UNCONTROLLED GLOBAL HiRISE MOSAIC. L. Plesea[1], T. M. Hare[2], [1]Esri, Redlands, California, LPlesea@esri.com, [2]U. S. Geological Survey, Astrogeology Science Center, thare@usgs.gov.

Introduction: The High Resolution Imaging Science Experiment (HiRISE) instrument [1] has been collecting images from the Mars Reconnaissance Orbiter (MRO) spacecraft since it began orbiting Mars in 2006. Spatially, HiRISE is the highest resolution instrument sent to Mars and continues to be a critical resource for science and helping to select safe landing sites for robotic exploration. Although, due to its' extreme spatial resolution, HiRISE has only covered about 3% of the martian surface.

Herein, we describe a work-in-progress, really a technology evaluation, to release all available HiRISE images at full resolution (0.25 m/p) as three composite mosaics, one equatorial and two polar, hosted on the cloud as a tiled service. This test has pushed our current limits for available processing clusters and data storage capabilities. We also used novel raster formats and compression techniques, and optimized distributed cloud capabilities.

HiRISE Data Volumes and Access: A typical HiRISE observation is ~20,000 pixels wide by ~120,000 pixel in length at a nominal resolution between 0.25-1 m/pixel (depending on the pixel binning mode) [2]. This makes a JPEG2000 compressed image file typically larger than 1 gigabyte in size. With more than 58,000 images collected, the compressed data volume is more than 50 terabytes on disk.

While the HiRISE Team has done an exceptional job making HiRISE images available, data volumes this large can be extremely difficult for users to efficiently access. Given this volume of data, the team has also excelled at rapidly releasing the data and providing science-ready (map projected) images, making them immediately useable in geospatial mapping applications (e.g. GISs). Essentially, these proposed mosaics would not have been possible if these products were not made available, due to the complicated processing pipeline, even with the ISIS3 [3] software freely available.

Spatial Limitations: Unfortunately, due to uncertainty in spacecraft location and pointing, the default placement on the martian surface for most HiRISE map projected images can be hundreds of meters off. Thus, these individual images, and hence our derived mosaics, cannot be considered foundational products as defined by a Planetary Spatial Data Infrastructure (PSDI). Recall that within a PSDI, as defined by Laura [4], foundational geospatial data products include: (1) geodetic control networks, (2) topography, and (3) **rigorously** controlled and orthorectified images tied to a standardized reference frame. Here, we do not yet attempt to fix any location inaccuracies, for example, using automatic image matching methods as proposed by [5, 6]. But, we feel the data access benefits for releasing as an efficient GIS web service, even with this spatial uncertainty, outweighs the disadvantages.

Meta Raster Format Introduction: The key to our ability to make a HiRISE mosaic is the Meta Raster Format (MRF, 7). In short, MRF is an image and tile storage format designed for fast access to imagery within a georeferenced tile pyramid at discrete resolutions. The format supports extremely large, tiled, multi-resolution and multi-spectral data. Sparse raster data, like the HiRISE mosaics, can be handled very efficiently. Lossy JPEG, lossless PNG and Limited Error Raster Compression (LERC, 8) are some of the supported tile compression formats. Grayscale, color, indexed (palette) color models are also supported. Depending on the tile compression used, all common bit types are supported (e.g. 8-bit to 64-bit floating point). Both the MRF format and LERC compression methods are open source and implemented within the Geospatial Data Abstraction Library (GDAL, 9). Lastly, the MRF format was originally developed at the NASA Jet Propulsion Laboratory and is currently used as the technology foundation for the EOSDIS GIBS project [10] and in multiple projects within Esri.

Mosaic Processing: As input, we used the HiRISE single-band, red-band-only (grayscale), 10-bit, map projected GeoJpeg2000 images as processed by the HiRISE Team and released by the Planetary Data System (PDS, 11). We used the 100+ node cluster available at the Astrogeology Science Center (ASC) and allocated 250 terabytes of temporary high-speed BeeGFS parallel storage for the processing steps.

First, to enable faster read access to the individual images, we reformatted the data from the native GeoJpeg2000 format to a tiled MRF using a lossless LERC compression. Input filenames were then sorted by resolution and separated into equatorial, north pole, and south pole files such that the highest resolution images were listed last. During the mosaic step, all images will be used to create the output mosaic, but when there is overlap, only the last image will be retained. Due to limited overlapping images, no other filtering was applied (e.g. phase or emission angle).

To build the equatorial mosaic, a specialized C++ routine, called *mbuilder*, was created using the GDAL library to initialize an empty 0.25 m/pixel resolution global mosaic as an MRF image with 83,886,080 by 41,943,040 pixels in total and divided into 512 by 512-pixel tiles. It is worth noting that as a 10-bit per pixel uncompressed file, this mosaic would use 4.4 petabytes of storage. Next, for each output file, the source HiRISE images that contributed are selected and reprojected on-the-fly to the output Simple Cylindrical map projection and combined to produce the output tile. Only the output tiles which contain data are created on disk. The MRF index file maintains each output tile's spatial location within the map projection.

This method of only writing the needed tiles to disk is the core and strength of the MRF format. It excels at sparse data sets like HiRISE, because the empty tiles are never created. The MRF format optionally supports maintaining all image overlaps into an output mosaic but it may greatly increase the overall size and complexity of the index file.

Because this is a visual image product with 10-bit fidelity, we chose a high-quality, but lossy, JPEG-Zen compression method for the final output mosaics. The JPEG encoder in GDAL can support up to 12-bit data, thus the range of the input HiRISE data was unchanged. Within GDAL, the MRF JPEG driver also supports the Zero Enhanced (Zen) JPEG extension. This extension, defines a 1-bit lossless zero-value mask. This allows for clean image edges not normally available with JPEG-compressed images. The Zen extension is applied by the MRF driver and is available in any client application that uses GDAL with MRF. When served as tiles, clients which do not support the Zen extension can read it as a normal JPEG, but with the usual JPEG edge artifacts.

To parallelize the processing of the mosaic across the cluster, we split the equatorial mosaic into 640 regions. Depending on the density of the HiRISE images per region, processing time was from 1 hour to more than a day. Total processing time on the ASC cluster for all quads approached one week. We then merged the 640 quads to a more manageable 8 hemisphere quarters of 90x90 degrees (quads) using the Python script *mrf_join.py*, which operates directly on the MRF file structure. These initial 8 quads currently range in size from 1.3 to 1.8 terabytes. Our goal was to keep the final quad size below 2 terabytes to make the file handling somewhat easier and to stay under the 5-terabyte single file size limit for our cloud provider.

The last step before staging, was to run *gdaladdo* to build lower-resolution versions (or pyramids) for each quad. The GDAL command *gdaladdo* offers a several down-sampling methods from nearest neighbor to lanczos. Fortunately, as part of the MRF driver implemented within GDAL, there is a specialized averaging method to down-sample, called *avg*, which efficiently ignores the empty tiles, thus greatly speeding up the processing of sparse mosaics.

The 8 quads will be loaded on a cloud object store and a GIS map tile server will seamlessly provide access to all the tiles as a single service. The tile server will be configured to convert the JPEG-Zen tiles to PNG with transparency, in both 10 and 8-bit formats to help support legacy applications.

A similar workflow is used to produce separate North and South Polar Stereographic map projection composite mosaics, from the HiRISE scenes, which were originally released in a polar map projection.

Conclusion: Once processed and served from a cloud provider, users will have immediate access to nearly all HiRISE images within a GIS mapping application or web mapping viewer (fig. 1). As this is an initial test to prove this process, there are still many tasks which will require further research, including the ability to better spatially register these images together and to a foundational base map. Long-term availability for this service is uncertain, but it will soon be publicly available for testing (see: http://bit.ly/HiRISE_mosaic).

References: [1] McEwen, A.S., et al., 2007. Mars Reconnaissance Orbiter's High Resolution Imaging Science Experiment (HiRISE). J. Geophys. Res., 112, E05S02, doi:10.1029/2005JE002605. [2] Sutton, S., et al., 2017. Correcting Spacecraft Jitter in HiRISE Images, ISPRS, https://doi.org/10.5194/isprs-archives-XLII-3-W1-49-2017. [3] Edmundson K.L., et al., 2012, Jigsaw: the isis3 bundle adjustment for extraterrestrial photogrammetry, ISPRS, doi: 10.5194/isprsannals-I-4-203-2012. [4] Laura, J., et al., 2017, Towards a Planetary Spatial Data Infrastructure, IJGI doi:10.3390/ijgi6060181. [5] Laura, J., et al., 2018, AutoCNet: A Python library for sparse multi-image correspondence identification for planetary data, SoftwareX, doi: 10.1016/j.softx.2018.02.001. [6] Sidiropoulos P. and Muller J.P., 2015, Matching of large images through coupled decomposition, IEEE, doi: 10.1109/TIP.2015.2409978. [7] http://bit.ly/MRF_User_Guide [8] http://github.com/Esri/lerc. [9] GDAL - Geospatial Data Abstraction Library, 2018, Version 2.3.0, Open Source Geospatial Foundation, http://www.gdal.org. [10] PDS, https://pds.nasa.gov.

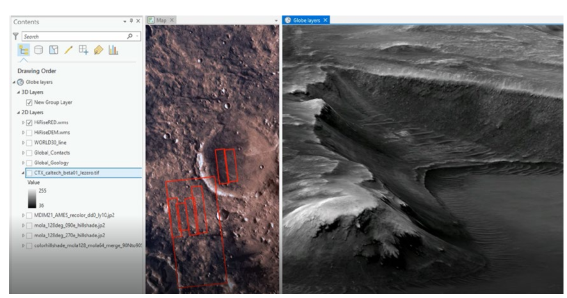

Figure 1. Within Esri's ArcGIS Pro, the left panel shows the colorized Mars Digital Image Mosaic (v2) over Jezero crater with one CTX and several HiRISE image footprints (red). The panel on the right, also within Jezero crater, shows an initial HiRISE MRF mosaic in 3D as drawn from a cloud service provider. Heights are from the released HiRISE digital terrain models also combined and served as a web-enabled MRF mosaic.