all the pixels along the spectral-dispersion dimension would be summed to obtain the value of the cross-correlation (plus background).

Such on-chip cross-correlation could be performed rapidly because the analytical function could be statically programmed into the APS array and the multiplications could be done simultaneously or nearly so. All of the additions could be done simultaneously by means of a single binning instruction. The charge wells of all the pixels could be connected simultaneously, collecting all the charge outputs from multiplication operations into one "superpixel," the single readout value of which would constitute the cross-correlation value for the given analytical function. For an instrument in which the APS rows were aligned along the spectral-dispersion dimension and in which the image of a spectrograph slit was aligned along the pixel columns and spanned multiple pixel rows, it would be possible to perform simultaneous cross-correlations for multiple target species by applying, to each pixel row, the analytical function corresponding to one of the target species. A separate readout would be needed for each target species.

In the other hardware implementation, cross-correlations would be computed externally to the APS array. The multiplications and additions would be performed in pipeline fashion. If the APS-array outputs were analog, then programmable analog signals representing the analytical functions would be synthesized in phase with the corresponding stream of analog APS-array outputs and the multiplications and additions would be performed by relatively inexpensive, commercially available analog mixing and filtering circuits, respectively. If the APS-array outputs were digital, the cross-correlations could be computed by a digital signal processor. Ordinarily, the analog approach would be preferable because the analog operations can be performed much more rapidly than can the corresponding digital multiplications and additions.

Prioritizing Scientific Data for Transmission

NASA's Jet Propulsion Laboratory, Pasadena, California

A software system has been developed for prioritizing newly acquired geological data onboard a planetary rover. The system has been designed to enable efficient use of limited communication resources by transmitting the data likely to have the most scientific value. This software operates onboard a rover by analyzing collected data, identifying potential scientific targets, and then using that information to prioritize data for transmission to Earth. Currently, the system is focused on the analysis of acquired images, although the general techniques are applicable to a wide range of data modalities. Image prioritization is performed using two main steps. In the first step, the software detects features of interest from each image. In its current application, the system is focused on visual properties of rocks. Thus, rocks are located in each image and rock properties, such as shape, texture, and albedo, are extracted from the identified rocks. In the second step, the features extracted from a group of images are used to prioritize the images using three different methods: (1) identification of key target signature (finding specific rock features the scientist has identified as important), (2) novelty detection (finding rocks we haven't seen before), and (3) representative rock sampling (finding the most average sample of each rock type). These methods use techniques such as K-means unsupervised clustering and a discrimination-based kernel classifier to rank images based on their interest level.

Determining Sizes of Particles in a Flow From DPIV Data

The same equipment would be used to measure sizes as well as velocities.

John H. Glenn Research Center, Cleveland, Ohio

A proposed method of measuring the size of particles entrained in a flow of a liquid or gas would involve utilization of data from digital particle-image velocimetry (DPIV) of the flow. That is to say, with proper design and operation of a DPIV system, the DPIV data could be processed according to the proposed method to obtain particle sizes in addition to particle velocities. As an additional benefit, one could then compute the mass flux of the entrained particles from the particle sizes and velocities.

As in DPIV as practiced heretofore, a pulsed laser beam would be formed into a thin sheet to illuminate a plane of interest in a flow field and the illuminated plane would be observed by means of a