A Millimetre Wave Radar Sensor for Autonomous Navigation and Landing

Graham Brooker
Australian Centre for Field Robotics
University of Sydney, Australia
gbrooker@acfr.usyd.edu.au

Tim Carter
Australian Centre for Field Robotics
University of Sydney
t.carter@acfr.usyd.edu.au

Abstract

Flight test results of a 94GHz dual FMCW/ICW imaging radar show sufficient contrast between selected features (including runways) and their surroundings to achieve a feature based navigation accuracy superior to that offered by GPS. Suitably projected images made at short range using the radar provide sufficient resolution and contrast for use as a landing aid during adverse weather conditions. A new light weight imaging radar under construction is to be the primary sensor for a small UAV and used to continue this development program.

1 Introduction

The primary objective of any terrestrial navigation sensor is to determine unambiguously its own position in space, either in a global sense based on a specific reasonably accurate model of the earth as provided by GPS, or locally with respect to nearby landmarks or features.

The evolution of more reliable and accurate navigation techniques has been one of the primary driving forces of modern technology, from the development of the chronometer to the advent of the now ubiquitous GPS.

From their debut half a century ago, radar ground-maps have been used to help pilots navigate, however, due to the limited resolutions and low contrasts available, the images produced were often ambiguous and difficult to interpret.

In range, improved resolutions could be achieved using wide bandwidth techniques (narrow pulses or large frequency chirps), but in azimuth the limited antenna aperture available on airborne platforms was a major restriction

2 Autonomous Navigation

Since the removal of the selective availability restriction earlier this year there still remains the concern that GPS may be rendered unusable by some form of interference or jamming. For this and other reasons, airborne vehicles generally rely on a combination of GPS and inertial navigation systems (INS). This is effective for short periods but in the longer term, residual errors and biases integrate up into unacceptably large position errors.

Under these conditions, millimetre wave radar can be applied in a number of ways to improve the navigation accuracy. It can be used as a Doppler sensor to improve the velocity estimation accuracy by measuring the ground speed, or it can be used to produce high resolution images that can either be matched autonomously to features in the region for navigation purposes or to offer a pilot an image of the runway to aid with landing.

3 Resolution Specification

Table 1. Radar Resolution & Target Identification

Feature to be Resolved	Resolution
Coastlines, large cities, mountains	150m
Major highways, variations in fields	20-30m
City streets, large buildings, small airfields	10-20m
Vehicles, houses, small buildings	2-3m

A study [Moore, 1979] has shown that it is the pixel area and not its shape that defines the effective resolution (out to a 5:1 aspect ratio). It is difficult to assess the actual resolution that is required to obtain a specific navigation accuracy as the relationship is strongly dependent on the target characteristics and the correction technique used.

As an autonomous navigation and landing system for use during extremely poor visibility conditions, the radar sensor would have to resolve the features that make up an airfield (a resolution of 15m) at a sufficiently long range to correct for the possible INS drift error.

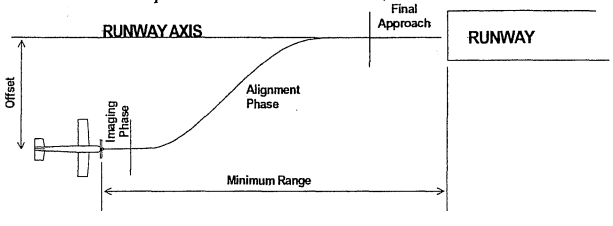

Figure 1: Plan View Showing Landing Scenario

The relationship between the maximum possible offset as determined from Figure 1 and the performance of the

aircraft in question determines the shortest acceptable imaging range. As a minimum requirement, images of the first 500m of runway and adjacent features must be produced as a coarse navigation update prior to landing.

A series of curves showing the maximum lateral displacement that can be achieved within a given range is shown in Figure 2. It is assumed that the imaging and processing to obtain a correction requires two seconds and that the aircraft must be well aligned 500m before the start of the runway.

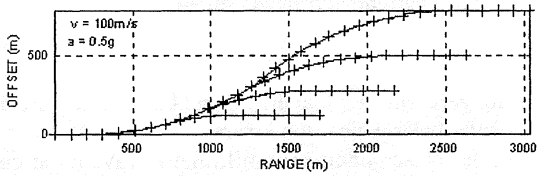

Figure 2: Maximum Displacement that can be achieved within a Given Range of the Runway

Images made during the landing phase that takes place during the last 500m would require a resolution of the order of 2 to 3m to resolve the beginning and edges of the runway. The image update rate would have to be significantly faster, up to 20Hz or more if a human pilot is in the loop.

3.1 Angular Resolution

Three different techniques [Currie and Brown, 1987] have been used to circumvent the angular resolution problem. The most obvious is to increase the aperture in the relevant axis by deploying a long thin antennas (typically 100 to 200 wave-lengths) along the length of the aircraft and using its forward motion to scan the beam across the ground, a technique known as side-looking airborne radar (SLAR). Increasing the frequency to produce a narrower beam from the same aperture is the second method that can be applied. In this instance the antenna is scanned over a limited sector about the front of the aircraft to produce an image. A third and more recent innovation [Wehner, 1987] is a technique known as Synthetic Aperture Radar (SAR) in which the forward motion of the aircraft is used to synthesise an effective aperture significantly larger than the actual antenna dimension in the direction of travel.

This paper deals with the application of the second technique to generate images with sufficiently good resolution for use as a visual landing aid and for autonomous navigation.

Table 2: Minimum Beamwidth as a Function of Frequency for a 400mm Antenna Aperture

Frequency (GHz)	Beamwidth (deg.)	Resolution (m)	
		For 15x15m equivalent at 3km	For 3x3m equivalent at 500m
10	5.15	275x0.8	45x0.2
35	1.5	80x2.8	13x0.7
94	0.56	30x7.7	5x1.8
140	0.375	20x11.5	3.25x2.75

Factors that must be considered include the available aperture for the antenna, the antenna scan angle and the size of the image that can be produced and analysed within the allocated time. For this application, the trade-off sets the minimum acceptable imaging range to be 3km with an associated lateral displacement of not more than 500m during the coarse alignment phase, and an error of less than 30m at the start of the landing phase.

Scanning requirements and various other mechanical restrictions limit the actual azimuth aperture available for the antenna to about 400mm. For a real aperture image the beamwidth and hence the cross-range resolution is then limited by the frequency of choice

To achieve both the equivalent of a 15x15m resolution at a range of 3km and comply with the aspect ratio requirement of 5:1 reduces the options to 94 and 140GHz. The former was selected as the atmospheric attenuation is lower and the technology more mature at the lower frequency.

3.2 Range Resolution

The required range resolution of 7.7m at long range and 1.8m at short range can be achieved using any of a number of techniques, each of which has certain advantages and disadvantages as outlined in Table 3.

Table 3: High Range Resolution Techniques

Technique	Advantage/Disadvantage
Short real pulse	Peak power limitations
FMCW	Linearity & antenna VSWR
Interrupted FMCW	Linearity & processing
Frequency step	Long sweep duration
Pulse Compression	Complex hardware & processing

A 7.7m range resolution requires a pulse width of about 50ns or a linear frequency chirp of at least 20MHz, For a resolution of 1.8m, this becomes 12ns and 85MHz. Solid state transmitters operating at 94GHz can produce 20W with pulse widths of less than 100ns, or 500mW continuously using a single IMPATT diode. Higher powers can be achieved at a price using a power combiner [Kuno and Fong, 1979] or using high efficiency frequency multipliers.

3.3 Waveform Selection

The maximum unambiguous PRF for a radar operating out to a range of 3km is about 50kHz which translates into a duty cycle of 0.25% and an average power (for 20W peak) of 50mW for a pulsed radar. Because this is a factor of 10 lower than the 500mW available at CW from an IMPATT oscillator, the latter is the obvious choice.

Due to temporal constraints, implementation of the stepped frequency method is not considered to be practical for this application [Currie and Brown, 1987]. This leaves only frequency modulated continuous wave (FMCW) or the interrupted version (FMICW) as potentially viable options.

The FMCW principle involves transmitting and receiving simultaneously. To perform this successfully requires extremely good isolation between the receiver and the transmitter as well as a radiated power of less than 50mW for a well designed radar.

Using the FMICW technique with a 50% duty cycle at the maximum operational range of 3km results in an average power of 250mW.

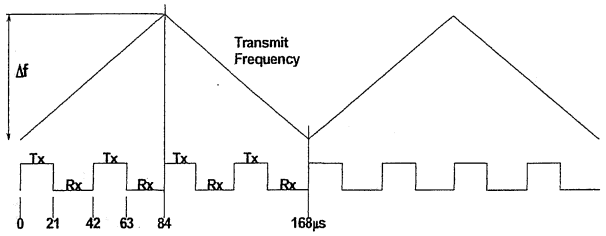

Figure 3: FMICW Timing for a 3km Operational Range

The swept bandwidth is selected such that the frequency chirp that occurs during the transmit time of 20μs exceeds 20MHz by a factor of at least 1.3 to compensate for the windowing function implemented on the received signal. A total chirp of about 140MHz in 84μs was in fact used.

In contrast, for short-range operation, round trip time restrictions, large chirp and the lower power requirements mitigate in favour of the FMCW option. The identical transmit waveform with a reduced average power of about 10mW is used. The receiver channel is open for the full 84μs.

4 Radar Configuration

4.1 Antenna

To achieve the high antenna gain critical to the performance of the radar, both the azimuth and elevation beamwidths are made as narrow as possible. Only the azimuth aperture determines the azimuth beamwidth, but in elevation the beamshape is determined from the ground range to be imaged and the height of the radar. The function that maintains constant reflected power received from each pixel irrespective of range conforms to a $\text{cosec}^2\cos^{1/2}$ gain [Skolnik, 1981]. For this application, a compromise pattern that accommodates both the long range and short range imaging requirements was defined, and a prime-focus antenna made from reinforced carbon-fibre with a layer of conductive paint was manufactured.

4.2 Front-end Overview

The hardware implementation of the dual FMCW/ICW radar front-end is shown schematically in Figure.4.

The transmit chain consists of a section that generates a sequence of linear chirps each spanning 140MHz at a centre frequency of around 94GHz. This is followed by a pulsed injection locked oscillator (PILO) that can be switched into the line using a waveguide-switch to accommodate both the FMCW and FMICW modes of operation. A standard homodyne configuration mixes the received signal with a portion of the transmitted one to generate a beat frequency proportional to the target range.

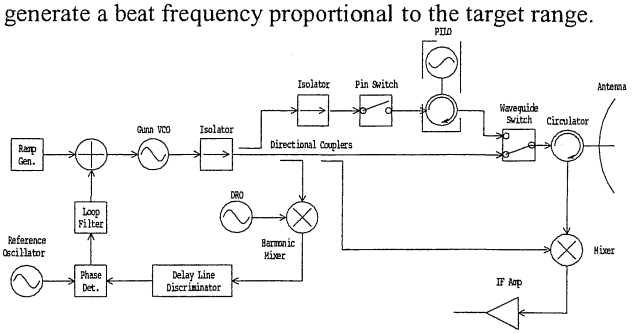

Figure 4: Block Diagram of the Radar Front-end

Because the antenna illuminates a large number of reflectors at different ranges, the received signal contains all of those frequencies with amplitudes proportional to their radar cross sections.

4.3 Chirp Linearisation

It can be shown that the range resolution is a function of both the swept bandwidth and the linearity of the chirp [Currie and Brown, 1987]. The chirp is produced by driving the varactor input of a Gunn Voltage Controlled Oscillator (VCO) with a voltage ramp that has been tailored to compensate for the gross non-linearity of the device. Linearisation of the rate-of-change of frequency is achieved by first differentiating the down converted chirp using a delay line discriminator to produce an output proportional to the chirp rate and then phase locking to a reference signal. The discriminator output spectrum is shown in Figure 5. The sharpness of the peak is an indication of the ramp linearity.

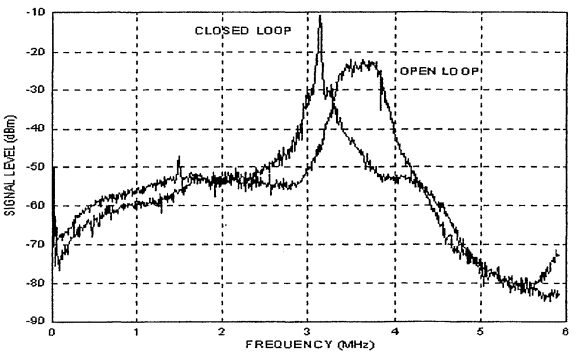

Figure 5: Delay Line Discriminator Output Spectrum showing the Effect of Closing the Linearisation Loop

4.4 Amplification for FMICW Operation

Injection locking an IMPATT oscillator involves the introduction a small signal of the required frequency into the tuned cavity containing the IMPATT diode to force the oscillator phase to match that of the injected signal. The

amount of gain that can be achieved depends on the tuning range that is required, and as the output power is fairly constant, the injected power level determines this.

Because this PILO is really an oscillator and not a true amplifier it is gated by switching the bias current.

4.5 Beat Frequency & Effective Resolution

As the range to the target decreases below 3km, the round trip time is decreased with the associated decrease in the available swept bandwidth. This leads to a degradation of the range resolution, but because of the simultaneous improvement in the cross-range resolution (due to the shorter range but constant azimuth beamwidth), the effective pixel area remains constant.

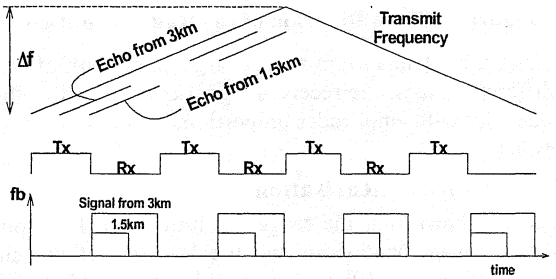

Figure 6: Beat Frequency and Gating the FMICW Radar

For short range FMCW operation this problem does not arise as the available swept bandwidth remains almost constant irrespective of the range.

4.6 Measured Front-End Performance

The performance of the front-end was determined with respect to a calibrated reflector at a known range.

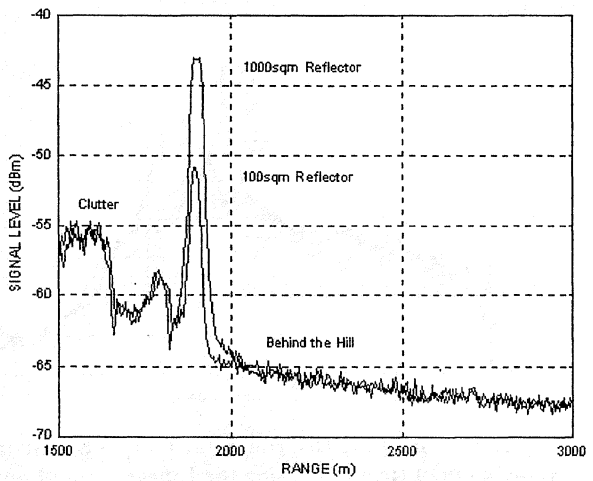

Figure 7: Measured Radar Performance for FMICW Operation at a Range of 1900m using Calibrated Reflectors

The theoretical signal level under these circumstances (1000m^2 reflector) is –40.8dBm, a closure to 1.8dB which is about as good as can be expected.

The range resolution determined from the 3dB bandwidth of the received signal was 173kHz compared to a predicted value of 77kHz, a factor of 2.25 too large. This discrepancy was traced to the PILO, which, for some as yet undetermined reason, widened, or even split the spectrum of the transmitted signal.

During FMCW operation, the PILO is not in circuit and so the system does not suffer from the same reduced range resolution.

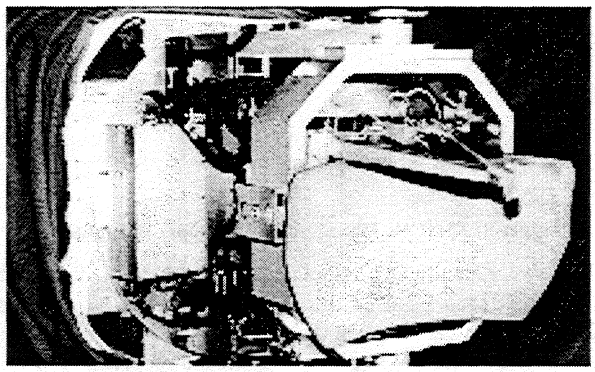

Figure 7. Photograph of the Front-end & Gimbals

4.7 Range Gating and Image Size

Operation in the 1.5 - 3 km region corresponds to a 15 - 30MHz IF. To circumvent the problems associated with sampling at 30MHz, the spectrum is divided into 12 blocks of 1.25 MHz and each spanned by 32 range gates.

The front-end is mounted on 3-axis gimbals that are stabilised. The antenna is scanned over 30° in 128 steps to build up a polar image of 384 × 128 pixels.

4.8 Spatial Integration

Image integration to reduce the effect of speckle is best performed in Cartesian space, so, as each row of 384 pixels is produced it is transformed from polar space and integrated onto this plane which has been "pinned" to the earth in the region of the target.

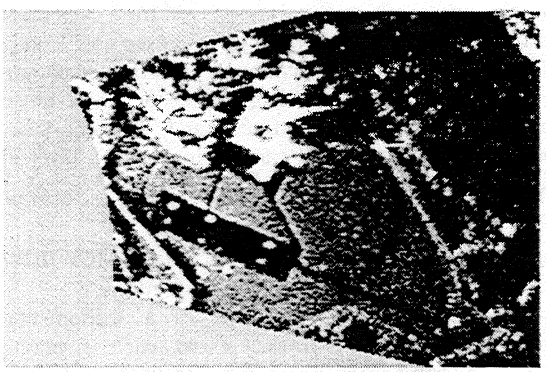

Figure 8. Long Range FMICW Image of an Airstrip

At this time due to the high processing requirements to perform the FNICW gating function a single 30° scan takes about 7s to complete. For a forward velocity of 30 m/s this

results in a displacement of only 200m on average before the same slice is re-illuminated. Is was found that 3 to 4 scans produced the best quality images.

The image section shows an airfield being approached at a slight angle to the long axis of the runway. The main features that are visible are a section about 130m wide and 500m long followed by a narrow section about 25m wide and 600m long. Two taxiways lead towards a cluster of hangars that are surrounded by a concrete apron. A wall and clusters of buildings are visible at the narrow end of the image.

The limited dynamic range of the printing process requires that the image be clipped with resultant loss of resolution for the bright objects (particularly the hangars.)

4.9 Short Range Operation

The signal processing approach adopted for short range FMCW operation more simple as the beat frequency at a range of 500m is only 2.5MHz. This allows direct sampling of the IF signal at 6.5MHz followed by windowing and a 512 point FFT to produce 256 range gates each with a range resolution of just over 2.5m.

The image data is logged in polar space for off-line processing and a subset superimposed in real time on a video image of the same scene as a pilot aid. Figure 10 shows a runway image made from the ground as the data from these flights has been lost.

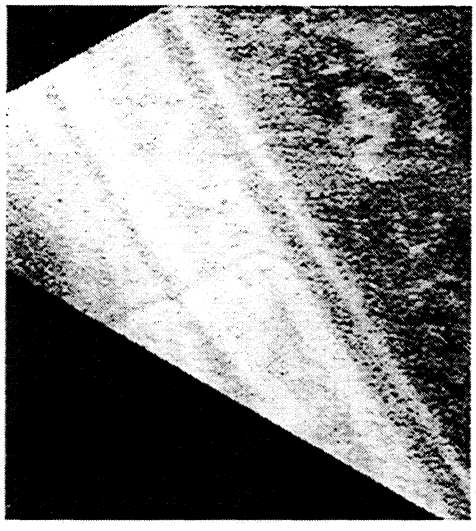

Figure 9. Short Range FMCW Image of an Airstrip

. The image shows a section of a tarmac runway about 15m wide and about 300m long running parallel to a narrow taxi-way and connected to it by an even narrower strip of tarmac

4.10 General Image Characteristics

These are some of the characteristics that are common to all mmWave images made at low grazing angles.

a) Smooth horizontal surfaces (runways etc.) are specular and reflect most of the power away, so appear dark.
b) Grass and cultivated fields scatter and show texture.
c) Scrub and trees are good scatters so appear bright.
d) Large trees, walls, banks etc. cast deep shadows.
e) Buildings and vehicles are usually good reflectors.

From a navigational perspective, large buildings such as hangars, or closely spaced clusters of buildings are useless, as it is not possible to predict exactly which part of the distributed structure is the primary reflector. Ideal candidates are groups of well-spaced point targets or homogenous structures sufficiently large to cover many pixels such as lakes or runways that are embedded in a uniform "sea" of vegetation.

5 Navigation Process

The image is built up on a Cartesian frame pinned to a specific latitude and longitude as determined by the GPS/INS based estimate of the radar position. Because of possible uncertainties in this estimate caused by the lack of GPS data, errors of up to 500m must be catered for.

The airfield area is first analysed using aerial photographs and any other available data to produce a series of features that will probably be visible on a millimetre wave radar image. These include point reflectors with dimensions larger than the resolution of the radar that are usually made up of an ensemble of closely packed reflectors, lines which generally consist of rows of point reflectors closely spaced and finally edges where a change of contrast occurs. The individual reference features are reduced to simple polar co-ordinate based structures defined relative to a fixed point on the image and stored in a database for the area.

The radar image which has been normalised by the pixel integration count is scaled and clipped to enhance the contrast for low signal to noise ratio areas before further processing. It is then passed through a filter matched to enhance the contrast of a specific feature type, before being re-normalised to allow for histogram based binarisation. Finally the binary image is processed to produce fact features which conform to the same polar format as the reference features do.

5.1 Navigation Error Estimation

Each reference feature of a specific type is compared to similar fact features within a given radius and with similar characteristics (length, orientation, area, etc.) to determine whether they match. A list of possible matches for each feature is stored.

The distance between all pairs of reference features and the distance between pairs of all their possible matches is then compared and rated according to some tolerance criteria. The rating of each fact feature accumulates asymptotically towards unity during the whole search process and finally the fact feature with the highest score is considered to correspond to its associated reference feature.

The distances between all the references and their associated fact features are then combined in some optimal

sense to produce an estimate of the navigation offset with respect to the GPS/INS based position estimate.

6 Results

The system performance was measured by flying in the direction of the runway, but offset from it and performing the image feature extraction and navigation processes before landing at a known reference. The position estimate generated by the GPS/INS system and that generated by the radar were logged. This process was repeated as often as possible to obtain a statistically significant sample, as shown in Figure 10:

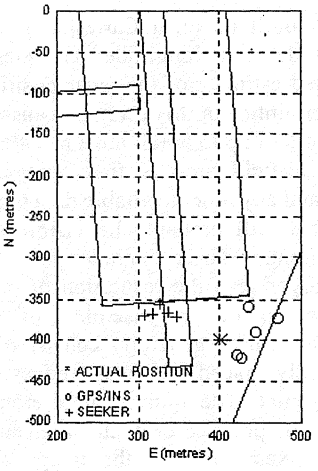

Figure 10. Measured Results of the Navigation Performance of the Radar

These results assume that there is no drift in the navigation system between making an image and landing which is patently not the case, also there were some scaling and alignment errors between the navigation system and the radar system that resulted in an offset. However, the standard deviation (σ) of the errors in each of the two tests was better than that provided by the navigation pack alone as shown in Table.3.

Table 3: Radar Image Based Navigation Correction Results.

Test flight	GPS/INS Navigation Error		Radar Navigation Error	
	σ East (m)	σ North (m)	σ East (m)	σ North (m)
1	21.7	26.6	15.8	7.6
2	29.8	43.2	20.1	20.0

7 Work in Progress

A mirror scanner attached to lightweight gimbals shown in Figure.12 is under construction at the Australian Centre for Field Robotics (ACFR). This scanner will accommodate either a millimetre wave radar or a laser front-end and will be flown as the primary sensor on an Unmanned Airborne Vehicle (UAV).

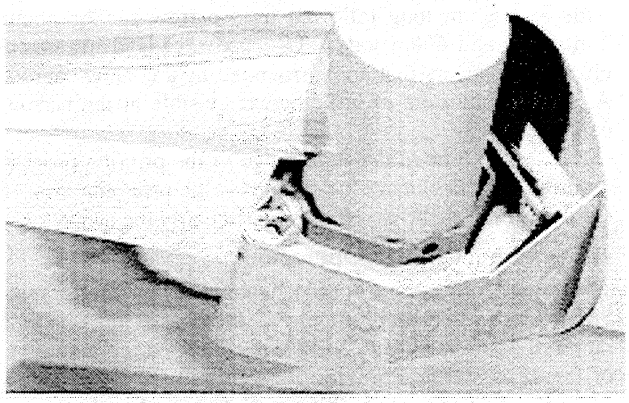

Figure.11 Radar/Laser Scanner for the UAV

In addition to its use as an imaging sensor to identify features as part of a BAe sponsored program to investigate Simultaneous Localisation and Mapping (SLAM), it is proposed to use the sensor to further investigate the use of radar vision as a landing aid during adverse weather conditions, and ultimately as the vision sensor for autonomous landing.

Acknowledgement

I would like to thank the people at AMS in South Africa who helped build and test the original radar system and for Prof. Hugh Durrant-Whyte for inviting me to come and work in Australia.

References

[Wehner, 1987] Wehner, D. *High Resolution Radar.* Artech House. Inc. 1987.

[Currie and Brown, 1987] Currie, N. and Brown, C. *Principles & Applications of Millimeter-Wave Radar.* Artech House, Inc. 1987.

[Moore, 1979] Moore, R. "Trade-off between picture element dimensions and non-coherent averaging in a side-looking airborne radar." *IEEE Trans. on Antennas and Electronic Systems,* Vol AES-15, No 5, September 1979.

[Skolnik, 1981] Skolnik, M. *Introduction to Radar Systems.* McGraw Hill. 1981.

[Kuno and Fong, 1979] Kuno, H. and Fong, T. "Solid State mmWave Sources and Combiners," *Microwave Journal,* June 1979.