

less of the nature or source of the input. This allows new scheduling inputs to be processed without the need to configure the external interface to a specific input format.

flexplan is also used in the scheduling operations of the LDCM (Landsat Data

Continuity Mission) spacecraft, also at GSFC, which is currently undergoing final mission and ground readiness testing to prepare for upcoming launch in January of 2013, and undergoing customization for operational use in the TDRS (Tracking and Data Relay Satel-

ites) Space Network upgrade project, SGSS, a joint venture between NASA GSFC and WSC.

This work was done by Assaf Barnoy and Theresa Beech of GMV USA for Goddard Space Flight Center. Further information is contained in a TSP (see page 1). GSC-15558-1

➤ Estimating Torque Imparted on Spacecraft Using Telemetry

Methodology is straightforward and does not involve the use of any complex supporting ground software.

NASA's Jet Propulsion Laboratory, Pasadena, California

There have been a number of missions with spacecraft flying by planetary moons with atmospheres; there will be future missions with similar flybys. When a spacecraft such as Cassini flies by a moon with an atmosphere, the spacecraft will experience an atmospheric torque. This torque could be used to determine the density of the atmosphere. This is because the relation between the atmospheric torque vector and the atmosphere density could be established analytically using the mass properties of the spacecraft, known drag coefficient of objects in free-molecular flow, and the spacecraft velocity relative to the moon. The density estimated in this way could be used to check results measured by science instruments. Since the proposed methodology could estimate disturbance torque as small as 0.02 N-m, it could also be used to estimate disturbance torque imparted on the spacecraft during high-altitude flybys.

When the expected value of torque imparted on the spacecraft is low and

within the control authority of the reaction wheel assemblies (RWAs), mission design engineers will use these RWAs to control the spacecraft attitude. Relative to thrusters, RWA can produce better pointing control and stability performance. To estimate the disturbance torque imparted on the Cassini spacecraft, the proposed methodology exploits the unique and known relation between the disturbance torque and the RWA-based attitude control error during an Enceladus or Titan flyby.

To estimate the disturbance torque imparted on the Cassini spacecraft, the unique and known relation between the disturbance torque and the attitude and attitude rate control errors during an Enceladus flyby (or a Titan flyby) on reaction wheels was used. The effectiveness of this methodology is illustrated using telemetry data obtained from the 50-km Enceladus-3 flyby. Results determined using this approach were compared with those determined using the "time rate of change of spacecraft angular momentum" approach. Results of

this flyby determined that using this new approach compared very well with that estimated using the angular momentum approach. In effect, density estimates made using these two independent engineering methodologies could cross check each other. Moreover, density estimates determined using these methods could also be used to cross check science-based density estimates.

This method could be used to estimate very small torque imparted on the spacecraft. The methodology is straightforward and does not involve the use of any complex supporting ground software. This methodology uses telemetry data that are available at high telemetry frequency, and the telemetry data involved (per-axis attitude control errors and per-axis attitude rate control errors) are floating-point data with high accuracy.

This work was done by Allan Y. Lee, Eric K. Wang, and Glenn A. Macala of Caltech for NASA's Jet Propulsion Laboratory. For more information, contact iaoffice@jpl.nasa.gov. NPO-47545

➤ PowderSim: Lagrangian Discrete and Mesh-Free Continuum Simulation Code for Cohesive Soils

John H. Glenn Research Center, Cleveland, Ohio

PowderSim is a calculation tool that combines a discrete-element method (DEM) module, including calibrated interparticle-interaction relationships, with a mesh-free, continuum, SPH (smoothed-particle hydrodynamics) based module that utilizes enhanced, calibrated, constitutive models capable of mimicking both large deformations and the flow behavior of regolith simulants and lunar regolith under condi-

tions anticipated during in situ resource utilization (ISRU) operations.

The major innovation introduced in PowderSim is to use a mesh-free method (SPH-based) with a calibrated and slightly modified critical-state soil mechanics constitutive model to extend the ability of the simulation tool to also address full-scale engineering systems in the continuum sense. The PowderSim software maintains the ability to address

particle-scale problems, like size segregation, in selected regions with a traditional DEM module, which has improved contact physics and electrostatic interaction models.

PowderSim provides answers with comprehensive cohesive-contact models and a new charge-spot model for electrostatic forces arising from localized charge patches on the surfaces and in the interiors of individual particles. For

systems that are too large to be simulated with a discrete element approach, PowderSim incorporates a continuum-based SPH module, which when considering the addition of a calibrated, cohesive, constitutive model (Lunar Regolith Constitutive Model (LRCM)), is a novel use of mesh-free methods. Because of the discrete and continuum methods implemented in the same framework, the software can capture dynamic particulate material behavior at a variety of

spatial scales from the coarse-grain scale (DEM) to the bulk scale (SPH). The DEM capability also supports clustering, which allows it to capture a rich variety of shape detail. Advanced contact models and charge spots capture many effects of contact plasticity and hysteresis, roughness, adhesion, and electrostatic interaction of particles. The SPH capability for bulk material behavior uses the LRCM to capture the critical-state behavior of cohesive lunar regolith.

This work was done by Scott Johnson, Otis Walton, and Randolph Settigast of Grain-flow Dynamics for Glenn Research Center. Further information is contained in a TSP (see page 1).

Inquiries concerning rights for the commercial use of this invention should be addressed to NASA Glenn Research Center, Innovative Partnerships Office, Attn: Steven Fedor, Mail Stop 4-8, 21000 Brookpark Road, Cleveland, Ohio 44135. Refer to LEW-18801-1.

Multiple-Frame Detection of Subpixel Targets in Thermal Image Sequences

This technique has applicability in fire detection, and tracking ships, ground vehicles, and aircraft.

NASA's Jet Propulsion Laboratory, Pasadena, California

The new technology in this approach combines the subpixel detection information from multiple frames of a sequence to achieve a more sensitive detection result, using only the information found in the images themselves. It is taken as a constraint that the method is automated, robust, and computationally feasible for field networks with constrained computation and data rates. This precludes simply downloading a video stream for pixel-wise co-registration on the ground. It is also important that this method not require precise knowledge of sensor position or direction, because such information is often not available. It is also assumed that the scene in question is approximately planar, which is appropriate for a high-altitude airborne or orbital view.

This approach tracks scene content to estimate camera motion and finds geometric relationships between the images. An initial stage identifies stable image features, or interest points, in

consecutive frames, and uses geometric relationships to estimate a “homography” — a transformation mapping between frames. Interest points generally correspond to regions of high information or contrast. Previous work provides a wide range of interest point detectors. In this innovation, SIFT (Scale Invariant Feature Transform) keypoints recovered by a difference of Gaussians (DoG) operator applied at multiple scales are used. A nearest-neighbor matching procedure identifies candidate matches between frames. The end result of this first step is a list of candidate interest points and descriptors in each frame.

An important benefit of SIFT detection is that the system permits absolute georeferencing based on image contents alone. The SIFT features alone provide sufficient information to geolocate a hot pixel. This suggests an initial characterization phase where the remote observer transmits high-contrast, SIFT descriptors along with images of the (fire-free) sur-

face. The ground system, with possible human assistance, would determine the SIFT features' geographic locations.

During regular operations, the system can query the database to find geographic locations of new observations. Any preferred single- or multiple-channel detection rule is applied independently in each frame with a very lenient threshold. Then, the algorithm matches consecutive detections across potentially large displacements, and associates them into tracks, i.e., unique physical events with a precise geographic location, that may appear in multiple frames. Finally, the system considers the entire sequence history of each track to make the final detection decision.

This work was done by David R. Thompson of Caltech and Robert Kremens of Rochester Institute of Technology for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP (see page 1). NPO-48129

Metric Learning to Enhance Hyperspectral Image Segmentation

NASA's Jet Propulsion Laboratory, Pasadena, California

Unsupervised hyperspectral image segmentation can reveal spatial trends that show the physical structure of the scene to an analyst. They highlight borders and reveal areas of homogeneity and change. Segmentations are independently helpful for object recognition, and assist with automated production of symbolic maps. Additionally, a good segmentation can dramatically re-

duce the number of effective spectra in an image, enabling analyses that would otherwise be computationally prohibitive. Specifically, using an over-segmentation of the image instead of individual pixels can reduce noise and potentially improve the results of statistical post-analysis.

In this innovation, a metric learning approach is presented to improve the

performance of unsupervised hyperspectral image segmentation. The prototype demonstrations attempt a superpixel segmentation in which the image is conservatively over-segmented; that is, the single surface features may be split into multiple segments, but each individual segment, or superpixel, is ensured to have homogeneous mineralogy.