The **Multi-Stage ATR System Architecture** incorporates a detection stage that first identifies potential ROIs where the target may be present by performing a Fast Fourier domain OT-MACH filter-based correlation.

tions process has been developed to adapt to various targets and datasets.

The objective was to design an efficient computer vision system that can learn to detect multiple targets in large images with unknown backgrounds. Because the target size is small relative to the image size in this problem, there are many regions of the image that could potentially contain the target. A cursory analysis of every region can be computationally efficient, but may yield too many false positives. On the other hand, a detailed analysis of every region can yield better results, but may be computationally inefficient. The multi-stage ATR system was designed to achieve an optimal balance between accuracy and computational efficiency by incorporating both models.

The detection stage first identifies potential ROIs where the target may be present by performing a fast Fourier domain OT-MACH filter-based correlation. Because threshold for this stage is chosen with the goal of detecting all true positives, a number of false positives are also detected as ROIs. The verification stage then transforms the regions of interest into feature space, and eliminates false positives using an artificial neural network classifier.

The multi-stage system allows tuning the detection sensitivity and the identification specificity individually in each stage. It is easier to achieve optimized ATR operation based on its specific goal. The test results show that the system was successful in substantially reducing the false positive rate when tested on a sonar and video image datasets.

This work was done by Tien-Hsin Chao, Thomas T. Lu, and David Ye of Caltech; Weston Edens of Butler University; and Oliver Johnson of Harvey Mudd College for NASA's Jet Propulsion Laboratory. For more information, contact iaoffice@jpl.nasa.gov.

The software used in this innovation is available for commercial licensing. Please contact Daniel Broderick of the California Institute of Technology at danielb@caltech.edu. Refer to NPO-47012.

Single-Receiver GPS Phase Bias Resolution

NASA's Jet Propulsion Laboratory, Pasadena, California

Existing software has been modified to yield the benefits of integer fixed double-differenced GPS-phased ambiguities when processing data from a single GPS receiver with no access to any other GPS receiver data. When the double-differenced combination of phase biases can be fixed reliably, a significant improvement in solution accuracy is obtained.

This innovation uses a large global set of GPS receivers (40 to 80 receivers) to solve for the GPS satellite orbits and clocks (along with any other parameters). In this process, integer ambiguities are fixed and information on the ambiguity constraints is saved. For each GPS transmitter/receiver pair, the process saves the arc start and stop times, the wide-lane average value for the arc, the standard deviation of the wide lane, and the dual-frequency phase bias after bias fixing for the arc. The second step of the process uses the orbit and clock infor-

mation, the bias information from the global solution, and only data from the single receiver to resolve double-differenced phase combinations. It is called "resolved" instead of "fixed" because constraints are introduced into the problem with a finite data weight to better account for possible errors.

A receiver in orbit has much shorter continuous passes of data than a receiver fixed to the Earth. The method has pa-

rameters to account for this. In particular, differences in drifting wide-lane values must be handled differently. The first step of the process is automated, using two JPL software sets, Longarc and Gipsy-Oasis. The resulting orbit/clock and bias information files are posted on anonymous ftp for use by any licensed Gipsy-Oasis user. The second step is implemented in the Gipsy-Oasis executable, gd2p.pl, which automates the

entire process, including fetching the information from anonymous ftp.

Ultra-Wideband Angle-of-Arrival Tracking Systems

UWB radio pulses afford temporal resolution needed for estimating angles of arrival.

Lyndon B. Johnson Space Center, Houston, Texas

Systems that measure the angles of arrival of ultra-wideband (UWB) radio signals and perform triangulation by use of those angles in order to locate the sources of those signals are undergoing development. These systems were originally intended for use in tracking UWB-transmitter-equipped astronauts and mobile robots on the surfaces of remote planets during early stages of exploration, before satellite-based navigation systems become operational. On Earth, these systems could be adapted to such uses as tracking UWB-transmitter-equipped firefighters inside buildings or in outdoor wildfire areas obscured by smoke.

The same characteristics that have made UWB radio advantageous for fine-resolution ranging, covert communication, and ground-penetrating radar applications in military and law-enforcement settings also contribute to its attractiveness for the present tracking applications. In particular, the waveform shape and the short duration of UWB pulses make it possible to attain the high temporal resolution (of the order of picoseconds) needed to measure angles of arrival with sufficient precision, and the low power spectral density of UWB pulses enables UWB radio communication systems to operate in proximity to other radio communication systems with little or no perceptible mutual interference.

The figure schematically depicts a simple system of this type engaged in tracking a single UWB transmitter on a plane. The system includes two UWB-receiver assemblies, denoted clusters, separated by a known length d. Within each cluster is a UWB receiver connected to two antennas that are separated by a length a that is much shorter than the
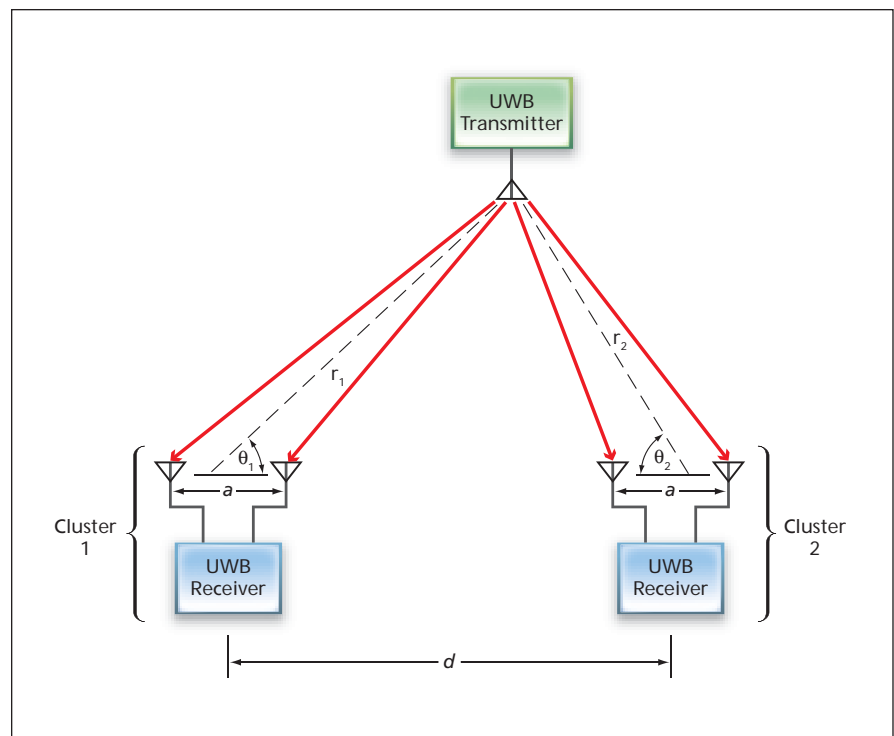
Angles θ_1 and θ_2 Are Estimated from differences between the times of arrival of UWB radio pulses at the antennas in each cluster. Then using these angles, the relative position of the transmitter is calculated by triangulation.

aforementioned length d. The signals received by the two antennas in each cluster are subjected to a process of cross-correlation plus peak detection to measure differences between their times of arrival. It is assumed that the distances (r_1 and r_2) between the clusters and the transmitter are much greater than a, as would usually be the case in most practical applications. Then the angles of arrival of the signals at the clusters are given by $\theta_1 \approx \arccos$ $(c\tau_1/a)$ and $\theta_2 \approx \arccos$ $(c\tau_2/a)$; where θ_1 and θ_2 are as shown in the figure; c is the

speed of light; τ_1 is the difference between the times of arrival of a pulse at the antennas in cluster 1; and τ_2 is the difference between the times of arrival of a pulse at the antennas in cluster 2. Then using θ_1 and θ_2, the two-dimensional location of the transmitter, relative to the known locations of the clusters, is calculated straightforwardly by use of the triangulation equations.

The processing of signals to determine the differences between their times of arrival, and the subsequent processing to determine the angles of ar-