



Remote Data Access with IDL

Goddard Space Flight Center, Greenbelt, Maryland

A tool based on IDL (Interactive Data Language) and DAP (Data Access Protocol) has been developed for user-friendly remote data access. A difficulty for many NASA researchers using IDL is that often the data to analyze are located remotely and are too large to transfer for local analysis. Researchers have developed a protocol for accessing remote data, DAP, which is used for both SOHO and STEREO data sets. Server-side side

analysis via IDL routine is available through DAP.

The tools allow normal DAP users to run IDL scripts on their data remotely via DAP. This powerful, user-friendly interface to DAP for IDL improved OPeNDAP bindings that fixed bugs in existing functionality, created a GUI client to explore data sets served with DAP, developed a pure IDL DAP implementation that provided complete DAP capabilities

along with a simple installation, improved network capabilities for GDL (the open source IDL alternative) and older versions of IDL, and modified the OPeNDAP Hyrax DAP server to process data on the server-side via a syntax in the DAP request.

This work was done by Michael Galloy of Tech-X Corporation for Goddard Space Flight Center. Further information is contained in a TSP (see page 1). GSC-16253-1

Data Compression Algorithm Architecture for Large Depth-of-Field Particle Image Velocimeters

Goddard Space Flight Center, Greenbelt, Maryland

A large depth-of-field particle image velocimeter (PIV) is designed to characterize dynamic dust environments on planetary surfaces. This instrument detects lofted dust particles, and senses the number of particles per unit volume, measuring their sizes, velocities (both speed and direction), and shape factors when the particles are large. To measure these particle characteristics in-flight, the instrument gathers two-dimensional image data at a high frame rate, typically >4,000 Hz, generating large amounts of data for every second of operation, approximately 6 GB/s.

To characterize a planetary dust environment that is dynamic, the instrument would have to operate for at least several minutes during an observation period, easily producing more than a terabyte of data per observation. Given current technology, this amount of data would be very difficult to store onboard a spacecraft, and downlink to Earth. Since 2007, innovators have been developing an autonomous image analysis algorithm architecture for the PIV instrument to greatly reduce the amount of data that it has to store and downlink. The algorithm analyzes PIV images and automatically reduces the

image information down to only the particle measurement data that is of interest, reducing the amount of data that is handled by more than 10^3 . The state of development for this innovation is now fairly mature, with a functional algorithm architecture, along with several key pieces of algorithm logic, that has been proven through field test data acquired with a proof-of-concept PIV instrument.

This work was done by Brent Bos, Nargess Memarsadeghi, Semion Kizhner, and Scott Antonille of Goddard Space Flight Center. Further information is contained in a TSP (see page 1). GSC-15960-1

Vectorized Rebinning Algorithm for Fast Data Down-Sampling

Applications include image processing, filter design, and anti-aliasing techniques.

Goddard Space Flight Center, Greenbelt, Maryland

A vectorized rebinning (down-sampling) algorithm, applicable to N-dimensional data sets, has been developed that offers a significant reduction in computer run time when compared to conventional rebinning algorithms. For clarity, a two-dimensional version of the algorithm is discussed to illustrate some specific details of the algorithm content, and using the language of image pro-

cessing, 2D data will be referred to as "images," and each value in an image as a "pixel." The new approach is fully vectorized, i.e., the down-sampling procedure is done as a single step over all image rows, and then as a single step over all image columns.

Data rebinning (or down-sampling) is a procedure that uses a discretely sampled N-dimensional data set to create a

representation of the same data, but with fewer discrete samples. Such data down-sampling is fundamental to digital signal processing, e.g., for data compression applications. Additional applications include image processing, filter design, and anti-aliasing techniques. Data rebinning is a computationally intensive procedure and thus the goal in this technology development is a more efficient algorithm

with reduced run times, as compared to existing rebinning approaches. This approach is able to take advantage of vectorized instructions such as Single Instruction Multiple Data (SIMD), to perform the rebinning operation.

The algorithm completely vectorizes the data rebinning operation, in the sense that a “single” arithmetic operation is applied simultaneously to multiply distinct data sets and is executed with the approximate run time of that operation applied to a single data set. For lower-level computer languages, such as C or assembly, vectorized operations can

be implemented using central processing unit (CPU) single-instruction, multiple-data (SIMD) capabilities, such as streaming SIMD Extensions 3 (SSE3) on x86 computer architecture or AltiVec on PowerPC processors. Thus, although the algorithm has been implemented using MATLAB, it is not fundamentally tied to MATLAB, and can be implemented using other programming languages.

The vectorized data rebinning (downsampling) procedure offers a reduced run time when compared with standard rebinning algorithms. In general, algorithms are often optimized by trading

decreased run time for increased memory, where the latter is needed for storing additional code, pre-computed results, or other ancillary data. However, the vectorized rebinning approach does not have increased memory requirements compared with conventional approaches. The underlying fundamental advantage to this technology is the utilization of vectorized instructions for the rebinning operation.

This work was done by Bruce Dean, David Aronstein, and Jeffrey Smith of Goddard Space Flight Center. Further information is contained in a TSP (see page 1). GSC-15949-1

➤ Display Provides Pilots with Real-Time Sonic-Boom Information

The impact of sonic booms can be controlled over populated areas.

Dryden Flight Research Center, Edwards, California

Supersonic aircraft generate shock waves that move outward and extend to the ground. As a cone of pressurized air spreads across the landscape along the flight path, it creates a continuous sonic boom along the flight track. Several factors can influence sonic booms: weight, size, and shape of the aircraft; its altitude and flight path; and weather and atmospheric conditions. This technology allows pilots to control the impact of sonic booms.

A software system displays the location and intensity of shock waves caused by supersonic aircraft. This technology can be integrated into cockpits or flight control rooms to help pilots minimize sonic-boom impact in populated areas. The system processes vehicle and flight parameters as well as data regarding cur-

rent atmospheric conditions. The display provides real-time information regarding sonic boom location and intensity, enabling pilots to make the necessary flight adjustments to control the timing and location of sonic booms. This technology can be used on current-generation supersonic aircraft, which generate loud sonic booms, as well as future-generation, low-boom aircraft, anticipated to be quiet enough for populated areas.

When fully deployed in real time, the display will leverage existing tools developed and enhanced by the U.S. Air Force and NASA to predict sonic boom parameters. The prediction data will be integrated with a real-time, local-area, moving-map display that is capable of displaying the aircraft’s current sonic

boom footprint at all times. The pilot will be able to choose from a menu of pre-programmed maneuvers such as accelerations, turns, or pushovers, and the predicted sonic boom footprint for that maneuver appears on the map. After fully developed and implemented, this will allow the pilot to select or modify parameters to either avoid generating a sonic boom or to place the sonic boom in a specific location. The system may also provide pilots with guidance on how to execute the chosen maneuver.

This technology will enable supersonic commercial flight without disturbing population centers on the ground.

This work was done by Ed Haering of Dryden Flight Research Center and Ken Plotkin of Wyle. Further information is contained in a TSP (see page 1). DRC-008-001

➤ Onboard Algorithms for Data Prioritization and Summarization of Aerial Imagery

Clustering/machine learning methods are used to structure data for prioritization, mapping, and downlinking.

NASA’s Jet Propulsion Laboratory, Pasadena, California

Many current and future NASA missions are capable of collecting enormous amounts of data, of which only a small portion can be transmitted to Earth. Communications are limited due to distance, visibility constraints, and competing mission downlinks. Long missions and high-resolution, multispectral imag-

ing devices easily produce data exceeding the available bandwidth. To address this situation, computationally efficient algorithms were developed for analyzing science imagery onboard the spacecraft. These algorithms autonomously cluster the data into classes of similar imagery, enabling selective downlink of represen-

tatives of each class, and a map classifying the terrain imaged rather than the full dataset, reducing the volume of the downlinked data. A range of approaches was examined, including k-means clustering using image features based on color, texture, temporal, and spatial arrangement.