cludes these features initially for zero cost. It is also assumed that the feature acquisition (FA) cost associated with each feature is known in advance, and that the FA cost for a given feature is the same for all instances. Finally, CFA requires that the base-level classifiers produce not only a classification, but also a confidence (or posterior probability).

CFA trains an ensemble of classifiers $M_0 \ldots M_f$ that use successively larger subsets of the features to classify instances. M_0 uses only the "free" (zero cost) features, and M_1 additionally in-corporates costly features F_1 through F_i. CFA reduces FA cost in that model M_i is trained only on instances that cannot be classified with sufficient confidence by model M_{i-1}. Therefore, values for feature F_i are acquired only for the instances that require it. At test time, each test instance is successively classified by M_0, M_1, M_2 . . . until its classification is sufficiently confident (i.e., until the confidence of the prediction reaches the confidence threshold). Again, features are acquired for the new instance only as required. In an empirical comparison with an existing method (Cost-Sensitive Naive Bayes) that makes acquisition decisions only during test time (and therefore requires that all training items be fully acquired), CFA achieves the same (or higher) level of performance at a much reduced cost (by at least an order of magnitude).

This work was done by Kiri L. Wagstaff of Caltech and Marie desJardins and James Mac-Glashan of the University of Maryland for NASA's Jet Propulsion Laboratory. For more information, contact iaoffice@jpl.nasa.gov. NPO-46886

Algorithm for Lossless Compression of Calibrated Hyperspectral Imagery

NASA's Jet Propulsion Laboratory, Pasadena, California

A two-stage predictive method was developed for lossless compression of calibrated hyperspectral imagery. The first prediction stage uses a conventional linear predictor intended to exploit spatial and/or spectral dependencies in the data. The compressor tabulates counts of the past values of the difference between this initial prediction and the actual sample value. To form the ultimate predicted value, in the second stage, these counts are combined with an adaptively updated weight function intended to capture information about data regularities introduced by the calibration process. Finally, prediction residuals are losslessly encoded using adaptive arithmetic coding.

Algorithms of this type are commonly tested on a readily available collection of images from the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) hyperspectral imager. On the standard calibrated AVIRIS hyperspectral images that are most widely used for compression benchmarking, the new compressor provides more than 0.5 bits/sample improvement over the previous best compression results.

The algorithm has been implemented in Mathematica. The compression algorithm was demonstrated as beneficial on 12-bit calibrated AVIRIS images.

This work was done by Aaron B. Kiely and Matthew A. Klimesh of Caltech for NASA's Jet Propulsion Laboratory. For more information, contact iaoffice@jpl.nasa.gov. NPO-46547

Universal Decoder for PPM of any Order

Complexity can be reduced and flexibility increased, at small cost in performance.

*NASA's Jet Propulsion Laboratory, Pasadena, California*

A recently developed algorithm for demodulation and decoding of a pulse-position-modulation (PPM) signal is suitable as a basis for designing a single hardware decoding apparatus to be capable of handling any PPM order. Hence, this algorithm offers advantages of greater flexibility and lower cost, in comparison with prior such algorithms, which necessitate the use of a distinct hardware implementation for each PPM order. In addition, in comparison with the prior algorithms, the present algorithm entails less complexity in decoding at large orders.

An unavoidably lengthy presentation of background information, including definitions of terms, is prerequisite to a meaningful summary of this development. As an aid to understanding, the figure illustrates the relevant processes of coding, modulation, propagation, demodulation, and decoding. An M-ary PPM signal has M time slots per symbol period. A pulse (signifying 1) is transmitted during one of the time slots; no pulse (signifying 0) is transmitted during the other time slots.

The information intended to be con-
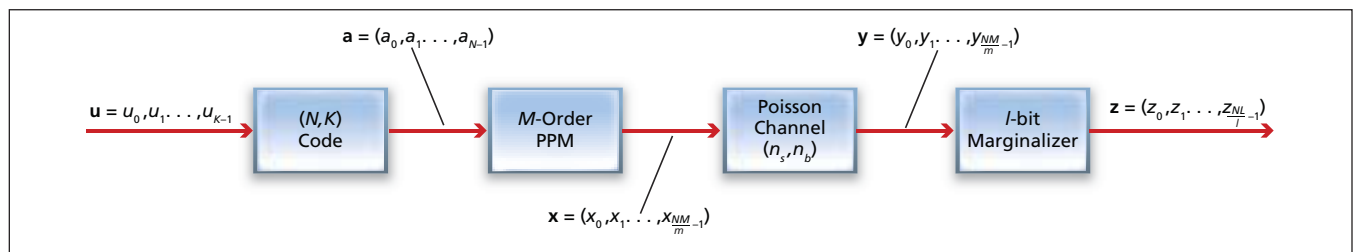
Processing of Information in an *M*-ary PPM communication system includes the sequence of steps depicted here. The *l*-bit marginalizer is a feature of the innovation reported here; the other features are typical of PPM systems in general.

veyed from the transmitting end to the receiving end of a radio or optical communication channel is a K-bit vector **u**. This vector is encoded by an (N,K) binary error-correcting code, producing an N-bit vector **a**. In turn, the vector a is subdivided into blocks of $m = \log_2(M)$ bits and each such block is mapped to an M-ary PPM symbol. The resultant coding/modulation scheme can be regarded as equivalent to a nonlinear binary code. The binary vector of PPM symbols, **x** is transmitted over a Poisson channel, such that there is obtained, at the receiver, a Poisson-distributed photon count characterized by a mean background count n_b during no-pulse time slots and a mean signal-plus-background count of n_s+n_b during a pulse time slot.

In the receiver, demodulation of the signal is effected in an iterative soft decoding process that involves considera-tion of relationships among photon counts and conditional likelihoods of m-bit vectors of coded bits. Inasmuch as the likelihoods of all the m-bit vectors of coded bits mapping to the same PPM symbol are correlated, the best performance is obtained when the joint m-bit conditional likelihoods are utilized. Unfortunately, the complexity of decoding, measured in the number of operations per bit, grows exponentially with m, and can thus become prohibitively expensive for large PPM orders. For a system required to handle multiple PPM orders, the cost is even higher because it is necessary to have separate decoding hardware for each order. This concludes the prerequisite background information.

In the present algorithm, the decoding process as described above is modified by, among other things, introduction of an l-bit marginalizer subalgorithm. The term "l-bit marginalizer" signifies that instead of m-bit conditional likelihoods, the decoder computes l-bit conditional likelihoods, where l is fixed. Fixing l, regardless of the value of m, makes it possible to use a single hardware implementation for any PPM order. One could minimize the decoding complexity and obtain an especially simple design by fixing l at 1, but this would entail some loss of performance. An intermediate solution is to fix l at some value, greater than 1, that may be less than or greater than m. This solution makes it possible to obtain the desired flexibility to handle any PPM order while compromising between complexity and loss of performance.

Algorithm for Stabilizing a POD-Based Dynamical System

Goddard Space Flight Center, Greenbelt, Maryland

This algorithm provides a new way to improve the accuracy and asymptotic behavior of a low-dimensional system based on the proper orthogonal decomposition (POD). Given a data set representing the evolution of a system of partial differential equations (PDEs), such as the Navier-Stokes equations for incompressible flow, one may obtain a low-dimensional model in the form of ordinary differential equations (ODEs) that should model the dynamics of the flow. Temporal sampling of the direct numerical simulation of the PDEs produces a spatial time series. The POD extracts the temporal and spatial eigenfunctions of this data set. Truncated to retain only the most energetic modes followed by Galerkin projection of these modes onto the PDEs obtains a dynamical system of ordinary differential equations for the time-dependent behavior of the flow.

In practice, the steps leading to this system of ODEs entail numerically computing first-order derivatives of the mean data field and the eigenfunctions, and the computation of many inner products. This is far from a perfect process, and often results in the lack of long-term stability of the system and incorrect asymptotic behavior of the model. This algorithm describes a new stabilization method that utilizes the temporal eigenfunctions to derive correction terms for the coefficients of the dynamical system to significantly reduce these errors.

Mission Reliability Estimation for Repairable Robot Teams

An analytical model demonstrates autonomous and intelligent control systems capable of operating distributed, multi-planetary surface vehicles for scouting or construction.

NASA's Jet Propulsion Laboratory, Pasadena, California

A mission reliability estimation method has been designed to translate mission requirements into choices of robot modules in order to configure a multi-robot team to have high reliability at minimal cost. In order to build cost-effective robot teams for long-term missions, one must be able to compare alternative design paradigms in a principled way by comparing the reliability of different robot models and robot team config-urations. Core modules have been created including: a probabilistic module with reliability-cost characteristics, a method for combining the characteristics of multiple modules to determine an overall reliability-cost characteristic, and a method for the generation of legitimate module combinations based on mission specifications and the selection of the best of the resulting combinations from a cost-reliability standpoint.

The developed methodology can be used to predict the probability of a mission being completed, given information about the components used to build the robots, as well as information about the mission tasks. In the research for this innovation, sample robot missions were examined and compared to the performance of robot teams with different numbers of robots and different numbers of spare components.