Article | Received 3 July 2026; Revised 11 August 2026; Accepted 4 November 2026; Published xx November 2026
https://doi.org/10.55092/esp20260006

Performance evaluation of chaos-based FM waveforms with optimized ambiguity functions against FMCW and noise radars in LPI applications

Amr S. Youssef

Aviation Science Department, Higher Colleges of Technology, Al Ain, UAE; E-mail: ayoussef@hct.ac.ae

Highlights:

- New IPCG metric captures replica-reuse resistance missed by the interceptability index.
- FMCW is coherently interceptable about 14 times farther than Lorenz-CBFM.
- A CNN recognizes all four waveforms but learns a reusable replica only for FMCW.

Abstract: Modern electronic-support measures (ESM) increasingly employ deep-learning classifiers that exploit the deterministic, periodic structure of conventional low-probability-of-intercept (LPI) waveforms such as linear frequency-modulated continuous-wave (FMCW) signals. This paper investigates chaos-based frequency-modulated (CBFM) waveforms, generated from the Lorenz and Rössler attractors, as a countermeasure. Using a fully simulation-based framework (2 GHz sampling rate, 10 μs pulse, 500 MHz bandwidth, time–bandwidth product 5000), we compare CBFM against calibrated FMCW and band-limited noise-radar references on spectral flatness, autocorrelation peak-sidelobe ratio (PSR), ambiguity-function shape, and a newly defined inter-pulse correlation gain (IPCG) that quantifies an interceptor's ability to reuse a previously captured pulse. Although FMCW attains the flattest spectrum (spectral flatness measure 0.97) and lowest sidelobes (PSR = −21.4 dB), it is perfectly periodic (IPCG = 1.0) and therefore fully exploitable by a replica-based interceptor. Lorenz-CBFM reduces the IPCG to 0.072, essentially the noise-radar floor (0.063), while delivering a thumbtack ambiguity function free of FMCW's range–Doppler coupling, at the cost of modestly higher sidelobes (−17.8 dB). A derivation linking IPCG to interceptor coherent gain shows that a replica-reuse interceptor detects an FMCW emitter from about 14 times farther away than a Lorenz-CBFM emitter. A convolutional-neural-network experiment confirms that all four families are recognizable by type, yet a learned next-pulse predictor recovers near-unity coherent gain only for FMCW. Detection is modulation-independent at equal pulse energy, while the thumbtack waveforms remove a systematic range–Doppler bias that displaces a 1500 m/s FMCW target by a full range-resolution cell. A 13-point parameter study identifies the Lorenz parameter ranges maximizing spectral flatness and minimizing sidelobe levels.

Keywords: low probability of intercept (LPI); chaotic frequency modulation; Lorenz attractor; Rössler attractor; ambiguity function; electronic support measures; deep learning; noise radar; cyclostationary analysis

Copyright©2026 by the authors. Published by ELSP. This work is licensed under Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium provided the original work is properly cited.

1. Introduction

Low-probability-of-intercept (LPI) radar seeks to detect targets while remaining undetected by hostile electronic-support measures (ESM) receivers [1]. The classical approach spreads the transmitted energy over a wide bandwidth and a long integration time so that the peak power density falls below the interceptor's detection threshold [2,3]. Linear frequency-modulated continuous-wave (FMCW) signals are the workhorse of this strategy because they are simple to generate and provide a large time–bandwidth product (BT). However, an FMCW sweep is strictly periodic and fully deterministic: once an interceptor estimates the (single) chirp-rate parameter, it can construct a matched replica and recover full coherent processing gain [1,4].

This vulnerability has become acute with the emergence of deep-learning-based ESM. Convolutional classifiers operating on time–frequency images or on the polyphase autocorrelation function now recognize and parameterize LPI modulations with high accuracy [5,6], and recent generative approaches explicitly learn to defeat fixed waveform libraries [7,8]. Such classifiers thrive on repeatable structure; a waveform that never repeats denies them the stationary features they rely on.

Chaotic systems offer exactly this property. A chaotic trajectory is deterministic—and therefore reproducible at a cooperative receiver that knows the initial conditions—yet aperiodic and broadband, with sensitive dependence on initial conditions [9–11]. Chaotic modulation has consequently been exploited for secure and spread-spectrum communications, and the Lorenz system in particular has been validated as a source of statistically random sequences that pass the standard randomness test suites of the National Institute of Standards and Technology (NIST) [12]. Passive interceptors, in turn, have moved well beyond energy detection: cyclostationary and higher-order-statistical processing extract cyclic spectral lines and modulation parameters from continuous-wave LPI emissions [1,13], and dedicated waveform-design methods now explicitly minimize the cyclic-spectral signature that such receivers exploit [13]. In radar and remote sensing, chaos-based frequency modulation (CBFM) has been studied for high-resolution imaging [14,15] and, more recently, for frequency-hopping LPI design [16]; the Lorenz and Rössler attractors in particular have been used to synthesize CBFM waveforms whose ambiguity functions approach a thumbtack shape [17–19], and chaos-based radar synchronization has been demonstrated for bistatic geometries [20].

Despite this activity, three gaps remain. First, prior CBFM studies rarely calibrate the chaotic waveform to the same occupied bandwidth and BT product as the baselines they are compared against, so reported advantages may simply reflect unequal bandwidth rather than any property of the modulation. Second, interceptability is usually argued qualitatively—the waveform "looks like noise"—rather than quantified with a metric that captures how an interceptor would actually exploit the signal. Third, the trade-offs against the two most relevant references—linear FMCW and stochastic noise radar—are seldom reported side by side on a common footing. A direct, metric-by-metric comparison under identical bandwidth and BT product is therefore still lacking.

Relation to prior chaotic-radar and LPI-detection work. Because chaos-based radar is an active field, we state precisely what is and is not new here. Prior work has established that Lorenz- and Rössler-derived CBFM waveforms possess thumbtack ambiguity functions and noise-like spectra [14,17–19], and has optimized chaotic sequences for radar and analyzed their synchronization [18,20]; complementary work on the interceptor side has designed LPI waveforms specifically to defeat cyclostationary analysis [13]

and has characterized higher-order-statistical intercept processing [1]. A parallel line of work makes the waveform itself adaptive, selecting parameters per pulse in response to the sensed environment [21]; the waveforms studied here use fixed chaotic parameters throughout, so that the properties of the modulation are isolated from those of any selection policy. Our contribution is not to re-establish that chaotic frequency modulation (FM) is noise-like. It is (i) a calibrated, metric-by-metric head-to-head placing Lorenz-CBFM, Rössler-CBFM, FMCW, and band-limited noise radar on an identical 500 MHz/BT = 5000 footing; (ii) the inter-pulse correlation gain (IPCG), a waveform-level metric for the replica-reuse threat that the classical radiometric index cannot see; (iii) an explicit derivation linking IPCG to interceptor coherent processing gain and interception range; and (iv) a deep-learning ESM experiment that tests, rather than asserts, robustness to a learning interceptor. Points (ii)–(iv) are, to our knowledge, not present in the works above.

These gaps motivate the present study, which addresses each in turn. First, we formulate Lorenz- and Rössler-based CBFM waveforms and calibrate all four waveforms to a common 500 MHz occupied bandwidth and BT = 5000, removing bandwidth as a confounding variable (closing the first gap). Second, we introduce IPCG and the peak cross-correlation between successive pulses—as a waveform-level interceptability metric that directly quantifies the replica-reuse threat, complementing the conventional link-budget interceptability index of Section 4.1 (closing the second gap). Third, we provide a reproducible, simulation-based comparison—spectral flatness, peak-sidelobe ratio (PSR), the Woodward ambiguity function, and a 13-point chaotic-parameter sweep—that places CBFM, FMCW, and noise radar on identical footing and quantifies precisely where chaotic FM helps and where it does not (closing the third gap).

2. Chaos-based FM waveform model

The general form of a chaos-based FM radar waveform is

$$s(t) = A\cos\left[2\pi f_c t + 2\pi k_f \int_0^t x(\tau)d\tau\right] \quad (1)$$

where A is the amplitude, f_c the carrier frequency, k_f the modulation (frequency-deviation) gain, and $x(\tau)$ a state variable of a continuous-time chaotic system. The instantaneous frequency follows the chaotic state directly,

$$f_i(t) = f_c + k_f x(t) \quad (2)$$

Figure 1 illustrates the construction: the normalized Lorenz state drives the instantaneous frequency, which in turn produces the aperiodic baseband waveform. Because $x(t)$ never repeats, neither does $s(t)$.

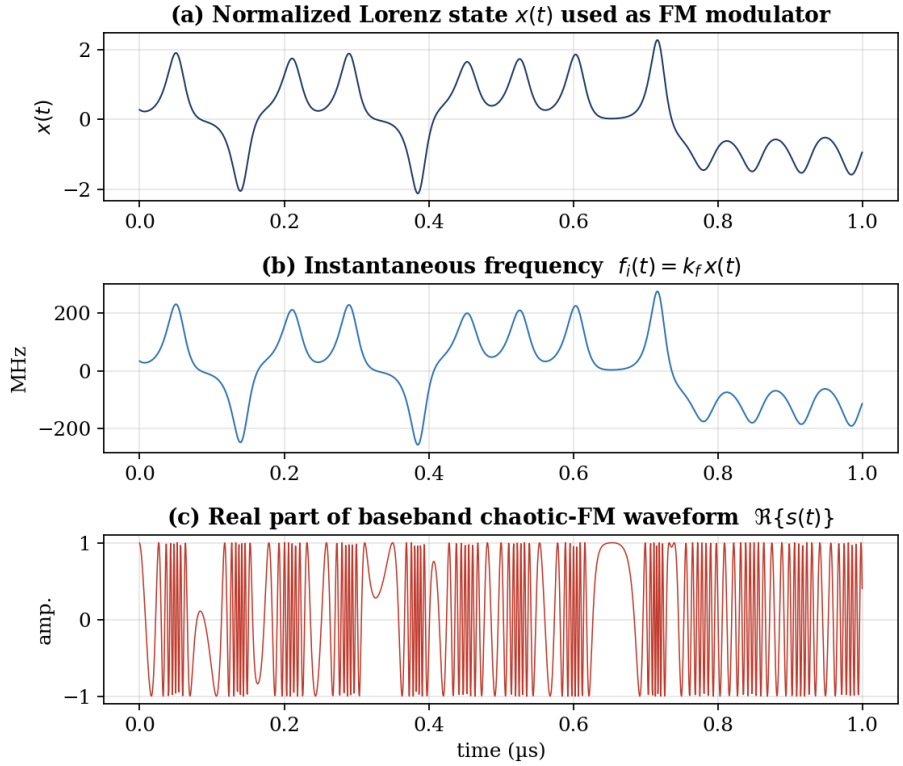


Figure 1. Construction of a chaotic-FM waveform. **(a)** Normalized Lorenz state x(t); **(b)** Instantaneous frequency $f_i(t) = k_f x(t)$; **(c)** Real part of the resulting baseband waveform Re{s(t)}.

The instantaneous-frequency excursion, and hence the occupied bandwidth, is set by the deviation gain k_f. Because the modulating state is normalized to zero mean and unit standard deviation ($\sigma_x = 1$), k_f carries units of hertz per unit of x—the frequency shift produced by a one-standard-deviation excursion of the chaotic state. To first order (Carson's rule), the root-mean-square (RMS) occupied bandwidth scales linearly with k_f,

$$B_{RMS} = 2k_f\sigma_x \tag{3}$$

where σ_x is the standard deviation of the modulating state. We quantify bandwidth throughout by the 99% occupied bandwidth—the contiguous band containing 99% of the signal energy. Unlike the −3 dB or null-to-null definitions, this measure is robust to the spectral ripple and slowly decaying tails of wideband FM and matches the convention used in spectrum-occupancy regulation, so it provides a fair, waveform-agnostic basis for equalizing bandwidth across the four waveforms. Figure 2 shows the instantaneous frequency for three values of k_f (legend in MHz per unit σ_x), confirming that larger k_f widens the excursion. Figure 3 validates Equation (3): the measured 99% bandwidth is exactly linear in k_f ($R^2 = 1.000$). The fitted slope (≈ 4.1) exceeds the RMS coefficient of 2 because the 99% bandwidth captures the peak excursions of the chaotic state (peak-to-RMS ratio ≈ 2) together with the finite modulation rate, in accordance with Carson's rule; the RMS estimate of Equation (3) is therefore a lower bound. We use this calibrated relationship to fix every waveform to 500 MHz ($k_f \approx 121$ MHz per unit σ_x).

Two canonical attractors are considered. The Lorenz system is

$$\dot{x} = \sigma(y - x), \qquad \dot{y} = x(\rho - z) - y, \qquad \dot{z} = xy - \beta z \tag{4}$$

with the canonical parameters σ = 10, ρ = 28, β = 8/3, for which the system is chaotic [22]. The Rössler system is

$$\dot{x} = -y - z, \qquad \dot{y} = x + ay, \qquad \dot{z} = b + z(x - c) \tag{5}$$

with $a = b = 0.2$, $c = 5.7$. Both systems are integrated with the classical fourth-order Runge–Kutta (RK4) method, chosen because its fourth-order local accuracy preserves the fine trajectory structure that defines the waveform; lower-order integrators introduce numerical dissipation that artificially smooths the spectrum (see Section 5).

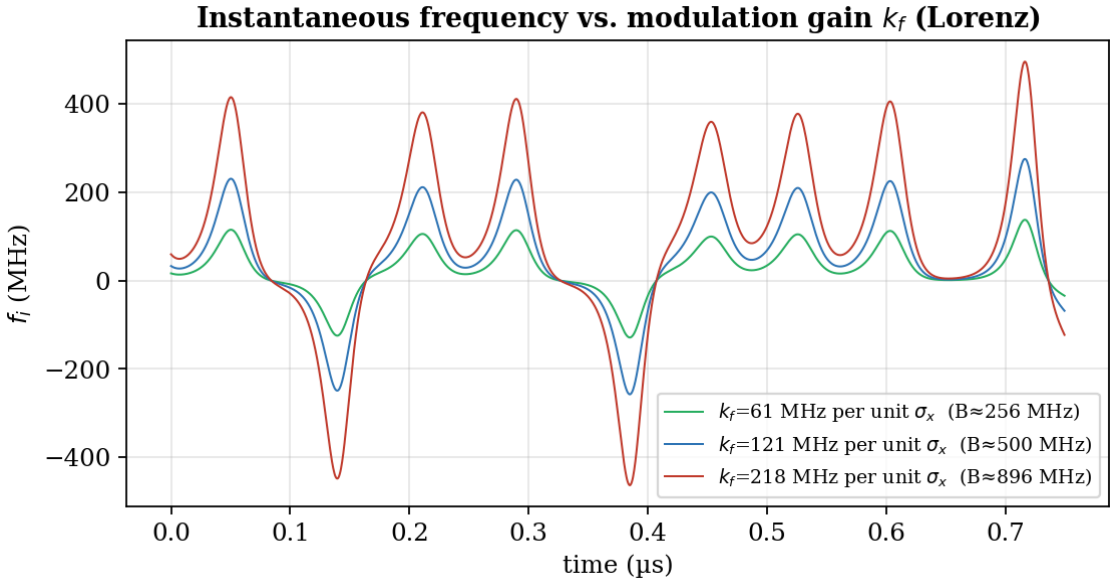


Figure 2. Instantaneous frequency of the Lorenz-CBFM waveform for three modulation gains k_f, expressed in MHz per unit standard deviation of the normalized chaotic state (σ_x = 1). Larger k_f widens the frequency excursion and hence the occupied bandwidth.

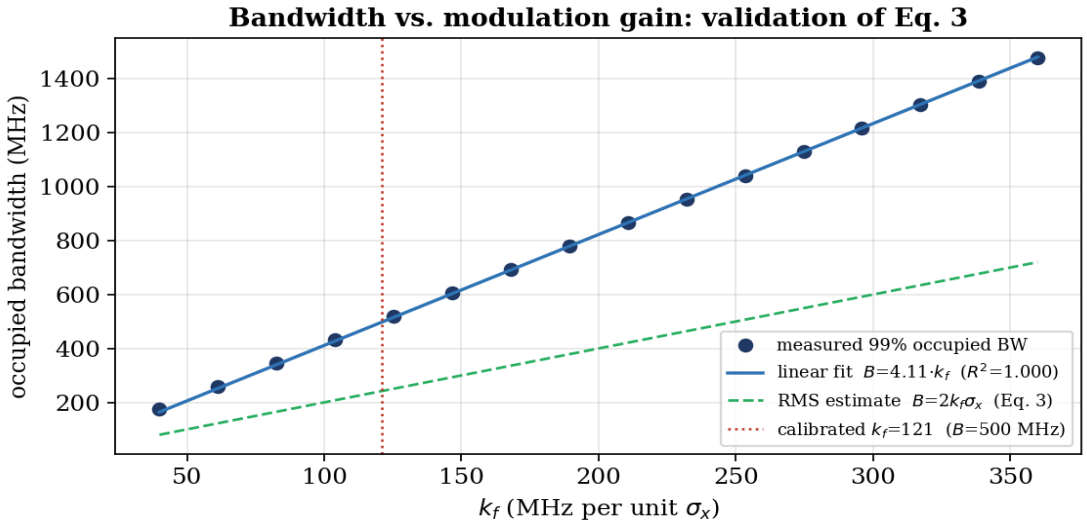


Figure 3. Validation of the bandwidth law Equation (3). Markers: measured 99% occupied bandwidth versus k_f (Lorenz). Solid line: least-squares linear fit, B = 4.11k_f (R^2 = 1.000). Dashed line: the RMS estimate B = 2k_f σ_x of Equation (3), a lower bound. Dotted line: the calibrated operating point (k_f≈ 121 MHz per unit σ_x, B = 500 MHz).

3. Baseline reference waveforms

To place CBFM in context we compare it against the two waveform families it most directly competes with.

3.1. Linear FMCW

The linear FMCW (chirp) waveform sweeps frequency linearly across the band,

$$s_{FMCW}(t) = A\,cos[2\pi f_c t + \pi\mu t^2] \tag{6}$$

where μ = B/T is the chirp rate. Its spectrum is nearly rectangular and its matched filter is well understood; its weakness is determinism, not spectral occupancy.

3.2. Noise radar

The band-limited noise-radar waveform replaces the chaotic state with a filtered Gaussian process n(t) [23,24]:

$$s_{noise}(t) = A\,cos\left[2\pi f_c t + 2\pi k_f \int_0^t n(\tau)d\tau\right] \tag{7}$$

Although Equations (7) and (1) are structurally identical, the two are fundamentally different. In noise radar, n(t) is a stochastic process: every pulse is an independent realization that cannot be reproduced, so the matched-filter reference must be obtained by capturing a copy of the transmitted signal. In chaotic radar, x(t) is deterministic: it is generated on demand at the receiver from the shared initial conditions and parameters, giving a noise-like waveform with a reproducible matched-filter replica. The price of this determinism is that the receiver must track the same trajectory as the transmitter; robust estimation of initial conditions [25] and re-synchronization under channel impairments have been demonstrated for chaos-based FM radar, including bistatic geometries [20,26]. Chaotic FM therefore seeks the interception resistance of noise radar while retaining the practical reference-generation of a deterministic waveform. Figure 4 contrasts the instantaneous frequency and spectra of the FMCW and noise references.

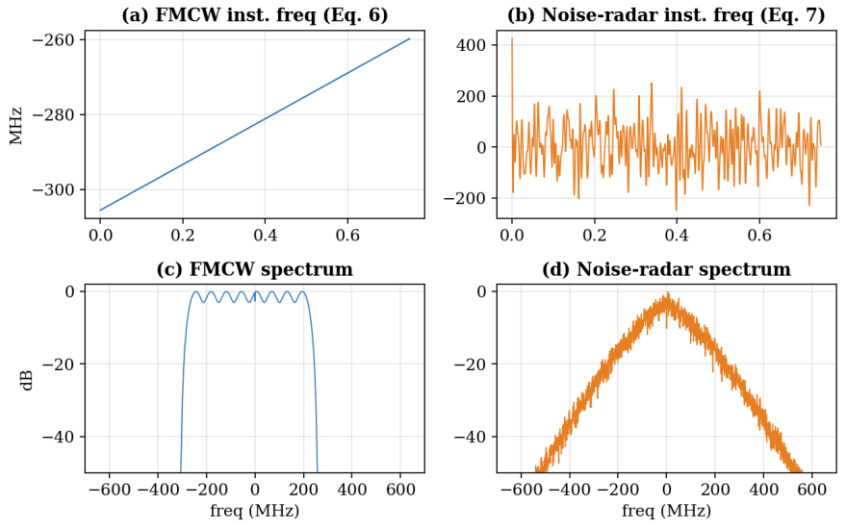


Figure 4. FMCW (Equation (6) *versus* noise-radar (Equation (7)) references: **(a**, **b)** Instantaneous frequency; **(c, d)** Power spectra. The FMCW frequency ramps linearly, whereas the noise waveform wanders randomly across the band.

All three waveform families share the transmitter architecture of Figure 5. For CBFM, a chaotic oscillator integrated by RK4 produces the modulating state, which is normalized, scaled by k_f, and applied to an FM modulator before up-conversion and amplification. The cooperative receiver regenerates the identical chaotic replica from the shared parameter set (x_0, σ, ρ, β).

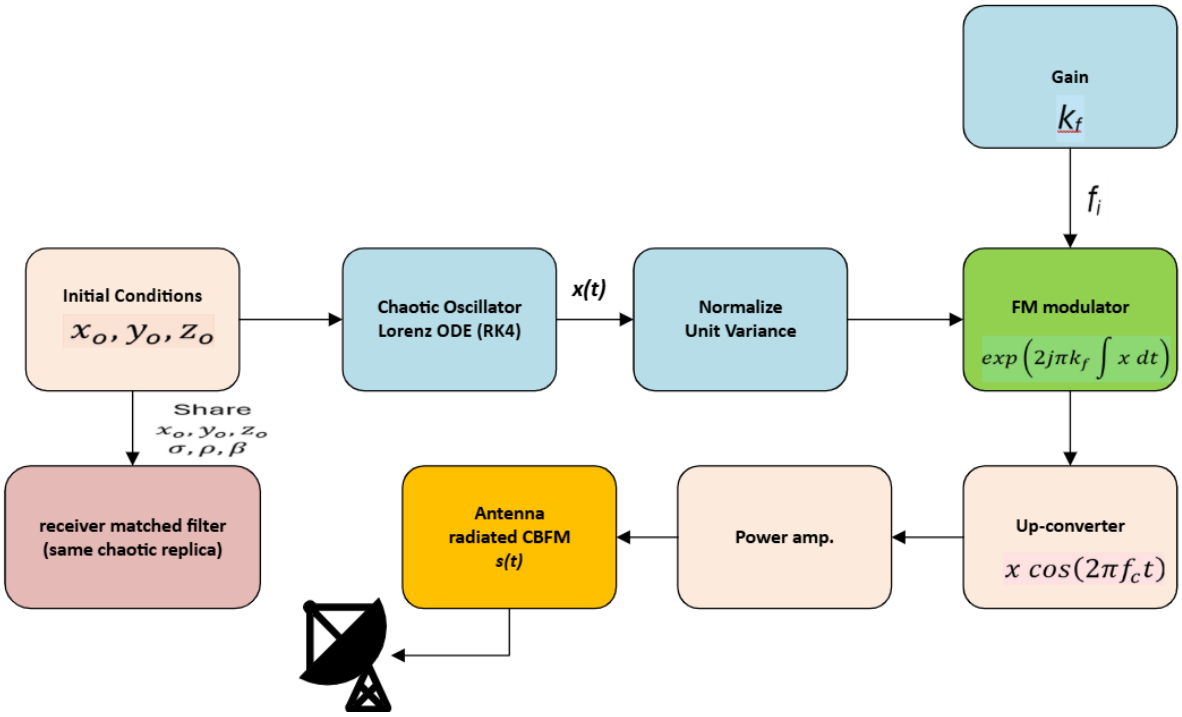


Figure 5. Chaos-based FM radar transmitter. The shared chaotic parameters allow the matched-filter replica to be regenerated at the cooperative receiver without transmitting a reference copy.

4. LPI and radar performance metrics

4.1. Interceptability index

The interceptability index *II* quantifies how easily a passive interceptor can detect the radar. Following the standard LPI formulation [1,3], it is the ratio of the signal-to-noise ratio (SNR) available to the intercept (ESM) receiver to the SNR the radar itself needs to detect its target,

$$II = \frac{SNR_{ES}}{SNR_R} \tag{8}$$

To make this operational, we compare two detection ranges. The radar detects a target over a two-way path, so its maximum detection range is

$$R_R = \sqrt[4]{\frac{P_t G_t G_r \lambda^2 \sigma\, G_p}{(4\pi)^3 k T_0 B_R F\, L\, SNR_R}} \tag{9a}$$

whereas the interceptor detects the radar's transmission over a one-way path,

$$R_{ES} = \sqrt{\frac{P_t G_t G_{ES} \lambda^2}{(4\pi)^2 k T_0 B_{ES} F\, L\, SNR_{ES}}} \tag{9b}$$

and the interceptability index is their ratio,

$$II = \frac{R_{ES}}{R_R} \tag{9c}$$

The symbols are defined as follows: P_t is the transmit power; G_t and G_r the radar transmit and receive antenna gains; G_{ES} the intercept-receiver antenna gain; λ the wavelength; σ the target radar cross-section; G_p the radar coherent processing gain (the product of bandwidth, dwell time, and number of integrated pulses); k Boltzmann's constant; T_0 the reference noise temperature; B_R and B_{ES} the radar and intercept bandwidths; F the noise figure; L the system losses; and SNR_R, SNR_{ES} the SNRs required for radar and intercept detection, respectively. The radar enjoys two-way (R^4) propagation loss but recovers it through the coherent processing gain G_p; the interceptor suffers only one-way (R^2) loss but, lacking the waveform, gains no coherent integration.

Reference scenario: To attach a concrete value, we adopt a representative covert-radar scenario: $P_t = 1$ W, $G_t = G_r = 30$ dBi, $G_{ES} = 0$ dBi (omnidirectional intercept antenna), $\lambda = 3$ cm (10 GHz), $\sigma = 1$ m², coherent integration over a coherent processing interval (CPI) of 10^4 pulses ($G_p = N_p \cdot BT \approx 5 \times 10^7$), $F = 3$ dB, $L = 3$ dB, and $SNR_R = SNR_{ES} = 13$ dB. These values give $R_R \approx 19.4$ km and $R_{ES} \approx 6.0$ km, hence $II \approx 0.31$. Because $II < 1$, the radar detects its target well before the interceptor detects the radar—the fundamental LPI condition is satisfied.

Why II cannot discriminate the four waveforms: Equations (9a)–(9c) contain no term that depends on the modulation: at equal transmit power, bandwidth, and geometry, II is identical for FMCW, noise, and both chaotic waveforms (the common value, 0.31, is reported for all four waveforms in Section 6.4). The radiometric index therefore measures the system (power, bandwidth, antenna, processing-gain) design, not the waveform. This is precisely why a complementary, modulation-sensitive metric is needed—the IPCG introduced next—which exposes the differences that II cannot. Transmit-power management can drive II still lower at the cost of radar range [3].

4.2. Inter-pulse correlation gain

The radiometric index of Section 4.1 is blind to the modulation: at equal transmit power, bandwidth and geometry it returns the same value for every waveform considered here. A capable interceptor, however, rarely relies on energy detection alone; it exploits structure. This subsection introduces a complementary, modulation-sensitive metric that captures the dominant structural exploit (Section 4.2.1), and then converts that metric into the quantity an operator actually cares about—the range at which the interceptor can detect the radar (Section 4.2.2).

4.2.1. Definition and interpretation

The strongest practical exploit against a repetitive waveform is replica reuse. The interceptor captures one clean pulse, stores it, and uses it as a matched-filter template against every pulse that follows. It requires no knowledge of the radar's design—only a recording. The value of the attack depends entirely on how similar consecutive pulses are. If the radar repeats itself, the stored template matches every future pulse, and the interceptor recovers the same coherent integration gain that the radar itself enjoys. If every transmitted pulse differs, the template is worthless against what follows and the interceptor gains nothing by listening longer.

We quantify this with the IPCG, defined as the peak normalized cross-correlation between two consecutive pulses,

$$IPCG = \frac{max\,|\int s_2(t)\, s_1{}^*(t-\tau)\, dt\,|}{\sqrt{E_1\, E_2}} \tag{10}$$

where $s_1(t)$ and $s_2(t)$ are consecutive transmitted pulses, E_1 and E_2 are their energies, and * denotes complex conjugation. The maximization over the delay τ makes the measure independent of any timing offset the interceptor applies, so IPCG reports the best alignment available to the attacker rather than a particular one. Normalization by the pulse energies confines IPCG to the interval [0, 1] and renders it independent of transmit power, so it characterizes the modulation alone.

The two limits set the scale. IPCG = 1 denotes a perfectly reusable waveform: consecutive pulses are identical, so one captured pulse matches every subsequent pulse and replica reuse is fully effective. IPCG → 0 denotes an aperiodic waveform for which a captured pulse carries no usable information about the next. Linear FMCW sits at the first limit by construction, since every sweep is a repetition of the last. Band-limited noise approaches the second, its non-zero floor set by the finite BT rather than by any residual determinism. Chaotic FM is of interest precisely because it is deterministic—and therefore regenerable at a cooperative receiver—while its trajectory never repeats, so it should approach the noise floor without inheriting the noise waveform's inability to reproduce a replica. Throughout this work, IPCG is evaluated between successive pulses of the same calibrated waveform under the calibrated conditions of Section 5, so that the values reported in Section 6.4 are directly comparable across the four waveform families.

4.2.2. From IPCG to interceptor coherent gain and interception range

Equation (10) is a property of the waveform alone: a dimensionless correlation that, by itself, does not state how much protection the waveform buys. We therefore convert it into interception range in three steps—the correlation between the interceptor's stored template and each received pulse, the coherent processing gain obtained by combining many such pulses, and the range that gain implies. Symbols introduced here are collected in Table 1; those already defined in Section 4.1 retain their meaning.

Table 1. Notation.

Symbol	Meaning
M	Number of pulses the interceptor captures and combines
$s_1(t)$	The pulse the interceptor captures first and stores as its template
$s_k(t)$	The k-th pulse transmitted by the radar, k = 1 … M
γ_k	Complex normalized correlation between pulse k and the stored template s_1
ρ	Shorthand for the IPCG of Equation (10); bounds $\|\gamma_k\|$ for $k \geq 2$
A	Complex amplitude of the pulse at the interceptor's correlator
n_k, σ_n^2	Receiver noise on correlator output k, and its power
$G_{int}(M)$	Coherent processing gain the interceptor obtains from M pulses
$R_{ES}(M)$	Interception range after combining M pulses; $R_{ES}(1)$ is the single-pulse value
II_{eff}	Effective interceptability index after combining, $R_{ES}(M)/R_R$

Correlation between template and received pulses: When the interceptor correlates the *k-th* received pulse against its stored template, the correlator output is a signal term plus receiver noise,

$$y_k = A\gamma_k + n_k \tag{11}$$

in which the factor γ_k is the normalized correlation between the transmitted pulse and the template,

$$\gamma_k = \frac{\int s_k(t).s_1^*(t)\,dt}{\sqrt{E_k\,E_1}} \tag{12}$$

This quantity answers a single question: how much of pulse k does the stored template actually capture? Two properties follow directly from the definition. For $k = 1$ the template is compared with the very pulse from which it was copied, so the match is exact and $|\gamma_1| = 1$. For every later pulse the match cannot exceed the peak correlation between two distinct pulses, which is precisely the quantity defined in Equation (10); writing ρ = IPCG, this gives $|\gamma_k| \le \rho$ for $k \ge 2$. A perfectly periodic waveform therefore has $|\gamma_k| = 1$ for all k, whereas an aperiodic waveform has $|\gamma_k| \approx 0$ beyond the captured pulse.

Coherent gain from combining M pulses. The interceptor holds M correlator outputs and combines them into a single decision. We grant it the optimal strategy, maximum-ratio combining, in which each output is weighted by the conjugate of its own signal coefficient so that strongly matching pulses contribute fully while poorly matching ones are down-weighted rather than injecting noise. The resulting figure is therefore an upper bound on what any replica-reuse attack can achieve, which keeps the comparison conservative in the radar designer's favour. Under maximum-ratio combining the output SNR is the sum of the individual branch ratios, the *k-th* branch carrying

$$SNR_k = \frac{|A|^2.|\gamma_k|^2}{\sigma_n^2} \tag{13}$$

Summing over the M captured pulses and referring the result to the interceptor working from its single captured pulse alone, for which $SNR_1 = |A|^2/\sigma_n^2$ because $|\gamma_1| = 1$, the coherent processing gain is simply the sum of the squared correlations. Substituting $|\gamma_1| = 1$ and $|\gamma_k| \le \rho$ for the remaining $M - 1$ pulses gives the central result,

$$G_{int}(M) = \sum_k |\gamma_k|^2 = 1 + (M-1).IPCG^2 \tag{14}$$

The structure of Equation (14) is easily read. The leading unity is the captured pulse, which always matches itself. The term $(M - 1)\cdot IPCG^2$ is the contribution of every other pulse, each entering only in proportion to the squared correlation with the template. The expression behaves correctly at both limits: for a perfectly periodic waveform (IPCG = 1) it reduces to $G_{int} = M$, the full coherent integration gain the radar itself would obtain, whereas for a perfectly aperiodic waveform (IPCG = 0) it reduces to $G_{int} = 1$, meaning the interceptor gains nothing whatever by listening beyond the first pulse.

Conversion to interception range: A processing gain is not yet a distance. The conversion follows from the one-way nature of the intercept link: the interceptor listens directly to the radar's transmission, so received power falls as $1/R^2$, in contrast to the radar's two-way $1/R^4$ path. Detection occurs when the received power reaches the receiver's sensitivity; if coherent processing lowers the required power by the factor G_{int}, then R^2 may grow by the same factor and the range itself by its square root,

$$R_{ES}(M) = R_{ES}(1).\sqrt{G_{int}(M)} \quad (15)$$

For large *M,* the leading unity in Equation (14) becomes negligible and $G_{int} \approx M \cdot IPCG^2$. Forming the ratio of the interception ranges obtained against a periodic waveform (IPCG = 1) and against an aperiodic one, the pulse count cancels and a design rule independent of dwell length emerges,

$$\frac{R_{ES}(\text{periodic})}{R_{ES}(\text{aperiodic})} = \frac{\sqrt{M}}{\sqrt{(M.IPCG^2)}} = \frac{1}{IPCG} \quad (16)$$

The factor by which a chaotic waveform shortens the interceptor's coherent detection range is thus simply the reciprocal of its IPCG, irrespective of how long the interceptor listens. This is the step that converts an abstract correlation into an operational quantity.

Reference distances: The distances used below are not assumed but computed: they are the outputs of Equations (9a) and (9b) under the reference scenario of Section 4.1, and are reproduced here so that every subsequent figure can be verified. Substituting P_t = 1 W, $G_t = G_r$ = 30 dBi (10^3 each), G_{ES} = 0 dBi, λ = 0.03 m, target cross-section 1 m², $G_p = 5\times10^7$, $B_R = B_{ES}$ = 500 MHz, F = L = 3 dB (2 each), $SNR_R = SNR_{ES}$ = 13 dB (19.95), $k_B = 1.381\times10^{-23}$ J/K and T_0 = 290 K, the two-way radar range of Equation (9a) evaluates as

numerator = $P_t \cdot G_t \cdot G_r \cdot \lambda^2 \cdot \sigma \cdot G_p = 1 \cdot 10^3 \cdot 10^3 \cdot (0.03)^2 \cdot 1 \cdot 5 \times 10^7 = 4.500 \times 10^{10}$

denominator = $(4\pi)^3 \cdot k_B \cdot T_0 \cdot B_R \cdot F \cdot L \cdot SNR_R = 3.156 \times 10^{-7}$

$R_R = (1.426 \times 10^{17})^{\wedge}(1/4) = 19433$ m ≈ 19.4 km

while the one-way intercept range of Equation (9b), for which the target terms and one factor of 4π drop out, evaluates as

numerator = $P_t \cdot G_t \cdot G_{ES} \cdot \lambda^2 = 1 \cdot 10^3 \cdot 1 \cdot (0.03)^2 = 0.900$

denominator = $(4\pi)^2 \cdot k_B \cdot T_0 \cdot B_{ES} \cdot F \cdot L \cdot SNR_{ES} = 2.511 \times 10^{-8}$

$R_{ES}(1) = (3.584 \times 10^7)^{\wedge}(1/2) = 5987$ m ≈ 6.0 km

These give *II* = 5987 / 19433 = 0.308, quoted throughout as 0.31. Because *II* < 1, the radar detects its target before a single-pulse interceptor detects the radar, and the LPI condition is satisfied—but only against an interceptor that does not exploit structure. It should be noted that the absolute distances depend on the assumed power, gains and dwell, whereas the conclusions of this subsection do not: Equations (15) and (16) involve only ratios of ranges, so the central result is independent of the particular scenario, which serves only to supply reproducible numbers.

Worked example: Evaluating Equations (14) and (15) at the reference dwell of $M = 10^4$ pulses—the same CPI already assumed for the radar in Section 4.1—and taking Lorenz-CBFM, whose IPCG is 0.072 (Section 6.4), the arithmetic runs

$IPCG^2 = 0.072^2 = 0.005184$

$G_{int} = 1 + (10^4 - 1) \cdot 0.005184 = 1 + 51.83 = 52.83$ (17.2 dB)

$R_{ES} = 6.0$ km · √52.83 = 6.0 · 7.269 = 43.6 km

The same three lines applied to FMCW, whose IPCG is exactly unity, give $G_{int} = 1 + 9999 = 10{,}000$ (40.0 dB) and R_ES = 6.0 · 100 = 600 km. The results for all four waveforms are collected in Table 2, with every intermediate quantity shown so that each row may be checked independently.

Table 2. Interceptor performance at the reference dwell M = 10^4 pulses.

Waveform	IPCG (ρ)	ρ^2	$(M-1)\,\rho^2$	G_{int}	G_{int} (dB)	$\sqrt{G_{int}}$	R_{ES} (km)	II_{eff}
FMCW	1.000	1.000000	9 999.0	10 000	40.0	100.0	600.0	30.9
Rössler-CBFM	0.207	0.042849	428.4	429.4	26.3	20.7	124.3	6.40
Lorenz-CBFM	0.072	0.005184	51.8	52.8	17.2	7.27	43.6	2.24
Noise radar	0.063	0.003969	39.7	40.7	16.1	6.38	38.3	1.97

$II_{eff} = R_{ES}(M) / R_R$ with R_R = 19.4 km.

The FMCW entry is the significant one. Over a dwell of 10^4 pulses a replica-reuse interceptor extracts 40 dB of coherent gain and extends its detection range to 600 km, an effective interceptability index of 30.9. Since $II_{eff} \geq 1$ means the interceptor sees the radar before the radar sees its target, the LPI condition of Section 4.1 is not merely eroded but comprehensively violated—a failure the radiometric index alone cannot reveal, since it reports 0.31 for FMCW exactly as it does for every other waveform.

Lorenz-CBFM holds the same interceptor to 17.2 dB and 43.6 km. Its effective index does rise above unity (2.24), so a patient interceptor still gains something: chaotic modulation reduces the replica-reuse threat by more than an order of magnitude rather than eliminating it. The resulting range advantage is 600 / 43.6 = 13.8×, precisely the 1/IPCG = 1/0.072 = 13.9× predicted by Equation (16), the small residual difference being the leading unity of Equation (14) that the large-M limit neglects. This calculation gives the waveform-level factor of roughly fourteen a concrete system-level meaning: it is the factor by which chaotic modulation shortens the range at which a structure-exploiting interceptor can coherently detect the emitter. Figure 6 shows both quantities as a function of the number of pulses the interceptor combines.

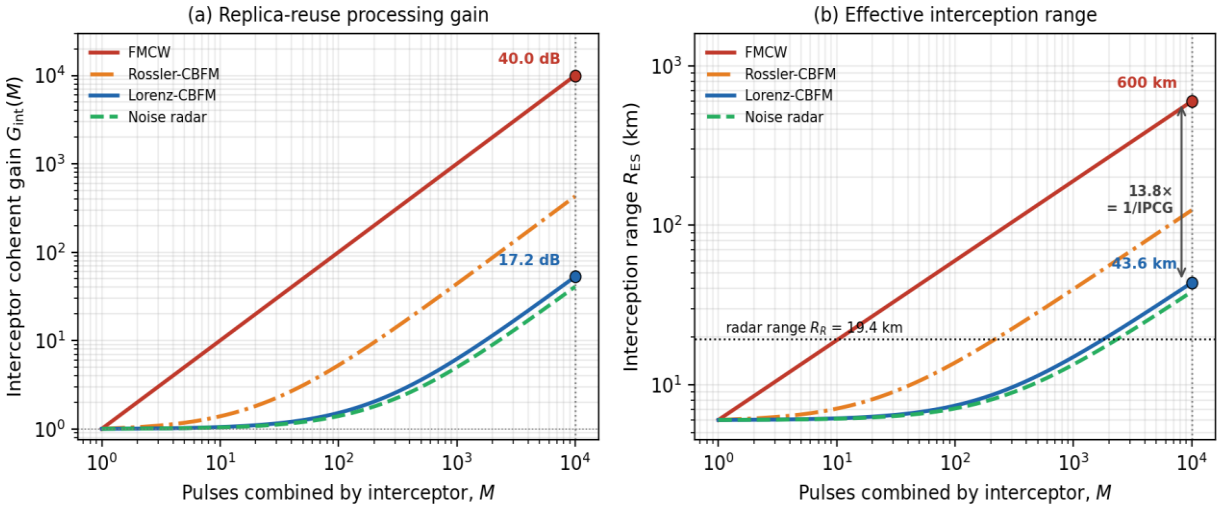


Figure 6. Quantitative bridge from IPCG to interception performance. **(a)** Interceptor coherent processing gain Gint(M) = 1 + (M − 1)·IPCG² of Equation (14), as a function of the number of pulses M captured and combined; **(b)** The corresponding effective interception range RES(M) = RES(1)·√Gint(M) of Equation (15); the horizontal dotted line marks the radar's own detection range RR = 19.4 km, above which the interceptor detects the radar before the radar detects its target. Filled markers on the vertical dotted line indicate the reference dwell M = 10^4. Over that dwell the interception range against FMCW exceeds that against Lorenz-CBFM by a factor of 13.8 (600 km *versus* 43.6 km), equal to 1/IPCG as predicted by Equation (16).

The remaining rows bracket the result as expected. Rössler-CBFM, whose smoother single-scroll trajectory yields a higher IPCG, permits a 124 km interception range; band-limited noise radar, the theoretical ideal, holds the interceptor to 38.3 km. Lorenz-CBFM at 43.6 km lies within 14% of that ideal while remaining a deterministic waveform whose matched-filter replica can be regenerated at the cooperative receiver from the shared parameter set.

Three qualifications delimit the model. First, it is an upper bound on the interceptor rather than a prediction: maximum-ratio combining presumes knowledge of the correlation magnitudes needed to weight the branches, which a real interceptor must estimate and will therefore approach only imperfectly. Second, it addresses one specific attack—replica reuse, the exploit the radiometric index cannot represent—and does not alter the radiometric budget itself, which remains modulation-independent; nor does it describe an interceptor that attempts short-horizon prediction of the chaotic state instead of reusing a stored replica, a threat evaluated experimentally in Section 7.1. Third, the IPCG values entering the table are simulated results obtained under the calibrated conditions of Section 5, so the model inherits the simulation-based limitations noted in Section 9.

4.3. Spectral flatness

A noise-like waveform should fill its band uniformly. We measure in-band spectral flatness (the Wiener entropy, equivalently the spectral flatness measure, SFM) as the ratio of the geometric to the arithmetic mean of the power spectral density over the occupied band; an equivalent normalized spectral entropy is

$$H_s = -\sum_i p_i \log_2 p_i \tag{17}$$

where $p_i = P(f_i)/\Sigma_k P(f_k)$ is the normalized power spectral density at frequency bin i (so that the p_i form a probability distribution summing to one). SFM = 1 corresponds to a perfectly white spectrum and SFM $\rightarrow$ 0 to a pure tone [27].

4.4. Autocorrelation peak-sidelobe ratio

When the matched filter processes the echo, the waveform's own autocorrelation determines how a strong target spreads in range. Range sidelobes are the off-peak lobes of the normalized autocorrelation (matched-filter output),

$$R(\tau) = \frac{\int s(t)\, s^*(t-\tau)\, dt}{\int |s(t)|^2\, dt} \tag{18}$$

where s^* denotes the complex conjugate and the denominator normalizes the mainlobe to unity at $\tau = 0$. PSR is the largest sidelobe relative to the mainlobe,

$$PSR = 20 \log_{10}\left[\frac{max\,|R(\tau)|}{|R(0)|}\right] \tag{19}$$

evaluated over all delays outside a narrow mainlobe guard region (here ±10 samples). A low PSR is important because a strong target's sidelobes can mask a nearby weak target at the same Doppler; for example, a PSR of −20 dB means sidelobes are a hundredfold below the mainlobe in power. Two practical points govern the measurement. First, the matched-filter response is computed over the finite 10 µs pulse, so a rectangular observation window would introduce spectral-leakage lobes that are an artifact of

truncation rather than of the waveform; we therefore apply a Hann (raised-cosine) window before spectral estimation, which tapers the record smoothly to zero at both ends and suppresses that leakage, leaving a sidelobe floor that reflects the waveform itself. Second, because the chaotic and noise waveforms are wideband and aperiodic, their sidelobe floor is governed by BT; the resulting PSR values are reported in Section 6.2.

4.5. Woodward ambiguity function

Joint range–Doppler behavior is described by the Woodward ambiguity function [28],

$$|\chi(\tau, f_d)|^2 = \left[\int s(t)\, s^*(t-\tau)\, e^{j2\pi f_d t}\, dt\right]^2 \quad (20)$$

which measures the matched-filter response to a target displaced by delay τ and Doppler f_d. An ideal LPI waveform exhibits a *thumbtack* ambiguity function: a single sharp central peak at ($\tau = 0$, $f_d = 0$) with a low, uniform pedestal elsewhere. This indicates high resolution in both range and Doppler simultaneously, with no range–Doppler coupling. By contrast, the linear-FMCW ambiguity function is a diagonal ridge: a target's delay and Doppler are coupled, so a Doppler shift masquerades as a range error. These surfaces are compared in Section 6.3.

5. Simulation setup

All experiments are simulation-based and fully reproducible. Waveforms are synthesized at complex baseband with the parameters of Table 3. The chaotic ordinary differential equations are integrated with the classical fourth-order Runge–Kutta scheme. Concretely, each state update advances the system by combining four slope evaluations (one at the interval start, two at the midpoint, one at the end) in the weighted average that gives RK4 its fourth-order accuracy; the integration step is $\Delta t = 0.5$ ns, matched to the 2 GHz sampling grid. After integration the state is resampled onto the radar time base, normalized to zero mean and unit variance, and scaled by the calibrated gain k_f. A Hann window is used for all power-spectral-density and autocorrelation estimates to control spectral leakage, as noted in Section 4.4.

Table 3. Simulation parameters.

Parameter	Symbol	Value
Sampling frequency	f_s	2 GHz
Pulse/integration time	T	10 μs
Samples per pulse	N	20,000
Target occupied bandwidth	B	500 MHz
Time–bandwidth product	BT	5000
RK4 integration step	Δt	0.5 ns
Lorenz parameters	σ, ρ, β	10, 28, 8/3
Rössler parameters	a, b, c	0.2, 0.2, 5.7
Spectral window	—	Hann

6. Results and analysis

6.1. Power spectral density

Figure 7 compares the power spectra of all four calibrated waveforms—FMCW, noise radar, Lorenz-CBFM, and Rössler-CBFM—occupying the same 500 MHz band. The FMCW spectrum is the flattest (SFM = 0.971) because a linear sweep dwells equally at every frequency; the chaotic spectra are slightly less uniform (SFM = 0.821 for Lorenz, 0.816 for Rössler) owing to the attractor's preferred regions, but remain markedly more noise-like than any tonal or stepped waveform. The noise-radar spectrum, although stochastically flat on average, is in fact the least uniform on a per-pulse basis (SFM = 0.537): a single realization exhibits strong spectral speckle that varies from pulse to pulse. A useful and perhaps counter-intuitive observation follows—the deterministic chaotic waveforms present a smoother, more repeatable in-band spectrum per pulse than a single noise-radar pulse, while still appearing noise-like to an interceptor.

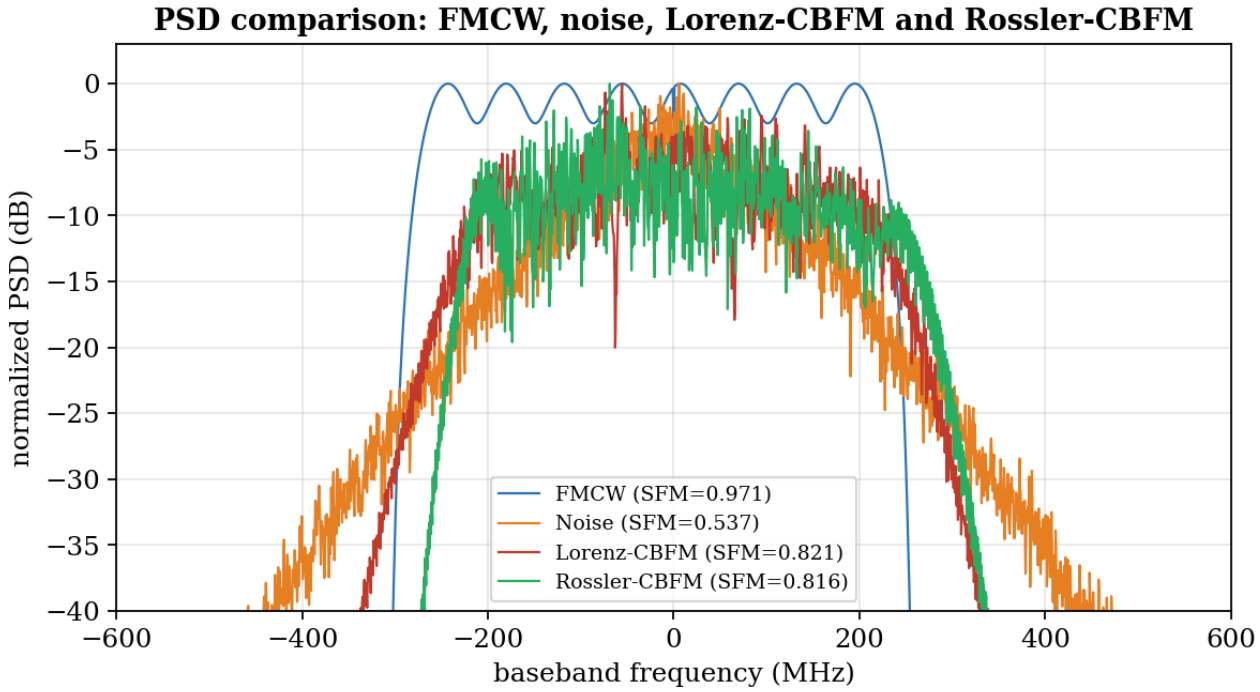


Figure 7. Power spectral density of all four waveforms—FMCW, noise radar, Lorenz-CBFM, and Rössler-CBFM—each occupying 500 MHz. In-band (SFM) is annotated; the single noise realization shows the strongest spectral speckle.

6.2. Autocorrelation and peak-sidelobe ratio

Figure 8 shows the matched-filter responses. FMCW achieves the lowest sidelobes (PSR = −21.4 dB), a direct consequence of its perfectly flat spectrum. Lorenz-CBFM follows at −17.8 dB and band-limited noise at −18.9 dB, while Rössler-CBFM is weakest at −14.1 dB because its single-scroll attractor produces a smoother, more correlated trajectory. Chaotic FM therefore does not improve range sidelobes relative to FMCW; its advantage lies elsewhere.

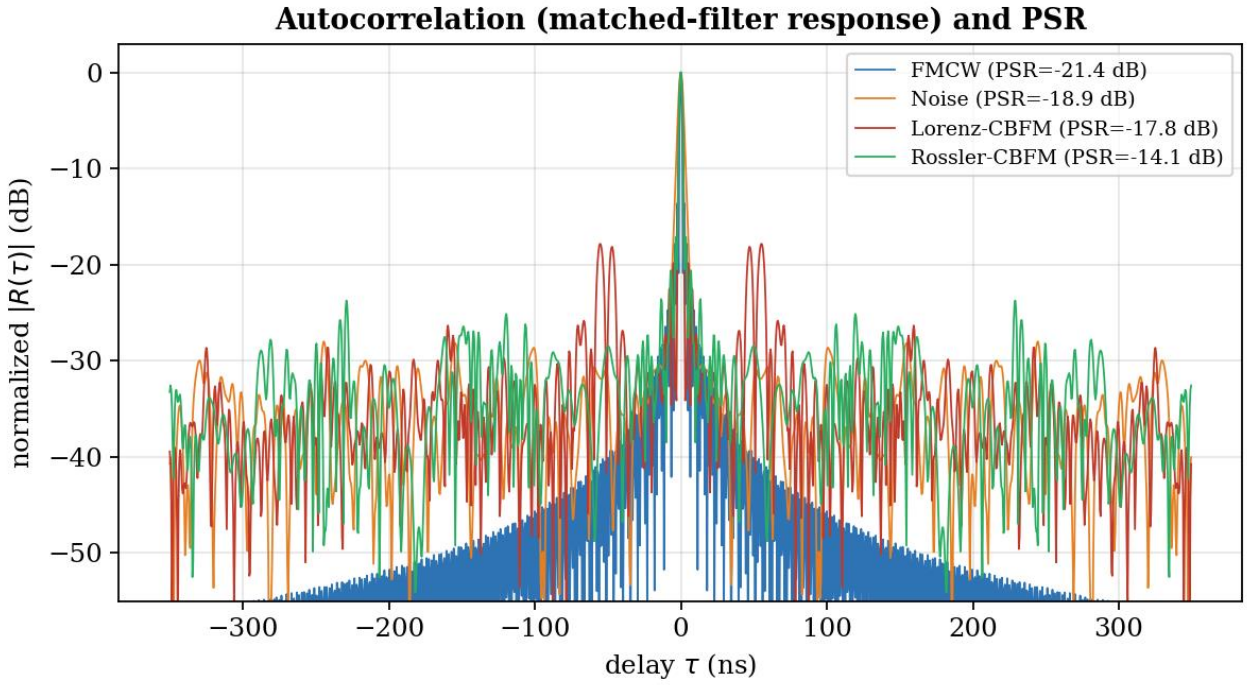


Figure 8. Autocorrelation (matched-filter) responses and peak-sidelobe ratios. FMCW has the lowest sidelobes; Lorenz-CBFM is within ~3.5 dB, and Rössler-CBFM is highest.

6.3. Ambiguity function

Figure 9 is where the waveform families diverge sharply. The FMCW ambiguity surface is a diagonal ridge: range and Doppler are coupled, so an uncompensated Doppler shift biases the range estimate and the operator must either scan a bank of Doppler hypotheses or accept a range–velocity ambiguity. The Lorenz-CBFM surface is a thumbtack—a single sharp central spike surrounded by a low, uniform pedestal. Several concrete benefits follow. (i) Range and Doppler are estimated independently and without bias, because the peak does not tilt with f_d; a fast target is measured at its true range. (ii) Resolution is simultaneously fine in both dimensions, set by the bandwidth in range and by the pulse length in Doppler, so the same waveform supports high-resolution ranging and accurate velocity estimation without a dedicated mode. (iii) The thumbtack is the matched-filter signature of a near-ideal noise-like waveform, which is exactly the property that yields the low inter-pulse correlation of Section 6.4—so the radar-performance advantage and the LPI advantage share a common origin in the waveform's aperiodicity. (iv) The low, diffuse pedestal spreads clutter and interference energy uniformly rather than concentrating it along a ridge, improving detection of small targets in clutter. The trade-off is that the pedestal raises the average sidelobe floor relative to a pulse-compression chirp, consistent with the PSR results of Section 6.2; for targets that are not Doppler-shifted, FMCW's ridge still yields lower sidelobes. The thumbtack is therefore the principal radar-performance benefit of chaotic modulation, and it is most valuable precisely when targets are fast or when Doppler is unknown.

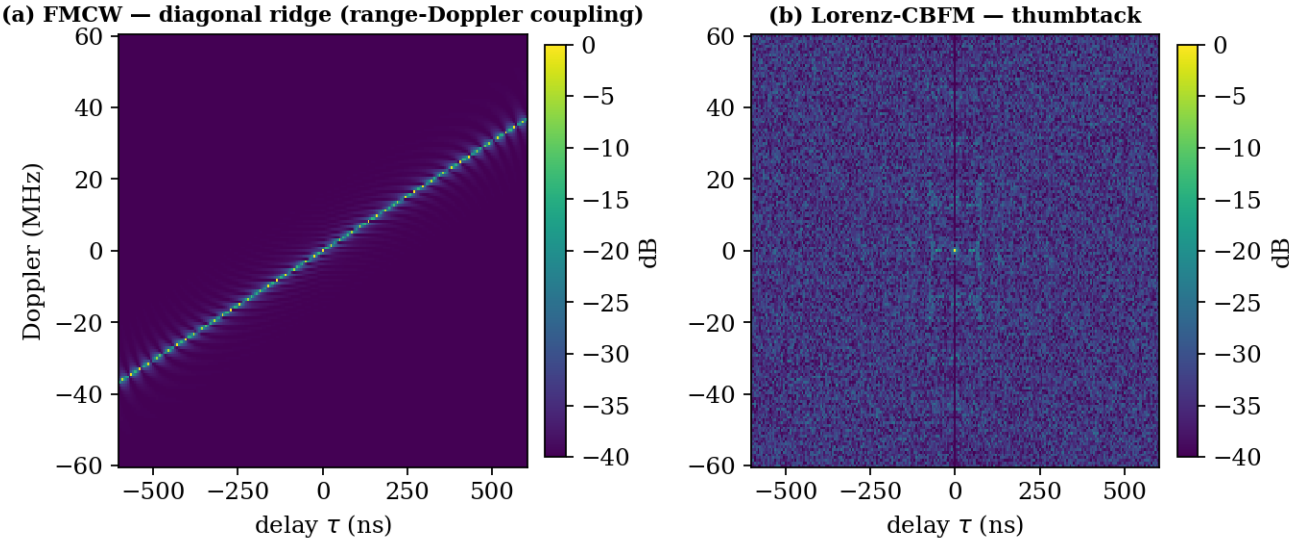


Figure 9. Ambiguity functions. **(a)** FMCW exhibits the characteristic range–Doppler ridge; **(b)** Lorenz-CBFM exhibits a thumbtack surface with no range–Doppler coupling.

6.4. Interceptability

Table 4 summarizes all metrics. Consider first the radiometric interceptability index *II* of Section 4.1. For the reference scenario it equals 0.31 for every waveform: against an energy-detecting interceptor, FMCW, noise, and both chaotic waveforms are equally (un)detectable, because *II* is fixed by power, bandwidth, antenna gain, and processing gain rather than by the modulation. This invariance is the key motivation for IPCG. Where *II* is blind, IPCG is decisive. FMCW is perfectly periodic: consecutive sweeps are identical, so a single captured sweep coherently matches every subsequent one (IPCG = 1.000). The chaotic waveforms never repeat—Lorenz-CBFM falls to 0.072 and Rössler-CBFM to 0.207—while band-limited noise reaches 0.063. Lorenz-CBFM thus approaches the interception resistance of true noise radar (within a factor of 1.1 of its IPCG) while retaining a deterministic, regenerable replica. Operationally, an interceptor that has captured a full Lorenz-CBFM pulse gains roughly 14 times less coherent processing against the following pulse than against FMCW; Section 4.2.2 and Figure 6 make this factor precise as a ratio of interception ranges.

Table 4. Waveform comparison (calibrated to 500 MHz, BT = 5000).

Waveform	Occupied BW (MHz)	SFM (flatness)	PSR (dB)	*II* (Section 4.1)	IPCG
FMCW	500	0.971	−21.4	0.31	1.000
Noise radar	500	0.537	−18.9	0.31	0.063
Lorenz-CBFM	500	0.821	−17.8	0.31	0.072
Rössler-CBFM	500	0.816	−14.1	0.31	0.207

The interceptability index *II* (radiometric, Section 4.1 reference scenario) is identical for all four waveforms because it is modulation-independent—hence it cannot rank them, which is precisely why IPCG is introduced. FMCW is best on flatness and sidelobes; noise and Lorenz-CBFM are best on interceptability (IPCG). Lorenz-CBFM offers the best overall balance.

6.5. Detection and parameter-estimation performance

The comparison presented so far has been deliberately one-sided: every metric introduced in Section 4 measures how well a waveform resists interception, and none measures how well it performs the radar's own job. A waveform that hid perfectly but detected poorly would be of no practical use. This subsection therefore evaluates the four calibrated waveforms on the two tasks a radar must actually perform—detecting a target and estimating its range—and reports where the modulation matters and where it does not.

Two distinct results in this paper share one underlying cause. The first is the behaviour of the interceptability index *II* of Section 4.1, the ratio of the range at which a hostile receiver detects the radar to the range at which the radar detects its target. *II* takes the identical value 0.31 for all four waveforms because Equations (9a)–(9c) contain transmit power, antenna gains, wavelength, target cross-section, bandwidth, noise figure, losses and processing gain, but no term whatever describing the shape of the modulation. The second is the detection performance measured below, which likewise proves identical across the four waveforms: a matched filter applied to a pulse of energy E in white noise of power spectral density N_0 delivers an output SNR of $2E/N_0$ [2,28], an expression depending on total energy and not at all on how that energy is distributed in time or frequency. Where the results below are described as consistent with the equality of *II*, the statement means precisely this—the same energy-only dependence that prevents the radiometric index from distinguishing the four waveforms also prevents their detection performance from differing. Chaotic modulation is neither a detection advantage nor a detection penalty; its cost appears in the sidelobe-limited dynamic range of Section 6.2 and its benefit in interception resistance (Sections 4.2.2 and 7.1) and in the estimation behaviour reported below.

Detection experiment: method. All four waveforms of Table 3 are generated at complex baseband with N = 20,000 samples per pulse (T = 10 μs, B = 500 MHz), giving a coherent processing gain of N, *i.e.* 43.0 dB. A single pulse is embedded in additive white complex Gaussian noise and passed through its own matched filter, detection being declared when the output magnitude exceeds a threshold. The per-sample input SNR ρ is swept from −46 to −26 dB in 1.5 dB steps, corresponding to a matched-filter output ratio $\chi = \rho N$ of −3.0 to +17.0 dB. The threshold is set analytically rather than estimated: under the noise-only hypothesis the filter output is zero-mean complex Gaussian, so its magnitude is Rayleigh-distributed, $P_{fa} = exp(-T^2/2\sigma_R^2)$, and $T = \sigma_R\sqrt{2\ln(1/P_{fa})}$ [29,30]. This matters, because a false-alarm probability of 10^{-4} cannot be estimated reliably from a few hundred Monte-Carlo trials. The threshold law was verified independently over 20,000 noise-only trials at three milder design points, where the statistics are sound, as Table 5 shows:

Table 5. Analytic detection threshold $T = \sigma_R\sqrt{2\ln(1/P_{fa})}$ verified against 20,000 noise-only matched-filter trials per design point. Design values are milder than the 10^{-4} used in Figure 10a, which would need of order 10^6 trials to verify empirically.

Design P_{fa}	Measured P_{fa}	False alarms observed
1×10^{-1}	0.0990	1980 of 20,000
1×10^{-2}	0.0092	185 of 20,000
1×10^{-3}	0.0011	23 of 20,000

Detection probability is then measured from 2000 independent trials per waveform per sweep point at $P_{fa} = 10^{-4}$, and compared against the exact non-fluctuating detection law rather than a Gaussian approximation. For a signal of known amplitude and unknown phase in complex Gaussian noise the envelope under the signal hypothesis is Rician, and the detection probability is the first-order Marcum Q-function [2,31],

$$P_d = Q1\left(\sqrt{2\chi}, \sqrt{2\ln(1/P_{fa})}\right), \qquad \chi = \rho N \tag{21}$$

where χ is the matched-filter output SNR. Setting $\chi = 0$ returns $P_d = P_{fa}$, as it must.

Figure 10a presents the outcome. The four measured curves are mutually indistinguishable and lie on the theoretical curve of Equation (21) across the entire transition region; the largest discrepancy between any measured point and theory, over all four waveforms and all sweep points, is 0.026 in probability, within the sampling error of 2000 trials. Primary detection is therefore not a discriminator between these waveforms: at equal transmitted energy an operator gains nothing and loses nothing in detection sensitivity by replacing a chirp with a chaotic waveform. This is a necessary result to establish rather than an incidental one, because it removes the most obvious objection to chaotic modulation—that a noise-like transmission must somehow detect less well than a deterministic chirp. It does not.

Range-estimation experiment: method. Detection establishes that a target is present; estimation establishes where it is. A target of radial velocity v imposes a Doppler shift $f_d = 2v/\lambda$, and a receiver correlating against a zero-Doppler replica may then report the wrong delay. Each waveform is multiplied by $\exp(j2\pi f_d t)$, correlated against the undistorted replica, and the correlation peak located to sub-sample precision by fitting a parabola to the peak sample and its two neighbours; the apparent range error is $c\hat{\tau}/2$. The simulation is noiseless, so the measured error is the systematic effect of Doppler alone, uncontaminated by noise-induced peak jitter. The Doppler shift is swept from 0 to 100 kHz in 5 kHz steps, the upper limit being one Doppler resolution cell (1/T = 100 kHz) and corresponding at 10 GHz to a radial velocity of 1500 m/s. For a linear-FM waveform the delay–Doppler coupling is classical [2,3,28]:

$$\Delta R = c \cdot f_d / (2\mu), \ \mu = B/T = 5 \times 10^{13} \text{ Hz/s} \tag{22}$$

Equation (22) predicts 30 cm of range error at the edge of the Doppler cell, a figure best appreciated against the range resolution of the waveform, $c/2B$ = 30 cm: the predicted bias is exactly one full range-resolution cell.

Range-estimation results: Figure 10b shows the measured bias with the theoretical law overlaid. The FMCW measurements fall exactly on that line at every point, rising linearly to 30.0 cm at 100 kHz. The three noise-like waveforms behave qualitatively differently: none follows a linear coupling law at all. Table 6 lists the measured bias at five Doppler shifts.

Three observations follow: First, the FMCW error is systematic and deterministic: it obeys a closed-form law, has a fixed sign set by the sweep direction, and is identical on every pulse, so a Doppler-unknown target is mis-ranged by a predictable amount that accumulates coherently over a processing interval rather than averaging away. At 750 m/s the target is displaced by half a range cell and at 1500 m/s by a full cell. Second, the thumbtack waveforms exhibit no comparable law: Rössler-CBFM and the noise reference hold the bias below 0.2 cm—two orders of magnitude below FMCW—across the useful Doppler band, the expected signature of an ambiguity surface whose peak does not tilt with Doppler [14,17]. Lorenz-CBFM shows a small residual drift, about three times smaller than FMCW over

the lower half of the Doppler cell (ratio 2.7 to 3.2, mean 3.0), which is not a coupling law, as shown next. Third, near the cell edge the thumbtack waveforms degrade differently: because their sidelobe floors are higher (Section 6.2), the mainlobe eventually falls far enough for the global maximum to jump to a sidelobe, so Rössler-CBFM (worst PSR at −14.1 dB) loses peak lock beyond 85 kHz and the noise reference beyond 90 kHz. This is a loss of lock rather than a bias and, unlike the FMCW coupling, is recovered by a modest Doppler filter bank.

Table 6. Range-estimation bias versus target Doppler; one range-resolution cell = 30 cm. "Lock lost" denotes that the ambiguity mainlobe has degraded far enough for the global correlation maximum to jump to a sidelobe.

f_d (kHz)	Velocity (m/s)	FMCW	Lorenz-CBFM	Rössler-CBFM	Noise radar
20	300	6.0 cm (0.2 cell)	1.9 cm (0.06 cell)	0.0 cm	0.0 cm
30	450	9.0 cm (0.30 cell)	3.0 cm (0.10 cell)	0.1 cm	0.0 cm
50	750	15.0 cm (0.50 cell)	5.6 cm (0.19 cell)	0.1 cm	0.0 cm
80	1200	24.0 cm (0.8 cell)	16.1 cm (0.54 cell)	0.2 cm	0.2 cm
100	1500	30.0 cm (1.00 cell)	29.2 cm (0.97 cell)	lock lost	lock lost

Origin of the residual Lorenz bias: The small Lorenz drift would otherwise appear to contradict the thumbtack claim of Section 6.3, so we tested whether it is intrinsic to Lorenz-CBFM or an artifact of the particular trajectory segment, by repeating the measurement for six initial conditions and recording for each the net linear trend of the normalized chaotic state across the pulse. Table 7 collects the results.

Table 7. Residual Lorenz-CBFM range bias at f_d = 50 kHz for six initial conditions.

Initial condition	Net trend of $x(t)$ across the pulse	Bias at f_d = 50 kHz
(0.5, −2.0, 3.0)	+0.181	1.68 cm
(3.0, 2.0, 1.5)	+0.677	3.95 cm
(−1.0, 0.5, 1.0)	−0.811	4.35 cm
(1.0, 1.0, 1.0)	+0.928	5.57 cm
(2.0, 1.0, 1.0)	+1.108	6.29 cm
(1.0, 3.0, 2.0)	+1.415	8.39 cm

The residual bias varies from 1.68 to 8.39 cm across realizations and is a strictly monotone function of the net frequency trend contained in the segment: a finite window of a chaotic trajectory generally has a small non-zero mean slope, which acts as a weak residual chirp and produces a correspondingly weak coupling. The effect is therefore incidental to the segment rather than intrinsic to the attractor, and it carries an operational consequence. Because consecutive pulses of a CBFM waveform are different segments of the trajectory—the same aperiodicity that produces the low inter-pulse correlation of Section 6.4—the residual trend, and hence the residual bias, differs from pulse to pulse and averages down over CPI, whereas the FMCW bias is identical on every pulse and accumulates. The distinction is between an error that integrates away and one that does not.

Taken together the two panels of Figure 10 answer the question a radar engineer would put to this paper: what does the LPI benefit cost in radar performance? For these calibrated waveforms it costs nothing in detection and yields a net gain in estimation. The detection result of Figure 10a closes off the

concern that a noise-like waveform must detect less reliably than a deterministic one, and the agreement with exact theory across four waveforms and a 20 dB sweep makes that a measured statement rather than an assertion. The estimation result of Figure 10b shows that the thumbtack surface of Section 6.3 is not merely a pleasing contour plot but removes a systematic, non-averaging range bias that reaches a full resolution cell for FMCW against a fast target—the benefit being largest exactly where LPI radars are most often deployed, against targets whose velocity is unknown or unresolved. The two results also close a logical loop: Section 4.1 could not distinguish the waveforms because the radiometric index depends only on energy, and detection is blind for the same reason. What separates the waveforms is not how much energy they deliver but how that energy is structured—which is what the IPCG measures for the interceptor, and what the ambiguity surface measures for the radar's own estimator.

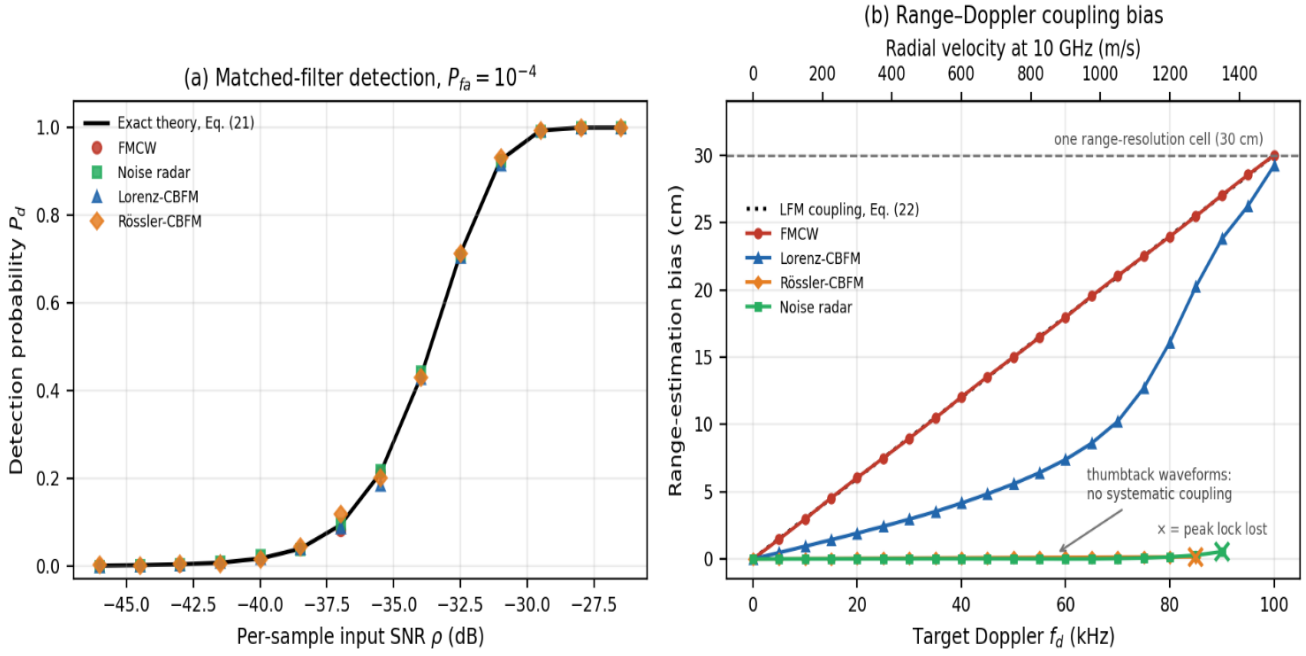


Figure 10. Detection and parameter-estimation performance **(a)** Matched-filter detection probability against per-sample input SNR at $P_{fa} = 10^{-4}$, from 2000 Monte-Carlo trials per point with an analytic Rayleigh threshold. All four waveforms coincide with the exact non-fluctuating detection law of Equation (21) (largest deviation 0.026), confirming that detection depends on pulse energy alone and not on the modulation; **(b)** Range-estimation bias against target Doppler, noiseless, with sub-sample parabolic peak interpolation; the upper axis gives the corresponding radial velocity at 10 GHz and the dashed line marks one range-resolution cell (30 cm). FMCW follows the linear-FM (LFM) coupling law of Equation (22) exactly, reaching a full resolution cell at the edge of the Doppler cell, whereas the thumbtack waveforms show no systematic coupling. Crosses mark the Doppler beyond which the correlation peak jumps to a sidelobe.

6.6. Sensitivity to chaotic parameters

Table 8 reports a 13-point sweep of the Lorenz parameters around the canonical operating point, varying one parameter at a time. Two regimes are evident. Below the chaos-onset boundary ($\rho \approx 24.74$ for $\sigma = 10$, $\beta = 8/3$) the system collapses toward a fixed point or limit cycle: at $\rho = 20$ the waveform degenerates (SFM = 0.099, PSR = −1.1 dB), and a similar collapse occurs at $\sigma = 5$. Inside the chaotic regime, spectral

flatness and sidelobes are robust, peaking near $\rho \in [28, 32]$, $\sigma \in [7.5, 10]$, and $\beta \in [2.0, 8/3]$. These ranges constitute the recommended design region.

Table 8. Lorenz parameter sensitivity (one parameter varied at a time).

Varied	σ	ρ	β	SFM	PSR (dB)	Regime
ρ	10	20	2.67	0.099	−1.1	Stable fixed point
ρ	10	24	2.67	0.734	−16.2	pre-chaos
ρ	10	28*	2.67	0.821	−17.8	chaotic
ρ	10	32	2.67	0.828	−16.9	chaotic
ρ	10	36	2.67	0.804	−15.9	chaotic
ρ	10	40	2.67	0.804	−17.3	chaotic
σ	5	28	2.67	0.215	−0.8	degraded
σ	7.5	28	2.67	0.823	−18.0	chaotic
σ	12.5	28	2.67	0.751	−15.7	chaotic
σ	15	28	2.67	0.743	−13.1	chaotic
β	10	28	1.0	0.703	−12.8	chaotic
β	10	28	2.0	0.828	−17.5	chaotic
β	10	28	3.5	0.728	−17.0	chaotic

*Canonical operating point used in Table 4 and all figures.

7. Deep-learning ESM: experiment and discussion

7.1. CNN recognition and learned-interceptor experiment

The deep-learning ESM threat set out in Section 1 is evaluated here experimentally. All four calibrated waveforms are transformed into 64 × 64 log-magnitude spectrograms by short-time Fourier transform at intercept-receiver SNRs from −14 to +6 dB, and a compact three-layer convolutional neural network (CNN) is trained to classify the waveform family. Recognition accuracy (Figure 11a) rises from 55.6% at −14 dB to 100% above −2 dB and is high for every family, the chaotic waveforms included. This outcome follows from what a spectrogram classifier actually measures: type recognition depends on gross time–frequency structure, not on whether successive pulses repeat. It is therefore consistent with the modulation-independence of the interceptability index *II* established in Section 4.1—chaotic frequency modulation conceals neither the presence of a wideband emission nor its modulation family.

The LPI benefit therefore cannot be, and is not claimed to be, concealment of type. To test what actually matters—whether a learning interceptor can build a reusable replica—we train a second network, a learned next-pulse predictor that maps the current pulse to the next and is scored by the coherence between its prediction and the true next pulse (a learned analogue of IPCG). Figure 11b contrasts this learned IPCG with the replica-reuse IPCG of Table 4. For FMCW the predictor is essentially perfect (learned IPCG = 0.999): a learning interceptor reconstructs the next sweep and recovers full coherent gain. For the chaotic and noise waveforms the learned predictor collapses to 0.14–0.20, close to the finite-length correlation floor, despite the classifier's ability to label them. In other words, recognizing that a pulse is Lorenz-CBFM buys the interceptor nothing it can coherently reuse against the next pulse. The learned floor (≈ 0.2) is set by the reduced 512-sample pulse length used to keep training tractable; the full-length IPCG of Table 4

resolves the ordering among the chaotic waveforms more finely. The qualitative separation—FMCW fully predictable, chaotic essentially unpredictable at the single-pulse horizon—is the experimental counterpart of the interception-range model of Section 4.2.2.

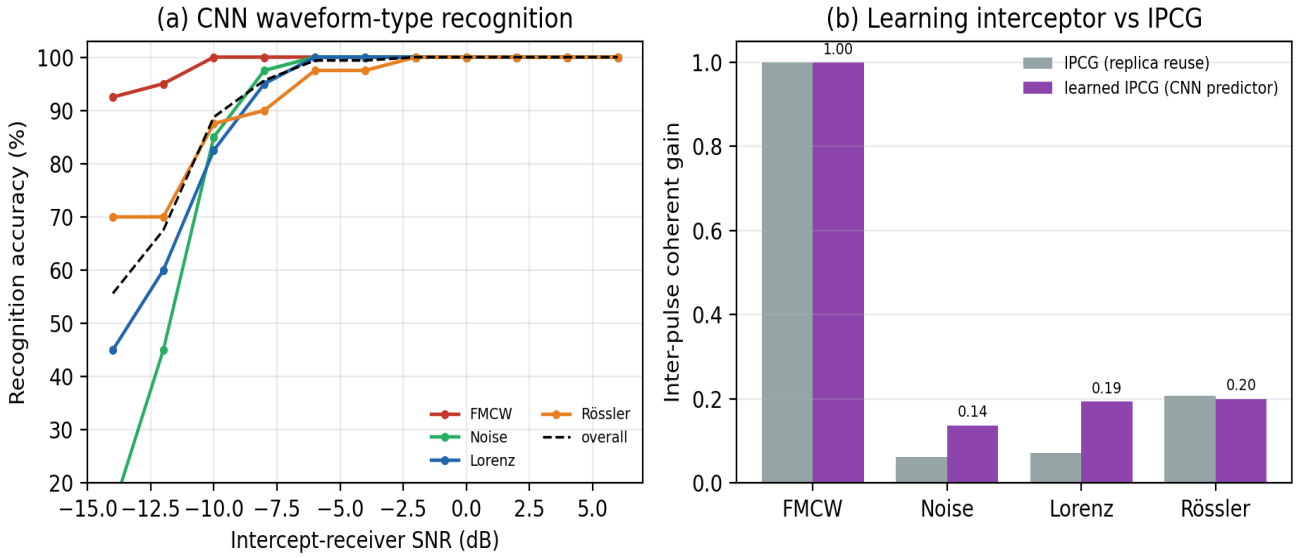


Figure 11. Deep-learning ESM experiment. **(a)** CNN waveform-type recognition accuracy versus intercept-receiver SNR: all four families are fully recognizable above −2 dB, so chaotic modulation does not hide the waveform type (consistent with equal *II*); **(b)** Learned IPCG (a CNN next-pulse predictor) *versus* the replica-reuse IPCG of Table 4: a learning interceptor recovers near-unity coherent gain for FMCW but only the noise-floor for the chaotic waveforms.

7.2. Discussion

The experiment clarifies why chaotic FM is relevant to the deep-learning ESM threat of Section 1. A learned classifier extracts repeatable features—chirp rate, cyclic spectral lines, or recurring time–frequency templates—from training examples [5,6,32,33]. FMCW supplies exactly such features: every sweep is identical, which is why its IPCG is unity and why a chirp-parameter estimator (human-designed or learned) succeeds. A chaotic waveform presents a different realization every pulse; its near-zero IPCG indicates the absence of the inter-pulse repeatability that feature-based and generative classifiers exploit [7]. Chaotic FM does not lower the radiometric interceptability of Equations (9a)–(9c)—no wideband waveform can, at fixed bandwidth, and our recognition experiment confirms the type is not hidden—but it removes the structural shortcut that lets a sophisticated interceptor outperform a radiometer, which Section 4.2.2 quantifies as a 1/IPCG reduction in coherent interception range. A caveat tempers this conclusion: chaotic trajectories are deterministic, and machine-learning predictors can forecast a chaotic state a limited number of steps ahead [34]. An interceptor that captured a long, clean segment might attempt short-horizon prediction rather than replica reuse; sensitive dependence on initial conditions bounds how far such prediction remains accurate, and the small margin by which the learned IPCG exceeds the replica IPCG in Figure 11b is a first measurement of that bound. Quantifying it against a stronger, sequence-modeling attacker remains important future work.

8. Design guidelines

Three practical recommendations follow from the study. (i) Operate well inside the chaotic regime: keep $\rho \gtrsim 28$ with $\sigma \approx 7.5$–10 and $\beta \approx 2.0$–$8/3$ for the Lorenz system; parameters near the chaos boundary degrade both flatness and sidelobes. (ii) Prefer the Lorenz (double-scroll) attractor over Rössler when sidelobe level and interceptability matter, as its richer trajectory yields a lower IPCG (0.072 *vs.* 0.207) and lower PSR. (iii) Accept the sidelobe trade-off knowingly: if minimum range sidelobes are the priority, FMCW remains superior, but if interception resistance and a thumbtack ambiguity surface are required, chaotic FM is the better choice. Mismatched or windowed matched filtering can recover part of the sidelobe gap if needed.

9. Limitations

This work is entirely simulation-based; hardware effects—power-amplifier nonlinearity, phase noise, and quantization—are not modeled and will raise the achievable sidelobe and IPCG floors. The interceptability analysis now considers three interceptor models—a radiometer (Section 4.1), a replica-reuse correlator (Sections 4.2.1–4.2.2), and a learned CNN interceptor (Section 7.1)—but a stronger sequence-modeling attacker that exploits short-horizon chaotic predictability and receiver synchronization sensitivity to initial-condition mismatch remains to be evaluated; the latter is critical for the deterministic-replica argument.

10. Conclusion

We compared Lorenz- and Rössler-based chaotic FM waveforms against calibrated FMCW and noise-radar references for LPI radar, with every waveform equalized to a 500 MHz occupied bandwidth and BT = 5000. The central finding is nuanced: FMCW retains the flattest spectrum (SFM = 0.971) and the lowest range sidelobes (PSR = −21.4 dB), and the radiometric interceptability index is identical (0.31) across all four waveforms because it is set by the system rather than the modulation. What distinguishes the waveforms is IPCG. FMCW's perfect periodicity (IPCG = 1.0) makes it fully exploitable by replica-based and learning-based interceptors, whereas Lorenz-CBFM reduces inter-pulse correlation to 0.072—essentially the noise-radar floor (0.063)—while providing a thumbtack ambiguity function free of range–Doppler coupling, at the price of about 3.5 dB higher sidelobes.

Three broader points emerge. First, fairly comparing LPI waveforms requires equalized bandwidth and a metric that captures how an interceptor actually exploits structure; the IPCG fills that role and, through the interception-range model of Section 4.2.2, connects a waveform property to a 1/IPCG reduction in the interceptor’s coherent detection range. Second, a deep-learning ESM experiment confirms that this protection is real against a learning interceptor: all four waveforms are recognizable by type, yet only FMCW yields a reusable learned replica. Third, the chaotic advantage is concentrated where it matters against modern electronic support: aperiodicity simultaneously delivers the low inter-pulse correlation that defeats replica reuse and the thumbtack ambiguity surface that decouples range from Doppler, so a single design choice improves both LPI protection and Doppler-tolerant estimation. For radars whose principal concern is resisting feature-hungry, learning-based ESM rather than minimizing sidelobes, chaotic FM with carefully chosen Lorenz parameters ($\rho \gtrsim 28$, $\sigma \approx 7.5$–10, $\beta \approx 2.0$–$8/3$) is a compelling choice. Future work will extend the comparison to hardware-in-the-loop measurements, to

stronger adversarial machine-learning interceptors and the predictability margin of chaotic sequences, and to a wider family of attractors.

Data availability statement

No supplementary or additional data were generated in this study.

Declaration of generative AI and AI-assisted technologies

The author did not use generative AI or AI-assisted technologies in the writing of this manuscript.

Acknowledgments

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflicts of interest

The author declares no conflicts of interest.

References

[1] Pace PE. D*etecting and Classifying Low Probability of Intercept Radar*, 2nd ed. Norwood: Artech House, 2009.
[2] Levanon N, Mozeson E. *Radar Signals*. Hoboken: Wiley–IEEE Press, 2004.
[3] Skolnik MI. *Introduction to Radar Systems*, 3rd ed. New York: McGraw-Hill, 2001.
[4] Galati G, Pavan G, Savci K, Wasserzier C. Counter-interception and counter-exploitation features of noise radar technology. *Remote Sens.* 2021, 13(22):4509.
[5] Wang C, Wang J, Zhang X. Automatic radar waveform recognition based on time-frequency analysis and convolutional neural networks. In *Proceedings of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, USA, March 5–9, 2017, pp. 2437–2441.
[6] Park DH, Jeon MW, Shin DM, Kim HN. LPI radar detection based on deep learning approach with periodic autocorrelation function. *Sensors* 2023, 23(20):8564.
[7] Ziemann MR, Metzler CA. Adaptive LPI radar waveform design with generative deep learning against learned electronic-support classifiers. *IEEE Trans. Radar Syst.* 2025, 3:417–429.
[8] Blunt SD, Mokole EL. Overview of radar waveform diversity. *IEEE Aerosp. Electron. Syst. Mag.* 2016, 31(11):2–42.
[9] Lorenz EN. Deterministic nonperiodic flow. *J. Atmos. Sci.* 1963, 20(2):130–141.
[10] Rössler OE. An equation for continuous chaos. *Phys. Lett. A* 1976, 57(5):397–398.
[11] Pecora LM, Carroll TL. Synchronization in chaotic systems. *Phys. Rev. Lett.* 1990, 64(8):821–824.
[12] Youssef AF, Sayed A. NIST test validation for Lorenz chaotic random generator CRN. *Eng. Lett.* 2023, 31(4):1526.
[13] Liu X, Zhang T, Yu X, Shi Q, Cui G, *et al*. LPI waveform design for radar system against cyclostationary analysis intercept processing. *Signal Process.* 2022, 201:108681.

[14] Flores BC, Solis EA, Thomas G. Assessment of chaos-based FM signals for range-Doppler imaging. *IEE Proc. Radar Sonar Navig.* 2003, 150(4):313–322.

[15] Venkatasubramanian V, Leung H. A novel chaos-based high-resolution imaging technique and its application to radar. *IEEE Signal Process. Lett.* 2005, 12(7):528–531.

[16] Zeng T, Chang S, Fan H, Liu Q. Design and processing of a novel chaos-based stepped frequency synthesized wideband radar signal. *Sensors* 2018, 18(4):985.

[17] Pappu CS, Beal AN, Flores BC. Chaos-based frequency modulation for joint monostatic and bistatic radar-communication systems. *Remote Sens.* 2021, 13(20):4113.

[18] Pappu CS, Beal AN, Blakely JN, Corron NJ. Analytic range-Doppler ambiguities for nonautonomous solvable chaos. *Chaos, Solitons Fractals* 2025, 197:116434.

[19] Pappu CS, Carroll TL. Quasi-FM waveform using chaotic oscillator for joint radar and communication systems. *Chaos, Solitons Fractals* 2021, 152:111449.

[20] Pappu CS, Flores BC, Debroux PS, Verdin B, Boehm J. Synchronisation of bistatic radar using chaotic AM and chaos-based FM waveforms. *IET Radar Sonar Navig.* 2017, 11(1):90–97.

[21] Gürbüz SZ, Griffiths HD, Charlish A, Rangaswamy M, Greco MS, *et al*. An overview of cognitive radar: past, present, and future. *IEEE Aerosp. Electron. Syst. Mag.* 2019, 34(12):6–18.

[22] Wolf A, Swift JB, Swinney HL, Vastano JA. Determining Lyapunov exponents from a time series. *Physica D* 1985, 16(3):285–317.

[23] Axelsson SRJ. Noise radar using random phase and frequency modulation. *IEEE Trans. Geosci. Remote Sens.* 2004, 42(11):2370–2384.

[24] Kulpa K. *Signal Processing in Noise Waveform Radar*. Norwood: Artech House, 2013.

[25] Youssef AS. Metaheuristic optimization for initial condition estimation in chaotic systems: toward robust chaotic communication networks. *J. Eng. Res.* 2026, in press.

[26] Govoni MA, Li H, Kosinski JA. Low probability of interception of an advanced noise radar waveform with linear-FM. *IEEE Trans. Aerosp. Electron. Syst.* 2013, 49(3):1351–1356.

[27] Bell MR. Information theory and radar waveform design. *IEEE Trans. Inf. Theory* 1993, 39(5):1578–1597.

[28] Woodward PM. *Probability and Information Theory, with Applications to Radar: International Series of Monographs on Electronics and Instrumentation*. London: Pergamon Press, 1953.

[29] Rice SO. Mathematical analysis of random noise. *Bell Syst. Tech. J.* 1944, 23(3):282–332.

[30] Richards MA. *Fundamentals of Radar Signal Processing*, 2nd ed. New York: McGraw-Hill Education, 2014.

[31] Marcum JI. A statistical theory of target detection by pulsed radar. *IRE Trans. Inf. Theory* 1960, 6(2):59–267.

[32] Huynh-The T, Hua CH, Doan VS, Pham QV, Kim DS. Accurate deep CNN-based waveform recognition for intelligent radar systems. *IEEE Commun. Lett.* 2021, 25(9):2938–2942.

[33] Hoang LM, Kim M, Kong SH. Automatic recognition of general LPI radar waveform using SSD and supplementary classifier. *IEEE Trans. Signal Process.* 2019, 67(13):3516–3530.

[34] Youssef AF, Sayed A. Predicting chaotic system for next-gen secure communication with machine learning. *IAENG Int. J. Comput. Sci.* 2025, 52(4):872.