

HUSIR Signal Processing

James V. Eshbaugh, Robert L. Morrison, Jr., E. Weber Hoen,

Timothy C. Hiatt, and Gerald R. Benitz

Existing imaging radars lack the high-resolution imaging capability needed to adequately characterize the recent spate of small satellites. The size of a typical microsatellite (tens of centimeters) is on the order of a single pixel at X band (10 GHz). Radar image resolution is inversely proportional to the bandwidth of the radar, driving radars to larger bandwidths to achieve the required resolution and, consequently, to higher frequencies to access appropriate spectrum regions. This increase in complexity has cascading effects on the hardware (generating the appropriate frequencies and bandwidths) and software of the radar system. The real-time software must process nearly an order of magnitude more data with every pulse. The image processing algorithms that have performed well for many years fail to generate high-quality images, requiring a ground-up review and rewrite of the code.



Satellite capabilities are increasing, while the size of satellites is decreasing. Owning and operating very capable satellites are now within the reach of all but the

most impoverished nations, and many commercial entities are launching very capable small satellites, with data available to anyone at very modest costs. The imaging radars that are part of the U.S. Space Surveillance Network are used to characterize satellites on orbit, and as the sizes of the satellites go down, the effectiveness of that characterization decreases because of the pixel size of the radar images. The only way to improve image resolution is to increase bandwidth, and the only way to obtain that bandwidth is to increase the center frequency of the radar's antenna. The Haystack Ultrawideband Satellite Imaging Radar (HUSIR) program was initiated to build a radar that would have the imaging resolution required to characterize small satellites.

The HUSIR program began with a number of looming technical challenges that had to be overcome. In addition to the large, highly accurate, agile antenna, gyrotron tube transmitter, and sensitive cryogenic receiver, the signal processing required to turn received target echoes into fully focused high-resolution radar imagery was a challenge in its own right. Drawing on Lincoln Laboratory's experience with the Radar Open Systems Architecture (ROSA), the HUSIR signal processing equipment is ROSA-based with modifications for use at W-band frequencies (92–100 GHz, or millimeter-wave, 3 mm wavelength).

In any radar system, the signal processing workload can be broken into two parts: real-time processing and post-processing (i.e., any processing done after the tar-

get echoes are collected). Real-time processing must be performed on a pulse-by-pulse basis to successfully track the target and collect and record radar target echoes. Traditionally, post-processing was performed well after the data collection, but that is no longer strictly necessary; radar images can be generated in near real time and still meet this definition for post-processing. The HUSIR signal processing chain's two logical parts, described in detail later in this article, rely heavily on mathematical modeling techniques.

Signal Processing

A mathematical model for the radar response to a point target must take into account all the possible variables along the radar's path. The model is used to understand the effects on the imaging process of uncompensated target motion and of nonideal waveform response (resulting from the combined path through the transmitter, atmosphere, and receiver). Such understanding is necessary to motivate the approaches to maintaining target tracking, range alignment, Doppler correction, transverse equalization correction, and autofocus.

Radar Data Collection Overview

The primary goal of any imaging radar is to keep the radar beam on the target so that good images can be generated during post-processing. The real-time processing must account for the theoretical distortions (e.g., target motion and time-bandwidth effects) in the target echoes. The HUSIR system is also strongly affected by atmospheric variations leading to mislocation of the target.

Imaging radars such as the Haystack Auxiliary radar (HAX), HUSIR, and the Millimeter-Wave radar (MMW) on Kwajalein Atoll, Marshall Islands, employ a linear frequency-modulated (LFM) waveform and stretch processing to allow low-rate sampling of the demodulated signal [1]. The stretch method mixes the received signal with a delayed copy of the transmitted signal, called the deramp waveform, to effectively perform a range-to-frequency conversion over a small swath (in time and distance) containing the target. For example, HUSIR sweeps from 92 GHz to 100 GHz in pulse widths varying from 51.2 μ s to 819.2 μ s, providing a range swath of 7 m to 120 m, depending on waveform.

In addition to keeping the radar beam on the target, the system needs to keep the range window centered

on the target and provide fully coherent output data. HUSIR's real-time tracker provides the required pointing and ranging information. The master timing subsystem uses the tracker output to generate the transmit and receive waveforms and uses related timing information to generate the coherent output.

Each output pulse is coherent in the sense that it is range and phase aligned to a hypothetical point target provided by the tracker. To achieve this alignment, the receive (deramp) waveform is adjusted to exactly match time and frequency of the expected backscatter, and the sampling of the demodulated signal is exactly centered at the round-trip time of the target. Following the pulse-compression fast Fourier transform (FFT), the hypothetical target appears in the center of the range window and has constant measured phase. If the hypothetical target happens to match the centroid of the actual target, the output pulses can immediately be used in Doppler processing and imaging. In practice, however, further post-processing is required to align the pulses for imaging.

Figure 1 presents an illustration of stretch processing for a point target where the deramp waveform is not exactly centered at the round-trip time of the target. Mixing the target echo with the deramp waveform results in a tone at which the frequency is proportional to the time delay. This tone indicates the range (delay)-to-frequency conversion by which scatterers at different ranges are resolved (i.e., the mechanism providing range resolution in inverse synthetic aperture radar [ISAR] images). The compressed pulse, illustrated in the rightmost portion of the figure, is generated by sampling the deramped signal and applying an FFT. The target appears offset in range because of the timing mismatch between the deramp waveform and the target echo. A phase shift on the target is proportional to the delay. The phase shift is due to the carrier frequency term in the transmitted waveform.

For radar imaging, two assumptions are implicit in the stretch processing scheme shown in Figure 1. First, the deramp waveform is assumed to match the centroid of the actual target. In practice, the target centroid is not known exactly, resulting in range walk and phase errors from pulse to pulse. These effects must be corrected prior to Doppler processing and imaging by using motion compensation. Second, the deramp waveform is assumed to be a matched replica of the transmitted chirp, i.e., an ideal LFM waveform with no distortions in magnitude




FIGURE 1. Illustration of stretch processing. When the round-trip time of the deramp waveform is not matched to the target centroid, the target appears offset in range (and is phase shifted) in the compressed pulse. Motion compensation is required to correct the range and phase shifts on each pulse prior to image formation.

and phase. This assumption does not hold well in practice, particularly for W band, requiring equalization of the pulses to ensure that the correct matched filtering operation is applied at the radar receiver. The equalization, which is derived from both calibration data and the target signature itself, can be applied as a signal multiplication in time prior to the FFT. The following sections present a mathematical description of motion compensation and equalization, indicating higher-order signal terms that can be ignored for lower-frequency radars but must be accounted for at W band.

Analysis of the Point-Target Response

If the response signal from a theoretical point target is not perfectly matched in phase to the real-time tracker data, corrections must be made. Pulses are assumed to be compensated as discussed above. They are recorded as the raw in-phase/quadrature (IQ) data in the real-time data recording (RTDR) file. The following analysis derives the response of a point target at range r and range rate \dot{r} as observed in the radar output. Definitions of the parameters are given in Table 1.

We begin with a transmitted waveform γ , centered at $t = 0$:

$$\gamma(t) = \exp\{j2\pi f_c t + j2\pi \mu t^2\}, |t| \leq \frac{T}{2}. \quad (1)$$

The chirp slope is denoted by μ , where $\mu = B/T$. Note that Equation (1) employs complex notation ($j = \text{sqrt}(-1)$) representing the “positive” frequency component. The received waveform from an ideal point target at range r with range rate \dot{r} is given by

$$s(t) = \gamma(\beta[t - \Delta]),$$

where $\Delta = 2r/c$ and $\beta = (1 - 2\dot{r}/c)$. Note that Δ is the time delay from the center of the transmit pulse to the center of

Table 1. Parameter Definitions

PARAMETER	DESCRIPTION
c	Speed of light, m/s
f_c	Center frequency, Hz
T	Pulse width, s
B	Bandwidth, Hz
r_0	Range at pulse center, m
\dot{r}_0	Range rate, m/s (>0 is outbound)
r	True range, m
\dot{r}	True range rate, m/s

the received pulse, and β is the time dilation factor caused by target motion. Acceleration effects are negligible, even for HUSIR's millimeter-wave signals.

The deramp waveform generated in the radar is effectively

$$z(t) = \gamma(\beta_0 [t - \Delta_0])^*$$

where $\Delta_0 = 2r_0/c$, $\beta_0 = (1 - 2\dot{r}_0/c)$, and * indicates complex conjugation. The demodulated output $d(\tau)$ can be defined as the product of $s(t)$ and $z(t)$, sampled at time values τ centered on Δ_0 , $\tau = t - \Delta_0, |\tau| \leq \frac{T}{2}$:

$$d(\tau) = s(\tau + \Delta_0) \cdot z(\tau + \Delta_0). \quad (2)$$

The factors in Equation (2) are expanded as follows:

$$\begin{aligned} s(\tau + \Delta_0) &= \exp\{j2\pi f_c \beta (\tau + \Delta_0 - \Delta) + j\pi\mu\beta^2 (\tau + \Delta_0 - \Delta)^2\} \\ z(\tau + \Delta_0) &= \exp\{-j2\pi f_c \beta_0 \tau - j\pi\mu\beta_0^2 \tau^2\}. \end{aligned}$$

Thus, the argument (phase "angle") of $d(\tau)$ is

$$\begin{aligned} \angle d(\tau) &= 2\pi f_c (\beta - \beta_0) \tau + 2\pi f_c \beta (\Delta_0 - \Delta) + \pi\mu(\beta^2 - \beta_0^2) \tau^2 \\ &\quad + \pi\mu\beta^2 [2(\Delta_0 - \Delta)\tau + (\Delta_0 - \Delta)^2]. \end{aligned} \quad (3)$$

The significance of each term in Equation (3) is described in Table 2. Recognize that linear terms in τ correspond to

range offset, quadratic terms cause defocus, and τ -independent terms are range-dependent phase offsets. Some of the terms introduce small perturbations that can be ignored as long as the tracker range, Δ_0 , doesn't exhibit large jitter, e.g., regularly causing the target to hop more than a meter within the range window.

Motion Compensation

The goal of motion compensation is to adjust the range and phase of each pulse to recenter it according to an improved model for the target motion. For the following calculations, the improved target motion model is assumed to be given. The ideal compensation is to apply the conjugate of all the terms in Table 2 with the deskew applied in the compressed domain to correct scatterers at all ranges. A more computationally efficient method is to apply only the first, second, and fourth terms, and to use the $\beta = 1$ approximation for the second and fourth.

The recommended motion compensation for HUSIR is

$$m(\tau) = \exp\left\{j\frac{4\pi}{c}(\mu\tau + f_c)(r - r_0) + j\frac{4\pi}{c}f_c\tau(\dot{r} - \dot{r}_0)\right\}. \quad (4)$$

The compensated pulse is simply $m(\tau) \cdot d(\tau)$. Observe that the first term in the exponent of Equation (4) is basic range alignment, and the second term is the new Doppler compensation required for HUSIR (applicable to MMW and HAX as well). The computational cost to include Doppler compensation is minor.

Table 2. Output Pulse Phase Contributions

PHASE CONTRIBUTION	SIGNIFICANCE
$2\pi f_c (\beta - \beta_0) \tau$	Doppler-induced range offset: Although negligible for other satellite imaging radars, this must be corrected for HUSIR because it can cause jitter on the order of a range bin, e.g., a 2 m/s tracker error causes a 1-bin range error for HUSIR.
$2\pi f_c \beta (\Delta_0 - \Delta)$	Target mid-band phase: This is always adjusted in motion compensation. The significance of $\beta = 1$ is sufficient as long as the tracker range jitter is small.
$\pi\mu(\beta^2 - \beta_0^2)\tau^2$	Chirp slope correction: This is small and can be ignored in HUSIR if the tracker velocity jitter is under ~2 m/s.
$2\pi\mu\beta^2 (\Delta_0 - \Delta)\tau$	Target range: This is the range-to-frequency conversion of stretch processing and is adjusted in motion compensation. As with the mid-band phase, using $\beta = 1$ is sufficient as long as the tracker range jitter is small.
$\pi\mu\beta^2 (\Delta_0 - \Delta)^2$	Skew phase: This results from the offset arrival time of the pulse. In synthetic aperture radar (SAR), it is always removed as part of the deskew process. (In deskew, the conjugate of this phase is applied to the compressed pulse.)

Because τ is not intuitive, it is useful to make the following transformation of variables. Define $f_{\text{RF}} = \mu\tau + f_c$, the instantaneous radio frequency (RF) of the samples, and $f_{\text{BB}} = f_{\text{RF}} - f_c$, the baseband instantaneous frequency. Then the motion compensation can be written as

$$m(f_{\text{RF}}, f_{\text{BB}}) = \exp\left\{j\frac{4\pi f_{\text{RF}}}{c}(r - r_0) + j\frac{4\pi f_{\text{BB}}}{c}\frac{f_c}{\mu}(\dot{r} - \dot{r}_0)\right\}. \quad (5)$$

From this form, it is easy to see how the range alignment is a true time-delay adjustment, while the Doppler compensation is a baseband alignment only, i.e., there is no mid-band phase offset. Interpreting τ as an instantaneous frequency is exactly true if deskewing has been performed because deramp followed by deskew effectively implements a Fourier transform, known as the chirp-z transform. In other words, stretch processing followed by deskewing transforms the data to the frequency domain.

Recall that the compensation in Equation (5) is an approximation. It is easy to verify that its error is minimal, e.g., only a few degrees of phase, by multiplying Equations (3) and (4) and evaluating the residuals.

Waveform Correction

To this point, the derivation has assumed an ideal open-loop transfer function describing the combined response of the transmitter chain, atmosphere, and receiver chain. In practice, however, the transmitter and receiver hardware and the atmosphere modify the magnitude and phase response of the waveform prior to deramp processing. To model this effect, we represent the open-loop response as $h(t)$. The modified waveform can then be expressed as

$$\tilde{s}(t) = s(t) * h(t),$$

where $*$ denotes convolution. Note that $s(t)$ can be expressed to first-order expansion in terms of its instantaneous frequency locally about some time $t = t_0$:

$$s(t) \approx s(t_0) \exp\{j2\pi f_1(t_0)(t - t_0)\} \text{ for } t_0 - \varepsilon \leq t \leq t_0 + \varepsilon,$$

where $f_1(t) = f_c + \mu t$ is the instantaneous frequency at time t , and ε is an arbitrary small value. By using the property that complex sinusoids are eigenfunctions of linear time-invariant systems, we have

$$\tilde{s}(t) \approx H(f_1(t_0))s(t_0) \exp\{j2\pi f_1(t_0)(t - t_0)\}$$

for $t_0 - \varepsilon \leq t \leq t_0 + \varepsilon$, where $H(f)$ is the Fourier transform of $h(t)$. Replacing the approximation for $s(t)$ with the original function yields

$$\tilde{s}(t) \approx H(f_c + \mu t) \cdot s(t).$$

Thus, we see that there is a linear mapping between the open-loop frequency response and each time sample of the waveform. A sample at time t is modified by a multiplicative gain $|H(f_c + \mu(t - \Delta_0))|$ and phase shift $\angle H(f_c + \mu(t - \Delta_0))$. In other words, the system response imparts an amplitude and phase modulation on the received waveform.

Waveform correction, or transverse equalization (TE), is the process of estimating the response $H(f_c + \mu t)$ and applying a corresponding magnitude and phase correction. This is usually accomplished by collecting data on a calibration sphere, which can be approximated as an ideal point target. The waveform correction

$$H_c(t) = \frac{H^*(f_c + \mu t)}{|H(f_c + \mu t)|^2}$$

is applied to the received signal to generate the corrected output

$$d(\tau) = \tilde{s}(\tau + \Delta_0) \cdot H_c(\tau + \Delta_0) \cdot z(\tau + \Delta_0).$$

Note that the previous expression is a mathematical representation only. Deramp processing is performed on the analog signal, while transverse equalization is applied after the signal has been digitized.

Waveform correction as described above assumes a point target model in which the radar return is exactly centered at the round-trip time of the pulse. For a distributed target with scattering centers at multiple ranges, the received signal $g(\tau)$ is a superposition of returns with different delays

$$g(\tau) \approx \sum_k \sigma_k \tilde{s}(\tau + \Delta_0 - \Delta_k),$$

where σ_k and Δ_k are the reflectivity and delay (relative to the “point source” round-trip time) for the k th scatterer, respectively. Note that the waveform correction $H_c(\tau + \Delta_0)$ is only valid for $\Delta_0 = 0$. Thus, applying the waveform correction to $g(\tau)$ is technically not correct. However, this approximation is reasonable when the delays are much smaller than the pulse width. For shorter pulses that have a large chirp slope, deskew cor-

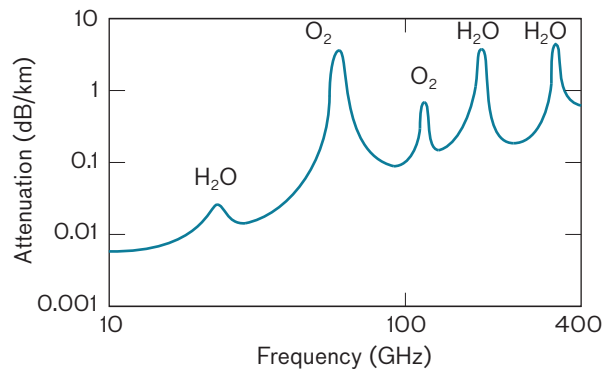


FIGURE 2. The oxygen and water-vapor absorption spectrum consists of a series of peaks and valleys. The peaks are due to molecular resonances that result in high absorption. W-band radar is in the valley between the two oxygen peaks. During humid days, the valley will fill in, raising the attenuation for the radar.

rection should be applied prior to transverse equalization. Deskew effectively centers the returns with different delays to the round-trip time, allowing a direct mapping between corrections in time and frequencies in the transmitted pulse. We shall denote this form of waveform correction as transmit TE correction. Figure 1 illustrates the transmit TE correction in the context of stretch processing for a hypothetical scatterer.

Atmospheric Challenges to W-Band Radar

The high operating frequency for HUSIR, while instrumental for increased resolution, does pose serious challenges because of the propagation medium. First, the atmosphere will attenuate the W-band radar signal many times more than for X- or Ku-band signals. This attenuation exists even though the W band (92–100 GHz) sits within a trough of the atmospheric oxygen absorption spectrum (see Figure 2). The level of the trough, however, is higher than that of the Ku- or X-band region and, more importantly, will “fill in” with the addition of water vapor. In fact, the attenuation at W band is predominantly due to water vapor, although water in liquid or solid form and, of course, oxygen gas will also attenuate the radar pulse.

There will be random changes in permittivity, the second atmospheric effect, along the radar line of sight because of turbulence of atmospheric gases (again dominated by water vapor and oxygen). These random fluctuations lead to changes in path length, i.e., the effective length of the microwave path through the atmosphere,

which manifests itself in phase variations. The path-length variations should be the same in X and Ku bands, but because of the smaller W-band wavelengths, the phase variations are much greater than those of X and Ku bands. These phase variations can have a host of effects, including a loss of antenna caused by spatial variations across the aperture, anomalous refraction (pointing errors) caused by a systematic phase tilt across the aperture, and lastly, a per-pulse phase variation from the tracking of the radar beam through a spatially and temporally variable medium.

During the construction of HUSIR, we investigated the various effects the atmosphere can have on the W-band beam in order to determine their severity and to develop mitigation strategies.

Attenuation

Models for atmospheric absorption—at W band, absorption primarily results from oxygen and water vapor—exist in the literature and have been experimentally validated [2]. Calculations based on measurements of ground-based meteorological conditions (pressure, temperature, and humidity) at Bedford, Massachusetts, combined with an assumed atmospheric scale height (~1.6 km) and a radar elevation angle (25°), produce values for the amount of “clear-sky” two-way loss throughout a typical year at W band, as shown in Figure 3.




FIGURE 3. The variation of the expected attenuation at 25° elevation from month to month has been calculated from ground-based meteorological data. As expected, during the hot, humid summer months, the attenuation seen by the radar will be at its highest. Month-to-month attenuation was calculated from local weather conditions.

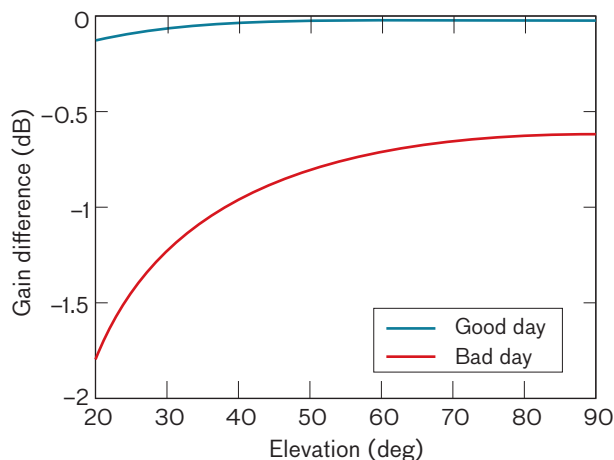


FIGURE 4. The antenna gain difference resulting from phase fluctuations is simulated over the aperture.

Obviously, summer months will have much higher amounts of atmospheric loss than winter months. Because of the high gain of the antenna, however, this amount of loss can be (and has already shown to be) overcome.

Clearly, clouds, rain, and snow will add to the amount of atmospheric attenuation. These conditions are more difficult to model, needing more input parameters than can be easily found [3]. However, the expectation that HUSIR will be operational even in moderate rain has been borne out by data collections during the integration and test phase.

Phase Fluctuations

LOSS OF GAIN

At a single point in time, spatial phase variations over the aperture will cause a drop in antenna gain. These variations are analogous to atmospheric seeing that limits the usable aperture size in optics. This aperture size, which is quantified by the Fried parameter r_0 (also called the coherence length), can be calculated for W-band radar as well and is found to range from ~60 m on very turbulent, wet days to > 400 m on calm, dry days for 25° elevation.

The difference between “good” and “bad” days for these simulations is quantified by C_n , the index of refraction structure constant. This parameter describes the “rockiness” of the turbulent atmosphere; it is obviously not a constant in time, nor is it a constant in space, especially vertically. However, like other researchers [4–6], we assume that C_n does not vary spatially. For the good days, we assume $C_n = 0.8 \times 10^{-7}$ and for bad days, 3×10^{-7} ,




FIGURE 5. The simulated angle-of-arrival (AoA) variations on both good and bad days are well below the HUSIR half-bandwidth.

which may be considered an upper limit for Massachusetts [7–9].

Using these two values of C_n , we simulate the loss of gain, shown in Figure 4. On good days, it is expected to be less than 0.25 dB, and on the worst days about 1.5 dB, again at 25°.

Angle of Arrival

Random fluctuations in the index of refraction can lead to slight changes in the apparent angle of arrival of the target. This phenomenon is also called “anomalous refraction.” We investigated how likely these fluctuations were and found that even on the worst days, a perceived angle-of-arrival deviation of greater than a half-beamwidth was exceedingly rare, as shown in Figure 5. Moreover, these large deviations have been reported to last only a few seconds [7]; thus, the HUSIR tracker would not be at risk of losing track of the target.

Per-Pulse Phase Variation

The most critical atmospheric effect is the addition of a phase shift on each pulse. Successive wavefronts from the in-orbit radar target as it is tracked across the sky will include this time-varying phase, even after compensating the wavefronts for the target trajectory. Nonlinear variation of this phase distorts the resulting ISAR image in the cross-range direction.

On any particular day, the phase error will depend to some degree on meteorological conditions, such as wind speed, humidity, and the movement of fronts, but

the precise amount per pulse is not readily predictable. Co-boresighted water-vapor radiometers have been used to predict the amount of water vapor along the line of sight of radio antennas [8], but because of HUSIR's higher frequency range and the intrinsic error in the radiometer measurement, this method is not expected to be accurate enough for HUSIR to reduce the phase error to levels that will not cause defocusing [10]. Therefore, compensation for this atmospheric phase by post-processing techniques (e.g., autofocus) is a necessity for HUSIR.

The amount of per-pulse phase error can be quantified in a number of ways. We choose to describe the atmospheric phase error by the coherence time T_c , the duration of time in which the coherent sum of unit-amplitude phasors falls to half the value of the noncoherent sum:

$$\left\langle \left| \frac{1}{T_c} \int_0^{T_c} e^{i\varphi(t)} dt \right| \right\rangle = 0.5. \quad (6)$$

Turbulence theory provides estimates of coherence times, and these calculations have been supported by in situ phase measurements at X, Ku, and W bands.

As can be seen in Figure 6, the simulated coherence time ranges from several seconds in the worst case to several hundred seconds in the best case. These values are consistent with W-band coherence times reported in the literature [9–11]. The typical time for HUSIR to form an image depends on the rotation rate of the satellite with respect to the radar; for low Earth orbits, this “delay” can range from about 10 to more than 100 seconds. This simulation implies that we will have to compensate for this phase degradation, not necessarily for every image but certainly for a significant number of them in order to form highly resolved images. The positive implication of this simulation is that, even in the most dire conditions, we will have a few seconds of W-band coherency, which will permit the HUSIR tracker, with a pulse-repetition frequency >200 Hz, to accrue significant integration gain.

Atmospheric Effects Measurements

Our calculations of coherence times needed to be corroborated with real data from the Haystack area. Ideally, we would use the residual phase from an in-orbit satellite having a constant scatterer, such as a sphere, as our phase error estimate. This approach would allow us to sample the atmosphere at higher elevations (both angle and height off the ground). The difficulty that arises with

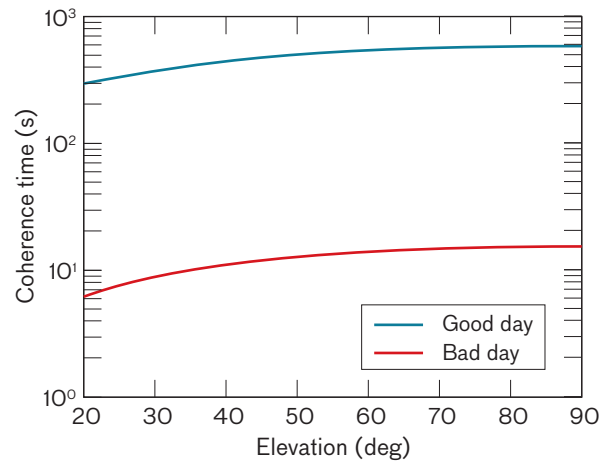


FIGURE 6. Simulated coherence times as a function of elevation. The coherence time varies from a few seconds to several hundred seconds depending on conditions.

this approach is resolving between ephemeris and atmospheric phase errors. One can reduce the ephemeris phase errors by fitting a low-order polynomial to the unwrapped phase, but from preliminary investigations using HAX and Haystack radars, we have found that significant higher-order ephemeris errors still remain. For this reason, we use an Earth-bound corner reflector.

HAX SUBSTITUTE FOR HUSIR VALIDATION

We performed a series of radar path-length experiments using radars on Millstone Hill (elevation 370 feet) and a 3-foot corner reflector on top of Mount Wachusett (elevation 2006 feet, see Figure 7), 36 km to the southwest. These experiments were predominantly performed at Ku band with the HAX radar, but also, for one week in April 2005, at W band with the University of Massachusetts at Amherst (UMass) mobile “tornado chaser” radar [12]. The phase of the return echo from the tower as a function of time was unwrapped to give a phase history, which was translated into a path-length variation.

UMASS W-BAND DATA COLLECTION

Dual Ku-band (16.7 GHz) and W-band (95 GHz) collection was performed in early April 2005 to validate the frequency proportionality assumption of the phase variation, i.e., that the atmospheric index of refraction is independent of frequency in the microwave regime. This validation was important because it allowed us to use scaled Ku-band data from HAX, which at the time were more readily available than were W-band data, to




FIGURE 7. In April 2005, the HAX radar and a small W-band radar from the University of Massachusetts at Amherst made simultaneous measurements of a corner reflector on top of Mount Wachusett, 36 km away. These measurements were used to validate frequency-scaling Ku-band data from HAX to W-band frequencies.

estimate the atmosphere at W band. The truck-mounted UMass W-band radar was situated in the parking lot near the HAX radar. The diameter of the UMass dish is 4 feet. There was a difference of approximately 20 m between the lines of sight to Wachusett for the HAX and UMass radars, although their beams had a 50% overlap halfway to Wachusett. For most of the week of the W-band experiment, the weather was excellent, with clear conditions, daytime temperatures averaging about 55°F, and rela-

tive humidity around 40%. As can be seen in Figure 8, the two radars, operating at two different frequencies and sharing no hardware (including timing sources), tracked the same atmospheric-dependent path-length variations. The root-mean-square difference between the two (one-way) path-length histories, over the 6500 seconds of measurements, is 0.4 mm, corresponding to one quarter of a phase cycle from two-way travel. Some of the residual difference may be due to unwrapping errors, as the UMass W-band data were relatively noisy, approximately 15 dB signal-to-noise ratio (SNR), primarily because of the small dish size and the attenuation from the 72 km round-trip distance. The match between Ku-/W-band path-length variation validates our assumption of frequency invariance of the atmosphere (at least between Ku band and W band), and therefore we could use the HAX radar Ku band to investigate the severity of the atmosphere’s effects at W band.

During the UMass experiment, the W-band radar was operated continuously during a three-day period (8–10 April). The HAX radar was not operational during this time. W-band phase and amplitude data were collected on the Wachusett tower. These data are plotted in Figures 9a and 9b, respectively, along with theoretical estimates from standard models (refractivity [13] and attenuation [2]).




FIGURE 8. When scaled to the same frequency, the phase variations observed by HAX and the UMass radar are strongly correlated, although the UMass data are noisier than the HAX data. These data were collected on 6 April 2005.

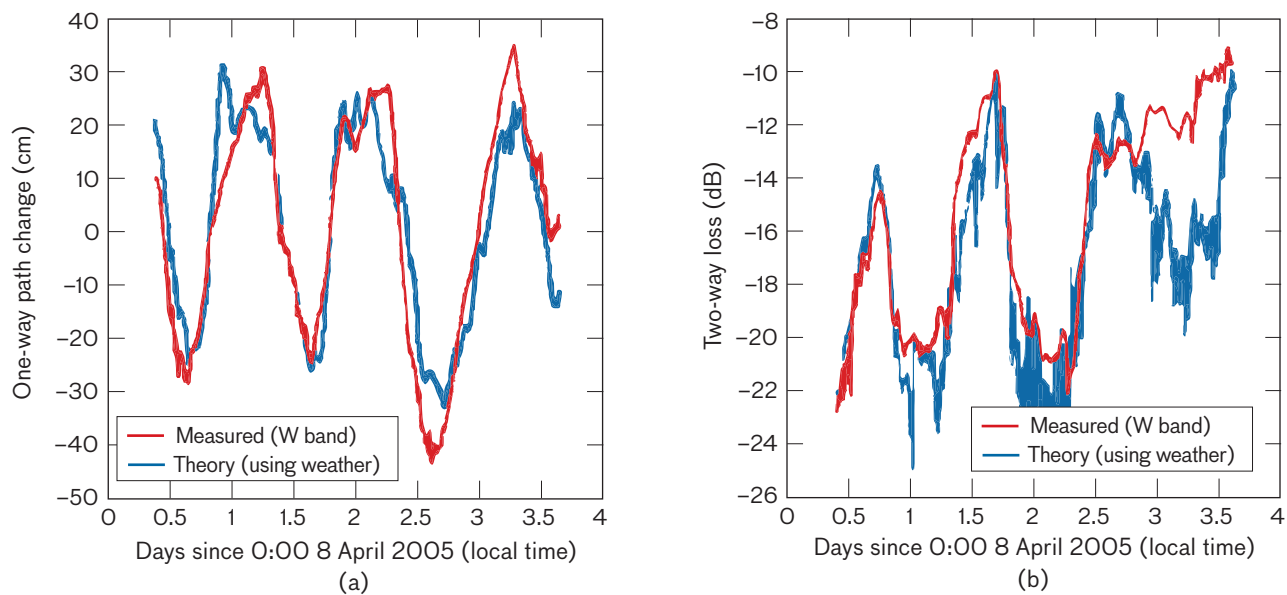


FIGURE 9. Measured W-band path length (a) and attenuation (b) over three days in April 2005. The delay and attenuation estimated from the local weather conditions are correlated with the radar measurements. The weather conditions used to generate the theory curve are from a single point and do not necessarily indicate conditions along the whole path, so some deviation should be expected.

In Figure 9, we can clearly see the diurnal variations in the path length, with early morning as the time of most path length and late afternoon as the least. Correspondingly, the loss variations match the path-length variations, with early morning (for this day) having the most attenuation along the path length, and the reverse for early afternoon. These results are consistent with calculated water-vapor density variations, as in the early morning the temperature cools, the column of water vapor contracts, and the density increases. The denser the water vapor is, the more path length is added and the more loss is engendered.

The theoretical curves for refractivity and attenuation match well with the measured data, except for the estimated loss curve around midnight on day 3. The meteorological measurements were collected at the site of the radar, so they may not indicate the conditions along the entire line of sight.

HAX “W-BAND” COHERENCE TIMES

HAX collected data returning from the Mt. Wachusett corner reflector on almost a weekly basis during 2005 to 2006. To calculate W-band coherence times from the HAX data, the Ku-band phase is scaled to W band, the path length is set to an elevation of 25° , and Equation (6)

is used. The coherence time associated with each observation is the mean coherence time. The coherence times from the complete HAX-Wachusett dataset are plotted in Figure 10. We can see that the range of values is consistent with the simulated coherence times in Figure 6, and that in the summer, coherence times of a few tens of seconds to even less than ten seconds are commonplace, whereas in the winter coherence times of a several hundred seconds are typical. By coupling these shorter coherence times with the increase in attenuation from greater water-vapor content, it is clear that the summer months will present the most challenging imaging environment for HUSIR.

Atmospheric Phase Correction

The Wachusett measurements by HAX and the associated weather conditions were used to simulate atmospheric conditions for HUSIR satellite observations. A per-pulse phase, derived from the HAX data and scaled to W band and appropriate path length, was injected, and the thermal noise floor was set by using the HUSIR sensitivity, range to the target, and weather conditions.

The simulated imaging chain uses compact-range data of a model satellite taken at W band. Figure 11 shows such an image that is using measured phase and weather




FIGURE 10. HAX data, when scaled to W band, were used to measure the coherence times of the atmosphere over more than a year. The coherence times vary from a high of 1000 seconds during the winter to a low of less than 6 seconds during the summer.


data from one of the “bad” days. Interpretation of the image shown in Figure 11a would be severely limited. To correct this phase distortion, an autofocus routine was utilized. There are several different autofocus algorithms in the SAR literature, such as phase gradient autofocus

(PGA) [14] or metric-optimization (such as image entropy [15]). Through a comparative study, we found that the most robust method, both in the ability to correct the images in the presence of phase noise and in computation time, is a variant of an entropy-minimization algorithm that makes use of the FFT to speed up the convergence time. This technique was developed by T. Kragh and G. Ushomirsky at Lincoln Laboratory [16]. The results of applying this algorithm to the image in Figure 11a are shown in Figure 11b. The image is essentially restored.


The performance of the entropy minimization and PGA routines as a function of SNR was explored. Here, SNR is on a per-pulse basis (i.e., before the cross-range FFT is applied) and is defined as the peak-to-noise floor difference. It was found that the routines are ineffective below 0 dB SNR, but by 10 dB SNR, the algorithms approach optimal performance, that is, nearly complete elimination of the effects of the applied “atmospheric” phase.

We note that the autofocus routines can eliminate the random phase that is applied to each pulse, but cannot remove the phase randomization of thermal noise, which is applied to each bin of every (range) pulse.

The SNR threshold of 10 dB should be attainable, provided that the radar can track the target. Typically,



(a)



(b)

FIGURE 11. By using W-band compact-range data, simulated atmospheric distortion (a) can be eliminated and corrections can be made using autofocus (b).

the tracker needs >10 dB on a single pulse, or an integration of pulses on a short time scale (<2 seconds), to keep the target in track. Thus, “If we can track it, we can image it” holds true for the HUSIR system, even with atmospheric phase noise.

Atmospheric Challenges Resolved

HUSIR faces several challenges in image formation caused by the atmosphere. One of the challenges, the loss in signal from atmospheric absorption, is overcome simply by the transmit power and the tremendous gain of the antenna. Another deleterious effect of the atmosphere is per-pulse phase error. We have found, from theoretical calculations and in situ measurements, that the expected W-band atmospheric phase error has the potential to cause significant distortions in the ISAR image. W-band correlation times of 10 seconds or less during the Massachusetts summer are possible, as shown by simulated and measured data. However, through autofocus routines a properly focused image can be achieved. The entropy-minimization algorithm developed at Lincoln Laboratory has already been incorporated into the HUSIR imaging chain for its excellent focusing ability and short computational time.

Real-Time Processing

The real-time processing performed by HUSIR takes place in a subsystem known as the digital pulse-compression subsystem (DPCS). At the time it was developed, this system pushed the capabilities of available off-the-shelf acquisition hardware to accomplish the processing required in real time. But before the DPCS can do any processing, the target must be maintained in the beam (tracked); this step is not as trivial as it sounds.

Tracking

The 120-foot-diameter antenna at W band yields a 5.7-millidegree (mdeg) diameter beam. This small beam size causes issues with tracking that are not typical with other radars. Typically, the angular motion of a satellite during the radar pulse time of flight is inconsequential because of the large size of the beam relative to that angular motion. For HUSIR, however, a satellite’s motion results in frequent track loss unless the standard tracking algorithm is altered. This solution results in an unavoidable loss of sensitivity but maintains track through the pass.




FIGURE 12. As a satellite flies over the radar site, its angular velocity relative to the radar changes with zenith angle. When the satellite is low on the horizon, it moves slowly, and as it approaches zenith, it moves faster. This angular rate, coupled with the time of flight of the pulses, results in an apparent separation of the receive echoes from the target and the position where the antenna needs to point to have the transmit pulse hit the target. For HUSIR, this separation is a large fraction of a beamwidth and resulted in frequent loss of track.

To illuminate the problem, assume a satellite is in a circular orbit. Neglecting the Earth’s rotation, the Earth-centric angular rate of the orbit is $\theta = 180/\pi \text{ sqrt}[\mu/r^3]$ (deg/s), where $\mu = 3.99 \times 10^5 \text{ km}^3/\text{s}^2$, a constant for all objects orbiting the Earth. For a zenith-crossing pass over the radar site, the maximum angular rate is given by

$$\frac{da(0^\circ)}{dt} \approx \frac{\theta(R_{\text{Earth}} + Alt)}{Alt},$$

where R_{Earth} is the radius of the Earth and Alt is the satellite altitude. The maximum angle change during the one-way time of flight of the radar pulse is

$$\Delta = \left(\frac{da}{dt}\right)\left(\frac{rng}{c}\right),$$

where rng is the range from the radar to the satellite and c is the speed of light. This time of flight results in an apparent tracking lead when viewing the monopulse response. That lead at zenith is fairly constant with orbit altitude, even over orbit altitudes from 200 km to 2000 km with the corresponding order of magnitude angular rate change. This tracking lead results in a transmit-pulse to

receive-pulse angular-separation curve, shown in Figure 12 for a zenith-crossing target at 1000 km altitude.

This beam-space angular separation is typically less than 3 mdeg for circular orbits, but can be worse for elliptical orbits. For all of the other space tracking radars, the beamwidths are sufficiently large that a 3 mdeg transmit-beam to receive-beam delta is insignificant; the receive echo will fall in the linear region of the monopulse response and poses no problem for the tracker. With HUSIR's 5.7 mdeg beamwidth, 3 mdeg is a very significant offset that can result in loss of track for almost any satellite track. For example, since the radar tracker, by default, attempts to keep the antenna pointed directly at the target, HUSIR loses sensitivity on both the transmit and receive signal paths, as shown in Figure 13.

Because the offset places the receive echo near the very edge of the monopulse response, it is in a nonlinear region, and any tracking noise (e.g., resulting from the servo loop or low signal to noise) can throw the target outside the capability of the monopulse system to maintain track. The solution to maintaining track is to modify the tracker so that it points the antenna in the direction

where the receive echo is expected. This is accomplished by introducing a tracking lag equivalent to the one-way flight time of the pulse, exactly compensating for the apparent lead of the return echo. This shift changes Figure 13 to look like Figure 14 and results in a loss of SNR of an additional 1.5 dB assuming perfect tracking. If there are tracking errors in addition, the loss of sensitivity could easily be as much as 5 dB. Fortunately, this loss is at a maximum when the target is at minimum range, and the SNR is at its largest, minimizing its overall impact to the system.

The Digital Pulse-Compression Subsystem

The HUSIR DPCS is a pulse-repetition-interval (PRI) rate digitizer and preprocessor of radar return pulses that feeds resulting signature data into the main radar computer (MRC). The MRC stores the data in a ring-buffer, making them available to the tracking and recording tasks of the HUSIR's real-time program (RTP). The DPCS is also responsible for assembling and synchronizing auxiliary data provided by all other radar subsystems, such as the master timing system (MTS), receiver control system, transmitter control system, and antenna

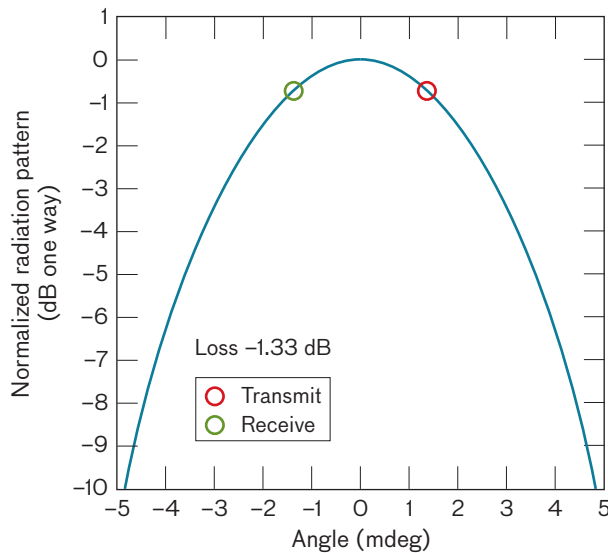


FIGURE 13. The default Radar Open Systems Architecture (ROSA) radar tracker attempts to always point the antenna at the current location of the satellite, thereby splitting the angular difference between the receive echo and the transmit pulse. For most radars, this angular separation has an inconsequential effect on sensitivity, but for HUSIR, it results in a measurable loss of sensitivity.

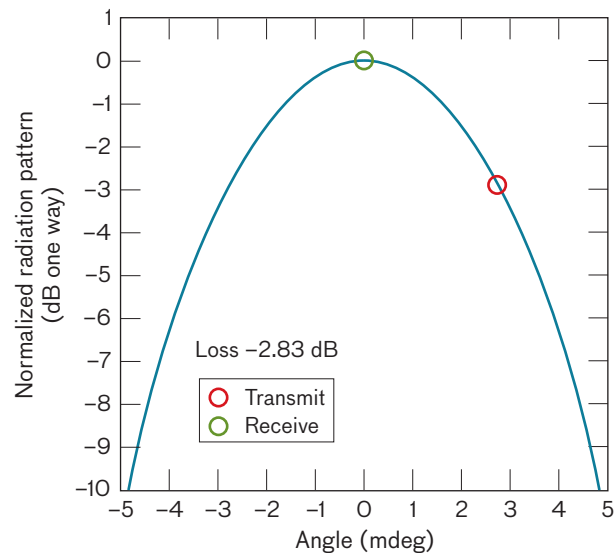


FIGURE 14. When the tracker is modified to keep the antenna pointed where the receive echo is expected to come from, there is an additional loss in sensitivity caused by the transmit pulse sliding further down the antenna pattern skirt, but the radar maintains track. There is a sufficient margin in the sensitivity budget to accept this increased loss to maintain track.




FIGURE 15. A ROSA subsystem consists of a VME chassis and a number of boards. Each subsystem, regardless of purpose, has a set of common boards, such as the VME controller, reflective memory board, and timing generator boards. Additional boards are added to perform the functions required. For the DPCS, Mercury Computer Systems DSP boards are installed to perform digital downconversion and pulse compression.

control system. The DPCS digitizes and downconverts three radar signature channels: the principal polarization return, along with the two angle error channels, traverse and elevation. A four-channel design has been implemented to provide a spare channel that may be used if the radar were ever upgraded to collect the opposite polarization target return.

The DPCS serves the same role as a standard, first-generation ROSA DPCS. However, it has been redesigned to meet a 4000 Hz pulse-repetition-frequency (PRF) requirement, which is twice that of the original ROSA design. The HUSIR design, developed around 2004, carries over many elements of the first-generation system, such as Versa Module Europa (VME) bus technology with a Motorola single-board computer as the main controller running the VxWorks operating system. The system also receives its commands via a reflective-memory interface looped to the MRC and all the other radar VME subsystems, but it uses a newer interface with 10 times the bandwidth of the original. A significant redesign of the signal processing portion of the DPCS was motivated by the loss of the commercial vendor responsible for the first-generation application-specific integrated circuit (ASIC)-based signal processing boards. Ultimately, a new

analog/digital (A/D) board with a field-programmable gate array (FPGA) direct digital receiver (DDR) section was selected to interface, via a RACE++ backplane, to multiple specialized compute-engine processing boards. Finally, proprietary serial front-panel data-port interfaces send the processed radar signature and auxiliary data at a PRI rate to the MRC. A picture of the DPCS VME chassis is shown in Figure 15.

Pulse-to-pulse waveform diversity was a critical requirement for the HUSIR DPCS and drove the selection of fast, highly configurable A/D, digital receiver, and signal processing hardware. Reconfiguration of A/D board parameters, filter settings, and compute-engine processing modes is possible pulse to pulse up to the 4000 Hz maximum PRF. The primary MVME6100 controller board is responsible for receiving waveform commands from the RTP and MTS via reflective memory. The main processing task on this board is interrupted at the PRF of the radar. The A/D and signal processing boards are then configured for the next receive pulse via VME bus commands from the main processor board. The A/D board selected is a four-channel Echotek ECV4-4, which contains FPGA firmware that implements four Graychip-like DDRs. The three radar intermediate-frequency inputs to the A/D board are 20 MHz bandwidth centered at 10 MHz, and are thus sampled with a 40 MHz clock. The internal DDRs in the ECV4-4 implement fast infrared filters with configurable decimation settings of 2, 4, 8, or 16, which are sufficient for most of the HUSIR tracking waveforms. Additional cascaded integrator-comb filtering is available within the ECV4-4 for higher decimations. Triggering of the A/D board is provided by a customized Lincoln Laboratory timing-generator board interrupted at receive time. The ECV4-4 board forwards processed data to signal processing boards via an eight-slot RACE++ serial interlink backplane. Results from each channel are simultaneously delivered to dedicated processors within Mercury MCJ6 boards, each with two dual PowerPC (PPC) compute engines. Custom signal processing of the radar data is performed within the compute engines as defined by precompiled and downloaded C code.

Pulse Processing

For a typical stretch tracking waveform, the processing consists of the following steps that are partially depicted in Figure 16:

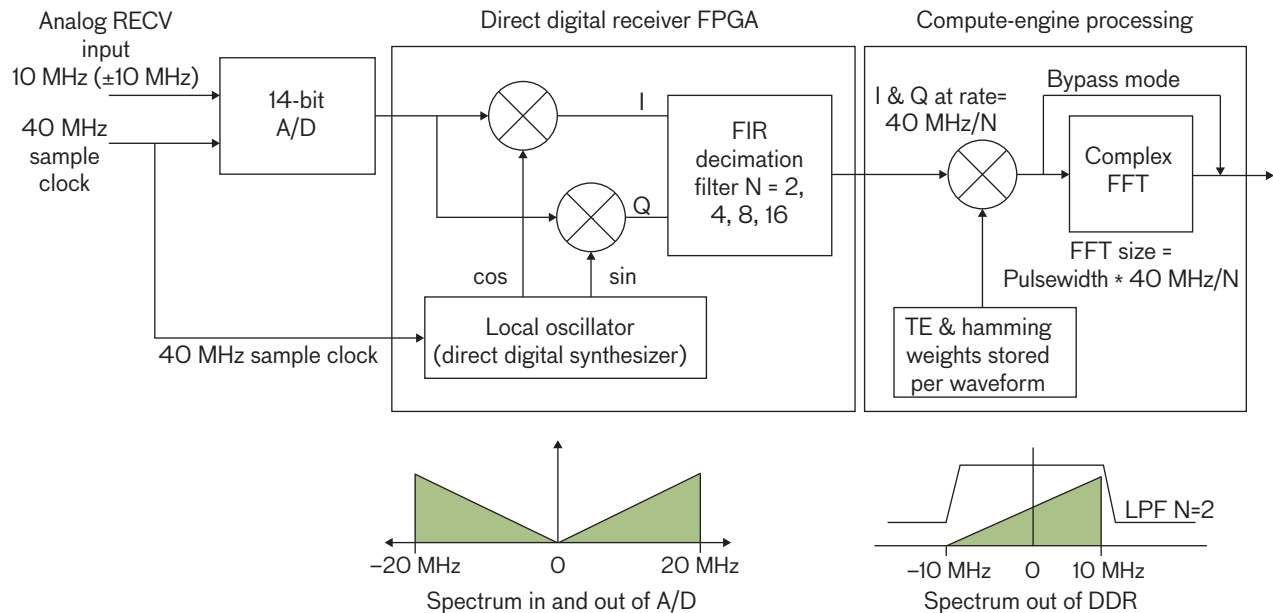


FIGURE 16. Single-channel radar channel processing performed by DPCS for a typical stretch waveform. Here, I and Q are the in-phase and quadrature data signals.

- Convert incoming intensity and phase time-series data to IEEE format
- Apply time-domain hamming and optionally transverse equalization weighting to pulse data
- Apply correction for randomized transmitter phase shifting (used to reduce internal DC leakage)
- Apply fine-range dithering corrections in time domain
- Perform FFT to frequency/range domain
- Apply fine-range dithering corrections in frequency domain
- Scale and format data for transfer to main radar computer

For pulsed continuous-wave processing, some of the above stages, such as the FFT step, are bypassed. Other waveform modes and custom filters can be implemented with C++ and the Mercury-provided signal processing libraries. The last Mercury board in the VME chassis is equipped with only one dual PPC board that is responsible for formatting all the processed radar data and forwarding it to the RINOJ serial fiber-optic interface, which is compatible with Curtiss-Wright FibreXtreme boards installed within the MRC. A diagram of the inter-board data flow is shown in Figure 17. Once the data reach the MRC, they are available for integration, detection processing, and recording to disk.

Post-Processing

HUSIR is a world-class sensor that is already producing exquisite imagery. However, factors affecting image quality that can be ignored for lower-frequency radars must be addressed for HUSIR to avoid significant artifacts. For example, W-band imaging is more sensitive to errors resulting from uncompensated target motion, nonideal transmitter and receiver system responses, and atmospheric effects. To realize the full high-resolution capability of HUSIR, these effects must be carefully accounted for in the image-formation process, thus necessitating a reevaluation of the existing framework for radar image processing. This reevaluation has motivated the development of a new image processing pipeline with algorithms designed to correct for the potential artifacts. The revised pipeline has been instrumental in forming best-possible images in the new W-band mode. This section describes the image-formation sequence, highlighting processing enhancements added for HUSIR.

The process of image formation involves three stages, as shown in Figure 18. The first is preprocessing, which attempts to correct for errors affecting the received pulses. These errors include uncompensated target motion and waveform defects introduced by the open-loop response of the radar and atmosphere. The second stage is image generation, in which an image is formed from the corrected

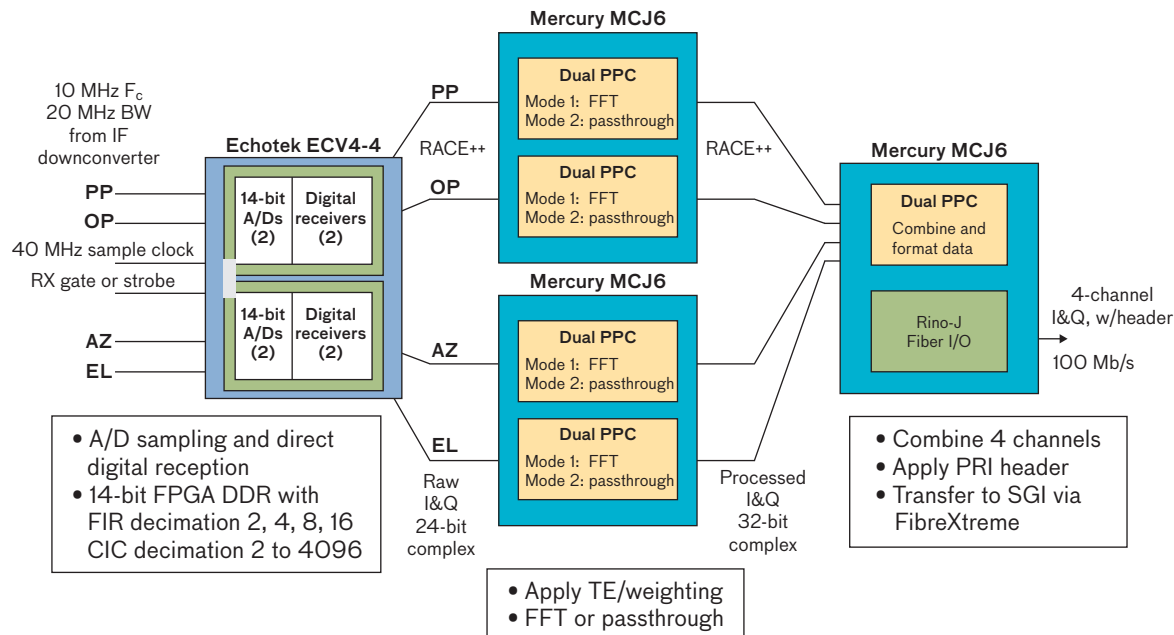


FIGURE 17. Four-channel, four VME board data flow. Data from each channel are routed through different processors to obtain the throughput required. In the figure, F_c is the center frequency, BW is bandwidth, FIR is finite impulse response, PRI is pulse rate interval, PP is principal polarization, OP is orthogonal polarization, TE is transverse equalization, and CIC is cascaded integrator-comb filter. AZ and EL are azimuth and elevation.

pulse data. The third stage is a “cleanup” stage, in which restoration algorithms are applied to the image to correct residual defocusing. When we were developing the revised pipeline, each stage had to be altered for HUSIR to accommodate effects associated with W-band operation.

Preprocessing

Each image formation stage is broken into subcomponents, detailing the end-to-end process of forming an image. The process starts with raw pulse-compressed radar data, which are subject to per-pulse range-shift and phase errors resulting from uncompensated motion, as described in detail earlier in this article. These errors cause the responses of point scatterers to spread into multiple resolution cells, resulting in image blurring. Range alignment, the first step in the preprocessing stage, uses the raw signal data to derive a refined estimate of the target motion for correcting the range and phase errors. New range-alignment approaches developed for HUSIR go beyond correcting time-of-flight range errors. These approaches compensate for inaccurate Doppler correction effects on range walk; these effects were not accounted for in previous processing and are critical at W band. In addition, the new algorithms benefit from exploiting the pulse-to-pulse phase informa-

tion, or coherence between pulses, to overcome signal fading. A Doppler-tracking algorithm then cleans up residual errors from range alignment. This new processing step was motivated by occasional poor alignment observed for some HUSIR datasets. Subsequent waveform correction modifies the magnitude and phase response of each pulse to compensate for signal defects introduced by the transmitter and receiver hardware and the atmosphere. For lower-frequency radars, data collected periodically from calibration spheres can be used for this purpose. However, at W band, dynamic effects caused by hardware make the use of calibration-sphere data alone insufficient for waveform correction, thus requiring a data-driven approach to the problem. Consequently, a new waveform-correction algorithm added to the pipeline for HUSIR optimizes an image focus metric to determine a magnitude and phase equalization function for each imaging interval. The output of this stage is a set of aligned and equalized pulses that can be used in the image-generation stage.

Image Generation

The first step in the image-generation stage is to perform band-limited interpolation in time to generate a reduced set of pulses that correspond to equal aspect angles. This

process increases the SNR of each pulse and enables the use of efficient image-formation techniques that rely on uniformly sampled input data. The use of band-limited interpolation was introduced for HUSIR to improve image fidelity; previous processing techniques used nearest-neighbor interpolation and coherent-pulse averaging as an approximation to this approach. The next step is to perform polar-to-Cartesian resampling, or polar reformatting, to generate a uniform grid of spatial-frequency samples. A weighted two-dimensional inverse FFT is then applied to produce the final image.

Post-Processing Cleanup

After the image-generation stage, post-processing algorithms are applied to correct fine-scale defocusing—the final cleanup stage. Image autofocus corrects per-pulse phase errors resulting from uncompensated motion not detected by the range-alignment stage and from atmospheric effects. The phase errors produce a distorted point spread function with high sidelobes that causes image defocusing in cross range. Autofocus algorithms derive a phase correction from the imagery to correct the defocusing. A new multistage optimization-based phase and magnitude autofocus routine that uses multiple focus metrics was developed for HUSIR to handle artifacts and phenomena specific to the new W-band mode. This routine produces exceptional results for a variety of phase errors, improving image focus and minimizing off-target sidelobe levels. In addition, an extended version of map-drift autofocus techniques is applied to the processed image to correct acceleration and higher-order motion errors [17]. The correction is applied as a two-dimensional phasor for improved fidelity, as opposed to the one-dimensional phase correction used in traditional autofocus techniques. Lastly, several processing options aimed at visually enhancing the images are available. These options tend to artificially sharpen images to improve their appearance and, when used alone, generally do not increase the image information content. However, the techniques can improve image interpretability when larger integration angles are utilized.

HUSIR and the Future

The future for HUSIR signal processing will see continued improvement in imaging algorithms, as the new imaging

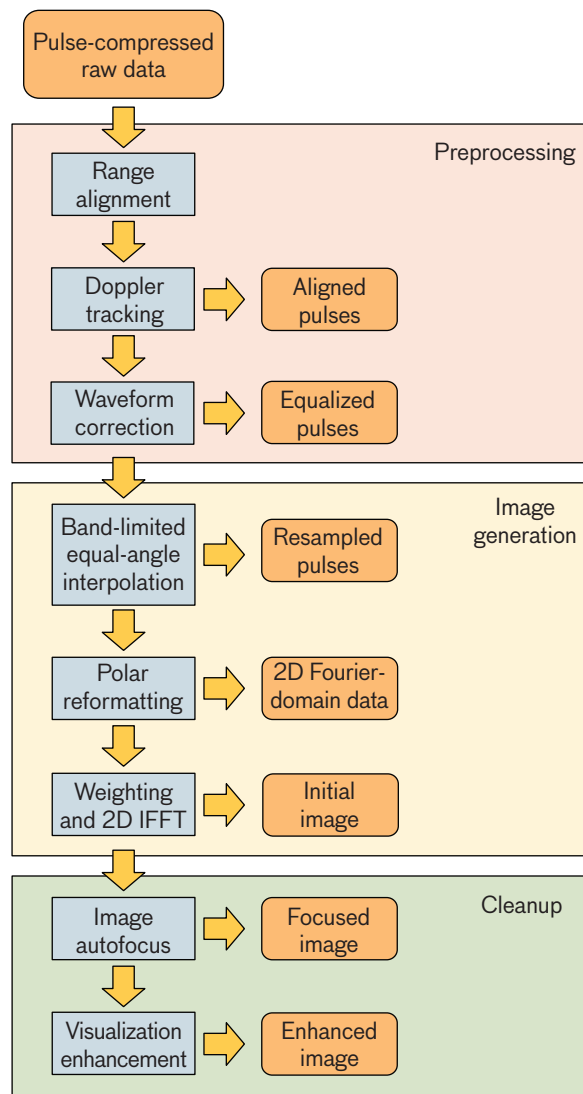


FIGURE 18. The HUSIR image-formation sequence as it has been implemented. This processing sequence is similar to standard synthetic aperture radar image-generation techniques. Here, IFFT stands for inverse fast Fourier transform.

chain is designed to be easily extensible, with the ability to insert new algorithms as they are developed. Improvements in the radar image exploitation arena enabled by the highly resolved, sharp images now being generated are all but assured.

Acknowledgment

The authors gratefully acknowledge Mark Dickson for his analysis of the tracking offset. ■

REFERENCES

1. J.R. Klauder, A.C. Price, S. Darlington, and W.J. Albersheim, "The Theory and Design of Chirp Radars," *The Bell System Technical Journal*, vol. 34, no. 4, 1960, pp. 745–808.
2. H. Liebe, "An Updated Model for Millimeter Wave Propagation in Moist Air," *Radio Science*, vol. 20, no. 5, 1985, pp. 1069–1089.
3. G. Brussard and P.A. Watson, *Atmospheric Modelling and Millimetre Wave Propagation*. London: Chapman and Hall, 1995.
4. R.N. Treuhaft and G.E. Lanyi, "The Effect of the Dynamic Wet Troposphere on Radio Interferometric Measurements," *Radio Science*, vol. 22, no. 2, 1987, pp. 251–265.
5. C. J. Naudet, "Estimation of Tropospheric Fluctuations Using GPS Data," *Telecommunications and Data Acquisition Progress Report*, vol. 42, no. 126, 1996, pp. 1–19.
6. R. Linfield, "The Effect of Aperture Averaging upon Tropospheric Delay Fluctuations Seen with a DSN Antenna," *Telecommunications and Data Acquisition Progress Report*, vol. 42, no. 124, 1996, pp. 1–7.
7. L. Olmi, "Systematic Observations of Anomalous Refraction at Millimeter Wavelengths," *Astronomy and Astrophysics*, vol. 374, no. 1, 2001, pp. 348–357.
8. L.P. Teitelbaum, R.P. Linfield, S.J. Keihm, M.J. Mahoney, and G.M. Resch, "A Demonstration of Precise Calibration of Tropospheric Delay Fluctuations with Water Vapor Radiometers," *Telecommunications and Data Acquisition Progress Report*, vol. 42, no. 126, 1996, pp. 1–8.
9. A.E.E. Rogers, A.T. Moffet, D.C. Backer, and J.M. Morgan, "Coherence Limits in VLBI Observations at 3-Millimeter Wavelength," *Radio Science*, vol. 19, no. 6, 1984, pp. 1552–1560.
10. M.A. Holdaway, "Atmospheric Coherence Times at Chajnantor," Main Alma Memo Series, no. 169, 1998.
11. J.H. Biegging, J.M. Morgan, W.W.J. Welch, S.N. Vogel, and M.C.H. Wright, "Interferometer Measurements of Atmospheric Phase Noise at 86 GHz," *Radio Science*, vol. 19, no. 6, 1984, pp. 1505–1509.
12. H.B. Bluestein and A.L. Pazmany, "Observations of Tornadoes and Other Convective Phenomena with a Mobile, 3-mm Wavelength, Doppler Radar: The Spring 1999 Field Experiment," *Bulletin of the American Meteorological Society*, vol. 81, 2000, pp. 2939–2951.
13. E.K. Smith, Jr., and S. Weintraub, "The Constants in the Equation for Atmospheric Refractive Index at Radio Frequencies," *Proceedings of the Institute of Radio Engineers*, vol. 41, 1953, pp. 1035–1037.
14. C. Jackowatz, D. Wahl, P. Eichel, D. Ghiglia, and P. Thompson, *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*. Boston: Kluwer Academic Publishers, 1996.
15. J.R. Fienup, "Synthetic Aperture Radar Autofocus by Maximizing Sharpness," *Optics Letters*, vol. 25, no. 4, 2000, pp. 221–223.
16. T.J. Kragh, "Monotonic Iterative Algorithm for Minimum-Entropy Autofocus," *Proceedings of the Adaptive Sensor Array Processing (ASAP) Workshop*, June 2006.
17. C.E. Mancill and J.M. Swiger, "A Map Drift Autofocus Technique for Correcting Higher Order SAR Phase Errors," *Proceedings of the 27th Tri-Service Radar Symposium*, 1981, pp. 391–400.

ABOUT THE AUTHORS



James V. Eshbaugh is a senior staff member of the Aerospace Sensor Technology Group. He joined Lincoln Laboratory in 2000 and has worked on several programs involving imaging radars. He served as the systems engineer and control system unit engineer during the HUSIR program. He received a bachelor's degree in electrical engineering and a doctorate

in microwave engineering from the University of Massachusetts at Amherst.



Robert L. Morrison, Jr. is a member of the technical staff in the Aerospace Sensor Technology Group. He received a bachelor's degree in electrical engineering from the University of Iowa in 2000 and master's and doctorate degrees in electrical engineering from the University of Illinois at Urbana-Champaign in 2002 and 2007, respectively. His graduate work focused

on data-driven image-restoration techniques for computer imaging systems, including advanced algorithms for synthetic aperture radar autofocus and parallel magnetic resonance imaging reconstruction. Since joining the Laboratory in 2007, he has been active in the areas of radar imaging, multistatic radar data exploitation, and radar systems analysis.



E. Weber Hoen is a staff member in the Aerospace Sensor Technology Group. His work focuses on advanced techniques for imaging of satellites. He joined Lincoln Laboratory in 2001. He received a bachelor's degree in physics from the University of North Carolina at Chapel Hill, and master's and doctorate degrees in applied physics from Stanford University. His

doctoral thesis concerned interferometric synthetic aperture radar measurements of the Greenland ice sheet.



Timothy C. Hielt is a member of the technical staff in the Aerospace Sensor Technology Group. He joined the MIT Haystack Observatory staff in 1999 after more than 10 years performing basic research in weather radar, infrared, and night-vision systems at the Air Force Research Laboratory. In 2009, he transitioned to Lincoln Laboratory to serve as the radar systems

manager for the Lincoln Space Surveillance Complex, which is responsible for operations, maintenance, and upgrades of the HUSIR, HAX, and Millstone Hill radars. He received bachelor's and master's degrees in microwave engineering from the University of Massachusetts at Amherst.



Gerald R. Benitz is a senior staff member in the Airborne Radar Systems and Techniques Group. He received a bachelor's degree in electrical engineering in 1979 from Purdue University, master's degrees in electrical engineering and mathematics, and a doctorate in electrical engineering in 1984, 1985, and 1987, respectively, from the University of Wisconsin at Madison, and then joined Lincoln Laboratory. His areas of research include analysis, development, and testing of algorithms for direction finding and waveform estimation with adaptive antenna arrays, and the application of adaptive processing to synthetic aperture radar (SAR) and other radar modes. His recent research includes applications of adaptive antenna arrays for SAR, moving-target detection, and terrain mapping with both monostatic and bistatic radar. He has published numerous papers and project reports on adaptive processing and superresolution, and their application to direction finding, SAR, and automatic target recognition. He received a patent for his work in SAR superresolution entitled "High-Definition Vector Imaging."



This aerial view shows the HUSIR construction site in August 2010 just before the installation of the new backstructure. The 400-foot-tall main crane towers over the truncated radome and the radome cap (upper left). In the foreground, a smaller crane lifts the quadrapod onto the backstructure.