

The Effect of Coarse Doppler Resolution on the Performance of A Vertical Synthetic Aperture Radar Navigation System

Tucker Haydon, *Sandia National Labs & The University of Texas at Austin*

Todd E. Humphreys, *The University of Texas at Austin*

BIOGRAPHY

Tucker Haydon Tucker Haydon is a navigation research engineer at Sandia National Labs where he designs, develops, and deploys navigation systems for prototype applications. He is also a PhD student at the University of Texas at Austin where his research is focused on vertical synthetic aperture radar navigation.

Todd Humphreys Todd E. Humphreys (B.S., M.S., Utah State University; Ph.D., Cornell University) holds the Ernest Dashiell Cockrell II Chair in Engineering in the department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin. He is Director of the Wireless Networking and Communications Group and of the UT Radionavigation Laboratory, where he specializes in the application of optimal detection and estimation techniques to positioning, navigation, and timing. His awards include the NSF CAREER Award (2015), the ION Thurlow Award (2015), the PECASE (NSF, 2019), the IEEE Walter Fried Best Paper Award (2012, 2020, 2023), the ION Kepler Award (2023), and the Royal Institute of Navigation Harold Spencer Jones Gold Medal (2025). He is a Fellow of the Institute of Navigation and of the Royal Institute of Navigation.

ABSTRACT

The trusted inertial terrain-aided navigation (TITAN) algorithm correlates vertical synthetic aperture radar (VSAR) range-Doppler measurements against a digital terrain elevation model to determine the position and velocity of an airborne radar without input from global navigation satellite systems (GNSS). Recent work has characterized the navigation performance of the TITAN algorithm when the radar Doppler resolution is assumed to be infinitely fine. Hyperfine Doppler resolution, however, comes at a cost: synthetic aperture processing requires the length and number of radar pulses in an aperture to increase with finer Doppler resolutions. Consequently, the size of VSAR images, the computational resources required to process these enlarged images, and the total energy required by the greater number of radar pulses all increase. Therefore, for remote navigation systems with limited computational processing and energy budgets, it is desirable to keep the Doppler resolution of a navigation VSAR coarse, provided that this coarsening does not significantly degrade navigation performance. Existing work has not characterized how the performance of the TITAN algorithm varies with Doppler resolution, so it is difficult to determine an appropriate Doppler resolution given navigation accuracy requirements. This paper addresses this gap and analyzes how TITAN's signal-to-noise ratio, ranging errors, and navigation accuracy are affected by Doppler resolution. The study concludes that there is little correlation between overall navigation performance and Doppler resolution when the number of navigation observables is fixed.

I. INTRODUCTION

Undoubtedly, there is a recent increased interest in technologies that can bolster, cross-validate, or entirely replace traditional GNSS-based navigation systems. Radar altimeter (RA) navigation technologies constitute a family of techniques that suit this need for airborne and spaceborne systems. RA navigation predates GNSS, and the technology was one of the first to enable remote, global, unambiguous navigation (Golden, 1980). Since the Global Positioning System (GPS) became fully operational in the mid-1990s, however, interest in and development of RA navigation systems has stagnated, and the development of RA navigation technology has not kept up with the general development of radar technology. Now, with increased interest in GNSS-independent navigation technologies, there is a resurgence in the development of RA navigation systems and an opportunity to upgrade the technology to the 21st century.

One such 21st-century technology is synthetic aperture radar (SAR). SAR technology greatly matured in the past two decades, fueled by the rapid improvement in and the miniaturization of computers in the digital age. This maturation, however, occurred during the period of stagnation in radar altimeter navigation technologies, so the use of SARs for altimeter-based navigation has only recently been explored. Most exploration has focused on interferometric SARs (InSAR) (Perrett and Krempasky, 2001; Lee et al., 2011; Oh et al., 2019; Kim et al., 2018) that leverage antenna arrays to make three-dimensional range and angle measurements to overflown terrain features. While the three-dimensional measurements of InSARs are attractive as they yield

unambiguous measurements of terrain features, the resource costs of these radars are high as the number of antennas, the number of channels that must be simultaneously sampled, and the number of digital samples that must be processed by a computer are all multiplied. Recently, the trusted inertial terrain-aided navigation (TITAN) system was proposed as a radar altimeter navigation technique leveraging a standard (i.e., non-interferometric / single-antenna) vertical synthetic aperture radar (VSAR) (Haydon and Humphreys, 2023). Importantly, because TITAN does not require an antenna array, its required hardware and software are much simpler, making it an attractive candidate for remote navigation systems with limited hardware and processing budgets.

Like most other range-based terrain-aided navigation techniques, TITAN measures the range to several overflown terrain features and correlates them against a local digital terrain elevation model (DEM) to estimate the position and velocity of the radar. Specifically, TITAN takes advantage of the Doppler processing within the VSAR and measures the range to the closest points along several prescribed iso-Doppler contours (Haydon and Humphreys, 2023). Recent work has characterized the navigation performance of the TITAN algorithm (Haydon et al., 2025), but this analysis made the simplifying assumption that the SAR's Doppler resolution was infinitely fine. This hyperfine assumption is not limiting as SARs can readily produce fine-resolution images on order of 20 cm (Doerry et al., 2004), but there are good reasons not to require hyperfine Doppler resolution and, in fact, prefer coarse-resolution SAR images. Fine Doppler resolutions require long synthetic apertures that emit and record many radar pulses (Carrara et al., 1995). As a consequence and for a fixed scene size, fine resolution SAR images often contain many pixels, and digital processing of these images is expensive and cumbersome. To mitigate this, SAR scene sizes can be reduced with a narrowed radar beam width, but this is antithetical to SAR navigation systems, which seek spatially diverse features to maximize navigation information. Furthermore, as synthetic apertures lengthen, the number of pulses and thus the amount of energy required by the radar increase, which can be taxing for remote navigation systems with limited batteries. Therefore, for computational and energy reasons, it may be preferable to maintain a large scene size and coarsen Doppler resolution, provided that this coarsening does not adversely affect navigation performance. The relationship between SAR Doppler resolution and navigation accuracy in the context of the TITAN system, however, is unknown, so requirements for the Doppler resolution are difficult to derive — this is the gap that this paper addresses.

The relationship between radar altimeter resolution and navigation accuracy is largely unstudied, so there is little prior experience to draw on. In the case of simple range-only radar altimeters used by algorithms such as the terrain contour matching (TERCOM) (Golden, 1980) or Sandia inertial terrain-aided navigation (SITAN) (Hollowell, 1990) algorithms, no synthetic aperture processing is invoked. Instead, these systems typically measure the returning radar power versus time and combine it with a barometric altimeter measurement to estimate the overflown terrain elevation prior to correlation against a DEM. Range resolution certainly affects these systems, but its effect cannot be differentiated from barometric pressure altitude and reference map elevation errors (Hollowell, 1990). These errors are often larger than radar range resolution errors, so no apparent individual analysis of range resolution errors has been performed (nor is there necessarily a need).

Vertical SAR (Choi et al., 2013; Haydon and Humphreys, 2023) and vertical InSAR (Perrett and Krempasky, 2001; Jensen, 2004; Lee et al., 2011; Kim et al., 2018; Oh et al., 2019) systems have recently been proposed and studied, but they, too, have assumed radar resolution and noise values and have not investigated the relationship between navigation performance and radar resolution. Importantly, it is unclear whether radar resolution effects can be simply lumped into an inflated modeled noise term as it is in range-only radar altimetry. The TITAN algorithm, for example, does not directly process SAR Doppler measurements within a Kalman filter, and instead uses them to filter minimum-range measurements prior to their processing within a Kalman filter (Haydon and Humphreys, 2023). Thus, in this application, the Doppler resolution only indirectly influences the navigation system, and it is unclear how, or if, *Doppler* resolution should be lumped into a *range* noise term. Similarly, Choi et al. (2013) employs a frequency modulation continuous waveform (FMCW) radar and Doppler sharpening to narrow the along-track beam width of the radar and identify the range to the closest point in the narrowed beam. In this application, changes to the Doppler resolution affect the along-track beam width and thus the possible identification of the minimum-range point. Again, it is not clear how indirect effects of changes to the Doppler resolution should be accounted for in this application. The vertical InSAR techniques of Perrett and Krempasky (2001); Jensen (2004); Lee et al. (2011); Kim et al. (2018); Oh et al. (2019) all process the range, Doppler, and interferometric angles directly within a Kalman filter, so it is likely that Doppler resolution effects in these applications could be lumped into an inflated Doppler noise term, though this conjecture is unstudied.

In the case of side-looking SAR systems, range and Doppler resolutions (often called downrange and azimuth/cross-range in this domain) are key factors as they directly relate to the pixel resolution in formed optical-like images. Bevington and Marttila (1990) performed some early analysis and characterized feature pixel registration errors in range and Doppler errors as uniformly distributed within a SAR resolution cell. Most side-looking SAR navigation publications assume SAR images of a given resolution and focus their analysis on the impact of navigation errors (Lindstrom et al., 2022) or the design of the navigation algorithm (Pogorelsky et al., 2022). Because of the close relationship between side-look SAR images and optical images, pixel resolution error analysis of optical navigation systems likely applies to side-looking SAR systems (see, for example, Lerner et al. (2004)), but such pixel analysis does not translate directly into the altimeter domain. While vertical and side-looking SARs leverage the same sensor, the change in imaging geometry yields an image product that must be utilized in two completely different ways: side-looking SARs measure the ground reflection coefficient and produce optical-like images, while vertical SARs measure range and angles and produce point-cloud-like measurements. Consequently, the effect of radar resolution on the

final image product and subsequent navigation solution is fundamentally different between the side-looking and nadir-looking SAR domains.

Ultimately, the relationship between navigation performance and SAR Doppler resolution for the TITAN system is not well understood, and there is no prior work from which to draw experience — this is the gap that this paper addresses. Specifically, this paper characterizes how TITAN’s signal-to-noise ratio (SNR), ranging errors, and ultimate navigation performance vary with SAR Doppler resolution. These quantities are studied using vertical SAR images captured aboard Sandia National Labs’ testbed radar platform, and the study concludes that, for a fixed number of range-Doppler measurements, the SAR Doppler resolution has little effect on the overall navigation performance.

The remainder of this paper is structured as follows. Section II briefly reviews the TITAN algorithm. Section III reviews important concepts related to SAR Doppler resolution. Section IV describes the experimental setup. Section V analyzes the relationship between Doppler resolution and SNR, ranging errors, and modeled navigation accuracy. Finally, Section VI compares the performance of a 15-state inertial navigation system aided by TITAN measurements of varying Doppler resolution and shows that real navigation performance is largely uncorrelated with Doppler resolution.

II. BRIEF TITAN REVIEW

The TITAN algorithm is not well known, so a brief review is offered to orient the results in the following sections. More specific details of TITAN (such as the signal processing or measurement prediction models), however, are largely unimportant to this publication, so interested readers are directed to the original TITAN publication for further information (Haydon and Humphreys, 2023).

The TITAN algorithm leverages a nadir-looking vertical synthetic aperture radar and operates on range-Doppler images of overflown terrain. For navigation features, TITAN determines the range to the closest ground point along several prescribed iso-Doppler contours. That is, the algorithm prescribes a set of Doppler shifts and, for each Doppler shift, determines the range to the closest point from the set of points that all have approximately that Doppler shift. These measurements are referred to as “minimum-range, prescribed-Doppler” (MRPD) measurements. The approximate nature of the Doppler shift in the MRPD measurement definition stems from the fact that the radar has a finite Doppler resolution and can only measure the Doppler shift of ground points to within one-half of a Doppler resolution cell of their true values. The purpose of this publication, of course, is to quantify how this Doppler resolution affects navigation accuracy.




Figure 1: (a) The TITAN concept in the Cartesian domain. The three MRPD measurements lie along quasi-rectangular regions on the ground spanning the cross-track direction. (b) The TITAN concept in the radar measurement domain. The three MRPD measurements lie where the returning radar power first transitions from noise (dark region) to signal (light region) in the three prescribed Doppler bins (columns).

Figure 1a and 1b illustrate the TITAN concept in the Cartesian navigation and radar measurement domains, respectively. In these figures, three Doppler bins have been illustrated, but readers should note that the algorithm is not limited to three and, in fact, many more are used in practice. For example 11 Doppler bins were used by Haydon and Humphreys (2023), and both 11 and 21 Doppler bins are used in this publication.

TITAN correlates the MRPD measurements against a stored, on-board digital terrain elevation model (DEM) to determine the

position and velocity of the radar. The correlation happens within a nonlinear Kalman filter where realized MRPD measurements are differenced against modeled MRPD measurements predicted with prior navigation information and the DEM. Previous work has suggested using the Cubature transform (Arasaratnam and Haykin, 2009) for the computation of the Kalman filter elements as its sampling-based navigation better handles terrain nonlinearities. More specific details of the terrain correlation process and Kalman filter are not important for this publication, so interested readers are again directed to Haydon and Humphreys' work for further information (Haydon and Humphreys, 2023).

III. DOPPLER RESOLUTION

This publication analyzes the relationship between Doppler resolution and navigation accuracy, so it is important to clarify what the Doppler resolution is and what factors influence it. Synthetic aperture radar images have two important resolutions: range resolution and Doppler resolution (also known as azimuth/cross-track resolution). Range resolution, as the name implies, captures the ability of the radar to distinguish between two neighboring ground feature points with similar ranges, and it is primarily a function of a SAR's linear frequency-modulated pulse bandwidth (Carrara et al., 1995).

Similarly, Doppler resolution captures the ability of the radar to distinguish between two neighboring ground feature points with similar Doppler shifts, and it is traditionally presented as a function of the distance traveled by the radar while emitting/collecting radar pulses (the synthetic aperture length) (Carrara et al., 1995). This resolution, however, can also be related to a more convenient factor for the analysis in this publication: the total number of pulses used to form a SAR image. For most aperture lengths, the radar's pulse repetition frequency (PRF) is usually constant because it is a function of the radar's speed and beam width, which do not change much over an aperture. Consequently, as the Doppler resolution is refined and the aperture length increases, so too does the total number of pulses (same pulse rate for a longer period of time equals more pulses). Therefore, the Doppler resolution is inversely proportional to the number of pulses within a synthetic aperture.

To illustrate this relationship, consider a vertical SAR application where a radar is flying parallel to the ground and looking straight down (this is the testing configuration in the publication). For short apertures, the ground-projected Doppler resolution ρ_a (the along-track width of the quasi-rectangles in Figure 1a) is given by $\rho_a \approx \lambda / (2 \cdot \delta\theta)$, where λ is the radar signal's wavelength and $\delta\theta$ is the angle subtended by the radar while emitting/collecting radar pulses (Carrara et al., 1995). If the subtended angle is approximated as the quotient of the along-track collection distance (the aperture length) x by the ground clearance h , $\delta\theta \approx x/h$, the along-track collection distance is approximated as the product of the radar's velocity v and collection time (the aperture time) Δt , $x \approx v \cdot \Delta t$, and the collection time is approximated as the quotient of the number of pulses N by the pulse repetition frequency PRF, $\Delta t \approx N/\text{PRF}$, then the ground-projected Doppler resolution can be written as a function of the number of pulses:

$$\rho_a \approx \frac{\lambda \cdot h \cdot \text{PRF}}{2 \cdot v \cdot N} \quad (1)$$

Hence, for a fixed radar wavelength (or frequency), altitude, speed, and PRF (which is a function of speed and beam width), the Doppler resolution is inversely proportional to the number of processed pulses. Importantly, for the experiments in the following sections, the Doppler resolution can be altered simply by processing a greater or fewer number of radar pulses. That is, data can be collected over a long aperture (corresponding to a fine Doppler resolution) and subsets of the pulses can be used to form coarser resolution radar images. This is precisely the approach taken in the following sections: the number of pulses taken from the center of an aperture N is chosen to achieve a specific ground-projected Doppler resolution ρ_a .




Figure 2: A range-Doppler image formed with three different ground-projected Doppler resolutions. The 10-meter image was formed with 5x fewer pulses than the 2-meter image, and the 50-meter image was formed with 5x fewer pulses than the 10-meter image.

Figure 2 illustrates three range-Doppler radar images processed using this pulse-reduction approach and the rectangular format image formation algorithm (Carrara et al., 1995). All three images are of the same piece of terrain and they correspond to ground-projected Doppler resolutions of 2, 10, and 50 meters. There are several important observations to draw from these images. First, note that the total Doppler bandwidth does not change, only the number of pixels/columns composing the bandwidth. That is, reducing the Doppler resolution does not impact the sensor's ability to 'see' the entire beam footprint, only its ability to distinguish between neighboring ground points. Second, note that the number of range pixels/rows has not changed, only the number of Doppler pixels/columns. That is, the range resolution and ability to detect the ground has not changed due to the reduction in the number of pulses. Finally, note the total number of Doppler pixels/columns in the 2-meter and 50-meter images: 630 and 25, respectively. This reduction in the number of pixels will have an impact on the number of possible MRPD measurements — this is explored later in the paper.




Figure 3: As the Doppler resolution is coarsened, the size of Doppler bins in the along-track direction widens.

Figure 3 illustrates the effect of a coarsening Doppler resolution in the Cartesian navigation domain. Ground-projected Doppler resolution describes the along-track width of Doppler bins, and coarsening the resolution amounts to increasing this width. As the bin width increases, fewer bins are needed to span the radar beam footprint, which is fixed by the radar's beam width and altitude. This is why the number of pixels in the radar images in Figure 2 also decreases: each column represents a fixed size on the ground.

Ultimately, the question to be answered in this study is the following: does coarsening the Doppler resolution like in Figures 2 and 3 affect the overall navigation performance?

IV. DATA COLLECTION

To answer the navigation performance question, this study adopted a data-driven approach instead of a simulation-driven approach because some of the questions to be answered by the study did not have the support of baseline theory that must inform simulations. For example, it is unknown how the distribution of the ranging errors changes as Doppler resolution is coarsened, so how could one simulate it? Lacking a proper baseline theory for simulations, this study jumped right to real data collection and analysis to fill in the missing theoretical gaps itself.

The following study was performed using data captured with Sandia National Labs' testbed radar aircraft. The testbed aircraft was a Twin Otter aircraft modified with an external radome and custom gimballed radar antenna housed therein. A high-fidelity motion measurement (MoMeas) system equipped with a navigation-grade Honeywell HG9900 inertial measurement unit and a Novatel GPS receiver with real-time TerraStar-C differential corrections was used to determine the ground-truth position of the aircraft to centimeter accuracy. This MoMeas system is similar to the one published by Kim et al. (2001) but with improved hardware.

A total of 77 vertical SAR images were captured by the aircraft, which was flying level at a ground clearance of approximately 3500 meters and speed of approximately 75 meters per second. To avoid unmapped ground features that are not included in the reference maps (such as buildings), the aircraft overflew a desert region of New Mexico near Cabezon peak. Nonetheless, several power lines were overflown during the flight, which caused minimum-range detections to return early. These power lines appeared in nine images, so these images were removed from the data set, leaving 68 images for analysis. Figure 4a depicts the testbed aircraft, and Figure 4b plots the ground track.




Figure 4: (a) Twin Otter testbed aircraft. (b) Aircraft ground track near Cabezon peak, New Mexico.

The radar was operated at Ku band with a 3 dB nominal beam width of 6.41 degrees. The range resolution was 1.8 meters, and sufficient pulses were captured for every image for a ground-projected Doppler resolution of better than 2 meters. The radar was gimballed towards nadir, but was tilted up by approximately 2.4 degrees to prevent gimbal lock (hence why the images in Figure 2 are not symmetric about zero Doppler). While the nominal 3 dB beam width was 6.41 degrees, the radar transmitted and received sufficient power that reliable detections were available out to about 10 degrees, so the 2.4 degree tilt did not impact the radar's ability to detect the nadir point. Images were captured approximately every 1.5 seconds at along-track offsets of approximately 115 meters.

A 10-meter resolution digital terrain elevation model issued by the United States Geological Survey (USGS) was used for terrain correlation and to compute the measurement residuals (U.S. Geological Survey, 2024). The vertical datum of the map is the North American vertical datum of 1988 (NAVD88), which is a height-above-geoid datum, so a geoid correction using the National Oceanic and Atmospheric Administration's (NOAA's) geoid 18 model (U.S. National Oceanic and Atmospheric Administration, 2025). The resulting {latitude, longitude, height-above-ellipsoid} DEM was then cropped to the flight region, projected onto a plane with a universal transverse Mercator (UTM) projection (Snyder, 1987), and linearly interpolated to 1-meter resolution for analysis.

In the following analysis, the Doppler resolution of each radar image was adjusted by selecting a subset of logged pulses from each aperture prior to image formation. Each radar image was formed with a ground-projected Doppler resolution of approximately {2, 3, 5, 10, 15, 20, 25, 30, 35, 40, 45, & 50} meters. To accomplish this, the number of required pulses for each resolution was calculated using Equation (1), and the pulses were taken from the center of each pulse set so that the 'center' of the aperture stayed the same. Each radar image was then formed with the rectangular format image processing algorithm

(Haydon and Humphreys, 2023; Carrara et al., 1995), and a 7-point, -50 dB Taylor window was applied to the data to suppress sidelobes (Doerry, 2017). The choice to start the resolution sweep at 2 meters was driven by a desire to initially match the range resolution of the radar (1.8 meters), while the choice to end the sweep at 50 meters was due to an eventual dearth in range measurements. At 3500 meters ground clearance and with a 10-degree beam width, the beam footprint in the along-track direction only spans approximately 450 meters. At 50-meter resolution, the number of available Doppler bins is approximately 9, which is far fewer than the desired 11 Doppler bins considered in this publication. The effects of this measurement scarcity are considered in the flight data analysis in Section VI.

V. EFFECTS ON SNR, RANGING ERRORS, AND DILUTION OF PRECISION

This section analyses how coarsening the Doppler resolution affects the radar’s signal-to-noise ratio (SNR), its ability to detect the ground, and TITAN’s theoretical navigation performance.

1. Effects on SNR

When the Doppler resolution is coarsened and the number of radar pulses reduced, it might be reasonable to assume that the signal-to-noise ratio of a SAR image decreases because the total energy emitted by the radar is reduced. Surprisingly, this is not the case, and the SNR of SAR imagery is, to the first order, unaffected by Doppler resolution (Doerry, 2006). The intuition for this is as follows. While the total energy emitted by the radar does decrease with shortened apertures and fewer pulses, the size of the Doppler resolution cells on the ground in the along-track direction increases as well (see Figure 3). Therefore, the total area of each resolution cell and thus the energy reflected by each cell increases. This decrease in total energy and increase in reflected energy per cell cancel to the first order, leaving the SNR unchanged.

Yet there are two assumptions of this analysis that should be noted and checked. First, Doerry’s (2006) analysis was performed for a traditional side-looking SAR under the assumption of a large standoff distance where ground-projected range-Doppler cells are approximately compact and rectangular (preferably square). In a vertical SAR application, the Doppler cells are long and approximately hyperbolic (see Figure 1a), while the range cells are sections of an iso-range sphere that extend down from the radar and intersect the terrain below. Consequently, in a vertical SAR application, a range-Doppler resolution cell on the ground is often not rectangular, nor is it necessarily compact. For example, a range-Doppler resolution cell near a hill often returns energy from two entirely separate sides of a hill (think of the intersection points of a circle and a negative parabola). Despite this difference in range-Doppler cell shapes, the overall area of a range-Doppler cell should still increase with Doppler resolution, so SNR should still stay constant. Second, Doerry’s (2006) analysis assumed a *uniform* scattering environment, i.e., that all points on the ground reflect radar energy with approximately the same attenuation factor. In the case of *discrete* scattering environment where there are large variations in the ground reflectance coefficients, the SNR may not be independent of Doppler resolution and may be influenced by the discrete scatterers (see, for example, Raynal et al. (2012) for a study on this topic). Fortunately, natural terrain (such as that found in the desert or on the surface of the moon) can reasonably be modeled as a uniform scattering environment, and concerns related to high-variance discrete scatterers can be neglected for the purpose of this study, which is focused on natural terrain environments.

Still, it is worthwhile to check the SNR of the vertical SAR imagery and confirm that it is uncorrelated with Doppler resolution. The SNR of the 68 vertical SAR images was calculated by first labeling the ‘signal’ and ‘noise’ pixels, and then by computing the quotient of the median of the magnitude-squared of the signal pixels by the noise pixels:

$$\text{SNR} = \frac{\text{median}[\mathcal{S}^2]}{\text{median}[\mathcal{N}^2]} \quad (2)$$

Here, $\mathcal{S}^2$ is the set of the magnitude-squared of all ‘signal’ pixels, and $\mathcal{N}^2$ is the set of the magnitude-squared of all ‘noise’ pixels. The median magnitude-squared was used for the SNR computation instead of the more traditional mean to mitigate the influence of unusually bright and dark pixels that occasionally arise due to SAR image speckle noise (Carrara et al., 1995). For signal/noise labeling, the 8x median threshold suggested by Haydon and Humphreys (2023) was employed. This method determines a pixel to be a ‘signal’ pixel if its magnitude exceeds 8x the median image magnitude, otherwise it determines the pixel to be a ‘noise’.

Figure 5a plots one vertical SAR image with the ground detections marked, and Figure 5b plots the SNR for several images and the mean SNR over all images as a function of ground-projected Doppler resolution. While there are in-run variations to the SNR, it is clear that there do not appear to be any significant trends. Certainly, the SNR did not degrade by 13 dB between 2- and 50-meter resolution as one might expect if the SNR were proportional to the number of pulses. This short experiment confirms for vertical SAR images what was already known about side-looking SAR images: that the SNR and thus the ability of the radar to sense the ground is not affected by Doppler resolution.




Figure 5: (a) Image 10 and its associated ground detections using the 8x median threshold. The 8x median threshold does a sufficient job of distinguishing between signal and noise pixels, mostly due to the high SNR. (b) SNR of several sample images and the average SNR across all 68 images as a function of Doppler resolution. The SNR trends are flat, indicating that SNR is mostly uncorrelated with Doppler resolution.

2. Effects on Ranging Errors

The previous subsection confirmed that the Doppler resolution does not impact the radar’s ability to sense the ground. This subsection studies whether the resolution affects the radar’s ability to determine the minimum range and TITAN’s ability to predict the minimum range.

To this end, minimum-range measurements were *extracted* from radar imagery by selecting the smallest-range detection from every column of each image, and minimum-range measurements were *predicted* using the TITAN measurement prediction algorithm from Haydon and Humphreys (2023), the recorded position, velocity, and attitude of the radar, and the 1-meter USGS digital elevation model. The measurement prediction algorithm is briefly summarized by the following four steps: (1) the radar’s beam is projected onto the DEM; (2) the range and Doppler shift to all illuminated map points are evaluated; (3) for each Doppler bin, any illuminated points whose Doppler shift exceeds the prescribed bounds are filtered out; and (4) for each Doppler bin, of the remaining Doppler-filtered points, the point with the minimum range is selected and its range is produced.

For the measurement prediction step, the radar beam width was modeled as a 10-degree beam width, but only range detections within the nominal 6.41 degree main along-track beam were tabulated. While detections are possible out to about 10 degrees degrees, the beam pattern starts to roll off and detections are become less certain and more noisy. It is for this reason that range measurements out to the edge of the beam pattern are often not used as TITAN observables. For example, prior work restricted the region of consideration to $\pm 2.5^\circ$ on a 10-degree beam width (Haydon and Humphreys, 2023). This along-track beam width restriction was maintained across all Doppler resolutions, so the measurements for each resolution were tabulated equally.

The realized minimum-range measurements were differenced by their associated predictions to form residuals. Figure 6a depicts the minimum-range detections and predictions for a single image, and Figure 6b plots the distribution of the range residuals over all the 2-meter and 50-meter resolution images. The histograms of other swept ground-projected Doppler resolutions (3, 5, 10, 15, 20, 25, 30, 35, 40, & 45 meters) all look similar and may be characterized as unimodal and approximately zero-mean. The standard deviations of the residuals associated with all the Doppler resolutions are {2.24, 2.28, 2.32, 2.41, 2.46, 2.58, 2.70, 2.45, 2.69, 2.74, 2.84, & 2.82} meters.

There appears to be a slight positive trend to the standard deviation of the residuals: it increases almost uniformly from approximately 2.2 meters to approximately 2.8 meters between 2-meter and 50-meter resolutions. This is an increase of 0.6 meters or approximately 27%. While a 27% increase may at first sound significant, note that the Doppler resolution changed by 2400%, so the relative percentage change is small. Another perspective: the standard deviation changed by 0.6 meters over a 48-meter increase in resolution, or a 1.2 cm increase in standard deviation per unit resolution meter. While the ranging errors did appear to increase slightly with Doppler resolution, the change certainly did not change the magnitude of the errors despite the magnitude change in resolution. For this reason, it is concluded that the Doppler resolution has little effect on the ranging errors.

The intuition for the small change in ranging errors is as follows: changing the Doppler resolution and the number of pulses has no effect on the range resolution and the ability of the radar to detect the ground. That is, in each case, the radar can still detect




Figure 6: (a) The detected and predicted range values overlaid on a 2-meter and 50-meter resolution image. Although the 50-meter image is much blurrier in the horizontal direction than the 2-meter image, the ranging performance in the vertical direction remains comparable. (b) The distributions of the range residuals over all 68 images for the 2-meter and 50-meter resolution cases. The two distributions are comparable, both in shape and magnitude.

the first return with the same degree of accuracy. The range resolution of a SAR is dependent on the linear frequency-modulated pulse bandwidth, which was held constant. If the range resolution were changed, the ranging errors would be expected to increase — but this was not the case in this experiment. Instead, changing the Doppler resolution simply changes the set of points that contribute to a range-Doppler resolution cell, and this set of points can still be accurately predicted using the TITAN measurement prediction algorithm.

3. Effects on Dilution of Precision

The last two subsections considered the ability of the radar to detect the ground and the accuracy of the associated detections. This subsection considers whether the information content in the range measurements is affected by Doppler resolution. To this end, the position dilution of precision (DOP) was evaluated for each set of TITAN MRPD measurements associated with each of the 68 radar images. The position DOP was computed as follows. Let $z = h[x^*, \theta] + w$ be the nonlinear measurement model for a set of N TITAN MRPD measurements. Here, $z \in \mathbb{R}^N$ is the MRPD measurement vector, $x^* \in \mathbb{R}^3$ is the true position of the radar, θ is a fixed parameter vector that includes additional information such as the radar's velocity, attitude, beam width, Doppler resolution, and the DEM, and $w \in \mathbb{R}^N$ is zero-mean Gaussian noise with covariance $R \in \mathbb{R}^{N \times N}$. Let $H := \partial h[x^*]/\partial x^*$ be the first order partial derivative of the measurement function with respect to the position x^* evaluated at the true position. The position dilution of precision is defined by

$$(H^\top H)^{-1} = C_\ell^g \begin{bmatrix} \sigma_{AT}^2 & \cdot & \cdot \\ \cdot & \sigma_{CT}^2 & \cdot \\ \cdot & \cdot & \sigma_{VT}^2 \end{bmatrix} C_\ell^{g\top} \quad (3)$$

where $C_\ell^g \in SO(3)$ is a coordinate transformation matrix that transforms vectors from the local navigation frame with orthonormal axes defined in the {along-track (AT), cross-track (CT), vertical (VT)} directions to the global navigation frame associated with x^* (e.g., the Earth-centered, Earth-fixed (ECEF) frame), and $\sigma_{AT} > 0$, $\sigma_{CT} > 0$, and $\sigma_{VT} > 0$ are the along-track, cross-track, and vertical dilutions of precision. These DOP values quantify the expected along-track, cross-track, and vertical positioning errors per unit range noise, and they are used as proxies to evaluate the 'information content' in a set of TITAN MRPD measurements.

The DOP was evaluated for each of the 68 radar images using 11 Doppler bins distributed uniformly throughout the radar's modeled 10-degree beam width. Note that because DOP is based on models and not real data, no beam width restriction was necessary, and the full 10-degree beam width model was invoked. Importantly, the prescribed Doppler shifts associated with

the 11 Doppler bins were maintained constant while the Doppler resolution was coarsened so that the Doppler bins for each resolution case were approximately centered on the same terrain. The choice to use 11 Doppler bins was motivated by prior work, which prescribed 11 Doppler bins and showed that 11 was sufficient for reasonably accurate navigation (Haydon and Humphreys, 2023). Other than the desire to align with prior work, the choice was somewhat arbitrary, and more Doppler bins could have been used. Prior work has shown that improved accuracy is achieved with more Doppler bins, albeit with diminishing returns (Haydon et al., 2025). The goal of this section, however, is to evaluate *trends* in the DOP and not its absolute magnitude, so the total number of prescribed Doppler bins is not important as long as system observability is maintained (which it was for every case with 11 Doppler bins).

The measurement Jacobian $\mathbf{H}$ was computed using the finite-difference Cubature transform method reported by Haydon and Humphreys (2023). Briefly, it works by perturbing the radar's position by a fixed amount, predicting the MRPD measurements with the 4-step process described in Section V.2, and computing a numerical finite-difference approximation of the Jacobian. This method is preferred by the authors over an analytical approach because the TITAN measurement function depends on the terrain map, which is discrete data, and the analytical Jacobian relies on several simplifying assumptions, which are not always met in practice (Haydon et al., 2025). A 3-meter independent standard deviation error was modeled on each position axis for the purposes of the finite-difference Cubature transform method.

Note that velocity was not included in the state nor in the computation of the DOP in Equation (3). While the TITAN measurement function is technically sensitive to changes in the radar's velocity (through translations in the position of a Doppler bin on the ground), it is much less sensitive to velocity perturbations than position perturbations. That is, the dominant terms influencing the instantaneous positioning accuracy are the position errors, not the velocity errors. Velocity errors were originally modeled in the DOP computation, but often perturbations to the velocity did not change the predicted ranges, so the columns of the $\mathbf{H}$ matrix associated with the velocity errors were all zero and the inverse in Equation (3) was undefined (velocity errors were instantaneously unobservable). For these reasons (small influence & inverse issues), the velocity was removed from the DOP consideration. Readers should thus take care with absolute interpretations of the position DOP (though it should be close to the true value).




Figure 7: (a) Sample TITAN position dilution of precision curves as a function of Doppler resolution. The trends in these curves are mostly flat, indicating that DOP and Doppler resolution are mostly uncorrelated. (b) Distributions of the slopes of the TITAN DOP curves. These distributions are all approximately zero mean, indicating that the DOP is, on average, uncorrelated with Doppler resolution.

Figure 7a plots the along-track, cross-track, and vertical DOP curves for three sample images. These images were specifically chosen to illustrate in-run variations in the DOP as the Doppler resolution was changed. Overall, the DOP curves appear generally flat (indicating a lack of correlation with the Doppler resolution), however the in-run trends and the different DOP magnitudes make it difficult to make a confident determination. Some curves, such as the cross-track DOP curve of Image 08, almost appear to have a positive trend.

To better understand whether this trend is true *on average* or whether it is simply an in-run variation, a line was fit to each DOP curve and the distributions of the slopes of the lines were evaluated. Figure 7b illustrates these distributions, which include data from all 68 radar images. The median value of the slopes fit to the along-track, cross-track, and vertical DOP curves is almost zero, indicating that the positioning performance of TITAN is *on average* independent of Doppler resolution. The variance of these distributions is due to the change in ranged features as the Doppler resolution is coarsened. For example, the ranged ground point in a 2-meter-resolution bin might differ from the range ground point in a 50-meter-resolution bin, despite the fact that the two bins are centered on the same Doppler shift. This difference in ranged ground points amounts to different information extracted from the measurements, which can affect performance, both potentially improving or degrading the instantaneous navigation accuracy.

The reason for why the DOP stays the same is less clear, and it likely merits further study in a controlled simulation environment. Until such analysis is produced, two hypotheses are offered, which might seed such future study. First, some background. Prior work has shown that the position DOP is proportional to the variance in the terrain normal vectors at the ranged ground points (Haydon et al., 2025):

$$(\mathbf{H}^\top \mathbf{H})^{-1} \propto \left(\sum_{i=1}^N \hat{\mathbf{o}}_i \hat{\mathbf{o}}_i^\top \right)^{-1} \quad (4)$$

Here, $\hat{\mathbf{o}}_i \in \mathbb{S}^2$ is the terrain normal vector at the location of the i th range ground point, and the summation is over N measurements in an instantaneous TITAN MRPD measurement set. See Figure 8 for an illustration of the terrain normals. Note that the outer product of $\hat{\mathbf{o}}_i$ produces a rank-one matrix, and that it is through the summation of $N > 3$ (in the case of position only) linearly independent matrices that the full matrix is invertible. Linear independence of the individual matrices results from variance in the terrain normal vectors $\hat{\mathbf{o}}_i$, and variance in the terrain normal vectors can come from two places: (1) variance in the terrain itself, and (2) forced variance due to the TITAN measurement constraints. It is this second condition that was studied in prior work, and it was shown that the Doppler constraint (the prescribed iso-Doppler bin) shapes the statistics of the terrain normal vectors at the MRPD ground points (Haydon et al., 2025).




Figure 8: Illustration of the terrain normal vectors $\hat{\mathbf{o}}$ at three MRPD ground points.

Specifically, the terrain normal vectors in the cross-track direction are forced to point back at the radar, while the same vectors in the along-track direction are free to vary. Haydon et al. (2025) showed that this freedom-to-vary is a direct result of the Doppler constraint, which primarily applies in the along-track direction, and it is this freedom-to-vary in the along-track direction that causes TITAN to have a better along-track than cross-track positioning performance. A reasonable hypothesis for why the positioning performance should worsen, then, is as follows: as the Doppler resolution coarsens, the Doppler constraint relaxes and the terrain normal vectors in the along-track direction should be forced more often to point back towards the radar like in the cross-track direction, lessening overall variance and worsening the positioning performance.




Figure 9: Empirical cumulative probability density function of the along-track and cross-track terrain normal angles from the look vector. As Doppler resolution is coarsened, the along-track angles tend towards the look vector, but the shift is not significant. Here, significance is measured against the cross-track angles, for which greater than 90% differ by less than 5 degrees from the look vector.

Analysis of the terrain normal vectors supported this line of thinking, but the magnitude of the change appeared insufficient to significantly affect the positioning performance. Let $\hat{o} \in \mathbb{S}^2$ be a terrain normal vector and $\hat{r} \in \mathbb{S}^2$ be the unit vector from the ranged ground point to the radar (the anti-look vector), and suppose these two vectors are expressed in the local {along-track, cross-track, vertical} frame of reference. Consider the difference in along-track and cross-track angles between these two vectors:

$$\Delta\theta_{AT} := \left| \tan^{-1} \left[\frac{\hat{o}_1}{\hat{o}_3} \right] - \tan^{-1} \left[\frac{\hat{r}_1}{\hat{r}_3} \right] \right| \quad (5)$$

$$\Delta\theta_{CT} := \left| \tan^{-1} \left[\frac{\hat{o}_2}{\hat{o}_3} \right] - \tan^{-1} \left[\frac{\hat{r}_2}{\hat{r}_3} \right] \right| \quad (6)$$

Here, the subscripts indicate the scalar components of the vectors, i.e., 1 is the along-track component, 2 is the cross-track component, and 3 is the vertical component. Figure 9 plots the empirical cumulative probability distributions of $\Delta\theta_{AT}$ and $\Delta\theta_{CT}$ for the 2-, 10-, and 50-meter resolution cases. There are several things to note. First, generally, $\Delta\theta_{CT}$ is smaller than $\Delta\theta_{AT}$ — this corroborates the conclusion offered by prior work that the cross-track angle deviations from the look vector are, on average, smaller in the cross-track direction than the along-track direction. Second, while there appears to be a general monotonic compression in the angle difference in the along-track direction towards zero as Doppler resolution is coarsened, the compression does not appear to be significant, and heavy tails remain. Significance, here, is measured against the cross-track angle differences, which are smaller than 5 degrees more than 90% of the time. Even in the 50-meter resolution case, the angle differences in the along-track direction only achieve a 90% cumulative probability at approximately 10 degrees difference. The heavy tails of the along-track difference also likely contribute to the overall positioning performance: a single large angular difference can greatly improve the overall variance of a set of terrain normal vectors. Here, the first hypothesis is offered for why the DOP does not significantly change with Doppler resolution: while the along-track angles do experience some compression, periodic hard activations of the Doppler bin constraint buttress the overall angular distribution of the terrain normal vectors in the along-track direction.

The second hypothesis requires less background and can be readily stated: distribution of the Doppler bins in the along-track direction implicitly adds variance to the terrain normal vectors approximately equal to the angular distribution of the Doppler bins. That is, even if the terrain normal vectors are forced to point at the radar due to relaxed Doppler constraints, the fact that

the Doppler bins are distributed throughout the beam width automatically inflates their along-track angle variance. For example, 11 bins distributed throughout a 10-degree beam width might add approximately 10 degrees of along-track angular variance to the terrain normals. Such automatic angular distribution, however, is not possible in the cross-track direction due to the lack of a cross-track constraint.

Unfortunately, these two hypotheses must remain as hypotheses for this study, as there is insufficient data to draw firm conclusions and there are many uncontrolled variables (e.g., the structure of the terrain may confound analysis). Future work could control for confounding variables with a controlled simulation (e.g., controlled, simulated terrain) and test these two hypotheses. Until then, just the observation stands: DOP does not appear to significantly vary with Doppler resolution.

VI. EFFECT ON TITAN-AIDED INERTIAL NAVIGATION SYSTEM

Section V examined the individual effects of Doppler resolution on SNR, ranging performance, and position dilution of precision and concluded that there is little correlation between these quantities and the Doppler resolution. These are all components that can affect the distributions of the innovations and the measurement model matrices in the Kalman filter. One might conjecture, then, that the performance of a TITAN-aided Kalman filter should then be unaffected by Doppler resolution. This section considers this hypothesis and evaluates the performance of a combined inertial navigation system and TITAN-aided Kalman filter.

A 15-state combined inertial navigation system (INS) and error-state Kalman filter (ESKF) was implemented using a navigation-grade Honeywell HG9900 inertial measurement unit (IMU), a Novatel GPS receiver, and the 68 vertical SAR images. The Novatel receiver was aided by real-time TerraStar-C differential corrections, and the combined INS-GPS measurements were filtered to form the ground-truth solution against which the TITAN-aided system was graded. The 15-state INS-ESKF estimated the 3D position, velocity, and attitude of the IMU in the Earth-centered, Earth-fixed (ECEF) frame as well as a constant bias on each of the three accelerometer and gyroscope axes. The 15-state INS-ESKF is the same standard filter than was implemented in prior work and it is not the focus of this publication, so, to avoid unnecessary repetition, interested readers are directed to that publication for specific implementation details (Haydon and Humphreys, 2023).

The filter components that are worth noting are the two nonstandard tuning choices. First, although Section V.2 showed that the ranging residuals have an approximate standard deviation of 2.5 meters, the TITAN measurements in this experiment were modeled as corrupted by independent identically-distributed zero-mean Gaussian noise with a 5-meter standard deviation. The authors have found that a conservative tuning of the measurement noise helps stabilize the nonlinear filter, especially in the presence of measurement outliers (for which there are many). Like in previous work (Haydon and Humphreys, 2023), to deal with these outliers, a snapshot innovations monitor was invoked that rejected any measurement whose normalized innovation $y := y_i / \sqrt{S_{(i,i)}}$ had a magnitude greater than 1.5. Here, $y_i \in \mathbb{R}$ is the scalar measurement innovation under consideration, and $S_{(i,i)}$ is the scalar component of the innovations covariance matrix $S \in \mathbb{R}^{N \times N}$ associated with the measurement. Normally, an innovations monitor rejects measurements at a magnitude greater than 3 (Groves, 2013), but since the modeled measurement error covariance was inflated, it was deemed appropriate to reject at a lower threshold. Second, the accelerometer and gyroscope biases were modeled as Schmidt/consider states, and the elements of the Kalman gain matrix associated with these states were zeroed out prior to the computation of the state correction (Zanetti and D'souza, 2013). In the filter, these states were modeled as constants (no process noise), so to prevent potential convergence of these states to an incorrect value due to system nonlinearities and errors, these states were prevented from converging altogether. Other than these two tuning choices, the filter evaluated in this publication is identical to the one evaluated in prior work (Haydon and Humphreys, 2023).

In this experiment, the navigation system was initially bootstrapped by the differential GPS measurements, which were active for the first 30 seconds of the experiment. Then, the GPS measurements were disconnected, and the system was forward-propagated for 90 seconds using only IMU measurements. At 120 seconds into the experiment, the vertical SAR images became available and the system proceeded using only IMU and TITAN measurements. The choice to let the system run free-inertial for 90 seconds was motivated by the desire to accrue some navigation error prior to the start of TITAN aiding so that a root-mean-square (RMS) error calculation would not be influenced by differential GPS aiding immediately prior to the RMS tabulation. The choice of a 90-second window allowed the system to accrue error to the approximate TITAN steady-state error level.

This experiment was repeated using vertical SAR images reduced to the $\{2, 3, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$ -meter resolutions. In each case, 21 Doppler bins were prescribed, spaced uniformly in the along-track direction and spanning 8 degrees. Here, 8 degrees was chosen based off the 10-degree modeled beam width and buffering 1 degree on either side so as not to prescribe a Doppler value right on the edge of the detectable region. The prescribed Doppler values were chosen based on the finest-resolution image (2 meters), and they were rounded to the nearest realized Doppler bin for each of the coarser images. Any duplicates due to rounding were discarded, which resulted in fewer measurements in some images. Above 25-meter resolution, there were fewer than 21 available Doppler bins, so significantly fewer than 21 measurements were used. The effect of this measurement reduction is discussed shortly. The choice to prescribe 21 Doppler bins was based on observed filter stability issues associated with measurement outliers when the default number of 11 bins was prescribed. Bumping the

number of measurements to 21 improved overall stability, and it allowed observation into one of the drawbacks of coarsening resolution: reduced measurement availability (discussed shortly).




Figure 10: (a) Along-track, cross-track, and vertical navigation errors and their associated 3σ error covariance bounds during TITAN-aided navigation for the 2-meter and 20-meter ground-projected Doppler resolution cases. Navigation performance is comparable in both cases. (b) Root-mean-square horizontal error and the median number of Kalman-filter-ingested measurements versus the Doppler resolution. As the number of filter-ingested measurements decreases, the RMS errors increase. Note that although 21 measurements were prescribed, a handful were, on average, rejected by the innovations monitor even when sufficient pixels were available.

Figure 10a plots the navigation error and associated 3σ error covariance bounds for the 2-meter resolution and 20-meter resolution cases when TITAN is active. While there are some in-run differences, which is expected due to different noise and measurement samples, the overall performance of the two cases is comparable. The RMS errors in the {along-track, cross-track, vertical} directions are {4.5, 4.8, 0.8} meters and {3.6, 5.6, 0.6} meters for the 2-meter and 20-meter cases, respectively. Both 3σ error covariance bounds are plotted in the figure, but they are nearly identical and almost indistinguishable. This result is impressive and worth repeating: the number of pulses was reduced by 10x and therefore the amount of required energy by 10x, yet the navigation errors barely changed. Similar results held for the other resolutions between 2 and 25 meters, but something changed at 25 meters.

Starting at 25 meters, the resolution became sufficiently coarse that there were fewer than 21 Doppler bins available ($21 \cdot 25 \text{ m} = 525 \text{ m}$, which is larger than the approximate 450 m beam footprint). In these cases, all available Doppler bins were used, though the total number of available measurements dropped significantly. Figure 10b plots the horizontal RMS errors and the median number of Kalman-filter-consumed measurements as a function of ground-projected Doppler resolution. Up until approximately 25 meters, both the median number of consumed measurements and the horizontal RMS errors are approximately constant. Starting at about 25 meters resolution, the average number of consumed measurements decreases, and the horizontal RMS errors increase. But this makes sense: one should expect worse performance with fewer measurements. This is the ultimate caveat of this resolution-coarsening work. While coarsening the resolution does not appear to change the navigation performance *when the number of measurements is fixed*, if such coarsening does affect the number of available/consumed measurements, then performance can decline. This conclusion is worth repeating: the cost of coarsening resolution is a reduction in the total number of available measurements, which can affect performance if all available measurements are used; however, if only a subset of the available measurements are used (say, only 11 out of 100 available), performance is largely unaffected by reducing the total number of measurements (say, from 100 to 50).

As a result of this analysis and final conclusion, one major recommendation is offered: during system design, the number of MRPD measurements should be set to achieve some required RMS navigation accuracy, and the Doppler resolution should be set to meet this required number of measurements; once this required number of measurements is met, further refinements to the Doppler resolution add little benefit to the overall system performance while greatly increasing the required computational

expenses and radar energy.

Some final notes. This work was unfortunately unable to extend the fixed-number-of-measurements analysis out to the 50-meter resolution case due to filter stability issues related to measurement outliers. Future work might revisit this analysis with fewer Doppler bins once improved outlier prefiltering has been developed. Alternatively, this analysis could be repeated at a higher altitude or with a larger beam width so that more Doppler bins are possible at coarser resolutions. Regardless, the trends established by this study were clear: ranging errors and DOP were approximately constant out to 50-meter resolution, and RMS navigation errors were approximately constant while the number of measurements were fixed, so it is fair to conjecture that the RMS navigation errors would have been approximately constant at 50-meter resolution as well. A final interesting note: in the limit as the number of pulses approaches unity and the ground-projected Doppler resolution approaches the beam footprint, the vertical SAR approaches a standard radar altimeter. That is, if only one pulse is emitted, the MRPD measurement is equivalent to a closest-overall-range measurement, which is what a standard radar altimeter produces (Vesely et al., 2025). It would be interesting for future work to make a system comparison.

VII. CONCLUSION

This publication considered the effect of coarse Doppler resolution on the performance of the trusted inertial terrain-aided navigation algorithm. The signal-to-noise ratio, ranging errors, and position dilution of precision were all individually evaluated as a function of Doppler resolution, and they were all determined to be largely unaffected by changes in the Doppler resolution. A flight experiment was performed and the root-mean-square navigation errors were also found to be largely uncorrelated with Doppler resolution when the number of measurements is held fixed. The caveat to this analysis is that the number of available measurements is greatly reduced by coarse Doppler resolution, and this reduction in measurement availability can diminish navigation performance if it causes a reduction in the number of Kalman-filter-consumed measurements. It is recommended, then, to set the number of measurements to meet some required root-mean-square navigation accuracy, and then set the Doppler resolution to meet this required number of measurements; further refinements to the Doppler resolution offer little improvement while requiring more resources.

ACKNOWLEDGMENTS

This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan <https://www.energy.gov/downloads/doe-public-access-plan>.

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

SAND2025-11844C

REFERENCES

- Arasaratnam, I. and Haykin, S. (2009). Cubature Kalman filters. *IEEE Transactions on automatic control*, 54(6):1254–1269.
- Bevington, J. and Marttila, C. (1990). Precision aided inertial navigation using SAR and digital map data. In *IEEE Symposium on Position Location and Navigation. A Decade of Excellence in the Navigation Sciences*, pages 490–496.
- Carrara, W. G., Goodman, R. S., and Majewski, R. M. (1995). *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*. Artech House.
- Choi, S., Chun, J., Paek, I., and Yoo, K. (2013). A new approach of FMCW-DBS altimeters for terrain-aided navigation. In *Conference Proceedings of 2013 Asia-Pacific Conference on Synthetic Aperture Radar (APSAR)*, pages 214–217.
- Doerry, A. (2006). Performance limits for synthetic aperture radar – second edition. Technical Report SAND2006-0821, Sandia National Labs.
- Doerry, A. W. (2017). Catalog of window taper functions for sidelobe control. Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States).
- Doerry, A. W., Gutierrez, V. D., and Wells, L. M. (2004). A portfolio of fine-resolution SAR images. In Trebits, R., Kurtz, J. L., Appleby, R., Wikner, D. A., and Salmon, N. N., editors, *Radar Sensor Technology VIII and Passive Millimeter-Wave Imaging Technology VII*, volume 5410, pages 28 – 35. International Society for Optics and Photonics, SPIE.
- Golden, J. P. (1980). Terrain contour matching (TERCOM): A cruise missile guidance aid. In Wiener, T. F., editor, *Image Processing for Missile Guidance*, volume 0238, pages 10 – 18. International Society for Optics and Photonics, SPIE.
- Groves, P. D. (2013). *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems, Second Edition*. Artech House, Boston, London.
- Haydon, T., Huang, A., and Humphreys, T. E. (2025). Analysis of the trusted inertial terrain-aided navigation measurement function. *NAVIGATION: Journal of the Institute of Navigation*, 72(3).
- Haydon, T. and Humphreys, T. E. (2023). Trusted inertial terrain-aided navigation (TITAN). In *Proceedings of the 36th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2023)*, pages 2088–2108.
- Hollowell, J. (1990). Heli/SITAN: A terrain referenced navigation algorithm for helicopters. Technical Report SAND89-1972C, Sandia National Labs.
- Jensen, D. (2004). PTAN’s potential. <https://www.aviationtoday.com/2004/07/01/ptans-potential/>. Accessed: 2025-02-18.
- Kim, T. J., Fellerhoff, J. R., and Kohler, S. M. (2001). An integrated navigation system using GPS carrier phase for real-time airborne/synthetic aperture radar (SAR). *Navigation*, 48(1):13–24.
- Kim, Y., Park, J., and Bang, H. (2018). Terrain-referenced navigation using an interferometric radar altimeter. *NAVIGATION*, 65(2):157–167.
- Lee, D.-T., Jung, H.-S., Yoon, G.-W., Kim, D.-R., and Sun, W. (2011). Extraction of accurate 3-D ground coordinates from interferometric radar altimeter (IRA). In *2011 3rd International Asia-Pacific Conference on Synthetic Aperture Radar (APSAR)*, pages 1–4.
- Lerner, R., Rivlin, E., and Rotstein, P. (2004). Error analysis for a navigation algorithm based on optical-flow and a digital terrain map. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I.
- Lindstrom, C., Christensen, R., Gunther, J., and Jenkins, S. (2022). Sensitivity of back-projection algorithm (BPA) synthetic aperture radar (SAR) image formation to initial position, velocity, and attitude navigation errors. *IET Radar, Sonar & Navigation*, 16(2):364–378.
- Oh, J., Sung, C.-K., Lee, J., Lee, S. W., Lee, S. J., and Yu, M.-J. (2019). Accurate measurement calculation method for interferometric radar altimeter-based terrain referenced navigation. *Sensors*, 19(7).
- Perrett, M. and Krempasky, J. (2001). Terrain aiding for precision navigation in heavy GPS jamming. In *Proceedings of the 14th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 2001)*, pages 924–931.
- Pogorelsky, B. S., Zanetti, R., Chen, J., and Jenkins, S. (2022). Synthetic-aperture-radar-based spacecraft terrain relative navigation. *Journal of Spacecraft and Rockets*, 59(5):1412–1424.

- Raynal, A. M., Bickel, D. L., Dubbert, D. F., Verge, T. J., Burns, B. L., Dunkel, R., and Doerry, A. W. (2012). Radar cross section statistics of cultural clutter at Ku-band. In *Radar Sensor Technology XVI*, volume 8361, pages 295–306. SPIE.
- Snyder, J. P. (1987). *Map projections—A working manual*, volume 1395. US Government Printing Office.
- U.S. Geological Survey (2024). USGS 1/3 arc second n36w108 20240416. <https://www.sciencebase.gov/catalog/item/6627416ed34ea70bd5efaeaba>. Date Accessed: 20250902.
- U.S. National Oceanic and Atmospheric Administration (2025). GEOID18 grids for conterminous united states (CONUS). <https://geodesy.noaa.gov/GEOID/GEOID18/downloads.shtml>. Date Accessed: 20250902.
- Vesely, M., Duník, J., and Béda, T. (2025). Barometer-free aircraft terrain-aided navigation system: Prototype design and flight-test validation. In *2025 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 1371–1379. IEEE.
- Zanetti, R. and D’souza, C. (2013). Recursive implementations of the Schmidt-Kalman ‘consider’ filter. *The Journal of the Astronautical Sciences*, 60(3):672–685.