PLANETARY SURFACE IMAGE GENERATION FOR TESTING FUTURE SPACE MISSIONS WITH PANGU

Iain M. Martin^{1*}, Martin N. Dunstan¹ and Manuel Sanchez Gestido², ¹University of Dundee, Dundee, UK, DD1 4HN, ² ESTEC, ESA, Noordwijk, The Netherlands. *[i.martin@dundee.ac.uk]

Abstract. Testing the image processing functions critical to the guidance systems of autonomous spacecraft landers for future space missions is a challenging task. One solution is to generate realistic synthetic images of the proposed planetary or asteroid surface to test the navigation and guidance systems in closed-loop tests. This paper presents a solution to this problem with the PANGU tool which creates multi-resolution models of planets and asteroids by enhancing the resolution of real terrain data with representative terrain and generates realistic, highquality, synthetic images in real-time using a custom GPU-based renderer which includes a parameterizable camera model.

Introduction.

Vision and LiDAR-based navigation systems can be used to support a wide variety of autonomous space applications including asteroid approach and landing missions¹, planetary landers, and surface rovers. Verifying and validating these mission critical systems and their underlying image-processing algorithms² is a challenging task that can be addressed by a test environment which can generate realistic images of the scenario to simulate the vision or LiDAR instruments used for autonomous navigation³.

It is possible to use physical mock-ups to generate simulated images for testing vision-guided landers such as a camera mounted on a robotic arm taking images of a physical model of the expected terrain or a wall-mounted high-resolution image such as the Testbed for Robotic Optical Navigation⁴ (TRON) facility. These physical mockups can be used to generate images in real-time for closed-loop testing, but they are expensive, inflexible and have difficulty simulating realistic lighting and atmospheric conditions. Virtual simulations are therefore an increasingly valuable tool to support testing and verifying future autonomous approach and landing systems⁵.

PANGU.

PANGU (Planet and Asteroid Natural Scene Generation Utility) is a software package designed to aid the testing and verification of the approach and landing of autonomous spacecraft missions by providing realistic simulations of onboard vision and LiDAR guidance sensors. PANGU runs on standard, modern PCs but will have better performance on powerful PCs with fast GPUs and plenty of memory. PANGU contains surface modelling tools, a custom-designed visualization tool, flexible integration options to interface with other simulators, and a NAIF/SPICE interface⁶ which can be used to accurately simulate historical events or future mission missions where SPICE data is available. It provides a high degree of realism while producing images from large models in real-time at frame rates expected of navigation cameras on a planetary lander. These images can be used in a variety of test setups including off-line open loop, closed-loop and hardwarein-the-loop. PANGU functionality and the realism of the simulated images have been extensively verified and validated in preparation for use in all phases of operational missions.



Figure 1: Multi-resolution PANGU lunar model with flight path.

Surface modelling.

PANGU models can be purely synthetic or instead based on real elevation data of whole planets, surface terrain patches and asteroid-type bodies. Patches of terrain can be based on imported DEMs (Digital Elevation Models) and whole planet models can be based on imported radial height DEMs where available. The landing site region can then be artificially enhanced or combined with a higher resolution DEM of the landing site region if one is available. Surface data can be imported in standard formats and the resolution of the models increased along the proposed flight path of the lander to enable the resolution range required to simulate realistic approach and landing from orbit to surface touch down. An example image from a large, lunar south pole model with enhanced terrain along a landing trajectory is shown in Figure 1.

The terrain resolution can be artificially increased through adding additional high-frequency terrain using standard fractal techniques. Small features such as realistic age-degraded impact craters, boulders, mounds and dunes can all be added in customizable size-density distributions as appropriate to simulate different surface types commonly found on rocky planetary bodies and Mars. Color maps can be imported to represent albedo variations across the model and spacecraft and rover models can be also imported and included in simulations if appropriate.

An important feature of PANGU that enables it to realistically enhance the surface of rocky planetary surfaces is the crater model⁷ which implements realistically overlapping, age-degraded craters which enables PANGU to create models that are accurate representations of cratered surfaces with crater size-frequency and age distributions created randomly from statistical distributions. The crater model was derived from Melosh's studies on impact crater⁸ and incorporates customizable crater degradation and rim irregularity distributions, models overlapping craters and can realistically integrate craters into existing terrain⁷.

Multi-resolution model.

An image from a section of the whole planet, multi-resolution model of Mars is shown in *Figure 2* with the low-resolution outer section defined from MOLA data which has been seamlessly stitched to the much higher resolution HRSC DEM⁹ and projected to match the MOLA frame of reference. A section of terrain around the landing site region has been synthetically enhanced by PANGU with horizontal resolution



Figure 3: Multi-resolution model of a Mars landing site

increased to around 40 cm. Distributions of craters and boulders smaller than the resolution of the base DEM data are generated for each resolution section so that the additional craters and boulders accurately represent Mars feature distributions across the multi-resolution model.

Small-bodies.

Small bodies such as asteroids are supported through either importing shape models in standard ICQ (Implicitly Connected Quadrilateral) or WaveFront OBJ formats or by generating entirely synthetic shape models from Perlin or Simplex noise, then increasing the model resolution of the imported shape model with added highfrequency terrain and small-scale surface features. Artificial or imported color maps can be used to represent surface color variations across the small body. *Figure* 3 shows an example of an imported shape model of Phobos¹⁰, with the resolution of the model artificially increased by PANGU with small craters and boulders added in realistic size-density distributions.



Figure 2: PANGU enhanced simulation of Phobos

Rendering.

Images are generated using a custom GPU-based renderer (Graphics Processing Unit) that includes simulated camera distortion through the integrated, parameterizable camera model. LiDAR simulations are also supported.

Image generation.

The PANGU renderer was custom designed to generate realistic images for simulating planetary surfaces. It can render images from large, multi-resolution models at fast frame rates and can render images from multiple planetary body objects and imported CAD models of spacecraft, satellites or surface rovers. PANGU surfaces and boulders can be rendered with different material properties, to allow different lighting models such as Lambertian, Oren-Nayar and Hapke to be applied to different surfaces and boulder types, tuned by specifying the BRDF (Bidirectional Reflectance Distribution Function) parameters. An example is shown in Figure 9 of boulders on a simulated lunar surface with different material properties to demonstrate boulders with different reflective properties for the different types of boulders and with different reflective properties to the surface.



Figure 4: Rolling shutter effect of a rotating camera

Image generation is split (internally) into two steps. First, a high-resolution pinhole camera radiance image is generated with user-controlled lighting and reflection models, shadows, material BRDF and atmosphere effects without any noise or detector effects. This is followed by an optional camera modelling step which samples the radiance images to model the chosen detector with different noise sources. An example of camera model distortion is shown in Figure 4 which shows the effect of apply rolling shutter effect of a rotating camera to a lunar surface model.

Camera Model

PANGU contains a sophisticated, integrated, parameterizable, camera model which runs in GPU shader code to enable complex camera model distortions to be applied without significantly lowering the image generation frame rate. The camera model simulates the transfer of photons through the optical system to the detector, the conversion to photo-electrons and then ADC (Analog-to-Digital) conversion to a digital number. It contains realistic models of photon and electron shot noise, read-out noise, radiation, saturation and blooming, and various per-pixel non-uniformities. The camera model uses the GPU on the graphics card for maximum performance. The high-level architecture of the PANGU detector model is shown in Figure 5.



Figure 5: Overview of PANGU detector model architecture

Radiance model

The effects of radiation, due to protons and electrons, is modelled in detail to support simulations in high (or low) radiation environments such as at Jupiter or around the Earth. For each particle (based on user-defined flux and power-law) the energy and direction is computed. The stopping power of silicon, the 3D dimensions of the pixel, and the number of electrons liberated per eV are used to determine the number of electrons which are added to each pixel. The model takes into account that as some of the energy of a particle is deposited in a pixel the stopping power changes and thus the energy available to be deposited along the rest of the path.

The inclusion of the camera model does not significantly increase the image generation time with a rendering rate of 10 Hz or better achievable on the large example missions scenario models included with PANGU which simulate a lunar south pole lander, a Mars approach and landing inside Gale crater and a Phobos approach and landing. Recent additions to the camera model include detector smear, scattered light and tangential distortion.

Detector smear

Detector smear is a camera artefact that may be observed in real cameras so simulating this feature has been added as an additional optional camera model feature. An example of the effects of smear are shown in the contrast-stretched image on the left of Figure 11 from Itokawa. Note that this is a scaled screen capture not a raw image, hence the pixilation and JPEG ringing artefacts. The image in Figure 11 is a PANGU model of Itokawa from a similar camera position with the PANGU v5 smear model enabled and configured using the reported Hayabusa/AMICA parameters.

Before an exposure begins, a camera CCD array is reset by repeatedly shifting the pixels down the columns while introducing zero charge to the first pixel. After the image has been exposed, the CCD shifts each pixel along to the readout area. Without a global shutter, light entering the detector will fall on each pixel as it shifts across the array causing smear. In PANGU the simulated smear operation is characterized by whether it applies (equally) to reset and readout or whether it applies to just one. These are referred to "same" and "different" modes respectively. Smear is also characterized by the readout direction (down/up/left/right).

For the "same" mode the down/up readout directions have the same smear contribution; similarly, the left/right directions are identical. For a given pixel in the readout column (or row) the smear is the sum of all radiances incident on that column multiplied by the standard factors such as exposure time during the reset/readout steps ($\Delta \tau$), quantum efficiency etc.

GPU shaders operate in "gather" mode which means they can read/compute over multiple input pixels to determine the value of the current (output) pixel. Since the contribution to each pixel in a given column (or row) is identical it would be inefficient for each pixel to sum of its column (or row) because all the other pixels in its column (or row) will repeat the calculation. For a 1024×1024 pixel detector this means that the smear computation would be repeated 1024 times more than necessary.

As an optimization, a "smear sum" pass is performed in which the detector image is divided vertically into k segments where k is a small value such as 8. Each pixel in the first row of each segment is assigned the sum of all the pixels in the same column of that segment. This means that each pixel in the input image contributes to one pixel in one of the k output "sum" rows. This smear sum can be used in the main smear shader both for the "same" and "different" modes.

For the "same" mode, the smear contribution of pixel (x, y) is the sum of the k "sum" pixels from the same column x. Although this causes each pixel in a column to perform the same summation as the other pixels in the column, the cost is k times the ideal minimum cost instead of a much a larger factor. Given the number of shader units in modern GPUs the impact is expected to be minimal. Also, it is possible to use k=1 for the minimum unit cost.

Examples of extreme read-out and reset smear on an Earth model are shown in Figure 13.

Scattered light

The multiple-weighted Gaussian PSF feature models scattered light within the optical system which adds a small amount of widely radiance spread over a wide area. The AMICA camera calibration report for the Hayabusa mission¹¹ states that the scattered light in the raw images can be modelled as a weighted sum of Gaussian PSFs with weights of the order of 10⁻³ and standard deviations of 8–710 pixels. This system can be modelled in PANGU to add scattered light to the "clean" radiance images and obtain images with representative scattered light.

Examples of real (AMICA) and PANGU simulation of scattered light is shown in Figure 12. The top left image is a real AMICA image of Itokawa while the top right image is a contrast stretched version showing the effects of scattered light. The bottom image images are from PANGU.

As with the existing PSF implementation, the ability to decompose a 2D Gaussian PSF into separate 1D Gaussian PSFs in the horizontal and vertical directions greatly reduces the runtime cost of the scattered light implementation. A 2D Gaussian with standard deviation of 710 pixels applied with a window of two standard deviations has dimensions 1420×1420 i.e. covers 2 million pixels. This is applied to each pixel of the radiance image: for a typical 4096×4096 radiance image with nearly 17 trillion pixels so a naïve 2D operation the Gaussian exponential must be evaluated 34 million times.

In contrast the 1D decomposition requires two passes with windows of 1420 pixels for a total of 48 billion operations i.e. $710 \times \text{less work}$.

Tangential distortion

Tangential distortion is unknown for any space-based camera studied for PANGU so no real-world examples of tangential distortion in space cameras are available. Unlike the radial distortion model which can represent (with different coefficients) the mapping in both directions between distorted and undistorted points, the asymmetric form of the tangential distortion correction model means that it cannot represent the inverse transformation which maps undistorted points to distorted points. It is also non-trivial to invert. However, PANGU only needs to map from distorted points to undistorted points so the inverse transformation can be ignored.

This means that although the PANGU radial distortion coefficients are for the undistorted-to-distorted direction, the tangential distortion coefficients must be for distorted-to-undistorted. This is not a major issue, but users will need to take care. However, in practice the tangential distortion coefficients will almost certainly never be used it is unlikely that the difference will be noticed. An example of a face-on image of a checkerboard with exaggerated tangential distortion is shown in Figure 6.



Running Simulations.

Open-loop simulations can be run with pre-defined trajectories and model dynamics in flight files. Closed-loop testing is supported through a server mode which will render images in response to TCP/IP socket commands. This enables PANGU to be integrated with

2nd RPI Space imaging Workshop. Saratoga Springs, NY. 28-30 October 2019.

external systems such as Matlab and SIMULINK or built into custom applications calling C or Java client APIs.

NAIF / SPICE

The NAIF/SPICE system⁶ has been integrated into PANGU to enable ephemerides to be used for accurate simulations of historical events as well as predicted future events. If a SPICE ephemerides library is available for a previous or future mission, SPICE can provide the position and orientation of the spacecraft camera and the relevant physical bodies such as the sun, a planet, any moons in view or an asteroid) from only the time of the event This lets us generate accurate renderings of given view from the time of the event. Examples of this are shown in Figure 7 recreating a Clementine UVVIS image, Figure 10 recreating an AMICA image of asteroid Itokawa, and Figure 14 where a NEAR/MSI image of Eros has been overlaid on a large PANGU image of the same view. In each case PANGU used SPICE to set the positions of the Sun, the target body and the camera.

Video generation

Video generation based on ffmpeg 4.0.2 (or similar versions) has been added to PANGU to allow users to generate high-quality MP4 videos from a sequence of previously saved PANGU images, or directly as the images are rendered. The ffmpeg system is freely available on different systems and is able to generate videos in a wide variety of formats.

Results.

An example that demonstrates the use of PANGU to enhance a low-resolution model into a more detailed model for simulating a spacecraft approach is shown in Figure 16 which contains three PANGU images of Deimos, the smaller Moon of Mars. The image on the top-left is the imported low-polygon OBJ shape model without any enhancements. On the top-right, is the result of PANGU enhancements and the image on the bottom is the enhanced model with camera model smear and cosmic ray radiation events enabled.

The imported Deimos shape model¹² was obtained in WaveFront (OBJ) format. It was imported using PANGU and converted to PANGU's internally used Implicitly Connect Quadrilateral (ICQ)¹⁰ shape model format. It has an approximate resolution of ~150 m between vertices in the shape model. The crater size-frequency distribution was defined as:

$$N_D = 0.1D^{-2.2}, 100 \le D \le 4000,$$

where N_D is number of craters per unit area with diameter greater than or equal to D. From this distribution, PANGU can generate an appropriate list of

craters for a given model surface area, specified as 495 km^2 in our model.

The low-resolution shape model has had the resolution synthetically enhanced by a factor of eight by adding high-frequency Perlin noise terrain, with an amplitude of 100m, to an interpolated shape model giving a final resolution of around 15m. A synthetic reddish albedo map generated by PANGU is applied, and boulders with different Hapke BRDF properties to the surface have been added with boulders defined in a similar way to the craters from boulder size-density and depth-density distributions.

The albedo map is generated synthetically using Perlin noise with colors in HSV (Hue, Saturation and Value) which allows us to constrain the color variation, saturation and brightness before converted the image back to RGB mode and applying it to the shape model. The initial shape model is 792KB in size and the PANGU enhanced model is 73MB, with each further doubling of resolution increasing the model size by approximately a factor of four. PANGU has been successfully tested with models up to 80GB in size on a high-end PC.

Validation, verification and performance.

PANGU has been validated throughout its development through modelling and rendering comparisons with real images and by review by independent planetary scientists who gave a detailed review of the realism of PANGU models of Mars, the Moon and asteroids¹³.

An example PANGU rendering comparison is shown in Figure 7 which shows a real image of the Moon (lue0222a.034, 1994-02-27T07:06:54.599) from the Clementine UVVIS¹⁴ instrument compared to the PANGU rendered image using the UVVIS camera model and a surface model created from a 120 m resolution digital elevation model (DEM) from Lunar Reconnaissance Orbiter¹⁵ centred at 85.68S, 6.06E. The images appear visually very similar and a pixel intensity histogram comparison (Figure 8) is also similar except for the very dark areas of the image which will contain scattered light in the shadowed area (not modelled by PANGU).

The target rendering speed for PANGU is 10 Hz on a modern PC with a dedicated graphics card. This is based on the real-time requirements for vision-based image processing systems that have used PANGU for testing. The achievable rendering performance depends heavily on the complexity and size of the scenario. Dynamic shadows (important for rapidly changing scenes such as asteroid rendezvous) and the GPU-based atmosphere model (important for Earth orbit, Mars and Titan) tend to Two examples of rendering performance are shown in Figure 16 and Figure 17. In Figure 16 the results of a recreation of the Hayabusa/AMICA orbit of Itokawa from on 10-Jan 2005 are shown. The PANGU orbit covers the period 03:44:57 to 18:57:59 with 1000 images (similar to Figure 10) taken every 50 seconds to obtain a realistic benchmark. The test was performed under Windows 7 on a PC with an Intel Xeon CPU E5-2630 v3 at 2.40GHz, an NVIDIA Quadro K5200 8GB graphics card and 128GB of RAM. The full camera model was enabled along with dynamic shadows which were necessary due to the changing Sun geometry. The average frame time was around 40 ms.

In Figure 17 the results of a recreation of the MSL/MARDI (Curiosity rover) descent sequence onto Mars are shown using Windows 7 on the same PC as the Itokawa orbit. The terrain model is the same as that shown in Figure 2 with the Hapke BRDF. The GPU-based atmosphere model was used along with MARDI camera model. Shadows were not enabled since the lighting does not change substantially during the 7 minute sequence. When the camera is high above the model with it filling the field of view images took 200–300 ms to generate i.e. slower than the desired 10 Hz. In the final stages of descent when the camera is near to the surface rendering speeds were much faster.

Acknowledgements.

PANGU was developed by the University of Dundee for ESA and is being used on many European activities aimed at producing precise, robust planetary lander and rover guidance systems. Further details can be found from the PANGU website portal¹⁶.

Conclusions.

PANGU has been shown to include a range of planet and asteroid modelling and features and image generation capability to simulate past and future spacecraft lander missions. The capability to generate large, multi-resolution models by combining real and synthetic data enables realistic simulations to be generated from orbit to surface approach and landing. The realism and performance of PANGU simulated images enables real-time, hardware-in-the-loop simulation to be used as a test and validation tool for future spacecraft lander missions.

The support for the NAIF/SPICE system enables accurate simulations of previous missions which has been used to verify the accuracy and realism of comparable PANGU images which are similar, but not identical to, the real images. However, they have been shown to be representative of real images so are suitable for testing and developing future spacecraft guidance, navigation, approach and landing missions.

References.

 A. Bidaux-Sokolowski, J. Lisowski, P. Kicam, O. Dubois-Matra and T. Voirin, "GNC Design for Pinpoint Landing on Phobos", In Proc. ESA Conference on Guidance Navigation & Control Systems (GNC), Salzburg, Austria. 2017.

[2] N. Rowell, S. Parkes, and M. Dunstan, "Image Processing for Near Earth Object based Vision Guidance System," IEEE Trans. Aeropspace and Electronic Systems, vol.49, no.2, pp.1057-1072, 2013.

[3] M. Dunstan, K. Hornbostel, "Image processing chip for relative navigation for lunar landing", 9th International ESA Conference on Guidance, Navigation, and Control Systems, Porto, Portugal, 2014.

[4] Krüger, H., Theil, S., TRON - Hardware-in-the-Loop Test Facility for Lunar Descent and Landing Optical Navigation,

IFAC Proceedings, Volume 43, Issue 15, 2010.

[5] R. Fisackerly, J. Carpenter, D. De Rosa, A. Pradier, C. Philippe, and B. Gardini and "The European Lunar Lander: Robotic Operations in a Harsh Environment", Moon, 2011.

[6] Acton, C., Bachman, N., Folkner, W. & Hilton, J. (2015). SPICE as an IAU Recommendation for Planetary Ephemerides, IAU General Assembly, 22, 2015.

[7] Martin, I., Parkes, S. & Dunstan, M., (2014). Modeling Planetary Surfaces with Real and Synthetic Terrain, IEEE Trans. Aerospace and Electronic Systems, vol. 50, no. 4.

[8] Melosh, H., Impact Cratering – A Geologic Process, Oxford University Press, 1989.

[9] Gwinner, K. et al., (2016). The High Resolution Stereo Camera (HRSC) of Mars Express and its approach to science analysis and mapping for Mars and its satellites, Planetary and Space Science, Vol. 126.

[10] R. Gaskell, "Phobos Shape Model V1.0", VO1-SA-VISA/VISB-5-PHOBOSSHAPE-V1.0, NASA Planetary Data System, 2011.

[11] Ishiguro, M., et al., The Hayabusa Spacecraft Asteroid Multi-Band Imaging Camera: AMICA, Icarus 207, 714–731, 2010.

[12] Deimos shape model, <u>https://3d-asteroids.space/moons/M2-</u> Deimos. Accessed on July 2019.

[13] Martin, I., Dunstan, M., Parkes, S., Sanchez-Gestido, M. and Ortega, G., "Simulating Planetary Approach and Landing to Test and Verify Autonomous Navigation and Guidance Systems", 10th International ESA Conference on Guidance, Navigation, and Control Systems, Salzburg, Austria, 2017.

[14] Malaret, E., Perez, L. and Taylor, H., "Clementine's in-flight calibration results (UVVIS camera)", ACT Final Report, NASW-5014, 1998.

[15] H. Riris et al, "The Lunar Orbiter Laser Altimeter (LOLA) on NASA's Lunar Reconnaissance Orbiter (LRO) Mission," in Conference on Lasers and Electro-Optics/International Quantum Electronics Conference, Optical Society of America, 2009, paper CFJ1.

[16] www.pangu.software

Supplemental images.



Figure 7: Clementine image (left) and the simulated PANGU equivalent image (right) from a 120m LRO DEM



Figure 8: Clementine / PANGU image comparison pixel intensity histogram



Figure 9: Boulder BRDF materials in PANGU



Figure 10: PANGU Itokawa model (I), with synthetic albedo (c) and real AMICA image (r)



Figure 11: COSPAR image of Itokawa (left) and the PANGU simulation demonstrating camera model smear



Figure 12: Real AMICA images of Itokawa (top) and PANGU images (bottom)



Figure 13: Examples of readout/reset smear in a PANGU image of the Earth



Figure 14: Still from PANGU-generated video with real image overlay and labels



Figure 15: Deimos imported OBJ (top left), the PANGU enhanced version (top right) and the PANGU enhanced version with camera model radiation effects and smear (bottom)



Figure 16: Rendering times for orbit of Itokawa model with dynamic shadows and full camera model (Windows 7)



Figure 17: Rendering times for MSL/MARDI descent sequence, full camera model, atmosphere, no shadows (Windows 7)