readers reference, the preprint version of this paper is included in the annexes together with other main authored publications.

HDPR was built between 2015 and 2017 as an outdoor testing and payload carrying platform. This six-wheeled medium size rover resembles MER in terms of dimensions and kinematic configuration. HDPR has been used in many field test activities, which are detailed in further sections below, due to its high autonomy and payload capacity. This is thanks to the power generator mounted on top and the powerful driving actuators. It has been used to embark several exteroceptive sensors such as optical stereo cameras, ToF cameras or LiDARs that fulfilled different purposes. These are mounted at the front truss structure or on top of the high mast which also includes a PTU. Camera sensors at the bench mounted on the PTU are placed at a height of 2 m from the ground, same as ExoMars Panoramic Cameras (PanCam). Finally, it also integrates a GNSS antenna and receiver module to provide ground truth position information. This is combined with the IMU and an additional single axis (*yaw*) Fiber-Optic Gyro sensor to add ground truth orientation information and provide full pose data. In Figure 3.5 a picture of HDPR during a field test campaign is shown where many of the mentioned sensors and components can be seen.

A general system description of HDPR can be found in (Boukas et al., 2016).



Figure 3.5: HDPR during a field test near ESA-ESTEC

REFERENCES

- Azkarate, Martin et al. (2021). «Concept, Development and Testing of Mars Rover Prototypes for ESA Planetary Exploration.» In: *IEEE Robotics and Automation Magazine (under review)*.
- Boukas, Evangelos et al. (2016). «HDPR: A mobile testbed for current and future rover technologies.» In: *13th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*.

3.1.2 Improving Rover Locomotion Capabilities with Wheel Walking

As already explained in [Chapter 1](#), the Martian surface is mostly similar to a desert-like terrain on Earth. There is no vegetation or humidity, and it comprises a mix of sand and soil of different granularity, and slopes, rocks and bedrock varying in distribution and abundance such as in the landscape seen in

Figure 1.2. Planetary missions are driven by science and their level of success is dependant on the scientific objectives they accomplish, which sometimes have led rovers to traverse significantly (and unpredictably) challenging terrain. The locomotion subsystem has always been seen as a critical part of the rover, since it provides mobility capabilities to the platform. The more capable a rover's locomotion system is, the more science objectives the mission will be able to achieve. In contrast, a failure in the locomotion performances can quickly lead a rover mission to an end. Motivated by the difficulties that MER rovers had traversing the Martian surface, even getting stuck in loose soil several times (Arvidson et al., 2010), and inspired by the peristaltic motion demonstrated on Lavochkin's planetary rover prototypes (Ehrenfreund et al., 1998), the PRL started a project in 2014 to implement and evaluate the Wheel Walking locomotion pattern.

Wheel Walking refers to a rover locomotion mode that synchronises the motion of the wheel driving motor with another motor that is connected through a lever or leg. This can be used to swing the wheel back and forth following a pattern that increases traction on soft soils. In the case of ExoMars, this second motor is referred to as deployment motor and, as already explained in Section 3.1.1, the Triple Bogie locomotion system has a total of 6 deployment motors, one for each wheel. The Wheel Walking locomotion mode can be implemented using more or less complex kinematic motion models, depending on the assumptions or simplifications made to model the motion of the platform over the terrain. Our Wheel Walking implementation uses the full 6-DoF kinematic relations based on the odometry model described in (Hidalgo-Carrio, Babu, and Kirchner, 2014) and inverts the Jacobian formula to compute at each control cycle the optimal command for each independent joint that shall produce the desired rover body motion. This may read like a complex control, yet, even more simple Wheel Walking implementations using open loop timely synchronised actuator commands have shown to be effective. The observed outcome has always been an improved tractive performance in challenging conditions, such as sandy terrains or high slopes, where the nominal roving motion is subject to high slip ratios.

In our experiments we tested the Wheel Walking mode with ExoTeR in a set of scenarios. One of them started with the rover bogged down in loose sand simulating a sank in a sand-trap condition and then commanded to wheel walk out of it. The same experiment was repeated by commanding the platform to simply rove using the standard wheel driving mode. In Figure 3.6 the sequence of images (or accessible video¹) shows the significant difference in traction performance between the two cases. The rover appears to be completely stuck

¹ <https://www.youtube.com/watch?v=qk0KzFq1SpY>

in the normal driving case showing little progress and a slip ratio higher than 95 %, while in the wheel walking case the rover manages to get out of it just immediately.

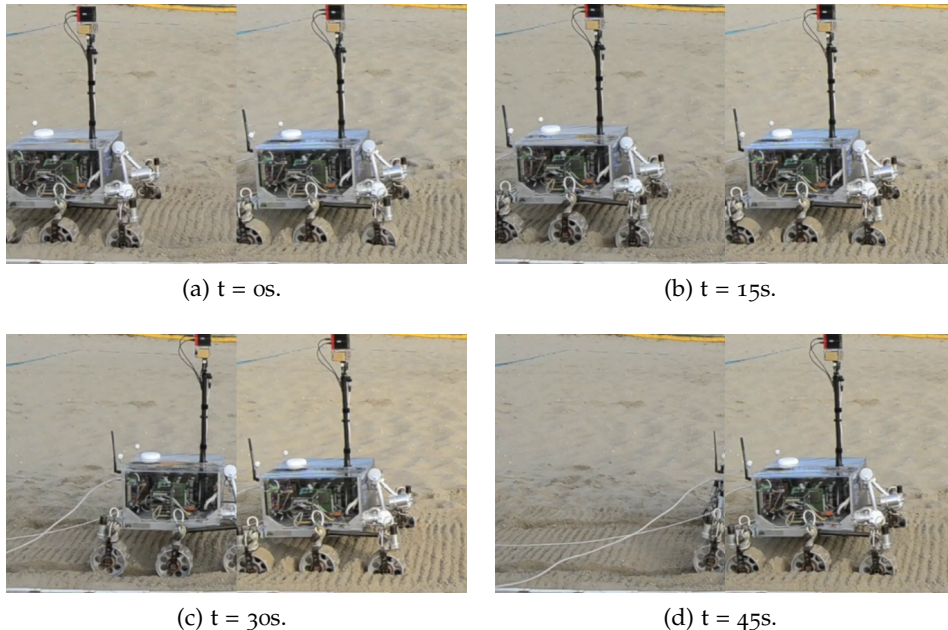


Figure 3.6: Wheel Walking test in loose sand trap.

Other scenarios focused on testing the gradeability capabilities of Wheel Walking. These run the rover up towards slopes of different angles and measured the slip ratio difference between the Wheel Walking and standard driving mode. The encouraging results, by which Wheel Walking outperformed the standard driving in all tested cases, motivated a second test campaign at the German Space Agency (DLR)'s Robotics and Mechatronics Center (RMC) facilities in Munich in early 2015 (see Figure 3.7), making use of a significantly larger test bed and different soil types available there, allowing us to validate our results in a wider set of test conditions. These campaigns were instrumental to demonstrate the need of this capability for the ExoMars mission, and motivated further research (Wiese, 2017) that eventually led to Wheel Walking being implemented as a locomotion mode on the ExoMars flight rover.

Further details on the Wheel Walking test campaigns and their experimental results can be found in (Azkarate et al., 2015).

Additionally, and following our work in the research line of Wheel Walking, this led to another recent publication in (Pérez del Pulgar et al., 2019), where a



Figure 3.7: ExoTeR at DLR performing a Wheel Walking test.

model to estimate the slip in different terrains and an approach to choose and switch between the best locomotion mode was presented.

REFERENCES

Azkarate, Martin et al. (2015). «First Experimental Investigations On Wheel Walking For Improving Triple-Bogie Rover Locomotion Performances.» In: *13th Symposium on Advanced Space Technologies for Robotics and Automation (ASTRA)*.

Pérez del Pulgar, Carlos J et al. (2019). «Choosing the Best Locomotion Mode in Reconfigurable Rovers.» In: *Electronics* 8.7. ISSN: 2079-9292.

3.1.3 Remote Rover Operations Field Test Campaigns

Approaching the original launch date for the ExoMars rover, the mission team needed to exercise the operational aspect of the rover, specially considering that ExoMars is the first rover that ESA is launching to space. The objective of the remote rover operation test campaigns was to assess the readiness level and adequacy of the procedures and decision making processes established for the future ExoMars ROCC. Two campaigns in consecutive years took place between CNES facilities in Toulouse and ESA facilities in European Space Research and Technology Centre (ESTEC). The ExoTeR rover was taken to the *Site d'Essai pour les Rovers Mobiles (SEROM)* field site in Toulouse (see Figure 3.8) while, in parallel,

a rover operations centre was temporarily arranged at [ESTEC](#), emulating the conditions of the [ROCC](#).




Figure 3.8: ExoTeR ready for Remote Operations tests at [SEROM](#) in [CNES](#).

In 2015, the campaign focused on the Egress phase of the mission. The objective was to validate whether the telemetry data coming from the rover together with the tools available at the control centre were sufficient to evaluate the potential hazards and decide on the egress direction in full situational awareness to minimise the risk of failure. Our team in the [SEROM](#) orchestrated up to five Egress scenarios adding several hazards and hidden traps for the operations team in [ESTEC](#). Some of the telemetry data used by the [ROCC](#) team comprised a set of panoramic images and local reconstructed [DEMs](#), such as the one seen in [Figure 3.9](#).

The campaign was a success, demonstrating that the operations team was capable of identifying any potential risk and managed to verify the procedural telecommands and telemetry checks to guarantee the safe egress of the rover.

In 2016, the focus was set on demonstrating post-Egress operations with the commissioning activities involving the rover and lander platform and a subsequent traverse towards the first scientific target and experiment cycle. Yet the campaign started with another rehearse of the Egress operations, which was considered adequate due to the confidence gained in the previous campaign. While the Egress itself was successful, it conditioned the rest of the operations considerably by setting the rover in a challenging situation for the upcoming activities, eventually failing to accomplish them, as it hit a rock while trying to escape across steep slope that produced significant slip (see [Figure 3.10](#)).




Figure 3.9: Example [DEM](#) data gathered from the rover's telemetry.




Figure 3.10: [ExoTeR](#) at post-egress challenging traverse. Left to the rover is the rock that it eventually hit.

Despite the unaccomplished objectives on the second, both campaigns provided many important lessons for the [ExoMars](#) mission rover operations team and served to validate numerous rover and control centre functionalities. More details on these campaigns can be found in (Azkarate et al., 2016).

REFERENCES

Azkarate, Martin et al. (2016). «Remote Rover Operations: Testing the Exomars Egress Case.» In: *13th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*.

3.2 DATASET ACQUISITION FIELD TESTS

The Remote Operations tests in CNES showed us how important it is the implementation of more complex navigation capabilities onboard the rover. If every single motion drive and safety check needs to be computed on Ground before it is commanded to the rover, the work of operators becomes quickly very tedious and the actual traverse of the rover highly limited. Similar to how the Sojourner rover was commanded during the pioneering Mars Pathfinder mission, which barely advanced a few meters at each operation cycle. Therefore, the PRL saw necessary to implement several functions to keep up with the needs of future planetary exploration missions. For example, having robust onboard localisation means, being able to follow a given path or to detect obstacles in front of the rover. In order to support the development of such functions and other core algorithms, and be able to eventually validate them in a representative environment, two field test campaigns were conducted focused on the acquisition of relevant sensor data. The gathered datasets could be used at a later stage to run and validate the algorithms in development. These field test campaigns were conducted using the HDPR rover in outdoor scenarios.

The first campaign was performed in the coastal area of Katwijk, a town close to our ESTEC work site. The test run along the beach shore added up to more than 2 km of traverse with several proprioceptive and exteroceptive sensor data logged. These include IMU and wheel encoder data, as well as front facing stereo LocCam and PanCam images, ToF camera images and LiDAR scans. Ground truth data was acquired with a Differential GNSS system for localisation and reference DEM and ORI maps collected with a surveying fix-wing professional drone. In Figure 3.11 the HDPR rover can be seen running this field test campaign.

All gathered data was finally processed and made publicly available and is nowadays hosted and accessible at our robotics group's servers². All pertinent details of this work are found in the accompanying publication (Hewitt, Boukas, Azkarate, et al., 2017).

After the successful experience at the Katwijk beach, it was decided to perform a second analogue campaign in an environment more representative of a plane-

² <https://robotics.estec.esa.int/datasets/>



Figure 3.11: HDPR during the field test campaign in Katwijk.

tary mission scenario. This time the tests took HDPR to Tenerife, in the Canary Islands, and particularly to the national park area where the Teide volcano is located. In Figure 3.12 HDPR can be seen at the field test site location. After several days of testing campaign, a total of 13 km were traversed with test runs executed at different times of the day, including some night-time traverses with artificial lighting onboard that could mimic a Lunar operation. The sensor suite for the dataset was very similar to the one used in Katwijk with some minor changes in the accommodation and the addition of another stereo cameras for navigation (NavCam). Ground truth localisation and reference maps were taken with the same equipment.



Figure 3.12: HDPR during the field test campaign in Tenerife.

The data gathered in Tenerife has been instrumental for the development and validation of several algorithms in our research group, including some of the work detailed in the sections hereafter. While the dataset is not yet publicly

available, our team is working on a similar publication to the one of the Katwijk beach dataset to make it accessible to the community.

REFERENCES

Hewitt, Robert A., Evangelos Boukas, Martin Azkarate, et al. (2017). «The Katwijk beach planetary rover dataset.» In: *The International Journal of Robotics Research* 37, pp. 12 –3.

3.3 EFFICIENT AUTONOMOUS NAVIGATION

In this section we describe an efficient solution for autonomous navigation of planetary rovers that can be exploited in relatively benign terrain with a low distribution of rocks or Cumulative Fractional Area (CFA). The principle behind this navigation approach is that for a rover navigating through a terrain where a low presence of hazards is foreseen, a roughly safe path can be calculated on Ground. This is done based on the available orbital imagery such as the maps provided by HiRISE of MRO. Despite the favourable terrain conditions, the relatively low resolution of orbital maps cannot guarantee the path to be completely safe. This is simply because rocks of small to medium size cannot be perceived beforehand at the given scale. However, in this conditions a considerably more simple approach to hazard detection can be taken. Computing a full DEM of few meters of terrain ahead to perform a path planning that avoids the hazards identified on it would most of the times become an unnecessarily resourceful task and burden. Moreover, this would also slow down the traverse speed. Instead, here we aim to implement a reactive hazard detection and avoidance navigation approach. This shall allow us to detect any imminent hazards in the close vicinity of the rover that would fall under the rover traverse unless reacted over them. Under the low CFA conditions, this would permit the rover to conduct a significant percentage of the traverse without stopping at all. This brings by default an increased traverse speed and an optimised use of computational resources onboard. In this scenario conditions, the rover executes the traverse following the pre-computed global path and only stops to re-plan locally the path whenever a hazard is detected. The re-planned path avoids the hazard and rejoins the global path trajectory again.

Three key algorithms were developed to support this autonomous navigation approach. All of them were implemented trying to minimise the computational resources required. These algorithms are described hereafter and the experimental results of a field test campaign implementing this navigation mode is presented in the end.

3.3.1 Trajectory Control

This is the core control algorithm that actively steers the rover platform to follow a given path as accurately as possible. As already explained in [Section 2.5](#), this function shall apply any corrective manoeuvres in order to drive along a sequence of waypoints and keep the rover within a predefined safety corridor. The implemented controller is a geometry based path tracker that extends on the concepts of the *pure-pursuit* algorithm. The name of this base algorithm comes from the fact that it constantly tries to reach a virtual *look-ahead* point in front at a defined distance from the rover. The selection of a geometry based approach is considered the most appropriate for the slow dynamics involved in rovers motion and, additionally, it simplifies the parameters configuration and tuning. For the development of this trajectory controller, named *conservative-pursuit* (or simply *c-pursuit*), the original algorithm is modified to ensure that the path tracking is bounded by a maximum allowable tracking error which corresponds to the predefined safety corridor. This is considered a critical safety feature for planetary exploration missions that the original *pure-pursuit* did not guarantee. The reactivity of *c-pursuit* is also increased by a dynamic factor proportional to the latest estimated tracking error that modifies the distance to the *look-ahead* point. This makes the algorithm more adaptive to the terrain conditions forcing it converge faster to the nominal path when the deviation is greater.

Additionally, *c-pursuit* automatically produces smooth trajectories, without requiring paths with closely interpolated waypoints. The distance between waypoints might be arbitrary and different for each waypoint and the controller smoothly transitions between segments connecting consecutive waypoints. The controller behaviour is represented in the plot below ([Figure 3.13](#)) that shows the vector field where the instantaneous trajectory followed by the rover at each possible coordinate is displayed.

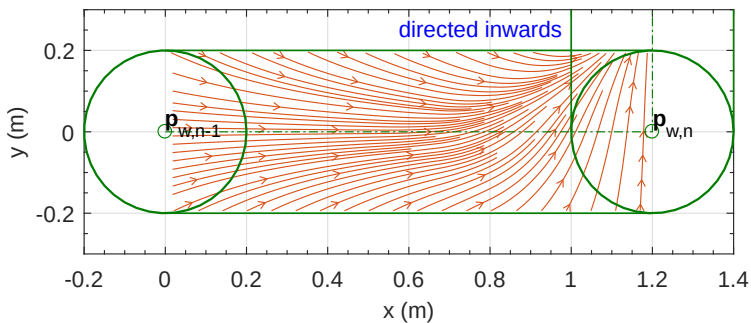


Figure 3.13: *c-pursuit* behaviour representation: Vector fields of direction to *look-ahead* point from all possible inside the safety corridor.

As the plot shows, the rover is always guided inwards within the safety corridor limits (these are represented by the green parallel lines in the figure) and its trajectory smoothly transitions between consecutive waypoints avoiding the execution of sharp turns or point turns. Finally, it is worth mentioning that the computational load of this algorithm is significantly low with only few 2D geometric computations needed at each cycle to compute the control command, which can easily run in real-time at sufficiently high frequency.

3.3.2 Hazard Detection

The implementation of a computationally efficient Hazard Detector is probably the key to this navigation approach. In fact, the detector runs in parallel to the rover motion, without requiring the rover to stop. It uses the input coming from the *LocCam* data to run the detector within a Region of Interest (*RoI*) contained in the near field of the rover. The advantage of delimiting the detection to the short-range in the cameras Field of View (*FoV*) is twofold. The first one is to guarantee that the detected hazard is actually falling onto the rover's trajectory and an avoidance manoeuvre is necessary. The second is to reduce the execution time of the detector by simply reducing the number of computations. This is illustrated in the [Figure 3.14](#) below.

As already discussed in [Section 2.2](#), most common approaches to the detection of hazards and identification of non-navigable areas are based on the generation and later analysis of a *DEM* of the terrain ahead, using images taken from the *NavCam* mounted on top of the mast. This is usually a relatively costly process with computations in the 3D space that can take up to several seconds, if not longer, to finish. The designed and implemented detector in this thesis takes a radically different approach since it is only based on stereo disparity calculations over a *RoI* in the *LocCam* images. Since the disparities of the *LocCam* images are already computed at the *VO* process for localisation, the detector does not require additional stereo processing to run, nor it builds a 3D map representation of the environment. This is a differential factor to significantly reduce the execution time of the hazard detector and enable its execution parallel to the rover motion.

The actual detection of hazards is achieved by comparing the disparity or depth values to certain threshold values calibrated in advance. Given the fixed position and orientation of the camera in the body frame, the depth corresponding to a nominal flat ground can be computed and stored in a calibration matrix. Then, the threshold values for the presence of an obstacle that exceeds a certain height can be computed. However, due to the perspective view of the camera the depth data cannot be directly subtracted by a constant value equivalent to the height limit for obstacles. Therefore, at this point




Figure 3.14: Hazard Detector Visualization on [LocCam](#) image: green means safe, red pixels are considered a hazard, and pixels without an overlay are not considered because they are outside of the [RoI](#) or because the stereo disparity could not be computed.

we make use of the projection angle and a simple trigonometric formula, to determine the depth values for each pixel in the [RoI](#) that the rover would see when an obstacle appeared in front. This is represented in [Figure 3.15](#), where $d_{\min,i}$ ($d_{\max,i}$) is the minimum (maximum) threshold distance for pixel i where an obstacle (ditch) of a height T_{near} (depth T_{far}) is detected.

After computing the threshold depth values for every pixel i in the [RoI](#), which can be done as part of the calibration procedure, every depth computed by the stereo disparity processing component can be compared to these thresholds to decide whether they are hazardous or not. The detector output is then run through a statistical filtering process that makes the detector robust to spurious outlier pixels (errors or noise in disparity). When the detector has positively identified a hazard in the camera frame, a two-dimensional perspective transformation is used to translate the location of the hazard onto the ground within the rover body reference frame. As for the computational complexity of the detector algorithm, note that most of the outlined computations are performed as part of the calibration and do not contribute to the execution time of the hazard detection at runtime. Once the system is set-up, this enables the detector to be run in real time at a sufficiently high frequency.

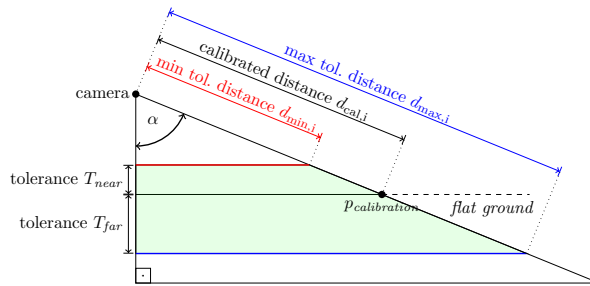


Figure 3.15: The underlying geometry on which the hazard detection is based. The tolerated distances from the camera are computed in such a way that hazards would be detected at certain heights.

3.3.3 Path Planner

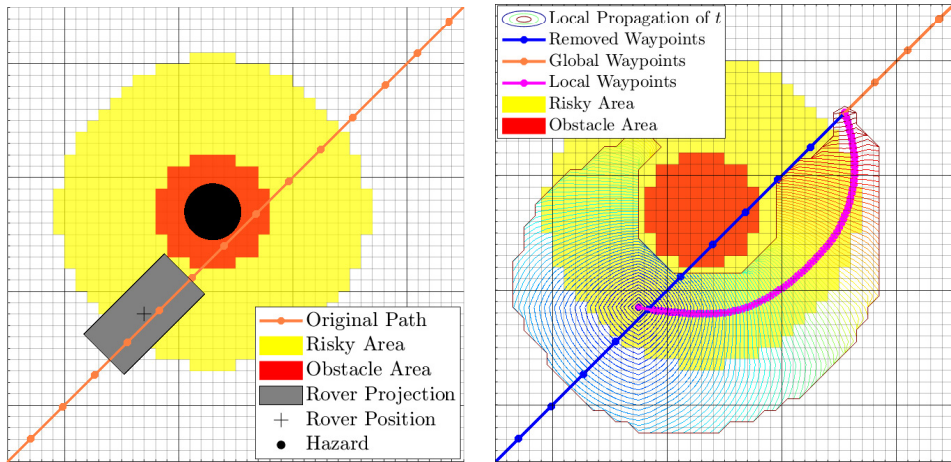
This component gives the rover the capability to react to detected hazards, and allow to continue the traverse instead of waiting for a new updated safe path to be provided from Ground. When a hazard is detected, the rover acquires more data describing the environment in which it is moving. The roughly safe path being followed up until the detection of the hazard, was previously created without considering this unknown information. Therefore, a local path planner is employed to dynamically update the path according to the newly received traversability data. The objective of it being to generate a new set of waypoints that will guide the rover around the detected hazard(s) while progressing along the original global path.

The implemented local path planner is based on the [FMM](#) approach introduced in [Section 2.4](#), in particular on its heuristic version (Heuristic Fast Marching ([FM*](#))). [FMM](#) is used to numerically compute the expansion of a wave through a continuous 2D surface represented by the traversability map. As a result, it provides a potential field from which the path connecting two points can be extracted by means of the gradient descent method. This results in a complete and optimal planner that despite working with traversability map grids has the capability to produce smooth solutions such as other *any-angle* planners like [Field D*](#). The resulting path is essentially a continuous single curve connecting two points in the 2D grid that integrate the least possible amount of cost, making it the optimal solution. Interestingly, the planner has a low computational complexity similar to more simple planners like [A*](#), which corresponds to $O(\zeta \log \zeta)$, where ζ is the total number of grid nodes. To further increase the speed of the computation at the expense of obtaining quasi-optimal paths, an heuristic function is used in [FM*](#), that prioritises the expansion of the wave towards the goal direction, reducing the number of visited nodes in the process. The decision of

using FM^* goes in line with the objective of limiting the computational resources used.

The cost map used by this method comes in the form of a 2D matrix field that indicates the difficulty of traversing each node and determines the rate at which the wave propagates over the surface at each point. The traversability grid is basically translated into this cost map where identified obstacles are assigned the highest possible cost values. The regions close to the obstacles and up to a certain defined distance, despite traversable, are considered risky area and all regions beyond this risk distance are considered safe area. Nodes in the safe area are assigned the lowest possible cost, however this has to be non-zero positive. Risky nodes are assigned a cost according to a gradient that goes from the highest cost of the obstacles nodes to the lowest cost of the safe nodes. The parameters for the risk distance and gradient can be tuned depending on the conservatism or margins to be taken from the hazard. This goes at the expense of introducing more or less pronounced turns at the planner's solution. Since the speed of the wave propagation is inversely proportional to the cost assigned to cells that it traverses, the wave expansion will follow the route that avoids the obstacles while minimising the length of traverse. The function that defines the relation between the wave expansion and the cost is known as *eikonal* equation. As a result of applying this equation along the wave propagation, the potential field is obtained. Two relevant features are added to the original FMM algorithm for the implementation of the path planner. First, is the capability to dynamically replan an already computed path, such as in the D^* algorithm. And second, the possibility to apply the planner simultaneously at different layers of resolution is included. These two improvements are necessary for its use case, where the path needs to be dynamically fixed at a local scale to avoid the unforeseen hazards at the global scale.

The behavior of the local path planner using the FM^* method is illustrated in [Figure 3.16](#). It describes a hypothetical case of replanning of a path due to the presence of a hazardous element. As seen in [Figure 3.16a](#), the grid is formed by square nodes placed regularly over the map. The original global path, represented by connected orange dots, is passing through nodes that now fall under obstacle or risky areas. The rover must compute a new set of waypoints to avoid the obstacle area nodes and eventually reconnect with the previous global path. To accomplish this, the propagation wave must start from the rover position and reach a certain global waypoint from the previous path. This waypoint is selected as the first one that is placed in the safe area immediately after those within the obstacle and risky areas. If a waypoint further away is selected, the resulting path becomes smoother, although more computation is required since the wave would have to propagate further.



(a) An obstacle is found by the rover. The surrounding area is represented on a regular square grid and divided into *Obstacle Area*, *Risky Area* and *Safe Area*.

(b) The local path planner finds an alternative path made up by local waypoints, which guides the rover to rejoin the original path while avoiding the obstacle.

Figure 3.16: Graphical example case where the rover finds a hazard on its way (a) and plans an alternate path to avoid it (b).

To reiterate, the total time spent on computing the FM^* wave is dependant on the number of nodes the wave expands to and therefore is dependant on the number, size and location of the obstacles found between the rover and the selected global waypoint. This dependency makes the algorithms computational time non-deterministic and it may therefore vary according to the particular situation encountered. However, given the short range of the local path re-planning, the computation time is kept significantly low. Furthermore, its impact can be considered negligible with regards to the overall traverse timescale due to the fact that the function is only executed when a hazard is detected. It is to be expected that this will be a sizeable number of times throughout the traverse. All details related to the path planner's algorithmic design and its implementation, named *DyMu*, can be found in (Sánchez-Ibáñez, Pérez-del-Pulgar, Azkarate, et al., 2019). The main author of that paper, R. Sánchez, did the path planner implementation as part of his own thesis studies, that were partially carried out in collaboration with the PRL and during his research stay at ESA within personal supervision.

REFERENCES

Sánchez-Ibáñez, J Ricardo, Carlos J Pérez-del-Pulgar, Martín Azkarate, et al. (2019). «Dynamic path planning for reconfigurable rovers using a multi-layered grid.» In: *Engineering Applications of Artificial Intelligence* 86, pp. 32–42.

3.3.4 Experimental Results

The Efficient Autonomous Navigation mode introduced in this section using the three described algorithms has been validated in a representative field test campaign executed in an analogue site near [ESA-ESTEC](#) facilities. The campaign consisted of several traverses of medium and long range which eventually demonstrated the viability of this navigation mode for planetary exploration missions. This included its robustness to varying conditions and difficulty.

The [HDPR](#) rover was used to run this test campaign and the stereo camera placed at the rim of the rover body and pitched down (see [Figure 3.5](#)) was used as the sensor for hazard detection, emulating the [ExoMars LocCam](#). After having validated the hazard detector in a set of varied rover orientation, environmental and lighting conditions, the test field was prepared by placing a significant amount of hazards spread around the analogue terrain. This can be seen in [Figure 3.17](#).



Figure 3.17: Panoramic view of the field test setup.

The rover was commanded to execute paths that would guide it across the test site, ensuring that it would encounter several hazards on its way. After several successful traverses, a final very long traverse was commanded to test the endurance and reliability of the navigation approach. The rover executed the path successfully replanning 21 avoidance manoeuvres on its way, with an overall traverse length of 191 m in 43 min. This is equivalent to a net traverse speed of 7.4 cm s^{-1} , which is a big achievement considering the rover's maximum driving speed was 10.0 cm s^{-1} . The trajectory executed for this test can be seen in the aerial view of [Figure 3.18](#). The rover can be seen at its starting position in the top right corner of the image and the traverse finishes at the center right of the

image after traversing through and around several craters and hazards on the field.




Figure 3.18: Efficient Autonomous Navigation test traverse. The red arrows represent the pre-planned waypoints while the thin blue line shows the actual traverse. The total planned traverse length was 171 m and the actual executed trajectory length 191 m.

Another important objective of the field test was to quantify the efficiency of the developed algorithms by estimating their execution time, in particular for the hazard detector and the path planner. The trajectory control was assumed to have negligible impact and actually to have a computational load similar to other path tracking algorithms. The hazard detector was proved to have a deterministic and very low execution time of 9 ms. Despite the fact that the processing module used in this test campaign was not a space grade unit, it is acknowledged that the algorithm could be ported to a more representative processor with enough margin for it to run in parallel to the rover's motion. As for the path planner, its execution time was measured and statistically analysed showing an average execution time of 75 ms and a standard deviation of 35 ms. Given these results, and recalling that the path planner is not periodically executed but only sporadically after the detection of a hazard, it is concluded that the impact of the planner on the overall stopped time of the rover cannot be considered significant.

Finally, other than that the algorithms execution time, the efficiency of the navigation mode is also demonstrated by the potential increase of length and net speed of traverse in benign terrain with no detailed *a priori* information of hazards.

More detailed information about the implementation of the three main functionalities and experimental results of the efficient navigation mode can be found in (Gerdes, Azkarate, et al., 2020).

It should also be highlighted that an additional field test experiment was executed, as an extension of the these efficient navigation test campaign, where an added feature to the path planner was tested. This referred to the capability of the planner to dynamically update the cost map of the terrain in real-time according to specific traverse performance metrics. This showed to be particularly interesting for round-trip traverses, such as the [SFR](#) use case. Performance statistics were gathered for the outbound traverse, so that the return inbound traverse path planning was executed using real-time informed data, making it clearly more efficient. These experiments led to a recent publication in (Paz-Delgado, Azkarate, Sánchez-Ibáñez, et al., 2020).

REFERENCES

- Gerdes, Levin, Martin Azkarate, et al. (2020). «Efficient autonomous navigation for planetary rovers with limited resources.» In: *Journal of Field Robotics* 37:7, pp. 1153–1170.
- Paz-Delgado, Gonzalo J, Martin Azkarate, J Ricardo Sánchez-Ibáñez, et al. (2020). «Improving Autonomous Rover Guidance in Round-Trip Missions Using a Dynamic Cost Map.» In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7014–7019.

3.4 SLAM FOR AUTONOMOUS PLANETARY ROVERS WITH GLOBAL LOCALISATION

In [Section 3.3](#) the localisation aspects were intentionally not addressed and the perception means optimally simplified under the benign terrain assumption. In this section the localisation function is explored, and a [SLAM](#) method for planetary rovers is proposed assuming that the traverse conditions would require a more comprehensive mapping of the terrain. The [SLAM](#) function tries to optimise the localisation and perception data processing modules taking a different approach to that of more classic sequential execution of these two pipelines. To further deal with the inherent drift in localisation of any dead-reckoning method a global localisation function that uses the output map of the [SLAM](#) solution is introduced. This module corrects the drift and reduces the localisation uncertainty at the global scale using orbital maps as reference.

Hereafter, the design of the [SLAM](#) and global localisation modules is described and the results of their experimental validation is presented.

3.4.1 SLAM

As already discussed in [Section 2.3](#), SLAM is a process that has hardly been adopted in the space robotics realm, except for few studies and work found in the literature that have tried to find different ways to reduce the computational load of SLAM to be able to embark it in space grade processors. The design approach of the SLAM method described here differs from other methods found in the literature on the fact that this one is made particularly fit for the planetary rover exploration case, and in particular for Mars rovers. It focuses on the specific needs and conditions that Mars exploration poses and solves the localisation and mapping problem that best allows for long-range autonomous navigation in an efficient way.

From the mapping side, the SLAM solution shall produce a volumetric elevation grid map (2.5D) similar to a DEM representation with an appropriately chosen grid resolution. This resolution is optimised according to the rover locomotion and traction capabilities. There is no point on generating highly dense 3D reconstructions of the environment, neither keeping a map of sparse visual features in the 3D scene. The elevation grid should be sufficient to later produce a traversability map that the rover can use to solve the local path planning and navigation problem. As such, the size of the generated map is bound and limited to a certain distance around the rover body. As the rover moves, the map is shifted along to always provide a local, rover-centered grid with the latest environment information.

As for the localisation, the objective shall be to provide an accurate pose estimation (relative to the map) while trying not to deviate from the globally planned path. The drift in localisation is corrected using the global localisation module explained after. However, it is worth mentioning that for the localisation process in SLAM, no history of poses is kept in order to identify loop closures, like in Graph SLAM approaches, since these could rarely happen and it would increase significantly the use of onboard memory and computational resources.

Taking all these conditions into account the designed method is based on a mix of the particle filter and the scan matching techniques. The particle filter uses odometry inputs for the pose prediction step while the pose update is done using perception measurements. The concept behind Iterative Closest Point (ICP) (Zhang, 1994) is used to assign the weights to each particle where the latest scan perception is matched to the map foreseen from the particle pose. This ensures consistency of the built local map and an accurate localisation of the rover pose relative to the terrain perceived. Data process optimisation is sought by pre-processing and filtering the perception data to minimise the memory usage and by balancing amount of particles used in the filter.

A conceptual schematic view of the SLAM design is shown in Figure 3.19.

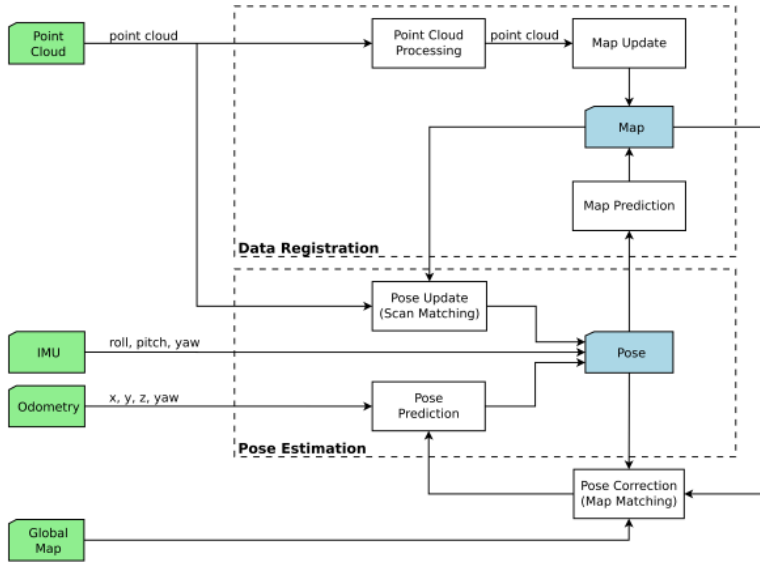


Figure 3.19: The high-level design diagram of the SLAM system.

Finally, a visual representation of this process is shown in Figure 3.20.

3.4.2 Global Localisation

As mentioned in the previous section the objective of the global localisation function is to correct the drift that the relative localisation builds along the traverse. Additionally, the intent is to achieve this by reusing the output of the SLAM process, again for efficiency purposes. As can be seen in Figure 3.19, this module uses the map produced by SLAM in order to match it against a reference global map provided by an orbiting satellite imagery.

The specific technique used is known as Template Matching (Jayanthi and Indu, 2016), a popular technique used in image processing. The local map is first processed to downsample it to match the global map's lower resolution. Then a gradient filter is applied to eliminate any offsets in absolute height values using a Sobel operator. At this point, the Template Matching process handles both the local and global maps as 2D images where the local (template) map image is iterated through the global (reference) map image sliding it over pixel-by-pixel and computing a matching score at each location. The Normalised Cross Correlation metric is used to compute this score which gives optimal results in cases where some pixels may have unknown elevation data. A visualisation of the map matching process is shown in Figure 3.21. The coordinates at which the




Figure 3.20: The particle cloud (red dots) of the particle filter distributed around the rover pose. The trajectory of the rover with respect to the initial state is represented by the orange line. The current local elevation map colors represent the range of height values from high (red) to low (purple).

correlation score is the highest are used for the pose correction and reduce the localisation drift.

This technique can quickly become a computationally expensive process. Template map matching requires the processing of a large amount of data, and a proper execution strategy has to be adopted to optimise its use. Two criteria are set that need to be validated before the process is executed. The first one checks that a certain threshold distance has been traversed since the last pose correction, to avoid running the process unnecessarily. The second one checks that the current local map environment is rich in elevation features so that the terrain has sufficient signatures for it to be matched with the reference map. These two criteria are meant to improve the efficiency and effectiveness respectively.

3.4.3 *Experimental Results*

The experimental validation of these two functions is performed using the dataset gathered at Teide volcano in the Canary Islands that has been presented in [Section 3.2](#). Data from the GNSS receiver is used as ground truth localisation and the maps gathered with the surveying drone are used as reference global maps. Different sensor data is used in the SLAM process. NavCam images are the main map perception input data while LocCam and IMU data are combined in the odometry input. Experiments are first run to evaluate the performance of SLAM




Figure 3.21: Graphical illustration of how template matching is performed between the local gradient image (template) and the global gradient image (reference).

by comparing its localisation pose output to the GNSS ground truth. Several traverses of different length and terrain characteristics are executed to see the robustness and sensitivity of the designed approach. In other experiments the performance of the global localisation is evaluated by triggering the map matching module once the pose output has drifted. Figure 3.22 is a 3D visualisation of the experimental process.

In Figure 3.23 one of the traverses used for the experiments to evaluate the relative localisation is shown. For each traverse taken from the dataset three independent runs are executed to give some statistical meaning of the results since the SLAM technique is partly based on a probabilistic method. The numerical results of the specific traverse shown in Figure 3.23 are gathered in Table 3.1.

The results show how the SLAM process manages to improve significantly the error of the odometry pose estimation, resulting on an average drift from ground truth of 1.34% of the total traverse length. On the reference provided at the end of this section several parameters impacting the performance results are analysed. This includes a set of experiments where the number of particles in the filter and its resampling frequency are changed to optimise the resources.




Figure 3.22: Perspective view of the *SLAM* process visualization. The trajectory of the rover is represented by the red line. The current local elevation map and global map (background) colors represent the relative separate range of height values of each map from high (red) to low (purple).



Figure 3.23: Traverses of experiments on relative localization. Black, blue and red lines represent the ground truth, odometry- and slam-estimated trajectories respectively.

As for the global localisation function, its performance is demonstrated in the experiment shown in [Figure 3.24](#). This is a relatively long traverse going over a

Trav	Run	Soil	Rocks	Slopes	Length	SLAM Error	Odom. Error
2	1	Loose	High	Med.	95m	0.82m (0.87 %)	2.21m (2.33 %)
2	2	Loose	High	Med.	95m	1.16m (1.22 %)	2.04m (2.14 %)
2	3	Loose	High	Med.	95m	1.84m (1.94 %)	2.79m (2.94 %)
2	Avg.	Loose	High	Med.	95m	1.35m (1.34 %)	2.48m (2.46 %)

Table 3.1: Summary table of results of Traverse 2

flat and somewhat featureless terrain including significant turns and changes in heading direction. This creates a noticeable drift in localisation of several meters. At the end of the traverse, where the presence of rocks and slopes increases, the global localisation module is able to find a matching correspondence with the orbital reference map and correct the rover pose.



Figure 3.24: Orthographic views of the environment and the rover's traverses in it. The black and red colors correspond to the ground truth and SLAM paths accordingly. The pose correction is visible in the end of the traverse.

A much more exhaustive set of experimental results and comprehensive analysis of the performances can be found in (Geromichalos, Azkarate, et al., 2020). The paper also describes all the underlying principles and design details of both the SLAM and Global Localisation components implementation.

REFERENCES

Geromichalos, Dimitrios, Martin Azkarate, et al. (2020). «SLAM for autonomous planetary rovers with global localization.» In: *Journal of Field Robotics* 37:5, pp. 830–847.

3.5 DESIGN OF A COMPLETE GNC ARCHITECTURE

In [Section 3.3](#) a solution for autonomous navigation is proposed targeted for benign terrain. However, planetary exploration rovers are commanded by their operations and science teams to reach more and more challenging locations motivated by the scientific interest that these can bring. As soon as more challenging and hazardous terrain is traversed the efficient navigation mode proposed could very quickly become dangerously inefficient, stopping constantly for path repairs and avoidance manoeuvres. In such terrain conditions, a more conservative approach that does not rely on a roughly safe path provided from Ground is needed. Instead, the rover should be capable of evaluating the conditions of the terrain ahead and deciding on the safest path at a local scale as it advances. In first instance, this would require the rover to have the perception means to model the environment and analyse it.

In this section we present a [GNC](#) architecture design that builds on top of the already presented efficient autonomous navigation adding a second layer or navigation mode targeted for more complex terrain where the [SLAM](#) and the global localisation functions presented in [Section 3.4](#) are used. The schematic design of the architecture can be seen in [Figure 3.25](#).

The concept behind this two layer architecture is to be flexible enough to the terrain conditions, maximising the net traverse speed and keeping always the rover safe. This [GNC](#) architecture design features two navigation modes which can be switched according to the current terrain conditions. In the lower part of the architecture the Efficient Navigation mode already described in [Section 3.3](#) is integrated. Within this mode the rover is commanded to follow a roughly safe path provided from Ground. This path is generated using global low resolution maps available from orbital imagery. As the rover moves, the hazard detector is checking for any obstacles on the way triggering the local path planner when a hazard is detected. The Full Navigation mode adds to the Efficient one the perception input from the [NavCam](#) images. These, together with the odometry inputs, are used to run the [SLAM](#) component which provides an improved relative localisation and the local elevation map of the rover surrounding to compute its traversability grid. Finally, this local elevation map is used at given

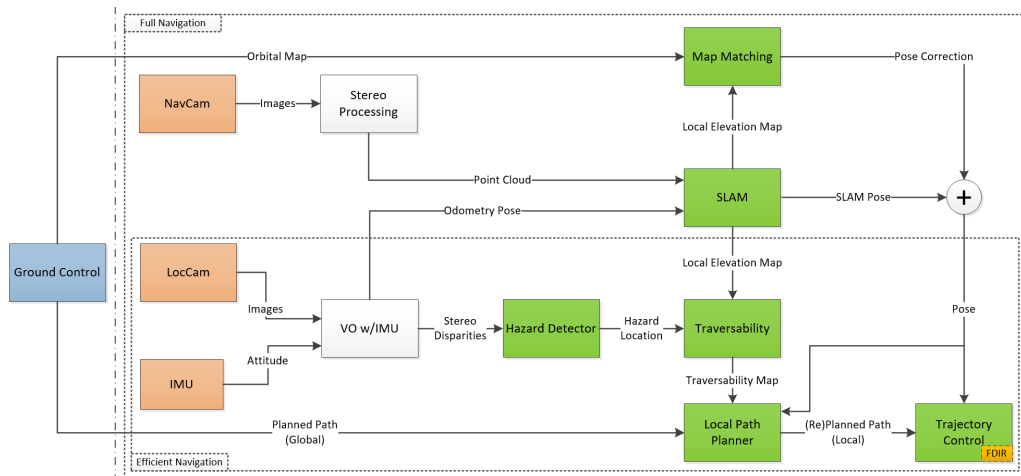


Figure 3.25: Schematic concept of the two-level GNC architecture for Autonomous Navigation of Planetary Rovers. Orange coloured boxes refer to sensors, green to developed components and greyed components to re-used from open-source community.

discrete times to match it against the orbital reference map in order to correct the pose and reduce the drift in the long range.

More details on the architecture design and applicability of the different navigation modes are given in (Azkarate et al., 2020).

REFERENCES

Azkarate, Martin et al. (2020). «A GNC Architecture for Planetary Rovers with Autonomous Navigation.» In: *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3003–3009.

3.6 CLOSING REMARKS

This closes the main core chapter of the thesis where all main contributions and publications made throughout the years of research have been briefly presented. It is believed that all contributions share the common high level objective of increasing the autonomous navigation capabilities of a planetary rover. Starting from the work previous to the thesis and progressing throughout the different developments, all contributions end up clearly supporting the definition of the final GNC architecture here presented. As such, in several occasions a former contribution is setting the foundations to enable a later one. In the next chapter,

the conclusions drawn out of all these contributions are discussed, and future proposed work is described.

CONCLUSIONS

The relevance of autonomous navigation in future planetary rover missions is acknowledged by the requirements set on these systems that are meant to accomplish more challenging tasks with less human intervention in every mission. [NASA MER](#) and [MSL](#) missions were the first ones to have implemented autonomous navigation capabilities. However, their net traverse speed was significantly reduced when running autonomously due to the amount of time needed to execute these functions. The [ESA ExoMars](#) rover mission is now embarking the first European *AutoNav* implementation, but again, the estimated distance it can cover per sol is shorter than when running on other not fully autonomous navigation modes. The [MSR](#) campaign has brought a shift of paradigm to this, and in particular the [SFR](#) mission, which requires a fast autonomous traverse of few hundreds of meters per sol in order to meet its mission target timeline. Therefore, the research work carried out in this thesis focused on studying methods to enhance the navigation capabilities of planetary rovers and explore the possibilities of making them more efficient, faster, navigating platforms. This objective is considered accomplished in this thesis. This comprised robotic systems design and integration first, followed by the development of algorithms for rover navigation, and moreover, the experimental validation of these in representative field test campaigns. Finally, the design of a full [GNC](#) architecture has been presented as well.

The knowledge and experience gained in the development and integration of three different planetary rover prototypes, and their preparation for different lab and field testing campaigns in support of [ESA](#) planetary missions have been instrumental for the research done throughout this thesis. Additional work dedicated to the execution of field test campaigns for acquisition of datasets from different sensors was also relevant for the development of algorithms for navigation. Aware of the limitations and difficulties for actual rover missions to embark capabilities for navigating autonomously, this thesis has studied and implemented novel methods for each of the core navigation functions. One of the main contributions of the thesis is the implementation and validation of

algorithms that are specifically designed for space planetary exploration and in particular the Mars rover mission scenario.

In a first step, three core components have been developed with the objective of increasing the net traverse speed of rovers in a navigation mode that makes an efficient use of resources onboard. These are: a trajectory control function based on the *Pure-Pursuit* algorithm with added safety features; a hazard detector that can effectively detect obstacles using only disparity values instead of generating a full 3D reconstruction of the scene; and a path planning function based on the *Fast Marching Method* with multi-layer and dynamic replanning capabilities. These functions are combined resulting in a lightweight navigation approach that is particularly efficient in relatively benign terrain. A field test campaign consisting of several autonomous traverses has demonstrated the significant increase in traversed distance per sol that this efficient navigation mode can bring. This is achieved by reducing the navigation stops to those strictly necessary to avoid the hazards encountered. This meaningful achievement could see an impact in the strategic planning and scientific return of real mission operations.

Once the rover is set to navigate autonomously for longer distances, a problem inherent to all dead-reckoning techniques becomes more relevant, i.e. the drift in localisation. In order to cope with the unavoidable increase on deviation and uncertainty over the rover's localisation, two more components have been developed. First, a **SLAM** component based on a mix of *Particle Filter* and *Scan Matching* techniques improves the results of commonly used relative localisation odometry functions. Additionally a local elevation map of the rover surroundings is produced that can be used to navigate effectively and smoothly across complex terrain with numerous hazards. A second component for global localisation based on *Template Matching* uses this same local elevation map to find the coordinates in the reference orbital map with highest correlation. The output corrects the localisation drift up to a range equivalent to the resolution of the orbital map grid. These two components combined together enable the possibility for long-range autonomous navigation with little supervision constraints. The evaluation of the performance and robustness of both components is done with extensive experiments using a dataset gathered in a field test campaign performed in an analogue terrain near the Teide Volcano in the Canary Islands. As one expects from the techniques and data products used, both components perform best in terrain with certain variation in elevation signatures, i.e., areas where the number of obstacles, rocks or outcrops, and slopes are more significant.

Finally, all developed components are put together in a **GNC** architecture design that features two navigation modes, each of them targeting different terrain difficulties. In areas of low to medium difficulty with relatively low **CFA**,

the efficient navigation mode based on the reactive hazard avoidance approach is used. This provides a fast navigation implementation that can significantly increase the distance traversed by a rover in a single sol. In areas with higher CFA and slopes, the full navigation mode is used which implements the SLAM and global localisation components. This allows for safer and more effectively aware navigation across more complex terrain where the straightforward reactive hazard avoidance approach would constantly be stopped and become less efficient. This is interestingly the type of terrain where the implemented SLAM and global localisation components perform best. The use of each navigation mode can be dynamically switched by onboard decision making using simple metrics that provide a sense of cost for traverse. Alternatively, this can be planned beforehand by operators from Ground making use of orbital maps that can be analysed to segment the areas by different terrain type signatures. Overall, it is concluded that the designed GNC architecture provides the needed navigation capabilities and flexibility for a rover mission with a clear ambition for fast traverse such as the SFR.

As for the proposed future work in continuation of the research line of the thesis, there are different topics that could be explored. Here, potential evolutions to the developments already done in the thesis could be mentioned, that eventually can lead to opening new lines of research.

One of the identified areas for potential enhancement is the process of traversability assessment that links the perception pipeline to the path planning component. The currently pragmatic process could be extended to include information about terrain features other than those of purely non-navigable obstacles. This could provide the path planner with an increased awareness of the terrain characteristics. By assigning an associated risk value to these features, different than the safely navigable flat area, the planner's solution could provide an alternative path with beneficial impact on the system overall wear or energy consumption.

A second improvement could be focused on the global localisation component to make it more usable for benign terrain conditions. Its performance is currently not providing complete robustness to certain terrain conditions, such as those that are found in benign terrain where the efficient navigation mode is used. If a large area of benign terrain would be encountered and traversed by a rover, the efficient navigation mode could potentially end up accumulating significant drift in localisation. In that case, using the global localisation module to correct the position would be an interesting capability. As such, and since the terrain would by definition have little elevation features (i.e. low CFA values), the component could be extended to match other features, such as the optical ones. An approach to this could be the exploitation of HiRISE plain image maps while comparing

them against projected and stitched images of the rover, i.e. Ortho-Rectified Image (ORI).

Other potential future work could be dedicated to the quantitative comparison of the efficiency of the two navigation modes in different terrain conditions. This analysis could support the proper selection of the most adequate navigation mode for each type of terrain. At the same time this could allow to identify corner cases where none of the navigation modes might be able to excel and trigger further developments of hybrid or alternative navigation modes targeting those scenarios.

The modular implementation of the work performed in the scope of this thesis, enables the aforementioned potential enhancements, and allows their verification in the same manner. The presented GNC architecture comprises such an end-to-end chain, allowing the configuration of the different functions and concepts developed in this thesis, and was used to verify these as a proof of concept.

BIBLIOGRAPHY

- Ali, Khaled S et al. (2005). «Attitude and position estimation on the mars exploration rovers.» In: *IEEE International Conference on Systems, Man and Cybernetics*. Vol. 1. IEEE, pp. 20–27.
- Amidi, Omead and Chuck E Thorpe (1991). «Integrated mobile robot control.» In: *Mobile Robots V*. Vol. 1388. International Society for Optics and Photonics, pp. 504–523.
- Andrakhanov, Anatoliy and Anton Stuchkov (2017). «Traversability estimation system for mobile robot in heterogeneous environment with different underlying surface characteristics.» In: *12th International Scientific and Technical Conference on Computer Sciences and Information Technologies (CSIT)*. Vol. 1, pp. 549–554.
- Arvidson, Raymond E et al. (2010). «Spirit Mars Rover Mission: Overview and selected results from the northern Home Plate Winter Haven to the side of Scamander crater.» In: *Journal of Geophysical Research: Planets* 115.E7.
- Avilés, M et al. (2018). «SPARTAN: Vision-Based Autonomous Navigation System for Fast Traversal Planetary Rovers.» In: *14th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS)*.
- Bailey, Tim and Hugh Durrant-Whyte (2006). «Simultaneous localization and mapping (SLAM): Part II.» In: *IEEE Robotics & Automation Magazine* 13.3, pp. 108–117.
- Biesiadecki, Jeffrey J, P Chris Leger, and Mark W Maimone (2007). «Tradeoffs between directed and autonomous driving on the mars exploration rovers.» In: *The International Journal of Robotics Research* 26.1, pp. 91–104.
- Biesiadecki, Jeffrey J and Mark W Maimone (2006). «The mars exploration rover surface mobility flight software driving ambition.» In: *IEEE Aerospace Conference*, 15–pp.
- Bora, L et al. (2017). «ExoMars rover control, localisation and path planning in an hazardous and high disturbance environment.» In: *14th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*, pp. 20–22.
- Boukas, Evangelos, Antonios Gasteratos, and Gianfranco Visentin (2018). «Introducing a globally consistent orbital-based localization system.» In: *Journal of Field Robotics* 35.2, pp. 275–298.
- Bousquet, P. (2011). «The CNES contribution to in-situ exploration: robotic implications.» In: *11th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.

- Cadena, Cesar et al. (2016). «Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age.» In: *IEEE Transactions on Robotics* 32.6, pp. 1309–1332.
- Carle, Patrick JF and Timothy D Barfoot (2010). «Global rover localization by matching lidar and orbital 3d maps.» In: *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 881–886.
- Carle, Patrick JF, Paul T Furgale, and Timothy D Barfoot (2010). «Long-range rover localization by matching LIDAR scans to orbital elevation maps.» In: *Journal of Field Robotics* 27.3, pp. 344–370.
- Carsten, Joseph et al. (2007). «Global path planning on board the mars exploration rovers.» In: *IEEE Aerospace Conference*, pp. 1–11.
- (2009). «Global planning on the mars exploration rovers: Software integration and surface testing.» In: *Journal of Field Robotics* 26.4, pp. 337–357.
- Castano, Andres et al. (2008). «Automatic detection of dust devils and clouds on Mars.» In: *Machine Vision and Applications* 19.5, pp. 467–482.
- Chen, Yuanpeng et al. (2018). «Optimization of Pure Pursuit Controller based on PID Controller and Low-pass Filter.» In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 3294–3299.
- Cheng, Yang, Mark Maimone, and Larry Matthies (2005). «Visual odometry on the Mars Exploration Rovers.» In: *IEEE International Conference on Systems, Man and Cybernetics*.
- Chiodini, Sebastiano et al. (2017). «Mars rovers localization by matching local horizon to surface digital elevation models.» In: *IEEE International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, pp. 374–379.
- Christian, John A and Scott Cryan (2013). «A survey of LIDAR technology and its use in spacecraft relative navigation.» In: *AIAA Guidance, Navigation, and Control (GNC) Conference*, p. 4641.
- Correal, Raul and Gonzalo Pajares (2011). «Onboard Autonomous Navigation Architecture for a Planetary Surface Exploration Rover and Functional Validation Using Open Source Tools.» In: *11th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.
- Coulter, R. Craig (1992). «Implementation of the Pure Pursuit Path Tracking Algorithm.» In: *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep.*
- Cozman, Fabio, Eric Krotkov, and Carlos Guestrin (2000). «Outdoor visual position estimation for planetary rovers.» In: *Autonomous Robots* 9.2, pp. 135–150.
- Daftry, Shreyansh et al. (2021). «Machine Vision based Sample-Tube Localization for Mars Sample Return.» In: *IEEE Aerospace Conference*, pp. 1–12.

- De Luca, Alessandro, Giuseppe Oriolo, and Claude Samson (1998). «Feedback control of a nonholonomic car-like robot.» In: *Robot motion planning and control*. Springer, pp. 171–253.
- Dijkstra, Edsger W (1959). «A note on two problems in connexion with graphs.» In: *Numerische mathematik* 1.1, pp. 269–271.
- Dimastrogiovanni, Mauro, Florian Cordes, and Giulio Reina (2020). «Terrain Estimation for Planetary Exploration Robots.» In: *Applied Sciences* 10.17, p. 6044.
- Dissanayake, Gamini et al. (2011). «A review of recent developments in simultaneous localization and mapping.» In: *6th International Conference on Industrial and Information Systems*, pp. 477–482.
- Durrant-Whyte, H. and T. Bailey (2006). «Simultaneous localization and mapping: part I.» In: *IEEE Robotics & Automation Magazine* 13.2, pp. 99–110.
- Ehrenfreund, Pascale et al. (1998). *Laboratory astrophysics and space research*. Vol. 80. Springer Science & Business Media.
- Elbanhawi, Mohamed and Milan Simic (2014). «Sampling-based robot motion planning: A review.» In: *Ieee access* 2, pp. 56–77.
- Ferguson, Dave, Maxim Likhachev, and Anthony Stentz (2005). «A guide to heuristic-based path planning.» In: *International workshop on planning under uncertainty for autonomous systems, international conference on automated planning and scheduling (ICAPS)*, pp. 9–18.
- Ferguson, Dave and Anthony Stentz (2005a). «The Field D* algorithm for improved path planning and replanning in uniform and non-uniform cost environments.» In: *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-05-19*.
- (2005b). «The delayed D* algorithm for efficient path replanning.» In: *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2045–2050.
- (2006). «Using interpolation to improve path planning: The Field D* algorithm.» In: *Journal of Field Robotics* 23.2, pp. 79–101.
- Furano, Gianluca et al. (2020). «Towards the Use of Artificial Intelligence on the Edge in Space Systems: Challenges and Opportunities.» In: *IEEE Aerospace and Electronic Systems Magazine* 35.12, pp. 44–56.
- Furgale, Paul, Pat Carle, and Timothy D. Barfoot (2010). «A comparison of global localization algorithms for planetary exploration.» In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4964–4969.
- Garrido, Santiago, David Álvarez, and Luis Moreno (2016). «Path planning for mars rovers using the fast marching method.» In: *Robot 2015: Second Iberian Robotics Conference*. Springer, pp. 93–105.
- Goldberg, Steven B, Mark W Maimone, and Larry Matthies (2002). «Stereo vision and rover navigation software for planetary exploration.» In: *IEEE Aerospace Conference*. Vol. 5, pp. 5–5.

- Gong, Weishu (2015). «Discussions on localization capabilities of MSL and MER rovers.» In: *Annals of GIS* 21.1, pp. 69–79.
- Gonzalez, Ramon and Karl Iagnemma (2018a). «Deepterramechanics: Terrain classification and slip estimation for ground robots via deep learning.» In: *arXiv preprint arXiv:1806.07379*.
- (2018b). «Slippage estimation and compensation for planetary exploration rovers. State of the art and future challenges.» In: *Journal of Field Robotics* 35.4, pp. 564–577.
- Graser, Evan et al. (2020). «Rimmed wheel performance on the Mars Science Laboratory Scarecrow rover.» In: *2020 IEEE Aerospace Conference*. IEEE, pp. 1–12.
- Grotzinger, John P et al. (2012). «Mars Science Laboratory mission and science investigation.» In: *Space science reviews* 170.1, pp. 5–56.
- Guinn, Joseph R (2001). «Mars surface asset positioning using in-situ radio tracking.» In:
- Gupta, S et al. (2014). «Fusion and Visualization of HiRISE Super-Resolution, Shape-from-Shading DTM with MER Stereo 3D Reconstructions.» In: *AGU Fall Meeting Abstracts*. Vol. 2014, P41D–3955.
- Hart, Peter E., Nils J. Nilsson, and Bertram Raphael (1968). «A Formal Basis for the Heuristic Determination of Minimum Cost Paths.» In: *IEEE Transactions on Systems Science and Cybernetics* 4.2, pp. 100–107.
- Helmick, D.M. et al. (2004). «Path following using visual odometry for a Mars rover in high-slip environments.» In: *IEEE Aerospace Conference*. Vol. 2, 772–789 Vol.2.
- Helmick, Daniel M et al. (2006). «Slip-compensated path following for planetary exploration rovers.» In: *Advanced Robotics* 20.11, pp. 1257–1280.
- Hidalgo-Carrio, Javier, Ajish Babu, and Frank Kirchner (2014). «Static forces weighted Jacobian motion models for improved Odometry.» In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 169–175.
- Hidalgo-Carrió, Javier, Pantelis Poulakis, and Frank Kirchner (2018). «Adaptive localization and mapping with application to planetary rovers.» In: *Journal of Field Robotics* 35.6, pp. 961–987.
- Hidalgo-Carrió, Javier et al. (2017). «Gaussian process estimation of odometry errors for localization and mapping.» In: *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5696–5701.
- Hwangbo, Ju Won, Kaichang Di, and Rongxing Li (2009). «Integration of orbital and ground image networks for the automation of rover localization.» In: *ASPRS Annual Conference*. Citeseer.

- Jayanthi, N and S Indu (2016). «Comparison of image matching techniques.» In: *International Journal of Latest Trends in Engineering and Technology* 7.3, pp. 396–401.
- Koenig, Sven and Maxim Likhachev (2002a). «D^{*} lite.» In: *AAAI Innovative Applications of Artificial Intelligence Conferences (IAAI)* 15.
- (2002b). «Improved fast replanning for robot navigation in unknown terrain.» In: *IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 1, pp. 968–975.
- Kostavelis, Ioannis et al. (2011). «SPARTAN system: Towards a low-cost and high-performance vision architecture for space exploratory rovers.» In: *IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*.
- Kostavelis, Ioannis et al. (2014). «SPARTAN: Developing a vision system for future autonomous space exploration robots.» In: *Journal of Field Robotics* 31.1, pp. 107–140.
- Kostavelis, Ioannis et al. (2016). «Stereo-based visual odometry for autonomous robot navigation.» In: *International Journal of Advanced Robotic Systems* 13.1, p. 21.
- LaValle, Steven M (2006). *Planning algorithms*. Cambridge university press.
- Lentaris, G. et al. (2018). «High-Performance Embedded Computing in Space: Evaluation of Platforms for Vision-Based Navigation.» In: *Journal of Aerospace Information Systems*.
- Lentaris, G. et al. (2019). «Single- and multi-FPGA Acceleration of Dense Stereo Vision for Planetary Rovers.» In: *Transactions on Embedded Computing Systems*.
- Lentaris, George et al. (2015). «Spartan/Sextant/Compass: Advancing Space Rover Vision via Reconfigurable Platforms.» In: *Applied Reconfigurable Computing*. Vol. 9040. Springer, pp. 475–486.
- Li, Rongxing et al. (2004). «Rover localization and landing-site mapping technology for the 2003 Mars exploration rover mission.» In: *Photogrammetric Engineering & Remote Sensing* 70.1, pp. 77–90.
- Likhachev, Maxim et al. (2008). «Anytime search in dynamic graphs.» In: *Artificial Intelligence* 172.14, pp. 1613–1643. ISSN: 0004-3702.
- Mac, Thi Thoa et al. (2016). «Heuristic approaches in robot path planning: A survey.» In: *Robotics and Autonomous Systems* 86, pp. 13–28. ISSN: 0921-8890.
- Maimone, M. (2017). «The Evolution of Autonomous Capabilities on NASAs Mars Rovers.» In: *Southern California Robotics Symposium*.
- Maimone, Mark W, P Chris Leger, and Jeffrey J Biesiadecki (2007). «Overview of the mars exploration rovers' autonomous mobility and vision capabilities.» In: *IEEE International Conference on Robotics and Automation (ICRA) - Space Robotics Workshop*.

- Maimone, Mark, Yang Cheng, and Larry Matthies (2007). «Two years of Visual Odometry on the Mars Exploration Rovers.» In: *Journal of Field Robotics*.
- Maki, J et al. (2012). «The Mars science laboratory engineering cameras.» In: *Space science reviews* 170.1-4, pp. 77–93.
- Maki, JN et al. (2020). «The Mars 2020 engineering cameras and microphone on the perseverance rover: A next-generation imaging system for Mars exploration.» In: *Space Science Reviews* 216.8, pp. 1–48.
- Malinowski, Mateusz (2019). «Hardware Accelerated Visual Localisation For Next Generation Mars Rovers.» In: *15th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.
- Mandelbaum, R. et al. (1998). «Real-time stereo processing, obstacle detection, and terrain estimation from vehicle-mounted stereo cameras.» In: *4th IEEE Workshop on Applications of Computer Vision (WACV)*, pp. 288–289.
- Marc, R. and P. Weclowski (2018). «Autonomous Multi-Mode Rover Navigation for Long-Range Planetary Exploration Using Orbital and Locally Perceived Data.» In: *International Astronautical Congress (IAC)*.
- Matthies, L. et al. (2007). «Computer Vision on Mars.» In: *International Journal of Computer Vision*.
- Matthies, Larry et al. (1995). «Mars microrover navigation: Performance evaluation and enhancement.» In: *Autonomous robots* 2.4, pp. 291–311.
- McEwen, A. S. et al. (2007). «Mars Reconnaissance Orbiter’s High Resolution Imaging Science Experiment (HiRISE).» In: *Journal of Geophysical Research: Planets* 112.E5.
- McManamon, Kevin, Richard Lancaster, and Nuno Silva (2013). «Exomars rover vehicle perception system architecture and test results.» In: *12th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*, pp. 15–17.
- Miles, Helen C., Matthew D. Gunn, and Andrew J. Coates (2020). «Seeing Through the “Science Eyes” of the ExoMars Rover.» In: *IEEE Computer Graphics and Applications* 40.2, pp. 71–81.
- Mishkin, Andrew H et al. (1998). «Experiences with operations and autonomy of the mars pathfinder microrover.» In: *IEEE aerospace conference proceedings (Cat. No. 98TH8339)*. Vol. 2, pp. 337–351.
- Moreno, Iñigo, Martin Azkarate, and Letizia Marchegiani (2021). *Fusion of velocity and relative pose measurements for localization of extraterrestrial rovers*. Tech. rep. European Space Agency.
- Moreno, Iñigo et al. (2021). *Improving Global Localization Algorithms for Mars Rovers with Neural Networks*. Tech. rep. European Space Agency.
- Moreno, S. (2013). «CNES robotics activities: towards long distance OB decision-making navigation.» In: *12th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.

- Murphy, Kevin P. (1999). «Bayesian Map Learning in Dynamic Environments.» In: *12th International Conference on Neural Information Processing Systems*.
- Naguib, Ahmed M et al. (2018). «Planetary Long-Range Deep 2D Global Localization Using Generative Adversarial Network.» In: *Journal of Korea Robotics Society* 13.1, pp. 26–30.
- Nilsson, Nils J (1980). «Principles of artificial intelligence.» In: *Tioga Pub Co., Palo Alto, CA* 476.
- Núñez, Pedro, Ricardo Vázquez-Martín, and Antonio Bandera (2011). «Visual Odometry Based on Structural Matching of Local Invariant Features Using Stereo Camera Sensor.» In: *Sensors* 11.7, pp. 7262–7284. ISSN: 1424-8220.
- Oettershagen, P. et al. (2019). «Investigation of specific wheel-terrain interaction aspects using and advanced Single Wheel Testbed test facility.» In: *15th Advanced Space Technologies for Robotics and Automation (ASTRA)*.
- Ollero, A and O Amidi (1991). «Predictive path tracking of mobile robots. Application to the CMU Navlab.» In: *5th International Conference on Advanced Robotics, Robots in Unstructured Environments (ICAR)*. Vol. 91, pp. 1081–1086.
- Ono, Masahiro et al. (2015). «Risk-aware planetary rover operation: Autonomous terrain classification and path planning.» In: *IEEE aerospace conference*, pp. 1–10.
- Patz, B. J. et al. (2008). «A practical approach to robotic design for the DARPA Urban Challenge.» In: *Journal of Field Robotics*.
- Peng, Huei and Masayoshi Tomizuka (1991). «Optimal preview control for vehicle lateral guidance.» In: *California Partners for Advanced Transit and Highways (PATH)*.
- Pham, T.H. et al. (May 2021). «Rover Relocalization for Mars Sample Return by Virtual Template Synthesis and Matching.» In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*.
- Pitkänen, Ville et al. (2017). «Path following controller for planar robots with articulated, actuated and independently steerable velocity-limited wheels.» In: *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2433–2440.
- Poulakis, Pantelis et al. (2015). «Overview and Development Status of the ExoMars Rover Mobility Subsystem.» In: *13th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.
- Rajamani, Rajesh (2011). *Vehicle dynamics and control*. Springer Science & Business Media.
- Rankin, Arturo et al. (2020). «Driving Curiosity: Mars Rover Mobility Trends During the First Seven Years.» In: *IEEE Aerospace Conference*, pp. 1–19.
- (2021). «Mars curiosity rover mobility trends during the first 7 years.» In: *Journal of Field Robotics*.

- Rusu, Alexandru et al. (2013). «Onboard Autonomous Path Planning for a Planetary Exploration Rover.» In: *AIAA Infotech@ Aerospace (I@A) Conference*, p. 4825.
- Samuel, Moveh, Mohamed Hussein, and Maziah Binti Mohamad (2016). «A review of some pure-pursuit based path tracking techniques for control of autonomous vehicle.» In: *International Journal of Computer Applications* 135.1, pp. 35–38.
- Sethian, James A (1999). «Fast marching methods.» In: *SIAM review* 41.2, pp. 199–235.
- Shaukat, A. et al. (2015). «Next-Generation Rover GNC Architectures.» In: *13th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.
- Shaw, Andrew et al. (2013). «Robust Visual Odometry for Space Exploration.» In: *12th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.
- Sheshadri, Aashish et al. (2012). «Position estimation by registration to planetary terrain.» In: *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pp. 432–438.
- Silva, N. et al. (2013). «ExoMars Rover Vehicle Mobility functional Architecture and Key Design Drivers.» In: *12th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.
- Snider, J. M. (2009). «Automatic Steering Methods for Autonomous Automobile Path Tracking.» In: *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep.*
- Stachniss, Cyrill, John J Leonard, and Sebastian Thrun (2016). «Simultaneous localization and mapping.» In: *Springer Handbook of Robotics*, pp. 1153–1176.
- Stentz, Anthony and Martial Hebert (1995). «A complete navigation system for goal acquisition in unknown environments.» In: *Autonomous Robots* 2.2, pp. 127–145.
- Sun, Xiaoxun, William Yeoh, and Sven Koenig (2009). «Efficient incremental search for moving target search.» In: *21st international joint conference on Artificial intelligence*, pp. 615–620.
- Tao, Y and J-P Muller (2016). «A novel method for surface exploration: super-resolution restoration of Mars repeat-pass orbital imagery.» In: *Planetary and Space Science* 121, pp. 103–114.
- Tao, Yu, Jan-Peter Muller, and William Poole (2016). «Automated localisation of Mars rovers using co-registered HiRISE-CTX-HRSC orthorectified images and wide baseline Navcam orthorectified mosaics.» In: *Icarus* 280, pp. 139–157. ISSN: 0019-1035.
- Thrun, Sebastian, Wolfram Burgard, and Dieter Fox (2005). *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press.

- Thrun, Sebastian et al. (2006). «Stanley: The robot that won the DARPA Grand Challenge.» In: *Journal of field Robotics* 23,9, pp. 661–692.
- Toupet, Olivier et al. (2020). «Terrain-adaptive wheel speed control on the Curiosity Mars rover: Algorithm and flight results.» In: *Journal of Field Robotics* 37,5, pp. 699–728.
- Townson, Daniel, Mark Woods, and Stuart Carnochan (2018). «ExoMars VisLoc - The Industrialised, Visual Localisation System for the ExoMars Rover.» In: *14th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (ISAIRAS 2018)*.
- Trawny, Nikolas and Stergios I Roumeliotis (2005). «Indirect Kalman filter for 3D attitude estimation.» In: *University of Minnesota, Dept. of Comp. Sci. & Eng., Tech. Rep 2*, p. 2005.
- Urmson, Chris et al. (2006). «A robust approach to high-speed navigation for unrehearsed desert terrain.» In: *Journal of Field Robotics* 23,8, pp. 467–508.
- Urmson, Chris et al. (2008). «Autonomous driving in urban environments: Boss and the urban challenge.» In: *Journal of Field Robotics* 25,8, pp. 425–466.
- Van Pham, Bach, Artur Maligo, and Simon Lacroix (2013). «Absolute map-based localization for a planetary rover.» In: *12th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*, pp. 1–8.
- Volpe, Richard et al. (2000). «Enhanced mars rover navigation techniques.» In: *IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 1, pp. 926–931.
- Ward, R. et al. (Aug. 2016). «Exomars VisLoc - The Visual Localisation System for the Exomars Rover.» In: *DASIA 2016 - Data Systems In Aerospace*. Ed. by L. Ouwehand. Vol. 736. ESA Special Publication, p. 17.
- Weclowski, Piotr and Róbert Marc (2019). «Capabilities of Long Range Autonomous Multi-Mode Rover Navigation System - ERGO Field Trials and Planned Evolution.» In: *15th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.
- Wiese, Tim (2017). «3D Kinematic Modeling and Evaluation of Rough-Terrain Locomotion Modes for an ExoMars-like Mobility Subsystem.» In: *Technische Universität München. MSc Thesis*.
- Williamson, T. and C. Thorpe (June 1998). «A specialized multibaseline stereo technique for obstacle detection.» In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 238–244.
- Winter, M. et al. (2015). «ExoMars Rover GNC System Design: a Detailed Description.» In: *13th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.

- Winter, M. et al. (2017). «Detailed description of the High Level Autonomy Functionalities Developed for the ExoMars Rover.» In: *14th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.
- Wu, Benjamin et al. (2019). «Absolute Localization Through Orbital Maps and Surface Perspective Imagery: A Synthetic Lunar Dataset and Neural Network Approach.» In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3262–3267.
- Yeomans, Brian et al. (2017). «MURFI 2016-From Cars to Mars: Applying autonomous vehicle navigation methods to a space rover mission.» In: *14th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*.
- Zhang, Zhengyou (1994). «Iterative point matching for registration of free-form curves and surfaces.» In: *International Journal of Computer Vision*.

Part II


MAIN AUTHORED PUBLICATIONS

EFFICIENT AUTONOMOUS NAVIGATION FOR PLANETARY ROVERS WITH LIMITED RESOURCES

L. Gerdes and M. Azkarate *et al.*

Journal of Field Robotics, 25 August 2020

DOI: 10.1002/rob.21981



ABSTRACT


Rovers operating on Mars require more and more autonomous features to fulfill their challenging mission requirements. However, the inherent constraints of space systems render the implementation of complex algorithms an expensive and difficult task. In this paper, we propose an architecture for autonomous navigation. Efficient implementations of autonomous features are built on top of the *ExoMars path following* navigation approach to enhance the safety and traversing capabilities of the rover. These features allow the rover to detect and avoid hazards and perform significantly longer traverses planned by operators on the ground. The efficient navigation approach has been implemented and tested during field test campaigns on a planetary analogue terrain. The experiments evaluated the proposed architecture by autonomously completing several traverses of variable lengths while avoiding hazards. The approach relies only on the optical Localisation Cameras (*LocCam*) stereobench, a sensor that is found in all current rovers, and potentially allows for computationally inexpensive long-range autonomous navigation in terrains of medium difficulty.

SLAM FOR AUTONOMOUS PLANETARY ROVERS WITH GLOBAL LOCALIZATION

D. Geromichalos and M. Azkarate *et al.*

Journal of Field Robotics, 28 February 2020

DOI: 10.1002/rob.21943



ABSTRACT

This paper describes a novel approach to **SLAM** techniques applied to the autonomous planetary rover exploration scenario in order to reduce both the relative and absolute localization errors, using two well proven techniques: particle filters and scan matching. Continuous relative localization is improved by matching high resolution sensor scans to the online created local map. Additionally, in order to avoid issues with drifting localization, absolute localization is globally corrected at discrete times, according to predefined event criteria, by matching the current local map to the orbiter's global map. The resolutions of local and global maps can be appropriately chosen for computation and accuracy purposes. Further, the online generated local map, of the form of a structured elevation grid map, can also be used to evaluate the traversability of the surrounding environment and allow for continuous navigation. The objective of this work is to support long-range low-supervision planetary exploration. The implemented **SLAM** technique has been validated with a dataset acquired during a field test campaign performed at the Teide Volcano on the island of Tenerife, representative of a Mars/Moon exploration scenario.

A GNC ARCHITECTURE FOR PLANETARY ROVERS WITH AUTONOMOUS NAVIGATION

M. Azkarate *et al.*

IEEE © International Conference on Robotics and Automation (ICRA), 31 May 2020

DOI: 10.1109/ICRA40945.2020.9197122



ABSTRACT

This paper proposes a **GNC** architecture for planetary rovers targeting the conditions of upcoming Mars exploration missions such as Mars 2020 and the **SFR**. The navigation requirements of these missions demand a control architecture featuring autonomous capabilities to achieve a fast and long traverse. The proposed solution presents a two-level architecture where the efficient navigation (lower) level is always active and the full navigation (upper) level is enabled according to the difficulty of the terrain. The first level is an efficient implementation of the basic functionalities for autonomous navigation based on hazard detection, local path replanning, and trajectory control with visual odometry. The second level implements an adaptive SLAM algorithm that improves the relative localization, evaluates the traversability of the terrain ahead for a more optimal path planning, and performs global (absolute) localization that corrects the pose drift during longer traverses. The architecture provides a solution for long-range, low supervision, and fast planetary exploration. Both navigation levels have been validated on planetary analog field test campaigns.

DESIGN, TESTING AND EVOLUTION OF MARS ROVER TESTBEDS FOR ESA PLANETARY EXPLORATION

M. Azkarate et al.

IEEE © Robotics & Automation Magazine, Submitted 1 May 2021

Accepted for Publication



ABSTRACT

This paper presents the system architecture and design of two planetary rover laboratory testbeds developed at the [ESA](#). These research platforms have been developed to provide early prototypes for validation of designs and serve [ESA's](#) Automation & Robotics Lab infrastructure for continuous research and testing. Both rovers have been built considering the constraints of Space Systems with the sufficient level of representativeness to allow rapid prototyping. They avoid strictly space-qualified components and designs that present a major cost burden and frequently lack the flexibility or modularity that the lab environment requires for its investigations. This design approach is followed for all the mechanical, electrical, and software aspects of the system. In this paper, two ExoMars mission-representative rovers, the [ExoTeR](#) and the [MaRTA](#), are thoroughly described. The lessons learnt and experience gained while running several research activities and test campaigns are also presented. Finally, the paper aims to give an insight on how to reduce the gap between lab R&D and flight implementations by anticipating system constraints when building these platforms. This allows to provide qualitative testing results that can eventually have an impact in real space missions.

Part III

RESUMEN

Brief summary of the thesis dissertation written in Spanish.

RESUMEN

SUMARIO

A partir de la experiencia obtenida en las misiones de rovers planetarios, como los Mars Exploration Rovers (**MER**) o Curiosity, y en vista de las futuras misiones de Exobiology on Mars (**ExoMars**) y el programa Mars Sample Return (**MSR**), esta tesis tiene como objetivo estudiar y mejorar las capacidades a bordo de los rovers para incrementar el retorno científico general de la misión. A medida que vemos misiones espaciales con mayores capacidades y funciones embarcadas, los rovers planetarios dependerán de las funcionalidades de navegación autónoma para cumplir con los objetivos de su misión. Las restricciones dadas por el entorno espacial y las limitaciones en las comunicaciones, que se dan en las misiones de exploración a Marte, hacen que la implementación de capacidades de navegación autónoma sea la solución más eficiente, o tal vez la única, para extender la distancia recorrida por sol (día marciano) más allá de las pocas decenas de metros.

La navegación autónoma es una capacidad compleja que se basa en la implementación de varias funcionalidades y la ejecución orquestada de éstas para su desempeño. Esas funcionalidades son, como mínimo, Localización, Percepción o Mapeo, Planificación de Ruta y Control de Trayectoria. En esta tesis se investigan todas estas funciones robóticas esenciales y se proponen nuevos métodos que se centran en el caso particular de los rovers planetarios. Éstos tienen en consideración las limitaciones y condiciones que se encuentran en las misiones a Marte, lo que las hace particularmente adecuadas y específicas para estos sistemas rover.

Se han estudiado diferentes formas de combinar y ejecutar estas funcionalidades, que en conjunto componen el subsistema de Guiado, Navegación y Control (**GNC**) de un rover que es capaz de navegar de forma autónoma. Se toma en consideración la potencial variación en las condiciones del terreno encontradas en una misión rover, adaptando el comportamiento y las funcionalidades ejecutadas por la arquitectura de control, con el objetivo de maximizar la distancia recorrida

por sol. Se propone un modo de navegación enfocado a los casos de terreno de dificultad media o baja. Este modo se basa únicamente en una cámara estéreo, que imita al sensor [LocCam](#) de [ExoMars](#), para localizar el rover y evitar los obstáculos a lo largo del camino. Esto se consigue sin requerir que el rover se pare periódicamente, lo que aumenta la velocidad efectiva. El modo de navegación se valida de forma experimental en una exhaustiva campaña de pruebas de campo que acumula más de 500 m de recorrido autónomo. Para los casos de terreno más dificultoso, se propone un modo de navegación que, basándose en el primer modo, añade nuevas capacidades al rover. Éste integra un nuevo componente de Localización y Mapeo Simultaneos ([SLAM](#)) desarrollado para rovers planetarios que permite mejorar la precisión de su localización y una detección más lejana de los peligros del terreno para poder eviatarlos de forma más fluida. El mapa producido por la función [SLAM](#) se utiliza posteriormente en un componente de Localización Global para corregir la deriva acumulada en travesías de largo alcance. Estos dos componentes son validados experimentalmente utilizando un conjunto de datos representativos adquiridos en una campaña de prueba de campo llevada a cabo en un terreno planetario análogo cerca del volcán Teide en Tenerife (Islas Canarias).

Finalmente, la tesis propone el diseño de la arquitectura [GNC](#) completa, que integra ambos modos de navegación y podría tener aplicación en futuras misiones a Marte como el Sample Fetch Rover ([SFR](#)) del programa [MSR](#).

INTRODUCCIÓN

Exploración Planetaria: Misión a Marte

Gracias a la exploración robótica planetaria, hoy en día los seres humanos han llegado a lugares remotos que de otra manera no podrían imaginar poder vislumbrar. Mientras emprendedores ricos y audaces compiten y prometen que pronto nos convertiremos en una especie multi-planetaria, podríamos decir que, en realidad, hace mucho tiempo que cruzamos esa línea. Antes de que los humanos lleguen a poner un pie en Marte, nuestro conocimiento del planeta rojo será tal que incluso podríamos saber que alguna vez en el pasado existió vida en Marte. Además, en cuestión de una década, podríamos estar a punto de traer una muestra de suelo Marciano a la Tierra. La misión [ExoMars](#) de la Agencia Espacial Europea ([ESA](#)) por un lado, y su colaboración con la National Aeronautics and Space Administration ([NASA](#)) en el programa Mars Sample Return ([MSR](#)) por otro, están allanando el camino para conseguir estos objetivos. De forma similar, misiones anteriores como Mars Science Laboratory ([MSL](#)), Mars




Figura E.1: Representación conceptual del programa Mars Sample Return (creditos: ESA)

Exploration Rover (MER), Phoenix o el pionero Mars Pathfinder, han permitido concebir una misión tan compleja como la MSR (ver Figura E.1).

No hay duda del valor inestimable de la ciencia que las misiones planetarias aportan a la sociedad, y del conocimiento que hemos adquirido sobre Marte desde que los módulos de aterrizaje de la misión Viking transmitieron sus primeras imágenes del planeta rojo. Misión tras misión, varias naves espaciales han contribuido, y continúan haciéndolo, al aprendizaje de la geología marciana y su potencial habitabilidad a partir de satélites en órbita o sondas de aterrizaje y finalmente rovers exploratorios. Con el tiempo, la exploración robótica actúa como un precursor de la exploración humana, más que como un reemplazo de ésta. Los instrumentos y mecanismos que llevan los rovers planetarios permiten realizar ciencia in situ antes de cualquier posible llegada humana y ayudan a preparar mejor los desafíos que podría implicar una futura exploración tripulada.

A continuación, se expone cómo se llevan a cabo estos primeros experimentos in situ y cómo se diseñan los rovers planetarios para que cumplan sus objetivos de misión y ejecuten sus tareas.

Recorriendo Marte: Los desafíos del planeta rojo

Tan difícil y emocionante pueden parecer las fases de Entrada, Descenso y Aterrizaje de una misión a Marte como los desafíos que vienen después de un aterrizaje seguro. Según explica Steve W. Squyres en su libro con el mismo título, *Recorrer Marte* es la gran aventura. Su entorno hostil con una atmósfera

delgada (1 % de la presión atmosférica de la Tierra) hace que la temperatura de la superficie oscile hasta 100 grados durante un día. Al mismo tiempo, la falta de un campo magnético expone la superficie del planeta directamente a toda la radiación cósmica proveniente del espacio, lo que lo hace aún más inhóspito. Además, hay clima en Marte y, sorprendentemente, el viento puede generar remolinos o incluso crear tormentas de polvo que aumentan la profundidad óptica atmosférica a valores tan altos que el flujo solar apenas se nota en la superficie marciana. Teniendo en cuenta que ésta es la principal fuente de energía de la mayoría de los rovers planetarios, excepto los que funcionan con un Generador Termoeléctrico de Radioisótopos nuclear (RTG), las tormentas de polvo son en realidad largas noches oscuras para los rovers que pueden poner fin a su misión. Además, la lejanía es probablemente el desafío más limitante, ya que la distancia de 2.66 au entre Marte y la Tierra (en sus coordenadas orbitales más lejanas) significa que una señal electromagnética puede tardar hasta 22 minutos en propagarse entre los dos planetas. Asimismo, una señal de este tipo requeriría una antena grande y potente de la que los rovers generalmente carecen debido a las restricciones del sistema en cuanto a masa y capacidad de potencia. En su lugar, los rovers dependen de uno o varios satélites en órbita de Marte para transmitir sus comunicaciones con la Tierra. Sin embargo, los rovers y sus compañeros en órbita no tienen visibilidad en todo momento y, por lo tanto, las ventanas de comunicación solo ocurren pocas veces al día, si las hay. Esto significa que los rovers deben pasar la mayor parte del día realizando sus tareas "por su cuenta". Finalmente, la superficie marciana es el último desafío a superar. Como se ve en la imagen del paisaje de la [Figura E.2](#), para poder navegar, los rovers necesitan encontrar su camino a través de obstáculos, pendientes y una combinación de características de terreno rocoso y arenoso. Éstos pueden llevar al rover a situaciones límite, como trampas en la tierra hechas de arena muy suelta, que solo pueden notarse cuando sea demasiado tarde.

En general, los rigurosos procesos de ingeniería que producen diseños que cumplen con todas esas duras restricciones han creado sistemas capaces de resistir el entorno marciano y moverse en él. El desafío de la ingeniería de sistemas es ampliamente reconocido en la robótica planetaria y la acomodación de todos los subsistemas y funcionalidades dentro de los límites ajustados de masa y potencia requiere de una obra maestra de ingeniería para ser resuelto. Sin embargo, el éxito que las misiones de exploración a Marte han tenido hasta ahora en proporcionar los datos científicos para los que fueron diseñadas, es debido principalmente a las características y capacidades autónomas a bordo de los rovers.




Figura E.2: Paisaje Marciano (creditos: [NASA](#))

Incrementando el retorno científico mediante la Autonomía

Un rover planetario es como una nave espacial en la superficie de un planeta. Como tal, se compone de varios subsistemas que se encuentran típicamente en cualquier nave espacial, como la generación, gestión y almacenamiento de energía, el control térmico, las estructuras y mecanismos, la protección contra la radiación, las comunicaciones y el sistema de manejo y almacenamiento de datos. Y, por supuesto, habría que añadir a todo ello los subsistemas de movilidad y navegación más relacionados con la robótica. La telemetría interna permite monitorear el estado de salud del sistema general en tierra. Sin embargo, dadas las limitaciones en las comunicaciones, es obligatorio tener un Sistema de Gestión de Misión a bordo que monitorice los recursos disponibles y pueda reaccionar a los eventos para garantizar la seguridad del rover y evitar que ejecute cualquier acción que pueda ponerlo en riesgo. A pesar de la implementación de estas funciones de autocontrol, todavía no podemos definir a los rovers como sistemas autónomos, ya que aún estamos lejos de ver un rover planetario con capacidades deliberativas embarcadas que puedan cambiar el curso de acción en sus tareas diarias. No obstante, vemos una creciente dependencia de las capacidades autónomas del rover para lograr los objetivos de la misión, lo que es una conclusión directa de los objetivos científicos para los que están contruidos.

Los rovers a menudo pueden verse como las naves de la carga útil científica o instrumentos que transportan, porque a diferencia de los módulos de aterrizaje, pueden atravesar la superficie planetaria para alcanzar objetivos científicamente interesantes que difícilmente serían accesibles desde el punto de aterrizaje

estático. Por lo tanto, cuanto más eficiente sea la travesía de un rover, mayor será el potencial retorno científico de la misión en general.

Por todas estas razones, la navegación autónoma está destinada a desempeñar un papel clave en las futuras misiones de rovers planetarios.

Navegación autónoma en misiones planetarias

Existen muchos ejemplos de aplicaciones de sistemas de navegación autónomos en la Tierra. Los [DARPA Grand y Urban Challenges](#) ya demostraron hace más de una década la viabilidad de los vehículos de conducción autónoma. Hoy en día, muchas empresas de alta tecnología en la industria de automoción están desarrollando vehículos totalmente autónomos capaces de transportar personas de manera segura en áreas de tráfico urbano. Pero, y a pesar de estos avances tecnológicos, rara vez encontramos ejemplos de sistemas de navegación autónoma en el espacio. No fue hasta hace poco que los rovers [MER](#) y [Curiosity](#) explotaron algunas capacidades y características de navegación autónoma, aunque solamente en escenarios limitados. Las próximas misiones están definiendo los requisitos de sus sistemas para incluir capacidades de navegación autónoma a bordo, pero aún así, los avances son lentos. Los desafíos y las limitaciones antes mencionados, impuestos por el entorno y los sistemas espaciales, suelen hacer que las soluciones adoptadas para las aplicaciones terrestres sean inviables. Los rovers en Marte no tienen conexión a servidores de datos ni a ningún Sistema de Navegación Global por Satélite ([GNSS](#)) como los que disponemos en la Tierra. Además, los rovers están atravesando superficies complejas y no estructuradas, como la que se muestra en la [Figura E.2](#), con conocimiento del terreno limitado, deducido de las imágenes de los satélites orbitadores. Aunque estas imágenes son útiles para la evaluación de posibles lugares de aterrizaje y zonas peligrosas, no pueden ofrecer el mismo nivel de detalle que ofrecen los mapas de ciudades de la Tierra. En consecuencia, habrá obstáculos que serán demasiado pequeños para ser capturados por los mapas del orbitador y, sin embargo, demasiado grandes para que el rover los supere y, por lo tanto, deben ser detectados y evitados por el rover de forma segura.

Otra consecuencia importante de las duras condiciones del entorno espacial, como el vacío, las temperaturas extremas y la radiación es que han creado una brecha tecnológica permanente para los componentes electrónicos y la aviónica disponibles para el espacio. Los últimos procesadores cualificados para el espacio son tan poco capaces como una computadora de hace 10 años en la Tierra. Además, los sensores como los escaners [LiDAR](#), ubicados en muchos de los coches autónomos en la Tierra, aún no han sido cualificados para el espacio, ya que sus espejos giratorios son incapaces de soportar las vibraciones

y shocks estructurales que suceden durante el lanzamiento. Las limitaciones en la masa y energía del sistema también juegan un papel importante en el diseño del sistema donde la selección de componentes debe ser altamente optimizada. Finalmente, debido a la naturaleza crítica y las oportunidades únicas que representan las misiones espaciales, se deben cumplir los estrictos requisitos de Fiabilidad, Disponibilidad, Mantenimiento y Seguridad (RAMS) durante todo su desarrollo. En particular, todo el software de vuelo se prueba exhaustivamente con requisitos RAMS y muchas de las librerías C++ de uso común que se ejecutan en aplicaciones de la Tierra no superan dichos tests. Todas estas restricciones hacen que la implementación de funciones de navegación autónoma y algoritmos de visión por ordenador (CV) para el espacio sea una tarea muy desafiante, donde la eficiencia y el coste computacional son dos de los factores más importantes que marcan su diseño.

CONTRIBUCIONES DE LA TESIS

En esta tesis se han explorado las posibilidades de aumentar el rendimiento y la eficiencia de las diferentes funcionalidades que se ejecutan a bordo de los rovers que se construyen para la navegación autónoma y se ha tenido como objetivo encontrar soluciones que aumenten el rendimiento general de la locomoción y la travesía de los mismos. Se ha enfocado tanto al diseño de algoritmos novedosos como a su validación experimental.

Las principales aportaciones conseguidas a lo largo de esta tesis se pueden resumir en:

- **Diseño e integración de sistemas robóticos para el I+D en exploración planetaria:**

Con el fin de facilitar los estudios esenciales de esta tesis y proporcionar plataformas de prueba donde se pueda realizar la validación experimental, una parte significativa del trabajo se centró en el desarrollo de prototipos y plataformas rover que podrían respaldar esta investigación, incluyendo la preparación y acondicionamiento final de las plataformas para realizar campañas de prueba de campo específicas. Como parte del mismo, se prestó especial atención al estudio y diseño de la cadena cinemática del chasis móvil y al control de la misma con el fin de mejorar las prestaciones locomotrices de los rovers y por tanto aumentar sus capacidades de navegación.

- **Adquisición y publicación de conjuntos de datos de sensores reales:**

La validación experimental fuera de entornos simulados se considera relevante para esta tesis. A su vez, se reconoce el esfuerzo que implica la ejecución de pruebas en plataformas reales, especialmente cuando se realizan fuera del entorno de laboratorio. Por lo tanto, la adquisición de un conjunto de datos de sensores reales realizada durante las campañas de prueba de campo en un terreno análogo representativo se considera una contribución valiosa de esta tesis. Estos conjuntos de datos no sólo han estado disponibles para el desarrollo de algoritmos en esta tesis, sino que también se han publicado y puesto a disposición de la comunidad.

- **Investigación de nuevos métodos para las funcionalidades que posibilitan la navegación autónoma en exploración planetaria:**

El núcleo del trabajo de esta tesis se centra en el estudio de los diseños algorítmicos y sus enfoques de diferentes funcionalidades que se ejecutan a bordo del rover relacionadas con su navegación. Se investigaron soluciones para el control de trayectorias, la planificación de rutas, la localización y la percepción, y se han desarrollado e introducido métodos novedosos para cada una de estas funciones enfocadas particularmente en las condiciones encontradas en las misiones de exploración con rovers planetarios a Marte.

- **Validación experimental de algoritmos implementados en condiciones análogas reales:**

Un objetivo importante de la tesis se centró en la validación del rendimiento de los algoritmos desarrollados en campañas de prueba representativas. Esto servirá para aumentar el nivel de preparación tecnológica de dichas implementaciones y para demostrar la idoneidad general de las contribuciones de la tesis para futuras misiones de rover a Marte. Los algoritmos desarrollados e implementados se han validado utilizando los conjuntos de datos ya mencionados adquiridos en campañas de prueba de campo anteriores o mediante campañas de prueba de campo explícitas dedicadas a su demostración.

- **Diseño de una arquitectura GNC para rovers planetarios con capacidades de navegación autónoma:**

El **GNC** de un rover es el subsistema encargado de implementar las funcionalidades de navegación y ordenar al rover para conducirlo hacia el siguiente objetivo de exploración. La última contribución de la tesis ha sido la definición de una arquitectura **GNC** que articula la ejecución de los algoritmos previamente diseñados y validados experimentalmente. Esta contribución se ha basado en todas las anteriores y abarca todos los

métodos desarrollados, aprovechando su uso desde un punto de vista global.

Cabe mencionar que el trabajo realizado y descrito en esta tesis se ha realizado en colaboración con otros miembros e investigadores visitantes del Laboratorio de Robótica Planetaria (PRL) de la ESA, por lo que las aportaciones son fruto de un esfuerzo conjunto. En cualquier caso, dado el rol del doctorando de director responsable de laboratorio durante los últimos siete años, el trabajo aquí presentado ha sido desarrollado por él mismo o realizado por compañeros a los que ha guiado y supervisado de cerca, y del cual siempre ha sido responsable.

CONTEXTO Y MOTIVACIÓN

Esta tesis se ha realizado en el contexto de los estudios de investigación sobre tecnología robótica para misiones con rovers a Marte realizados en el PRL de la ESA y en colaboración con el Laboratorio de Robótica Espacial de la Universidad de Málaga. Los intereses y campos de investigación comunes han llevado a ambos laboratorios a tener una colaboración continuada desde 2016.

Los resultados de esta tesis son una demostración de la evolución de los estudios del PRL en una de las principales áreas de la robótica espacial y para la que el laboratorio da soporte a misiones planetarias reales como ExoMars o el futuro Sample Fetch Rover (SFR). La tesis también se suma al desarrollo continuo y al aumento de la experiencia y saber hacer del laboratorio en este importante campo de investigación.

Esta investigación fue motivada principalmente por la incorporación de capacidades de navegación autónoma a la misión del rover ExoMars, que se reconsideraron para ser incluidas en la misión después de haber sido explícitamente suprimidas durante algunos años. Además, la relevancia de las investigaciones y los resultados de este estudio también se destacan por los requisitos de navegación rápida establecidos para la misión SFR, que debe atravesar largas distancias en un tiempo limitado para cumplir con los objetivos de la misión, lo que convierte a SFR en el rover de navegación autónoma más rápido hasta el momento.

ESTRUCTURA DE LA TESIS

La tesis se divide en cuatro capítulos, en los que se incluye un primer capítulo introductorio y el último dedicado a las principales conclusiones y descripción de líneas de trabajo futuras.

El segundo presenta un análisis de la vanguardia o estado del arte en los temas centrales de esta tesis. En concreto, se profundiza en los pilares de la navegación autónoma y las distintas funcionalidades que la componen. Para cada una de las funciones analizadas proporciona una explicación de los principios teóricos y los enfoques de diseño encontrados en la literatura.

El tercero describe las principales contribuciones de la tesis. Se divide en apartados donde cada uno de ellos destaca y explica brevemente un aporte singular que forma una parte significativa del desarrollo para la tesis. Dado el formato de la tesis por compendio de publicaciones, las secciones describen los principios y conclusiones más importantes de cada contribución sin pretender profundizar en los detalles de implementación. Para ello, cada apartado va acompañado de una o varias publicaciones relevantes donde se puede encontrar información detallada sobre cada contribución y se explican en profundidad los fundamentos teóricos y los resultados experimentales. Por tanto, el capítulo sirve de resumen y visión de conjunto de las principales contribuciones de tesis.

Por último, la tesis está estructurada y escrita de manera concisa, para permitir al lector captar rápidamente los conceptos de las principales contribuciones y se basa en las referencias dadas a los artículos de revistas publicados para proporcionar una extensión a cada uno de estos conceptos.

CONCLUSIONES

La relevancia de la navegación autónoma en futuras misiones con rovers planetarios es reconocida por los requisitos establecidos en estos sistemas que están destinados a realizar tareas más desafiantes con menos intervención humana en cada misión. Las misiones [MER](#) y [MSL](#) de la [NASA](#) fueron las primeras en implementar capacidades de navegación autónoma. Sin embargo, su velocidad de desplazamiento neta se reducía significativamente cuando era dirigido de forma autónoma, debido a la cantidad de tiempo necesario para ejecutar estas funciones. La misión del rover [ExoMars](#) de la [ESA](#) también lleva embarcada la primera implementación europea de *AutoNav* (Navegación Autónoma), pero nuevamente, la estimada distancia que puede cubrir por día es más corta que cuando se ejecuta en otros modos de navegación no completamente autónomos. La campaña [MSR](#), y en particular la misión [SFR](#), ha generado un cambio de paradigma en este asunto, ya que requiere recorrer de forma rápida y autónoma unos pocos cientos de metros por sol para alcanzar el objetivo de la misión.

Por tanto, el trabajo de investigación realizado en esta tesis se ha centrado en estudiar métodos para mejorar las capacidades de navegación de los rovers planetarios y explorar las posibilidades de convertirlos en plataformas más eficientes y rápidas. Este objetivo se considera cumplido en esta tesis fruto de

los años de trabajo de investigación. Éstos comprendieron primero el diseño y la integración de sistemas robóticos, seguido del desarrollo de algoritmos para la navegación de rovers y la validación experimental de éstos en campañas de prueba de campo representativas. Por último, también se ha presentado el diseño de una arquitectura **GNC** completa.

El conocimiento y la experiencia adquiridos en el desarrollo e integración de tres prototipos de rover planetarios diferentes y su preparación para diferentes campañas de pruebas de laboratorio y de campo en apoyo de las misiones planetarias de la **ESA** han sido fundamentales para la investigación realizada a lo largo de esta tesis. El trabajo adicional dedicado a la ejecución de pruebas de campo para la adquisición de conjuntos de datos de diferentes sensores también ha sido relevante para el desarrollo de los algoritmos de navegación. Consciente de las limitaciones y dificultades de las misiones de rover reales para embarcar capacidades para navegar de forma autónoma, esta tesis ha estudiado e implementado métodos novedosos para cada una de las funciones de navegación esenciales por separado. Una de las principales contribuciones de la tesis es la implementación de algoritmos diseñados específicamente para la exploración espacial planetaria y en particular el escenario de las misiones con rovers a Marte.

En primer lugar, se han desarrollado tres componentes o funcionalidades con el objetivo de aumentar la velocidad de desplazamiento neta de los rovers en un modo de navegación que hace un uso eficiente de los recursos a bordo. Éstos son:

- una función de control de trayectoria basada en el algoritmo *Pure-Pursuit* con características de seguridad adicionales;
- un detector de obstáculos que puede identificarlos de manera efectiva utilizando sólo valores de stereo-disparidad en lugar de generar una reconstrucción 3D completa de la escena;
- y una función de planificación de rutas basada en el *Fast Marching Method* con capacidades de replanificación dinámica y multicapa.

Estas funciones se combinan dando como resultado un modo de navegación ligero que es particularmente eficiente en terrenos relativamente benignos. Gracias a una campaña de pruebas de campo que ejecutó varios recorridos autónomos se ha demostrado el aumento significativo en la distancia recorrida por sol que este modo de navegación eficiente puede lograr reduciendo las paradas de navegación a las estrictamente necesarias para evitar los peligros encontrados. Este logro significativo podría tener un impacto en la planificación estratégica y el retorno científico de las operaciones de una misión real. Una vez

que el rover se dispone a navegar de forma autónoma durante distancias más largas, hay un problema inherente a todas las técnicas de localización relativa que se vuelve más relevante: la deriva o error en la localización. Para hacer frente al inevitable aumento de la desviación y la incertidumbre sobre la localización del rover, se han desarrollado dos componentes más:

- un componente de **SLAM** basado en una combinación de técnicas de *Filtro de Partículas* y *Scan Matching* que mejora los resultados de las funciones comunes de odometría de localización relativa. Además, produce un mapa de elevación local de los alrededores del rover que se puede usar para navegar de manera efectiva y fluida a través de terrenos complejos con numerosos obstáculos, y
- un segundo componente para la localización global basado en *Template Matching* que utiliza este mismo mapa de elevación local para encontrar las coordenadas en el mapa orbital de referencia con la correlación más alta. La salida corrige la desviación de la localización hasta un rango equivalente a la resolución de la cuadrícula del mapa orbital.

Estos dos componentes combinados permiten la posibilidad de una navegación autónoma de largo alcance con pocas necesidades de supervisión. La evaluación del rendimiento y la robustez de ambos componentes se ha realizado con experimentos exhaustivos utilizando un conjunto de datos recopilados en una campaña de prueba de campo realizada en un terreno análogo cerca del Volcán Teide en las Islas Canarias. Como se puede esperar de las técnicas y tipos de datos utilizados, ambos componentes tienen un mejor rendimiento en terrenos con ciertas variaciones en las características de elevación, es decir, áreas donde el número de obstáculos, rocas o afloramientos rocosos es más significativo.

Finalmente, todos los componentes desarrollados se combinan en un diseño de arquitectura **GNC** que presenta dos modos de navegación, cada uno de ellos diseñado para diferentes dificultades de terreno. En áreas de dificultad baja a media con una abundancia rocosa o **CFA** relativamente baja, se utiliza el modo de navegación eficiente basado en el modo de evitación de obstáculos reactivo. Esto proporciona una implementación de navegación rápida que puede aumentar significativamente la distancia recorrida por un rover en un solo sol. En áreas con mayor **CFA** y pendientes, se utiliza el modo de navegación completa que implementa los componentes de localización global y **SLAM**. Esto permite una navegación más segura y consciente a través de terrenos más complejos donde el modo directo de evitación de obstáculos reactivo se detendría constantemente y se volvería menos eficiente. Curiosamente, este es el tipo de terreno donde los

componentes de localización global y [SLAM](#) implementados funcionan mejor. El uso de cada modo de navegación se puede seleccionar dinámicamente mediante la toma de decisiones a bordo utilizando métricas simples que proporcionen una evaluación de coste en su proceso. Alternativamente, ésto puede ser planeado de antemano por operadores de Tierra haciendo uso de un mapa orbital que puede ser analizado para segmentar las áreas por diferentes características de tipo de terreno. En general, se concluye que la arquitectura [GNC](#) diseñada proporciona las capacidades de navegación y la flexibilidad necesarias para una misión de exploración planetaria con una clara ambición de avance rápido como el [SFR](#).

En cuanto al potencial trabajo futuro como continuación a la línea de investigación de la tesis, existen diferentes temas que podrían explorarse. Aquí podríamos mencionar posibles mejoras o ampliaciones a los desarrollos ya realizados en la tesis, que eventualmente pueden conducir a la apertura de nuevas líneas de investigación.

Una de las áreas de mejora identificadas es el proceso de evaluación de la transitabilidad que vincula la función de percepción con el componente de planificación de la ruta. El actualmente simple proceso podría ampliarse para incluir información sobre características del terreno distintas de las de los obstáculos puramente no navegables. Esto podría proporcionar al planificador de rutas una mayor conciencia de las características del terreno. Al asignar a estos elementos un valor de riesgo asociado, diferente al asignado a zonas planas seguras y navegables, la solución del planificador podría proporcionar una ruta que se desvíe del terreno y pueda afectar al desgaste general del sistema o el consumo de energía.

Una segunda mejora podría centrarse en el componente de localización global para hacerlo más útil en condiciones de terreno benignas. Actualmente, su rendimiento no es tan robusto para determinadas condiciones del terreno como las que se encuentran en terrenos benignos donde se utiliza el modo de navegación eficiente. Si un rover encontrara una vasta área de terreno benigno y la atravesara, el modo de navegación eficiente podría terminar acumulando una deriva significativa en la localización. En ese caso, usar el módulo de localización global para corregir la posición sería una capacidad interesante. Como tal, y dado que el terreno por definición tendría pocas características de elevación (es decir, valores bajos de [CFA](#)), el componente podría extenderse para encontrar correspondencias con otras características, como las ópticas. Un posible enfoque para solucionar ésto podría estar basado en utilizar los mapas de imagen simple de [HiRISE](#) y compararlos con las imágenes proyectadas y unidas del rover, es decir, imágenes orto-rectificadas.

Por último, otro potencial trabajo futuro podría dedicarse a la comparación cuantitativa de la eficiencia de los dos modos de navegación en diferentes

condiciones del terreno. Este análisis podría ayudar a la correcta selección del modo de navegación más adecuado para cada tipo de terreno. Al mismo tiempo, esto podría permitir identificar casos extremos en los que ninguno de los modos de navegación fuera particularmente efectivo y desencadenar nuevos desarrollos de modos de navegación híbridos o alternativos dirigidos a esos escenarios.

Autonomous Navigation of Planetary Rovers deepens the field of autonomously-driven robotic systems and its application to planetary rovers. Based on the experience of previous successful Mars missions such as the Mars Exploration Rovers or Curiosity, it focuses on the challenges that future European missions such as ExoMars or the Sample Fetch Rover shall face. It explores the state-of-the-art in different functions that together compose the capability of Autonomous Navigation. Solutions for localisation, perception, SLAM, path planning and trajectory control are studied, and novel designs and implementations of these are proposed. Altogether, an architecture of the complete Guidance, Navigation and Control subsystem of a rover is designed, and its implementation tested to demonstrate the enhancements that it can bring to future rover missions.

