EXOMARS ROVER VEHICLE PERCEPTION SYSTEM ARCHITECTURE AND TEST RESULTS

Kevin McManamon [1], Richard Lancaster [2], Nuno Silva [3]

[1] Astrium Ltd, Gunnels Wood Road, Stevenage, SG1 2AS, UK, Email: kevin.mcmanamon@astrium.eads.net
[2] Astrium Ltd, Gunnels Wood Road, Stevenage, SG1 2AS, UK, Email: richard.lancaster@astrium.eads.net
[3] Astrium Ltd, Gunnels Wood Road, Stevenage, SG1 2AS, UK, Email: nuno.silva@astrium.eads.net

ABSTRACT

The 2018 ESA ExoMars mission will land a rover on the Martian surface to search for signs of life, past or present. In order to maximise the number and variety of sites of scientific interest visited during the mission, the rover is designed to autonomously traverse previously unseen terrain with minimal assistance from Earth. To achieve this, the rover contains a highly autonomous navigation system capable of planning a safe and drivable path across the terrain. The ExoMars Rover Module is the responsibility of Thales Alenia Space. Astrium Ltd is responsible for the Rover Vehicle and the development of its navigation system. The front end of the navigation system is the *perception system* capable of analysing stereo images in order to create 3D models of the terrain in front of the rover. The source images are provided by stereo cameras mounted on the rover's mast. This paper presents the architecture of the rover's perception system and results from Astrium Mars Yard and numerical simulation testing.

1. INTRODUCTION TO THE ROVER MOBILITY SYSTEM

The purpose of the ExoMars rover (Fig. 1) is to transport scientific payloads across the Martian terrain to sites of scientific interest. Communication between the rover and the ground is via a Mars orbiter. However, the orbiter is only visible to the rover twice per day. Therefore, controlling the rover directly from the ground would severely limit the amount of terrain traversed by the rover and hence the total science return from the mission. To enable more rapid progress, the ExoMars Rover's Guidance, Navigation and Control (GNC) system uses vision-based methods to autonomously traverse Martian terrain without the aid of ground controllers. From a GNC perspective, the rover has the following key features:

- Total rover mass: 300kg
- Six independently driveable and steerable wheels
- Mast with a pan/tilt mechanism
- A pair of 'Navigation Cameras' placed on the mast's pan/tilt mechanism, approximately 2 metres above the terrain
- A pair of 'Localisation Cameras', approximately 1 metre above the terrain
- Three axis gyro
- Redundant three axis accelerometer
- Sun sensor

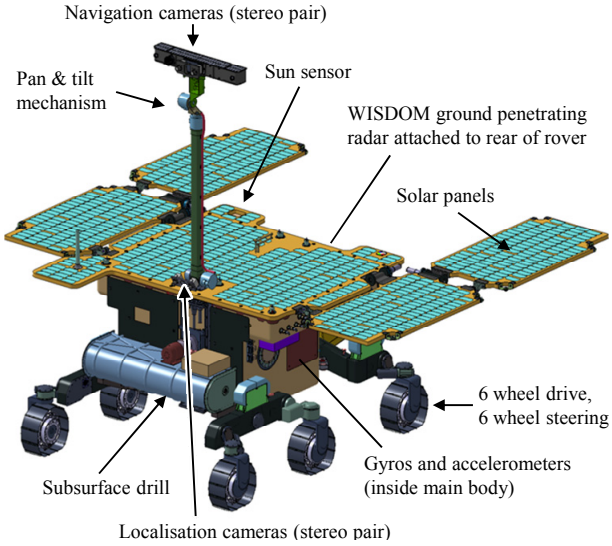

Figure 1. The ExoMars Rover Vehicle

Each day (sol) the ground controllers can command the rover to travel to a target that is approximately 70 metres away. The rover's GNC system then autonomously drives the rover to the target during the course of the day. It achieves this using its Navigation and Path Following functions.

1.1. Navigation Function

While stationary, the navigation cameras on the mast take stereo image pairs of the terrain in front of the rover. The rover's perception system then analyses these images to help generate a 3D model of the next 6 metres of observable terrain. The 3D terrain model is then analysed to determine which areas of the terrain are unsafe for the rover to traverse. For those areas of the terrain that are safe, it determines how challenging they will be for the rover. From this analysis it produces a "navigation map" which details where it is safe and unsafe for the rover to drive and the difficulty of the safe areas. The navigation system then plans a safe path, which is approximately 2 metres long, across the navigation map in a direction that should get the rover closer to its final target.

1.2. Path Following Function

The rover then drives along the planned path, its trajectory control system compensating for any small obstacles such as rocks or slopes that push it off the

path. Once it reaches the end of the path, it stops, and then repeats the process of navigation and path following until it reaches its target.

2. INTRODUCTION TO THE PERCEPTION SYSTEM

The role of the perception system is to analyse a pair of stereo images and produce a disparity map. A disparity map describes the apparent shift in corresponding pixels between a pair of stereo images, as illustrated in Fig. 2. Pixels corresponding to objects close to the cameras will exhibit a larger disparity than pixels corresponding to objects farther away (note that within Fig. 2 features closer to the cameras appear brighter within the disparity map, denoting larger disparity values). For each disparity map pixel, the magnitude of the disparity may be used to transform, through triangulation, the 2D pixel location into a 3D point in space. A 3D model of the terrain can therefore be generated in the later stages of the navigation pipeline (Fig. 3). The rover stereo cameras are separated by 150mm (the 'stereo baseline'). Each camera provides an image with a resolution of 1024 by 1024 pixels, and has a square field of view of 65 degrees.

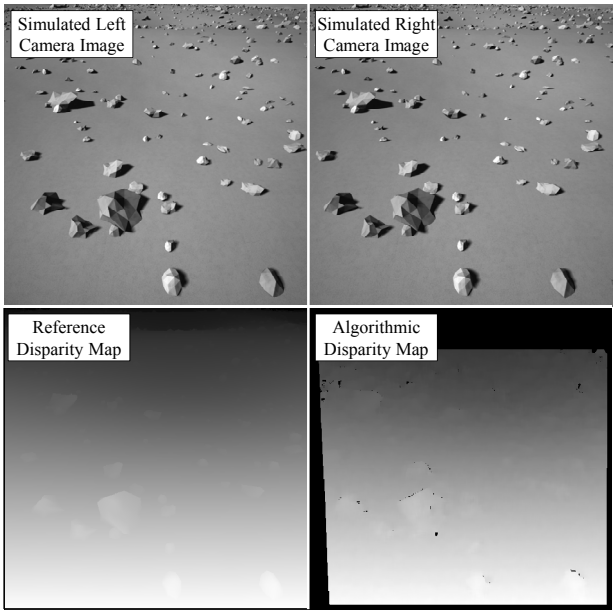

Figure 2. Reference and algorithmic disparity maps

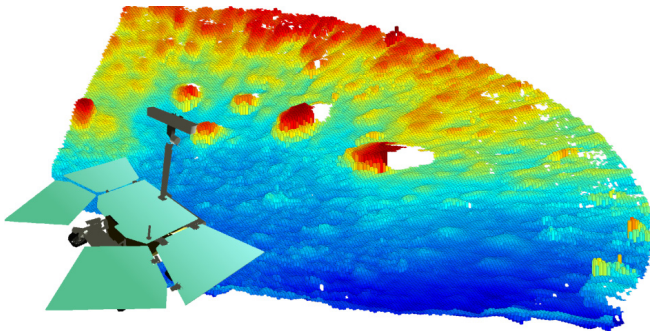

Figure 3. A 3D terrain model from breadboard testing

The rover perception system must satisfy the following requirements:

- The perception system must produce a disparity map accurate enough to adequately represent a specified region of terrain such that the navigation system can determine its safety. Approximately 6 metres of terrain in front of the rover must be sufficiently characterised by the output disparity map.
- The perception system must execute fast enough on the flight hardware in order to meet system level requirements for the distance the rover must be able to drive during one sol. Each call of the perception system must take less than 20 seconds to run on the ExoMars Rover 96MHz LEON2 co-processor.

The remainder of this paper will describe the design and algorithmic decisions made in order to meet these requirements.

3. EXOMARS PERCEPTION SYSTEM OVERVIEW

Fig. 4 provides an overview of the perception system architecture.

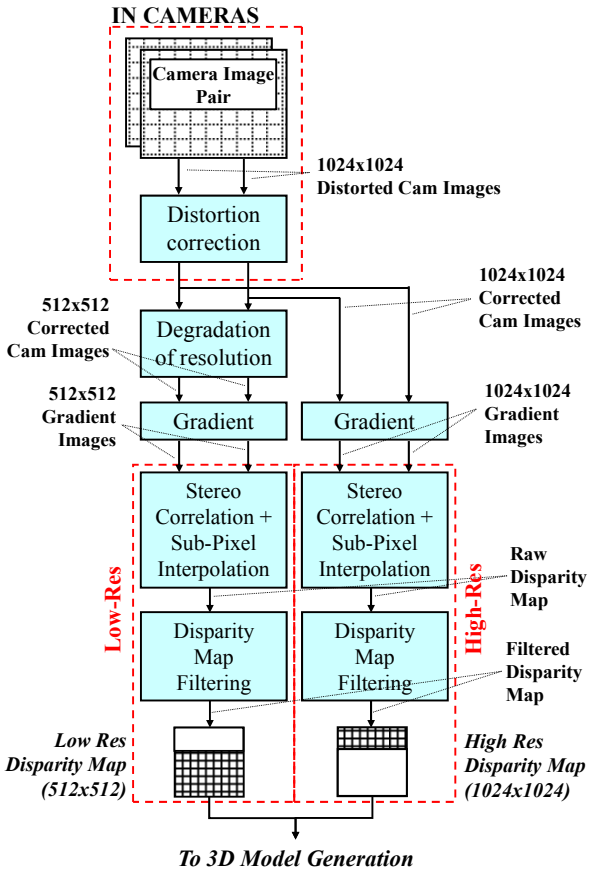

Figure 4. Perception system architecture overview

The following points summarise some key features of the perception system:

- Optical distortion effects are rectified within the rover cameras and are therefore not considered within the perception system.
- A multi-resolution approach is used to maximise the amount of terrain covered by the disparity maps whilst mitigating the adverse processing time implications of using high resolution images.
- When calculating the low-resolution section of the disparity map, the 1024 by 1024 pixel camera images are reduced to 512 by 512 pixels.
- A Laplacian of Gaussian (LOG) pre-processing filter is used to produce gradient images before the stereo-correlation algorithms attempt to calculate disparity values.
- A sum of absolute differences (SAD) correlation algorithm is used to scan the stereo images horizontally to detect disparity values to the nearest pixel.
- A linear sub-pixel interpolation method is used to calculate decimal disparity values.
- Disparity map filtering is used to remove any erroneous values from the raw disparity map to produce the final, filtered, multi-resolution disparity map.

3.1. Multi-Resolution Disparity Map

A disparity map with a resolution of 512 by 512 pixels has shown, through breadboard testing and simulation, to provide insufficient coverage in far-field regions of the visible terrain. However, solely increasing the total resolution to 1024 by 1024 pixels increases the processing time to unacceptable levels. Furthermore, it is unnecessary to increase the resolution for near-field regions, only far-field regions. Therefore, a multi-resolution approach is employed where regions of the disparity map covering the far-field use a higher resolution than the remainder of the disparity map, as illustrated in Fig. 5.

The border at which the algorithm switches between resolutions changes dynamically depending on the observed terrain. The stereo correlation algorithm initially scans row-by-row from the bottom of the 512 by 512 images towards the top. Once an entire row has been scanned and disparity values assigned, the algorithm calculates the average disparity value for that row. When the row's average disparity value drops below a predefined threshold, the algorithm switches to the higher 1024 by 1024 resolution. Using a similar approach, a further predefined threshold is used to stop the algorithm analysing regions of the images where the observed terrain is too far away to be significant for the navigation system. This appears as the 'Ignored' region within Fig. 5.

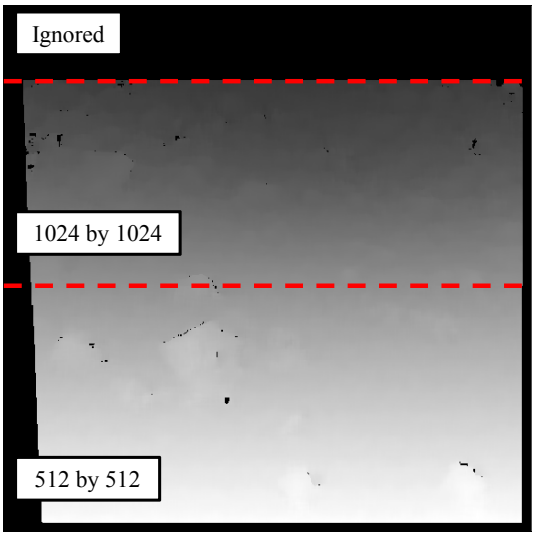

Figure 5. Multi-resolution disparity map

3.2. Gradient Image Filtering

In order to improve the robustness of the stereo correlation algorithms, a Laplacian of Gaussian (LOG) pre-filter is applied to the camera images. This filter reduces the influence of image noise on the correlation algorithms. It also helps to reduce matching errors caused by differing camera properties which would otherwise produce significantly dissimilar pixel values between the left and right stereo images. In addition to the LOG filter, further processing is carried out in order to improve the results of the subsequent stereo correlation process. Typically, rock edges produce large gradient values which tend to dominate the stereo correlation process, leading to significant errors around the edges of rocks. In order to mitigate these errors, further processing of the gradient images is carried out in an attempt to reduce the influence of these large gradients.

3.3. Stereo Correlation

In order to calculate the output disparity map it is necessary to find for each point P on the left image the corresponding point on the right image, as illustrated in Fig. 6. The method uses a correlation window, W, centred on the point P in the left image and tries to match this window to the corresponding area within the right image [1]. It does this by scanning horizontally across the right image assessing how well each right image correlation window, W_R, correlates with the original left image correlation window, W_L. The scan is constrained by defined minimum and maximum disparity values, d_{min} and d_{max}. A Sum of Absolute Differences (SAD) algorithm is used to assess how well two correlation windows match. Eq. 1 demonstrates how correlation windows are used to calculate a correlation criterion, C.

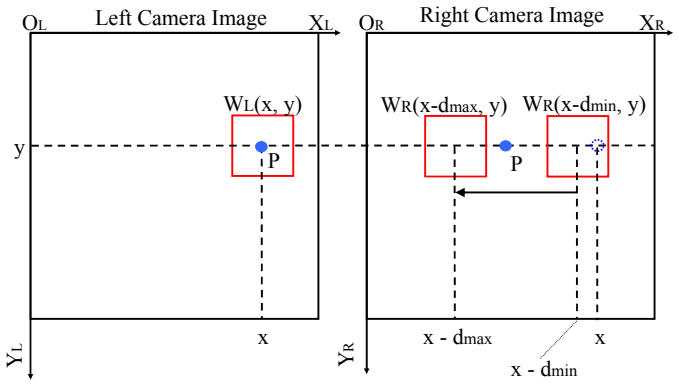

Figure 6. Stereo correlation process

Using SAD, lower values of C correspond to greater similarity between the two correlation windows.

$$C(x,y,d) = \sum_i \sum_j \left| I_1(x+i, y+j) - I_2(x-d+i, y+j) \right| \quad (1)$$

In Eq. 1, d is the disparity value, $I_1(x,y)$ and $I_2(x,y)$ are the intensity values in the left and right gradient images respectively at position (x,y), i and j are the column and row indices for the correlation windows.

By calculating the correlation criterion for each correlation window position in the right image (positions $W_R(x\text{-}d_{min}, y)$ to $W_R(x\text{-}d_{max}, y)$), a 'correlation function' is obtained, as illustrated in Fig. 7.

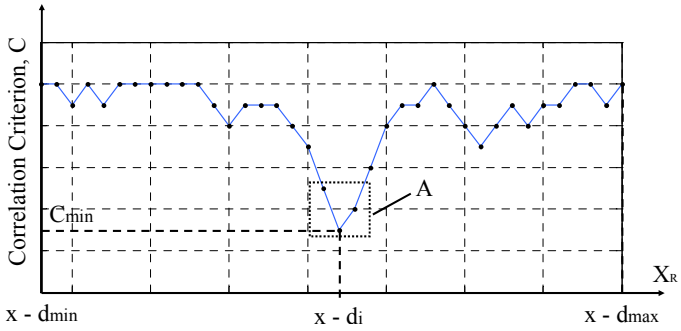

Figure 7. Correlation function

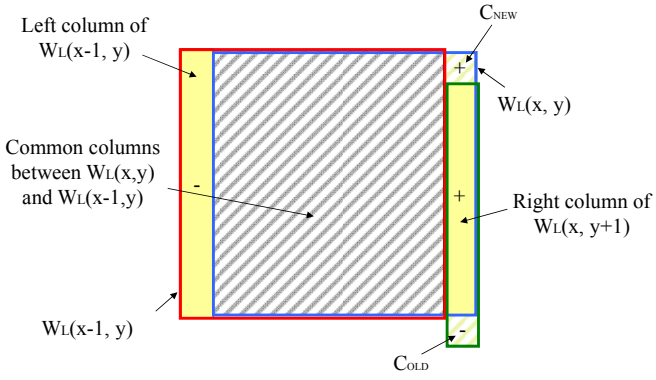

Figure 8. Possible correlation optimisation [1]

Note that for most correlation criteria, optimisations are available to speed up the calculation of C. The use of previous correlation window calculations allows the evaluation of Eq. 1 to be far more efficient [1], an example of which is illustrated in Fig. 8.

Fig. 8 shows how the correlation calculation for window $W_L(x, y)$ for a given disparity d uses the previously analysed column's correlation value ($W_L(x\text{-}1, y)$), the left column correlation value of $W_L(x\text{-}1, y)$, and the previously analysed row's right column correlation value ($W_L(x, y+1)$). The evaluation of C may therefore be expressed in an optimised form, as shown by Eq. 2.

$$C(x,y,d) = C(x-1,y,d) - C_{leftCol}(x-1,y,d) + C_{rightCol}(x,y+1,d) + C_{NEW} - C_{OLD} \quad (2)$$

For a 9 by 9 pixel correlation window, Eq. 2 requires two pixel correlation values to be calculated (C_{NEW} and C_{OLD}) compared to the 81 pixel correlation values required for Eq. 1.

3.4. Dynamic Disparity Search Range

For most stereo correlation algorithms of this type, the disparity search range (defined by d_{min} and d_{max} within Fig. 6) remains fixed across the entire image. The user must set appropriate values of d_{min} and d_{max} in order to correlate features within a desired range of the cameras. This usually means that a relatively wide disparity search is carried out for each pixel in the image. In order to improve the efficiency of the algorithms, the ExoMars perception system employs a dynamic disparity search range, where each row of the image defines its own values for d_{min} and d_{max}. This means that a far narrower search range is required for each image row, resulting in far fewer correlation criterion calculations. The use of this dynamic search range is one of the main algorithmic features which allows a multi-resolution disparity map to be calculated with acceptable processing times.

3.5. Sub-pixel Interpolation

The correlation function of Fig. 7 provides a minimum correlation value, C_{min}, for a certain disparity value, d_i. However, this disparity value is only accurate to the nearest integer pixel. Therefore, sub-pixel interpolation of the correlation function is carried out in order to obtain decimal values of disparity, d [5]. Note that the interpolation process only takes into account the three correlation values contained within region 'A' of Fig. 7. To illustrate, take the disparity values to the left and right of d_i within Fig. 7 to be d_{i-1} and d_{i+1} respectively. Also take the corresponding correlation criteria for d_{i-1} and d_{i+1} to be C_L and C_R respectively. The correlation function is interpolated linearly around its minimum

value C_{min}, by two lines with opposite slopes but with the same gradient magnitude. Fig. 9 illustrates the concept for two cases where $C_L > C_R$ and $C_L \leq C_R$. The linear sub-pixel interpolation relies on the assumption that $\alpha_1 = \alpha_2$ for both cases. The decimal disparity value d may be calculated by Eq. 3.

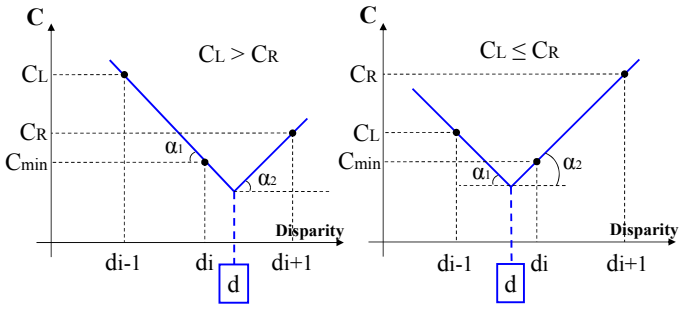

Figure 9. Linear sub-pixel interpolation [5]

$$d = \begin{cases} d_i + \dfrac{C_L - C_R}{2(C_L - C_{min})} & \text{if } C_L > C_R \\ d_i + \dfrac{C_L - C_R}{2(C_R - C_{min})} & \text{if } C_L \leq C_R \end{cases} \quad (3)$$

3.6. Disparity Map Filtering

When the stereo correlation and sub-pixel interpolation algorithms are applied to each pixel of the left camera image, a raw disparity map is generated. Erroneous disparity values will inevitably exist due to incorrect detection of C_{min}. Therefore the raw disparity map is filtering in an attempt to remove any erroneous disparity values and produce the final output disparity map. The filtering process works by assessing whether the calculated disparity values are consistent with each other and make physical sense. For example, if a small patch of the disparity map contains relatively large disparity values, and this patch is surrounded by a large region containing low disparity values, then this would represent some patch of ground 'floating' above the surrounding terrain. This would not make physical sense and so the smaller patch of larger disparity values would be invalidated.

3.7. Disparity Map Rover Masking

One final feature of the perception system is its ability to ignore any parts of the rover that are visible within the stereo camera images. If, for instance, the rover wheels were visible, the perception system may inadvertently represent them as obstacles immediately in front of the rover. To overcome this, the perception system receives a dynamically calculated disparity map mask indicating which pixels correspond to parts of the rover. The perception system then ignores these pixels during the disparity calculation process, allowing the subsequent 3D model generation to properly represent the terrain immediately in front of the rover.

4. PERCEPTION SYSTEM ACCURACY RESULTS

The ExoMars rover perception system has been tested using three different sources for the input stereo images:

1. Simulated stereo images originating from the Planet and Asteroid Natural scene Generation Utility (PANGU) software developed by the University of Dundee.
2. Real stereo images originating from the Locomotion Performance Model (LPM) breadboard rover in the Astrium Mars Yard [3].
3. Real stereo images originating from the NASA Mars Exploration Rovers (MER) taken from the surface of Mars [7].

The use of PANGU allows reference disparity maps to be produced which may then be compared to the disparity maps originating from the perception algorithms. This allows the accuracy of the perception system on simulated images to be quantified. No such reference data is available for either the LPM or MER stereo images, and therefore these tests are carried out purely to increase confidence that the system provides reasonable results on real images. The use of MER images is particularly important to test the algorithms on images taken in the true Martian environment. The relevant configuration parameters for the LPM, MER, and ExoMars rovers are provided in Tab. 1 for reference.

Table 1. Configuration parameters for the LPM, MER, and ExoMars rovers

Configuration Parameter	LPM	MER	ExoMars
Camera Height	2m	1.54m	2m
Focal Length	0.0048m	0.01467m	0.00426m (TBC)
Horizontal Field of View	77.3°	45.4°	65°
Vertical Field of View	65.2°	45.4°	65°
Stereo Baseline	0.1m	0.2m	0.15m
Pixel Size	6 x 6µm	12 x 12µm	5.3 x 5.3µm (TBC)
Image Resolution	1280 x 1024	1024 x 1024	1024 x 1024

4.1. Numerical Simulation Environment Results

PANGU is capable of providing high-resolution, stereo images of a simulated Martian terrain environment. It is designed to model rocks, slopes, craters, shadows, optical depth effects, fog, terrain texture, sky texture, lighting conditions, and surface reflectance. The images provided are free of any camera effects, such as noise, distortion, or lens artefacts. Therefore, a dedicated camera model is used to add realistic camera effects to the images originating from PANGU before passing them to the perception system. Using this simulation environment, it is possible to compare the algorithmic disparity maps produced by the perception system to

reference disparity maps originating from PANGU (an example reference disparity map is shown in Fig. 2). Such comparisons allow the errors to be viewed visually in the form of a so-called Error Disparity Map, which is calculated as the difference between the algorithmic and reference disparity maps. The numerical simulation environment allows the performance of the perception system to be quantified quickly based on numerous test cases. Quantifying the performance of the perception system at the disparity map level (i.e. in terms of pixel accuracy) allows the subsequent users of the perception data to propagate the errors. Therefore, the algorithms that generate the 3D terrain models can use the quantified perception errors along with other system errors (such as the accuracy of the pan and tilt unit) to determine the expected accuracy of the final terrain model.

There are two types of disparity error to consider: *mismatch* errors; and *estimation* errors. Mismatch errors are where the SAD correlation algorithms have incorrectly identified the integer disparity value (i.e. the value of C_{min} does not correspond to the true disparity). Typically the filtering algorithms invalidate these kinds of errors; however in some cases (notably at the edges of rocks) these errors may be present in the final output data. Estimation errors refer to cases where the correct integer disparity is calculated, but small errors are introduced due to the linear sub-pixel interpolation process. Tab. 2 presents the average test results from 100 random rover locations within the simulated terrain environment, including the standard deviation and absolute mean of the estimation errors for both the low and high resolution disparity maps. The relevant configuration parameters are provided in Tab. 3. These parameters have been derived from previous analysis and are expected to be similar to the final flight rover configuration. Note that the estimation errors tend to exhibit a normal distribution, as shown in Fig. 10.

Table 2. Average test results from 100 cases

Low Resolution			**High Resolution**		
% mismatch pixels	Estimation error σ (pixels)	Absolute mean estimation error (pixels)	% mismatch pixels	Estimation error σ (pixels)	Absolute mean estimation error (pixels)
1.35	0.0796	0.0605	2.02	0.2238	0.1592

Table 3. Perception testing configuration parameters

Configuration Parameter	**Values**
Height of Cameras Above Ground (m)	2
Field of View (degrees)	65
Stereo Baseline (m)	0.15
Higher Resolution (pixels)	1024 by 1024
Lower Resolution (pixels)	512 by 512
Low Res Correlation Window Size (pixels)	11 by 11
High Res Correlation Window Size (pixels)	25 by 25

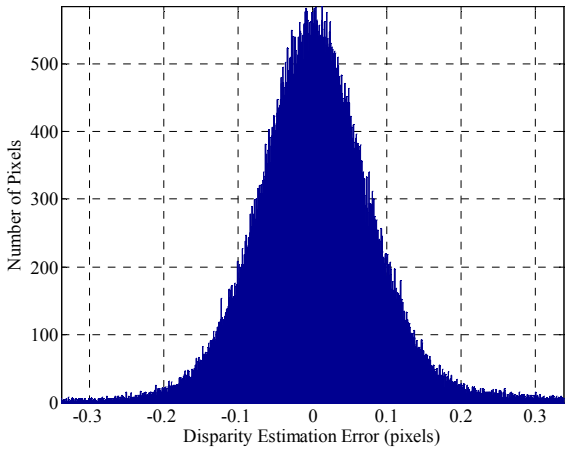

Figure 10. Normal distribution of estimation errors

The errors present within the perception system for simulated images have been propagated through to the 3D terrain model generation process. This propagation has shown that the perception system is capable of providing the 6 metres depth of terrain data necessary to meet the higher system-level requirements. Unfortunately, the 10cm stereo-baseline of the LPM breadboard rover cameras means that this cannot be tested on real images in the Mars Yard. However, future Mars Yard testing with a 15cm stereo baseline is due to take place in the near future, and testing with a 10cm baseline shows the expected performance.

4.2. Astrium Mars Yard Testing

The stereo images originating from the LPM breadboard rover navigation cameras are used for confidence testing of the perception system. It is currently not possible to obtain truth data regarding the terrain observed by the LPM cameras. Therefore, the algorithmic disparity maps are checked through inspection by comparing the 3D models produced by the rover GNC system to the measured physical characteristics of the Mars Yard (e.g. heights and spacing of rocks). Such testing has demonstrated that the perception system provides an accurate representation of the actual perceived terrain to the accuracy that it can be measured. The results from the LPM cameras are also compatible with the test results originating from simulated images. For instance, the types of terrains produced for rock fields are similar to those of simulated environments; the multi-resolution disparity map, the dynamic disparity search range, and the termination of the correlation algorithm when enough terrain has been analysed, all function correctly. An example of a terrain model generated from LPM breadboard images is given in Fig. 3. In addition to the specific testing in the Mars Yard, the perception system is also regularly exercised in the form of system-level testing, where the entire suite of GNC algorithms are executed together in closed-loop.

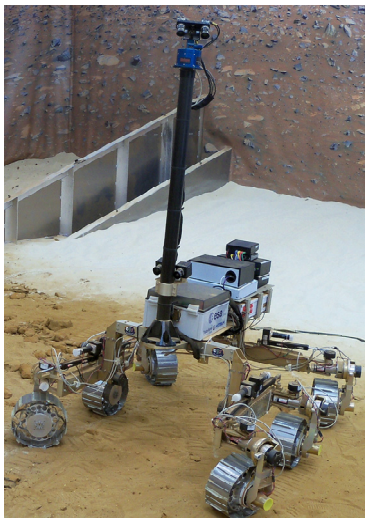

Figure 11. LPM breadboard rover in the Mars Yard

4.3. Mars Exploration Rover Image Testing

Similarly to the LPM testing, the use of MER stereo images helps to build confidence in the system's ability to cope with real image data. Reference [7] contains a vast catalogue of MER mission data, including distortion-corrected navigation camera images. Although it is not possible to derive even rudimentary truth data regarding the terrain observed by the MER cameras, their use allows the perception system to be tested on images acquired in the true Martian environment. A number of scenarios have been tested with MER image data, such as cases with high optical depth, low optical depth (leading to very harsh, dark shadows), cliff cases, and cases when the rover shadow covers a large portion of the acquired images. The 3D models derived from the disparity maps again demonstrate sensible and realistic terrain reconstruction. This again builds confidence in the perception system algorithms due to the presence of real Martian shadows, illumination, terrain texture, optical depth, etc. An example algorithmic disparity map from the MER stereo images is shown in Fig. 13, with the corresponding 3D terrain model shown in Fig. 12.

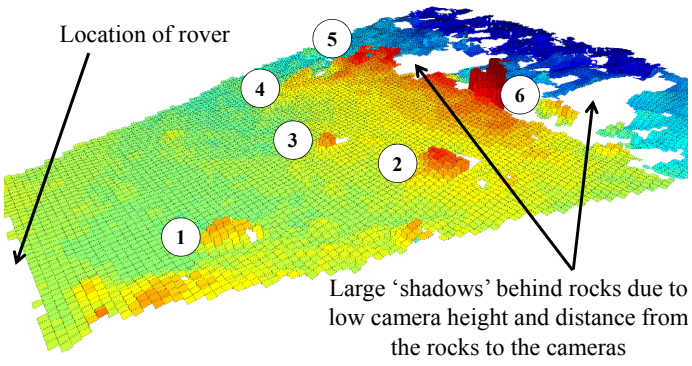

Figure 12. 3D terrain generated from MER images

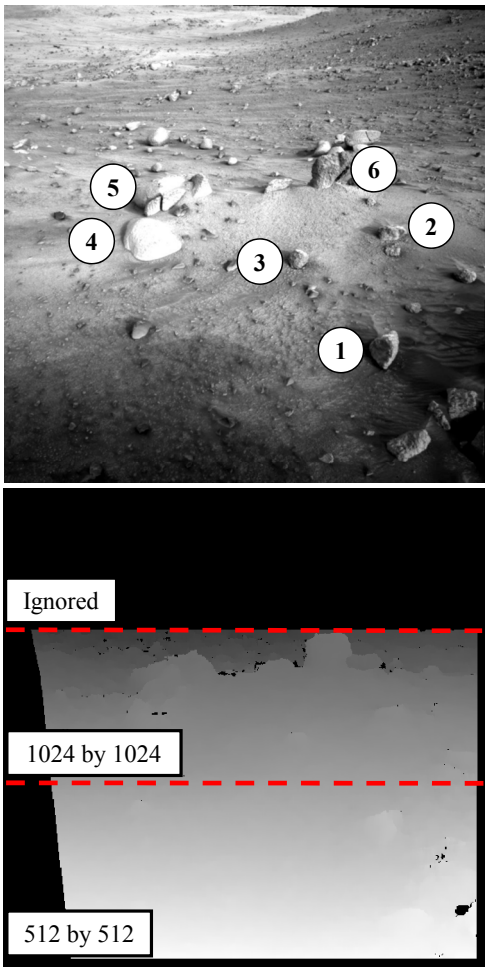

Figure 13. Disparity map from MER image testing

5. PERCEPTION PROCESSING TIME RESULTS

In flight, the perception system algorithms will execute on the rover's 96MHz LEON2 co-processor module. The execution speed of the perception system algorithms on this target hardware is important to the system level performance of the rover. The greater the time the rover spends stationary performing perception calculations, the less distance the rover can travel during a sol, hence the less useful science the rover can perform over the life of the mission.

As part of the ExoMars Rover program, a breadboard processing development environment based on a Pender Electronic Design GR-CPCI-XC4V board has been developed. The development environment includes a 70MHz LEON2 processor, 256Mb of SDRAM and version 4.6.2 of the RTEMS operating system. This makes it a suitable platform for benchmarking the execution speed of the perception algorithms, because although it does not exactly match the specifications of the rover's flight co-processor module, its design is similar enough that it is possible to perform analytical scaling of the execution time results based on the specification differences between the two platforms.

Each call of the perception system is required to take less than 20 seconds on the flight co-processor to produce a multi-resolution disparity map from the input stereo images. The perception system algorithms have been implemented in the C programming language. They have been cross compiled for the LEON2 processor on the Pender board, and wrapped in test harness code that runs on the RTEMS operating system. This test harness code allows the execution speed of the algorithms to be evaluated on the Pender board for given stereo pair inputs and configuration parameters, as illustrated in Fig. 14.

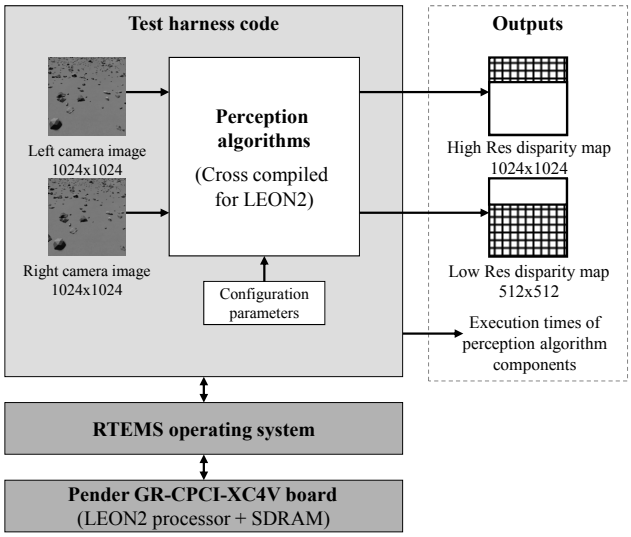

Figure 14. Execution time assessment environment

Real stereo images are used to assess the processing time performance of the perception system. Typical execution time results are provided in Tab. 4. These results have been scaled to represent the predicted processing time requirements on the 96MHz flight co-processor.

Table 4. Perception execution time (flight co-processor)

Perception Process		Time (s)
Perception Initialisation		0.20
Degradation of Image Resolution	Left Image:	0.08
	Right Image:	0.08
Create 512 by 512 Gradient Image	Left Image:	0.36
	Right Image:	0.37
Create 1024 by 1024 Gradient Image	Left Image:	0.61
	Right Image:	0.61
Stereo Correlation	Low Res:	1.85
	High Res:	5.48
Disparity Map Filtering	Low Res:	0.44
	High Res:	1.25
	TOTAL:	**11.35**
	Target:	20.00

6. CONCLUSIONS

This paper has summarised the design and architecture of the ExoMars rover perception system and presented results from tests carried out with simulated and real camera images. The disparity errors observed within the simulation test environment result in acceptable levels of accuracy to provide the 6 metres depth of terrain data necessary to meet the higher system-level requirements. The use of a multi-resolution disparity map and a dynamic disparity search range allows the perception system to provide greater coverage of the observed terrain whilst adhering to the system's processing time requirements. Furthermore, the use of real camera images in testing provides high confidence that the algorithms work outside of the simulation environment. In particular, tests on images taken from the MER rovers demonstrate that the algorithms produce sensible terrain models from data acquired within the true Martian environment. The perception algorithms are now being prepared for inclusion into the ExoMars rover's flight software.

7. REFERENCES

1. Faugeras, O. et al. (1993). Real Time Correlation-Based Stereo: Algorithm, Implementation and Applications. INRIA (Institut National De Recherche En Informatique Et En Automatique).

2. Silva, N., Lancaster, R. & Clemmet, J. (2013). ExoMars Rover Vehicle Mobility Functional Architecture and Key Design Drivers. *ASTRA Conference 2013.*

3. Meacham, P., Silva, N. &Lancaster, R. (2013). The Development of the Locomotion Performance Model (LPM) for the ExoMars Rover Vehicle. *ASTRA Conference 2013.*

4. Lancaster, R., Silva, N., Davies, A. & Clemmet, J. (2011). ExoMars Rover GNC Design and Development. *ESA GNC Conference 2011.*

5. Shimizu, M., & Okutomi, M. (2005). Sub-Pixel Estimation Error Cancellation on Area-Based Matching. *International Journal of Computer Vision 63(3), 207–224, 2005*.

6. Goldberg, S., Maimone, M. & Matthies, L (2002). Stereo Vision and Rover Navigation Software for Planetary Exploration. *IEEE Aerospace Conference Proceedings, March 2002.*

7. PDS Geosciences Node (MER Planetary Data Database). Online at *http://an.rsl.wustl.edu/mer/* (as of 12 April 2013).