Utilizing the Open Sun Grid Engine in the OSIRIS-REx Altimetry Pipeline. R. C. Espiritu[1], A. H. Nair[1], L. Nguyen[1], O. S. Barnouin[1], E. Palmer[2], B. Gaskell[2], J. Weirich[2], M. Daly[3], J. Seabrook[3], M. E. Perry[1] C. L. Johnson[4,2], M. Al Asad[4], H. C. M. Susorney[4], J. H. Roberts[1], M. C. Nolan[5], D. S. Lauretta[5], and the OSIRIS-REx Team, [1]The Johns Hopkins University Applied Physics Laboratory, Laurel, MD, USA (Raymond.Espiritu@jhuapl.edu); [2]Planetary Science Institute, Tucson, AZ; [3]The Centre for Research in Earth and Space Science, York University, Toronto, Ontario, Canada; [4]Department of Earth, Ocean, and Atmospheric Sciences, University of British Columbia, Vancouver, Canada; [5]Lunar and Planetary Laboratory, University of Arizona, Tucson, AZ, USA.

Introduction: The Origins, Spectral Interpretation, Resource Identification, and Security–Regolith Explorer (OSIRIS-REx) spacecraft arrived at the asteroid (101955) Bennu in December 2018. Imaging and laser altimetry are being used to construct shape models of the asteroid for use in identifying a tag site at which material will be sampled from the surface of Bennu and returned to Earth [1]. The Altimetry Working Group (AltWG) is tasked with the generation of the shape models as well as other ancillary products created using gravity modeling algorithms and statistical analysis of tilts for individual facets in the model. Shape models and ancillary products are generated at various resolutions in order to facilitate coarse site selection as well as detailed site analysis. Finally, products must also be generated in a timely manner for use by other Working Groups.

The Sun Open Grid Engine was used to parallelize computations of gravitational models and production of data products in order to meet the processing load and time constraints.

Products: The AltWG is responsible for producing local and global digital terrain maps (DTMs) for science objectives [2] and for Natural Feature Tracking (NFT) used to guide OSIRIS-REx to the sample location on Bennu's surface [3]. The products are generated from two independent processes and datasets: initial global shape models created by stereophotoclinometry (SPC) processing of images [4] taken by the OSIRIS-REx Camera Suite (OCAMS) [5] and by the analysis of range data taken by the OSIRIS-REx Laser Altimeter (OLA) [6]. Five initial models are created with a fixed number of facets (see Table 1). The initial models are converted into (DTMs) in FITS file format. The DTMs are represented within the FITS file as a set of 2D image arrays where the pixel position corresponds to a point on the DTM surface and the pixel value is the data value at that point.

AltWG also generates a suite of ancillary data associated with each DTM. These ancillary data include gravity vector, gravitational potential, slope, elevation, and tilt maps [7]. The ancillary data are stored as binary tables in FITS format, with table columns storing the various parameters of each facet of the associated shape model. Figure 1 shows an image of an elevation ancillary fits file generated from a global shape model with 80cm ground sample distance.

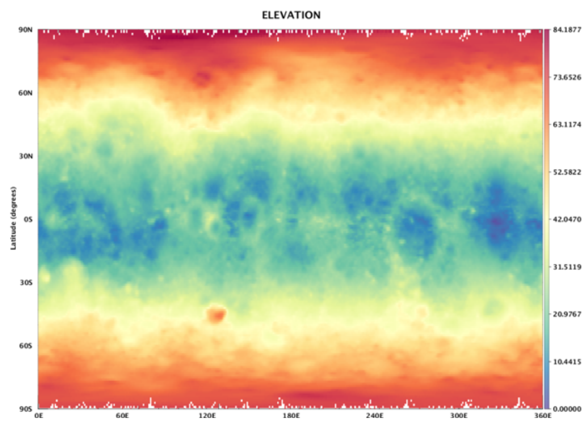

Figure 1 Elevation (m) for highest resolution global shape model

Pipeline: The Open Sun Grid Engine is utilized to perform certain tasks in the AltWG product generation pipeline. The tasks generally fall into two categories. The first category is the use of the grid engine to split a processing task consisting of a large number of iterations into several smaller tasks, where each task is performing the same type of iteration over different intervals. An example of this is in the computation of the gravitational acceleration and potential at each facet of the shape model. The total number of facets N in the shape model is divided by the number of maximum simultaneous processes M, or jobs, that can be supported by the grid engine, rounded down. This is the number of facets F on which an individual instance would perform computations. The grid engine manager generates separate instances of the gravitational algorithm and sends them to various nodes for processing by the grid engine. At each instance there is a range of facets over which to calculate the gravitational values, which are then output to an ascii file. The first instance iterates from facet 1 to facet F, the second instance iterates from facet F+1 to facet 2F, and so on. If necessary the last instance operates over the modulus of N/M. Once all jobs have been completed, the pipeline concatenates the results and uses it to append gravitational planes to the DTM or create ancillary fits files.

The second category is known as an array job. It consists of a series of workload segments that can all be run in parallel but are completely independent of each other. Each workload segment, or task, is identical except for the data sets on which they operate. An example of this is the creation of local DTMs tiled to provide global coverage of the Bennu surface. Each DTM contains the same number of image planes and has the same pixel size but is created at different latitude/longitude centers. The grid engine is given an array of jobs where every job contains a call to the same software tool that creates a DTM but with different latitude/longitude centers.

Discussion: Usage of the Sun Open Grid engine has substantially decreased the processing time required to generate the full set of AltWG products required by the Project. Table 1 shows a comparison of the processing time to calculate the gravitational acceleration and potential at each facet for global shape models at various resolutions. using a grid engine with 100 nodes vs. a single server with 16 cores. The grid engine has a longer processing time compared to a single server for the lowest-resolution shape model. This is due to the time delays between the submission of a job, recognition of the job submission and placement into the grid queue, assignment of the job to a node with appropriate resources, and job completion and release back to the main pipeline process. This illustrates that the grid engine should not be used in situations where a single server would suffice.

Table 1 Gravitational algorithm processing time, 16 core single server vs. 100 node grid engine

Number of facets	one server processing time (sec)	grid engine processing time (sec)
12,288	19	60
49,152	271	120
196,608	4396	360
786,432	71,493	1620
3,145,728	629,145	18,120

Future Work: Other stages of the pipeline may also benefit from being reworked to utilize the grid engine. Some of these stages are recent additions to the pipeline and may have to be refactored to support both operations in a grid engine or single server environment. In addition, generation of local DTM products from OLA shape model data will also need to be optimized when higher-resolution OLA-based shape models are produced in 2019.

Finally, work is being done to generalize pipeline operations to work with other grid engine networks to further parallelize operations.

References:

[1] Lauretta D. S. et al. (2019) *Nature, 568,* 55-60. [2] Lauretta D. S. et al. (2017) *Space Sci. Rev.*, 212, 925-984. [3] Jackman C. D. et al. (2017) *Spaceflight Mechanics*, 160, 3337-3354. [4] Weirich J. R. et al. (2017) *LPS XXVII*, Abstract #1700. [5] Rizk B. et al. (2018) *Space Sci. Rev.*, 214: 26. [6] Daly M. G. et al. (2017) *Space Sci. Rev*., 212, 899-924. [7] Barnouin O. S. et al. (2019) *Nature Geosci.*, 12, 247-252