The TRADEX Multitarget Tracker

Glenn W. Meurer, Jr.

■ The Multitarget Tracker (MTT) is a real-time signal processing and data processing system installed in the TRADEX radar at the Kiernan Reentry Measurements Site (KREMS) on Kwajalein Atoll in the Marshall Islands. The TRADEX radar is a high-power, high-sensitivity instrumentation radar that was originally designed to track and gather signature data on a single target. The MTT is designed to detect and track as many as 63 targets within the beam of the radar. It provides data necessary for determining the angular locations and ranges of all of these targets, as well as signature data necessary for target identification. The TRADEX MTT is unique because it utilizes a large, mechanically steered, pencil-beam antenna, whereas other MTT systems generally rely on electronically steered antennas or rotating antenna platforms. The MTT system automatically processes received signals, reports targets, initiates and maintains target track files, and presents target information to the radar operators through real-time interactive graphical displays. This information is given to the KREMS Control Center and from there is made available to other systems in the test range. This article presents an overview of the TRADEX MTT system and discusses its implementation, application, and operation.

THE MULTITARGET TRACKER (MTT), an integral part of the Target Resolution and Discrimination Experiment (TRADEX) radar system, is the most recently added significant capability in a long series of improvements to TRADEX. The TRADEX radar system is one of four instrumentation radars located at the Kiernan Reentry Measurements Site (KREMS) on Kwajalein Atoll in the Marshall Islands [1]. The KREMS facility is part of the Kwajalein Missile Range (KMR), which is the terminus of the Western Test Range. The KMR performs the technical functions of the United States Army–Kwajalein Atoll (USAKA).

The KMR sensors and the KREMS radars are ideally located to observe, detect, and track orbiting objects as well as reentering objects over the Western hemisphere. The sensors lie at the same latitude as Panama and the same longitude as New Zealand; this location is near the equator and approximately 2100 miles southwest of Hawaii. Figure 1 shows the location of Kwajalein Atoll in a map of the area.

The four KREMS radars (ALTAIR, TRADEX, ALCOR, and MMW) are located on the island of Roi-Namur, which is at the northeastern tip of Kwajalein Atoll. Facing the northwest, these radars can view vehicles launched from the Western Pacific nations and can determine intended orbit and system type (such as an ICBM test or satellite deployment). The radars also track orbiting objects such as satellites or space shuttles. Figure 2 shows the locations of the four KREMS radars on the island of Roi-Namur. The Range Operations Center is located on the island of Kwajalein, which is at the southern tip of Kwajalein Atoll. Many additional range sensors are located throughout Kwajalein Atoll.

Test launches of intercontinental ballistic missiles from Vandenberg Air Force Base in California impact in the vicinity of Kwajalein Atoll. Improved data-collection capabilities for these multiple-reentry-vehicle test launches were one of the main reasons for the development of the MTT. These test firings are performed to measure the atmospheric effects of reentry, the spatial deployment, dynamics, and targeting of reentry vehicles, and a variety of other radar observables. The MTT tracks the missile-system components (including stages, deployment debris, instrumentation packages, reentry vehicles, decoys, and associated objects) beginning at mid-course and until they leave the beam of the radar. The MTT also independently tracks the single target that TRADEX tracks through reentry into the Earth's atmosphere and until impact.

The TRADEX Radar

The TRADEX radar has had a long and varied history. Conceived in 1959 and made operational in 1962 as a dual-frequency UHF and L-band radar, TRADEX has been in nearly continuous service at KREMS for 30 years. In 1965 a VHF capability was briefly added, and a major upgrade in 1972 brought it to its present dual-frequency L-band/S-band configuration. In 1983 the Multistatic Measurement System (MMS) was added to give TRADEX a dual-bistatic signature capability and allow it to trilaterate for improved tracking metrics. A coherent integration capability was added in 1986 to provide the radar with the gain needed to see reentry-vehicle-sized objects as they appear on the horizon. Work on the MTT began in 1988, and its basic operational capabilities were brought on line in 1991. Figure 3 is a photograph of the antenna and pedestal of the TRADEX radar. The pedestal rests on top of the building that contains the TRADEX radar equipment and the KREMS Control Center (KCC).

TRADEX is a large, mechanically steered, pencil-beam tracking radar system. Its parabolic reflector antenna is 25.6 m in diameter. Its two-way 6-dB beamwidth is 10.6 mrad, or 0.61°. For comparison, the moon is approximately 0.5° in angular width when viewed from Earth. This relatively narrow

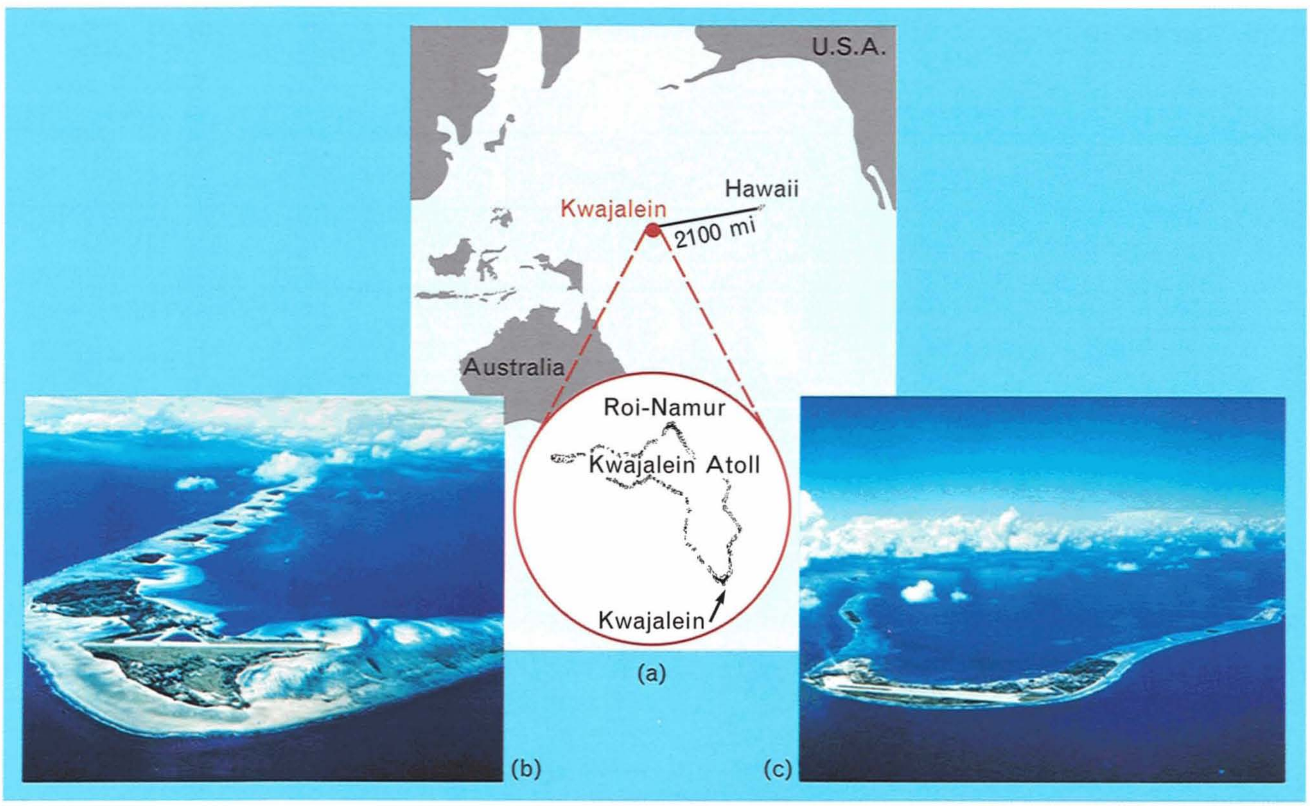

FIGURE 1. (a) The location of Kwajalein Atoll in the South Pacific. (b) The four KREMS radars, including TRADEX, are located on the island of Roi-Namur. (c) Range operations are located on the island of Kwajalein.

FIGURE 2. Aerial view of the KREMS facility on the island of Roi-Namur. The TRADEX radar facility is in the right foreground.

beam, which concentrates a significant amount of power, contributes to the long acquisition range of TRADEX; the antenna gain at L band is 48.2 dB. The beam is broad enough at long ranges, however, to illuminate a large cross-range extent, and hence a number of objects simultaneously.

The static weight of the antenna is 279,000 pounds, or 140 tons. It is fixed to an elevation-over-azimuth pedestal that can steer through 290° in azimuth and 180° in elevation. Thus TRADEX is capable of pointing to any location in its hemisphere of view. Steering TRADEX, however, requires a lot of power. Each of the two drive mechanisms (azimuth and elevation) is powered by a three-phase, 125-horsepower motor driving a pump for four hydraulic motors, each of which drives the antenna through a 624:1 (elevation) or 324:1 (azimuth) gear train. Even with this much drive power, the dynamic motion of the antenna is still restricted. The velocity and acceleration limits for TRADEX are 218 mrad/sec (12.5°/sec) and 230 mrad/sec^2 (13.2°/sec^2), respectively. These upper limits, which are imposed to protect the antenna from being torn apart by inertial forces, necessarily restrict its agility. This restriction is what makes the TRADEX MTT unique; radars that track a multitude of distinct objects usually rely on scanning platforms or electronically steered beams to detect, resolve, and track multiple targets. The TRADEX MTT operates only on multiple targets visible in its beam.

TRADEX utilizes a dual-frequency focal-point feed,

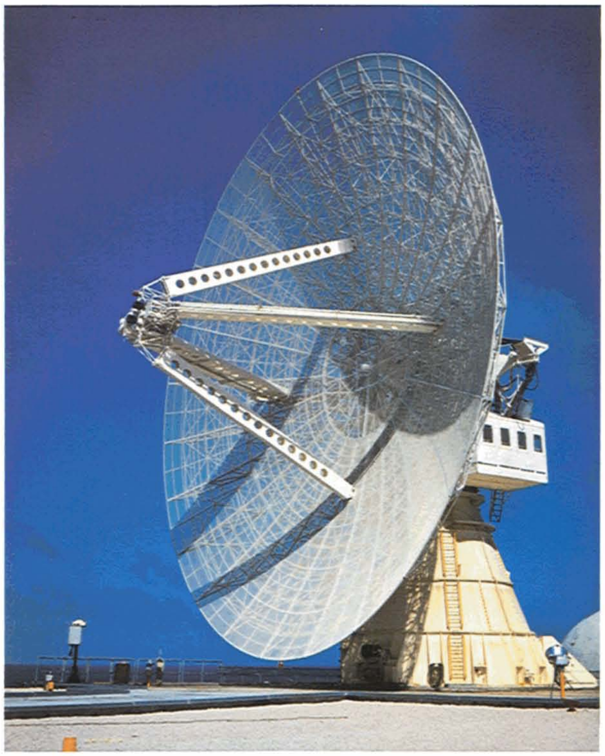

FIGURE 3. The TRADEX parabolic antenna is 25.6 m in diameter. Its two-way 6-dB beam width is 10.6 mrad, or 0.61°, which is slightly larger than the angular width of the moon.

which includes a five-horn monopulse feed for L band, and a coaxial S-band feed that rests inside the L-band reference horn. Because angle metrics are possible only at L band (no angle-error horns are associated with the S-band feed), the MTT is designed to support only the L-band part of the TRADEX radar.

The Multitarget Mission

Sophisticated missile test systems and launches are costly. In an effort to mitigate the costs of a reentry test, several experiments are often flown simultaneously. Economics force this approach because the cost of a launch often outweighs the costs of the individual experiments. As a result, the radar operators at KREMS see an increasing number and variety of radar-observable objects in these test missions. In addition to test reentry vehicles, missile flight systems often carry support vehicles such as telemetry packages and observation platforms. Each object that is deployed in a test can also contribute deployment debris, such as exploded bolts, de-spin weights, springs, shields, and thrust plates. The primary objective of the mission, however, must still be accomplished. This objective is to find, identify, and track particular test vehicles, while gathering the radar signature data needed to determine certain behaviors of these vehicles. The MTT was designed to assist the TRADEX operators in accomplishing this requirement.

Figure 4 illustrates a typical multitarget test mission. The multistage rocket is launched from Vandenberg Air Force Base in California. The total travel time until impact near Kwajalein Atoll is thirty minutes. At approximately 480 seconds after lift-off the first four reentry vehicles (RV1, RV2, RV3, and RV4) are released from the bus vehicle. One minute later, RV5 and RV6 are released, and one minute after that RV7 and RV8 are released. The different release times result in range and angular separation of the groups of targets as seen by the radar, and cause the RVs to impact in different locations. At approximately eighteen minutes into the test mission (twelve minutes to reentry), the bus vehicle and the RVs appear above the horizon and can then be seen by the TRADEX radar. The last four targets released from the bus vehicle (RV5 through RV8) form the first target group to be seen by the radar; the first four targets released (RV1 through RV4) form the second target group.

The Purpose of TRADEX

TRADEX plays several roles during the course of a mission. From the earliest acquisition the radar operators start to count the number of observable objects in the target group. This role, which is called *complex evaluation*, determines whether the target group has been deployed as expected. The operators then attempt to establish single-target tracks on as many individual objects as they can. These tracks can be used by the narrower-beam systems (e.g., ALCOR and MMW) for acquisition. TRADEX also collects individual intervals of metric data and signature data. Finally, TRADEX tracks a single test object through reentry and until impact.

The TRADEX radar usually acquires the large bus vehicle approximately twelve minutes to reentry, when it appears above the horizon; this point is called the *horizon break*. At this time, the TRADEX radar beam illuminates a cross-range distance of approximately

45 km. TRADEX first establishes track on the bus vehicle, which is much larger than the RVs and can be seen at greater ranges. Without the MTT, TRADEX tracks only this one target. TRADEX operators can observe the signature of this target, but they cannot observe the signatures of the other objects in the first target group (e.g., RV5 through RV8). These other objects are in the beam of the radar but are not tracked.

In this single-target tracking mode, the TRADEX operators systematically move the radar beam from one target to the next as the targets become visible (i.e., have sufficient signal to noise). The operators track as many of the objects as they can find and use the signature data and metric data to identify each object in turn. The state vectors obtained for these objects are stored by the computer for later use in reacquiring any of these objects.

At a predetermined time, the TRADEX radar operators stop tracking objects in the first target group (RV5 through RV8 and the bus) and move the beam of the antenna to cover the second target group (RV1 through RV4). The time-consuming object-by-object track process is then repeated for the second target group. TRADEX typically needs ten to twelve minutes to track and identify all of the objects in this type of test mission. At two minutes before reentry, TRADEX reacquires a particular reentry vehicle from the first target group, and tracks that vehicle until impact. If time permits, the TRADEX operators reacquire a particular reentry vehicle from the second group, and track it until impact. The reentry phase for a target vehicle typically lasts thirty seconds.

The Purpose of the MTT

When TRADEX is in single-target mode, and when

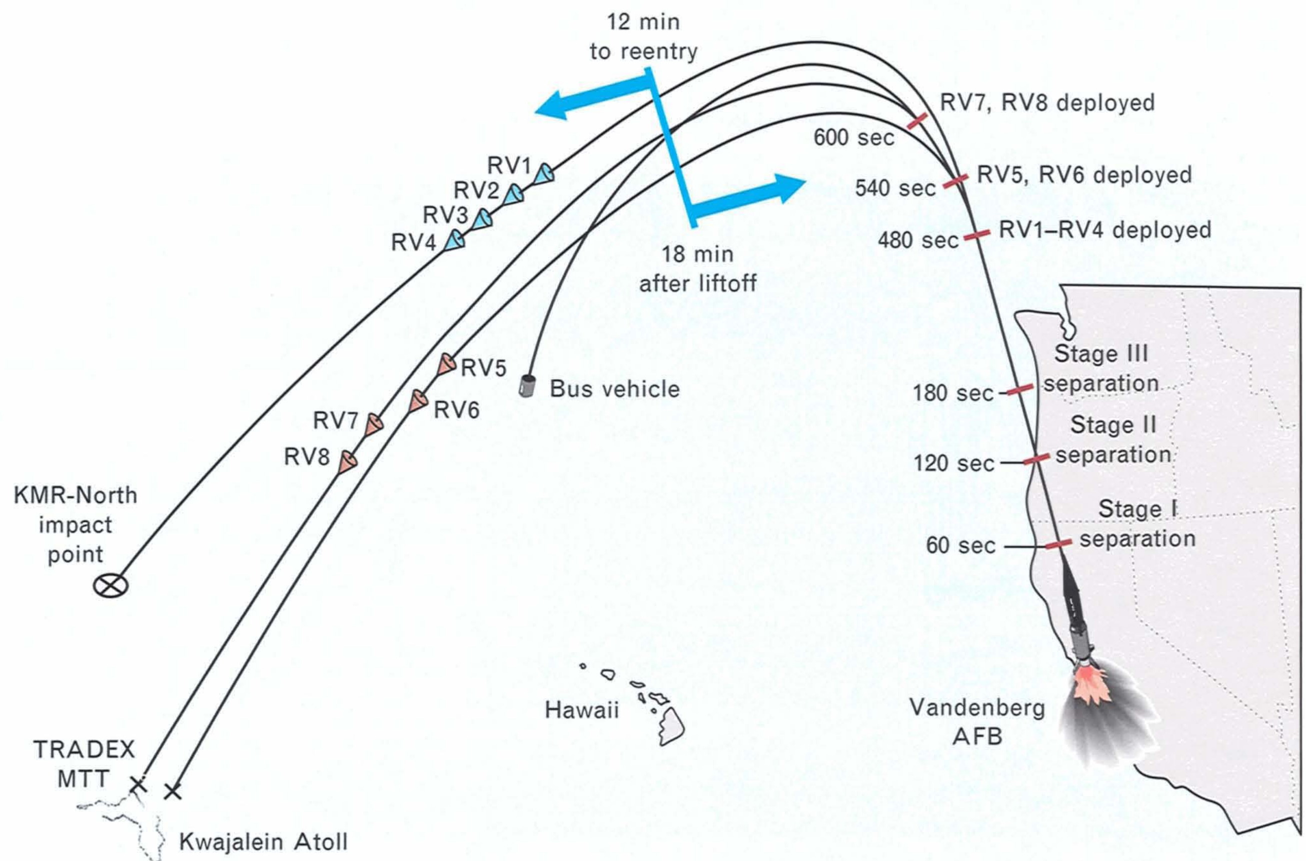

FIGURE 4. A typical multitarget test mission. A rocket is launched from Vandenberg Air Force Base in California and the multiple reentry vehicles impact in the vicinity of Kwajalein Atoll approximately thirty minutes later. The TRADEX radar first sees the bus vehicle and the reentry vehicles approximately twelve minutes before reentry.

the number of targets is greater than presented in this uncomplicated example, or when the launch deployment is non-nominal, the operators require more time to acquire, track, identify, and collect data on multiple objects than is available before reentry. The MTT reduces the need for the radar to hop from one object to the next, which is a labor-intensive task for the operators and which results in reduced data coverage. Longer or continuous data coverage improves the real-time and post-mission data-processing capabilities, which results in better trajectories and better data for studying subtle radar observables.

The MTT frees the operators to focus on the data collected, and provides continuous tracking of all objects in the target group. The MTT also allows the radar to allocate its recording resources more efficiently by focusing data recording only on objects and regions of interest. Focusing data collection on actual targets reduces the data-sampling throughput and allows the radar to operate at higher pulse repetition rates, which makes possible the analysis of higher-bandwidth observables on multiple objects.

With the addition of the MTT, all visible objects in the beam are automatically acquired and tracked while the TRADEX radar tracks its single target. The MTT initiates a track file, maintains state vectors, and buffers the radar cross section (RCS) of each visible object. By using information shown on the MTT graphical-data displays the radar operators can observe the signature and metrics of all the targets in a target group simultaneously. This information allows the operators to identify particular objects more quickly; it can also be shared with the KCC, where it is compared to other sensors' files and sent to other test-range systems to assist with acquisition, identification, and data collection.

The MTT can track typical targets from horizon break until loss of signal (i.e., when the target leaves the beam, burns up, or splashes). This capability gives the radar operators enough time to observe all visible targets, evaluate the target complex, and decide which target is their most important objective.

A recent and ongoing upgrade to the KCC facility automates much of the process described above. When the automated KCC is fully established, the MTT will send all of its objects to the KCC, where the data-fusion and target-identification functions will be performed algorithmically by a computer rather than by human operators. This new automated system changes the role of the MTT analyst, because the automation software will ultimately find and identify the appropriate targets. This automated system frees the MTT operators to concentrate on the data collected rather than on the mechanics of target acquisition and identification.

How the MTT Works at TRADEX

Figure 5 illustrates the operation of the TRADEX MTT. It shows the 10.6-mrad radar beam and depicts a complex of targets illuminated by the beam. At a nominal acquisition range of 4200 km, the radar beam covers a cross-range extent of approximately 45 km. This cross-range extent normally encompasses all vehicles within a single group, and in some cases illuminates multiple groups. As shown in the figure, the TRADEX sample pattern is designed to cover only the prime target, while the MTT sample pattern covers a range interval of up to 655 km.

Each target seen by the MTT is first resolved in range. For each resolved target, the angles relative to the boresight are then observed. Although multiple targets are illuminated and their backscattered energy is received by the radar, the off-axis targets are not fully illuminated and are not optimally received. This effect is known as *beam loss.* In general, because the maximum response of the TRADEX antenna is at the center of its beam, or *boresight,* only the on-axis target is optimally illuminated and received. With calibration of the monopulse system, the radar's angle-channel observations can be converted into known angle offsets from the boresight. Once the angle offsets are determined, the signature data from off-axis targets are corrected for beam loss. This process poses a new challenge, though, as the monopulse systems of these large, boresight-tracking radars have heretofore been well calibrated only in the area near the boresight. The full angle-calibration method and results are described later in this article.

MTT System Overview

The TRADEX MTT is a signal processing and data processing adjunct that operates in parallel with the

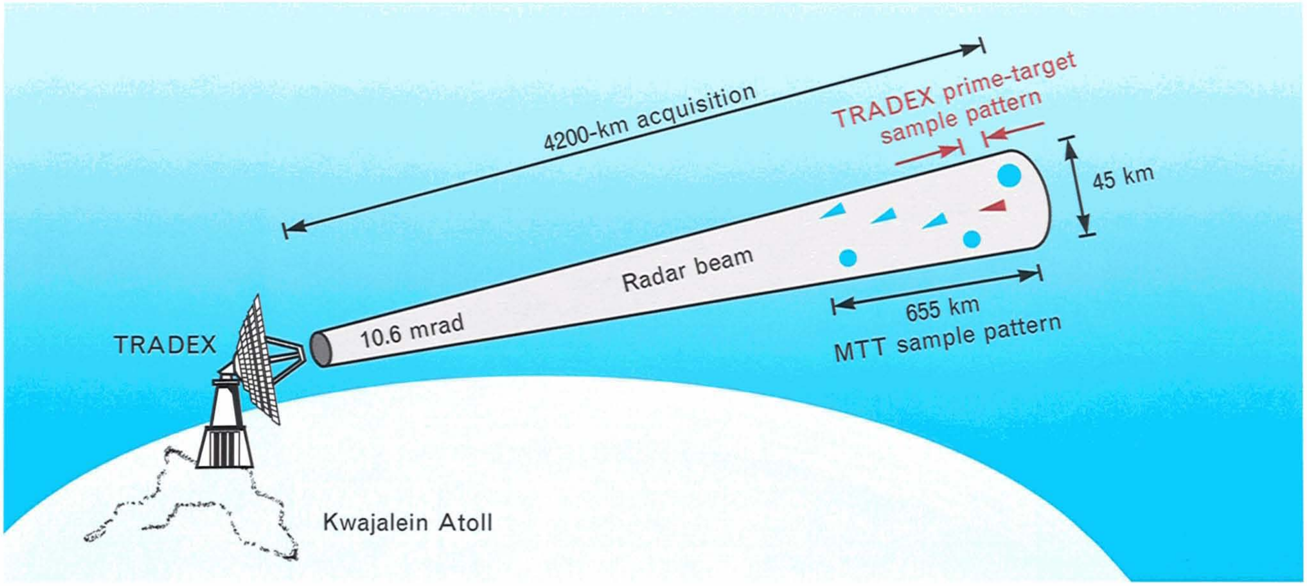

FIGURE 5. Geometry of the TRADEX MTT system. The TRADEX beam is centered on the prime target, and the radar collects data from a relatively small range interval around that target. At a nominal acquisition range of 4200 km, the radar beam covers a cross-range extent of approximately 45 km, which illuminates many other targets. The MTT collects data from a relatively large range interval that includes all visible targets in the beam.

normal TRADEX radar signal processing and data processing systems. The MTT is not specifically designed to replace existing TRADEX capabilities, such as primary track and initial target acquisition, but it could accomplish these tasks if necessary. The MTT was designed to add significant new functionality to the TRADEX system. In particular it generates track files on all radar-observable objects. It makes these track files available to TRADEX for control of independent data-collection sample sets and to the KCC for pointing of other range instruments.

To maximize the number of targets seen by the MTT, the radar operators pick a target near the angular center of the target group (as seen by the radar) and make this the primary, or *prime*, target. Alternatively, they can direct the radar to illuminate and sample a point in space that represents the geometrical center of the group of targets. Directing the TRADEX radar in this manner maximizes the number of targets seen by the MTT.

Figure 6 shows a block diagram of several important TRADEX radar subsystems, including the TRADEX host computer and its Real-Time Program (RTP), the timing subsystem, the L-band receivers, the L-band data-recording system, and the high-speed data recorders. The TRADEX host computer consists of two GOULD SEL 9780 mainframe computers sharing a common database. These computers, which are connected to each other in tandem, provide the computational power to support the control of all radar systems.

The host computer communicates with the timing subsystem every 100 msec, sending a *frame data block* containing all information necessary for control of TRADEX during the next 100 msec. The timing subsystem parses this frame data block and, in turn, sends the appropriate control words and strobe signals to other radar subsystems such as the S-band and L-band transmitters, the S-band and L-band recording systems, and the MTT.

Important systems added to the TRADEX radar as part of the MTT include the Multitarget Integrator Box (MIB), the MASSCOMP Display and Tracking Computer (MADTraC) and its RTP, and the various system interfaces. Figure 7 shows these systems and their interdependencies.

The Multitarget Integrator Box

The MIB is a custom high-throughput digital signal processing system. It consists of seven analog-to-digi-

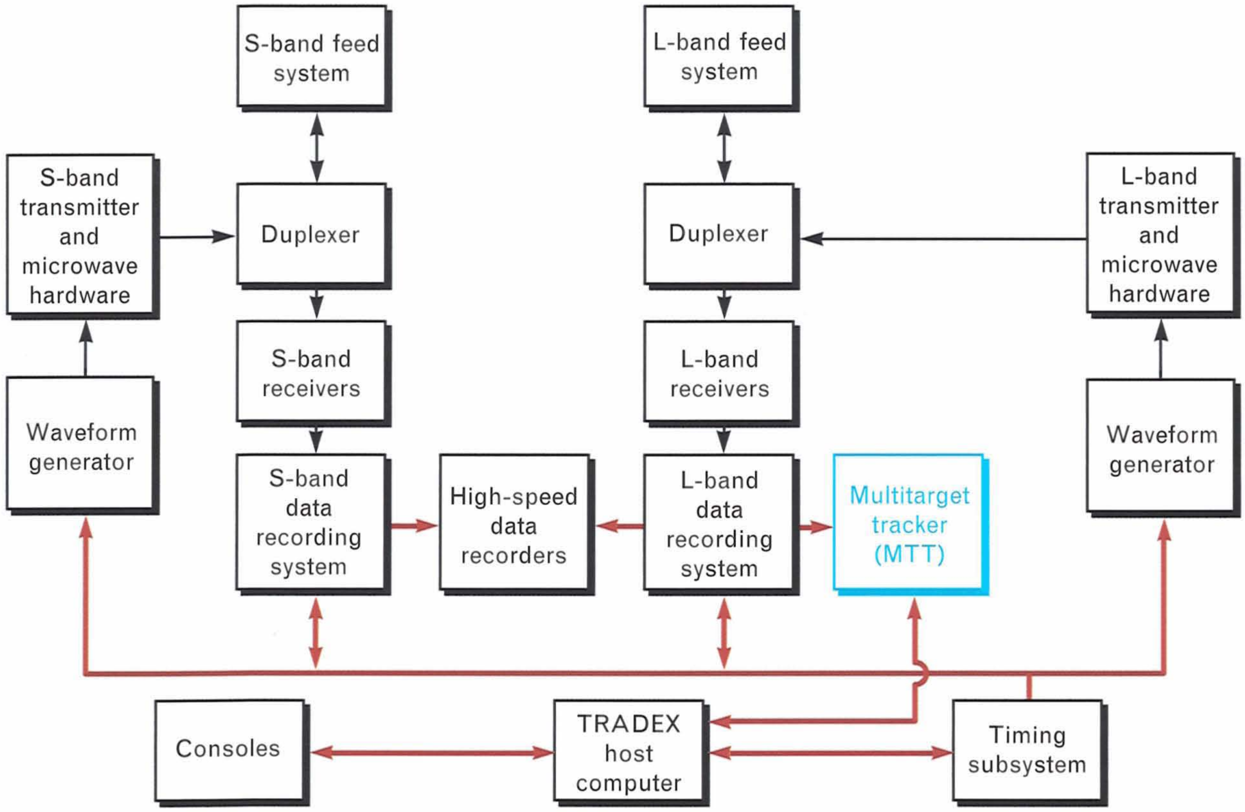

FIGURE 6. Block diagram showing TRADEX subsystems. The MTT sits between the L-band data-recording system, which provides the analog video signals from the radar, and the TRADEX host computer. The TRADEX host computer and its real-time program provide data processing and control for the entire radar system.

tal converters and a Motorola 68020-based VMEbus single-board computer with an interface to seven custom processing boards. The boards include a two-to-one data-slowdown multiplexer, two high-speed triple-buffered integrators, a constant false-alarm-rate (CFAR) processor, and a range-time-intensity (RTI) pixel processor.

The MIB processes data in batches called the *data frame.* A data frame, which lasts for 100 msec, usually contains between 18 and 300 radar pulses. The data frame synchronizes the MIB to the processing interval of the TRADEX host computer, which updates its database ten times per second (all of the KREMS computer systems operate in synchronization with a site-standard reference of ten updates per second).

During one special radar mode, which is called the *extended-range mode,* the processing interval of the MIB is changed to a multiple number of data frames. During extended-range operations, the TRADEX array-processor system collects a number of data samples (up to 1024, or 1k) for the purpose of coherent integration. During this extended processing interval the MIB integrates continuously, which provides additional noise variance reduction (noncoherent integration). The returns from as many as 511 radar pulses can be integrated. At a radar pulse repetition frequency of 100 pulses/sec, this extended-range interval results in a processing cycle of 5.1 sec.

The MIB quantizes the seven analog video signals provided by the TRADEX receivers, calculates the monopulse phase, integrates the returns in each range cell over an interval of 100 msec, and then automatically detects the presence of targets. It collects all of the data necessary for target detection, marking, and tracking (reference and monopulse angle-channel signals, range, and signal-to-noise ratio) and formats it into a target report. The MIB processes up to 128k range samples per radar pulse, and up to 300 radar

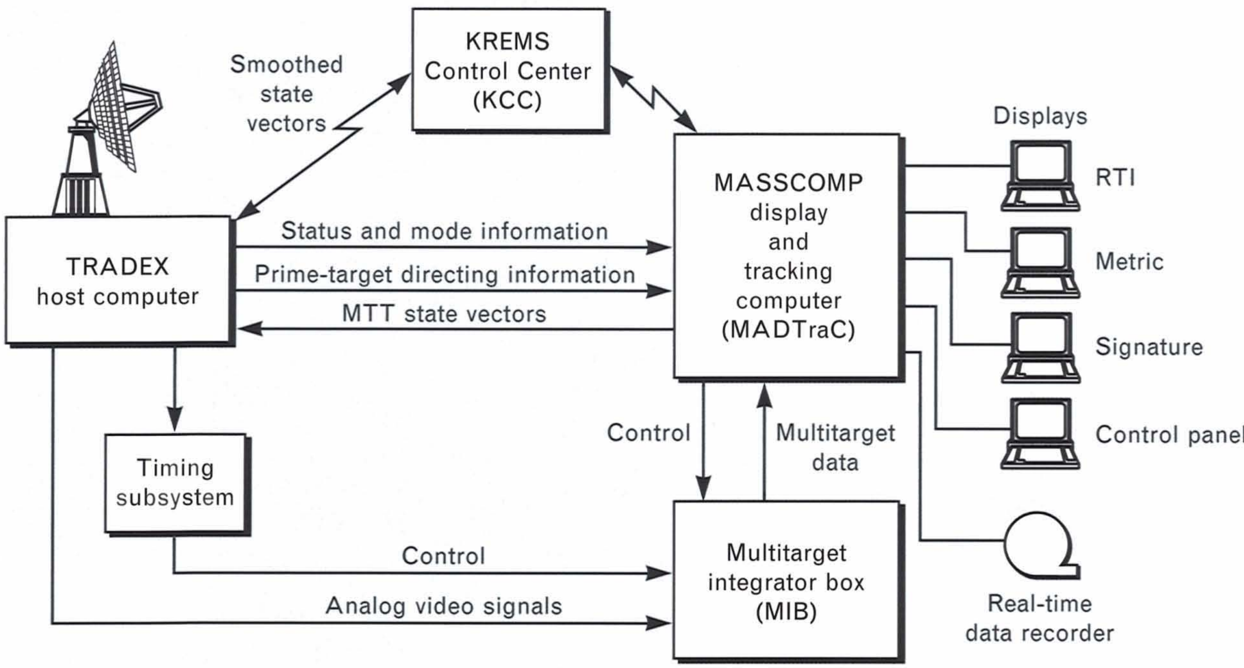

FIGURE 7. TRADEX MTT system diagram. The MTT consists of two hardware systems, the MASSCOMP Display and Tracking Computer (MADTraC) and the Multitarget Integrator Box (MIB). The MADTraC includes four interactive graphics displays and a real-time data recorder.

pulses per 100-msec processing interval. It continuously integrates the incoming current radar data while it simultaneously processes and formats the data from the previous 100-msec interval.

The MIB also processes the data for the RTI display in real time. (The RTI display is discussed in detail later in this article.) The range axis of the RTI display, which has 1152 pixels in the range dimension, shows the integrated video return over range extents from a few meters (one sample) to the full 655 km (128k samples). The MTT metric analyst selects the desired range interval on the RTI display and sends the appropriate commands to the MIB to control the processing. Processing consists of either compressing or expanding the selected interval of range samples into 1152 pixels, where the color of each pixel represents the amplitude of a range subinterval (i.e., 1/1152 of the selected range interval).

The MADTraC Display and Tracking Computer

The MADTraC is a MASSCOMP (currently Concurrent Computer Corporation) model 6600 computer system with three independent high-speed central processor units (CPU) and four independent graphics processor units (GPU). The seven processors all operate independently but reside on a common backplane and share system memory and peripherals. The CPUs are 33-MHz Motorola 68030 microprocessors, each with a Weitek floating-point accelerator; the GPUs are 16-MHz Motorola 68020 microprocessors. The operating system is Real-Time UNIX (RTU), which is a superset of UNIX with additions to support real-time intertask communications and shared memory. RTU provides a more deterministic interrupt response than UNIX.

Two custom high-speed parallel interfaces connect the MADTraC to the MIB and to the TRADEX host computer. These interfaces exchange data and control messages ten times per second during MTT operations. The MIB interface provides up to 63 target reports and preprocessed RTI display data for the MADTraC. The MADTraC returns data to the MIB that controls the processing for target detection and the RTI display data.

The TRADEX computer interface provides the MADTraC with radar system status, radar modes, and a MTT pedestal file consisting of the pointing angles of the antenna and the reference range and time of validity of the MIB sample set. The MADTraC returns status information and six target state vectors that are used by the TRADEX radar to control as many as five auxiliary sample sets. Thus the TRADEX host computer and the MADTraC are effectively connected through a relatively small shared portion of their databases.

The MADTraC also communicates with the KCC at a rate of ten times per second. The MADTraC provides the KCC with as many as ten target state vectors that can be sent to other sensors or used internally by the KCC to aid in target identification. The KCC in turn sends the MADTraC as many as ten target state vectors (which can be other sensors' tracks, nominals, or postprocessed state vectors) that are used by the MTT operators to aid in identification and sampling of targets.

The MTT operator's console consists of four independent graphics displays. Two of the displays, the RTI display and the metric display, show the locations of the targets for the MTT metric analyst. Two other displays, the signature display and the control-panel display, show the radar cross section of each target for the MTT signature analyst. These displays also allow the signature analyst to identify individual targets, to provide TRADEX with state vectors that it uses to control pulse-by-pulse sampling, and to select the MTT-generated target files to be sent to the KCC. Figure 8 shows these four displays in the MTT operator's console; the four displays are discussed in-

FIGURE 8. The MTT operator's console. The range-time-intensity (RTI) display and the metric display, on the right side of the console, show the locations of the targets to the MTT metric analyst. The signature display and the control-panel display, on the left side of the console, show the radar cross section of each target and allow the signature analyst to identify individual targets and provide MTT state vectors for the TRADEX L-band data-recording system.

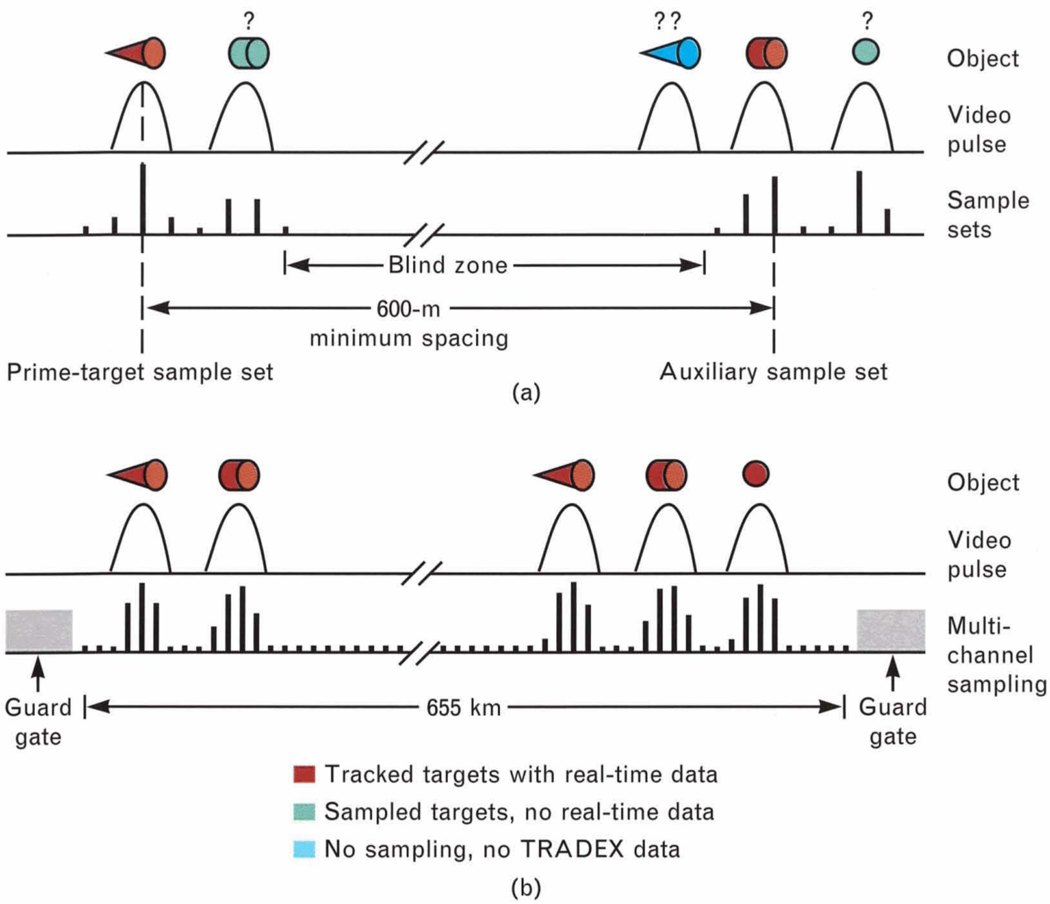

FIGURE 9. Data sampling by the TRADEX radar and the MTT. (a) The TRADEX radar samples only the portion of the range interval around the prime target, as shown by the prime sample set on the left side of the figure. The auxiliary sample set illustrated on the right side of the figure is one of as many as five auxiliary sample sets designated to one of the state vectors provided by the MTT. The blind zone occurs because of timing limitations that prohibit placement of any sample set within approximately 600 m of any other sample set. (b) The MTT samples the entire ambiguous range interval continuously, which allows it to track every object in the beam.

dividually in greater detail later in this article.

TRADEX Data Sampling

The TRADEX radar, when it is in a single-target mode, moves its antenna to keep the target in the center of the radar beam, and controls the sampling of data to cause the return from the target to appear in a particular sample in range. This target, as mentioned earlier, is called the prime target. The TRADEX radar continues to track (i.e., direct the antenna and data sampling to follow) the prime target while the MTT detects and tracks all other targets in the beam. Figure 9 illustrates how the TRADEX radar and the MTT each sample data from the targets in real time.

Figure 9(a) shows the prime target sampled so that its peak return is captured in the third sample of the prime sample set. The relatively small TRADEX sampling pattern is necessary because the TRADEX system can process raw return data from only a limited number of data samples in the immediate vicinity of the prime target in real time. Another target appears in the prime sample set (at the sixth sample) but its data are not available to TRADEX in real time. A second TRADEX sample set is shown at the right in Figure 9(a); this auxiliary sample set is one of up to five additional TRADEX sample sets linked to an MTT state vector. The radar allocates these additional sample sets to specific MTT-tracked objects.

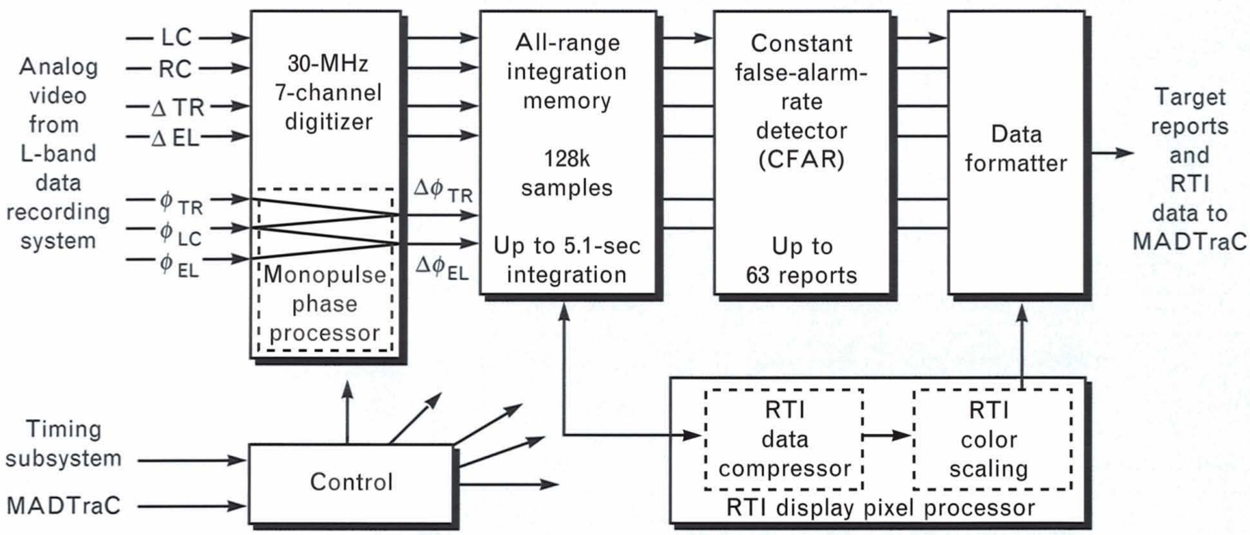

FIGURE 10. MIB function block diagram. The MIB is composed of a seven-channel digitizer, an integrator, a false-alarm detector, a data formatter, a pixel processor, and a controller.

This sampling scheme is designed to use the TRADEX radar's relatively limited data-recording resources most efficiently.

MIB Data Sampling

Figure 9(b) depicts the data sampling performed on multiple objects by the MIB. The MIB sample set is slaved to the TRADEX prime sample set because the location of the prime target is the reference for the MTT. In contrast to the relatively small TRADEX prime sample set, the MIB samples the entire ambiguous range interval continuously. The ambiguous range interval is the time between two adjacent transmit pulses when the radar emits pulses at a rate faster than the travel time of the pulse to the target and back; TRADEX normally operates in this ambiguous-range-interval mode. Receiver isolation switches place guard gates around the transmit pulses, which protects the receiver from saturation or burnout and determines the usable range interval.

The TRADEX system, because of its limited recording throughput, does not sample continuously, but covers multiple smaller range intervals. The MIB sample set covers the entire ambiguous interval, which allows the MTT to collect data on targets not sampled in the TRADEX prime sample set or auxiliary sample sets. The MIB independently samples the prime target and all other targets that appear in the ambiguous interval. The MIB contains its own analog-to-digital converters, and it samples the range interval independently of the TRADEX system, which allows it to operate without impacting the normal sampling and data-recording operations of the TRADEX radar. The timing subsystem provides the MIB with control words and data-sampling strobe signals independently of the TRADEX sampling.

MIB Signal Processing and Detection

The diagram in Figure 10 summarizes the processing sequence in the MIB. The MIB accepts seven analog video signals from the L-band data-recording system, along with control words and data-sampling strobe signals that identify the beginning and end of the processing interval, the beginning and end of the current pulse repetition interval (PRI), and the individual sample times from the timing subsystem.

Analog Signal Processing and MIB Quantization

The seven analog video signals that are output by the TRADEX receivers are first digitized by the MIB. The signals are down converted from the radar's radiated frequency of 1320 MHz through an intermediate frequency at which pulse compression is performed [2]. The TRADEX system employs linear frequency-modulated waveforms that are compressed by using dispersive delay devices such as surface-acoustic-wave filters. The relatively long radiated pulse is swept across a frequency bandwidth of 75 kHz or 20 MHz, with the bandwidth depending on the desired resultant resolution. The surface-acoustic-wave filters delay the received signal by an amount proportional to the instantaneous frequency of the signal, which, in effect, stacks up the return into a relatively narrow time interval and results in a larger amplitude signal occupying a much shorter time. Range resolution is determined by the time width of this compressed pulse.

Following compression, the amplitude pulses are logarithmically scaled, which results in greater signal dynamic range, and then detected. The detected signals are voltage pulses that are proportional to the power received from each target in each of the four amplitude channels: the left circular channel (LC), the right circular channel (RC), the traverse-error channel (ΔTR), and the elevation-error channel (ΔEL). Phase channels associated with the LC, ΔTR, and ΔEL are also sampled by the MIB. These three phase channels—the left circular phase ϕ_{LC}, the traverse-error phase ϕ_{TR}, and the elevation-error phase ϕ_{EL}—are digitally subtracted to produce two voltage pulses proportional to the phase of the return signal relative to the transmitted (reference) signal.

Monopulse Phase Processing

The traverse-error phase and elevation-error phase that are detected and sampled by the MIB are used for the purposes of angle marking in the monopulse system. They can be integrated along with the amplitude channels because they have been subtracted from the reference channel. Except for the contribution of noise and the practical limitations of a real antenna feed, the quantities

$$\Delta\phi_{TR} = \phi_{TR} - \phi_{LC}$$

and

$$\Delta\phi_{EL} = \phi_{EL} - \phi_{LC} \,,$$

which are called the *monopulse phases*, are step functions. In the noise-free ideal case the value of the monopulse phase is $\pi/2$ for negative angle offsets, and the value is $-\pi/2$ for positive angle offsets. Thus for a target that remains at a relatively stationary offset within the beam of the radar, the average phase difference of n such samples is simply the value $\pm\pi/2$. For a target that is directly in the center of the beam, the average of n such samples has an expected value of zero.

Integration

In the next step of MIB processing, the six data channels (LC, RC, ΔTR, ΔEL, $\Delta\phi_{TR}$, and $\Delta\phi_{EL}$) are integrated. During the first PRI of the data frame, the data samples are written into one of the triple-buffered memory banks of the MIB. The data samples from all subsequent PRIs are aligned with the first PRI (so that the return from the prime target falls in the same sample), and then they are added to the previously collected samples and the result is written into the second buffer. The integration continues with the data bouncing between these two buffers for the duration of the data frame.

At the end of the data frame, the results of the integration are stored in one of the three buffers. The MIB then switches to the other two buffers and begins the integration process again for the next data frame. The data in the third buffer (the result of the integration of the previous data frame) are now used by the CFAR detector. This integration of returns occurs at a sample rate of 30 MHz or 2 MHz (depending on the choice of waveform) across the entire ambiguous interval (655 km, or 128k samples) and for the duration of the data frame.

Integration in the MIB is said to be noncoherent because the absolute reference-channel phase information is discarded from each radar pulse during the logarithmic compression and subsequent integration of the target returns. Target detection is based on the sum of the logarithmically compressed pulses. Such integration, while not ideal, is well suited for systems such as the MTT. A simple and straightforward implementation requires only accumulators and memory for each range gate, because integration is performed on the fly.

The noncoherent integrator of the MTT integrates the sampled returns from 128k range gates, and reports the results at a rate of 10 Hz. This noncoherent process is in contrast to coherent integration processors that often require more memory, perform processing in batches, and have a significant computational overhead. The coherent integration processor in the main TRADEX system, for example, is capable of processing no more than 1024 complex data samples (which is a relatively small window size) with less than one update per second.

While ideal coherent integration of radar signals results in an improvement in signal-to-noise ratio proportional to *n* (the number of pulses integrated), noncoherent integration achieves its improvement through the reduction of the variance of the noise, which is reduced by a factor of the square root of *n*. This flattening of the background noise is a perfect match for the CFAR detector implemented in the MIB. Its function is reduced to the simpler task of detecting signals that appear above a relatively flat noise background.

The Constant False-Alarm-Rate Detector

Target detection in the MIB is performed by a weighted, cell-averaging, hardware-implemented constant false-alarm-rate (CFAR) detector [2]. The CFAR detector tests each integrated LC data sample in the MIB buffer (the *test sample*) and compares it to a cell preceding and a cell following the test sample. Up to 16 samples on each side of the test sample are weighted and summed to form the leading-cell average and the trailing-cell average, as shown in Figure 11.

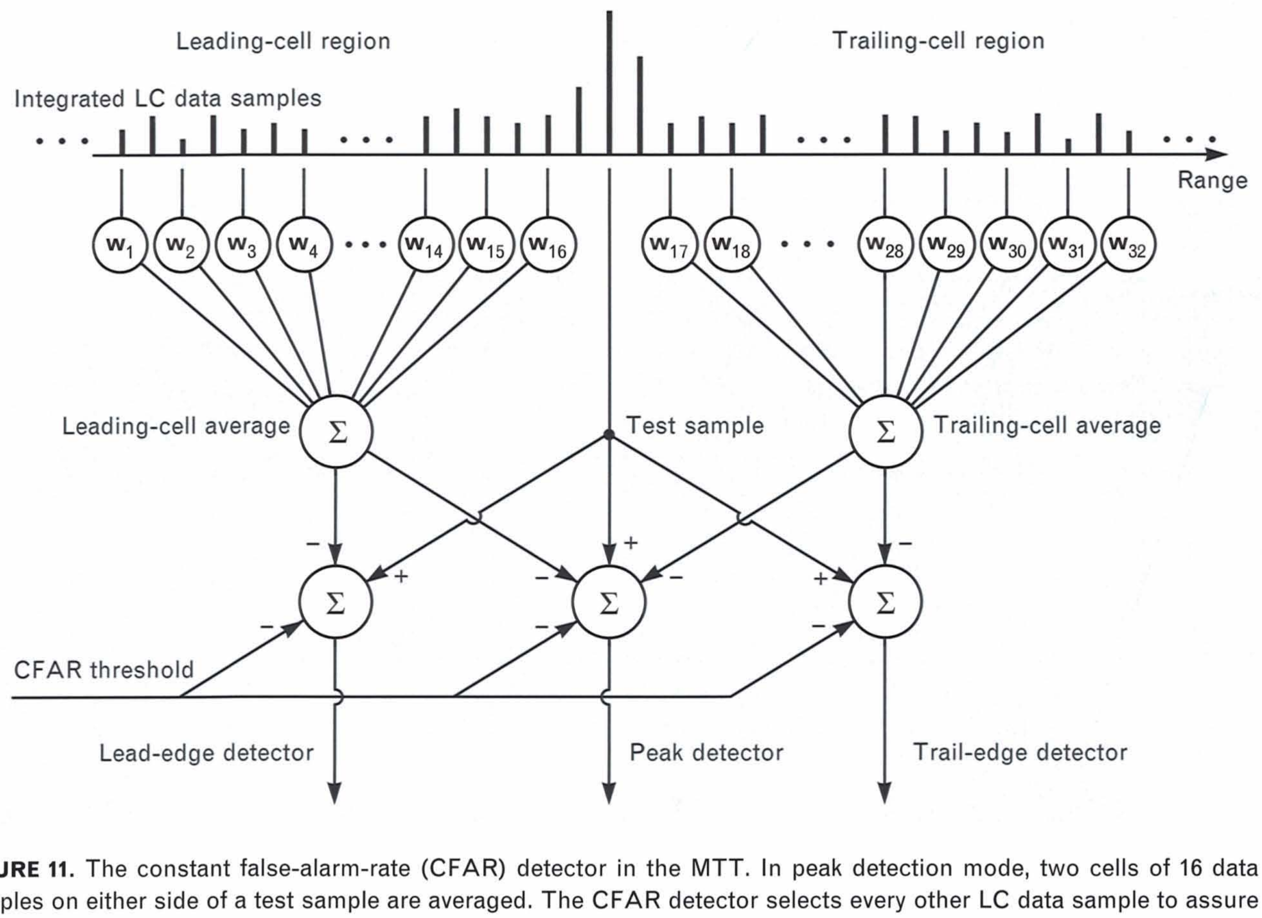

FIGURE 11. The constant false-alarm-rate (CFAR) detector in the MTT. In peak detection mode, two cells of 16 data samples on either side of a test sample are averaged. The CFAR detector selects every other LC data sample to assure that the cell averages are based on statistically independent samples. This process results in an estimate of the noise plus clutter in the vicinity of the test sample. A CFAR detection threshold is added to the average noise level, and then compared to the test sample. If the output of the detector is greater than zero, a detection is declared. Different weights are used for lead-edge detection mode and trail-edge detection mode. The CFAR detector performs this process for every LC data sample in the entire MIB buffer (up to 128k data points) during each data frame (100 msec).

In the *peak detection* mode, both cells are averaged together to arrive at an estimate of the noise plus clutter in the vicinity of the test sample. TRADEX typically operates in a low-clutter environment, so clutter seldom competes with targets for detection. If the return from the test sample exceeds the estimate by a specified offset, and the amplitude of the test sample exceeds the amplitudes of both neighboring samples, a detection is declared. During every data frame, the CFAR detector searches the entire LC data buffer in this manner, and buffers the addresses of the first 63 such detections. The peak detection mode is typically used to detect exoatmospheric targets.

Two other modes are also available for the detection of specific target events. The *lead edge* mode zeros the weights used for the trailing cell, which causes the detector to ignore signals on the trailing-edge side of the test sample when determining the noise level. This mode causes a detection to be made on the leading edge of a distributed target, such as a waking reentry vehicle or a vehicle separation, during the interval of time that the two pieces are unresolved. Similarly, the *trail edge* mode zeros the weights used for the leading cell, which forces the detector to ignore the leading edge of the target return and causes a detection on the trailing edge of the distributed target. The sum of all weights is unity for each of the three detection modes.

RTI Display Pixel Processing

Also operating on the integrated LC data in the third buffer is the RTI pixel processor. The RTI pixel processor sweeps across the entire LC buffer and compresses as many as 128k range samples (up to 655 km in range extent) into the 1152 pixels of the RTI display, where each pixel represents the peak amplitude of a small range interval. The RTI display itself is discussed in greater detail later, but a brief description of display pixel processing is required here.

Figure 12 illustrates the operation of the RTI pixel processor, showing a number of integrated data samples and a few of the 1152 pixels. The metric analyst uses information shown in the RTI display to select the desired range interval for display, and to select the center of the display relative to the prime target. The MIB takes these parameters and computes the number and locations of the range samples that must be considered for each of the 1152 pixels. For each group of samples associated with a pixel, the MIB selects the sample with the highest amplitude as the amplitude to represent the pixel. This *greatest of* selection process is repeated across the range interval, which scales the interval to the RTI display window.

The MIB starts with 8-bit samples but because of the integration, which is an accumulative process, ends up with data that require 16 bits. To minimize the processing required by the MADTraC, the MIB scales these 16-bit quantities after integration by dividing each value by the number of samples integrated. This normalization reduces the data-storage requirement from 16 bits back to 8 bits. These 8-bit quantities are reported directly to the MADTraC, where they are used to generate a new line of the RTI display. The MADTraC then dynamically maps each 8-bit value to one of 32 colors in the RTI display rainbow.

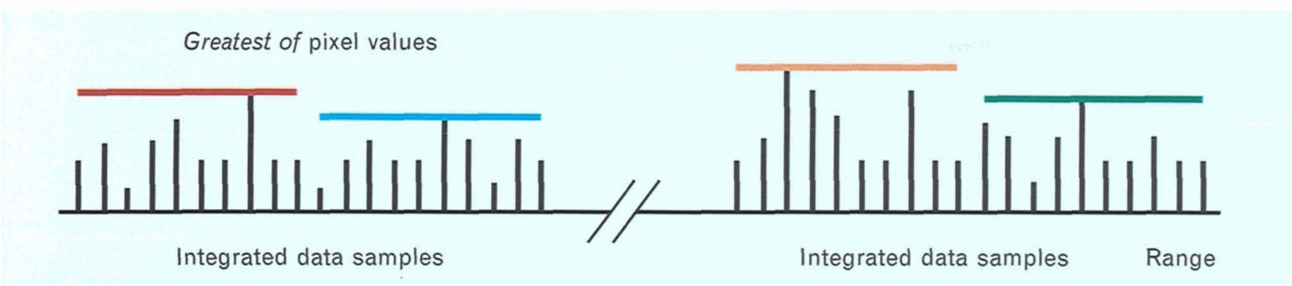

FIGURE 12. RTI display pixel processing. The metric analyst selects a desired range width and center of the RTI display, relative to the prime target. The MIB then compresses the range data samples to the display width by grouping samples and selecting the highest amplitude sample within a group (the *greatest of* sample), and that amplitude is then represented as a colored pixel on the RTI display.

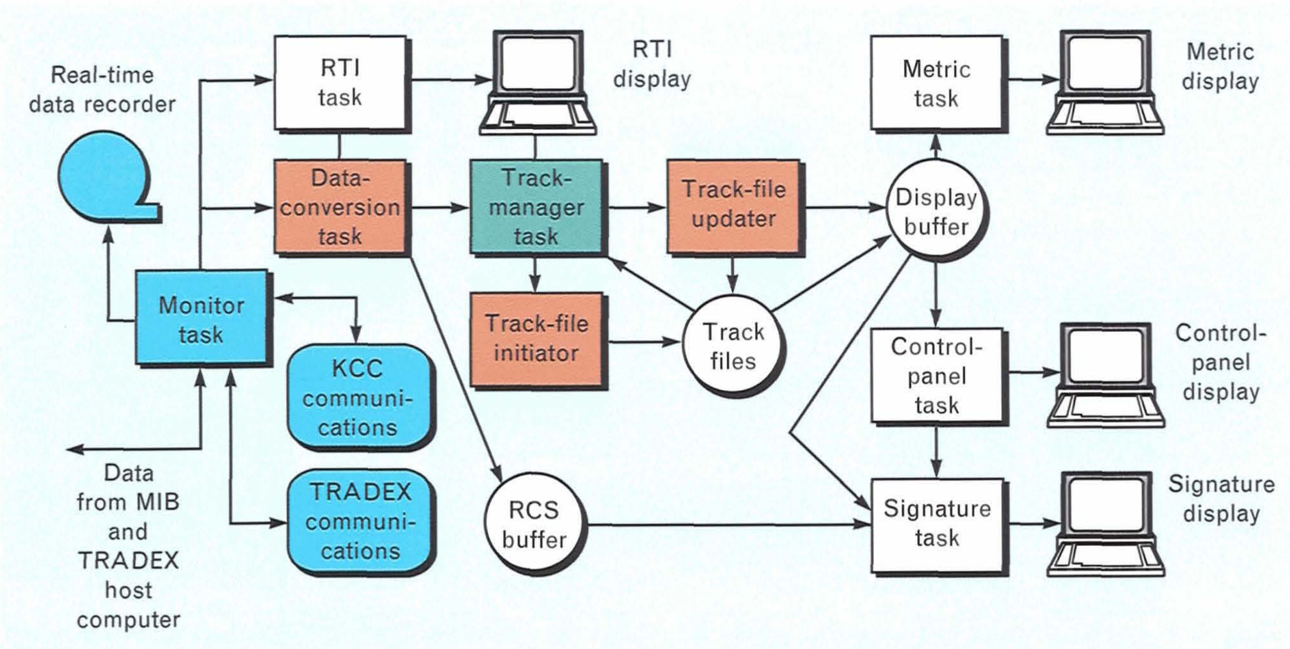

FIGURE 13. The MADTraC Real-Time Program (RTP). The colors indicate the real-time processing tasks.

Data Formatting

Finally, the MIB formats data into a message to be sent to the MADTraC for further processing. The MIB collects its target detections (as many as 63 of them), and assembles the data from all six channels (LC, RC, ΔTR, ΔEL, $\Delta\phi_{TR}$, and $\Delta\phi_{EL}$) for each detection into a target report. For the LC data, two samples—one from either side of the detection (the *early gate* and the *late gate*)—are also included. These two extra samples are used in a range discriminator (discussed below) to obtain a more accurate range mark. The target reports and the RTI display data are packed into a data message and sent to the MADTraC, which updates the RTI display and processes the target reports further to search for trackable targets.

Tracking and Displays

The MADTraC RTP operates in synchronization with both the TRADEX host computer and the MIB. It is synchronized to them by the exchange of data over two real-time, dedicated, high-speed parallel interfaces. The MADTraC RTP is a data-interrupt-driven, multitasking, real-time program. Like the TRADEX host computer, it performs a complete database update at the data-frame rate of 10 Hz. It is essentially a batch processor—it exchanges data messages with both the MIB and the TRADEX host computer at the beginning of its cycle, and then processes the inputs, which results in the update of the track-file database and all display databases.

The MADTraC RTP manages a shared-memory database along with three real-time processing tasks and four near-real-time display tasks that share this database. The executive, or *monitor*, task sets up the shared memory and in turn initiates all the other MTT tasks. It then goes into its real-time mode. The block diagram in Figure 13 illustrates the functionality of the MADTraC RTP.

The monitor task handles all input and output between the MADTraC RTP and the MIB, the TRADEX host, and the KCC. It also handles the real-time data recording of the MTT database for post-mission analysis and playback. Synchronization of the MADTraC RTP with the MIB and the TRADEX host computer is achieved through I/O completions. The MADTraC RTP posts read requests to both the MIB and TRADEX interfaces. Upon successful completion of these reads, the MADTraC RTP begins its processing cycle, which it completes before the next I/O request is satisfied. When a MIB read request is satisfied, the monitor task responds by returning the control message to the MIB. On receiving the TRADEX host message, the monitor task responds by sending a message containing up to six target state vectors that the TRADEX main system uses to collect additional independent datasets on a pulse-by-pulse basis.

On completing each of these writes, the monitor task again posts reads, and then sends an asynchronous system trap (a software intertask interrupt) to the data-conversion task and sets a flag for the RTI display task indicating that new data are available for update of the display. Thus begins another MADTraC-RTP cycle. The monitor task continues, responding to Ethernet messages from the KCC and handling the real-time tape I/O buffering. Subsequent messages from the MIB or TRADEX host computer interrupt the monitor task, which switches its operation to handle the data transfers for the next data frame.

The RTI Display Task

The RTI display task is designed and coded to execute entirely on one of the four independent GPUs in the MADTraC without real-time support from any of the three CPUs. Between display updates the RTI display task spins on a shared memory flag. When the flag is raised, the RTI display task copies the new display dataset from shared memory directly into its display memory buffer. It checks for and acts on operator button pushes, then lowers the memory flag and returns to its spin state, awaiting the next data update.

The Data-Conversion Task

Upon activation (in response to the asynchronous system trap sent by the monitor task), the data-conversion task begins its job of transforming the raw data (such as counts, addresses, and offsets) from the MIB into engineering units (such as meters, microradians, and dBsm). It is constrained at initialization to run on its own CPU, and therefore it begins processing simultaneously with and independently of the monitor task.

Each target report includes a detection range that is relative to the reference range of the MIB sample set. The detection range reported by the MIB for a target report is the number of memory locations the target is from the reference range. The memory locations are separated in time by $1/f_s$, where f_s is the MIB data-sampling rate. This time corresponds to a radar range separation between adjacent memory locations of $c/(2f_s)$, where c is the speed of light. The MIB range of the prime target is generally zero. The data-conversion task calculates the following range for each target report:

$$R_i = R_{\text{MIB}} + \Delta R_i + \Delta r_i \, ,$$

where R_i is the radar apparent-space range of the ith target report, R_{MIB} is the apparent-space range of the MIB's reference sample, ΔR_i is the relative range of the ith target report from the reference sample, and Δr_i is the fine range correction determined by the range discriminator.

The range discriminator uses the preceding (early gate) sample and the trailing (late gate) sample in the target report and a calibration of the compressed pulse shape to estimate the actual location and amplitude of the pulse. The relative range ΔR_i is found by multiplying the number of memory locations the target is from the reference range (reported by the MIB) by the sampling interval $c/(2f_s)$. *Apparent space*, in our context here, refers to observations made by the radar that are uncorrected for pointing angle bias, atmospheric aberrations such as tropospheric refraction, and other known and characterizable discrepancies. *True space* refers to such measurements when all known and characterizable errors have been corrected. Measurements used internally in the radar are left in apparent-space coordinates; these measurements are used to point the radar. Measurements sent to other systems are corrected for known measurement errors and are hence in true space; this correction is done because the radar system receiving the data usually knows nothing about the characterization of the source radar, and simply wants to know where the object is, not where it appears to be.

Figure 14 shows a video pulse as seen by the analog-to-digital quantizers, and it shows the locations of the sample that leads to a detection (the peak sample) along with its early and late gates. The peak sample misses the actual peak by a small amount; as a result its range does not represent the actual range of the target and its amplitude does not represent the actual

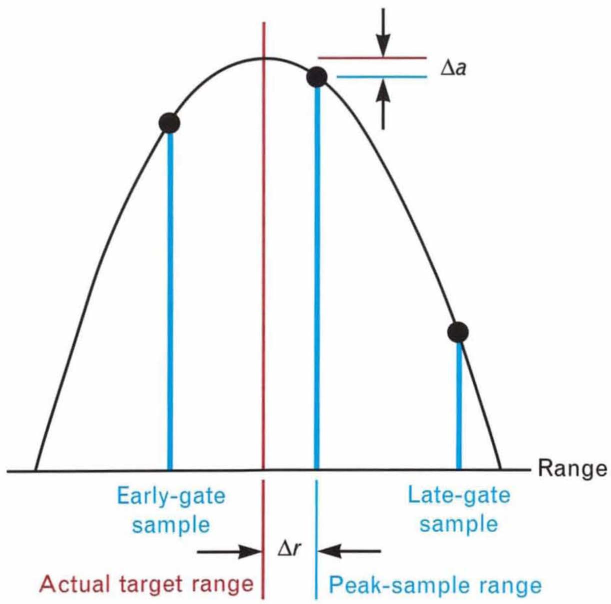

FIGURE 14. Video pulse and MIB data samples for the range discriminator. In general, a constant sampling interval causes the actual target range (the peak of the pulse) to be missed by a small error Δr. An amplitude sampling error Δa is associated with this range sampling error.

amplitude of the target. Because we assume that the target is unresolved, and hence approximates a point scatterer in range, the pulse shape is relatively constant and characterizable. This assumption is reasonable because the best resolution waveform supported by the MTT has a resolution of 15 m. Most targets of interest are much smaller than 15 m in extent.

Given a symmetric pulse shape centered on the peak sample, the early-gate and late-gate samples would have identical amplitudes. As the pulse shifts to either side of the center on the peak sample, the relative amplitudes of the early-gate and late-gate samples change, and the difference between the early-gate and late-gate amplitudes varies monotonically with offset. This variation forms the basis for a fine range correction based on a characterization of the pulse shape. Figure 15(a) shows the range error as a function of the difference in counts between the early gate and the late gate. Figure 15(b) shows the amplitude error as a function of range error as determined by the range discriminator.

After the range data are converted to engineering units, the data-conversion task sends an asynchronous system trap (an intertask interrupt) to the track-manager task. The track-manager task also runs on a dedicated CPU, and hence begins operation immediately and independently of the other two tasks (the monitor task and the data-conversion task). The data-conversion task continues conversion of the remainder of the target-report data.

Target-report angles are calculated by using com-

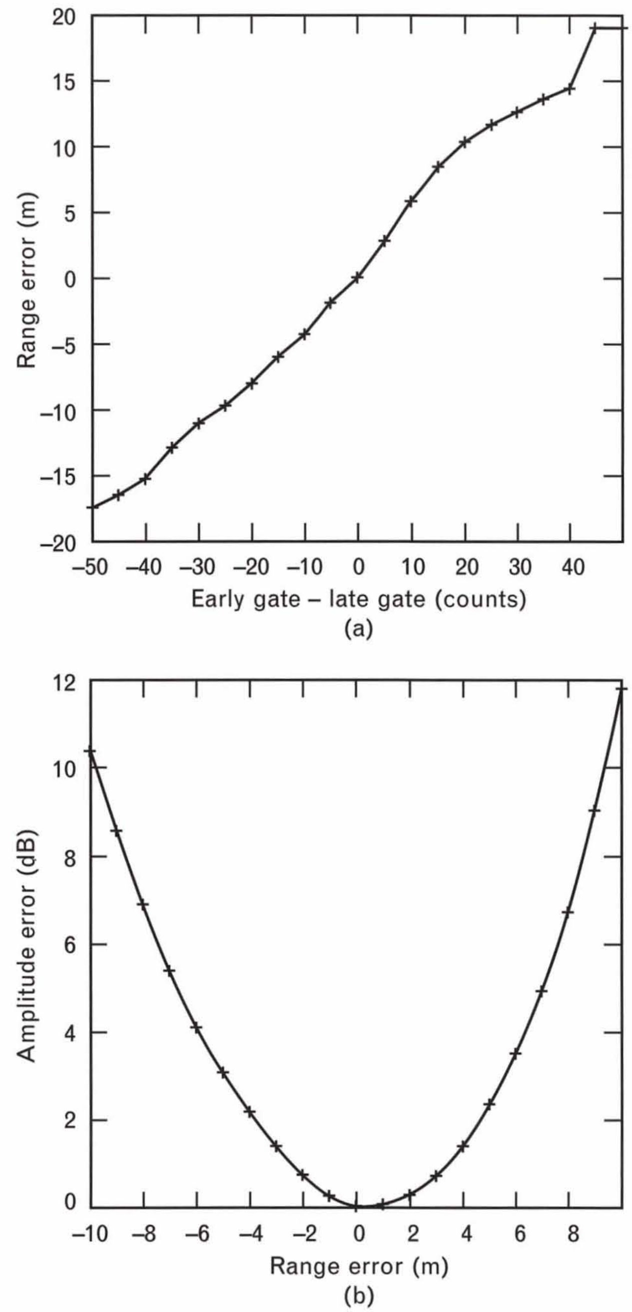

FIGURE 15. (a) Range-discriminator calibration curve. (b) Amplitude-discriminator calibration curve.

puted quantities known as normalized voltage ratios, which are defined as

$$\mathrm{TRVR}_i = \frac{\Delta \mathrm{TR}_i}{\mathrm{LC}_i} \sin(\Delta\phi_{\mathrm{TR}\,i})$$

and

$$\mathrm{ELVR}_i = \frac{\Delta \mathrm{EL}_i}{\mathrm{LC}_i} \sin(\Delta\phi_{\mathrm{EL}\,i}),$$

where TRVR_i and ELVR_i are the normalized traverse voltage ratio and the normalized elevation voltage ratio, respectively, for the *i*th target report. The quantities $\Delta\mathrm{TR}_i$, $\Delta\mathrm{EL}_i$, LC_i, $\Delta\phi_{\mathrm{TR}_i}$, and $\Delta\phi_{\mathrm{EL}_i}$ are the converted MIB values for the *i*th target report.

The normalized voltage ratio is a monotonic function of angle offset across the usable region of the beam of the antenna, and is therefore invertible to obtain angle offsets as functions of measured voltage ratio. These voltage ratios are calculated for each target report, and they are used jointly to evaluate the calibrated angle offset (relative to the center of the beam of the antenna) of each target reported.

The data-conversion task next calculates the calibrated angle offset of each target. These offsets are

$$\mathrm{EL}_i = \mathrm{EL}_{\mathrm{ped}} + f_1(\mathrm{TRVR}_i,\ \mathrm{ELVR}_i)$$

and

$$\mathrm{AZ}_i = \mathrm{AZ}_{\mathrm{ped}} + \frac{\left(f_2(\mathrm{TRVR}_i,\ \mathrm{ELVR}_i)\right)}{\cos(\mathrm{EL}_i)},$$

where EL_i and AZ_i are the calibrated elevation and azimuth of the *i*th target, $\mathrm{EL}_{\mathrm{ped}}$ and $\mathrm{AZ}_{\mathrm{ped}}$ are the elevation and azimuth pointing directions of the pedestal of the antenna, and f_1 and f_2 are the angle-calibration characteristics (described below) that are joint functions of the two normalized voltage ratios. In these equations, azimuth (which is an element of a spherical coordinate system) is distinguished from traverse. Both measure a target motion or offset independent of elevation, but they represent different quantities. Traverse motions or offsets are measured relative to the range vector. Azimuth is the angle of the target vector projected to the zero-elevation plane. The MTT maintains its target state vectors in range-azimuth-elevation coordinates but measures angle offsets in traverse and elevation coordinates.

The next step in the data-conversion process is the calculation of the radar cross section, or RCS, for each target report. The RCS of a target is the radar apparent size of the target, measured in units of meters squared, or dBsm (decibels relative to a square meter). For an unresolved target, such as most TRADEX targets, the RCS can be thought of as a measure of the area of the scattering surface presented to the radar. The RCS varies as the target rotates and precesses, presenting different aspects to the radar. The variations of the RCS over time provide a radar signature for the target; this signature allows the trained RCS analyst to infer the shape and motion of the target. When properly calculated, the RCS of a given target is independent of the range of the target from the radar. The RCS must also be calculated so that it is independent of the location of the target within the beam.

Expressions for the RCS of a target can be found by permuting the radar range equation [2] and accounting for characterizable gains and losses. In the MTT, the LC and RC RCS are calculated as follows:

$$\begin{aligned}\mathrm{RCS}_{\mathrm{LC},i} &= \mathrm{LC}_{i,\mathrm{dB}} + \Delta a_i + \Delta b_i + \mathrm{C}_{\mathrm{LC}} \\ &\quad + \mathrm{P}_{\mathrm{TX}} + \mathrm{L}_{\mathrm{LC}} - 40\log_{10}(R_i)\end{aligned}$$

and

$$\begin{aligned}\mathrm{RCS}_{\mathrm{RC},i} &= \mathrm{RC}_{i,\mathrm{dB}} + \Delta a_i + \Delta b_i + \mathrm{C}_{\mathrm{RC}} \\ &\quad + \mathrm{P}_{\mathrm{TX}} + \mathrm{L}_{\mathrm{RC}} - 40\log_{10}(R_i),\end{aligned}$$

where $\mathrm{LC}_{i,\mathrm{dB}}$ and $\mathrm{RC}_{i,\mathrm{dB}}$ are the calibrated LC and RC amplitudes of the *i*th target report converted to dB, Δa_i is the amplitude correction from the range discriminator, Δb_i is the amplitude correction due to beam loss (because the target is off boresight and hence inefficiently illuminated and non-optimally received), C_{LC} and C_{RC} are the LC and RC calibration constants, P_{TX} is the transmitted power in dB, L_{LC} and L_{RC} are the LC and RC receiver losses (gains minus system losses), and R_i is the range of the target.

This computation of the RCS for each target report completes the conversion of the target-report data from the MIB. The data-conversion task now waits for the completion of the target correlation process of the track manager.

The Track-Manager Task

The track-manager task is responsible for much of the automation of the MTT. The track-manager task takes the set of converted target-report ranges as input, and builds and maintains the database of track files for all of the objects visible in the beam of the radar. This task is responsible for cycle-to-cycle correlation of target reports, automatic track initiation, target-report-to-track-file correlations, and removal of track files from the database when the objects are no longer visible.

To understand the track-manager task, we must look at the target observations made by the MTT. Foremost in the track-manager process is the realization that all targets observable by the MTT as individual objects are range resolved. In other words, the MTT is capable only of handling—as individual targets—objects that are separated in range by at least the resolution of the waveform in use. This situation is in contrast to the coherent processor in the TRADEX main system, which can resolve targets either in range or velocity (Doppler). Targets in the MTT that are not range resolved are handled either as a single target or as two crossing targets (if both targets were previously resolved and in track).

The crossing-target criterion in the MTT is actually more stringent than just stated. The detected video pulses from two closely spaced targets exhibit a mutual interference that causes corruption of both pulse amplitude and phase as well as apparent range as the targets pass by each other. Untreated, this interference often results in an exchange of track files between the two objects, which negatively affects the primary mission task to track and collect data on a particular object. Thus the MTT decides that targets are crossing if the range between them falls below a minimum standoff distance called the *crossing-target neighborhood.*

The MTT also assumes during the track initiation phase that all targets visible in the TRADEX beam have essentially the same velocity and acceleration. On typical multiple-object test flights, all test vehicles originate from the same launch vehicle. To remain in the beam they must follow similar trajectories and they must all have approximately the same velocity. Thus the MTT will not be tasked with detecting and tracking targets with high relative velocities (relative to the prime target) or targets that move rapidly across the beam. Targets in the beam during one data frame are most likely to be there—essentially in the same place—during the next data frame.

The process of target correlation in the MTT consists of either (1) finding range matches between new target reports and existing track files, (2) finding range matches between new target reports and previous target reports, or (3) keeping track of new target reports for consideration in the following data frame. The target reports for any given target pass all of these stages during the acquisition and track initiation phases. By making the above assumptions and by considering only range correlations we greatly simplify the target correlation process for the MTT.

There are two parameters of importance to the track correlation process: the crossing-target neighborhood and the *correlation-gate half width.* Figure 16 illustrates the relationship of these parameters to their targets. The crossing-target neighborhood is a static parameter that depends only on the waveform in use. Track files whose crossing-target neighborhoods overlap are said to be crossing, and are coasted (i.e., not correlated) to prevent track-file exchange between the two targets. The correlation-gate half width, which is a variable determined by the error between the filtered state-vector range and a short history of observed target ranges, can vary in value from a few meters up to the value of the crossing-target neighborhood.

Discussing the track correlator is difficult without also discussing the data structure with which it is implemented. The database for the track correlator is implemented in a doubly linked list, which is a structured data type that has, among its elements, pointers to each of its neighbors. Each node in the list is linked to two others or to one other and the head or the tail of the list. Linked lists are routinely applied if ordering of some data element is desired. For the MTT, target range is the critical data element that must remain ordered.

Each linked-list node holds a collection of data for a given target report or a track file. Among these data items are the last time of validity of the node, the

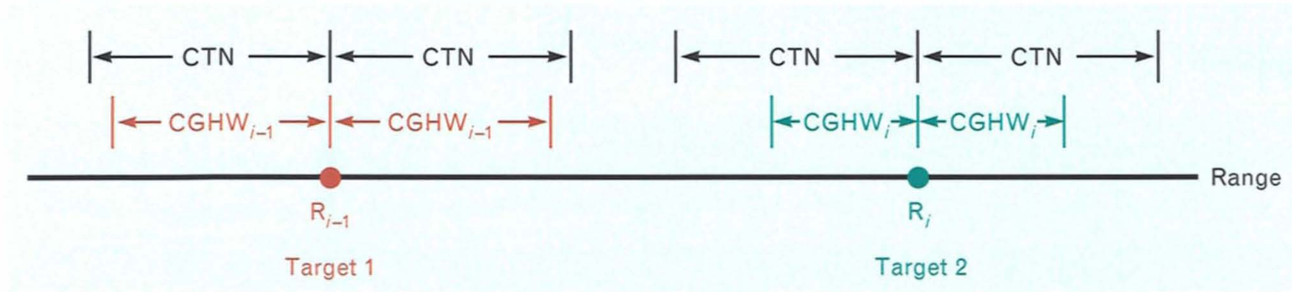

FIGURE 16. Relationship of the correlation-gate half width (CGHW) and crossing-target neighborhood (CTN) for two target positions. The CTN is independent of target position; the CGHW, which is determined by dynamic tracking error, is different for each target.

node range, the correlation-gate half width, the number of hits, the number of misses, the target-file number, the accumulated range error, the accumulated time interval, and pointers to the previous and next nodes in the list. The target linked list is completely updated during each cycle of the MTT RTP, which results in the association of all target reports with a linked-list node and, in cases where it exists, a track file.

Figure 17 illustrates target acquisition. Figure 17(a) shows the target linked list before any target reports are received. At this point in the target-acquisition scenario the list is empty, consisting of only a head and a tail. At the start of the next data frame, three target reports are received, as shown in the vertical range axis of Figure 17(b). Because no target nodes existed before this point, the three target reports generate three target nodes. The hit counter in each node is set equal to one and the miss counter is set equal to zero (1/0). The time of validity of each node is set to the time of validity of the current MIB dataset, and the range of each node is set to the range of the target report. Figure 17(c) shows three target reports received in the next data frame, each of which correlates (or *hits*) a node and is counted. Figure 17(d) shows two target reports received in the next data frame. The first target report hits the first node and is counted. Three hits have occurred within five cycles, so a track file (TF1) is initiated. The second node is missed, because no target report occurs in its range neighborhood; node two does not have enough hits to be declared a target yet. The second target report hits the third node and is counted, which also results in the initiation of a track file (TF2).

Correlation is completed by three passes (or *walks*) through the list. During the first pass, all target nodes are flown forward in time in preparation for correlation with the new target reports. As discussed previously, we assume that all targets seen by the MTT have approximately the same velocity and acceleration as the prime target. This assumption is important because the expected node ranges for the times of validity of the new dataset can be calculated by using the velocity and acceleration of the prime target. This calculation is done with the equation of motion

$$R_{n+1,i} = R_{n,i} + v_{\text{prime}} \cdot \Delta t + \tfrac{1}{2}\, a_{\text{prime}} \cdot \Delta t^2 \,,$$

where $R_{n+1,i}$ and $R_{n,i}$ are the expected range and previous range of the of the ith target node respectively, v_{prime} and a_{prime} are the velocity and acceleration of the prime target, and Δt is the time elapsed between the previous and current time of validity. This equation flies the range of a node forward in time so that it represents an expected range for the current time interval. If a target file exists for a given node, the filtered state-vector range components of velocity and acceleration for that target file are used in place of the corresponding components of the prime target. Thus once an object is acquired, it is correlated and updated based on its own velocity and acceleration rather than those of the prime target.

As the range of each node is flown forward, the node is sorted and relinked in the list (if necessary) to maintain range order. Range order of the target reports and correlator nodes is critical for the optimized execution of the correlator. The reordering is necessary because once a track file is initiated, it flies at its

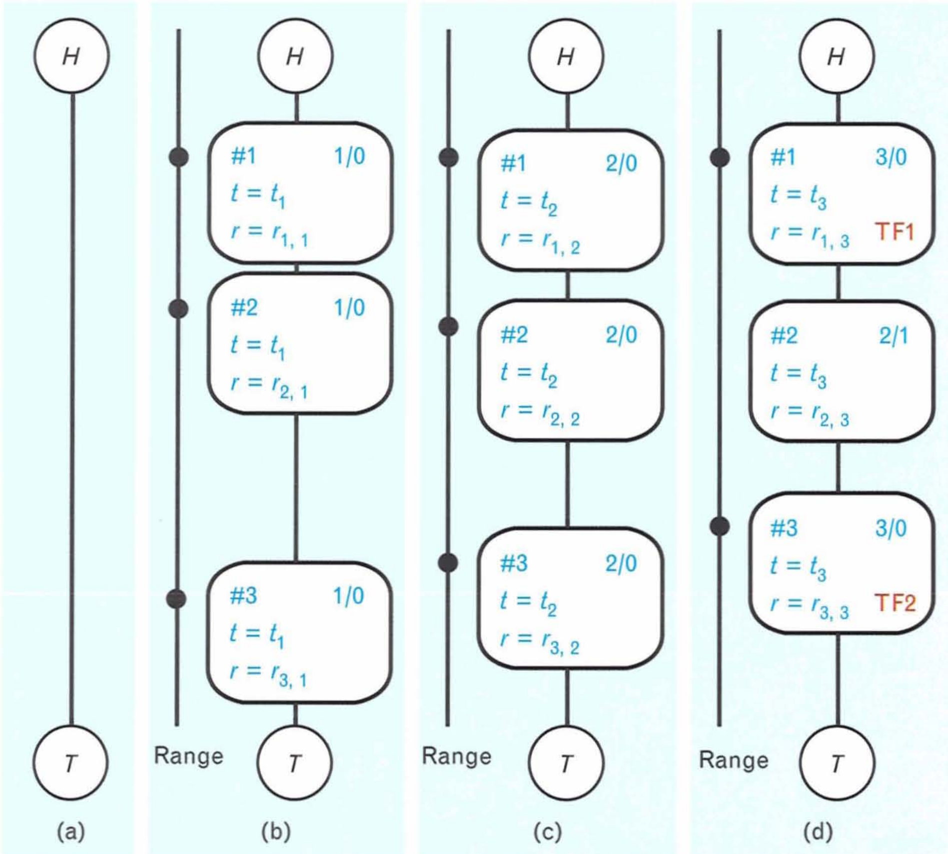

FIGURE 17. The track correlation process showing track-file formation for two targets: (a) In the first data frame, the target linked list is empty before any target reports are received. (b) In the second data frame three target reports are received. (c) In the third data frame three target reports are received, each of which hits a node and is counted. (d) In the fourth data frame two target reports are received. The first target report hits the first node and is counted. Three hits have occurred within five cycles, so a track file (TF1) is initiated. The second node is missed because there is no target report in its range neighborhood. The second target report hits the third node and is counted, and a track file (TF2) is initiated.

own rate and acceleration and can overtake or slip behind another target in range.

During this first pass through the target linked list, nodes whose range neighborhoods overlap are also handled. Overlapping nodes can result in ambiguous correlations (a target report that falls in two nodes). If both nodes have track files associated with them, they are crossing targets and will be marked accordingly during the second pass. If only one of them has a track file, the other is discarded. If neither has a track file, the one with fewer hits is discarded.

All nodes with track files whose crossing-target neighborhoods overlap each other are marked during the second pass through the target linked list. Nodes marked as crossing during the second pass are not correlated during the third pass, because any target reports that could match them are not confidently range resolved. These nodes and their track files are coasted and reexamined for crossing status during the next data frame.

Target correlation is performed during the third pass. The target linked list is walked from head to tail, while the list of range-ordered target reports is considered one report at a time. At this point the program fetches an updated node range, the correlation-gate half width, and the target-report range, and determines one of four cases as depicted in Figure 18. The first case is determined by the condition

$$R_{\text{target report}} > (R_i + \text{CGHW}_i)\,,$$

where R_i is the range and CGHW_i is the correlation-gate half width of the ith node. This condition indicates that no target report correlates to this node, and hence the node has been missed. If n cycles have gone by without m hits, the node is treated as a false detection and is dropped from the list.

The second case, which describes a hit, is determined by the condition

$$(R_i - \text{CGHW}_i) \le R_{\text{target report}} \le (R_i + \text{CGHW}_i)\,.$$

That is, the range of the target report falls within the neighborhood (as described by the correlation-gate half width) of the expected location of a previous target report or track file. When the correlation is made, the number-of-hits parameter is incremented. If a track file already exists, the track-file number is put in the correlation table along with the index of the correlated target report. If no track file exists, a check of the number of hits is made to see if there have been m hits in the last n cycles (where m and n are static parameters—typically three and five, respectively—used to tune the acquisition performance of the system). If m hits out of n cycles is satisfied, a track file is initialized with the range and angles of the target report and the velocities and accelerations of the prime target. The initial velocity given to the track file is actually the sum of the velocity of the prime target and an initial-velocity estimate that is calculated as $\Sigma(\Delta r)/\Sigma(\Delta t)$, where $\Sigma(\Delta r)$ is the accumulated range error from the correlation process and $\Sigma(\Delta t)$ is the elapsed time from node formation to track-file initiation. This initial-velocity correction term gives the velocity of the track file a kick in the right direction.

The third case, which indicates a target report that does not fall within any existing nodes, is satisfied by the condition

$$(R_{i-1} + \text{CGHW}_{i-1}) < R_{\text{target report}} < (R_i - \text{CGHW}_i)\,.$$

In this case a new node is linked into the list before the current node, and the target report is treated as a new potential target.

The fourth case falls under the heading of exception handling. This case, in which the range of the target report falls within the neighborhood of the previously hit node, indicates a multiple correlation. It is satisfied by the condition

$$R_{\text{target report}} \le (R_{i-1} + \text{CGHW}_{i-1})\,.$$

This case is currently handled by taking the target report with the highest signal-to-noise ratio or best range match as the target report to use for updating the track file.

When track correlation is complete, the correlation table contains a list of track file-to-target-report

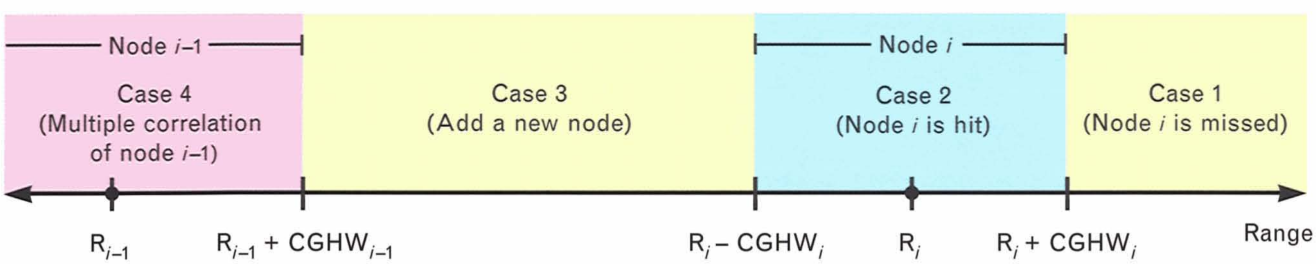

FIGURE 18. The track-file manager determines which of four cases describes the target report currently under consideration. If the range of the target report is beyond (at greater range than) the current node's neighborhood (Case 1), the node is missed (i.e., no other target report will be within its neighborhood). If the target report is within the neighborhood of the node (Case 2), a hit is declared and the target report will be used by the track updater to update the target's state vector. If the target-report range is before (at lesser range than) the current node (Case 3), the node is missed, but the next target report could hit it. A new node is linked in before the current node, and the next target report is processed. If the target report lands within the neighborhood of the previous inrange node (Case 4), a multiple correlation is declared. The multiple correlation could result from a separation event or from a false target report.

associations. The track-manager task raises the spin flag of the *track-file updater* (the lower half of the data-conversion task) to indicate that it can update the track files with the new target-report information. The track-manager task then begins collecting and buffering significant portions of the database for recording to tape.

Track-File Updater

The track-file updater is the lower half of the data-conversion task. It is responsible for filtering and updating all the state vectors of the MTT. A state vector in the MTT is the portion of a track file that represents the location of the target in space. The MTT state vector contains ten elements: position, velocity, acceleration (each with three coordinates—range, azimuth, and elevation), and time of validity. The state vector is a complete description of the location of a target, and it can be flown forward or backward to estimate the location of the target at other times.

The track-file updater implements a polynomial track filter, or *tracker* (also known as a Helms tracker [3], α-β-γ tracking filter, or fading-memory filter) that provides smoothed position, smoothed velocity, and smoothed acceleration, as well as prediction estimates of the same quantities for the next data frame. The development of the tracker in the MTT closely follows the development and notation presented in three references [3, 4, 5]. The MTT employs three independent trackers—one each for range, azimuth, and elevation. Each target in the MTT is treated independently by the tracker. Targets are tracked in radar-apparent, TRADEX-centered, spherical coordinates of range, azimuth, and elevation.

Three sets of weights are implemented in the tracker. The first set, known as the *acquisition weight set*, is designed to minimize lock-on transients by placing greater weight on the initial values of velocity and acceleration. This weighting is done because the TRADEX radar system provides well-filtered (i.e., smooth) and known velocities and accelerations in the state vector of the prime target. Recall that initial estimates of velocity and acceleration for a target report are obtained from the state vector of the prime target. Also, as previously mentioned, an initial range-velocity estimate for the target is available from the track correlator. This acquisition weight set is used for the first k cycles in the track-file update process, where k is a parameter in the configuration file of the RTP used to tune the lock-on performance of the tracker.

The second weight set is designed for optimal track performance as defined in a paper by H.D. Helms [3]. These weights are used after the k cycles mentioned above, and until a third weight set, which is tuned for reentry, is brought into the calculation. The second weight set is derived by estimating the *jerk*, or da/dt (the rate of change of acceleration), of the target, setting an allowable estimation error (which relates to the correlation neighborhood used by the track-file manager), computing the smoothing time constant, and calculating the filter coefficients (weights) by using Helms's method.

The third weight set is determined in the same manner as the second weight set, except the jerk and estimation-error values are modified to include early reentry effects. This modification shortens the smoothing time, which allows the tracker to hold track through the greater dynamic errors associated with reentry. The third weight set is used by the tracker when the prime target reaches an altitude of approximately 100 km.

The weight sets can be represented as follows:

$$\begin{bmatrix} \mathbf{w}_{1,\mathrm{R}} \\ \mathbf{w}_{1,\mathrm{A}} \\ \mathbf{w}_{1,\mathrm{E}} \end{bmatrix}, \quad \begin{bmatrix} \mathbf{w}_{2,\mathrm{R}} \\ \mathbf{w}_{2,\mathrm{A}} \\ \mathbf{w}_{2,\mathrm{E}} \end{bmatrix}, \quad \begin{bmatrix} \mathbf{w}_{3,\mathrm{R}} \\ \mathbf{w}_{3,\mathrm{A}} \\ \mathbf{w}_{3,\mathrm{E}} \end{bmatrix},$$

where

$$\mathbf{w}_{i,\mathrm{X}} = \left[\alpha_{i,\mathrm{X}}, \beta_{i,\mathrm{X}}, \gamma_{i,\mathrm{X}}\right]$$

are the individual weights for each tracker. The number i refers to one of the three weight sets (1 = lock-on, 2 = normal, and 3 = reentry) while X refers to range, azimuth, and elevation.

The tracker provides smoothed estimates of state (position, velocity, and acceleration) as well as a prediction of state to be used for correlation of the next set of observations. The prediction equations are given by

$$\mathbf{x}_p(k) = \mathbf{x}_s(k-1) + \Delta t\, \mathbf{v}_s(k-1) + \frac{\Delta t^2}{2} \mathbf{a}_s(k-1)$$

$$\mathbf{v}_p(k) = \mathbf{v}_s(k-1) + \Delta t\, \mathbf{a}_s(k-1)$$

$$\mathbf{a}_p(k) = \mathbf{a}_s(k-1)\,,$$

where $\mathbf{x}_p(k)$, $\mathbf{v}_p(k)$, and $\mathbf{a}_p(k)$ are the predicted position, velocity, and acceleration vectors for the kth interval, and $\mathbf{x}_s(k-1)$, $\mathbf{v}_s(k-1)$, and $\mathbf{a}_s(k-1)$ are the smoothed position, velocity, and acceleration vectors from the previous observation interval. Each vector consists of three elements—range, azimuth, and elevation.

The smoothed estimates of state for the kth interval are

$$\mathbf{x}_s(k) = \mathbf{x}_p(k) + \alpha_n\left[\mathbf{x}_o(k) - \mathbf{x}_p(k)\right],$$

$$\mathbf{v}_s(k) = \mathbf{v}_p(k) + \frac{\beta_n}{\Delta t}\left[\mathbf{x}_o(k) - \mathbf{x}_p(k)\right],$$

$$\mathbf{a}_s(k) = \mathbf{a}_p(k) + \frac{\gamma_n}{\Delta t^2}\left[\mathbf{x}_o(k) - \mathbf{x}_p(k)\right],$$

where $\mathbf{x}_o(k)$ is the kth position observation (from the converted target report), and Δt is the time interval between observations.

The prediction equations and the equations for the smoothed estimates of state can be combined by substituting the prediction equations into the smoothed estimates of state. We can also reverse the order of calculation so that the most recently filtered velocities and accelerations are used. Thus the update of a single target can be represented by four operations:

$$\Delta\mathbf{x}(k) = \mathbf{x}_o(k) - \mathbf{x}_p(k)\,,$$

$$\mathbf{a}_p(k+1) = \mathbf{a}_p(k) + \frac{\gamma_n}{\Delta t^2}\Delta\mathbf{x}(k)\,,$$

$$\mathbf{v}_p(k+1) = \mathbf{v}_p(k) + \Delta t\,\mathbf{a}_p(k) + \frac{\beta_n + \gamma_n}{\Delta t}\Delta\mathbf{x}(k)\,,$$

$$\mathbf{x}_p(k+1) = \mathbf{x}_p(k) + \Delta t\,\mathbf{v}_p(k) + \frac{\Delta t^2}{2}\mathbf{a}_p(k) + \left(\alpha_n + \beta_n + \frac{\gamma_n}{2}\right)\Delta\mathbf{x}(k)\,.$$

This set of four equations represents the filter-and-update process for each track file. The coasting of an unmatched target (a target file for which no correlated target report exists) is simply a special case of the above update, with the tracking-error term $\Delta\mathbf{x}(k)$ set equal to zero.

Calibration of Angles for Off-Axis Tracking

The TRADEX system was originally designed and built as a single-target boresight tracker, and it was operated as such for over 25 years. When the goal of tracking is to keep the beam centered on a single target, calibration of the angle channels is relatively straightforward. The radar is scanned in azimuth (at zero elevation offset) and in elevation (at zero azimuth offset) across a calibration target, while reference-channel data and angle-channel data are recorded. These are called *principal axis* scans, or *cuts*, and are illustrated in Figure 19.

Voltage ratios are calculated for each angle offset along the individual scans, and the resulting function is inverted to yield calibrated angle offsets as a function of the assumed independent voltage ratios. For angle offsets near the center of the beam of the radar, the traverse and elevation channels behave independently, and because the antenna is actively driven to keep the beam centered on the target, the measured angle offsets are typically small.

The TRADEX MTT endeavors to track all the observable objects in the beam, regardless of their location, to the same level of accuracy as the TRADEX radar when it is boresight-tracking a single target. When the target is off axis, however, (i.e., not near the center of the beam) the angle-channel responses are not independent. An offset in azimuth, for example, results in a change in elevation voltage ratio, which if not calibrated properly can appear as additional elevation offset. This fact clearly necessitates a different approach to calibration of angle offsets; for the MTT the scan for calibration data collection must cover the full angle extent of the beam.

For proper calibration of the monopulse system over the entire beam, we must collect calibration data over the entire region to be calibrated. Thus the principal-axis calibration scans illustrated in Figure 19 are no longer adequate. Because angle calibrations are

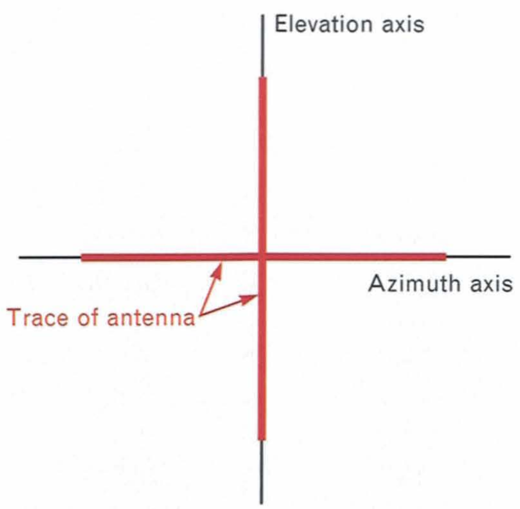

FIGURE 19. Principal-axis calibration scans. These scans are used to calibrate the centermost region of the beam. The TRADEX radar is scanned in azimuth (at zero elevation offset) and in elevation (at zero azimuth offset) to determine the single-target angle-channel calibration data. The principal-axis calibration assumes that the traverse and elevation channels are independent in the small center region of the beam.

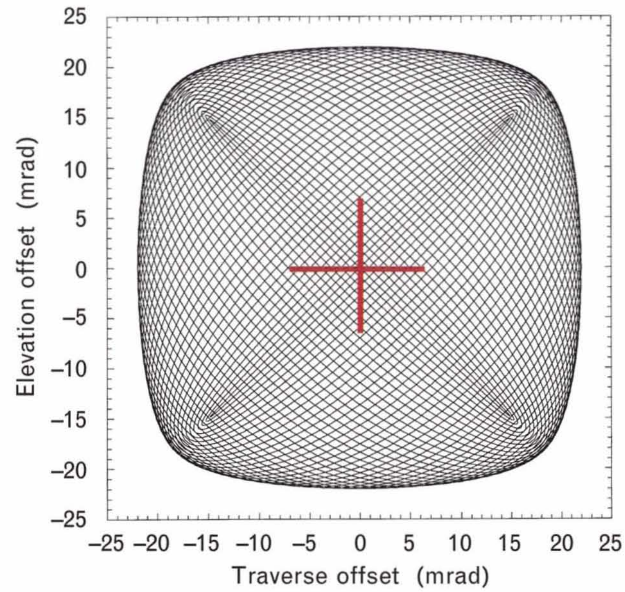

FIGURE 20. Full-angle calibration scans. The antenna is scanned around a Lissajous trace that is designed to cover uniformly the angle extent of the beam to be calibrated. The cross in the center of this trace represents the extent and coverage of the principal-axis scan shown in Figure 19.

performed regularly (typically before every mission), they can contribute significantly to the wear and tear on the antenna system. Angle scans must therefore be performed in a manner that minimizes sharp discontinuous motions while still providing the necessary coverage. Of the many scans considered, the Lissajous trace shown in Figure 20 most effectively satisfies these scan requirements. It is based on simultaneous sinusoidal modulations in both traverse and elevation, so the position, velocity, and acceleration of the antenna are all smooth and continuous throughout the scan. The density of sampling along the scan is controlled by the radar pulse repetition frequency and the antenna scan rate. The density across the scan is determined by the number of scan cuts, which are determined by the sinusoidal frequencies of the Lissajous trace.

The antenna is scanned around the trace shown in Figure 20, with the trace centered on an aluminum calibration sphere. The sphere is tracked by another sensor, typically the MMW radar or the ALCOR radar, which provides a state vector to TRADEX. TRADEX scans the antenna so that its boresight follows the Lissajous trace with the angle offsets as indicated in Figure 20. The TRADEX radar transmits and receives as usual as the calibration scan progresses, and it records reference and angle-error-channel data on a pulse-by-pulse basis. The traverse voltage ratio and the elevation voltage ratio are calculated for each radar pulse, and associated with the antenna angle offset relative to the sphere in much the same way that the principal-axis scans are recorded and processed for a single-target angle calibration. Thus, after preprocessing, we have an ordered quadruple consisting of the traverse voltage ratio, elevation voltage ratio, traverse angle offset, and elevation angle offset. Figures 21 and 22 show the traverse voltage ratio and elevation voltage ratio as functions of these angle offsets. The inverse of the monotonic regions of Figures 21 and 22 are the desired calibration functions.

Processing of the full-angle scan data is significantly different from that performed on the principal-axis cuts. The voltage-ratio/angle-offset double pairs are first gridded on regular angle offsets. Gridding is a process whereby the data samples randomly distributed in both traverse and elevation are systematically adjusted to coincide with desired equidistant angle steps. Gridding as implemented here uses the n nearest neighbors to the desired angle and estimates the expected value at the desired angle. The voltage ratios are then fitted with bicubic splines and the well-behaved monotonic region is determined. Bicubic splines are third-order polynomial splines that are fitted in both dimensions (traverse and elevation) simultaneously [6, 7]. Splines have the most desirable features for this application; the function and its first two derivatives are continuous, and they have compact support, which means they fit data only in their immediate area. These features produce a smooth, regular, and noise-free model of the underlying characteristic of the fitted surface. The characterizable region is that area of angle offsets over which the voltage ratios are monotonic. This region is then inverted, which yields angle-offset pairs as joint functions of smoothed voltage-ratio pairs. The angle offsets are then fit with bicubic splines, which results in a smoothed, calibrated angle-offset surface parameterized by pairs of voltage ratios. Figures 23 and 24 show the calibrated angle offsets as functions of the traverse and elevation voltage ratios. These calibration surfaces are implemented (as f_1 and f_2) in the data-conversion task of the MTT.

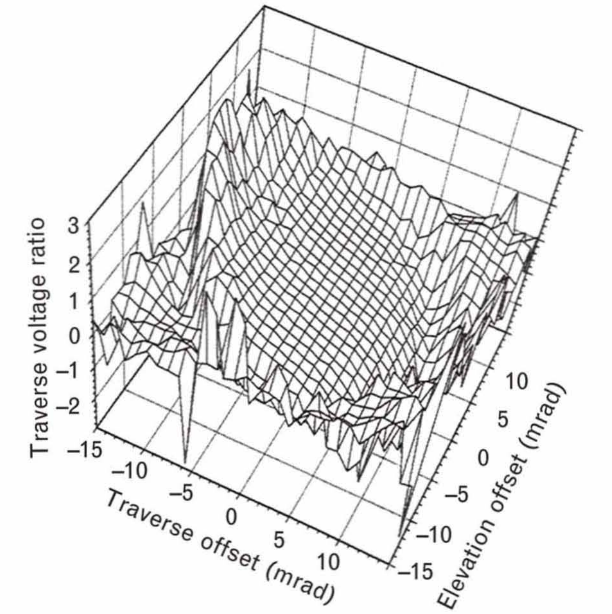

FIGURE 21. Three-dimensional plot of the traverse voltage ratio as a function of angle offset.

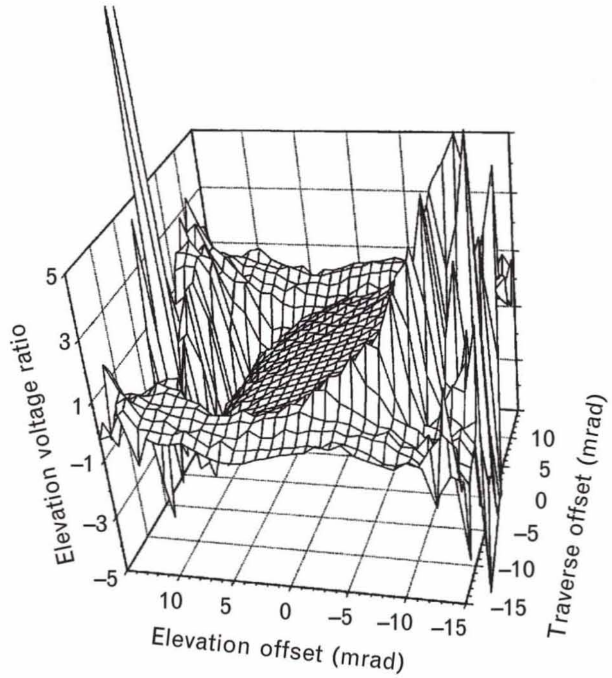

FIGURE 22. Three-dimensional plot of the elevation voltage ratio as a function of angle offset.

MTT Graphical User Interface

The MTT graphical interface consists of four interactive data displays, arranged for two operators, as shown in Figure 8. The displays present target metric and signature data graphically, with limited and carefully chosen alphanumeric data as needed. The operators can control MTT system performance and operation through these graphical interfaces. The operators each use a mouse to interact with display symbols, icons, and function buttons to control MIB processing, the display parameters, and the flow of data with TRADEX and the KCC.

The MTT was designed to be operated by two engineers, the *metric analyst* and the *signature analyst* (TRADEX operations historically have been divided between metric and signature operators). The metric analyst, who is concerned with the actual and relative locations of targets, compares the location of an unidentified target with a predetermined estimate (called a *nominal*) to look for agreement that would form the basis for an identification. With a known object (such as a reentry vehicle with a transponder), the metric analyst looks at predetermined range and angle offsets

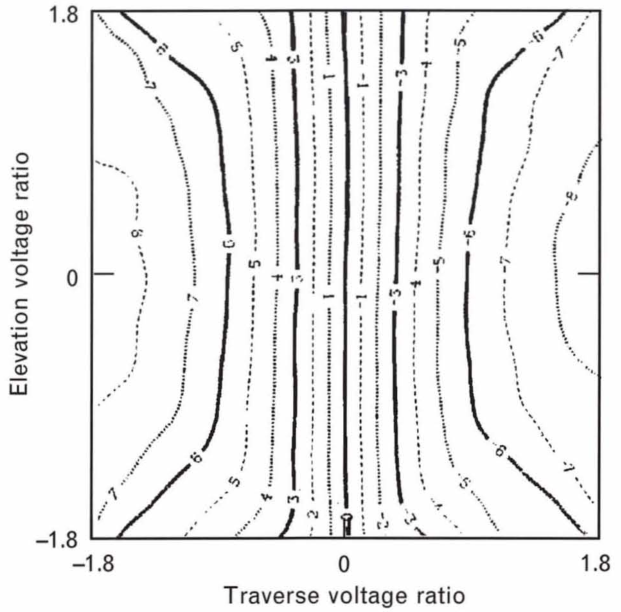

FIGURE 23. Three-dimensional plot of the traverse angle offset as a function of traverse and elevation voltage ratio. This plot is a smoothed inverse of the monotonic regions of Figures 21 and 22.

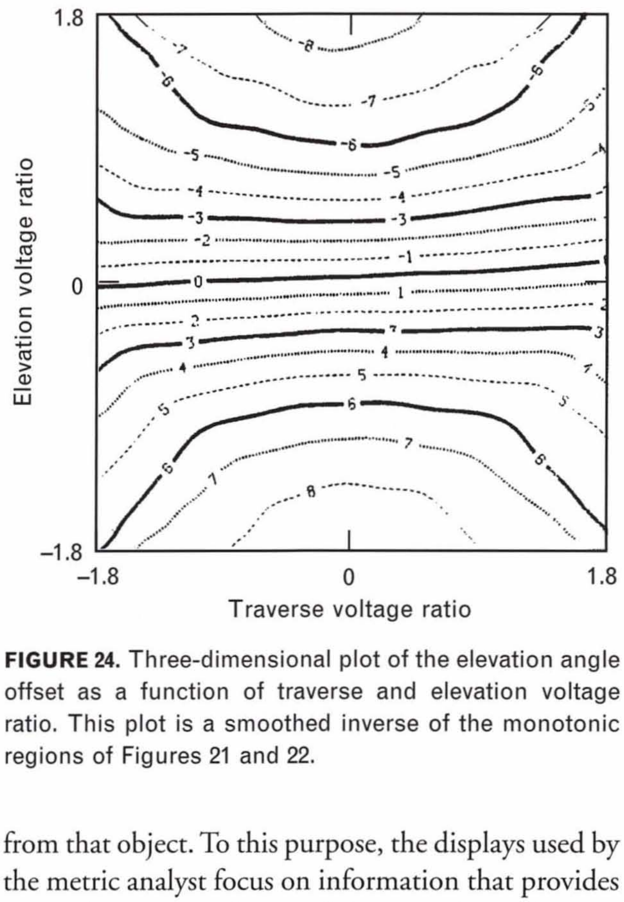

FIGURE 24. Three-dimensional plot of the elevation angle offset as a function of traverse and elevation voltage ratio. This plot is a smoothed inverse of the monotonic regions of Figures 21 and 22.

from that object. To this purpose, the displays used by the metric analyst focus on information that provides a complete metric picture of all targets in track.

The signature analyst, who is concerned with the RCS signature and absolute RCS of the targets, compares the signature of a target to a nominal or expected signature determined from static patterns taken of the vehicle prior to launch. The signature of a target is the time history of its RCS. Reentry vehicles, which are typically spin stabilized, exhibit the spin and associated precession as variations of the RCS. Because of this historical ordering to the identification process (i.e., agree metrically, then check with signature) the signature analyst also posts target identifications for MTT targets.

The Range-Time-Intensity Display

The range-time-intensity (RTI) display shown in Figure 25 gives a time history of the ranges of multiple targets. The horizontal axis depicts range, with *inrange* (closer to the radar) to the left and *outrange* (farther from the radar) to the right. The vertical axis is time, with the most recent time at the bottom of the display and the oldest time at the top. With each cycle, or data frame, a new line of data is inserted at the bottom of the display, which causes the display to scroll up the screen. The target amplitudes are displayed in color.

The word "intensity" in RTI is a holdover from the early days of radar when target amplitudes were displayed by intensity on a monochromatic display. The original name has prevailed, even though RTI displays are usually in color today. Colors on the left side of the rainbow (located on the lower part of the display) indicate low amplitude (weak) scatterers; colors on the right side of the rainbow indicate high amplitude (strong) scatterers. Note that the colors in this display show amplitude and not RCS. The amplitude of the return is directly proportional to received power, which is inversely proportional to the fourth power of range, and which also diminishes as a target moves off boresight in the beam of the antenna. The RCS of a target, as defined and calculated in the TRADEX MTT, is independent of the target's range and location relative to the center of the beam. Even large targets at long range, or on the edge of the beam of the radar, appear as weak scatterers on the RTI display.

The RTI display depicts the all-range (pulse-to-pulse ambiguous interval) integrated return signal or any subset of this interval. Target traces appear on the RTI display at signal amplitudes several dB below the level at which they can be tracked. Hence the RTI display is an invaluable tool for evaluating the target complex at long ranges, where weak signals, and hence poor tracking, are the norm.

The RTI display screen in Figure 25 shows seven targets. The prime target appears in the center of the display. The orange hashbar below the prime target indicates the range extent of the TRADEX prime-target sample set and the location of the tracking gate within the sample set. Two targets are found inrange of the prime target. From the display we can see that these two targets have approximately equal velocities (they are nearly parallel), which are slightly greater than the velocity of the prime target (indicated by their increasing separation from the prime target). The two targets are weak scatterers that exhibit some variation in RCS, which indicates relatively small, complex (irregularly shaped) objects. By consulting the metric display, the analyst can determine that the objects are not subjected to much beam loss.

The two objects immediately outrange of the prime target exhibit the most notable RCS variations. These objects are somewhat larger than the two inrange targets. The object just outrange of the prime target is undergoing a regular rotation or precession. Its velocity is slightly lower than the velocity of the prime target. The two most outrange objects are similar in size and shape to the two inrange objects. They have approximately equal velocities that are somewhat lower than the velocity of the prime target.

The range extent and tracking gate of each of the TRADEX auxiliary sample sets (up to five of them) appear in the area to the right or left of the marker for the prime-target sample set. The range locations for the state vectors received from the KCC are depicted just below this area, along with the ID of the object and the source for the state vector. (No auxiliary sets or KCC files are shown in this figure.)

The metric analyst can cause the RTI display to show any range extent within the ambiguous range interval by clicking the buttons on the lower part of the display or clicking and dragging the mouse through the active mouse region. As the RTI extents are changed, the range axis is updated to show the radar apparent range at the center of the display, and the ranges of the two display edges relative to the center. The analyst can either center the display anywhere within the ambiguous interval or center the display on the prime target by pointing and clicking with a mouse.

The Metric Display

Figure 26 shows the metric display, which is the second graphical display used by the metric analyst. It presents a complete metric picture of the tracked targets in the MTT, and shows the relative ranges and angles of all targets currently in track as well as the relative ranges and angles of targets received from the KCC (the KCC targets are not shown in this figure). Although not all targets visible on the RTI display appear on the metric display, all targets on the metric display also appear on the RTI display.

The lower display window of the metric display shows the ranges of all tracked targets. The lower of the two range axes always depicts the entire ambiguous range interval, with inrange to the left; the left and right edges of the lower range axis are the ranges of the receiver-protection guard gates. Each tracked target is depicted as a vertical stick. Actively updated targets (with correlated target reports) are shown in yellow; coasting targets (with missing target reports) are shown in green. This display indicates a total of six targets in track. These six targets are the prime target and two inrange and three outrange targets, as depicted on the RTI display. The target just outrange of the prime target (as shown on the RTI display) is too close to be reliably tracked independently (the two targets are just within the crossing-target neighborhood of each other at the latest display time).

The operator controls the range over which the CFAR detector operates by dragging two sets of indicators (shown as yellow and blue triangles) across the lower range axis to enable and disable a desired range of detection. The operator can thus interactively select two range regions over which detection is performed and can use these controls to exclude visible clutter regions or extended targets (e.g., chaff clouds). Returns from these excluded regions still appear on the RTI display, however, which is independent of

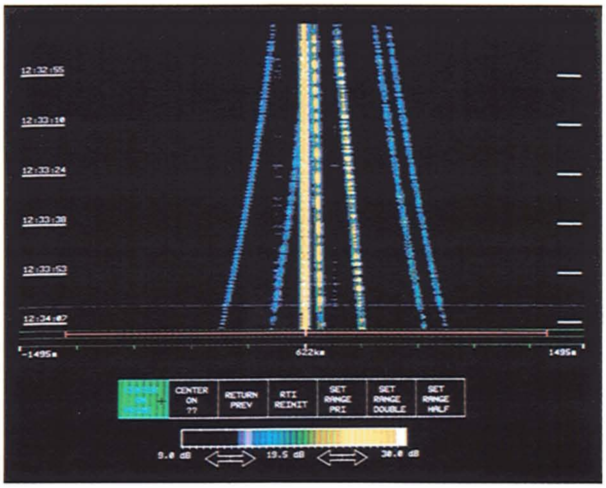

FIGURE 25. The MTT RTI display.

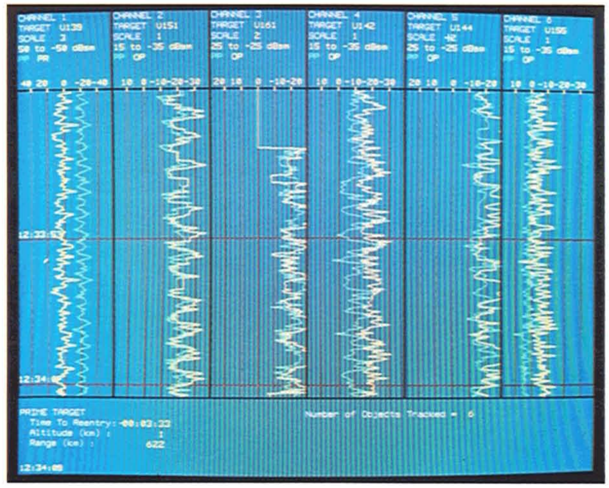

FIGURE 27. The MTT signature display.

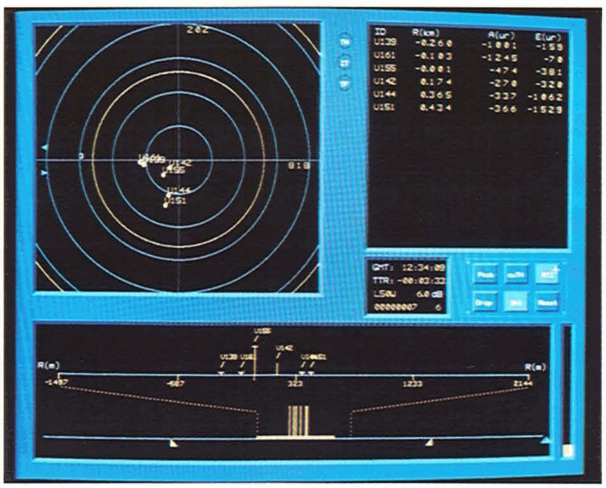

FIGURE 26. The MTT metric display.

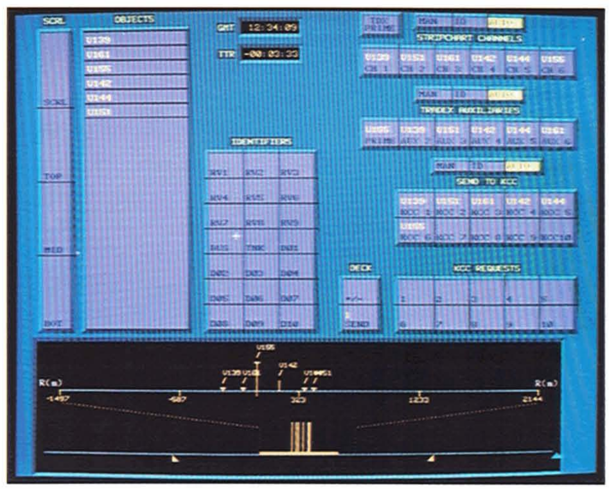

FIGURE 28. The MTT control-panel display.

tracked targets. The orange bar on the lower range axis of the metric display provides the operator with a graphical indication of the selected range coverage of the RTI display.

The operator of the metric display can zoom in on portions of the ambiguous interval by clicking and dragging the dashed vertical lines. The zoomed extent is shown on the upper range axis, which has range labels relative to the prime target. Targets again appear as vertical sticks, color coded as before, and with a height proportional to their current RCS. Large RCS targets are tall sticks; small RCS targets are short sticks. The ID of the target, if it has been identified, is displayed above the target. Targets that are unidentified show their acquisition number as Uxxx, where xxx is a serial number. In this example, U155 is the prime target, as indicated by the short orange tick mark below it.

The upper left window in Figure 26 is the boresight display. This window depicts the angles of tracked targets relative to the center of the beam. The prime target normally appears in or near the center of this display. The boresight depicts the locations of targets as seen through the beam of the radar, as if the radar were a telescope or gun sight. Clicking the small blue triangles at the left in this window control the displayed beam angle. The inner and outer yellow circles represent the TRADEX S-band and L-band 6-dB two-way beamwidths, respectively. The boresight display shows only those targets selected to appear on the zoomed range axis at the bottom. All of the targets in this example lie near the boresight.

The window in the upper right in Figure 26 contains a table of the relative ranges and angles of all targets in the boresight display (relative to the prime target), along with the ID of each target. The operator uses this table when quantitative results (such as the precise location of an object relative to the prime target) are needed. The small buttons in the right center of the metric display allow the operator to select the CFAR detection mode (peak, lead, or trail) and set the zoomed range axis to display the range extent shown in the RTI display. The SET and RESET buttons allow the operator to store and recall display settings. The elevator bar on the lower right allows the operator to set the CFAR detection threshold manually. The numerical value of the CFAR threshold (6.0 dB) is shown in the small window to the left of the push buttons, along with the waveform identifier (in this case L50W) and the system status code (00000007). The status code indicates seven target reports were found by the MIB for this cycle and no errors (shown by the zeros) were detected. The "6" in this window indicates that six track files are being maintained by the MTT. The GMT indicates the current time, and the TTR indicates the time to reentry for the prime target. Indicator lights (TN, IT, and TF) in the center of the display warn the operator of MIB overload conditions.

The Signature Display

The signature display shown in Figure 27 is one of two displays used by the signature analyst. It depicts a time history of calibrated dual-polarization RCS values for six different targets. The signature analyst uses this time history to estimate the shape and dynamics of the target, and then compares the displayed signatures with predetermined patterns based on known target shapes and nominal deployments. When a target vehicle is poorly deployed, the radar signature often provides the only clues to the cause and nature of the deployment problem.

The signature display plots up to thirty seconds of buffered RCS data for each tracked object. The red horizontal graticule lines in the display match the time marks on the RTI display. The zero-value data at the top (oldest times) of channel 3 indicate this target was not in track during those times. The most recent time is at the bottom of the display, and the display is updated every few seconds, which produces a plot of the RCS up to the latest time buffered. Both received polarizations of the reference channel (LC and RC) are plotted simultaneously. The signature analyst can select to display LC (principal polarization, or PP) and polarization ratio (PR) data by clicking either the PP label or the orthogonal-polarization (OP) label in the box above each stripchart channel (which changes the labels in the display to PP and PR), as has been done on channel 1. Clicking on the PP label or the PR label then switches that channel back to PP data and OP data. The polarization ratio is defined at TRADEX as the ratio of the LC RCS to the RC RCS. With adequate signal-to-noise ratio, the PR data more readily convey information pertaining to the motion of the target. This information can be misleading, however, if either the LC channel or the RC channel exhibits low signal-to-noise ratio.

The signature display can show RCS at three different scalings. By clicking one of the three mouse buttons on the SCALE label above each channel, the analyst can select one of the three scalings for that channel. This ability to select appropriate scaling allows the analyst to view the RCS of a target at a scale commensurate with its extent.

The ID of the target displayed in each channel is shown in the box at the top of the channel; these IDs are posted by the signature analyst on the control-panel display. Additional information pertaining to the TRADEX radar and MTT system operation appears in the box at the bottom of the signature display. The TTR (time to reentry) is coordinated with other MTT and TRADEX system displays, which allows the radar operators to keep pace with events occurring during the flight. The ALT (altitude of the prime target) and RNG (range of the prime target) are also used for this purpose.

The Control-Panel Display

The control-panel display shown in Figure 28 provides the signature analyst with the functions necessary to post IDs on particular targets. The analyst also controls which targets are shown on the signature display, along with the flow of track-file information to the TRADEX main system and to the KCC. Most of these features are automated, but they also have manual overrides.

As targets are automatically acquired, they appear in the OBJECTS list, labeled with their acquisition number. This acquisition number, or U-number, serves as the identifier for the target on the three other displays until an official identifier mnemonic is assigned. A mission-configurable list of IDENTIFIERS appears in the middle button panel. The signature analyst selects an OBJECTS button and an IDENTIFIERS button in either order to post the ID. The ID code is written into the track file, and it thereafter appears as the ID for the target on the other displays. At this point, the OBJECTS list displays the ID mnemonic rather than the specific U-number. (Note that none of the objects in Figure 28 have been assigned an ID.)

Control of the contents of the signature display is accomplished by click-selecting the object (from either the OBJECTS button panel or the IDENTIFIERS button panel) and then clicking on one of the STRIPCHART CHANNELS buttons. The ID is moved into the button, and the contents of the appropriate stripchart channel are redrawn to show the appropriate target.

The analyst repeats these actions to send the state vector of a target file to the TRADEX main system for sample set control, except that the operator clicks on one of the five TRADEX AUXILIARIES buttons rather than one of the STRIPCHART CHANNELS buttons. This action again deposits the ID in the button, and results in the appropriate state-vector information being copied to the TRADEX output message. Track files to be sent to the KCC are similarly selected by clicking on one of the KCC FILES buttons.

The lower window in the control-panel display is an exact duplicate of the range window from the metric display. This window is provided for the signature analyst, who must maintain a conversation with the metric operator regarding targets and IDs. The window also provides the signature analyst with visual feedback of the target ID as it is posted to the appropriate object.

Summary

The MTT provides TRADEX and KREMS with valuable new automatic capabilities. It allows the TRADEX system to acquire and track all targets observable in the L-band radar beam. Targets tracked by the MTT are available to the TRADEX main system for control of auxiliary sample sets, which greatly improves the pulse-by-pulse data-recording capabilities of the radar. The MTT targets are also available to the KCC for use as designation sources for other sensors, and to provide additional targets for track-to-track correlation. Because of these capabilities, the MTT is an important data source for the ongoing automation of KCC functions.

Real-time interactive graphical data displays provide the two MTT analysts with an easily assimilated view of the mission, as seen through the eye of the radar. The metric display, with its boresight window and range plots, provides this point of view most effectively. The signature display provides the analysts with data that allow them to infer the shape and dynamics of any of the targets in track. The analysts control MIB processing and target-data flow by interacting with the graphical displays in a natural and intuitive manner.

Acknowledgments

The author thanks the leaders of the Field Systems group, the Radar Measurements division, and TRADEX and KREMS for their support, encouragement, and confidence throughout the development of the MTT system. The author also extends appreciation to Fred Rose and the late Jay Walton, both of Lincoln Laboratory, and George Maas of General Electric, who freely shared their wisdom and advice early in the development of the system.

The TRADEX MTT is the result of the efforts and contributions of many talented individuals over a period of four years. The author thanks these individuals—too numerous to mention—who offered thoughtful suggestions and constructive criticisms, and who shared their invaluable KREMS mission experience and insight. Special appreciation goes to the people who made this development effort happen. Mike Shields of Lincoln Laboratory and Hal

Pschunder of General Electric contributed much to the conceptualization, design, and specification of the MTT. Their efforts helped lay the foundation that made the MTT the system it is today. Thanks are also due to Daniel LeVine of Loral Training and Technical Services, who developed the mathematical methods for the full-angle calibration models.

The MIB was designed and built by General Electric (formerly RCA) in Moorestown, New Jersey. A number of talented and dedicated people contributed significantly to its completion. The author thanks Carl Specht, William Mays, Mike Hammond, Steve Fink, and Paul O'Neill, who formed the core of the team that successfully built and delivered the finished product.

The MTT would not work without software. The author thanks the software team for their diligent efforts, their creative spirit, and most of all their patience. David Schmechel of Lincoln Laboratory designed and wrote the MTT RTP executive and the testbed/playback software, and was responsible for the integration and delivery of all the software. Dick Holmes of General Electric designed and wrote the data-conversion, track-update, and track-manager tasks. Shaun Woodburn of General Electric completed the RTI display task and designed and wrote the real-time recording routines, as well as an invaluable database access package. He followed the MTT system to TRADEX, where he continues to support it. Margery Harrison of General Electric designed and wrote the signature-display and control-panel-display tasks, and contributed to the ergonomic design of the displays. Janice Rejto of General Electric designed and wrote the metric-display task and also contributed to the ergometric design of the displays. She also followed the MTT system to TRADEX, where she is now responsible for the ongoing development and refinement of the MTT system software. Adam Kearns of Lincoln Laboratory wrote the Quick-Look graphical playback program, repeatedly updated the database description, and contributed in many other ways too numerous to list here.

Finally, thanks go to Steve Blackmon of Lincoln Laboratory, who provided some of the most important support of all—keeping the computer system alive, maintained, backed up, and available to the software developers through several operating systems and hardware upgrades.

REFERENCES

1. K.R. Roth, M.E. Austin, D.J. Frediani, G.H. Knittel, and A.V. Mrstik, "The Kiernan Reentry Measurements System on Kwajalein Atoll," *Linc. Lab. J.* **2**, 247 (1989).
2. M. Skolnik, *Introduction to Radar Systems* (McGraw-Hill, New York, 1962).
3. H.D. Helms, "Maximally Reliable Exponential Prediction Equations for Data-Rate-Limited Tracking Servomechanisms," *The Bell Sys. Tech. J.* **44**, 2337 (1965).
4. S.S. Blackman, *Multiple Target Tracking with Radar Applications* (Artech House, Norwood, MA, 1986), chap. 2.2.
5. P.R. Kalata, "The Tracking Index: A Generalized Parameter for α-β and α-β-γ Target Trackers," *IEEE Trans. Aerosp. Electron. Syst.* **20**, 174 (March 1984).
6. L.L. Schumaker, *Spline Functions: Basic Theory* (Wiley, New York, 1981), chap. 12.
7. P.M. Prenter, *Splines and Variational Methods* (Wiley, New York, 1975), p. 131.

GLENN W. MEURER, JR.
is a staff member in the Field Systems group. His research specialty is in radar systems development. He received a Bachelor of Music degree with a specialty in music education, and B.S. and M.S. degrees in electrical engineering from the University of Massachusetts at Amherst. Before coming to Lincoln Laboratory in 1987, Glenn was employed at the Microwave Remote Sensing Laboratory (MIRSL) at the University of Massachusetts at Amherst, where he helped develop the Stepped-Frequency Delta-*k* Radar for remote sensing of ocean surface currents. He is a member of Tau Beta Pi, Eta Kappa Nu, and the IEEE.