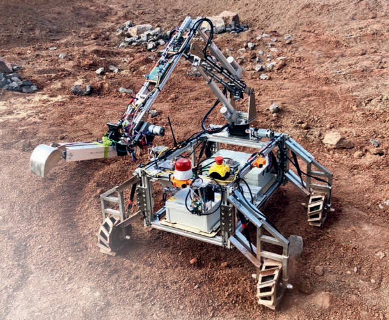

Autonomy Challenges for the Next Generation of Mars Rovers

Victor Covasan
Department of Aerospace Engineering
University of Sheffield, UK
vcovasan1@sheffield.ac.uk

Viktor Fedun
Lyudmila Mihaylova
Marwan Taher
Department of Automatic Control and Systems Engineering
University of Sheffield, UK

Joycelyn Fontanilla
Department of Chemical Engineering
University of Sheffield, UK

Hari Krishna
Department of Computer Science
University of Sheffield, UK

Abstract**—**Achieving full autonomy for a planetary explorer is the main requirement in rendering feasible missions when the communication time with the ground station does not allow real-time operation and monitoring. The design process involved in building the MarsWorks rover illustrates the challenges to be addressed in a typical surface exploration mission. This paper presents the main stages necessary to achieve rover autonomy in Mars-analogue environments. The focus is on two key areas: rough terrain navigation and autonomous manipulation with a six degree-of-freedom robotic arm. The first topic covers fundamental data fusion and Kalman filtering methods that estimate the current pose, as well as displacements from the starting position by means of visual-inertial odometry. An approach to guidance and control is then presented from the perspective of the dynamic window technique. Subsequently, autonomous grasping with increasing levels of automation is presented: from the low-level proportional–integral–derivative (PID) control to inverse kinematics, motion planning, computer vision, and automatic target recognition. Finally, onboard data handling, fusion of the sensor data used for scientific sample analysis, and communication with the ground station are briefly discussed. Each section presents future ambitions and possible ways of optimising individual subsystems of the MarsWorks rover.

INTRODUCTION

Planetary science missions are arguably the most emblematic accomplishments of the space industry, many of them (Curiosity, Perseverance, Europa Lander) being categorised as flagship-class missions, implying both immense efforts, technically and economically, but also colossal contributions to our understanding of the Universe. Flagship mission proposals are called for based on the Planetary Science Decadal Survey [1] published by the United States National Research Council. Such surveys consider the most relevant scientific questions of the decade and reflect the public interest in the exploration of certain celestial bodies. This combination of public enthusiasm and scientific return often motivate space agencies to organise student rover competitions that simulate innovations towards the next generation of planetary missions. Part of this series of contests and contenders is MarsWorks, an interdisciplinary team of students from the University of Sheffield, UK, a team who dedicated a substantial part of the academic year to designing and building a Mars Rover autonomous vehicle. The vehicle participated in the European Rover Challenge (ERC). Figure 1 shows the MarsWorks rover at ERC2019, which will be the subject of this article.

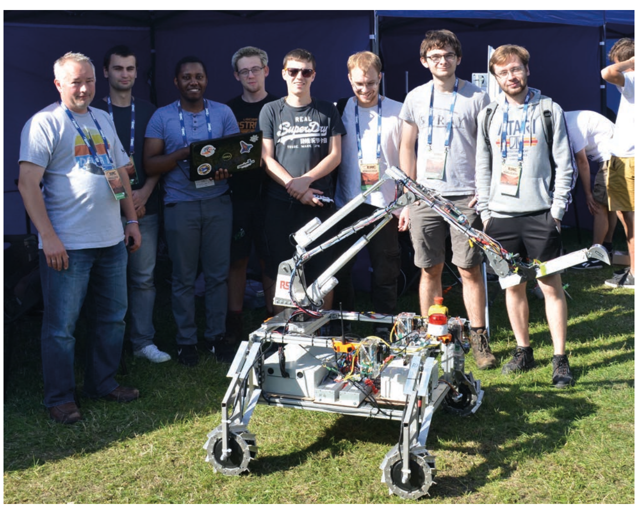

Figure 1
MarsWorks representatives at ERC2019.

The first planetary rovers, Lunakhod [2], were launched in the early 1970s and focused on extreme terrain mobility and small body/microgravity mobility. Navigation and control were difficult since computers were bulky and slow, thus Lunakhod was a teleoperated mission. The first rover on Mars was the National Aeronautics and Space Administration's (NASAs) Sojourner, a small 11.5 kg rover which explored the area within site of the Pathfinder Lander's camera, taking measurements of surface properties, imaging rocks, and obtaining their elemental composition. The Mars Exploration Rovers were the first to use advanced navigation methods such as visual odometry, a technology reaching maturity in the space sector which was also employed on the MarsWorks rover. This progression reveals the evolution of scientific objectives for Mars Rovers, starting with mobility demonstration, to search for water and life and more recently the investigation of in-situ resource utilisation potential or habitability [3].

The European Space Agency (ESA) ExoMars rover Rosalind Franklin is planned for launch in July 2020, with a focus on astrobiology. The mission will revolve around searching for past life on Mars, investigating gases and their sources and, by doing this demonstrating capabilities for a Mars sample-return mission in the future. The results delivered by the ExoMars rover will be complementary with the measurements taken by the Trace Gas Orbiter (TGO), which maps the distributions of hydrogen and methane in the atmosphere of Mars. More importantly, the satellite also serves as a communication link between future landers and the Earth, so rovers only need to uplink their measurements to TGO, which will relay them to the ground station. Sample-return missions have been known to have a great scientific return on investment as analysis is freed from the time, budget, or space constraints of spacecraft sensors; therefore, they are projected to become increasingly relevant over the next decades. Another motivation for sample-return research is the potential for asteroid capture and exploitation of resources from bodies located in the near vicinity of the Earth.

As a result, the tasks in ERC are based on the needs of a sample-return mission as well. The rovers involved shall be capable of autonomous navigation to a desired site and deep surface drilling once at the location, much like Rosalind Franklin [4], but also autonomous detection and collection of predefined targets in the Martian field. Additionally, rovers must be able to conduct onboard scientific analysis of collected soil samples and to operate a control panel in order to aid astronauts with maintenance of a Mars base. The concept of astronaut-aiding robots is motivated by the need to reduce the exposure time of astronauts to high radiation environments. As opposed to Earth, Mars lost its magnetic field with the cooling of its metallic core; therefore, the noxious solar radiation is not deflected around the planet but reaches the surface, which makes it a significant hazard to surface explorers, but especially to future human settlements. This is why it is crucial for most outdoors activities to continue to be performed by rovers and robots even with sustained human presence on Mars.

Achieving the highest degree of autonomy is also a continuing endeavour for current planetary explorers since the communication time between the ground station and spacecraft is much longer than the time available to respond to hazards. Specifically, the communication time to Mars is about 20 minutes; therefore, if a violent sandstorm starts to develop it might be too late for the storm warning to reach the Earth, then the "take shelter" command to be received and executed by the rover, because by that time the storm might have gained threatening proportions. However, the control algorithms employed also need to prove robust and predictable enough to be certified for space applications, this being a reason why inherently black-box approaches such as artificial neural networks are not suitable for safety-critical scenarios. Conversely, advanced model predictive approaches have been widely adopted by the space industry, especially for attitude control of orbiters, since they guarantee the desired performance and stability margins. On the other hand, increased autonomy drives the need for quantitatively more sensor measurements but also adequate accuracy of the estimated states, which in turn prompts the requirement for more computational power as well as more advanced data fusion mechanisms to handle the diverse range of sensors and increased data volume. Current rovers rely on a suite of sensors—infrared cameras, accelerometers, gyroscopes—for trajectory planning and attitude of the vehicle, but also spectrometers (to analyse the composition of sampled materials), atmospheric analysers, and radiation detectors. Often, measurements from all these transducers are requested by the same module of the rover for navigation, guidance, or control tasks, this requirement pointing back to the need for robust onboard data handling and fusion to ensure data compatibility.

MarsWorks is a student-led project at the University of Sheffield scoped with designing and building a Mars Rover for the European Rover Challenge (ERC). This project is part of Sheffield Space Initiative (SSI), a highly cross-disciplinary space technology platform that is now developing a real heritage of success for the University of Sheffield and our science, technology, engineering, and mathematics (STEM) students in particular. SSI was founded in 2017 to further engage the University of Sheffield students in the science and engineering challenges involved in space exploration. MarsWorks has its origins in Project MoonWorks, as the first rover was dedicated to the fabrication of a miniature lunar vehicle which could retrieve ice samples from the depths of lunar craters. The team has a broad range of activities. A successful participation in the national competition organised by UK Students for the Exploration and Development of Space in 2018 led to a prize award for the best innovation for the developed advanced scooping mechanism. The project called MarsWorks, moved forward to design and build a fully autonomous Mars Rover to participate in the European Rover Challenge (ERC), the biggest space robotics event in Europe.

With these considerations in mind, the article aims to demonstrate how control and data processing algorithms are implemented in a complex real-world system such as a Mars surface explorer with a view to achieving the highest degree of autonomy. The rest of the paper is organized as follows: the systems architecture is presented, the robotic arm control system is described, followed by the guidance, navigation, and control system. The latter sections focus on the sensor data fusion approaches used for the scientific measurements, and conclusions and the key lessons learned are given.

THE SYSTEM ARCHITECTURE

In modern space systems engineering, integration is primarily achieved in software, implying that all subsystems and components must communicate with a central computing element. As a result, it is worth starting the discussion with the high-level system breakdown together with the interfaces between

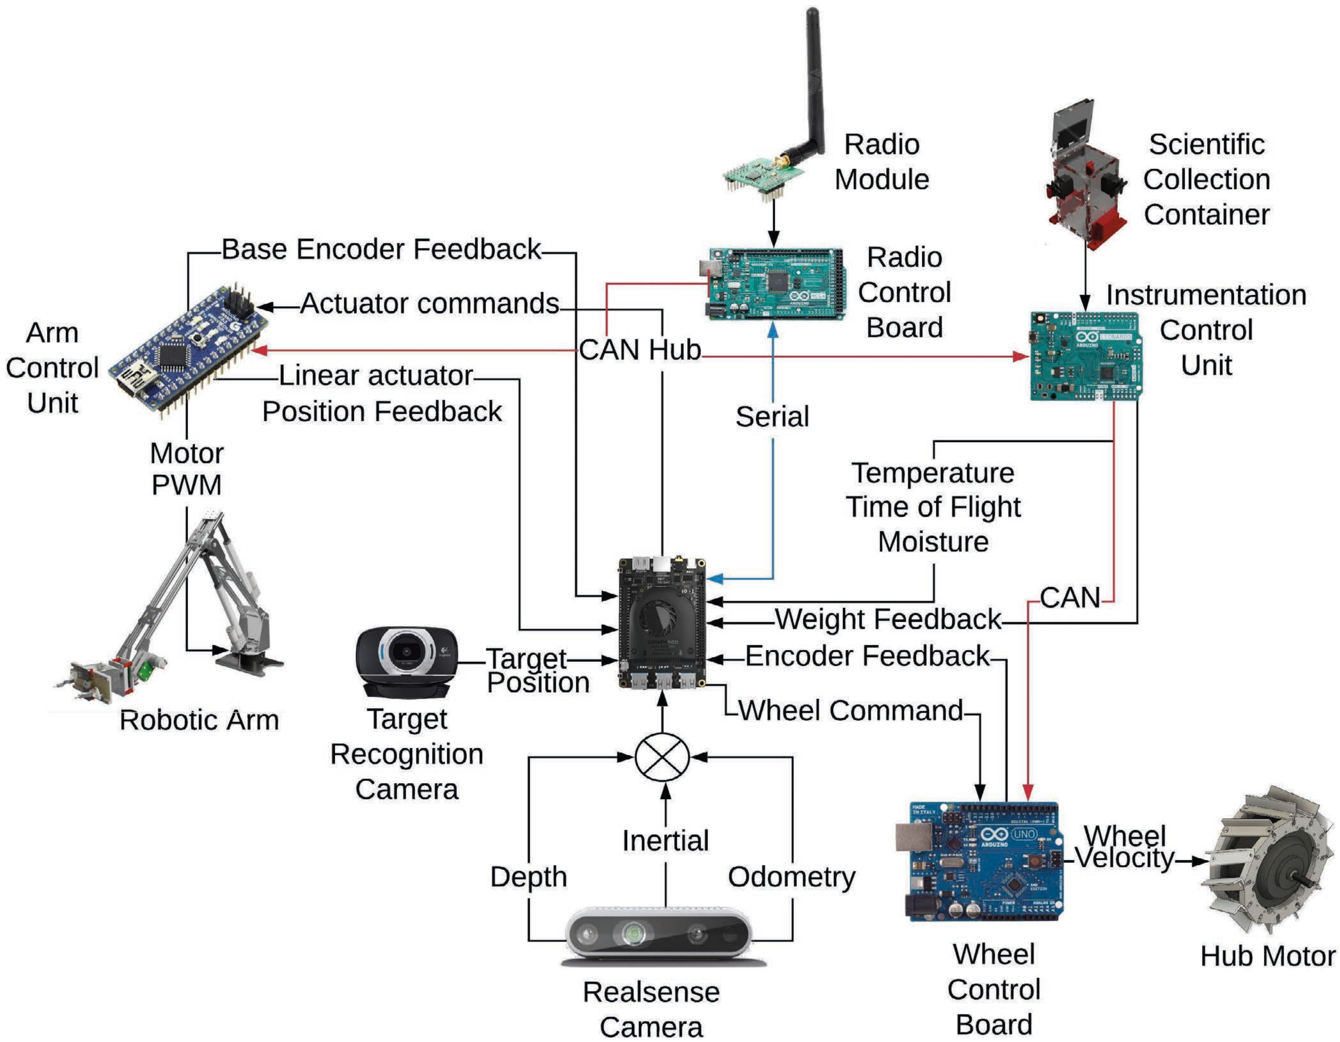

Figure 2
System breakdown structure of the main control instances.

the individual subsystems and the selection of the main onboard computer. Due to the computationally demanding nature of the image processing tasks involved in autonomous navigation and motion planning, the Advanced Reduced Instruction Set Computer (RISC) Machines architectures were considered unable to satisfy the mission requirements. As a result, a LattePanda Alpha development board was found more appropriate. The Alpha serves as the Rover Compute Element (RCE), being the highest authority in the control hierarchy of the system. Figure 2 shows the final control architecture of the MarsWorks rover.

Secondary computing units were designed around popular microchip controllers (ATmega32u4, ATmega328, and ATmega328p [5]) to handle specialized tasks such as motor control, radio communication, and signal processing. This way, the computational load on the RCE is reduced by implementing real-time operating systems together with custom scheduling algorithms on each microcontroller. It also guarantees that no control task will be interrupted by lower priority tasks, hence avoiding instability. These secondary computers are: four Instrumentation Control Units which process the sensor measurements of the collected soil sample (mass, temperature, humidity, and time-of-flight), four Motor Control Boards to run the Proportional–Integral–Derivative (PID) speed control algorithms for each wheel, one Arm Control Board to drive the six arm actuators, one Radio Controller, and one Emergency Stop Controller. The main protocol used to communicate between microcontrollers was Controlled Area Network (CAN), which is the preferred standard at ESA and the other space agencies for most spacecraft data handling applications due to its decentralised architecture.

The framework used to coordinate all the processes running on the craft is the Robot Operating System (ROS) [6], hosted on the RCE. ROS was considered the ideal middleware between low level motion control and autonomous navigation or guidance since it provides built-in capabilities for fusing sensor measurements coming from different subsystems of the rover, state estimation, visual odometry, as well as multiple open-source packages for interfacing with the RealSense cameras or even Arduino. For example, rough terrain traversal is achieved by having the RCE analyse the output of two infrared and depth cameras to detect the next site to be reached as well as determining its own position and pose (navigation). Next the RCE analyses the picture frames for clear paths towards the target, loops through the results to select the shortest one, and computes the trajectories needed to get there (guidance). Finally, these values are transmitted to the individual wheel controllers; they are then converted to Pulse-Width Modulated values which are sent to the direct current motors to drive the wheels (control), all while the same motor controllers read the encoder outputs from the motors to ensure that the desired speed was achieved without offset (PID).

Due to the nature of the competition tasks, the arm and locomotion dynamics are relatively decoupled for most path planning scenarios. Consequently, the detailed analysis will follow each subsystem individually on a component level.

NAVIGATION AND LOCOMOTION CONTROL SYSTEM

The main task to be completed by the navigation stack is autonomous traversal: finding and following the shortest clear path leading to the previously localised target in the odometric frame. To achieve that, depth and stereo data from the RealSense cameras was used to create a two-dimensional occupancy grid (local costmap) of the immediate environment. Then, the costmap is searched for clear paths towards the target using the Dynamic Window Approach algorithm [7], which is traditionally used in robotics for collision avoidance. Then, the algorithm loops through the results found and selects the shortest path. Once the trajectory corresponding to the chosen path is computed, ROS will generate the control actions required to drive the wheels and guide the craft along the path. The control signal is transmitted from the RCE to the Radio Controller through serial, and from there to the individual Wheel Controllers through CAN bus.

For navigation purposes, Visual-Inertial Odometry (VIO) was the preferred approach over Simultaneous Localization and Mapping (SLAM) due to the ease of incorporating wheel-encoder feedback into the visual position data, which together with the displacement estimates obtained from the Inertial Measurement Units (IMUs) provide a very robust system that accounts for slippage and many other uncertainties in hazardous terrains [8]. Furthermore, VIO also proved to be more computationally efficient than SLAM especially in scenarios when the navigation algorithms have to run in parallel with the arm motion planning within ROS. All these contributed to VIO being the preferred navigation methodology for the NASA Mars Exploration Rovers, where it delivered unprecedented performance (97% convergence on Spirit and 95% on Opportunity) [9].

One of the fundamental questions posed by VIO is estimating the distance travelled by the vehicle from the origin of the odometric frame. This is accomplished within ROS by feeding all the measured displacements (encoder ticks, doubly integrated accelerations from the IMU, and stereo camera distances) to a state estimation algorithm (the Extended Kalman Filter function), which weights and fuses the measurements according to the reliability gains specified in the measurement covariance matrix. The output of this function is an optimum distance estimate.

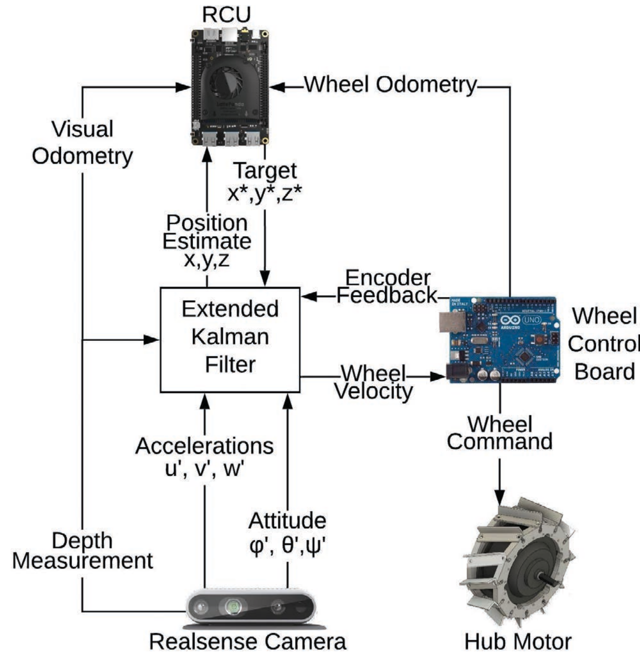

Figure 3
Navigation system—main components.

A major goal for the next design iteration is being able to incorporate a multirotor unmanned aerial vehicle (UAV) to the rover, which would be able to perform extensive mapping of the environment and relay the data in real time to the rover. Using this approach, the rover would become capable of planning the traversal task potentially kilometres ahead and learn about the feasibility of certain routes way advance, therefore making the most out of its time in commission. This would arise fascinating challenges both in terms of translating an aerial perspective to a terrestrial planner but also from the point of view of autonomous docking of the UAV with the rover for recharging, which is a great opportunity to experiment with model predictive control approaches.

ONBOARD SENSOR DATA FUSION

Sensor data fusion [10] as a process of knowledge extraction from multiple sensors plays an important role in the navigation of the Mars Rover. The data from multiple cameras are fused with data from other sensors, such as encoders. As part of the scientific task, basic properties of the collected soil sample have to be measured (temperature, density, humidity), which inevitably involves the fusion of signals from very different transducers, but also filtering, interpolation, or downsampling to enable the variables to be manipulated together in calculations. One basic example is the density determination, which involves combining the readings from a load cell (mass) with the time of flight data (volume occupied); therefore, the sampling rates of the two had to be matched and different low pass filters had to be implemented according to the known transducer dynamics.

For the simulated Martian environment in which the rover was designed to operate, the collected sand was stored in three identical scientific containers equipped with sensors that measure the temperature, humidity, proximity/time-of-flight, and weight of the sample. Each container is equipped with an Instrumentation Control Board, which performs sampling, signal amplification, analog-to-digital conversion (especially for the low voltages coming from the load cells), as well as preliminary filtering before transmitting the data through CAN bus.

Two types of moisture-sensing devices were present onboard, this form of redundancy offering the opportunity to apply more advanced averaging and voting techniques. Namely, the DH11 temperature sensor [11] also provides relative humidity measurements; however, these are not as reliable as the ones returned by the dedicated Grove NE555DR [12] capacitive moisture sensor. As a result, the dual sensor was given a reliability weight of 30% as compared to a weight of 70% for the dedicated moisture sensor. It was also observed that upon collection, the moisture measurement provided by DH11 sensor presented a settling time of about five seconds, hence a moving average filter was included in the loop as a smoothing mechanism for the initialisation period. To further improve the estimation accuracy of the measurements, a simple averaging-and-voting system was implemented. This was based on the assumption that for a given digging site, no radical leap in humidity should occur within a depth of five or less centimetres; therefore, in the event that any of the three containers returns a humidity value 10% or more greater than the other two, the measurement should be excluded from the estimation and the other two are averaged instead.

PID control was also implemented for collection to ensure that the amount of soil poured by the arm into the containers matches the perfectly desired quantity. To achieve this, the weight measurement is used as a feedback signal for the robotic arm upon releasing the sample. It was therefore critical to ensure that no spikes are present in the load cell signal (which often proved to be the case during the first seconds of auto-calibration) as those would lead directly to violent control actions in the robotic arm. To account for that, a low pass filter with a cut-off frequency of 0.5 Hz was introduced in the loop. Low pass filtering is usually resorted to in industrial applications for noise reduction in signal processing or for lowering pixel contrast in image editing, however for the given application it also helped with smoothing the sharp transitions occurring at auto-calibration.

Handling the proximity (time of flight) data was comparatively straightforward. The proximity sensor was used to determine the distance from the lid to the sand surface inside the container, therefore ON-OFF control was implemented to receive distance data only once the lid was closed. Having the distance to the sand surface, one can determine the volume occupied by the sample, and together the measured mass the average density of the collected soil can be estimated.

The decentralised architecture of the data handling system allows for modularity and communicational efficiency, the data being digitised and filtered locally, therefore only relayed through the network once clean and reliable. One immediate advantage is that signal attenuation and contamination are avoided, no analog signal having to traverse a region of the vehicle with potential for electromagnetic interference before analog-to-digital conversion. This approach also reduces the computational load on the RCE and ensures that compact data packets are being transmitted to other subsystems, which is both bandwidth-efficient for transmission and prevents control hazards arising from raw signal contamination. As a result, the scientific data leaves the Instrumentation board already at Data Processing Level 3 [13], before reaching other segments of the rover or the ground station.

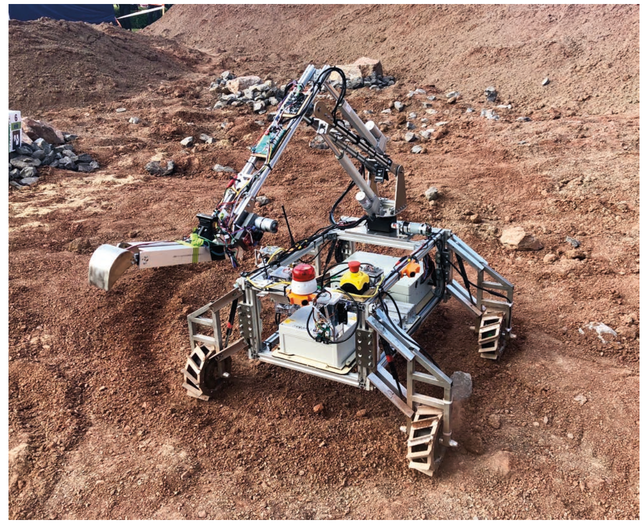

Figure 4
The MarsWorks rover in excavator configuration at ERC2019.

REFLECTIONS AND CONCLUSIONS

As access to space becomes more affordable and launch frequency grows due to space commercialization, ground-in-the-loop guidance of spacecraft will become prohibitively expensive because of scheduling conflicts but also increases in maintenance and labor costs. This will only contribute to an increasing probability of a human error. Automation can prevent such outcomes, enabling greater numbers and types of missions to operate concomitantly while improving robustness, reducing risk, and hence increasing future commercial and scientific return from space.

The ERC prepares young engineers for the challenges posed by the next generation of Mars Rovers, which are expected to have an increasingly high scientific pay-off. The competition itself is akin to a test at an analogue site, achieving striking similarity with the actual martial terrain. The development of autonomous navigation, manipulation, sample collection, and return are the core technology gaps to be addressed for these sample retrieval missions. The benefits of development in autonomous space technologies would also spill into other industries, for example a major potential for these types of rovers would be in farming and agriculture. Most noticeably, agriculture and food production are important sectors for any country, particularly the UK. The UK-Robotics and Autonomous Systems network stated that the Agri-Food industry is the largest manufacturing sector in the UK. This industry has a high manual labour demand which results in low productivity. The UK government recently committed to investing £90 m to boost productivity through automation and process monitoring. A key capability for agricultural robotic vehicles is that they must be able travel on uneven terrain, without damaging crops, which parallels the requirements for Mars Rovers. In the future, more investment is to be expected in this area due to the reduction in the number of manual labour workers and rising wages.

Although building a Mars Rover proved to be an exceedingly rigorous and demanding technical endeavour, some of the most valuable lessons learnt in the process concern the human aspect of the project, revealing that the transfer of knowledge between teams, good communication structures within the organisation, and having dedicated people for systems integration are just as critical to project success as the science and engineering underpinning it. The key lessons generally agreed by the MarsWorks team after ERC2019 are:

- Testing and verification should be carried at every stage as opposed to only happening at the end of the development process.
- Systems integration should happen as soon as possible in order to avoid collisions or incompatibilities close to delivery.
- Independent external reviews from separate University-based teams are a great opportunity to exchange good practices and motivate better documentation.
- Every two subteams should share at least one member. This promotes cohesion between subsystems and makes the management aware of problems at the interface between technical areas or subteams.
- Team members with a background in systems engineering and general knowledge of each technical area tend to be the most effective leaders.

ACKNOWLEDGEMENTS

Funding: MarsWorks was made possible with support from the Automatic Control and Systems Engineering and Mechanical Engineering Departments together with the School of Mathematics and Statistics from the University of Sheffield as well as our Alumni Fund, KE Fund, EEE, AMRC Sheffield, and the UK Mars Society.

Author Contributions: L.M. prompted the motivation for the study and provided invaluable comments on the interpretation of the sensor data. V.F. coordinated the overall project and contributed to the production of the manuscript text. V.C. analyzed the data and wrote the first draft of the manuscript. J.F. presented the management strategies and development process involved in designing and manufacturing the rover. H.K. covered the computer vision methodology applied for cache recognition. M.T. provided documentation on the navigation algorithms developed for the traversal task.

Additional thanks are extended to Dr. Li Sun, Dr. Gary Verth, and Dr. Istvan Ballai for helping establish the scientific foundation on which the system was subsequently designed and engineered.

REFERENCES

1. United States National Research Council. Planetary science decadal survey. Available: https://solarsystem.nasa.gov/science-goals/about/, last access Mar. 2020.
2. National Aeronautics and Space Administration. Solar System Exploration, Missions. Lunakhood 01. Available: https://solarsystem.nasa.gov/missions/lunokhod-01/in-depth/, last access Oct. 2019.
3. National Aeronautics and Space Administration. Mars exploration rover history. Available: https://mars.nasa.gov/mars-exploration/missions/mars-exploration-rovers/, last access Oct. 2019.
4. European Space Agency. ExoMars. Available: https://www.esa.int/Our_Activities/Human_and_Robotic_Exploration/Exploration/ExoMars/ESA_s_Mars_rover_has_a_name_Rosalind_Franklin/, last access Oct. 2019.
5. Microchip. ATmega328, ATmega328P, ATmega328PBp. Available: https://www.microchip.com/wwwproducts/en/ATmega328, last access Oct. 2019.
6. Robot Operating System. ROS kinetic. Available: http://wiki.ros.org/kinetic, last access Oct. 2019.
7. Fox, D. The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine*, Vol. 4, 1 (1997), 23–33.
8. Pavone, M. Spacecraft autonomy challenges for next generation space missions. *Lecture Notes in Control and Information Sciences*, Vol. 460, (2016), 1–34.
9. Maimone, H. *Two Years of Visual Odometry on the Mars Exploration Rovers*. Hoboken, NJ: Wiley, 2007.
10. Fourati, H. (Ed.). *Multisensor Data Fusion: From Algorithms and Architectural Design to Applications*. CRC Press, Taylor & Francis Group, Aug. 26, 2015.
11. Waveshare. DHT11 temperature-humidity sensor. Available: https://www.waveshare.com/wiki/DHT11_Temperature-Humidity_Sensor, last access Oct. 2019.
12. Grove. Capacitive moisture sensor. Available: http://wiki.seeedstudio.com/Grove-Moisture_Sensor/, last access Oct. 2019.
13. National Aeronautics and Space Administration. Data processing levels. Available: https://science.nasa.gov/earth-science/earth-science-data/data-processing-levels-for-eosdis-data-products, last access Oct. 2019.