Long Range, Low SWaP-C FMCW Radar

William A. Lies, Lakshay Narula, Peter A. Iannucci, and Todd E. Humphreys

Radionavigation Laboratory
The University of Texas at Austin
Austin, TX, USA

Abstract—A method is developed, analyzed, and tested to adapt low-cost, automotive-grade radar chipsets for long-range sensing. These disruptive chipsets offer impressive performance at low size, weight, power, and cost (SWaP-C) that could benefit applications with tight SWaP-C budgets such as urban air mobility and urban air logistics. The short range of these radars currently prevents their deployment in long-range applications, so this paper employs extended measurement intervals coupled with sophisticated signal processing to significantly extend their range. After deriving the optimal maximum likelihood estimator, the paper presents suboptimal, more efficient techniques for target range estimation that are robust to target motion uncertainty. These techniques are validated in simulation and demonstrated via experiment. The results show that low SWaP-C radar chipsets are capable of operating at low SNR to perform long-range sensing when augmented with this paper's signal processing techniques. This potent combination of low SWaP-C hardware and advanced signal processing will drive innovation in urban air mobility, urban air logistics, and other areas in need of long-range sensing.

Index Terms—FMCW; urban air mobility; model reduction; maximum likelihood; low SWaP-C; low-cost sensing; long-range sensing; range bin migration.

I. INTRODUCTION

The recent emergence of low-cost radar chipsets presents new opportunities in the realm of low-footprint sensing. These monolithic microwave integrated circuits (MMICs) achieve impressively low size, weight, power, and cost (SWaP-C), though they have shorter range than their high SWaP-C counterparts. For example, Analog Devices offers a chipset for the 24 GHz band which costs about $70 in volume [1]. These MMICs are targeted toward the automotive industry, where relatively short range (on the order of $100\,\mathrm{m}$) is adequate. To exploit this disruptive technology for high-sensitivity ranging of distant targets, this paper leverages sophisticated signal processing and long dwell intervals that boost the radar's signal-to-noise ratio (SNR) for distant targets. These techniques are developed with urban air mobility (UAM) and urban air logistics (UAL) in mind, though they are directly applicable to any frequency-modulated continuous-wave (FMCW) radar platform.

UAM and UAL are the subject of an intense worldwide research and development effort. Existing aerospace companies and start-ups alike are developing small, vertical take-off and landing aircraft (VTOLs) with the goal of offering short-distance air transportation as an alternative to ground transportation for both people (UAM) and products (UAL). These aircraft will eventually be autonomously piloted to reduce cost [2], [3].

Detect and avoid (DAA) is an essential capability for achieving large-scale UAM and UAL [4]. Autonomously-piloted aircraft must be able to reliably see and avoid airborne objects such as conventional aircraft, VTOLs, small delivery drones, and even large birds. This capability is essential when a threatening object is not broadcasting its position (non-cooperative) or during a malfunction of cooperative collision avoidance systems. DAA systems must match or exceed a human pilot's ability to visually avoid obstacles [5], [6]. Current research is focused on camera-based systems, attractive for their low SWaP-C. The fundamental limitation of such visual DAA systems is their inability to accurately resolve range. Thus, cameras alone are not sufficient for reliable DAA. A capable radar system with low SWaP-C would provide a fundamental complement to visual sensing for autonomous DAA, since radar has excellent range discrimination.

A. Related Work in DAA for UAM

A growing literature explores camera-based solutions for DAA [7]–[11]. However, much of the literature does not consider the problem of range estimation; it focuses instead on improving the detection and false alarm probabilities of the vision algorithms. The solutions for camera-based range estimation that do exist either perform poorly or impose unrealistic constraints:

1) Contrast-to-Noise Ratio: The authors of [12] propose a DAA solution that leverages an array of narrow-field-of-view cameras. It uses a time series of contrast-to-noise ratios to estimate range, but the errors appear to be at least $\pm 1\,\mathrm{km}$, which is unacceptably large [12]. Additionally, this method can be expected to be sensitive to atmospheric conditions such as fog and precipitation.

2) Stereo Camera Triangulation: Accurate depth sensing is difficult to achieve at long range because camera separation is limited to the size of the vehicle, which introduces severe geometric dilution of precision at long range. This difficulty is compounded by airframe flexibility [13]. Consider, for example, stereo cameras separated by $10\,\mathrm{m}$ with a focal length of $4\,\mathrm{mm}$ and 4K resolution (3840×2160). Assume a target at $500\,\mathrm{m}$. For these conditions, the range error per degree of camera misalignment is $440\,\mathrm{m/deg}$, not accounting for noise in the cameras' bearing estimates. These properties prompt the authors of [13] to dismiss stereo cameras as infeasible for DAA.

3) Bearings-Only 3D Location Estimation: Range can be observed under bearings-only tracking if the ego-aircraft "out-maneuvers" the target. But the large maneuvers required of

the ego-aircraft, and the long filter convergence time (5-30s), render bearings-only unacceptable as a primary tracking method [13].

4) Apparent Size: Estimating range based on an object's apparent size in an image requires knowledge of the physical size of the target object, which is generally not available [12], [13].

B. Deficiency in Existing Radar Solutions

Although radar would be useful in a variety of sensing applications, the poor SWaP-C characteristics of traditional systems prevent its widespread use. UAM and UAL are among those hesitant to adopt commercial radar, which is particularly SWaP-C prohibitive [13]. For example, the EchoFlight Airborne Radar from Echodyne costs $20,000 per unit, consumes 40 W, and has a mass of 800 g [14]. Moreover, three or four separate units would be required to obtain an adequate field of view on an aircraft. Thus, even if volume pricing reduced the unit price significantly, such products would impose a prohibitive strain on a VTOL's SWaP-C budget. A long-range radar solution with low SWaP-C is highly desirable for UAM, UAL, and a variety of robotics and surveillance applications.

C. Adapting Low SWaP-C Radar for Long Range Applications

Due to their low power, radar systems built around low SWaP-C radar MMICs typically have ranges only up to 200 m when operating in a traditional scanning mode. It is possible to increase their range sensitivity beyond 1 km with greater transmit power and application of weak-signal correlation techniques in which the frequency-modulated continuous wave (FMCW) radar takes measurements over an extended interval (e.g., up to two seconds) at a chosen azimuth and elevation. Such measurement intervals make traditional scanning too slow for many applications, but can be effectively used to determine range to a target whose azimuth and elevation is already known.

This paper proposes a two-step process for low SWaP-C radar range extension. First, a machine vision algorithm operates on visible light images to detect a target of interest and determine its azimuth and elevation relative to the ego-vehicle. Second, radar returns arriving from the detected direction are constructively combined to obtain a range estimate. The ego-vehicle then has enough information to infer the relative position of the obstacle. In the following sections, the signal processing of this technique is described in detail and developed for robustness against target motion uncertainty. Experimental results with a low-cost radar chipset demonstrate the effectiveness of these techniques.

D. Contributions

This paper's contributions are the concept, theory, implementation, and experimental testing of MMIC-based low SWaP-C radar for long-range sensing. After an exposition of FMCW radar and related challenges, Section V derives and analyzes the optimal maximum likelihood estimator for FMCW radar systems. Section VI presents two parametric models for target motion which guide the paper's search algorithms. Section VII develops computationally efficient, practical radar processing algorithms. The simulation results in Section VIII reinforce the preceding theory and show the efficacy of the methods from Section VII. Finally, Section IX demonstrates the viability and performance of low SWaP-C radar by way of a small-scale experiment. In this way, the paper lays the groundwork for the extension of emergent low SWaP-C radar MMICs to long-range sensing applications.

A preliminary version of this paper was published in [15]. The current version significantly expands the estimation discussion, introduces new methods for carrier phase modeling, and accounts for range bin migration in the FMCW processing.

II. FMCW RADAR OVERVIEW

FMCW radar determines range to an object by emitting and receiving "chirp" signals, which are sinusoidal waveforms whose frequency increases linearly with time. While transmitting a chirp, the radar simultaneously receives the echo from that chirp. Because the chirp frequency increases linearly with time, the difference in frequency of the transmitted and received signals is proportional to the time-of-flight of the electromagnetic wave. This difference is therefore proportional to the distance to the reflecting object. This is illustrated in Fig. 1.
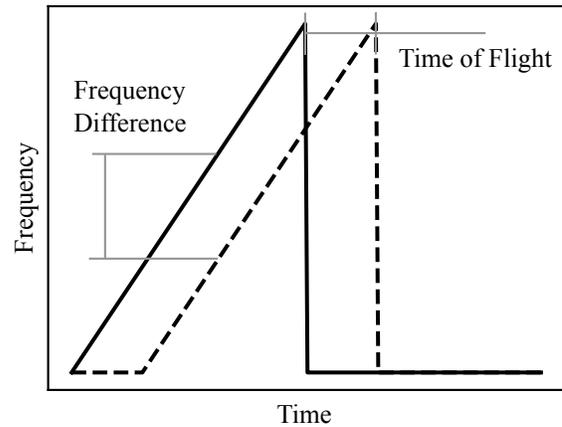
Fig. 1: FMCW chirp (solid) with corresponding echo (dashed). Longer time-of-flight corresponds to a greater frequency difference between the transmitted and received signals.

Ignoring amplitude factors and noise, the transmitted and received signals for a single target are given by

$$s_{\text{TX}}(t) = \cos\left(\frac{\pi\beta}{T_c}t^2 + 2\pi f_0 t\right) \tag{1}$$

$$s_{\text{RX}}(t, t_d) = \cos\left(\frac{\pi\beta}{T_c}(t - t_d)^2 + 2\pi f_0(t - t_d)\right) \tag{2}$$

where

t_d is the round-trip time-of-flight to a single target (s)
T_c is the chirp duration (s)
f_0 is the nominal carrier frequency (Hz)

β is the chirp bandwidth (Hz)

To measure the frequency difference depicted in Fig. 1, the transmitted and received signals are mixed and lowpass-filtered to obtain the so-called intermediate frequency (IF) signal. Many FMCW front-end systems only handle the in-phase (real) baseband component. But, as explained in [16], a complex downmixing and sampling architecture avoids the aliasing of image frequency noise into the signal passband, improving the signal-to-noise ratio compared with a real-valued downmixing and sampling architecture. This paper assumes such a complex downmixing architecture, in which case, a single target's baseband signal is given by

$$s_{\mathrm{bb}}(t, t_d) = \exp\left(j2\pi\left(\frac{\beta t_d}{T_c}t - \frac{\beta t_d^2}{2T_c} + t_d f_0\right)\right) \quad (3)$$

where j is the imaginary unit.

The baseband signal is sampled by the radar processor. Its frequency, which is proportional to the target range, is estimated in the frequency domain [17]. The discrete frequency-domain baseband signal will be referred to as the "range spectrum" because the frequency-domain bins correspond to discrete target ranges.

III. EXTENDING FMCW RADAR RANGE

The goal of this paper is to maximize the range of the radar by maximizing the effective signal-to-noise ratio (SNR) in the FMCW range spectrum. The judicious choice of FMCW parameters and the initial processing help accomplish this SNR gain without increasing transmit power.

A. Chirp Interval

For a fixed sampling rate, SNR increases linearly with the length of the chirp interval [17]. The size of hardware buffers can limit the maximum chirp interval. For example, the Inras RadarLog development kit [18] has a maximum buffer size of 10,240 samples. If the sampling rate is set at 2.5 Msps, this buffer size limits the chirp interval to a maximum of 4.096 ms.

B. Filtering

In addition to the signal of interest, the baseband signal contains wideband noise. If there are known constraints on the frequency content of the signal of interest, then SNR can be increased by filtering out some of the noise. For example, because the baseband signal is usually sampled by an analog-to-digital converter (ADC), the signal of interest is known to reside below the Nyquist frequency of the ADC sampling rate. In this case, a lowpass filter can remove all of the noise power above the Nyquist frequency before the final sampling stage. For this reason, a low sampling rate is preferable for high-sensitivity FMCW radar.

C. Coherent Integration of Multiple Chirps

The range spectra from multiple chirps can be summed coherently (retaining phase information in a complex sum) to increase SNR if the target signal's phase is consistent or predictable across the chirps. This causes the signal to add constructively, while the noise averages out to its mean. The result is that SNR increases by a factor of N_c, where N_c is the number of recorded chirps [19].

D. Noncoherent Integration of Multiple Chirps

When coherent integration is not possible due to phase uncertainty, noncoherent integration may be performed by summing the magnitudes of range spectra from multiple chirps. This discards the phase information, and is therefore less efficient than coherent integration. Noncoherent integration increases SNR by a factor of N_c^γ where $\gamma \in (.5, 1)$ [19].

E. Angle of Arrival Selection

If there is prior knowledge of a target's azimuth and elevation, then the FMCW processor may use a phased array to ignore signals arriving from all but the target direction. In the DAA application, this direction is determined by cameras. This relaxes the computational demands of the signal processing, allowing more time for the other SNR-increasing techniques mentioned above [15].

IV. LIMITATIONS ON INCREASING SNR

The long measurement interval, which is central to this paper's SNR-increasing technique, makes the radar more sensitive to the relative motion between the radar platform and the target. Relative target motion is characterized by changes in slant range and bearing over the measurement interval. Changes in bearing may be easily detected by the camera system and compensated for via simple beam steering. However, changes in range are more problematic.

A. Phase Drift

The phase of a target's radar echo depends on the modulus of the round-trip distance and signal wavelength. Due to the short wavelength of MMIC-based FMCW radars (millimeters), the phase of the corresponding peak in the range spectrum is extremely sensitive to changes in target range [17]. This has two effects:

1) The phase of a target's peak in the range spectrum is inconsistent between consecutive range spectra, which makes coherent integration impossible without phase correction.
2) Targets with nonzero radial velocity reflect a Doppler-shifted waveform. This results in a range estimation error because distance is determined by the frequency of the radar echoes. However, the radar processor can compensate for this Doppler shift by associating radar echoes across an up-chirp and down-chirp [17]. The Doppler problem is therefore not treated in this paper.

B. Range Bin Migration

Low SWaP-C FMCW chirps are configured to be short enough (e.g., $1\,\text{ms}$) that, for the range rates encountered in practice, a target's signal is concentrated in only one range bin of a single FMCW chirp's range spectrum. Standard FMCW processing assumes that targets remain in the same range bin over all the chirps in a measurement interval. However, non-constant slant range may cause a target's signal to migrate between range bins over an extended observation interval and thus invalidate this standard assumption. This so-called range bin migration (RBM) complicates the process of chirp combination because it forces the radar processor to determine not only the initial range to the target, but also account for the increased or decreased range at each successive chirp.

V. MAXIMUM LIKELIHOOD ESTIMATION

When approaching the problem of target motion estimation, it is natural to ask whether the maximum likelihood (ML) estimator is feasible to implement. In the DAA case, a visual detector has identified a target for which the range function is unknown. Thus, the object of estimation is the slant range of the target, which has many continuous degrees of freedom. To specify an ML algorithm, one must define bounds on each degree of freedom and determine an appropriate quantization for each estimated parameter. At low SNR, which is the regime in which this paper's algorithms must operate, ML estimation then becomes a search over the discretized motion parameter space. Let $\bar{\rho}(t, \boldsymbol{\alpha})$ approximate the true slant range $\rho(t)$ of the target as a linear combination of basis functions. Each basis function is a template pattern that represents an additive contribution to the target's slant range.

$$\rho(t) \approx \rho(t, \boldsymbol{\alpha}) = \sum_{k=0}^{M} \alpha_k f_k(t) \qquad (4)$$

where

$\boldsymbol{\alpha} = [\alpha_0, \alpha_1, \ldots, \alpha_M]$ is a vector of motion parameters
$M + 1$ is the number of parameters
The set $F = \{f_0(t), f_1(t), \ldots, f_M(t)\}$ spans the possible target trajectories

Specific choices of F (motion models) will be introduced in Section VI. The ML estimator for $\boldsymbol{\alpha}$ will now be derived.

Consider a radar that collects N_s noisy baseband samples for each of N_c chirps in the presence of a single target. For simplicity, assume that the SNR is constant at all ranges. The baseband FMCW signal may be modeled as a complex, discrete function [20]

$$s[n_s, n_c] = \tilde{s}[n_s, n_c] + w[n_s, n_c] \qquad (5)$$

where

$n_s \in \{0, \ldots, N_s - 1\}$ is the fast time index
$n_c \in \{0, \ldots, N_c - 1\}$ is the slow time index
$\tilde{s}[\cdot]$ is the ideal (noiseless) baseband signal
$w[\cdot]$ is zero-mean complex Gaussian noise

Fast and slow indices are used to indicate time relative to elapsed chirps and elapsed samples within a chirp: If T_s is the sampling interval, and T_h is interval between the start of consecutive chirps, then the n_s^{th} sample within the n_c^{th} chirp is taken at time $n_s T_s + n_c T_h$.

Let $g[n_s, n_c, \boldsymbol{\alpha}]$ be a function that models normalized, noiseless baseband samples given the motion model parameters $\boldsymbol{\alpha}$. This function is simple to construct based on (3). Simply substitute for t and t_d:

$$g[n_s, n_t, \boldsymbol{\alpha}] = s_{\text{bb}}\left(n_s T_s \,,\, \frac{2}{c}\bar{\rho}(n_s T_s + n_c T_h, \boldsymbol{\alpha})\right) \qquad (6)$$

where c is the speed of light. Assume that

$$g[n_s, n_c, \boldsymbol{\alpha}] = \tilde{s}[n_s, n_c] \ \ \forall\, n_s, n_c \qquad (7)$$

The ML estimator is derived from the minimization of the negative log-likelihood function

$$\hat{\boldsymbol{\alpha}}_{\text{ML}} = \underset{\boldsymbol{\alpha}}{\arg\min} \sum_{n_s=0}^{N_s-1} \sum_{n_c=0}^{N_c-1} \left| s[n_s, n_c] - g[n_s, n_c, \boldsymbol{\alpha}] \right|^2 \qquad (8)$$

The summands in (8) can also be written as

$$s^*[\cdot]s[\cdot] - s^*[\cdot]g[\cdot] - s[\cdot]g^*[\cdot] + g^*[\cdot]g[\cdot] \qquad (9)$$

where $(\cdot)^*$ denotes complex conjugation.

The sum across $s^*[\cdot]s[\cdot]$ is constant because it is independent of $\boldsymbol{\alpha}$. Because $g[\cdot]$ is normalized, the sum across $g^*[\cdot]g[\cdot]$ is also constant for all $\boldsymbol{\alpha}$. Therefore, the ML estimator may be written

$$\hat{\boldsymbol{\alpha}}_{\text{ML}} = \underset{\boldsymbol{\alpha}}{\arg\max} \sum_{n_s=0}^{N_s-1} \sum_{n_c=0}^{N_c-1} \text{Re}\left\{ s[n_s, n_c]g^*[n_s, n_c, \boldsymbol{\alpha}] \right\} \qquad (10)$$

At low SNR, optimization techniques such as gradient ascent cannot be applied to quickly find $\hat{\boldsymbol{\alpha}}_{\text{ML}}$ because the landscape of the objective function is suffused with noise-induced local maxima. Instead, a search throughout the vector space of $\boldsymbol{\alpha}$ is required [21]. Such a search presents a serious computational burden: its complexity grows exponentially with $M + 1$, and, because only the real part of the complex product is summed, the argmax operand is sensitive to small phase inaccuracies in the model $g[\cdot]$. Therefore, even though the search is trivially parallelizable, real-time applications may require more efficient suboptimal estimators.

A. ML Search Quantization

How coarsely may the parameter space for ML search be quantized before one runs the risk of "missing" the global optimum? To formalize the notion of missing the optimum, suppose a lower bound is placed on *coherence*, a measure of alignment between modeled and true motion. Let θ_k be the average phase difference over the k^{th} chirp between the approximate motion model and the true model. The coherence bound is then given by

$$C_{\min} \leq C = \text{Re}\left\{ \frac{1}{N_c} \sum_{k=0}^{N_c-1} \exp(j2\pi\theta_k) \right\} \qquad (11)$$

The coherence between $\bar{\rho}(t, \boldsymbol{\alpha})$ and a particular $\bar{\rho}(t, \boldsymbol{\alpha} + \boldsymbol{\varepsilon})$ is approximated by

$$C \approx \text{Re}\left\{ \frac{1}{T_m} \int_0^{T_m} \exp\left(j \sum_{k=0}^{M} \varepsilon_k \nu_k(t) \right) \text{d}t \right\} \qquad (12)$$

where

$$\nu_k(t) = \frac{4\pi f_0}{c} f_k(t) \tag{13}$$

and ε_k is the parameter error. Expanding (12) to second order in ε_k yields

$$C \approx 1 - \frac{1}{2} \sum_{k=0}^{M} \sum_{\ell=0}^{M} \varepsilon_k \varepsilon_\ell \langle \nu_k(t)\nu_\ell(t) \rangle \tag{14}$$

$$= 1 - \frac{1}{2} \boldsymbol{\varepsilon}^\mathsf{T} \boldsymbol{N} \boldsymbol{\varepsilon} \tag{15}$$

where $\langle \cdot \rangle$ denotes the mean over $t \in [0, T_m]$, and $\boldsymbol{N}$ is the Gram matrix with elements $\langle \nu_k(t)\nu_\ell(t) \rangle$. Now

$$\boldsymbol{\varepsilon}^\mathsf{T} \boldsymbol{N} \boldsymbol{\varepsilon} \lesssim 2(1 - C_{\min}) \tag{16}$$

When the error is in dimension k, solving for ε_k gives an expression for the quantization required to maintain a given coherence. A coherence of .5 is considered adequate for most radar applications, in which case the quantization requirement becomes

$$|\varepsilon_k| \lesssim \frac{\sqrt{2(1 - C_{\min})}}{\sqrt{\langle \nu_k(t)^2 \rangle}} \tag{17}$$

(Note: there is a second definition of coherence that is appropriate for *detection* rather than estimation; it replaces Re with the absolute value in 11.)

VI. Motion Models

The radar estimator must make assumptions about target motion to handle the effects of measurement noise, phase drift, and RBM. These assumptions are manifested in the motion model introduced in (4). This section presents two possible motion models, i.e., two possible choices for F: the polynomial model and the principal component model.

A. Polynomial Model

The polynomial model is convenient because it has an intuitive physical interpretation: the elements of $\boldsymbol{\alpha}$ relate directly to position, velocity, acceleration, etc. Let the slant range to the target aircraft $\rho(t)$ be an infinitely-differentiable function of time. Chirps are collected over a measurement interval $[0, T_m]$. Over this time, $\rho(t)$ may be approximated by a truncated Taylor series expanded about $t = T_m/2$. In terms of (4), this is expressed as

$$\alpha_k = \frac{\rho^{(k)}\left(\frac{T_m}{2}\right)}{k!} \tag{18}$$

$$f_k(t) = \left(t - \frac{T_m}{2}\right)^k \tag{19}$$

$$M = M_p \tag{20}$$

where M_p is the order of the Taylor polynomial.

1) Determining Polynomial Order (M_p): The measurement interval T_m is closely related to the required polynomial order M_p. Intuitively, a longer T_m demands a Taylor series that is accurate over a longer interval, resulting in a larger M_p.

Define $d_l \geq |\rho^{(l)}(t)|$ as an *a priori* upper bound on the magnitude of the l^{th} derivative of $\rho(t)$ for all $t \in [0, T_m]$.

The Lagrange error bound on $\Delta(t) = |\bar{\rho}(t, \boldsymbol{\alpha}) - \rho(t)|$ guarantees that

$$\Delta(t) \leq \frac{d_{M_p+1}}{(M_p + 1)!} \left| t - \frac{T_m}{2} \right|^{M_p+1} \tag{21}$$

for all $t \in [0, T_m]$.

Therefore, the maximum unmodeled distance error is,

$$\Delta_{\max} = \max_{t \in [0, T_m]} \Delta(t) \tag{22}$$

Under this framework, a designer can determine the required M_p by specifying d_l, T_m, and a maximum allowable $\Delta_{\max}$.

2) Design Example: Consider an FMCW radar system whose range bins are $15 \, \text{cm}$ wide and whose targets are guaranteed to not exceed a maximum acceleration of $10 \, \text{m/s}^2$ and a maximum jerk of $5 \, \text{m/s}^3$. The designer determines that $T_m = 1 \, \text{s}$ is required for adequate SNR. The criteria are as follows:

$$\Delta_{\max} \leq 0.15 \, \text{m}$$
$$d_2 = 10 \, \text{m/s}^2$$
$$d_3 = 5 \, \text{m/s}^3$$
$$T_m = 1 \, \text{s}$$

For $M_p = 1$, (21) and (22) yield $\Delta_{\max} = 1.25 \, \text{m}$. This violates the first criterion, so the designer must move to $M_p = 2$. In this case, $\Delta_{\max} = 0.104 \, \text{m}$. Thus, it will always be possible to find a second-order motion model accurate to within 1 range bin of the actual target position over the entire measurement interval.

3) Bounding the Search Space: As shown in Supplement A, bounds for each polynomial coefficient may be derived from a constraint on the throttle function of the target. This provides an intuitive means of limiting the parameter search space.

B. Principal Component Model

Under the hypothesis that the range evolves as a Gaussian process, and under the minimum-mean-square-error criterion, principal component analysis (PCA) gives a basis of search directions that minimizes truncation error, i.e., RMS residual range, for any given dimension of search space. Gaussian range models include nearly-constant-position, -velocity, and -acceleration models, as well as Fourier or polynomial models with jointly Gaussian coefficients.

Chirps are collected over a measurement interval $[0, T_m]$. In each of N_c chirps, the target reflects the radar signal with some phase shift. Suppose that the vector of phase samples is drawn from some normal distribution $\mathcal{N}(\boldsymbol{\mu}, P)$. For any given P, PCA gives a basis of eigenfunctions f_k and corresponding eigenvalues σ_k^2 such that the greatest part of the variance of the distribution of phase functions $\mathcal{N}(\boldsymbol{\mu}, P)$ is captured by variability along the axes defined by the first M_{PCA} eigenfunctions.

As noted in Section V, it is not sufficient to specify the axes of the ML search grid: one also requires bounds on these axes and a quantization scale. The quantization argument of Section V-A applies here without alteration. This section will therefore focus on the PCA procedure itself, bounds on the axes, and bounds on the number of axes required to achieve a given degree of coherence.

1) PCA for Selected SDE Models: An important category of Gaussian processes are the (Markov) Itô linear stochastic differential equation (SDE) models. These include nearly-constant-position, -velocity, and -acceleration models as special cases. For these models, PCA works out to the following procedure: (1) construct the linear covariance operator; (2) write out the Fredholm integral equation for the eigen-problem of the covariance operator; (3) convert the integral equation to an ordinary differential equation (ODE) using the Fundamental Theorem of Calculus; (4) find the general solution of the ODE and substitute back into the integral equation; (5) solve for conditions on the existence of eigenvalues.

For the nearly-constant-position model, the k^{th}-largest eigenvalue and the corresponding eigenfunction are given by

$$\lambda_k = \omega_k^{-2} \qquad f_k(t) = \sqrt{\frac{1}{2T_m}}\sin(\omega_k t), \text{ where} \quad (23)$$

$$\omega_k = \frac{k\pi + \frac{\pi}{2}}{T_m} \quad (24)$$

Note that the eigenvalues fall off with index k like $1/k^2$. This is because the nearly-constant-position model varies less in ways that correspond to high-frequency motion.

The results for the nearly-constant-velocity model follow $\lambda_k = \omega_k^{-4}$ with $\omega_k = \phi_k/T_m$ determined by the positive roots of

$$\cos(\phi) + \text{sech}(\phi) = 0 \quad (25)$$

Note that the nearly-constant-velocity eigenvalues drop even more rapidly than nearly-constant-position because nearly-constant-velocity is a double-integrator system, and high-frequency trajectory components are strongly suppressed in probability. The reader is directed to Supplement B for a detailed derivation of these expressions.

2) Coefficient Bounds: No finite search grid will include the entire Gaussian distribution. Let ϵ_b be an upper bound on the probability that the true range function is not in the search grid. Then, by application of the Gaussian tail bound, the magnitude of each coefficient in the search grid may be limited to

$$|a_k| \leq \sqrt{\lambda_k}\,\Phi^{-1}\left((1-\epsilon_b)^{\frac{1}{M_{\text{PCA}}}}\right) \quad (26)$$

where $\Phi()$ is the cumulative distribution function of the normal distribution.

3) Determining M_{PCA}: The final aspect of the PCA model is the number of coefficients required to achieve a certain accuracy. For a minimum absolute-value coherence of .5, the asymptotic solution ($T_m \to \infty$) for the nearly-constant-velocity case is given by

$$M = \frac{1}{\pi}\sqrt[4]{8\ln\left(\frac{1}{\epsilon}\right)T_m{}^2\sigma_0{}^2 T_h} - \frac{1}{2} \quad (27)$$

with

$$\sigma_0 = 2\pi f_0 \nu_{\text{RMS}} \quad (28)$$

where

 T_m is the measurement interval
 T_h is the time between the start of consecutive chirps
 ν_{RMS} is a nearly-constant velocity model parameter

The reader is directed to Supplement B for a detailed derivation of these expressions.

C. Choosing a Motion Model

Both motion models presented in this section are useful in different senses. The polynomial model's intuitiveness makes it preferable for demonstration, so it is used in the remainder of this paper where computational complexity is of secondary importance. The principal component model is optimal in the sense that it minimizes the dimension of the search space, so it would be the better choice for practical implementations.

VII. Computational Considerations

As shown in Section V, the optimal ML estimator may be computationally intractable for many applications. This is driven by the high-resolution search over the motion parameters. Intuitively, the ML estimator must find an $\hat{\alpha}_{\text{ML}}$ that matches both the RBM and phase drift of the true motion model over the measurement interval. Suboptimal techniques will now be examined that decouple RBM and phase drift. In doing so, the algorithms exploit the relative coarseness of the range bins to significantly reduce the search space.

These techniques all operate in the range-time space. This requires a discrete Fourier transform, typically implemented via the fast Fourier transform (FFT) and referred to hereafter as such, of the baseband signal $s[n_s, n_c]$ along the fast-time dimension to obtain $S[r_s, n_c]$, where $r_s \in \{0, 1, \ldots, N_s{-}1\}$ is the range bin index.

A. RBM Mitigation

RBM mitigation concentrates the target signal energy into a single range bin, yielding a considerable increase in effective SNR. For example, let SNR_{c} be the SNR of a single fast-time chirp. After the range FFT, the signal is given by $S[r_s, n_c]$. In the worst case, the target signal is in a different bin in each chirp's range spectrum. In this case, the slow-time signal in each range bin has an SNR of at most $\text{SNR}_{\text{b}} = (\text{SNR}_{\text{c}})(\text{N}_{\text{s}}/\text{N}_{\text{c}})$. After RBM mitigation, all the signal energy resides in a single range bin. The SNR of the corresponding slow-time signal is given by $(\text{SNR}_{\text{c}})(\text{N}_{\text{s}})$.

Several algorithms exist for RBM mitigation in pulsed-Doppler radar applications [22], [23]. The narrow sampling bandwidth and dechirp-on-receive architecture of low SWaP-C FMCW systems prevent the direct application of these pulsed-radar techniques. Motion compensation algorithms for FMCW synthetic-aperture radar offer promising insight, but they assume that the platform motion is known relative to the target scene [24], [25]. Reference [26] offers two methods for

linear ($M_p = 1$) RBM compensation. These involve changing each chirp's bandwidth β or duration T_c such that linear RBM is naturally cancelled. These techniques are not considered in this paper due to the constraints they impose on the radar hardware.

The most straightforward way to accomplish RBM mitigation for FMCW radar is to search over target motion hypotheses. This search over target trajectories is much coarser than the ML search because it only needs to match the true trajectory to within a range bin, whereas the ML search requires phase-level accuracy. Let $\bar{S}_i[r_s, n_c]$ denote the range-time baseband signal where each range spectrum has been shifted according to the i^{th} RBM hypothesis. Under the best RBM hypothesis index i^*, all the signal energy will reside in a single range bin of $\bar{S}_{i^*}[r_s, n_c]$. The best hypothesis will be chosen as that which maximizes the detection statistic of a signal detector applied to each $\bar{S}_i[\cdot]$. Signal detectors are examined below.

B. Signal Detection for $M_p = 1$

For a given range bin index r, $\bar{S}_i[r, n_c]$ is a sequence of complex values along the slow-time dimension. The radar processor simply needs to find the range bin index r^* containing the target's signal. An ML search over velocity would accomplish this. Because the phase progression of $\bar{S}_{i^*}[r^*, n_c]$ is linear for $M_p = 1$, the ML search may be efficiently implemented as an FFT along the slow-time dimension. This runs in $O(N_s N_c \log(N_c))$ time. Conveniently, the highest peak in the resulting range-Doppler space indicates ML estimates for range and Doppler [17], [27].

C. Signal Detection for $M_p > 1$

For $M_p > 1$, the phase progression of $\bar{S}_{i^*}[r^*, n_c]$ is a polynomial of order M_p. Again, the optimal solution for parameter estimation is an ML search [21]. Because this requires an $(M_p + 1)$-dimensional search in each range bin, it quickly becomes infeasible to do in real time.

Many techniques exist for estimating the parameters of polynomial-phase signals. For second-order polynomial-phase signals, [28] and [29] offer search-based algorithms that approach optimal performance at high SNR. Reference [30] presents a fast solution for third-order polynomial-phase signals. All of these algorithms would require adaptation to be useful in the present scenario because the radar processor is not really interested in estimating the parameters of the polynomial-phase signal. Rather, it need only distinguish between range bins that contain signal and range bins that don't. Therefore, this paper considers the following three detectors for $M_p > 1$: the discrete polynomial transform (DPT), noncoherent integration, and spectral entropy. Each of these is examined below and simulated in Section VIII-A.

D. Discrete Polynomial Transform (DPT)

The DPT is a parameter-estimation algorithm suitable for any polynomial order that is also useful for detection [21]. It is computationally efficient with $O(N_s N_c \log(N_c))$ runtime, but requires relatively high SNR to be effective for detection.

E. Noncoherent Integration

Described in Section III-D, this technique is attractive for its simplicity and $O(N_s N_c)$ runtime, though it sacrifices performance by discarding phase information.

F. Spectral Entropy

The signal-bearing bin contains a slow-time complex signal whose phase and amplitude correspond to the target motion. Bins without signal simply contain Gaussian noise. Spectral entropy offers an efficiently-calculated quantity that measures the "randomness" of the slow-time signals in each bin [31]. Signal-bearing bins exhibit lower spectral entropy than those that contain only noise. Like the DPT, this method requires $O(N_s N_c \log(N_c))$ time. The spectral entropy of a complex signal $z[n]$ is most easily computed as follows:

1) Compute the power spectrum $P_z[f] = |\mathcal{F}\{z[n]\}|^2$.
2) Normalize $P_z[f]$ so it can be treated as a probability mass function (PMF).
3) Compute the Shannon entropy of this PMF, which is the spectral entropy.

The algorithm used in this paper is similar to the more advanced technique used in [31], which involves a windowed short-time Fourier transform to compute the power spectrum.

VIII. SIMULATION RESULTS

A. Signal Detection Comparison

The FFT, DPT, noncoherent integration, and spectral entropy signal detectors were tested against each other in simulation. To permit a focus on the detection problem unencumbered with RBM mitigation, the simulation's motion model was such that the target remained in a single range bin over the measurement interval. Results are presented in Fig. 2.

B. Full Simulation for $M_p = 1$

This simulation applied a nearly-constant velocity model that gave rise to significant RBM which had to be compensated. Results are presented in Fig. 3.

C. Full Simulation for $M_p = 2$

This simulation applied a nearly-constant acceleration model that gave rise to significant and rate-increasing RBM which had to be compensated. Results are presented in Fig. 4.

IX. EXPERIMENTAL RESULTS

This paper's range-extending technique was implemented and tested with the Inras RadarLog, a development platform for 77 GHz radar [18]. Note that the techniques explored in this paper are transferrable to the 24 GHz band, which is preferred for aviation in the United States due to FCC regulations [33]. To simulate in-flight conditions, the RadarLog was placed face-up on the ground in an open field. An unmodified DJI Mavic 2 Pro drone [34] was flown at various discrete altitudes up to $120\,\mathrm{m}$ directly above the RadarLog to act as the target. See Fig. 5 for a photo of the experimental hardware. At each

(a)

(b)

Single-Chirp SNR (dB)

(c)
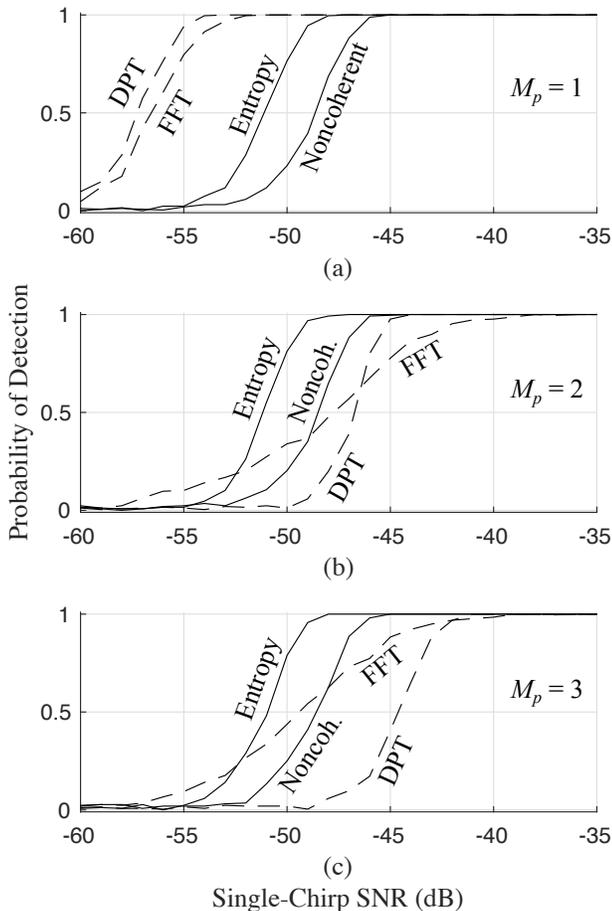
Fig. 2: Simulation results for four detectors: DPT, FFT, noncoherent integration, and spectral entropy. The simulator randomly generated 256 target trajectories for the $M_p = 1$, $M_p = 2$, and $M_p = 3$ cases. The vertical axes reflect the percentage of these 256 trajectories whose range was estimated to within $1\,\mathrm{m}$ of the true range. Each trajectory was generated such that the target remained in a single range bin over the entire measurement interval. 1024 chirps were simulated for each trajectory, and each method was tested on the simulated baseband signals.

The FFT and DPT are optimal and equivalent for $M_p = 1$ [21], [27]. The DPT uses an internal zero-padded FFT for greater resolution, which accounts for the slight difference in the plots. For $M_p > 1$, the DPT requires relatively high SNR, and the FFT performs poorly because it implicitly assumes that the slow-time signals have linear phase. The FFT was included in the $M_p > 1$ simulations as a point of interest, though it is only appropriate for $M_p = 1$.

Spectral entropy and noncoherent integration are invariant to M_p because neither method makes any assumptions about the order of the slow-time phase. Spectral entropy consistently outperforms noncoherent integration at the cost of greater computational complexity.
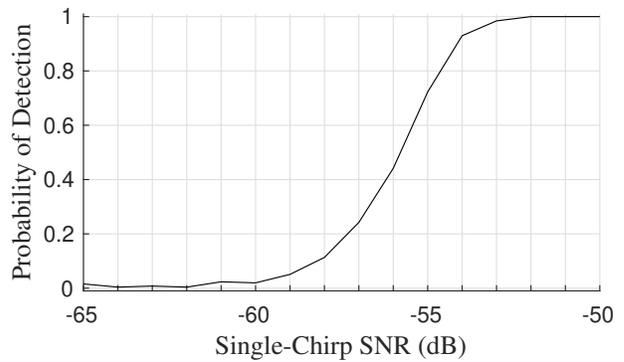
Fig. 3: Simulation results for $M_p = 1$. The simulator randomly generated 256 target trajectories using the nearly-constant-velocity model presented in [32]. The vertical axis reflects the percentage of these 256 trajectories whose initial range was estimated to within $1\,\mathrm{m}$ of the true initial range. The target velocities ranged from $-5.8\,\mathrm{m/s}$ to $5.8\,\mathrm{m/s}$, and the velocity hypotheses for RBM mitigation covered the same range. 1024 baseband chirp signals were simulated for each trajectory, and these were processed using range bin shifting and a slow-time FFT as described in Sections VII-A and VII-B.
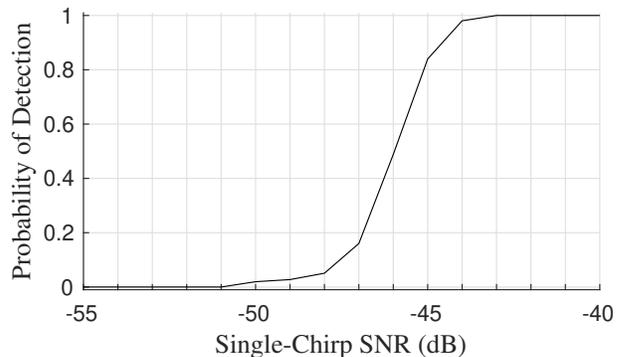
Fig. 4: Simulation results for $M_p = 2$. The simulator randomly generated 256 target trajectories using the the nearly-constant-acceleration model presented in [32]. The vertical axis reflects the percentage of these 256 trajectories whose initial range was estimated to within $1\,\mathrm{m}$ of the true initial range. The target velocities ranged from $-2.9\,\mathrm{m/s}$ to $2.9\,\mathrm{m/s}$, and the target accelerations ranged from $-2.8\,\mathrm{m/s^2}$ to $2.8\,\mathrm{m/s^2}$. Hypotheses for RBM mitigation covered the same ranges. 1024 baseband chirp signals were simulated for each trajectory, and these were processed using range bin shifting and noncoherent integration as described in Sections VII-A and VII-E. Noncoherent integration was applied for detection instead of spectral entropy, which as shown in Fig. 2 is superior for the $M_p = 2$ case, to reduce the computation demands of the simulation.

altitude, the drone was commanded to hover. The drone would have been detectable beyond $120\,\mathrm{m}$, but FAA regulations do not allow remote pilots to operate above that altitude without a waiver.

Chirp parameters were chosen based on the theory in Section III. A low sampling rate of 2.5 Msps and long
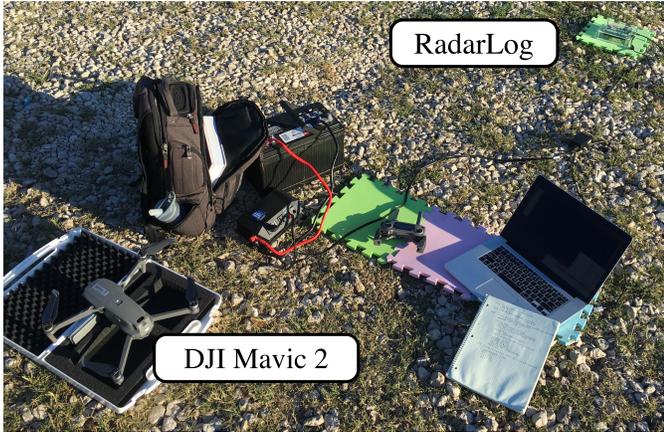
Fig. 5: Experimental hardware: the RadarLog unit rests on a ground pad in the upper right corner of the photo. An unmodified DJI Mavic 2 was flown up to an altitude of 120 m to act as a target.

chirp durations up to 4.096 ms were chosen to maximize the effective SNR. See Table I for a list of all configurations. These configurations were tested on the target UAV at 20 m increments up to an altitude of 120 m. The transmit power was held constant at 10 dBm for the entire experiment. The results of this experiment are summarized in Fig. 6. Only data points for which the drone was detectable are plotted.

TABLE I: Chirp configurations tested at each altitude. Note that the number of chirps for each configuration was adjusted to maintain a constant measurement interval. In other words, the same amount of data was collected for each chirp configuration. The chirps were summed noncoherently to produce the results in Fig. 6.

Chirp Interval (μs)	Buffer Length	Number of Chirps	Total Measurement Interval (s)
4096	10240	256	1.049
2048	5120	512	1.049
1024	2560	1024	1.049
512	1280	2048	1.049
256	640	4096	1.049

Notice that, for the most part, the shorter chirp intervals yield lower SNR despite the constant measurement interval. This happens because an FMCW chirp is intrinsically coherent, and therefore more efficient at raising SNR than the noncoherent integration used to combine separate chirps. In general, a few long chirps will yield higher SNR than many short chirps over equivalent measurement intervals. In practice, chirp length is limited by RBM.

A. RBM and Phase Drift

Although the target drone was commanded to hover at each altitude, the experimental data exhibits mild RBM and significant phase drift which are depicted in Figs. 7 and 8, respectively. These figures are representative of the RBM and phase drift for all configurations listed in Table I. Linear range bin shifting and noncoherent integration were used to mitigate these effects.
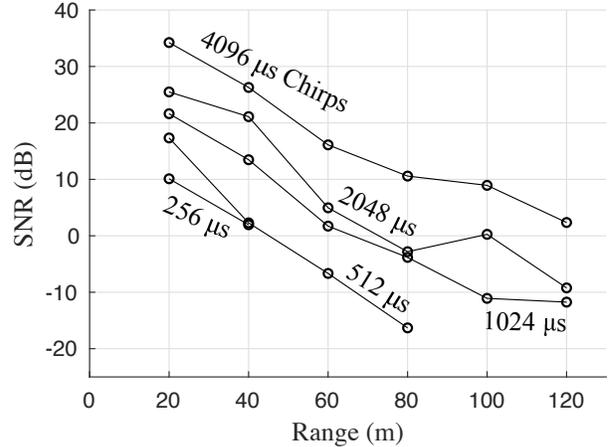
Fig. 6: Effective SNR of the radar return after range bin shifting (assuming $M_p = 1$) and noncoherent integration. Noise power was estimated from control data which was recorded with no target present. This noise estimate was used to calculate SNR.
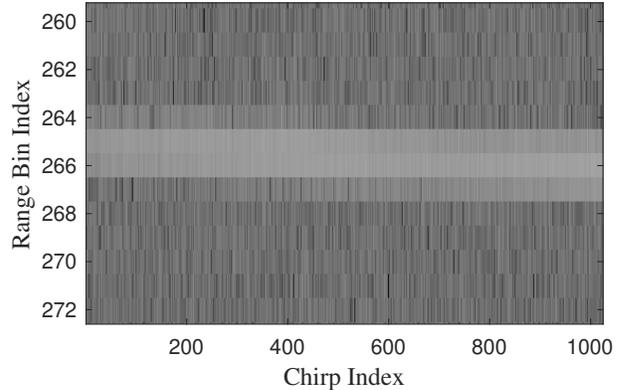
Fig. 7: RBM in experimental data. The magnitudes of the signal-containing range bins are shown for each chirp in the measurement interval. These data were collected at 20 m range with $N_c = 1024$ and $N_s = 2560$. The bright horizontal band indicates that the target dwelled in at most four range bins over the measurement interval. Note that at longer ranges, the signal's bright band would not be visible in such magnitude plots.

X. EXTRAPOLATION OF RESULTS

The foregoing experimental results can be extrapolated to infer the performance of a real system designed with the same techniques. Consider how target radar cross section (RCS) and transmit power would change upon application of this paper's techniques to UAM:

1) Radar Cross Section: Radar cross section is a measure of a radar target's reflectivity. The signal power in watts that a radar receives from a target is given by [17]

$$P_{\text{RX}} = \frac{P_{\text{TX}} G_{\text{TX}} \sigma A_{\text{RX}}}{(4\pi)^2 d^4} \tag{29}$$

where

9

Fig. 8: Phase drift in experimental data. Gray points indicate the raw phase of the slow-time signal in the target's range bin $(120\,\mathrm{m})$. Each phase measurement is vertically replicated every 2π radians to reveal contours in the plot. These data were collected with $N_c = 1024$ and $N_s = 2560$.

P_{TX} is the transmit power (W)

G_{TX} is the transmit antenna gain (dimensionless)

σ is the RCS of the target (m^2)

A_{RX} is the effective aperture of the receive antenna (m^2)

d is the distance to the target (m)

An object's RCS is often difficult to calculate because it depends on many factors, including shape, material, size, and orientation. For the purpose of this analysis, assume that RCS scales linearly with geometric cross section. The DJI drone in this experiment has overall horizontal dimensions of $0.322\,\mathrm{m}\times 0.242\,\mathrm{m}$ [34].

Joby Aviation, a leader in UAM technology, recently unveiled their VTOL airframe, which has a wingspan of $10.7\,\mathrm{m}$ [35]. Assume the height is $3\,\mathrm{m}$. Under these assumptions, the ratio of geometric cross sections (front of Joby's airframe to the bottom of the DJI drone) is about

$$\frac{\Sigma_{\mathrm{Joby}}}{\Sigma_{\mathrm{DJI}}} \approx \frac{10.7\,\mathrm{m} \times 3\,\mathrm{m}}{0.322\,\mathrm{m} \times 0.242\,\mathrm{m}} \approx 412$$

According to (29), this increase in RCS translates to a range increase by a factor of

$$\sqrt[4]{\frac{\Sigma_{\mathrm{Joby}}}{\Sigma_{\mathrm{DJI}}}} \approx 4.5$$

The RadarLog detected the DJI drone at $120\,\mathrm{m}$ using only linear range bin shifting and noncoherent integration, so it can be expected to detect a typically-sized VTOL aircraft at $540\,\mathrm{m}$ under similar conditions with no modification to the hardware.

2) Transmit Power: The RadarLog currently transmits at $10\,\mathrm{dBm}$. With a more robust transmit amplifier, the power could be increased by 10 dB with little change to SWaP-C. This would increase overall performance and allow detection of targets in more challenging conditions.

XI. Conclusions

Recently-developed low-cost radar chipsets are enjoying widespread use in the automotive industry, but their application to other industry segments is limited by their relatively short range. If their range could be extended, these emergent chipsets would be readily adopted for use in long-range applications. They would be particularly useful in detect-and-avoid systems for urban air mobility (UAM), where they would provide a fundamental complement to existing camera-based systems.

A range-extending solution for low-cost radar was developed, analyzed, and tested. The solution employs SNR-maximizing parameters, extended measurement intervals, and additional signal processing to maximize range. Using the maximum likelihood estimator as a starting point, this paper developed more efficient techniques for target range estimation that are robust to target motion. The performance of these techniques was quantified in simulation and demonstrated via experiment. The results show that low-cost radar can be adapted for range-extended applications with the method presented in this paper. These techniques will enable UAM and other robotics applications to enjoy the benefits of the latest low-cost radar hardware.

References

[1] J. Morrissey and P. Walsh, "High performance integrated 24 GHz FMCW radar transceiver chipset for auto and industrial sensor applications." https://www.analog.com/en/technical-articles/integrated-24-ghz-fmcw-radar-transceiver-chipset.html, Jan. 2020.

[2] D. P. Thipphavong, R. Apaza, B. Barmore, V. Battiste, B. Burian, Q. Dao, M. Feary, S. Go, K. H. Goodrich, J. Homola, *et al.*, "Urban air mobility airspace integration concepts and considerations," in *2018 Aviation Technology, Integration, and Operations Conference*, p. 3676, 2018.

[3] J. Holden and N. Goel, "Uber elevate: Fast-forwarding to a future of on-demand urban air transportation," *Uber Inc., San Francisco, CA*, 2016.

[4] W. C. Barott, E. Coyle, T. Dabrowski, C. Hockley, and R. S. Stansbury, "Passive multispectral sensor architecture for radar-EOIR sensor fusion for low SWaP UAS sense and avoid," in *2014 IEEE/ION Position, Location and Navigation Symposium-PLANS 2014*, pp. 1188–1196, IEEE, 2014.

[5] J. James, J. J. Ford, and T. L. Molloy, "Learning to detect aircraft for long-range vision-based sense-and-avoid systems," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4383–4390, 2018.

[6] S. Smearcheck, S. Calhoun, W. Adams, J. Kresge, and F. Kunzi, "Analysis of alerting performance for detect and avoid of unmanned aircraft systems," in *2016 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pp. 710–730, IEEE, 2016.

[7] J. James, J. J. Ford, and T. L. Molloy, "Below horizon aircraft detection using deep learning for vision-based sense and avoid," *arXiv preprint arXiv:1903.03275*, 2019.

[8] C. M. Geyer, D. Dey, and S. Singh, "Prototype sense-and-avoid system for UAVs," Tech. Rep. CMU-RI-TR-09-09, Carnegie Mellon University, Pittsburgh, PA, May 2009.

[9] J. James, J. J. Ford, and T. L. Molloy, "Quickest detection of intermittent signals with application to vision-based aircraft detection," *IEEE Transactions on Control Systems Technology*, vol. 27, no. 6, pp. 2703–2710, 2018.

[10] J. Lai, J. J. Ford, L. Mejias, and P. O'Shea, "Characterization of sky-region morphological-temporal airborne collision detection," *Journal of Field Robotics*, vol. 30, no. 2, pp. 171–193, 2013.

[11] D. Bratanov, L. Mejias, and J. J. Ford, "A vision-based sense-and-avoid system tested on a scaneagle uav," in *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 1134–1142, IEEE, 2017.

[12] C. Minwalla and K. Ellis, "Experimental evaluation of PICAS: An electro-optical array for non-cooperative collision sensing on unmanned aircraft systems," in *AIAA Information Systems-AIAA Infotech@ Aerospace*, p. 0907, 2017.

[13] A. Mcfadyen and L. Mejias, "A survey of autonomous vision-based see and avoid for unmanned aircraft systems," *Progress in Aerospace Sciences*, vol. 80, pp. 1–17, 2016.

[14] Echodyne, "EchoFlight airborne radar." https://echodyne.com/products/echoflight/, Sept. 2019.

[15] W. A. Lies, L. Narula, P. A. Iannucci, and T. E. Humphreys, "Low SWaP-C radar for urban air mobility," in *Proceedings of the IEEE/ION PLANSx Meeting*, 2020.

[16] K. Ramasubramanian and T. Instruments, "Using a complex-baseband architecture in fmcw radar systems," *Texas Instruments*, 2017.

[17] S. Rao, "Introduction to millimeter-wave sensing: FMCW radars," *Texas Instruments (TI) Millimeter-Wave Training Series*, 2017.

[18] INRAS, "Products." http://www.inras.at/en/products.html, Jan. 2020.

[19] M. A. Richards, "Noncoherent integration gain, and its approximation," *Georgia Institute of Technology, Technical Memo*, 2010.

[20] MathWorks, "Radar data cube." https://www.mathworks.com/help/phased/gs/radar-data-cube.html, Sept. 2020.

[21] S. Peleg and B. Friedlander, "The discrete polynomial-phase transform," *IEEE Transactions on Signal Processing*, vol. 43, no. 8, pp. 1901–1914, 1995.

[22] F. Uysal, "Comparison of range migration correction algorithms for range-doppler processing," *Journal of Applied Remote Sensing*, vol. 11, no. 3, p. 036023, 2017.

[23] J. Chen, K. Jin, S. Yu, T. Lai, and Y. Zhao, "Radar coherent detection for maneuvering target based on product-scaled integrated cubic phase function," *International Journal of Antennas and Propagation*, vol. 2019, 2019.

[24] G. Xue, J. Yang, and P. Liu, "Modified range migration algorithm integrated with motion compensation for FMCW SAR," in *IET International Radar Conference*, IET, 2013.

[25] R. Wang, Y.-H. Luo, Y.-K. Deng, Z.-M. Zhang, and Y. Liu, "Motion compensation for high-resolution automobile FMCW SAR," *IEEE Geoscience and Remote Sensing Letters*, vol. 10, no. 5, pp. 1157–1161, 2013.

[26] F. Roos, D. Ellenrieder, N. Appenrodt, J. Dickmann, and C. Waldschmidt, "Range migration compensation for chirp-sequence based radar," in *2016 German Microwave Conference (GeMiC)*, pp. 317–320, IEEE, 2016.

[27] D. Rife and R. Boorstyn, "Single tone parameter estimation from discrete-time observations," *IEEE Transactions on information theory*, vol. 20, no. 5, pp. 591–598, 1974.

[28] S. Ma, J. Jiang, and Q. Meng, "A fast, accurate and robust method for joint estimation of frequency and frequency rate," in *2011 International Symposium on Intelligent Signal Processing and Communications Systems (ISPACS)*, pp. 1–6, IEEE, 2011.

[29] T. J. Abatzoglou, "Fast maximnurm likelihood joint estimation of frequency and frequency rate," *IEEE Transactions on Aerospace and Electronic Systems*, no. 6, pp. 708–715, 1986.

[30] P. O'shea, "A fast algorithm for estimating the parameters of a quadratic fm signal," *IEEE Transactions on Signal Processing*, vol. 52, no. 2, pp. 385–393, 2004.

[31] MathWorks, "Pentropy: Spectral entropy of a signal." https://www.mathworks.com/help/signal/ref/pentropy.html, July 2020.

[32] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley and Sons, 2001.

[33] Federal Communications Commission, "47 CFR Part 15.253: Operation within the bands 46.7–46.9 GHz and 76.0–77.0 GHz." Federal Register, Oct. 2000.

[34] DJI, "MAVIC 2 specifications." https://www.dji.com/mavic-2/info, Jan. 2020.

[35] Vertical Flight Society, "Joby S4." https://evtol.news/aircraft/joby-aviation/, Jan. 2020.