THE GLOBAL CTX MOSAIC OF MARS: LESSONS FOR THE CONSTRUCTION AND DISSEMINATION OF MASSIVE IMAGING DATA SETS. J. L. Dickson[1], B. L. Ehlmann[1,2], L. Kerber[2], C. I. Fassett[3], T. M. Hare[4], D. P. Quinn[5], L. Plesea[6], D. Noss[7], [1]California Institute of Technology, Division of Geological and Planetary Sciences, Pasadena, CA, 91125 (jdickson@caltech.edu), [2]Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr, Pasadena, CA 91109. [3]NASA Marshall Space Flight Center, Huntsville, AL 35805. [4]U.S. Geological Survey, Astrogeology Science Center, Flagstaff, AZ, USA. [5]University of Wisconsin, Madison, Department of Geoscience, Madison, WI, USA. [6]ESRI, Redlands, CA. [7]Mars Space Flight Facility, 201 E Orange Mall, Arizona State University, Tempe, AZ 85287 USA.

Introduction: The global CTX mosaic of Mars [1], given its unprecedented scale (> 8 trillion pixels), has presented challenges for both its construction and dissemination. We generated and released a proof-of-concept version of the mosaic ("beta01") in order to (1) determine what unforeseen challenges were involved with constructing a mosaic of this size, (2) engage the scientific community as to how best to construct the final mosaic, and (3) develop techniques for efficient access to the mosaic for the scientific community. While this preliminary version was meant primarily to be a technology demonstration, the mosaic has proven useful to an array of scientific projects and has entered the planetary science literature [2-8].

Here, we present a summary of our findings through this process and discuss our plans for production of an improved global CTX mosaic of Mars that addresses the immediate needs of the Mars science community.

Background and Motivation: CTX [9] has imaged >97% of the surface of Mars with high S/N at resolutions that typically range from 5.0-6.0 m/px, providing the highest resolution remote sensing data of >90% of the martian surface. As such, the CTX data set has extensive value as a *primary* data source, not simply as a *supporting* product that gives context for more targeted instruments like HiRISE and CRISM, in spite of its name. Thus, preserving metadata from each orbit pertaining to illumination conditions and spacecraft pointing properties is critical to prevent proper geomorphologic interpretation of surface features.

Mosaicking of imagery at mid- and high-latitudes on Mars is complicated by (1) poor illumination conditions and (2) seasonal processes (most commonly frost deposition) that dramatically alter the reflective properties of the surface as a function of solar longitude (L_s). This, in effect, makes the blending of overlapping CTX orbits at high-latitudes frequently impossible to be done in a seamless manner. At present, coverage of CTX in the mid-latitudes is insufficient to produce complementary, L_s-filtered mosaics of the surface at multiple seasons, thus some abrupt seams between images is unavoidable.

These two concerns – accessing original data from the derived mosaic and pixel-for-pixel documentation of image seams – have been solved by using non-destructive image processing [1,11], which is an alternative to traditional iterative processing common to geospatial data processing architectures (USGS ISIS, GDAL, etc.). Non-destructive processing is built to be information-preserving and has proven successful at (1) blending orbits in a way that minimizes visible seams, (2) tracing all seams between images, and (3) producing vectorized polygons that retain attributes from the original data, such that all imaging conditions for each pixel in the mosaic are available to the user.

Further, non-destructive processing greatly improves the processing efficiency of the pipeline. Thus, the CTX mosaic can be made with less hardware and more efficient software, providing a potential model for future geospatial data processing architectures.

Dissemination: The "beta01" CTX mosaic has prompted innovation in how to provide efficient access to an imaging data at a colossal scale. Of importance is that users may prefer and/or need to engage with the product in a variety of manners depending on (1) their access to hardware/software resources and (2) their specific scientific needs. Thus, we have developed multiple ways by which users can access the mosaic.

1. Individual tiles. The mosaic is stored as 3,960 individual 4°x4° tiles, with a GeoTiff for the raster data and a shapefile for the polygon seam maps. While the mosaic can be downloaded in full, more common use cases that we have seen involve accessing a batch of tiles over a specific region of interest. For this case, we have developed a new "Mars Data Browser" (murray-lab.caltech.edu/Mars-Data-Browser/) that allows for the dynamic selection of tiles and the production of a link list for data download.

Engaging with a large volume of tiles, or the entire mosaic as a whole, is best conducted, in our experience, via the "Mosaic Dataset" feature within ArcGIS. Within this framework, large collections of individual tiles function as if working with one raster.

2. Streaming. The development of the CTX mosaic coincides with the advancement in the ability to remotely stream large volumes of tiled geospatial data over the web [12]. The CTX mosaic has been used by multiple platforms to stretch the capacities of these services. We have successfully produced streaming versions of the "beta01" release through (1) ESRI's

Map Viewer web interface and (2) the ASU JMARS standalone platform (jmars.asu.edu) [13]. The ESRI product can also be accessed as an online data resource through ArcMap. For QGIS, a pointer file stored by the Murray Lab (murray-lab.caltech.edu/CTX/) can be loaded to stream the mosaic, as well.

Our findings are that the streaming products and the individual tiles are highly complementary and are frequently used in conjunction as a function of the scale and the nature of the scientific problem being addressed.

Image Registration: The "beta01" version of the CTX mosaic used SPICE pointing information for estimated image positioning, but did not further correct for pointing uncertainties. Feedback from the science community has focused mostly on image-to-image offsets, which are rarely recognizable to the eye due to our non-destructive blending technique, but do exist and hinder measurements made across multiple orbits.

We have begun the construction of the successor to "beta01", which will be semi-controlled (CTX images are registered to each other, but not to ground). As of this writing, we have automatically co-registered and manually validated 51,074 images with each other between 68°S and 68°N (Fig. 1). Mean offset between two images before co-registration is 150.6 m. 94.1% of images are natively offset by less than 300 m (the distance between two sequential MOLA points) and 98.6 % are offset by less than 463 m (the size of a non-polar MOLA gridded pixel). These offsets are strictly image-to-image and *not* relative to ground.

The relatively small scale of offsets of non-polar CTX images relative to the scale of the global MOLA DEM suggests that a semi-controlled version of the mosaic should satisfy the vast majority of scientific applications for the mosaic, based upon feedback from the scientific community. Our efforts are in conjunction and collaboration with a parallel effort [14] to geodetically control all viable CTX imagery to provide higher-precision engineering information as a product of the mosaicking process.

Acknowledgements: Thanks to a NASA PDART grant B.L.E./J.L.D. (#80NSSC19K0422) for funding the semi-controlled CTX mosaic effort.

References: [1] Dickson J.L. et al. (2018) *LPSC 49,* 2480. [2] Pan L. and Ehlmann B.L. (2018) *JGR-P*, 10.1029/2017JE005461. [3] Kremer C.H. et al. (2019) *Geology*, 10.1130/G45563.1. [4] Horgan B.N.H. et al. (2019) *Icarus*, 10.1016/j.icarus.2019.113526. [5] Day M. (2019) *JGR-P*, 10.1029/2019JE006226. [6] Shahrzad S. (2019) *GRL*, 10.1029/2018GL081402. [7] Mandon L. et al. (2020) *Icarus*, 10.1016/j.icarus.2019.113436. [8] Pan L. et al. (2020) *Icarus*, 10.1016/j.icarus.2019.113511. [9] Malin M.C. et al. (2007) *JGR-P*, 10.1029/2006JE002808. [10] Edwards C.S. et al. (2011) *JGR-P*, 10.1029/2010JE003755. [11] Dickson J.L and Ehlmann B.L. (2019) *4th Planet. Data Workshop*, 7109. [12] Plesea L. and Hare T.M. (2019) *LPSC 50*, 1986. [13] Dickenshied S. et al. (2019) *4th Planet. Data Workshop*, 7108. [14] Robbins S.J. et al. (2019) *LPSC 50*, 1678.

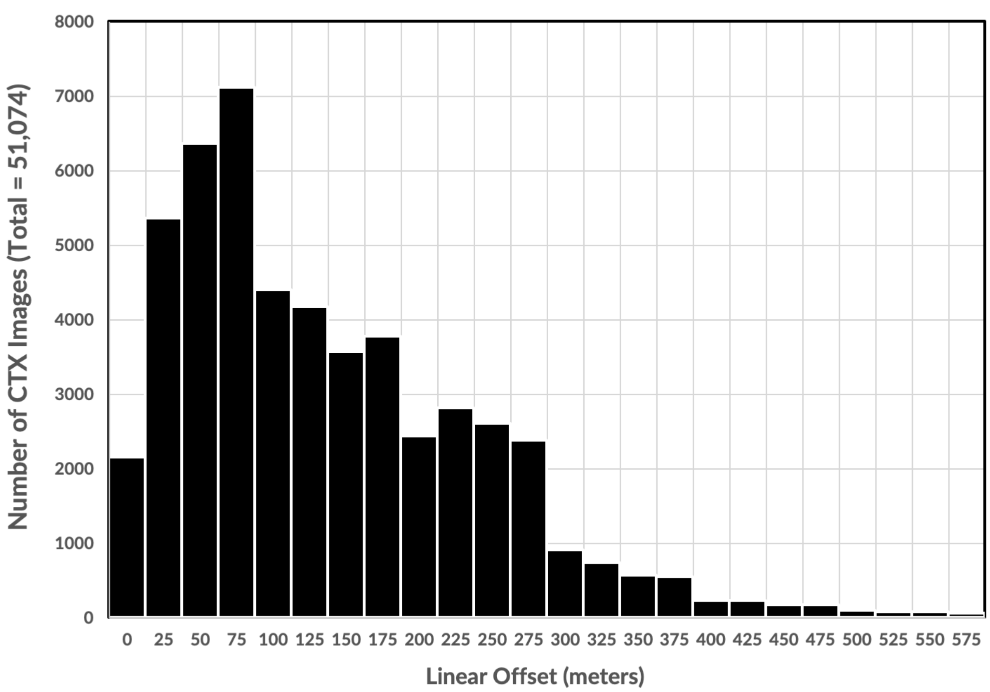

Figure 1. Distance of linear shift required to achieve image-to-image registration of multiple CTX images.