Unified representation for collaborative visualization of planetary terrain data

DigitalFish: *Daniel Herman, Dr. Nathan Litke, Douglas Epps* **CMU:** *Dr. Lorenzo Flückiger, Dr. David Lees*

Current and future NASA planetary missions are generating ever-increasing volumes of terrain data from orbital and surface-based assets at vastly different resolutions.[1] Conventionally, Digital Elevation Maps (DEMs) or triangle meshes have been used to encode planetary terrain, but these methods suffer from numerous limitations. We have applied an alternative technology, subdivision surfaces, coupled with a volumetric reconstruction process, to help manage and present high-fidelity mesh representations of the disparate range of terrain data collected by rovers and satellites.

This work was supported by NASA STTR 2014 Phase I contract NNX14CA59P. The authors thank our NASA and CMU program managers and partners.

http://www.digitalfish.com/mars-terrain

Digital terrain models of Mars

Current state-of-the-art tools for terrain visualization, such as Google Earth, only manage terrain represented as DEMs (not full 3-D), and integrate high-resolution imagery, but not high-resolution or 3-D terrain geometry. In addition, multiple sources of information are presented as image layers, not as a unified surface.

Subdivision surfaces (*subdivs*) are a compact, multiresolution representation for 3-D surface geometry in computer graphics. An advantage of subdivs over other 3-D formats such as triangle meshes and NURBS is that subdivs support arbitrary topology: the subdiv base mesh may freely intermix faces and vertices of different valence. This topological flexibility is an advantage in the presence of large local variations in mesh density.

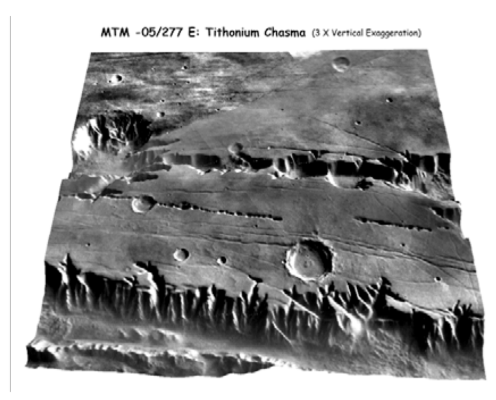

An example DEM of Martian terrain, showing typical features challenging for height-field representation, including sheer vertical walls, overhangs, and high variation in spatial frequency that leads to oversampling in conventional DEMs. [USGS image]

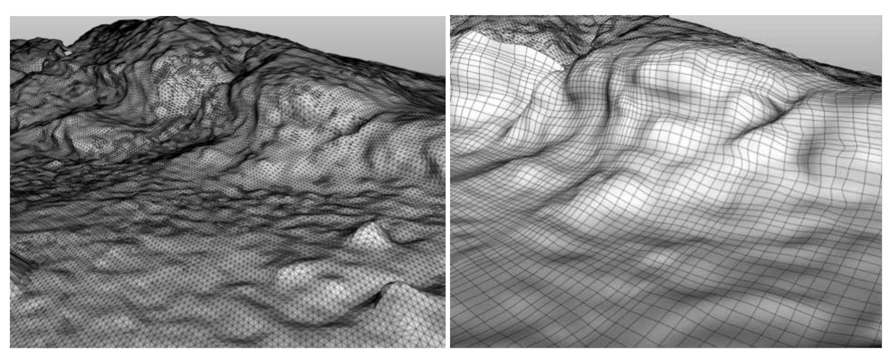

Conventional triangle mesh (left) and approximating subdivision surface with creases (right). Note the sparse use of control vertices to capture gross shape in the subdiv, with dense vertex spacing concentrated in areas of high spatial frequency. [Mesh data courtesy of NASA Ames/CMU]

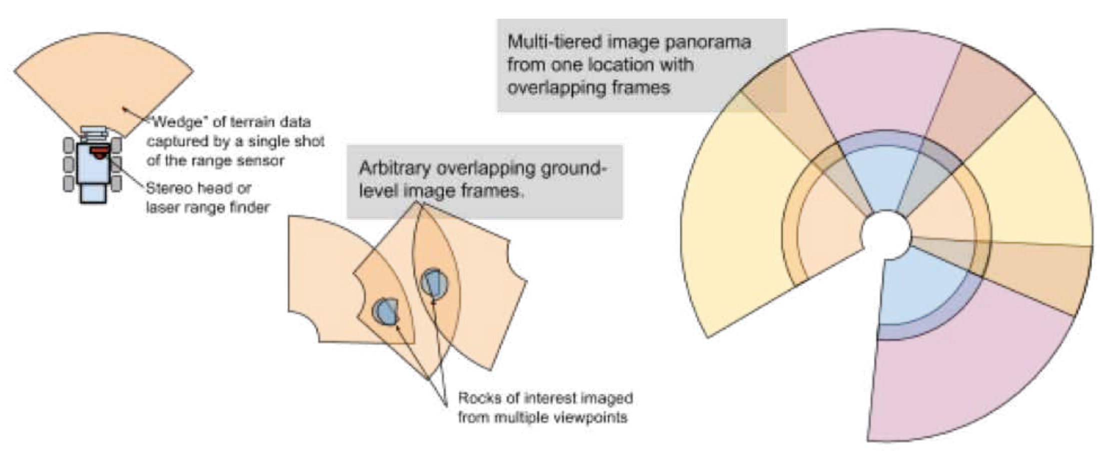

Three scenarios of terrain assembly: a rover captures multiple instances of arbitrary overlapping wedges and multi-tiered panoramas to be presented within a large orbital DEM.

Subdivision surfaces on the GPU

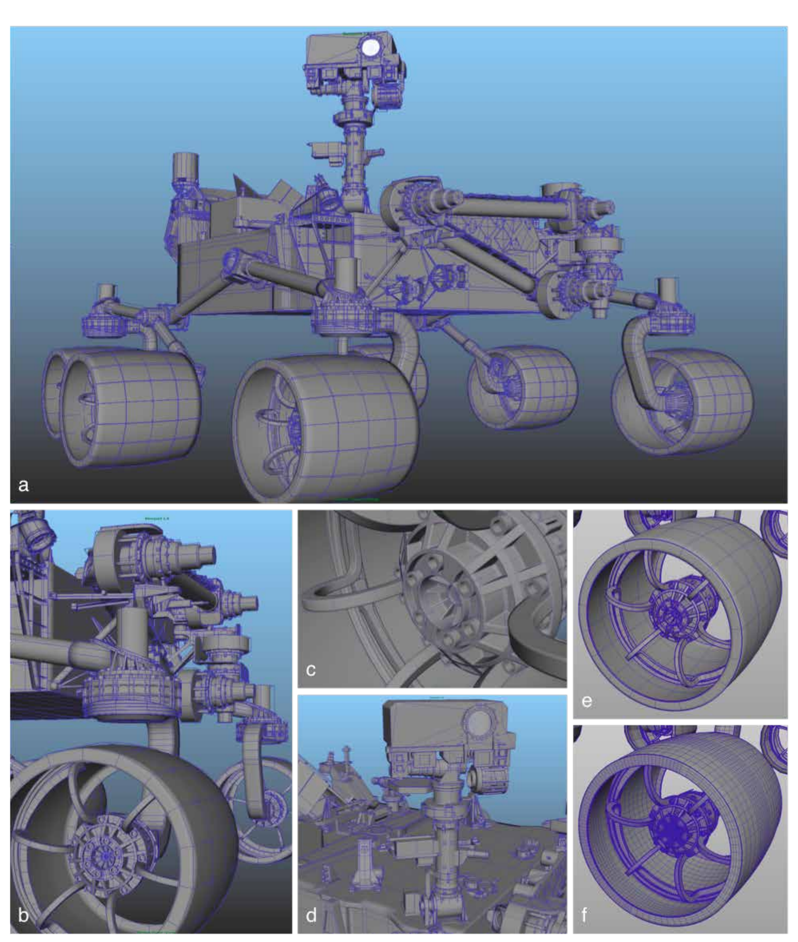

Subdivision surfaces are supported in a recently released open-source library, Pixar's OpenSubdiv (OSD)[2]. A key advancement within OSD is the reformulation of the core subdivision algorithms for fast evaluation on large-scale Single Instruction, Multiple Data (SIMD) architectures, including on Graphics Processing Units (GPUs) found in desktop and mobile computing devices.

LEFT: *Our demonstration tool displays the Mars Science Laboratory (MSL) Curiosity rover for context. We built this model of Curiosity as subdiv geometry using COTS software in order to demonstrate the utility of the subdiv representation for man-made forms. These screenshots of the model in the COTS authoring tool illustrate the tremendous economy from using subdivs (a-e), compared with the high vertex count needed to represent this geometry with similar fidelity using a polygon mesh (f).*

Terrain reconstruction from volumes

A central contribution of this work is a procedure for converting terrain data from disparate sources into a volumetric representation from which a unified mesh is reconstructed. This method provides fine-grain control over the accuracy of the reconstructed surface and supports any number of sources with overlapping terrain features. We use DreamWorks' OpenVDB for volume processing currently. We plan to extend the processing to run on scaled GPU clusters such as NVIDIA Grid.

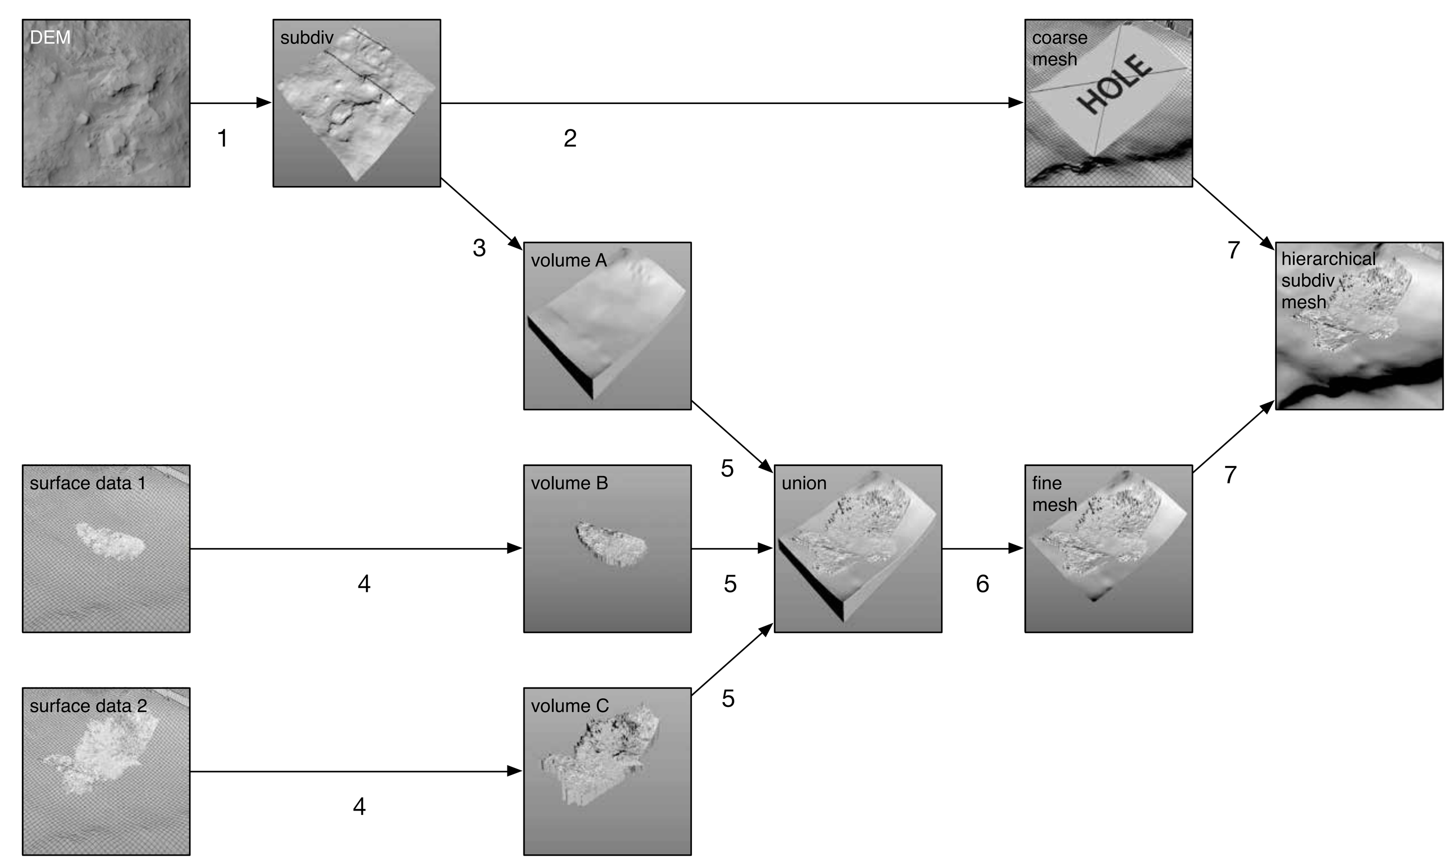

1. The DEM is converted to a subdiv that smoothly interpolates the DEM sample points.
2. The coarse mesh is prepared by assigning "hole" tags to subdiv faces in the region where multiple input datasets overlap.
3. A volumetric representation of the targeted region of orbital mesh is defined by a signed distance function (SDF).
4. A SDF is constructed from the surface data.
5. The SDF for the fine mesh is the union of the SDFs for the orbital terrain and the two surface-detail wedges.
6. The fine mesh is reconstructed from the SDF via the Dual Contouring method[3] and the sides and bottom that were extruded to create a volume are discarded.
7. Two rings of quadrilaterals are sampled from the coarse mesh and stitched to the fine mesh to extend its boundary.

Unified mesh representation

We augment Pixar's subdivision surfaces by allowing nested control meshes within the subdivision hierarchy. The resulting hierarchical subdiv mesh is a graph-theoretic tree that defines a subdivision surface at each node. Each child node is a finer resolution than its parent and occupies a "hole" which it smoothly interpolates along their shared boundary to guarantee a crack-free surface. Because each hole is simply represented by metadata in the parent's subdivision hierarchy, no mesh surgery occurs that could disturb the terrain data and introduce error.

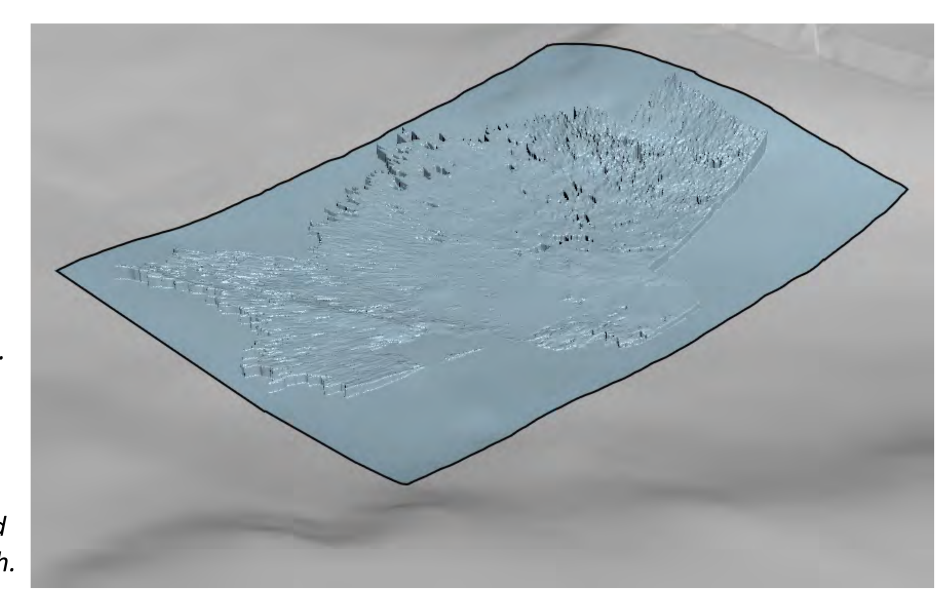

RIGHT: *Our unified mesh representation is defined by a hierarchy of nested subdiv meshes. Here the fine mesh (highlighted) occupies a "hole" embedded within the subdivision hierarchy of the coarser mesh. For visualization the fine mesh can be toggled on and off by alternately showing or hiding the fine mesh and the hole-tagged subdiv faces in the coarse mesh.*

REFERENCES

1. Nelson, S.V., Arvidson, R.E., Slavney, S., & Springer, R.J., Mars Exploration Program: Expected Data Volumes and Data Access Requirements for Research and Public Engagement. April 2001.
2. Nießner, M., Loop, C., Meyer, M., & Derose, T. "Feature-adaptive GPU rendering of Catmull-Clark subdivision surfaces." *ACM Trans. Graph.* 31, 1 (February 2012): 6:1-6:11.
3. Ju, T.J., Losasso, F., Schaefer, S., & Warren, J., "Dual Contouring of Hermite Data." *ACM Trans. Graph.* 21, 3 (July 2002), 339-346.

Reconstruction from multiple sources

A user wants to see registered views of a feature from multiple viewpoints, including:

- Overlay texture data after object merge for science analysis
- Share a manually registered merge quickly with remote users

After downlinking image data from Mars, stereo pairs are processed to produce range maps from the navigation camera (NAVCAM) images. Our automated pipeline converts these range images to 3-D models of the local terrain around the rover and stores them for further processing. One or two members of the science team manually co-register images of an interesting rocky outcrop taken from two different rover positions, merge them into a 3-D subdivision surface and store them for future use.

Later, a remote member of the science team downloads the processed subdiv surface and merges it with her global terrain model. She then applies high-resolution color Mastcam images as a texture on the 3-D terrain of the outcrop so she can plan a new imaging sequence. Using the merged data, she plans a MAHLI (microscopic imager) sequence to get detailed images of small features nested in the surface of the rock outcrop.

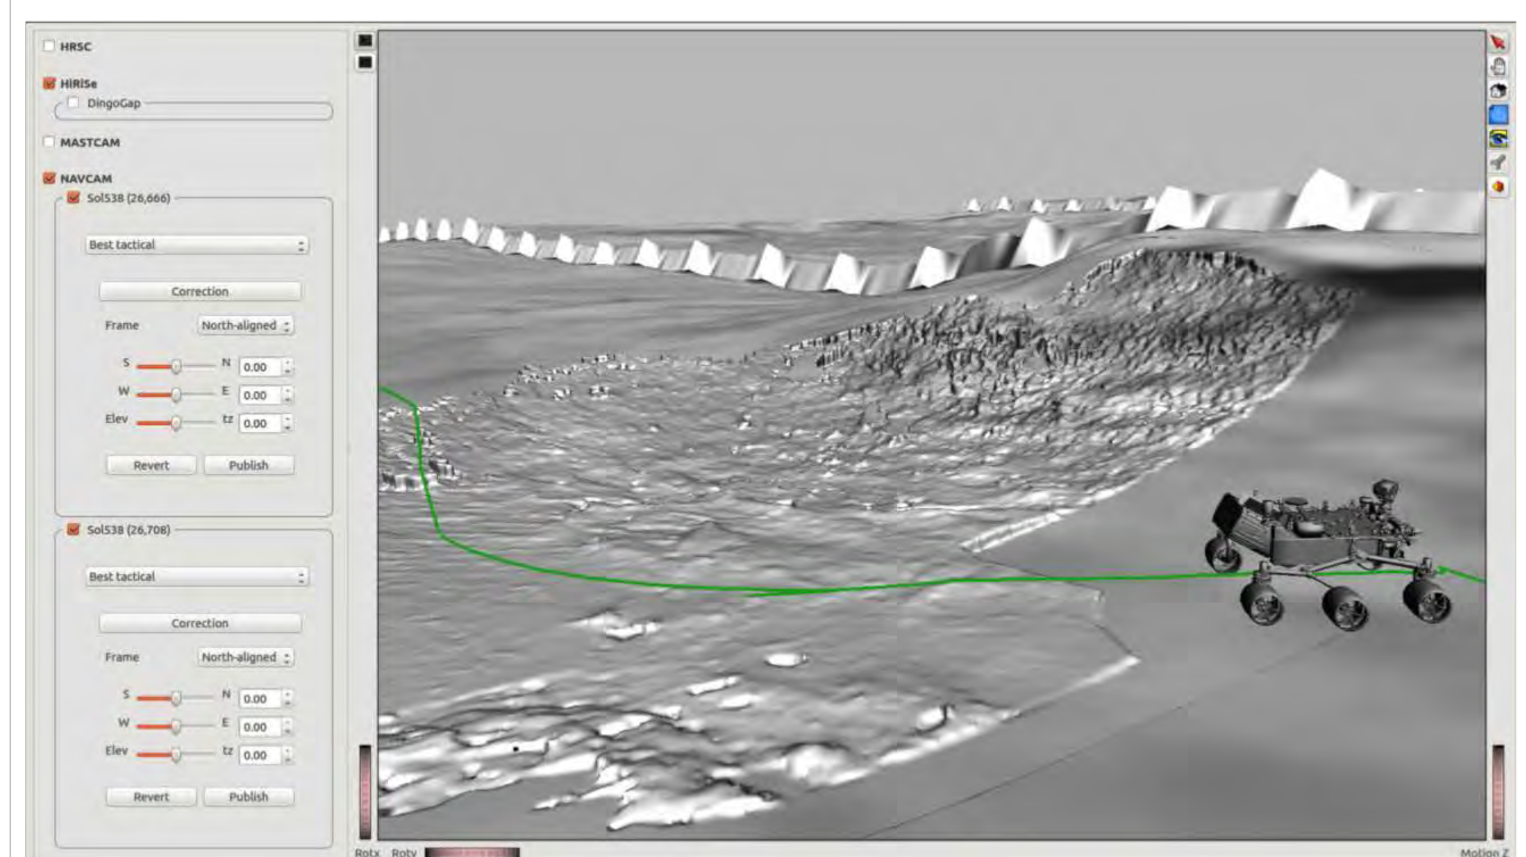

A screenshot of our demonstration application, showing merged terrain from MSL Sols 538 and 539 in the vicinity of Dingo Gap. The sharp ridge in the background is an artifact in the input data where two orbital DEMs with inconsistent registration were combined.

FEATURE	BENEFITS
Features/benefits of subdivision surfaces and OpenSubdiv	
Full 3-D mesh representation	Retain full detail extracted from captured range data or stereo pairs Represent 3-D features not expressible in a height map
Arbitrary mesh topology	Freedom of vertex placement and mesh-density distribution, allowing a sparse mesh with fine detail only where needed
Features/benefits of unified mesh representation	
Multiresolution representation	Natural mechanism for scalable level of detail (LOD) Dynamic view-dependent LOD without offline pre-processing Progressive-refinement based LOD preserves high-res data
Arbitrary nesting of subdivision surfaces	Represent geometry with large dynamic range of feature scales, from planetary to microscopic
Features/benefits of volumetric processing and OpenVDB	
Volumetric merge abstracts away the details of input-terrain datasets	Integrate input data of varied formats (DEM, mesh, point cloud, range data, *etc.*), on inconsistent coordinate grids and of widely varied spatial resolutions Supports source data in original format without a need for additional preprocessing
Volumetric remeshing creates subdiv control meshes	Represent 3-D features of complex geometry not expressible with high fidelity in a substantially planar mesh that would be directly derived from a DEM, including annuli or deep in-pockets such as are found in lava tubes and other "cave-like" structures
Highly parallelized volume operations	Well suited to implementation on GPU arrays, *e.g.* NVIDIA GRID