THE CHAMELEON FIELD TRIAL: TOWARD EFFICIENT, TERRAIN SENSITIVE NAVIGATION

Mark Woods [(1)], Andy Shaw[(1)], Iain Wallace [(1)], Mateusz Malinowski[(1)]

[(1)] SCISYS, 23 Clothier Rd., Bristol, BS45SS, UK, Email: mark.woods@scisys.co.uk

ABSTRACT

This paper outlines an adaptive, software based approach to autonomous navigation. The Chameleon project set out to understand if a remote rover could dynamically self-select an autonomous navigation mode based on visual characteristics of the local terrain. Underlying support functions for autonomous navigation such as localisation and mapping can be implemented in a variety of ways. Each mode has different accuracy characteristics and an energy cost. Currently mode selection is fixed a priori at planning time by ground operators.

For long range traverses this may be sub-optimal. By allowing a rover to autonomously select the navigation mode based on expected energy costs and mission constraints we might expect to make energy savings which could increase the overall range of the rover. In the course of the project an experimental architecture was developed and tested during dedicated field trials in the Atacama Desert. The paper summarises the results of this work.

1. BACKGROUND

The exploration of various solar system bodies using mobile, robotic platforms is at an interesting juncture. Existing or near-term missions such as NASA's MER Rover, Curiosity and ESA's ExoMars have settled on a relatively stable but conservative operations paradigm consisting of short-range (typically <100m) traverse implemented as a mix of mainly manual (blind) and semi or fully autonomous drives [1], [2] and [3]. At the time of writing little information is available on the Chinese Lunar rover called Yutu but published data indicate that is very similar to American and European developments.

Over the longer-term the Mars Sample Return (MSR) concept will rely on rovers which must meet much more difficult operating requirements. At the heart of the MSR concept lies the Sample Fetch Rover (SFR). ESA has conducted two Phase A studies looking at Fetch Rover feasibility [4] and [5].

The landing error ellipse for SFR has a semi-major axis value of 7.5km which means that the rover must be able to traverse a minimum straight line distance of 7.5 km x 2 = 15 km. However in practice the rover will not travel in a straight line given obstacle avoidance requirements which means that a target distance (one-way) of around 10.5 km or more is expected. In total therefore the rover should aim to achieve a traverse of around 21 km. The proposed landing date is Ls 133 in Sept 2025. The mission is therefore characterised mostly by a decreasing pattern of solar energy (worst case) given that the target landing is most likely to occur in the Northern Hemisphere (anticipated range between -5 deg and 25 deg latitude). The mission will last for 180 sols however only approximately 112 sols are available for travel and cache retrieval. It is anticipated that 73 sols will be available for return-trip travel. Additional characteristics of the operational period include local dust storms, a solar conjunction and diminishing power availability as the mission progresses [4].

These studies show that this mission is difficult to achieve using state of the art technology. Various hardware solutions have been proposed to address some key challenges but remain at a low technology readiness level. The feasibility assessments also rely on a number of additional assumptions which are tenuous, leading to a high probability of mission failure.

The key technical challenge for this mission therefore is to develop a low-mass rover (60kg order) which can traverse 21km in 79 sols with significant power and computational constraints.

The total range which the rover can achieve in a given period of time is a function of the power available to support navigation and locomotion. This in turn is the function of a complex mix of variables such as; Power Generation; Sensor Choices; Algorithm Complexity; Navigation Modes; Computational Platforms and Locomotion Subsystems.

This work looked at the software and related sensor aspects of the challenge. It sought to determine whether or not greater rover navigation efficiency could be achieved by dynamically selecting navigation modes (and by implication the required sensing and algorithms) based on an assessment of the current terrain.

2. APPROACH

This research sought to explore the possibility that a Mars rover can significantly maximise its range per fixed level of resource by determining which sensors it should use to safely navigate over unknown terrain. Sensors produce data products which are used by various algorithms to support localisation, mapping and path planning or obstacle avoidance. The nature of the environment, for example terrain or time of day, will determine the performance and reliability of these data products. The performance of the algorithms which use these products is also dependent on various parameters such as number of features tracked and resolution of traversibility maps. An intelligent rover could in principle alter the sensor, algorithm and control parameter configuration of the overall navigation subsystem to suit real-time environmental conditions. If we consider a menu of possible configurations or navigation "modes" an intelligent rover could make an optimal mode choice based on its classification of the environment and an understanding of how the modes perform in different conditions i.e. a mode-environment performance (resource, time, memory)/risk model or mapping.

If the choice of mode makes a significant difference to resource use (as currently observed on existing missions) and the resource cost realised by making the choice itself is relatively low then the overall range of a rover could be significantly increased for a given power budget versus the current manual and non-adaptive approach. The key point is that this could be achieved by software means alone i.e. zero or low additional mass cost versus other mechanical approaches.

In short, the rover should locally be able to ask and answer a questions such as "which navigation/sensor mode is the cheapest in resource terms for the given risk/range setting for this sol?" and act accordingly.

1.1. Lab Tests and Sensor Modelling

To explore this question of sensor choice we conducted a series of lab based experiments in order to build a baseline model which captured the relationship between various sensor products, their quality, resources required to generate (performance) and their veracity when key environmental or internal parameters were varied. The frequency of sensor acquisition was also investigated as it directly impacts power consumption in processing data and is dependant on platform speed relative to pose estimates and the accuracy of mapping data for navigation.

To build our model we considered platform speed as the main variable, which combined with evaluating mapping methods for quality and crucially range gives a minimum frequency (and thus power) per-method that maps must be built. Knowing the size of maps and speed of the platform we can then extract the minimum required pose update frequency for a rover to safely navigate with a specified safety margin to avoid obstacles.

For example, the LRF operates well for mapping over a

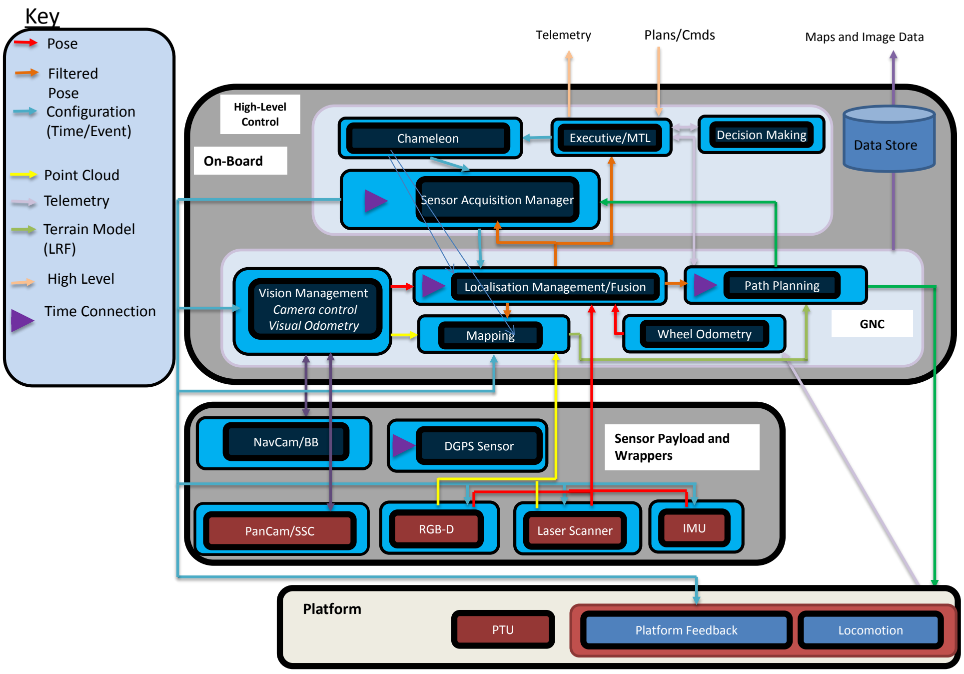

Figure 1: Chameleon System Architecture

2m range compared to 5m for low resolution Bumblebee images, so although building 3D maps from images may be more costly, it only must happen half as often. These lab results would allow us to assess such trade-offs.

3. ARCHITECTURE

The Chameleon system was designed to provide a frame work for autonomous rover control for navigation, sample identification and response. An implementation of this iteration of the SCISYS Overseer on-board architecture was developed and tested in field trials.

These trials took place using rover and a UAV platform (to provide simulated orbital data) in the Martian analogue Atacama region of Chile.

At a high level the system comprised of:

- A rover platform.
- Appropriate sensors and actuators hosted on the platform.
- Computer vision and autonomy software applications running on-board the platforms where possible.
- A mission control centre with planning, monitoring and control software
- External data inputs i.e. satellite and UAV imagery.

3.1. Robot Platform

The SCISYS SOLO (Figure 2) autonomous robotic platform is a four wheeled passive suspension chassis with all wheels capable of both drive and steering, providing a high degree of freedom allowing it to emulate the motion of a host of planetary rover designs. It is lightweight and high capacity, with the base platform being under 30kg in typical configurations yet being able to host up to 40kg of payload, drive up to 0.5m/s and support all-day operations with over 1kWh of li-ion battery capacity. For the Chameleon trial SOLO was configured with a 1.6m main mast.

Figure 2 SCISYS SOLO Robotic Platform

3.2. Sensing

Key to the Chameleon project is capturing and processing a host of sensor data to evaluate different approaches. Figure 3 shows the mast sensor configuration of SOLO. In addition to these some sensors were mounted on the rover body giving a total payload as follows:

- High-resolution SCISYS Stereo Camera (SSC) on PTU.
- Light Meter.
- Novatel DGPS.
- Front and rear facing Bumblebee XB3 stereo cameras.
- Asus Xtion structured IR light sensor.
- Wheel Encoders.
- Low power Hokuyo scanning Laser Range Finder (LRF).
- Inertial Measurement Unit (IMU)
- Wheel Encoders
- Platform power telemetry.

Figure 3; SOLO Sensing (Partial)

3.3. GNC

The Chameleon architecture supported evaluation of a large variety of GNC configurations, using different sensors to provide pose and mapping data in different combinations. Table 1 summarises these options showing which sensors were used to provide pose estimates, mapping data for navigation or just used for data capture. Tests were carried out with every combination of pose and mapping data, including no map data at all. Some tests were run with unused sensors capturing data for future use and some using only the required sensors to validate lab based power experiments.

Table 1 - Sensors and Data Product GNC Use

Sensor	**Pose**	**Map**	**Capture**
SSC Low-Res (1Mp)	X	X	X
SSC High-Res (2.3Mp)	X	X	X
Bumblebee Images (0.3Mp)	X	X	X
LRF		X	X
Xtion		X	X
Wheel Odometry	X		X
DGPS			X
IMU			X
Light Meter			X
Power Telemetry			X

3.4. High-Level Autonomy

In addition to the usual sensor processing and GNC components the Chameleon architecture also adds components to perform additional higher level autonomy functions as part of plan execution.

The two key components are the Chameleon component and the Sensor Acquisition Monitor (SAM). The SAM controls the various pose and map generation algorithms, allowing the robot to dynamically switch between GNC configurations and update rates to control power consumption. The Chameleon component is responsible for managing when these mode changes occur, as a result of image and terrain map analysis.

3.5. Operations

Plans were created, dispatched and monitored during execution with the Overseer INTERACT tool, previously deployed for mission operations for the SAFER trial [6]. This allows for 3D timeline planning and execution monitoring combined with a real-time co-registered view of collected data products.

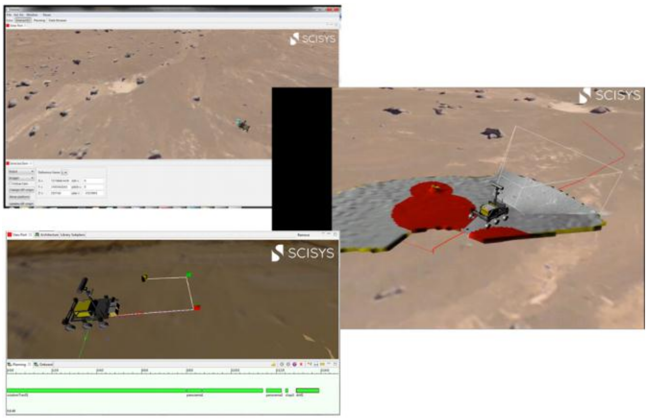

Figure 4: Overseer INTERACT Rover Operations Tool

4. FIELD TRIAL

To further validate the lab based results and test the approach using real trial data, SCISYS led a team deployment to the Atacama Desert in the Northern Province of Antofagasta, Chile in Sept/Oct 2014. The ESO Paranal Observatory served as a local base for the team who were otherwise self-sufficient. All equipment and supplies were shipped to the region to facilitate field trial execution.

During a previous reconnaissance to the area several sites were evaluated in order to select a final test area. SEEKER Valley, which was used by the team in previous work [7], was selected for the trials.

The purpose of the trials was to use the highly representative analogue Martian terrain found in the Atacama to extend the sensor data product characterisation begun in the lab. A key variable which affects sensor product quality is the visual nature of the terrain. This in turn is largely influenced by the terrain texture, morphology, viewing angle and solar illumination. Understanding the cost of producing and using such products is key to determining which modes suit which terrain types in terms of energy cost which in turn affects range. This will form the basis for any autonomous mode switching.

A series of approximately 100 m x 100 m sub-regions or patches within the large (4 km x 2 km) test area were chosen to provide a spectrum of terrain variables which could also be tested using alternate viewing angles through sensor placement. The majority of tests were executed using platform mounted sensors in order to maximise the completeness of the tests in the least possible time.

For the purposes of the trials 5 unique terrain types were identified:

1. Sand or small grains
2. Shale
3. Dried out “lakes”
4. Mixed sand and shale with small pebbles
5. Large boulder fields

Figure 5; Example terrain patches. Moving clockwise these are Small Grains, Shale, Mixed and Lake. iPhone 5 used for scale in some cases.

During the course of the field trial, sensors were moved across the terrain patches using various attack angle patterns in order to support product evaluation. Longer km order trajectories were also executed in order to further validate these conclusions. The results of this work are presented in section 5.

Understanding how parameters such as terrain type and viewing angles affect product usage, navigation cost and ultimately rover range is a key aspect of the dynamic navigation concept. In order to realise the complete concept it is also essential to have a component which can recognise mode change trigger points. These trigger points are largely based on terrain visual texture and 3D morphology. As part of field trial work we also evaluated a terrain recognition approach based on our research in the area of science autonomy. The field trial location provided an excellent opportunity to test the terrain recognition abilities of the software in an analogue environment given the prevalence of hard and soft (graduated) terrain type boundaries. We executed a number of trials with this in mind and also assessed the run-time cost associated with this function.

4.1. UAV and Data Collection

During field trials UAV’s can be used to substitute orbital data for offline planning during operations. Fixed wing UAV’s can provide coarse resolution wide area coverage but in order to experiment with future robot long range guidance concepts higher resolution data, more accurately targeted, is required. As part of the Chameleon work SCISYS commissioned an adapted quadcopter design in order to optimise data acquisition in the exceptionally low humidity and moderately high altitude conditions found at the Atacama operating site.
Image data acquired using the platform has been used to produce high-resolution (millimetre scale) orthographic and 3D models of the underlying terrain at patch sizes of 200 m x 200 m. This data is now being used to support offline simulation of rover autonomy concepts.

Figure 6: SOLO Rover visualised in Overseer INTERACT using 5mm resolution terrain model reconstructed from UAV data.

Figure 7: SCISYS UAV Platform

5. RESULTS

The Chameleon trial produced an exceptional wealth of data for on-going study and analysis:

- 6,875m of autonomous traverse.
- 680 Minutes of rover operations.
- 1.25 Million Images.
- 225.5 thousand 3D point clouds.
- 2300 high resolution UAV images.

This is in addition to data collected during lab trials to characterise sensors.

We are well positioned to investigate the performance and power consumption of individual algorithms and components of a rover system, as we have a full stack rover system from the wheels on the ground to the planning an execution monitoring software. This means that in lab tests we were able to isolate data capture and processing stages, measure their execution time and system power consumption and derive energy costs. Table 2 shows power consumption to produce one instance of a particular data product as measured in our lab trials. These results are produced from prototype algorithms on COTs hardware, but still provide us with informative results and indicate potential avenues of future investigation. The results are broadly as expected with a few points worthy of comment. Complex operations on larger data products, such as filtering high resolution dense point clouds consume the most energy, with small data products and simple operations such as reading data from an IMU consuming the least. The SSCs allow us to assess the energy impact of generating the same data products from the same sensor at different resolutions; we see the cost of the higher resolution to be ~3x the low resolution option. Of course energy consumption is just one part of the equation – the quality of data produced and therefore frequency required is also relevant. A particularly notable result is that applying pre-trained machine learning algorithms is relatively cheap – on par with a Visual Odometry (VO) pose estimate. This suggests that certain machine learning techniques for advanced autonomy may be possible on flight hardware.

Table 2: Data product size and energy cost from lab trials. *Numbers are based on reported CPU/Memory subsystem power consumption of COTs hardware running prototype quality algorithms and only provide indicitive results.*

Data product or Test	data product size (Mb)	energy cost over idle (joules)
Inertial Measurement Unit Pose	0.00008	0.0010494
SSC Rectified Image pair (lower res.	0.3	0.113608
3D Map to Occupancy Grid for nav	0.05	0.19635
Bumblebee Rectified image pair (low res, 0.3Mp)	0.81	0.24576
SSC Rectified Image pair (High res.	1.12	0.37576
Xtion Point Cloud, low detail	0.69	0.51374
Xtion Point Cloud, High detail	4.6	0.79515
Convert Laser Scan to Point Cloud	0.009	1.4028
SSC Visual Odometry Estimate (low res,	0.0003	1.53695
Bumblebee Visual Odometry Estimate	0.0003	1.59258
Machine Learning Based Terrain Analysis (1 frame)	0.0001	1.66152
Bumblebee Point Cloud (low detail, 7m)	2.6	1.7652
SSC Visual Odometry Estimate (High	0.0003	3.783
SSC Low Resolution Point Cloud, (7m)	7	4.731
Evaluate 3D map	0.0002	4.7312
generating map SSC low res	0.175	13.365
generating map SSC high res	0.14	39.183
Bumblebee Point Cloud, SOR filter	2.3	54.02
Bumblebee Point Cloud, ROR filter	2.6	84.723
SSC High Resolution Point Cloud, (30m)	26	92.064
SSC Low Res Point Cloud SOR filter	6.4	149.688
SSC Low Res Point Cloud, ROR filter	7	230.308
SSC High Res Point Cloud SOR filter	26	663.564
SSC High Res Point Cloud ROR filter	26	3940.5

With this data, we can then predict the power consumption required for particular navigation modes and sensor configurations – the necessary input to any system dynamically switching between them. Table 3 presents initial results from analysing a selection of the captured field data. We use the above power table combined with the sensor configuration (sensors used, mode, triggering frequency, platform speed) to estimate the power consumption and then compare it with the measured power consumption over real drives. The results show that all estimates except for the wheel Odometry (WO) and LRF fall within one standard deviation of the measured results. This difference is not significant either, as it is easily explained. The calculated power uses the lab result for processing Xtion maps as the Xtion was running in this trial – but it was set to capture only and not process so the estimate is over. This close correlation between our power model and field trial verified runs is encouraging, as it validates the approach.

Table 3: Lab Power Estimates Compared to Field Trial Data

Navigation Mode	Power Estimate (from Lab, W)	Measured Mean Power (W)	Measured Std. Dev.
SSC low res. VO & Maps	22.6	24.5	4.15
Bumblebee VO & Maps	25.8	30.3	6.9
WO & LRF Maps	16.6	14.1	1.7
Bumblebee VO & LRF Maps	17.1	17.2	6.15

Taking these results and applying them to an analysis of optimal traverse modes for the Solo platform lets us investigate the potential gains from switching navigations sensor modes. Figure 8 presents a simple comparison of the gains possible by showing the difference in possible traverse distance for SOLO on flat ground in different sensor modes. We compare “traditional” rover navigation using visual methods to the cheapest method of blind drive on wheel Odometry and the cheapest “safe” method using a LRF to build navigation maps.

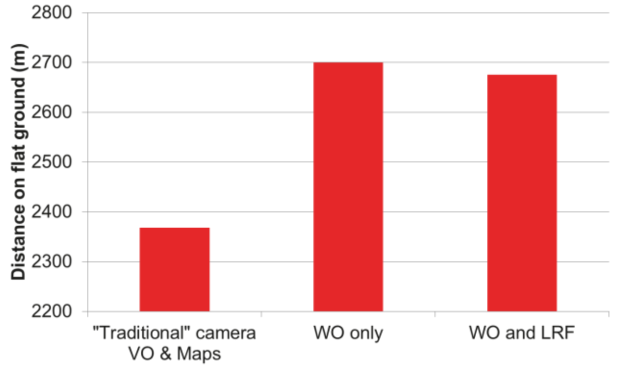

Figure 8: Comparing Traverse Modes. *Analysis of possible traverse distance on 100Whr at 0.3m/s extrapolated from field verified power consumption results.*

Analysing the UAV data captured we can only present preliminary results of processing, but they are encouraging. Figure 6 shows a 5mm resolution terrain model we were able to reconstruct from a flight at 5m altitude – this highlights the excellent model fidelity that would be impossible with a fixed wing aircraft.

6. CONCLUSIONS

The Chameleon trial represented a first in terms of the exploration of dynamic terrain sensitive rover navigation techniques for future rover missions. During the course of the work a large volume of ground and aerial sensor data and derived products such as high resolution maps were generated. The preliminary results indicate the feasibility of the concept indicating how dynamic selection of various navigation modes can significantly increase rover range.

Other fundamental research work in the area of science autonomy carried out by the team has also shown that the Chameleon algorithms could be used to also detect important science targets. For future work we seek to combine these two strands and develop an "eagle-eyed", but highly efficient rover through software and algorithmic means alone. This would enable a rover to both travel further in terms of metres per watt and detect science targets as it travels. This will require fundamental but high impact research and dedicated trials.

Another major issue is the organisation of large volumes field trial data for later use. The Chameleon trials have included a first-level ontology which we would like to develop in future trials to support greater use of the data. In addition to data organisation we are also considering new ways to visualise and interact with the various products to ensure greater understanding from an operational perspective.

7. REFERENCES

The authors would like to thank and acknowledge UKSA who funded this work under the CREST 2 programme.

8. REFERENCES

[1] J. J. Biesiadecki, P. C. Leger, and M. W. Maimone, "Tradeoffs Between Directed and Autonomous Driving on the Mars Exploration Rovers," *Int. J. Robot. Res.*, vol. 26, no. 1, pp. 91–104, Jan. 2007.

[2] G. Webster, "Mars Science Laboratory: NASA's Curiosity Mars Rover Approaches 'Cooperstown,'" *NASA's Curiosity Mars Rover Approaches "Cooperstown,"* 29-Oct-2013. [Online]. Available: http://mars.jpl.nasa.gov/msl/news/whatsnew/index.cfm?FuseAction=ShowNews&NewsID=1530. [Accessed: 09-Jan-2014].

[3] N. Silva, R. Lancaster, and J. Clemmet, "ExoMars Rover Vehicle Mobility Functional Architecture and Key Design Drivers," in *Proc. 12th Symposium on Advanced Space Technologies in Robotics and Automation*, ESTEC, Noordwijk, The Netherlands, 2013.

[4] A. Merlo, J. Larranaga, and Falkner, "SAMPLE FETCHING ROVER (SFR) FOR MSR," in *ASTRA*, ESTEC, Noordwijk, The Netherlands, 2011.

[5] E. Allouis, T. Patel, and A. Ratcliffe, "Sample Fetching Rover - Lightweight Rover Conceptsfor Mars Sample Return," in *ASTRA*, ESTEC, Noordwijk, The Netherlands, 2011.

[6] M. Woods, A. Shaw, I. Wallace, Malinowski, Mateusz, and Rendell, Phillip, "Simulating Remote Mars Rover Operations in the Atacama Desert for Future ESA Missions," presented at the 13th International Conference on Space Operations, 2014.

[7] M. Woods, A. Shaw, E. Tidey, B. Van Pham, U. Artan, B. Maddison, and G. Cross, "SEEKER-AUTONOMOUS LONG RANGE ROVER NAVIGATION FOR REMOTE EXPLORATION," in *I-SAIRAS 2012*, Torino, Italy, 2012.