Presented at the SPIE Defense and Security Symposium, April 12-16, 2004, Orlando, Florida

An X-Band Radar Terrain Feature Detection Method for Low-Altitude SVS Operations and Calibration Using LiDAR

Steve Young[a], Maarten Uijt de Haag[b], Jacob Campbell[b]

[a]NASA Langley Research Center, Hampton, VA, USA 23681
[b]Ohio University, Athens, OH, USA 45701

ABSTRACT

To enable safe use of Synthetic Vision Systems at low altitudes, real-time range-to-terrain measurements may be required to ensure the integrity of terrain models stored in the system. This paper reviews and extends previous work describing the application of x-band radar to terrain model integrity monitoring. A method of terrain feature extraction and a transformation of the features to a common reference domain are proposed. Expected error distributions for the extracted features are required to establish appropriate thresholds whereby a consistency-checking function can trigger an alert. A calibration-based approach is presented that can be used to obtain these distributions. To verify the approach, NASA's DC-8 airborne science platform was used to collect data from two mapping sensors. An Airborne Laser Terrain Mapping (ALTM) sensor was installed in the cargo bay of the DC-8. After processing, the ALTM produced a reference terrain model with a vertical accuracy of less than one meter. Also installed was a commercial-off-the-shelf x-band radar in the nose radome of the DC-8. Although primarily designed to measure precipitation, the radar also provides estimates of terrain reflectivity at low altitudes. Using the ALTM data as the reference, errors in features extracted from the radar are estimated. A method to estimate errors in features extracted from the terrain model is also presented.

Keywords: digital elevation model, integrity, x-band radar, monitoring, LiDAR, terrain shadowing

1. INTRODUCTION

Flight in Instrument Meteorological Conditions (IMC) requires pilots to manipulate flight controls while referring to a Primary Flight Display (PFD). The PFD indicates aircraft attitude along with, in some cases, many other state variables such as altitude, speed, and guidance cues in the form of Flight Director (FD) symbology. The FD is a predictor that indicates to pilots where to steer the aircraft. Current PFDs provide no information about the terrain environment in the vicinity of the aircraft. In IMC, terrain avoidance depends on a well-designed flight path for the FD to follow, diligent chart referencing by the crew, and/or sufficient ATC monitoring and communication. In an aircraft equipped with a Terrain Awareness and Warning System (TAWS), proximity to a terrain model is monitored and warnings are generated [1]. While TAWS has improved aviation safety tremendously, the FAA has deemed its products "advisory-only" primarily due to the fact that the Digital Elevation Models (DEMs) used by TAWS do not have proven integrity.

With a Synthetic Vision System (SVS), the traditional information provided on a PFD will overlay a scene depicting the location of terrain and other geo-spatial features [2]. Because of the compelling nature of an SVS PFD function, it is expected that pilots will use all of the information presented to perform the navigation function, particularly in IMC, when there is little or no visibility. This leads to a fundamental question: How can an SVS PFD function be designed such that a pilot can have confidence that the "picture" is not lying to him? In terms of avionics systems engineering, this question translates to a problem of reducing the likelihood of providing Misleading Terrain information (MTI) to flight crews to an acceptable level. A solution to this integrity problem is likely required to enable certification of an SVS PFD function as anything other than an "advisory-only" system such as TAWS.

One assumption with respect to this integrity problem should be stated here: It is assumed that overlaying high integrity PFD-like functions, such as FD guidance, on top of lower integrity data, such as a terrain model, will result in lower system integrity than a traditional PFD. When information with different levels of performance (e.g. integrity) is presented on a display, there may be cases in which the information depiction results in a conflicting image. For

example, FD guidance may intersect the virtual depiction of terrain. When this happens, human nature may lead to unexpected results; whereas without the terrain depiction, the presumed "high integrity" guidance would be followed without question. This assumption, along with the fact that there are cases where FD guidance may be unavailable, or of questionable integrity, has helped motivate the research.

Previously, the authors have presented an approach to achieving DEM integrity for SVS by implementing a real-time monitor [3][4][5]. A functional block diagram of the monitor described in [5] is shown in Figure 1. The monitor is an algorithm that consists of two parallel threads that execute in real-time providing feature information to a disparity checking function. Each thread consists of a Shadow Detection and Extraction (SHADE) function that extracts terrain-related shadow features from independent sources and translates them into a common reference domain. This enables disparity checking to determine statistical agreement. The two threads of SHADE will be discussed separately.

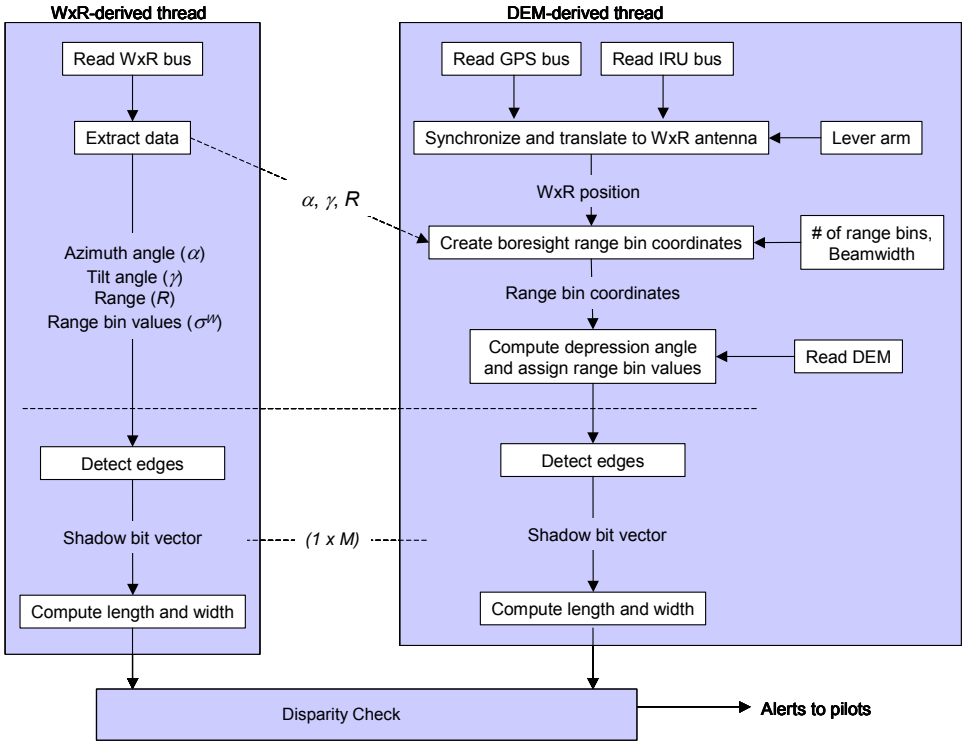

Figure 1. Integrity monitor algorithm

1.1. Terrain feature extraction from ground-mapping radar

The left thread in Figure 1 extracts terrain-related features from measurements provided by a scanning x-band weather radar (WxR) sensor located in the nose radome of the aircraft. Each radar scan, *S*, consists of *N* radial measurements, each at a different azimuth pointing direction, α_i, where $i = 1,\ldots,N$. Furthermore, each radial measurement consists of *M* range bins, or intervals, with a corresponding Z-level that represents the energy, or radar cross-section, seen by the radar in that range interval. For example, standard commercial radars used in civil aviation provide four Z-levels (0..3) across 512 range bins [6].

Range bins are identified by an integer number, *j*, and a range bin size, ΔR. The range interval covered by a range bin is $(j-1)\cdot\Delta R < R < j\cdot\Delta R$. Measurements for each WxR scan can be represented by the set given in equation (1), where $\sigma_{i,j}^{W}$

represents the received energy or cross-section for a particular radial and range bin. One WxR scan therefore represents a set of $N \cdot M$ measurements.

$$S = \{\sigma_{i,j}^{W} \mid i = 1,...,N; j = 1,...,M\} \tag{1}$$

For the current implementation, the terrain features of interest are shadowed regions. For the WxR measurements, these features are identified as those segments along each radial where all of the range bin values equal zero, indicating a reflected energy value below the noise floor of the radar. Range bins whose values are non-zero, before and after these segments, are identified as shadow edges. To improve confidence in the shadow detection, a feature is only classified as a shadow feature if both a front and back edge is detected, or if a shadow extends to the range setting of the radar.

Mathematically, the result, H_i^W, of the WxR shadow detection thread for each radial measurement *i* can be represented as (2) using Finite State Machine (FSM) notation [7]. Three states are defined for the FSM. 'A' is the initial state and is appropriate when a radial measurement begins with a series of zero's. State 'B' indicates the range bin value is greater than zero. State 'C' indicates that the range bin value is zero. State transitions are based on Condition G: range bin value equals zero. The values assigned to $\eta_{i,j}^W$ are '-3', '1', and '0', for states 'A', 'B', and 'C', respectively.

$$H_i^W = \{\eta_{i,j}^W = FSM(\sigma_{i,j}^W) \mid i = 1,...,N; j = 1,...,M\} \tag{2}$$

1.2. Terrain feature extraction from a digital elevation model

Whereas the WxR shadow detection thread can produce results directly from the WxR measurements, the DEM-derived thread requires a model of the WxR measurement mechanism. Figure 2 shows the model that has been applied for the proposed shadow detection scheme. The WxR antenna pattern is modeled using the three dB beam-width, θ^{β}, of the main lobe. It is assumed that when energy hits terrain that has any surface component facing the antenna, a portion of the energy will be reflected back to the antenna resulting in non-zero reflectivity values in the corresponding range bins. However, for regions of terrain that are hidden from the perspective of the antenna (i. e. not in line-of-sight), little or no energy will be reflected and zero reflectivity values can be expected for the corresponding range bins. For example, according to the model shown in Figure 2, zero reflectivity values in range bins 9 through 16 would be expected.

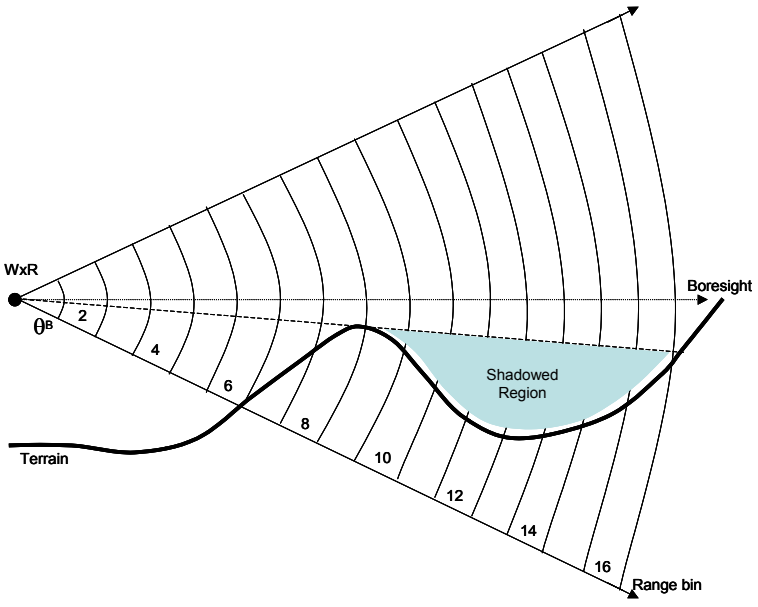

Figure 2. WxR shadow feature model

To apply the model, the proposed algorithm uses aircraft position (latitude, L^G, longitude, λ^G, and height, h^G) determined from GPS receiver measurements, aircraft attitude (pitch, θ, roll, ϕ, and heading, ψ) from an inertial reference unit (IRU), the antenna pointing direction (azimuth, α, and tilt, γ) from the WxR, the above-mentioned beam model, and a DEM. Because traditional ray tracing approaches require excessive computational cycles that make a real-time implementation problematic, an alterative method has been implemented that uses the depression angle from the aircraft to the DEM as the basis for determining the shadow locations.

The first step in finding shadow features involves the retrieval of the elevation values from the DEM at each range bin location along the bore-sight pointing direction of the antenna. *M* points are generated along radial *k* expressed in the aircraft body reference frame (*x, y, z*) and emanating from the antenna along its bore-sight pointing direction. To retrieve the corresponding DEM elevations, these coordinates are transformed to the geodetic reference frame used by the DEM (latitude, longitude, height). To accomplish this transformation, IRU attitude (pitch, roll, and heading) information is required along with aircraft position information from GPS. The origin of the aircraft body reference frame is established at the WxR antenna. For position information, this requires a translation from the GPS antenna location to the WxR antenna location. This fixed lever arm can be measured prior to flight.

Once the range bin coordinates for a radial have been computed in the geodetic frame, the corresponding terrain elevations can be retrieved from the DEM. Bi-linear interpolation [8] is performed when the coordinate lies between DEM points. Based upon the height of the WxR antenna, h_i^W, and the DEM-derived elevations along the bore-sight radial, h_j^D, the depression angles, $\phi_{i,j}^d$, from the WxR antenna to the DEM can be computed using (3).

$$\phi_{i,j}^d = -\tan^{-1}\left(\frac{h_i^W - h_j^D}{R_j}\right) \tag{3}$$

R_j is the range from the WxR antenna to the center of the j^{th} range bin and is given by $R_j = j \cdot \Delta R$.

Next, the computed depression angles are compared to the expected range of angles that define the main beam for a particular measurement. Based on our assumed shadow model (Figure 2), the range of angles where the strongest WxR reflectivity estimates are expected would be within the three dB beam-width, θ^β. This range of angles, $\phi^{dmin,0}$ to $\phi^{dmax,0}$, is γ_i minus/plus $\theta^\beta/2$. The result of the DEM-derived thread for each radial measurement *i* can be represented by (4) using Finite State Machine (FSM) notation [7]. The states listed in Table 1 are defined and state assignment is based on three conditions: (A) $\phi_{i,j}^d > \phi_i^{dmin,0}$, (B) $\phi_{i,j}^d > \phi_{i,j+1}^d$, and (C) $\phi_{i,j+1}^d > \phi_i^{dmax,0}$.

$$H_i^D = \{\eta_{i,j}^D = FSM(\phi_{i,j}^d) \mid j = 1,\ldots,M\} \tag{4}$$

Table 1. State definitions for DEM-derived radial measurements

State	$\eta_{i,j}^D$	Definition
D	**-3**	$\phi_{i,j}^d$ **outside range of antenna three dB beam-width, initialized**
E	**1**	$\phi_{i,j}^d$ **inside range of antenna three dB beam-width**
F	**2**	$\phi_{i,j}^d$ **local maxima inside range of antenna three dB beam-width**
G	**0**	$\phi_{i,j}^d$ **after local maxima and until greater than local maxima**

As discussed previously, the state assignment algorithm for the DEM-derived thread of SHADE depends directly on the computation of the depression angle. However, because the FSM is designed to search for local maxima located within the antenna beam, it is valid to omit the inverse tangent function of (3) and replace the monotonically decreasing function for depression angle with the monotonically decreasing function given by (5).

$$\delta_{i,j} = -\frac{h_i^W - h_j^D}{R_j} \tag{5}$$

Figure 3 shows an example of the shadow environment perceived by SHADE for the two threads. In Figure 3, black indicates shadowed range bins and green indicates non-shadowed range bins.

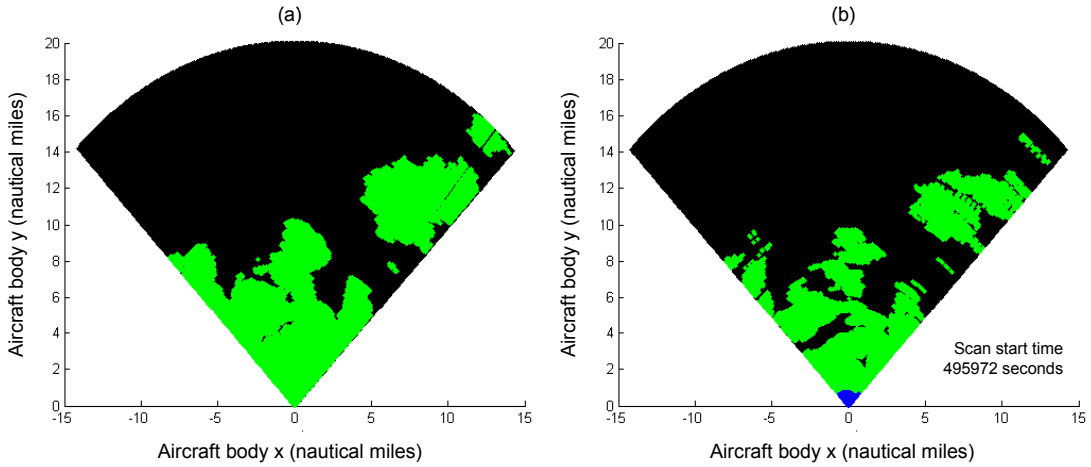

Figure 3. Example SHADE results – (a) WxR-derived thread, (b) DEM-derived thread

1.3. Disparity checking

The disparity checking function in Figure 1 consists of computing a consistency measure using the outputs of both threads and evaluating this measure against a pre-defined threshold. If the threshold is exceeded during the flight, then integrity cannot be assured within the design guideline. Operationally, it is recommended that pilots be alerted when this condition exists. This section will describe the statistical basis for choosing a consistency metric, the method for computing values for this metric, and the method for establishing appropriate thresholds.

As previously discussed, the monitor function produces feature vectors that represent terrain shadow environments that are derived from two independent sources. These vectors are in the same reference domain; therefore, ideally, differences should be zero as the sensor scans the terrain environment. However, due to nominal random error components that exist with the data sources, statistically valid inferences with respect to agreement cannot be drawn from individual point comparisons. Specifically, we know that the WxR, the GPS receiver, the IRU, and the DEM will provide data that includes a random error component. These error components will be manifested in the final products of the two threads. The monitoring approach assumes that a nominal amount of random error is acceptable for a particular operation. This leads to a null hypothesis H_0: $D \sim N(0,\sigma_D^2)$ and alternate hypotheses such as H_1: $D \sim N(\mu_D,\sigma_D^2)$ [9]. H_0 denotes that the "fault-free" behavior is characterized by disparities that are normally distributed with zero mean and a standard deviation, σ_D. H_1 represents a hypothesized "faulty" behavior that includes a bias error, μ_D. H_0 and H_1 assume that a Normal probability density function (PDF) can be used to represent the true PDF. Note, other alternate hypotheses can also be posed.

As was discussed, shadow feature radials for two independent threads are provided to the disparity checker. Shadow feature disparity, D_k, can be found radial-by-radial, using shadow center points, C_k.

$$D_k = C_k^W - C_k^D \tag{6}$$

The mean, μ_D, and variance, σ_D^2, for this disparity is given by:

$$\mu_D = \mu_{C^W} - \mu_{C^D} \tag{7}$$

$$\sigma_D^2 = \sigma_{C^W}^2 + \sigma_{C^D}^2 \tag{8}$$

Equation (10) assumes that C_k^W and C_k^D are independent and therefore, their covariance is zero. To find the PDF for C^W and C^D, we focus on individual shadow radials. For either thread, the center point is found using (9) where F_k and B_k are the front and back edge range bins for shadow k.

$$C_k = \frac{F_k + B_k}{2} \tag{9}$$

If we assume that F and B are independent zero-mean normal distributions with variances of σ_F^2 and σ_B^2, respectively, then the PDF of C can be defined by (10) with a variance given by (11).

$$C \sim N(0, \sigma_C^2) \tag{10}$$

$$\sigma_C^2 = \frac{1}{4}\left(\sigma_F^2 + \sigma_B^2\right) \tag{11}$$

Using (6), (7), (10), and (11), we arrive at the PDF for disparity, D, given in (12) where σ_D^2 is found using (8). This confirms that our null hypothesis, H_0, is valid under the specified assumptions.

$$D \sim N(0, \sigma_D^2) \tag{12}$$

Based on this derivation of the expected error distributions, a metric for quantitatively and statistically assessing disparity can be postulated. Previous work suggests the mean-square difference (MSD) of disparity is an appropriate metric [10][11][5]. Alternatively, the ratio of the number of features seen by the radar to the number of features seen by any means has been suggested as a metric [12]. In order to quantify integrity bounds in terms of PDFs and PDF parameters, the approach described in [10] is selected. The resulting disparity test-statistic, T, is given by (13).

$$T = \frac{1}{N_D}\sum_{x=1}^{N_D} D_x^2 \tag{13}$$

where N_D is the number of measurements used to establish the disparity test-statistic and D is found using (6). The variable x is used, instead of i, or k, to allow us to down-sample the raw measurements. As the horizontal beam-width of the antenna is often greater than the angular resolution, correlation can be expected among in-beam measurements.

To determine the PDF of T given equation (13) two assumptions are made: (1) the PDFs of C^D and C^W are normal with zero mean and standard deviations of σ_{C^D} and σ_{C^W}, respectively; (2) C^D and C^W are independent, uncorrelated, random variables. The first assumption follows from the derivations of the PDF for C given previously. The second assumption is based on the fact that C^D and C^W are produced using independent sources of information. The only exception is the radar pointing angle, (α, γ). Two failure modes are possible. If the mechanical scanner has failed, α will not vary. This is easily detectable and a flag can be raised. If the scanner is misaligned, the reported (α, γ) will not represent the true pointing direction of the antenna. This can also occur if the radar is gimbaled and attitude inputs to the radar are biased

or otherwise in error. In either case, both threads will still be looking at the same terrain coverage area and the terrain model for the area will be independent of the radar ground clutter sensed for the area.

Based on these assumptions, we know that the resulting PDF for T will be a Chi-squared distribution with N_D degrees of freedom [10]. Further, by following the approach presented in [10], we can now establish thresholds for T appropriate for user-defined probabilities of false alarm (P_{FA}) and missed detection (P_{MD}). Based on these configuration parameters, minimum detectable biases (MDBs) can also be computed. MDBs represent the minimum amount of error that can be observed and can be used to support the definition of safe operational limits for SVS. Table 2 provides thresholds for the test-statistic, T_D, as well as MDBs for various user-defined configurations. Note, for Table 2, for a typical WxR, an N_D of 60 corresponds to a scan range of –45 degrees to +45 degrees and back, where the sampling interval is three degrees. The two σ_D's used in Table 2 correspond to σ_C's of two and four range bins, respectively. The values computed for MDB and T_D in Table 2 are in units of range bins. Table 2 illustrates the trade-off between false alarms and missed detections. When the detection threshold is lowered, more false alarms can be expected, however, there will be fewer missed detections (i.e. higher integrity). Also, as we increase the acceptable amount of disagreement (i.e. σ_D), we will lose fidelity in our detection capability (i.e. MDB increases).

Table 2. Sample monitor configurations

Configuration	N_D	σ_D	P_{FA}	P_{MD}	*MDB*	T_D
1	60	2.83	10^{-1}	10^{-5}	3.53	74.40
2			10^{-3}	10^{-3}	3.70	99.61
3			10^{-5}	10^{-1}	3.38	118.58
4	120	5.66	10^{-1}	10^{-5}	5.64	140.23
5			10^{-3}	10^{-3}	5.92	173.62
6			10^{-5}	10^{-1}	5.43	197.83

Based on the derivation, the primary factor in determining thresholds for the proposed monitor is the distribution of nominal errors for the terrain features used by the disparity checker. Specifically, for the proposed feature of interest, terrain shadows, an estimate of the PDF for shadow center points, in the presence of only the nominal errors, is required. One method of obtaining this PDF is through calibration. This process is discussed for the two threads in the next two sections.

2. X-BAND RADAR TERRAIN FEATURE CALIBRATION

In order to determine appropriate values for σ_{C^W}, an experimental system was installed on a DC-8 aircraft. DC-8 test flights were flown from July 28 through August 4, 2003. The testing originated at Edwards Air Force Base, Edwards, CA. Experiment-specific trajectories were flown over areas near Modesto, CA, and Reno, NV. Each provided a unique terrain environment to evaluate the performance of a commercial-off-the-shelf WxR. Specifically, the WxR for the DC-8 test was a WxR-2100 model manufactured by Rockwell-Collins [13]. The WxR antenna installation on the DC-8 is shown in Figure 4(a). Interface was via the standard data bus [6].

For the WxR thread, shadow edge locations are determined by looking at the transitions from/to non-zero reflectivity values to/from zero reflectivity values. [14] provides a detailed discussion of the effects of terrain shadowing on radar sensitivity. Figure 5, from [14], illustrates how the PDF for surface reflectivity can be affected by shadow regions. In Figure 5, a measured PDF is shown for the surface reflectivity over a given spatial region. This PDF is constrained by a minimum (x1) and maximum (x6) measured clutter/signal strength. Because the noise floor and saturation ceiling can be range dependent, clipping of the PDF may be further constrained by the strongest noise level (x3) and the minimum saturation strength (x4). Somewhere between these bounds will be the actual PDF limits (x2, x5) for a particular set of measurements. Figure 5 also shows how the PDF will be affected by increasing amounts of noise, shadow, and saturation. With respect to disparity checking, increasing levels of shadowing will "push up" the left side of the PDF while "pulling down" the right side. For example, a region that is almost completely shadowed would push up the left side of the PDF to nearly 100% and the right side would flatten to near zero percent.

(a) (b)

Figure 4. DC-8 equipment - (a) WxR installation, and (b) ALTM installation

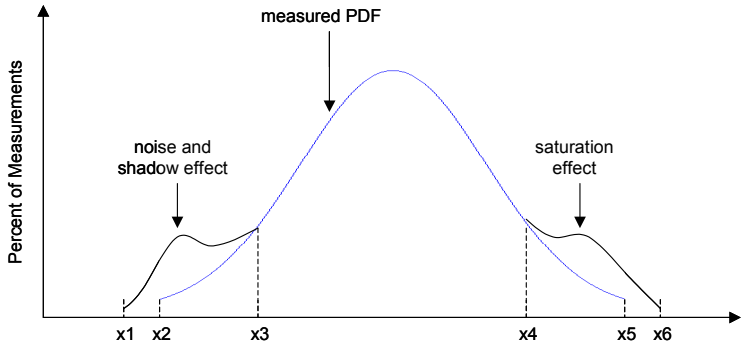

Figure 5. Theoretical WxR measurement PDF [14]

To avoid this non-deterministic reality, theoretical discussions of radar clutter PDFs tend to use clutter-only patches. Because the monitor is looking for boundaries between clutter and "noise", the PDFs derived in texts such as [14], [15], and [16] do not directly apply. In fact, it appears that PDFs for these boundaries have not been derived previously. For these reasons, WxR calibration is suggested as a method that empirically derives the nominal behavior of a radar with respect to terrain-related features such as shadowing. This procedure can also be used to test the assumption of normality previously discussed.

Data collected during the DC-8 flights has been used to compare WxR-derived shadow edge locations to reference shadow edge locations over a region of moderate terrain. Reference edge locations were established using the DEM-derived thread of the monitor, along with a high-quality reference DEM and kinematically-derived position and attitude solutions. The reference DEM was generated for a region near Modesto, California. Terrain elevations are stored in this DEM as a grid with five meter post-spacing and a vertical accuracy less than one meter (1σ). The reference DEM represents over 44 million elevation points. The data required to generate this DEM was collected on the DC-8 using an Airborne Laser Terrain Mapping (ALTM) system manufactured by OpTech, Inc. [17]. This LiDAR-based sensor was installed in the cargo bay of the DC-8 as shown in Figure 4(b).

The ALTM system consists of three sensing elements: a scanning laser, an IMU, and a GPS receiver. The laser emits 33,000 laser pulses per second while scanning +/- 20 degrees, with respect to nadir, at a nominal rate of 30 Hz. While scanning, ALTM measures the time between transmission and reception of laser pulses to determine range to the reflecting surface. The ranging error for ALTM is a function of the pulse rise time and signal-to-noise ratio and is on the

order of 2-5 cm (1σ). ALTM scan angle errors have been estimated at 0.0007 degrees (1σ). ALTM IMU attitude errors have been estimated as 0.008 degrees (1σ) for pitch and roll, and 0.015 degrees (1σ) for heading. Kinematic GPS positioning errors have been estimated as 5-15 cm (1σ) [18]. ALTM processing tools use data collected from the entire sensor suite to produce a DEM with a vertical accuracy of 15-30 cm (1σ) for all elevation values. Vertical accuracy for ALTM systems is determined by flying calibration segments over surveyed control points.

A reference trajectory for the aircraft flight path during WxR operation is also required. This reference trajectory is produced using GPS and IMU data collected by ALTM, as well as GPS base station data collected by a Cooperative Reference Station (CORS) [19]. By using the PosPac toolset from Applanix [20], a Smoothed Best-Estimate Trajectory (SBET) can be generated. This SBET provides position and attitude for the flight at 50 Hz and synchronized to GPS time. A specific flight maneuver to be used for the WxR calibration was planned for the DC-8. This maneuver was designed to bi-sect the ALTM DEM coverage so that +/- 45 degrees of WxR scanning, to a range of 10 nautical miles, would be included. The SBET generated for this maneuver is shown in Figure 6. The quality of the SBET during this maneuver is shown in Figure 7. Notice that the aircraft is flying low initially (~1000 ft above ground level), then over a reservoir, and toward, and over, moderately severe terrain. This maneuver was selected as a good candidate for capturing a representative shadow environment.

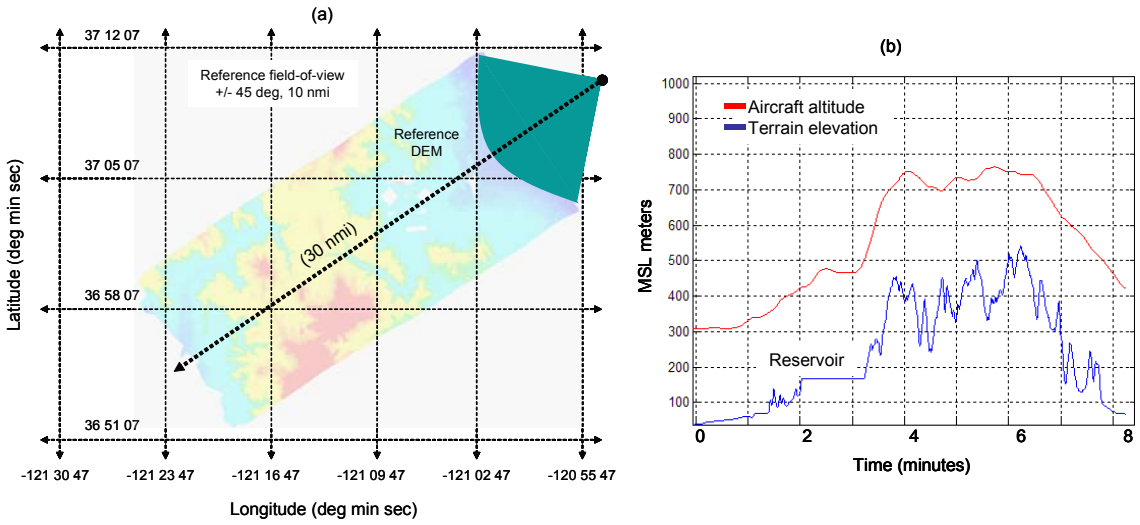

Figure 6. SBET for WxR calibration flight – (a) horizontal and (b) vertical view

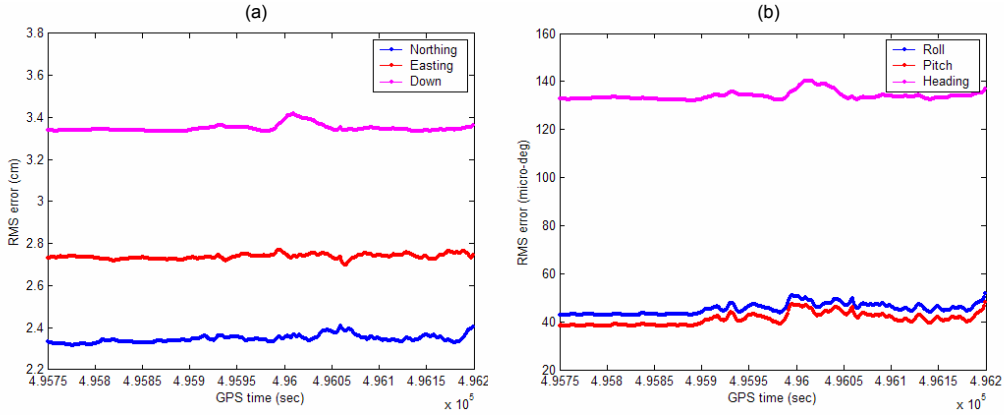

Figure 7. SBET quality during WxR calibration maneuver – (a) position, (b) attitude

For this test, the radar was configured to provide data using six bit encoding of the reflectivity values over 256 range bins and out to 20 nautical miles (nmi). This resulted in an effective range bin size of 144.8 m. The antenna scanned in azimuth over a +/- 90 degree sector at a nominal 40 degrees per second rate. Radial measurements were produced every 0.5 degrees during scanning. Although physically the antenna is continuously scanning in both directions, the WxR-2100 provides data that updates in one direction, from –90 to +90 degrees, repeatedly. The WxR accumulates and processes the raw data to produce this desired effect for the pilot's radar indicator. As is common for commercial installations, the WxR-2100 uses attitude inputs from the IRU and a gimbaling technique to stabilize the radar such that pitch and roll remain near zero with respect to local-level. Finally, and most importantly, all radar data received via the standard interface was stored along with a GPS-referenced timestamp. This was essential for the analysis in order to determine aircraft position and attitude from the SBET appropriate for each WxR measurement.

At this point, as long as we use enough independent samples, all of the required information is available to establish a PDF for shadow edges, and center points, as seen by the WxR. To this end, 40 scans have been selected that cover the +/- 45 degree sector in front of the aircraft during the low-level WxR calibration maneuver. The 40 scans encompass 450 seconds (7.5 minutes) of flight. Two scans representative of this group are shown in Figure 8 where the black areas are range bin values equal to zero, and the non-black areas are range bin values greater than zero. From these measurements, we use the WxR thread to extract all shadows seen by the radar. Similarly, we use the DEM-derived thread to extract shadows observed using our reference DEM and SBET. The result is our truth-reference for shadow features that should be seen by the WxR during the calibration maneuver. The residual difference of these two (measurement – reference) is the expected nominal shadow feature error.

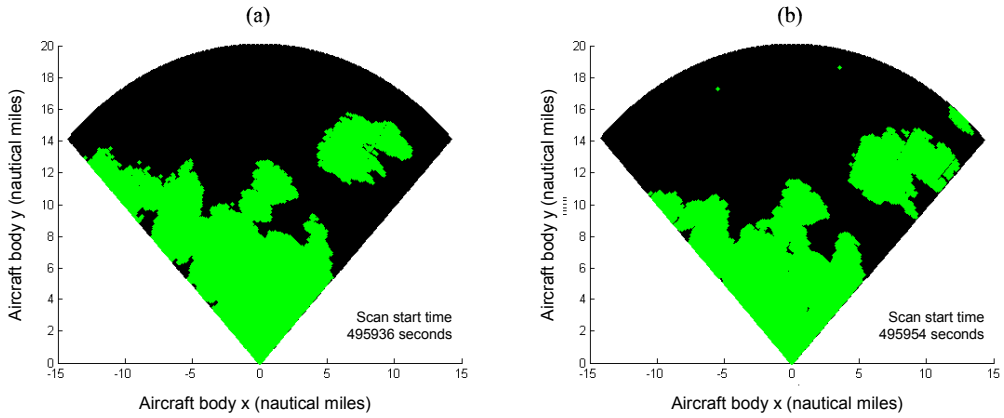

Figure 8. Selected WxR scans during calibration manuever

For example, Figure 9(a) shows the shadow center point residuals using a range constraint of five nmi. Figure 9(b) shows the distribution of these residuals. Note the PDF for the residuals appears normal. The resulting estimate for σ_{C^W} is 3.01 range bins.

As described previously, the nominal disparity variance, σ_D^2, will depend on constraints that are established for range and shadow size. To confirm this observation, Table 3 provides values for σ_{C^W} derived from the calibration data presented here and using a five nmi range limit. The relationship between σ_{C^W} and σ_D^2 is given by (8). Note. Small mean values (less than one range bin) were observed in the experimental data. These biases are most likely due to alignment errors of the WxR, as installed on the DC-8. A rigorous alignment during installation is recommended. Alternately, calibration testing, similar to what is presented here, can be used to measure these biases and remove them.

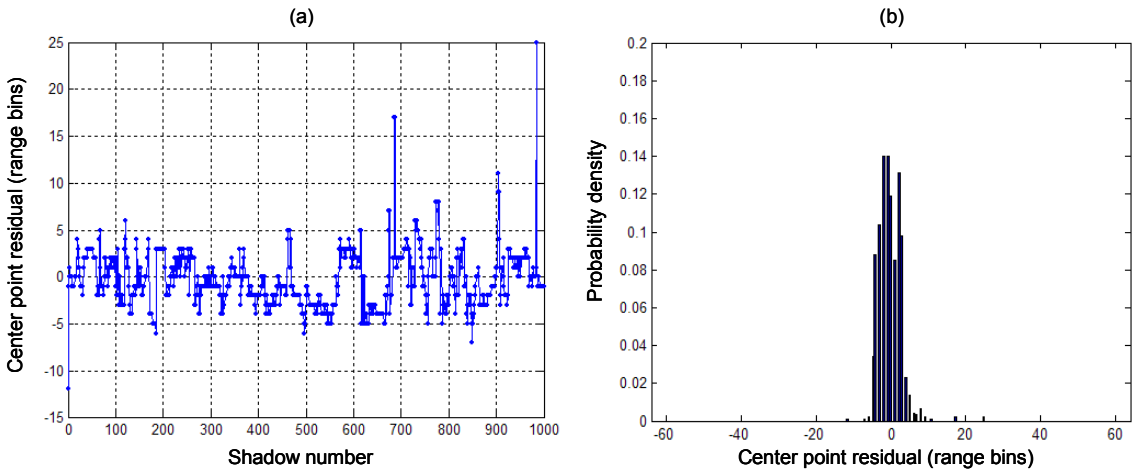

Figure 9. (a) Residuals and (b) probability distribution for shadow feature center points using five nmi range limit

Table 3. Shadow center point variability using five nmi range limit

Shadow size constraint (RBs)	σ_{C^W}
1-64	3.01
1-32	2.68
1-16	2.38
1-8	1.04

3. DIGITAL ELEVATION MODEL FEATURE CALIBRATION

To establish nominal shadow feature errors seen by the DEM-derived thread, a similar approach is taken. A reference DEM is used to represent the true height of the terrain at each post location. Shadow features are extracted using the reference DEM and reference aircraft state information. Subsequently, the reference DEM and aircraft state information are corrupted by known, nominal, random errors. Shadow features are extracted again, this time using the "noisy" versions. As with the WxR-derived shadow features, the nominal error distributions are estimated using the residual differences.

For example, a reference DEM is shown in Figure 10. This DEM has a one arc-second (~23 m) grid-spacing and is centered at the Reno-Tahoe International Airport (RNO) in Reno, Nevada. The coverage area is 100 nmi x 100 nmi. The average elevation for this DEM is 1720 m with respect to Mean Sea Level (MSL). Minimum and maximum elevations for the DEM are 1149 and 3280 m (MSL), respectively. The standard deviation of the elevations is about 73 m.

Reference aircraft state information is given in Table 4. The reference aircraft lateral position is at the RNO airport reference point. Reference altitude will be varied to show how changes in altitude affect shadow feature errors. Reference azimuthal pointing directions of the WxR are varied to cover all angles and provide sufficient measurements to establish a PDF. The error distribution assumed for the "noisy" cases are also listed in Table 4. To establish residual PDFs, a Monte Carlo simulation approach is used. Each simulation run is at a specified altitude and includes 360 azimuthal pointing directions. 256 range bins of size of 144.8 m are assumed. Range is limited to 10 nmi and only the first shadow detected along each radial is used to establish the PDFs of the residuals. Figure 11 illustrates the residual PDF for shadow feature center points using an aircraft altitude of 2000 m (MSL). Note that the PDF appears Normal. Table 5 provides the PDF parameters at various altitudes. These numbers were generated using a fixed range of 10 nmi and a minimum shadow size of three and a maximum shadow size of 126 range bins. As expected, altitude relative to

terrain, has a significant impact on the perceived shadow environment (i.e. the number of shadows decreases with increasing altitude).

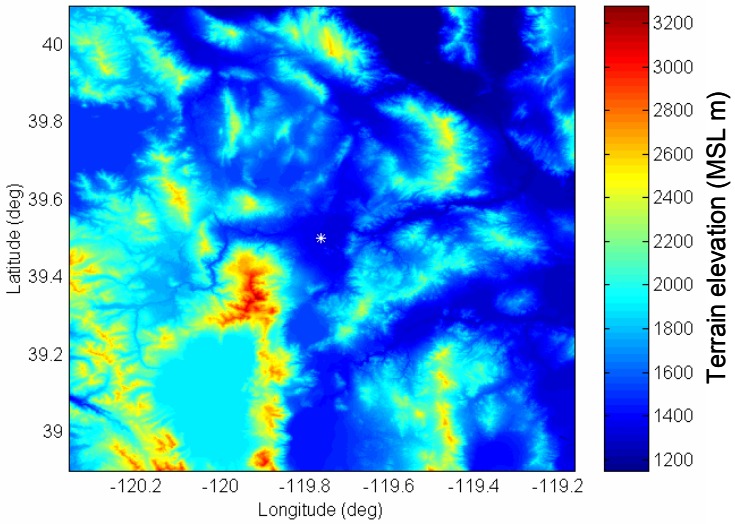

Figure 10. Reference DEM (Reno, Nevada)

Table 4. Reference aircraft state information and assumed error distributions

State variable	Reference value	Error PDF
Latitude (L^G)	39.50 deg	$N(0,20^2)$ μdeg
Longitude (λ^G)	-119.77 deg	$N(0, 20^2)$ μdeg
Altitude (h^G)	Varies	$N(0,5^2)$ m
Pitch (θ)	0.0 deg	$N(0,0.01^2)$ deg
Roll (ϕ)	0.0 deg	$N(0,0.01^2)$ deg
Heading (ψ)	0.0 deg	$N(0,0.01^2)$ deg
Azimuth (α)	[0..359] deg	$N(0,0.01^2)$ deg
Tilt (γ)	0.0 deg	$N(0,0.01^2)$ deg

Table 5. Effect of aircraft altitude on DEM-derived edge detection performance

Altitude (m)	N	μ_{E^D}	μ_{B^D}	μ_{C^D}	σ_{E^D}	σ_{B^D}	σ_{C^D}
1600	12618	-0.01	0.00	-0.00	1.80	3.74	2.09
1800	17727	-0.26	-0.59	-0.41	3.34	5.04	3.58
2000	11813	0.04	-0.18	-0.11	1.72	4.77	2.52
2200	7028	-0.14	-0.05	-0.10	2.20	2.51	1.64
2400	4890	0.21	-0.06	0.04	1.66	2.12	1.43

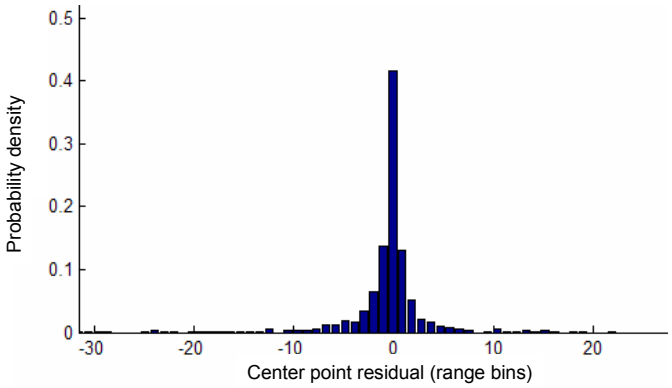

Figure 11. DEM-derived shadow feature center point PDF (Range = 10 nmi)

Another example calibration was done using the Modesto DEM, SBET, and WxR pointing angles as references and the same nominal errors for the sensors given in Table 4. The resulting σ_{C^D}'s were 4.06, 2.34, and 1.24 for range limits of 20, 10, and 5 nmi, respectively. In each case, the PDF was Normal as in Figure 11 with many more shadows used in the statistics (N > 100,000). In general, it is recommended that the calibration of the DEM-derived shadow features be based on the planned operational environment and the nominal behavior of the specific equipment installed.

Note. The PDFs that have been derived in this section will be physically constrained (i.e. discretized) by the radar interface [6]. For example, when using a five nmi range limit and a shadow minimum size of three, the domain of F will be integers in the range *(1..60)*, while the domain of B will be integers in the range *(5..64)*. The resulting domain of C will be integers in the range *(3..62)*. For all three, this constrains the domain of residuals (i.e. errors) to integers in the range *(-59..59)*. For both threads, the actual PDFs will be spread across these discrete value domains.

4. CONSIDERATIONS

For an operational implementation, the following issues should be considered:

a) Whenever multiple shadows are seen along a radial, longer shadows should have priority when computing the disparity test statistic. Longer shadows are more likely to be seen by both threads and when detected, are more likely to represent the same spatial region. In addition, longer shadows reduce the potential for spatial correlation between the front and back shadow edges. Lastly, for longer shadows, the variability of disparity will be smaller thereby leading to a smaller MDB. These observations also hold for shadow width (i.e. using wider shadows will lead to better performance). The negative effect of constraining shadow size will be reducing availability. Larger shadows may not occur in some operational environments and therefore the disparity checking function would not produce results (i.e. higher integrity).

b) Due to angular resolution, closer shadows should have priority when computing the disparity test statistic. Specifically, shadow edges seen by the radar at long range will be less accurate than edges seen at close range. This is due to the fact that the spatial volume represented by a single range bin will grow with range. As with the shadow size constraint, the negative effect of constraining range will be reduced availability in some operational environments. Short-range shadows may not always been seen.

c) If the aircraft is equipped with a radar altimeter, it may be beneficial to use this sensor to trigger operation. The operational concept suggests DEM integrity is only needed at lower altitudes. Typical operating range for commercial radar altimeters is zero to 2500 feet above ground level (AGL). In addition, the AGL height of the radar, h_{AGL}^{W}, could be used to determine range constraints for the disparity checking function. For example, if the terrain is flat, the range, R, to the point where the bottom of the radar main beam intersects the surface is given by (14). This value could be used to determine the starting range bin for shadow searches.

$$R = \mathrm{abs}\left[\frac{h_{AGL}^{W}}{\tan(\gamma - \theta^{\mathrm{B}}/2)}\right] \tag{14}$$

d) At low altitude, radar-reported range bin values that are below the noise floor are not always attributable to terrain shadowing. The most common examples are bodies of water such as lakes. For large N_D, small water-body features should not corrupt performance. However, in regions of large water bodies, a feature database containing water body boundaries may be required.

e) To improve integrity, tracking shadow features over time is being considered. This process can help to reduce the number of false alarms and missed detections that are generated due to misdiagnosis, by either thread, of a non-shadow as a shadow. The assumption is that if a feature is observed in successive scans and moving at a rate consistent with the aircraft track and speed, then the likelihood that the feature is indeed a shadow increases.

5. SUMMARY

An algorithm for detecting and identifying terrain shadowing effects, from the perspective of an aircraft flying toward and over terrain, has been specified. The algorithm has been designed to function in real-time and to provide results for measurements derived from two sources: an active radar and a stored terrain model. A statistical method for assessing the consistency of shadow features derived from both sources has also been presented. To establish quantitative agreement bounds, nominal feature error distributions have to be estimated prior to flight. Obtaining these estimates requires calibration of the radar and the DEM with respect to these features. For the radar, an example calibration using flight-test data and commercial-off-the-shelf radar and LiDAR-based sensors has been presented. For the DEM, a general approach to shadow feature calibration has been presented and illustrated by example. Ultimately, DEM shadow feature calibration should be tailored to the planned operational environment and the nominal behavior of the specific equipment installed.

This research has application to synthetic vision integrity monitoring, terrain navigation, image fusion, and terrain model accuracy assessment.

ACKNOWLEDGEMENTS

The authors would like to thank the DC-8 flight crew from NASA's Dryden Flight Research Center. Jonathon Sayre's efforts at Ohio University developing the WxR data collection software and hardware is also very much appreciated. Bill Gutelius from Optech provided much needed expertise and support while installing and operating the ALTM. Donna Gallaher from NASA Langley Research Center led the equipment installation activities on the DC-8, her talents and dedication are to be commended as well. The work presented in this paper was supported through the NASA Langley Research Center under Cooperative Agreement NCC-1-351.

REFERENCES

1. "Terrain Awareness and Warning Systems," United States Federal Register - Rules and Regulations, Federal Register Document 00-7595, Vol. 65, No. 61, pp. 16735-16756, March 27, 2000.
2. Prinzel, L. J., Comstock, J. R., Glaab, L. J., Kramer, L. J., and Arthur, J. J., "The Efficacy of Head-Down and Head-Up Synthetic Vision Display Concepts for Retro-Fit and Forward-Fit of Commercial Aircraft," *International Journal of Aviation Psychology*, Vol. 14, No. 2, 2004, pp. 53-77.
3. Uijt de Haag, M., Sayre, J., Campbell, J., Young, S., and Gray, R., "Flight Test Results of a Synthetic Vision Database Integrity Monitor," *Proceedings of the 15th Annual International Symposium on Aerospace/Defense Sensing, Simulation, and Controls*, AeroSense, SPIE, Orlando, FL, April 16-20, 2001.
4. Uijt de Haag, M., Young, S. D., Sayre, J., Campbell, J., and Vadlamani, A., "DEM Integrity Monitor Experiment (DIME) Flight Test Results," *Proceedings of the International Society for Optical Engineering (SPIE) Enhanced and Synthetic Vision 2002*, edited by J. G. Verly, Vol. 4713, SPIE, 2002, pp. 72-83.

5. Young, S., Uijt de Haag, M., and Sayre, J., "Using X-band Weather Radar Measurements to Monitor the Integrity of Digital Elevation Models for Synthetic Vision," *Proceedings of the International Society for Optical Engineering (SPIE) Enhanced and Synthetic Vision 2003*, edited by J. G. Verly, Vol. 5081, SPIE, 2003, pp. 66-76.
6. "Airborne Weather Radar with Forward Looking Windshear Detection Capability", ARINC Characteristic 708A-3, ARINC, November 15, 1999.
7. Young, S., and Uijt de Haag, M., "A Shadow Detection and Extraction Algorithm using Digital Elevation Models and X-band Weather Radar Measurements," draft submission to the International Journal of Remote Sensing, February, 2004.
8. Lim, J. S., *Two-Dimensional Signal and Image Processing*, Prentice-Hall, 1990.
9. Duda, R. O., Hart, P. E., and Stork, D. G., *Pattern Classification*, 2nd ed., John Wiley & Sons, Inc., New York, NY, 2001.
10. Gray, R. A., "In-flight Detection of Errors for Enhanced Aircraft Flight Safety and Vertical Accuracy Improvement Using Digital Terrain Elevation Data with an Inertial Navigation System, Global Positioning System and Radar Altimeter," Ph.D. Dissertation, Ohio University, Athens, Ohio, June 1999.
11. Uijt de Haag, M., Young, S., and Gray, R., "A Performance Evaluation of Elevation Database Integrity Monitors for Synthetic Vision Systems," *Proceedings of the 8th International Conference on Integrated Navigation Systems*, Saint Petersburg, Russia, May 28-30, 2001.
12. Morici, M., "Aircraft Position Validation using Radar and Digital Terrain Elevation Database", U. S. Patent Number 6,233,522, U. S. Patent Office, May 15, 2001.
13. "WxR-2100 MultiScan™ Radar", Data Sheet 523-0780587-00211J, Rockwell Collins, www.rockwellcollins.com, October, 2001.
14. Billingsley, J. B., *Low-Angle Radar Land Clutter Measurements and Empirical Models*, William Andrew Publishing, New York, NY, 2002.
15. Long, M. W., *Radar Reflectivity of Land and Sea*, 3rd ed., Artech House, Boston, MA, 2001.
16. Wehner, D. R., *High-Resolution Radar*, 2nd ed., Artech House, Norwood, MA, 1995.
17. http://www.optech.on.ca
18. Campbell, J., Uijt de Haag, M., Vadlamani, A., and Young, S., "The Application of LiDAR to Synthetic Vision System Integrity," *Proceedings of the 22nd AIAA/IEEE Digital Avionics Systems Conference*, AIAA/IEEE, Indianapolis, IN, October, 2003.
19. http://www.ngs.noaa.gov/CORS/cors-data.html
20. http://www.applanix.com/html/products/prod_pac_index.html