Technology Demonstration by the Onboard Signal and Data Processor

Carl G. Pfeiffer and Brian L. Masson

The Midcourse Space Experiment (MSX) satellite carries as one of its experiments the Onboard Signal and Data Processor (OSDP). The OSDP's purpose is to demonstrate real-time detection and tracking of targets in space, using long-wave infrared data from MSX's Spatial Infrared Imaging Telescope III sensor. Hughes Aircraft built and delivered the OSDP flight unit. Calling upon the experience gained from developing the signal processor for the Airborne Surveillance Testbed Program, OSDP implements improved and simplified algorithms that allow single-scan acquisition functions and multiscan tracking functions to be merged in a single-board, radiation-hard processor. This unified capability will be required for future satellite surveillance programs, where multiscan data will be used to improve single-scan detection probability and radiometric and goniometric accuracy. Multiscan object tracking will also allow a scanning acquisition sensor to hand over track state vectors to an onboard staring sensor with a small field of regard, which can detect, track, and discriminate objects in the threat complex in a different waveband. This article describes the onboard signal and data processing functions that the OSDP demonstrates.

INTRODUCTION

The Onboard Signal and Data Processor (OSDP) was designed to process the long-wave infrared (IR) data from the Spatial Infrared Imaging Telescope III (SPIRIT III) sensor on the Midcourse Space Experiment (MSX) satellite. The OSDP flight unit was supplied by Hughes Aircraft, the builder of the "second generation" signal processor for the Airborne Surveillance Testbed Program. The OSDP is the "third generation" signal processor, and it is a pathfinder for the technology that will be needed for future operational surveillance systems. It implements improved and simplified algorithms that promise to reduce the computational load while substantially enhancing onboard processing capability.

The OSDP can track targets autonomously using two SPIRIT III focal plane assembly (FPA) wavebands (colors A and D) in the mirror-scan mode (Fig. 1). It will support a number of technology demonstrations, both on-orbit and off-line, to show that current state-of-the-art onboard signal and data processing technology can

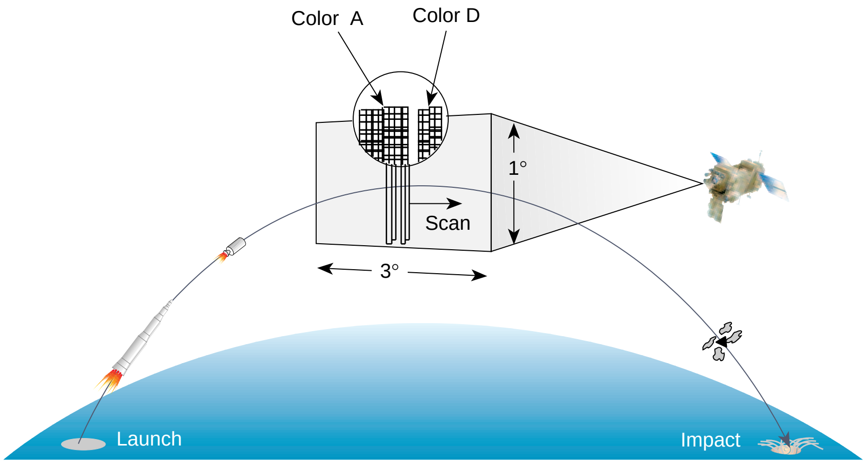

Figure 1. The Onboard Signal and Data Processor. Targets are tracked autonomously using two wavebands (colors A and D) of SPIRIT III's focal plane assemblies in mirror-scan mode.

meet the requirements of the next generation of space-based optical sensors. Some of the key objectives of these demonstrations are

- Tracking of up to 100 objects
- Background-adaptive thresholding
- Track initiation
- Scan-to-scan correlation of objects in track
- Acquisition and tracking of a designated object
- Jitter correction
- Object-velocity-corrected signal processing
- Birth-to-death tracking of object clusters

This article describes the functional architecture of the OSDP and the key functions of its two main components, the time-dependent processor (TDP) and the object-dependent processor (ODP).

OSDP HARDWARE

Figure 2 shows the OSDP flight hardware. The TDP is a combination of very large scale integration (VLSI), application-specific integrated circuit (ASIC) chips and field-programmable gate arrays. Its processing is called time dependent because the TDP operates upon all of the digitized detector samples that are output from the color A and color D focal plane assemblies to detect the presence of target objects. The ODP is implemented in Honeywell's radiation-hardened, generic, very-high-speed, spaceborne computer (GVSC), which hosts software programmed in Ada. Its processing is called object dependent because the ODP operates only upon data sets (strip reports) surrounding the target objects identified by the TDP. The OSDP hardware is fully redundant. The unit consumes 28 W, weighs 18 kg, and occupies a volume of 0.024 m^3.

OSDP FUNCTIONAL ARCHITECTURE

Figure 3 shows the OSDP's two main components, the TDP and the ODP, and their interfaces with each other, the sensor, and the spacecraft. The primary input from the sensor is the stream of digitized samples from the analog-to-digital converter. Also input are the sensor's scan mirror angles and times. The TDP performs gamma (spike) event circumvention, detector responsivity correction, background

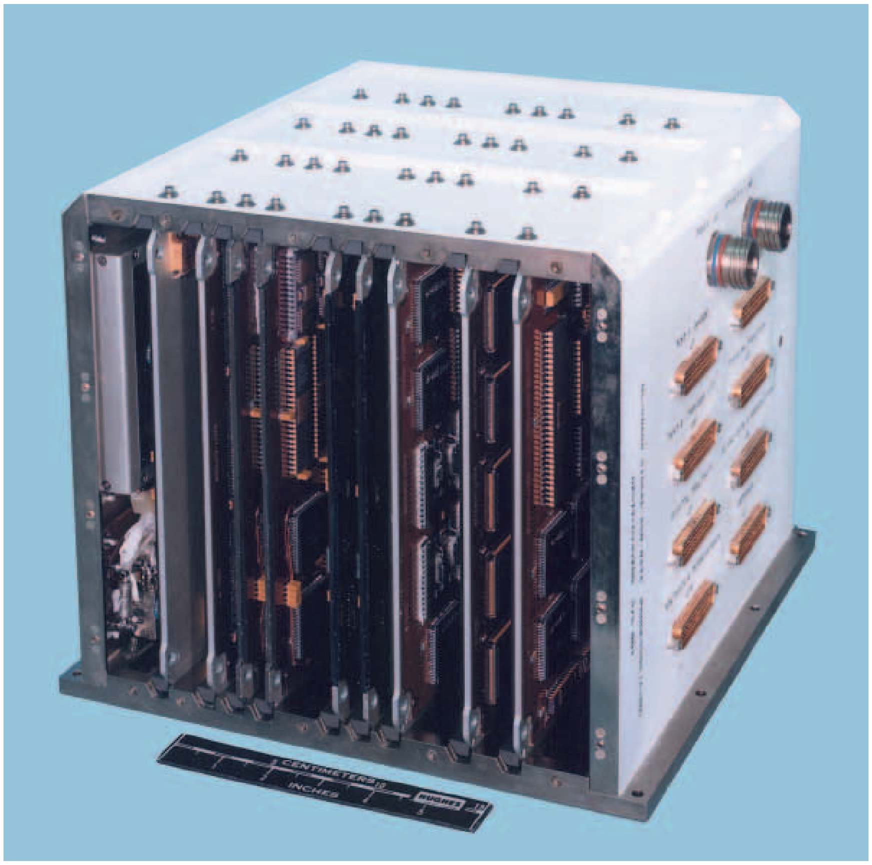

Figure 2. The flight hardware unit of the Onboard Signal and Data Processor.

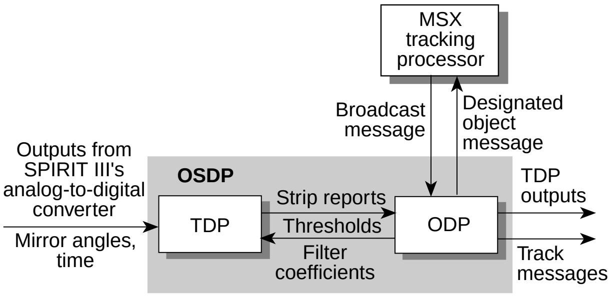

Figure 3. Time-dependent processor (TDP) parameters are fed back from the object-dependent processor (ODP) to improve signal processing performance.

subtraction, time delay and integration (TDI), matched filtering, and thresholding. The gamma pulses (spike events) are caused by such phenomena as cosmic ray and proton collisions with the FPA. The process of TDI involves boosting object signal-to-noise ratio (SNR) by coherently adding the outputs of detectors aligned in the scan direction of the sensor. Using the known scan rate, and assuming that object inertial motion is negligible, the data from each detector are stored in a TDI chain for the time delay required to cause each detector to observe the same spatial position. (The instantaneous field of view of each detector is approximately 90 × 90 μrad.) Matched filtering is used to test for the presence of an object by convolving the TDI output with the known pulse shape (i.e., the detector-convolved, optical point spread function). Thresholding is adaptive, with a local estimate of background standard deviation (σ) calculated in the TDP and a global threshold σ multiplier supplied by the ODP.

The ODP is capable of initiating and maintaining tracks on objects in the sensor's field of view. The availability of these tracks gives rise to a number of spin-off functional enhancements, such as (1) feedback of object velocity estimates to correct for velocity distortions in array correlation and matched filtering; (2) windowing, a convenient and computationally efficient form of scan-to-scan correlation; and (3) designated object tracking, where the MSX tracking processor provides (in the broadcast message) a startup state vector for an object to be tracked. The ODP then selects from its track file the object whose estimated position and velocity most closely match that of the designated object, and repeatedly returns accurate updated tracks based on fresh long-wave IR sensor measurements. Attitude estimates from the FPA, based on multiple star measurements, will be demonstrated in the OSDP data processing center.

The TDP–ODP interface and the manner in which the TDP passes thresholded data to the ODP are key to the OSDP design. The TDP hands off matched filter threshold exceedances to the ODP in groups called strip reports, constructed in a manner that maintains the integrity of clumps extending across many pixels. This function is called clump processing. Figure 4 illustrates the structure of strip reports and the rationale for clump processing. The figure shows the pixel space mapped out by the FPA in a portion of the scan where several pixels exceed threshold. The horizontal axis is the scan direction and the vertical axis is the cross-scan direction. A clump appears as a group of connected threshold exceedances. Ideally, the TDP would report the clump's centroid, individual threshold exceedances, and extent to ODP. Because this process would require excessive storage and logic in the TDP, however, the clump processing task has been partitioned such that the TDP does cross-scan processing and the ODP does in-scan processing. The TDP's criterion for constructing the cross-scan strips is that it reports not only threshold exceedances, but also every pixel that has a threshold exceedance as an immediate neighbor. This can result in strips containing no threshold exceedances, such as strips 14 and 15 in Fig. 4. This approach ensures that the ODP algorithms have adequate data for accurately computing a clump's centroid and moments and for deciding when to split and merge clumps.

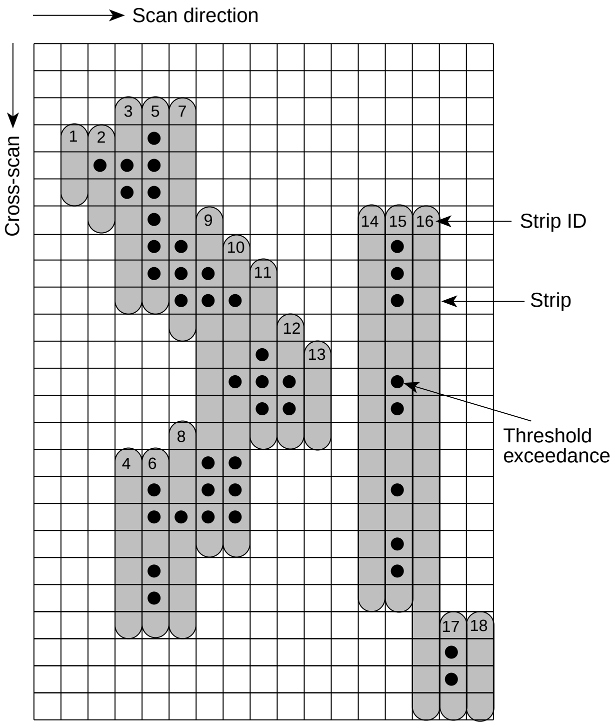

Figure 4. The time-dependent processor passes data in strips to the object-dependent processor.

THE TIME-DEPENDENT PROCESSOR

The TDP processes the samples separately from the color A and color D FPAs, performing the key functions described in the following paragraphs.

Gamma Circumvention

The gamma circumvention logic operates on the stream of contiguous samples from an individual detector before time delay and integration (Fig. 5). To be rejected as a gamma event, a sample must differ in amplitude from the preceding sample by a threshold value. In addition, the sign of the difference must differ from the sign of the preceding difference (between the preceding sample and *its* predecessor). This ensures that samples from the leading edge of a very bright object will not be falsely identified as a spike.

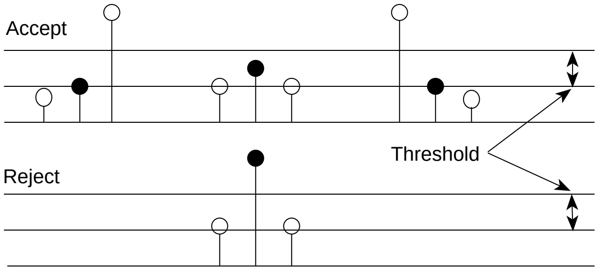

Figure 5. Spike-contaminated samples are identified by the magnitude and sign of amplitude changes.

Responsivity Correction

The amplitude of the signal from the detector is a nonlinear function of the number of photons incident in an integration period (Fig. 6). In the OSDP, three constants (C_0, C_1, C_2) are used to linearize the detector output as follows:

$$X = C_0 + C_1 Z + C_2 Z^2 ,$$

where Z is the detector signal and X is the linearized output. There is one set of unique coefficients for each detector for each of three gain ranges. These coefficients are stored in an electrically erasable, programmable read-only memory. They can be loaded on orbit.

Background Filtering

The background filter is the first stage of a two-stage matched filter designed to detect targets in the presence of noise and clutter. The optimal (maximum SNR) matched filter is

$$\begin{aligned} a_i &= \mathbf{f}^{\mathrm{T}}[R + \sigma^2 I]^{-1} \mathbf{z}_i \\ &= \mathbf{f}^{\mathrm{T}}[\sigma^2 I + R - R][R + \sigma^2 I]^{-1} \mathbf{z}_i / \sigma^2 \\ &= \mathbf{f}^{\mathrm{T}}[\mathbf{z}_i - b_i] / \sigma^2 , \end{aligned}$$

where a_i is the maximum likelihood estimate of peak target amplitude, assuming that the target is at FPA position i; $\mathbf{z}_i$ is the vector formed by concatenating relevant data samples from the two-dimensional array surrounding position i; $\mathbf{f}^{\mathrm{T}}$ is the corresponding target pulse shape vector; R is the clutter autocorrelation matrix associated with $\mathbf{z}_i$; σ^2 is the data noise variance; I is the identity matrix; and

$$\begin{aligned} b_i &= R\,[R + \sigma^2\, I]^{-1}\, \mathbf{z}_i \\ &= \text{maximum likelihood estimate of background.} \end{aligned}$$

Thus, the optimal target detection filter in clutter can be decomposed into a background estimation-and-subtraction filter, followed by a target detection filter matched to the noise-free pulse shape.

This approach offers two advantages: first, flexibility in adapting to the background with digital-design-friendly approximations to the background estimation filter; and second, reduced TDP memory requirements, since the target filter need span only the length of the target, instead of the autocorrelation length of the background. The background filter in OSDP is implemented as a first-order, recursive digital filter in the scan direction, with a programmable time constant. Like the functions that precede it, background filtering is performed on a single detector basis, thereby eliminating the need for detector offset correction. Convergence of the background estimate in the noise-free case is assured if the ith background estimate B_i for a given detector is

$$B_i = B_{i-1} + (1-K)(Z_i - B_{i-1}) ,$$

where Z_i is the data and K is the programmable constant. The background data residual now takes the simple form

$$\Delta B_{i+1} = (Z_{i+1} - B_i) = K\,\Delta B_i + (Z_i - Z_{i-1}) .$$

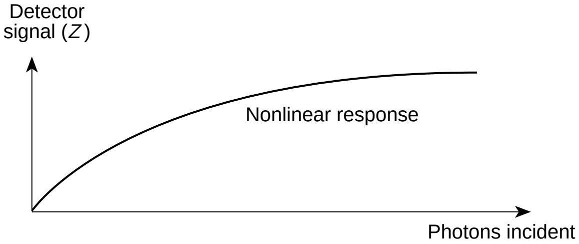

Figure 6. The amplitude of the signal from the detector is a nonlinear function of the number of photons incident in an integration period.

The computations are performed with three extra bits of precision to minimize the quantization noise generated by the filtering process. These three least significant bits are rounded at the output of the background residual filter.

Spike-Adaptive TDI

Time delay and integration are performed on the four-detector TDI sets of color A and the two-detector TDI sets of color D. The TDI function counts the number of good samples (not spike contaminated) that are accumulated during the integration of a TDI set, and the accumulated sum is normalized by the number of good samples.

Programmable Transversal Filter

The transversal (or matched) filter is a five-tap, finite-impulse-response digital filter used to process sequences of time delayed and integrated samples in the in-scan direction. The taps are programmable and can be changed by a ground-based user. To reduce the TDP-to-ODP data rate, filter outputs are forwarded every other sample time. The resulting effective TDP sampling rate is two samples per target pulse width, which is approximately at the knee of the SNR-versus-sample rate curve. The ODP uses a nonlinear position and amplitude interpolation algorithm to correct the phase error associated with the two-sample-per-pulse rate.

Adaptive Thresholding

The OSDP provides two threshold control modes. One is commanded; i.e., the detection threshold for TDP is set to a constant. The other is adaptive; i.e., the detection threshold is set to maintain a specified rate of false alarms or TDP strip reports using a local estimate of noise standard deviation. (A false alarm is an observation that failed either to correlate with an existing track or to initiate a track.) The detection threshold is a tabulated function of this value for each column of detector space. Thus, the detection threshold is determined by a combination of two factors: (1) the standard deviation of background noise in the local portion of the scene and (2) longer term (global) noise statistics or object observation rates, as determined by the lookup table stored in the TDP.

Clump Processing

The TDP constructs strip reports in the cross-scan direction and forwards them to the ODP.

THE OBJECT-DEPENDENT PROCESSOR

The ODP algorithms are divided into single-scan and multiscan functions as shown in Fig. 7. The single-scan functions serve three main purposes. First, they associate the incoming TDP strips with existing object tracks and stars. Second, they reconstruct the two-dimensional objects from the strips so that their characteristics can be determined. Third, they classify the objects as resolved objects, closely spaced object (CSO) pairs, and clumps and then determine whether clumps should be split or merged. The multiscan functions are functions that require data from multiple scans, or history from previous scans, to do their job. Examples are track initiation and update. The multiscan functions provide two classes of service. They first create data products such as object tracks that are used by downstream processing functions (e.g., discrimination) and external users. Then they feed back parameters such as object velocity estimates to the earlier processing functions, both TDP functions and ODP single-scan functions, enabling those functions to perform with higher fidelity.

The ODP Single-Scan Functions

The ODP single-scan functions support target tracking primarily by associating the observations from the TDP with windows representing objects in track and reconstructing the shape of these windows in preparation for the next scan. This process, which is key to the OSDP design, is illustrated in Fig. 8. The upper left portion of the figure shows a number of TDP strips being matched with a window predicted from the previous scan. Each of the four strips touches the window,

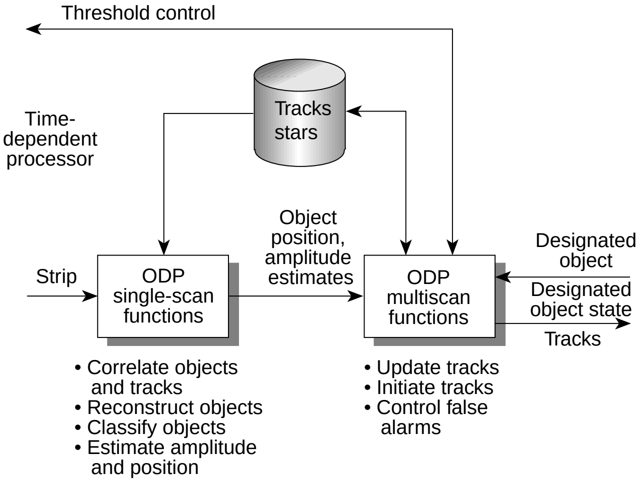

Figure 7. The algorithms of the object-dependent processor (ODP) are divided into single-scan and multiscan functions.

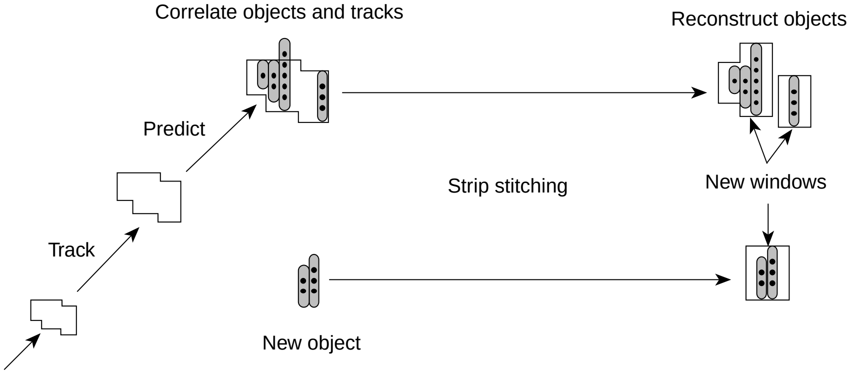

Figure 8. The single-scan functions minimize association processing. Because prediction accuracy is better than optical resolution, miscorrelation is unlikely.

so each is tentatively declared to belong to the object that the window represents. Each strip is independently tested for overlap with the window, since at this point in the processing the ODP knows nothing of adjacencies between strips. Next, the objects are reconstructed via strip stitching. All strips in the scan are tested for adjacency with other strips, regardless of their association with windows. Any strip then found to be adjacent to a strip associated with an object is declared part of the object. In this way, a clump can grow from one scan to the next as its member targets separate. Similarly, a change in an object's shape is recognized and its window adjusted accordingly. The breakup of a clump is identified by excessive space between strips associated with the clump, and a new object and window are created, as shown in Fig. 8.

This approach to correlating observations with tracks is based upon the understanding that sensors of the SPIRIT III class can track better than they can resolve; the 3σ track prediction uncertainty is less than the object resolution capability, so that multiple objects within a prediction window always show up as an unresolved clump. The OSDP design, therefore, circumvents one of the well-known problems in multitarget tracking: the potential for multiple observations to fall within a track prediction window, and the attendant processing burden of testing multiple, hypothetical combinations of observations and tracks. The OSDP prediction window is tailored to the shape of the object as seen in the last scan, thereby treating both resolved objects and clumps.

In addition to object/track correlation and object reconstruction, the single-scan functions include telescope attitude interpolation, computation of amplitude moments, object typing, and position and amplitude estimation. A brief description of the single-scan functions follows.

Telescope Attitude Interpolation

The line-of-sight vector of an object in scan mirror coordinates at the time of crossing the centerline of the scan mirror, as determined by the signal processor, is transformed into the inertial line of sight by knowledge of the inertial orientations of the mirror axes at that time. Interpolation is used to calculate these vectors, using mirror azimuth angles and the telescope attitude and attitude rates supplied by the MSX broadcast message. The interpolated telescope attitude and rate are also used to predict object windows for the next scan.

Window Prediction

To associate the incoming TDP data with objects already in track, the expected current-scan position of an object is first predicted in inertial coordinates and converted to FPA coordinates. A window is then constructed about this position in pixel space. This window consists of the last-seen pattern of TDP strips (see Fig. 4), plus a matching border strip on each side, plus a detector spacing on the top and bottom of each strip. As stated earlier, the 3σ prediction accuracy of the centroid of the window is expected to be less than the optical resolution capability of the sensor, so no assignment algorithm is needed to resolve window ambiguities.

Window Matching

An attempt is made to associate each TDP strip with a window. Each TDP strip in turn is compared with the window strips in the same column of pixel space, looking for overlap. If overlap is found, the TDP strip is tagged with the identification word (ID) of the window's object. If no overlap is found, the TDP strip is declared to be a component strip of a new object. The "stitching" process that follows associates it with adjacent strips in neighboring columns that are probable members of the same object.

Strip Stitching

The stitching process examines all TDP strips for adjacency to their fellow strips in neighboring columns, linking those that are connected. In general, the process results in a group of linked TDP strips tagged with the same object ID for each object in the track file, as well as some number of groups of linked strips with no

ID, representing objects not yet seen in enough scans to establish track. This process effectively reconstructs the objects as seen in the current scan, recognizing such effects as changes in the apparent shape of a clump due to reorientation of its member objects and breakup of a clump.

Amplitude Moments

The color A and color D FPAs each have a companion FPA with detectors offset in the cross-scan direction by one half of a detector to produce adequate cross-scan object resolution. The two halves were brought approximately into spatial alignment in the TDP by the data delay from the leading half. This alignment is now adjusted in the ODP on an object-by-object basis to compensate for object velocity variation across the FPA, as determined by multiscan tracking. These corrections result in a properly indexed, two-dimensional array of amplitude samples at half-detector spacing, referenced to the time of observation of the trailing half-FPA. Velocity-corrected, two-dimensional amplitude moments are now calculated, and properly indexed data can later be fetched for amplitude and position interpolation of single objects and CSO pairs.

Object Typing

The shape information from the computation of moments is used to classify the data set as one of four possible object types according to the following rules:

- *Gamma event:* streak in scan direction, equal to matched-filter impulse response function
- *Resolved object:* small circular object with width of a single (resolved) object
- *CSO pair:* width of a resolved object, length greater than one but less than two resolved objects
- *Clump of more than two objects:* none of the above

The classification process uses decision thresholds that will be tuned as part of the experiment. The classification algorithm is applied independently to each of the two colors. Conflicting assessments of object type are resolved by an empirically derived table look-up that is based on which colors favor which object types. The CSO pair position estimates are the weighted average of the individual color values, where the weighting factor is inversely proportional to the squares of the SNRs. The object amplitudes for each color are independently calculated. The classification assigned to a data set determines what happens to it next. Gamma events are dropped. Resolved objects are processed to establish the peak position and amplitude. The CSO pairs undergo special processing to estimate the peak of each object. Clumps do not undergo position and amplitude estimation; they are characterized in subsequent processing by their centroid position and peak amplitude (as determined from their moments).

Single-Object Position and Amplitude Estimation

The brightest matched filter output in the data set is selected, along with its four immediate neighbors in the in-scan and cross-scan directions. These values are used in a nonlinear algorithm that fits a two-dimensional quadratic surface to the five points, solving for two position components, two pulse shape parameters, and the peak amplitude. The output is the equivalent of a two-dimensional velocity adaptive, interpolated matched filter estimate of object position and amplitude.

CSO Pair Position and Amplitude Estimation

The separation vector of the objects is estimated from the amplitude moments, assuming that the objects are of equal amplitude. This quantity is reported, along with the position of the CSO centroid and estimated amplitude.

The ODP Multiscan Functions

Track Update

The track update process is the first of the multiscan functions in the processing chain. It is illustrated in Fig. 9. First, the line-of-sight vectors for all objects successfully matched with existing tracks by single-scan processing are used to update the line-of-sight state vectors (unit vector angular position, rate, and acceleration) from the track file. Until the track reaches a minimum angular or temporal length (“min1” in Fig. 9), a sequential batch method is used; thereafter, a Kalman filter is used. When the line-of-sight track reaches another specified minimum length (min2), a metric track initiation is attempted; i.e., the range and range rate are estimated using a modified version of Laplace orbit determination. If algorithm convergence is obtained, the metric track is maintained. Convergence depends upon viewing geometry.

Track Initiation

The track initiation process exploits the concept of “birth-to-death tracking,” where many-object threats are first seen as a few small clumps of unresolved targets. Tracking the clump centroids produces a velocity that can later be used to initialize the tracks of individual objects and subclumps that split off from the main clump. Track initiation of the initial clumps, and most other new objects, is not computationally expensive,

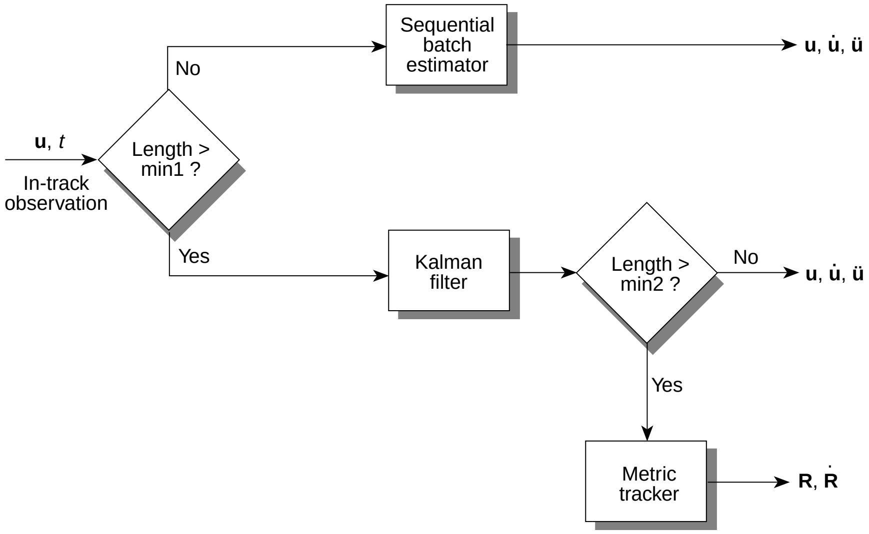

Figure 9. Track update is the first multiscan function in the processing chain. Track length determines update method. ($\mathbf{u}$ = line-of-sight vector in Earth-centered inertial coordinates at time t; $\dot{\mathbf{u}}, \ddot{\mathbf{u}}$ = derivatives; $\mathbf{R}$ = position vector; and $\dot{\mathbf{R}}$ = velocity vector.)

because new objects will be relatively few and their amplitude moments can be used for scan-to-scan association. A group velocity is calculated for objects that cannot be distinguished by their amplitude moments. Star tracks are initiated by postulating a zero inertial velocity. Each new track velocity is added to a list of "golden velocities," which can then be used for subsequent track initiation.

Track initiation is used to process all observed objects that failed to match the window of an existing track. The algorithm operates on unmatched measurements from successive scans, performing a sequence of tests to identify them as either stars, the designated object that the MSX tracking processor has requested OSDP to track, or new objects. Figure 10 illustrates the process. Individual uncorrelated observations from two successive scans are tentatively matched, and velocity is computed for each pair thus formed. If the computed velocity matches an entry in the golden velocity list within a tolerance, a window is created and an attempt is made to initiate a track with observations from the subsequent scans. Initially, the only entries in the golden velocity list are star velocity (zero) and the designated object, if any. Processing begins after the second scan of data has been collected. An observation from scan 1 is paired with one from scan 2. If the pair's velocity matches an entry in the list (star or designated object, at this point), a match is declared; if not, the scan 1 observation is tested with another scan 2 observation. After the scan 1 observation has found a match or the supply of unmatched observations from scan 2 is exhausted, the process is repeated with the second unmatched observation from scan 1 and the remaining unmatched observations from scan 2.

The process continues in this way until 20 matches have been found or all possible pairs of observations have been tested. If 20 matches have not been found, the shapes of any unresolved observations (clumps and CSO pairs) in the two scans are tested for similarity by examining their moments, and similarly shaped observations are identified as tentative matches. A group velocity is then computed from the centroids of all remaining unmatched observations in the two scans and entered in the golden velocity list. All matched observation pairs are then carried as candidate tracks. The last step in the scan 1/scan 2 processing is to make predictions for all candidates and compute windows for scan 3. If an observation is later not found in a candidate track's window in one of the next two successive scans, the scan 1 observation is discarded and the scan 2 observation is released into the pool of unmatched observations. After scan 3 is complete, the process just described is repeated with data from scans 2 and 3, and so on. When a candidate track has correlated with three observations, its velocity is added to the golden velocity list; thus, the list quickly grows over a period of a few scans to include all significant objects in the field of view. The expected result of this processing is that all objects of interest in the field of view, including stars, will be brought into track at the rate of approximately 20 per scan interval. All objects whose velocities match that of the designated object

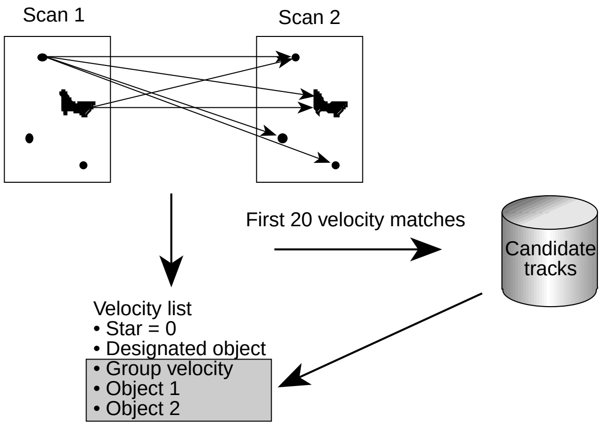

Figure 10. Tracks are initiated by velocity tests and scan-to-scan comparison of moments. The process repeats for scans 2 and 3. Moment comparison is used for closely spaced object pairs and clumps that fail velocity matching.

in the velocity list are reported to the MSX tracking processor as companions of the designated object; also reported is the average of all their state vectors.

CONCLUSION

The OSDP technology demonstrations on MSX will show that both single-scan acquisition functions and multiscan tracking functions can be merged in a single-board, radiation-hard computer. This capability will be needed for future satellite surveillance programs, where multiscan data will be used to improve single-scan detection probability and radiometric and goniometric accuracy. Multiscan object tracking will also allow a scanning acquisition sensor to hand over track state vectors to an onboard staring sensor with a small field of regard, which can detect, track, and discriminate objects in the threat complex in a different waveband. The next-generation matched filter algorithm will also implement clutter-adaptive background subtraction in the scanning sensor and streak detection (track-before-detect) in the staring sensor. These innovations will build upon the demonstrated OSDP design.

ACKNOWLEDGMENT: The MSX mission is sponsored by the Ballistic Missile Defense Organization. This work was supported under contracts C836410 and C903764 with the Space Dynamics Laboratory of Utah State University.

THE AUTHORS

CARL G. PFEIFFER is a Principal Scientist/Engineer in the Electro-Optical Engineering Division of Hughes Aircraft Company. He designed the signal processing algorithms and logic for the Forward Acquisition System, Airborne Surveillance Testbed, and Onboard Signal and Data Processor programs, and has patented hardware implementations of signal processing functions. He is currently in charge of Onboard Signal and Data Processor design for the Space and Missile Tracking System (formerly Brilliant Eyes) Flight Demonstration System. He holds a B.S. degree in mechanical engineering from Cornell University. Mr. Pfeiffer's e-mail address is cgpfeiffer@ccgate.hac.com.

BRIAN L. MASSON is a Scientist/Engineer in the Sensors and System Engineering Center at Hughes Aircraft Company. He is currently working on electro-optical targeting and navigation systems for tactical weapon platforms, including fixed-wing aircraft. His past experience includes the Space and Missile Tracking System (formerly Brilliant Eyes), the Space Surveillance and Tracking System, and the Follow-on Early Warning System. He led the sensor design team on the Early Midcourse Option study that culminated in the definition of the Space and Missile Tracking System concept. He holds an M.S. degree in electrical engineering from Syracuse University. Mr. Masson's e-mail address is blmasson@ccgate.hac.com.