A GNU Radio SDR Framework for Bistatic Sensing Using 5G NR Zadoff-Chu and SSB Waveforms

R Krishna Kanth KROKKAM@NCSU.EDU
Khawaja Wahab WKHAWAJ@NCSU.EDU
Ali C. Gurbuz ALIGURBUZ@NCSU.EDU

Department of Electrical and Computer Engineering,
North Carolina State University, Raleigh, NC 27606 USA

Abstract

We present a real-time bistatic sensing chain built entirely in GNU Radio that uses 5G New Radio (NR) synchronization waveforms as the sensing signal. The approach exploits two properties of the 5G synchronization signal block (SSB): (i) its frame structure is fully specified, so a receiver can reconstruct the transmitted waveform locally as a matched-filter reference, and (ii) its synchronization sequences are short and fully deterministic. In 5G NR the PSS is a length-127 binary m-sequence rather than a Zadoff-Chu (ZC) sequence, so for the sensing reference itself we use the ZC family, which 5G NR employs for PRACH preambles and whose constant-amplitude zero-autocorrelation (CAZAC) property yields a range response free of correlation sidelobes. Because the standard SSB is transmitted at a low duty cycle, a few percent at most in the configuration considered here, we transmit the synchronization pilot continuously and use an extended ZC reference, trading standards compliance for sensing energy. The processing chain comprises adaptive least-mean-squares (LMS) direct-path and near-range clutter cancellation, FFT-based matched filtering, and an exponential-average moving-target indicator (MTI). It is implemented as streaming GNU Radio out-of-tree blocks and runs in real time at 30.72 MS/s on NI USRP-2901 hardware. MATLAB simulations of a drone target at four bistatic excess ranges (19.5–195.3 m) give target-to-clutter ratios above 35 dB with correct range and Doppler recovery. Over-the-air experiments in a 38 m indoor corridor track a walking human continuously over 4.5 - 40.4 m with a range RMSE of 3.53 m and a radial-velocity RMSE of 0.26 m/s using the extended ZC waveform, and we compare this directly against the 5G NR PSS radiated as a continuous pilot.

Draft manuscript of the 16^{th} GNU Radio Conference

1. Introduction

Integrated sensing and communications (ISAC) has become a central theme for 5G-Advanced and 6G, with the goal of enabling communication infrastructure to provide sensing capabilities alongside data transmission (Liu et al., 2022). A closely related concept is *passive coherent location* (PCL) (Howland et al., 2005), in which a receiver detects targets from reflections of a signal transmitted for another purpose by using a reference copy of the transmitted signal for correlation. Both approaches rely on a similar receiver architecture that includes a reference channel, a surveillance channel, suppression of the direct path, and coherent range-Doppler processing.

5G NR signals are particularly attractive for this class of sensing systems because their physical layer is rigorously specified by 3GPP. The Synchronization Signal Block (SSB) is transmitted periodically and contains the Primary Synchronization Signal (PSS), Secondary Synchronization Signal (SSS), and Physical Broadcast Channel (PBCH). In 5G NR the PSS is a length-127 binary m-sequence and the SSS a Gold sequence formed from two m-sequences. Both are fully specified, so a receiver can regenerate them locally without decoding user data and use them as a correlation reference (Abratkiewicz et al., 2023). Zadoff-Chu (ZC) sequences, which 5G NR employs for PRACH preambles and which formed the PSS in LTE, additionally possess the constant-amplitude zero-autocorrelation (CAZAC) property i.e. for a prime length the cyclic autocorrelation is impulse-like with zero sidelobes. That property is what a matched-filter sensing reference most wants, and it is why the sensing waveform in this work is drawn from the ZC family rather than from the NR PSS itself.

The main challenge is real-time processing. At 5G NR sample rates, a sensing receiver must continuously suppress a direct-path signal that can be orders of magnitude stronger than the target echo, perform matched filtering on the residual signal, and reject static clutter while keeping pace with the incoming sample stream. GNU Radio (GNU Radio Project) provides a suitable framework for this task because its streaming scheduler and modular block architecture allow each processing stage to operate continously and to be reused across different system configurations.

In this work, we deliberately separate the *processing chain* from the *signal source*. The processing chain follows a standard PCL receiver architecture and is independent of the particular transmitter. The experiments in Section 6, however, use a transmitter under our control together with a wired reference connection, and continuously transmit the synchronization pilot rather than embedding it within a standard-compliant 5G frame. The resulting system therefore provides a *cooperative bistatic* validation of a PCL-style sensing chain rather than opportunistic sensing of an operational commercial base station. This distinction is maintained throughout the paper, while sensing with ambient gNB transmissions is left as a natural direction for future work. This experimental configuration also allows us to investigate a fundamental design question concerning 5G-derived sensing waveforms, namely how the waveform should be structured to provide sufficient sensing energy while retaining the desirable properties of 5G synchronization signals.

Contributions of this paper:

- We implement a complete streaming GNU Radio sensing chain for 5G NR synchronization waveforms. The chain is implemented as three reusable out-of-tree blocks: (i) an *LMS Canceller* for adaptive direct-path and near-range clutter suppression, (ii) a *Fast Cross Correlator* for FFT-domain matched filtering, and (iii) a *ZC Real-Time MTI* for exponential-average clutter rejection. The chain runs in real time at 30.72 MS/s.
- We quantify the duty-cycle limitation imposed by the standard SSB for sensing and adopt a continuously transmitted, extended ZC pilot, and we show in simulation why the ZC-only reference gives a cleaner range-Doppler map than the full SSB burst.
- We characterise the chain in MATLAB against a drone target at four bistatic excess ranges (19.5–195.3 m), obtaining target-to-clutter ratios above 35 dB with correct range and Doppler recovery.
- We validate the implementation over the air on an NI USRP-2901 bistatic testbed, tracking a walking human continuously across a 38 m corridor, and we compare the extended ZC pilot against the 5G NR PSS radiated as a continuous pilot.

The rest of the paper is organised as follows. Section 2 reviews the 5G NR SSB and ZC properties relevant to sensing and motivates the continuous-pilot waveform. Section 3 develops the bistatic signal model and presents each processing stage together with the GNU Radio block that implements it. Section 4 covers the hardware, flowgraph, and real-time budget. Section 5 reports MATLAB simulation results, and Section 6 the over-the-air experiments. Section 7 concludes the paper.

2. 5G NR Waveforms for Sensing

2.1. SSB Structure and the PSS

A 5G NR downlink frame lasts 10 ms and is divided into subframes, slots and OFDM symbols. The Synchronization Signal Block (SSB) occupies four consecutive OFDM symbols over 240 subcarriers and carries three physical-layer signals:

- Symbol 0: Primary Synchronization Signal (PSS)
- Symbol 1: Physical Broadcast Channel (PBCH)
- Symbol 2: Secondary Synchronization Signal (SSS) and PBCH
- Symbol 3: PBCH

The PSS and SSS are the components of interest for sensing: both are deterministic and can be reconstructed locally without decoding any user data, which is exactly what a matched filter requires. In NR the PSS is a length-127 binary m-sequence generated by the seven-stage recursion $x(i+7) = [x(i+4) + x(i)] \bmod 2$, with the three cell-identity values distinguished by cyclic shifts. The SSS is a Gold sequence built from two such m-sequences. Neither is a Zadoff-Chu sequence, a point that matters for the waveform choice made in Sec. 2.3.

2.2. Zadoff-Chu Sequences

Zadoff-Chu sequences are used in 5G NR for the PRACH preambles and as DMRS base sequences, and were the basis of the PSS in LTE (Lin et al., 2019). They are not the NR PSS. We adopt them here for the sensing reference purely on the strength of their correlation properties. A ZC sequence of length N_{ZC} with root index u is defined as

$$s[n] = \exp\left(-j\pi \frac{u\,n(n+1)}{N_{ZC}}\right), \quad n = 0, 1, \ldots, N_{ZC}-1. \tag{1}$$

ZC sequences exhibit CAZAC properties when the sequence length and root satisfy the required conditions. For a prime sequence length, the cyclic autocorrelation has an impulse-like response with zero sidelobes. This property is particularly useful for sensing because a weak target echo must be detected in the presence of a potentially much stronger direct-path component.

The 127-sample length of the NR synchronization sequences provides limited processing gain and a short correlation interval for sensing. We therefore take the ZC construction of Eq. (1) and extend it to $N_{ZC} = 16381$ (the largest prime below 2^{14}, chosen to preserve the zero-sidelobe guarantee, which holds for prime lengths) with root $u = 347$. This extended ZC sequence is the primary sensing waveform used throughout the paper.

2.3. Duty Cycle and the Continuous Pilot

An SSB burst repeats with a period T_b and only a small fraction of that period carries signal. NR allows $T_b \in \{5, 10, 20, 40, 80, 160\}$ ms, with 20 ms recommended for cell-defining SSBs, and up to L beam-swept SSBs per burst set ($L_{max} = 4$ or 8 in FR1, 64 in FR2). We take the single-SSB, $T_b = 5$ ms case as the reference configuration throughout: the PSS alone then occupies one 75 μs OFDM symbol, a duty cycle of 1.5%, and the full four-symbol SSB burst occupies 300 μs, or 6.0%. Fig. 1 shows this timing over five burst periods. Beam sweeping makes the picture less favourable rather than more: only one of the L swept beams illuminates a given target, so the duty cycle usable for sensing in a fixed direction is the single-SSB figure, below 2% at the recommended 20 ms periodicity.

The low duty cycle is an important practical limitation when standard 5G synchronization signals are considered for sensing. A sensing receiver can integrate information over multiple SSB transmissions, but the available sensing energy is still determined by the fraction of time during which the waveform is transmitted. Consequently, a standard SSB transmission provides only 6% of the energy that would be available from continuous transmission over the same time interval when the full four-symbol SSB is considered.

The objective of this work is to characterize the processing chain rather than to perform opportunistic sensing of a live commercial base station. We therefore transmit the synchronization pilot *continuously*, at 100% duty cycle, in both the simulations of Section 5 and the experiments of Section 6. Successive N_{ZC}-sample sequences are radiated back to back, so the pulse repetition interval is $T_{PRI} = N_{ZC}/f_s = 0.533$ ms. The resulting waveform is derived from the 5G NR signal set and retains the ZC construction that NR uses for PRACH, but it does not constitute a standards-compliant 5G NR frame. This configuration intentionally trades standard compliance for increased sensing energy. For comparison we also radiate the NR PSS itself as a continuous pilot, the length-127 m-sequence repeated ten times per transmitted block, which keeps the standard sequence but inherits its much shorter correlation length.

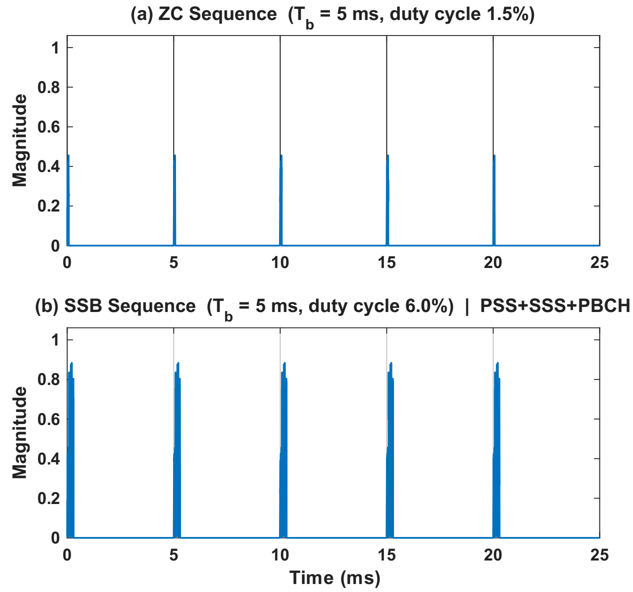

Figure 1. 5G NR PSS and SSB burst timing over five repetitions of the $T_b = 5$ ms period. (a) PSS-only: 75 μs burst, 1.5% duty cycle. (b) Full SSB burst (PSS+SSS+PBCH): 300 μs burst, 6.0% duty cycle.

3. Signal Model and Processing Chain

This section develops the bistatic model and then presents each processing stage, together with the GNU Radio block that implements it and the parameter values used.

3.1. Bistatic Geometry and Received Signal

Fig. 2 shows the geometry, with the transmitter at $\mathbf{p}_{TX}$ and receiver at $\mathbf{p}_{RX}$ separated by a baseline d, and the target at $\mathbf{p}_{tgt}$ with velocity $\mathbf{v}$. The bistatic excess range is

$$R_{bi} = |\mathbf{p}_{tgt} - \mathbf{p}_{TX}| + |\mathbf{p}_{tgt} - \mathbf{p}_{RX}| - d \tag{2}$$

and the bistatic Doppler shift is

$$f_d = \frac{\mathbf{v} \cdot (\hat{\mathbf{r}}_{TX} + \hat{\mathbf{r}}_{RX})}{\lambda}, \qquad \lambda = \frac{c}{f_c}, \tag{3}$$

where $\hat{\mathbf{r}}_{TX}$ and $\hat{\mathbf{r}}_{RX}$ are unit vectors from the target toward TX and RX, and their sum gives the rate of change of total path length (Willis, 2005). For a near-boresight quasi-monostatic geometry this approaches $2v/\lambda$: at $f_c = 2.4\,\text{GHz}$ ($\lambda = 12.5\,\text{cm}$), a 10 m/s drone gives $f_d \approx 160\,\text{Hz}$ and a 1 m/s walking human gives $f_d \approx 16\,\text{Hz}$. These two figures differ by an order of magnitude and, as shown in Sec. 3.4, they require different clutter filter settings.

Two range conventions are used in this paper and it is worth separating them explicitly, because they differ by a factor of two. A matched filter of bandwidth B resolves delay to $1/B$, which at $B = f_s = 30.72\,\text{MHz}$ corresponds to $c/B = 9.76\,\text{m}$ of *total path length*, i.e. of the bistatic sum $|\mathbf{p}_{tgt} - \mathbf{p}_{TX}| + |\mathbf{p}_{tgt} - \mathbf{p}_{RX}|$. The simulation of Sec. 5 reports bistatic excess range R_{bi} directly, so its bin spacing is $\Delta R_{bi} = c/B = 9.76\,\text{m}$.

The over-the-air geometry of Sec. 6 is quasi-monostatic: the TX and RX horns sit on a 45 cm baseline aimed along a 38 m corridor, so $|\mathbf{p}_{tgt} - \mathbf{p}_{TX}| \approx |\mathbf{p}_{tgt} - \mathbf{p}_{RX}| \approx R$ and the bistatic sum is $\approx 2R$. The path length therefore changes twice as fast as the target range, and the detector, which reports one-way target range, has a bin spacing of

$$\Delta R = \frac{c}{2B} = 4.88\ \text{m}. \tag{4}$$

The 38 m corridor thus spans approximately 7.8 range bins.

The surveillance antenna receives the direct path, static clutter, the target return and noise:

$$\begin{aligned} y_{surv}[n] = &\underbrace{A_{dp}\, s[n]}_{\text{direct path}} + \underbrace{\sum_k A_{cl,k}\, s[n - \tau_{cl,k}]}_{\text{clutter}} \\ &+ \underbrace{A_{tgt}\, e^{j2\pi f_d (n/f_s)}\, s[n - \tau_s]}_{\text{target}} + \underbrace{w[n]}_{\text{AWGN}} \end{aligned} \tag{5}$$

where R_{bi} follows Eq. (2), f_d Eq. (3), $\tau_s = \lceil R_{bi} f_s / c \rceil$ is the target delay in samples, $A_{dp} \gg A_{tgt}$, and $w[n]$ is complex AWGN. The reference channel carries essentially the direct path alone,

$$y_{ref}[n] = A_{dp} s[n] + w_{ref}[n]. \tag{6}$$

3.2. Stage 1: Adaptive Cancellation (*LMS Canceller*)

The direct path is suppressed with a complex LMS adaptive filter (Widrow & Stearns, 1985) of K taps, driven by the reference channel:

$$\hat{y}[n] = \textstyle\sum_{k=0}^{K-1} w_k[n]\, y_{ref}[n-k], \tag{7}$$

$$e[n] = y_{surv}[n] - \hat{y}[n], \tag{8}$$

$$w_k[n+1] = w_k[n] + \mu\, y_{ref}[n-k]\, e^*[n]. \tag{9}$$

The residual $e[n]$ of Eq. (8) is the canceller output. The `LMS Canceller` block implements Eqs. (7)–(9) sample-by-sample, with K and μ exposed as GRC parameters.

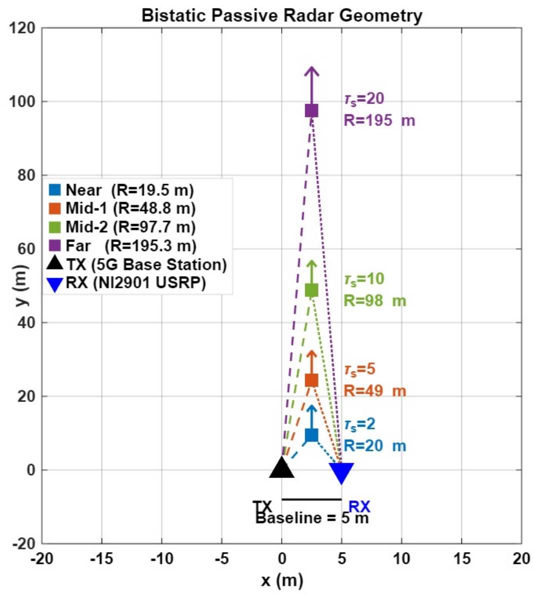

Figure 2. Bistatic geometry used in simulation. Squares mark the simulated drone positions and arrows the $+y$ velocity direction.

The tap count K has a direct physical meaning that governs how the block is configured. With a small step size the filter converges toward the *stationary* component of the surveillance channel that is linearly predictable from the reference, over a span of K delay bins, i.e. over $K \cdot \Delta R_{bi}$ of bistatic excess range. Increasing K therefore suppresses not only the zero-lag direct path but also static clutter at the first K range bins. A moving target within that span carries a time-varying Doppler phase, is not well modelled by a fixed weight vector, and survives with limited attenuation.

We use this property differently in the two settings reported here. The MATLAB simulation study of Section 5 uses $K = 1$, restricting the filter to the zero-lag direct path so that the simulated targets at $\tau_s \geq 2$ are provably untouched and the cancellation behaviour can be assessed in isolation. The over-the-air system of Section 6 uses $K = 5$, spanning delays of 0 to 4 samples. In the quasi-monostatic corridor geometry that is roughly the first 20 m of target range (Eq. 4), or about half the corridor, and it is chosen to reject the strong static returns near the antennas. Targets beyond that span lie outside the filter entirely; within it, detection relies on the stationary-versus-moving argument above. The residual attenuation of a moving target inside the cancellation span has not been separately quantified and is a useful subject for a controlled K-sweep.

3.3. Stage 2: Matched Filtering (*Fast Cross Correlator*)

The residual is correlated against the locally generated ZC reference by FFT-domain multiplication:

$$R[\tau] = \mathcal{F}^{-1}\{\mathcal{F}\{e[n]\} \cdot \mathcal{F}^{*}\{s[n]\}\}\,. \tag{10}$$

With $N_{ZC} = 16381$, an FFT size $N_{FFT} = 32768$ satisfies $N_{FFT} \geq 2N_{ZC} - 1$ and so yields linear rather than circular correlation. The `Fast Cross Correlator` block precomputes $\mathcal{F}^{*}\{s[n]\}$ at initialisation and, for each arriving N_{ZC}-sample block, zero-pads, transforms, multiplies and inverse-transforms. The output is retained as a complex range profile, with no magnitude operation applied, so that the pulse-to-pulse phase needed for Doppler estimation is preserved.

3.4. Stage 3: Clutter Rejection (*ZC Real-Time MTI*)

Residual stationary clutter is removed by an exponential-average moving-target indicator operating independently in each range bin:

$$b_m[\tau] = (1-\alpha)\, b_{m-1}[\tau] + \alpha\, R_m[\tau], \tag{11}$$

$$\tilde{R}_m[\tau] = R_m[\tau] - b_m[\tau], \tag{12}$$

where m indexes pulses. The `ZC Real-Time MTI` block maintains the complex baseline $b[\tau]$ persistently across scheduler calls so that it is not reset at buffer boundaries, and exposes α as a GRC parameter. The output is kept complex so that the pulse-to-pulse phase is preserved for the downstream Doppler estimation of Sec. 3.5, on which both the simulation's Doppler FFT and the over-the-air pulse-pair estimator depend.

Eqs. (11) - (12) form a single-pole high-pass filter in slow time with 3 dB cut-off

$$f_{3dB} = \frac{\alpha}{2\pi\, T_{PRI}}, \tag{13}$$

so Eq. (13) shows that α must be chosen against the expected target Doppler, and this is why the two experiments in this paper use different values. For the drone simulation, $\alpha = 0.10$ gives $f_{3dB} = 29.9$ Hz, far below the $f_d \approx 160$ Hz drone return, and the baseline converges within the $M_{CPI} = 256$-pulse batch, since $(1-\alpha)^{M_{CPI}} = 1.9 \times 10^{-12}$. Reducing α to 0.01 in that batch setting would leave $(1-\alpha)^{M_{CPI}} = 7.6 \times 10^{-2}$, i.e. only about 22 dB of baseline convergence. For the walking human of Section 6 the situation reverses: at $f_d \approx 16$–24 Hz the target would fall *below* the 29.9 Hz cut-off and be attenuated, so the over-the-air system uses $\alpha = 0.01$ ($f_{3dB} = 2.98$ Hz). Convergence is not a constraint there because the streaming block adapts continuously rather than over a fixed batch.

3.5. Stage 4: Range-Doppler Estimation

After M_{CPI} coherent pulses the range-Doppler map is formed by a Doppler FFT across slow time:

$$\mathrm{RD}[\tau, f_d] = \sum_{m=0}^{M_{CPI}-1} \tilde{R}_m[\tau]\, w_{Ham}[m]\, e^{-j2\pi f_d m \cdot T_{PRI}}, \tag{14}$$

with $w_{Ham}[m]$ a Hamming window to suppress Doppler sidelobes. The Doppler resolution is $\Delta f_d = 1/(M_{CPI}T_{PRI})$, which for $M_{CPI} = 256$ gives 7.33 Hz, equivalently $\Delta v = \lambda \Delta f_d/2 = 0.46$ m/s.

Target range and velocity are estimated from the location of the correlation peak. It is important to distinguish *resolution* from *accuracy* here. Range resolution is fixed by bandwidth at $c/2B = 4.88$ m (Eq. 4) and no amount of post-processing improves it. Two targets closer than one bin cannot be separated. Range *accuracy*, the precision with which a single dominant peak is located, is limited instead by signal-to-noise ratio and can be made far finer than one bin.

The true peak generally falls between samples, and the three magnitude samples around the peak bin k carry the fractional offset in their asymmetry. Fitting a parabola through $|R[k-1]|$, $|R[k]|$, $|R[k+1]|$ gives

$$\delta = \frac{1}{2}\,\frac{|R[k-1]| - |R[k+1]|}{|R[k-1]| - 2|R[k]| + |R[k+1]|}, \qquad |\delta| \leq \tfrac{1}{2}, \tag{15}$$

and the refined range is $(k+\delta)\,\Delta R$. If the target lies exactly on a bin the neighbours are equal and $\delta = 0$. If it lies midway between bins they are maximally asymmetric and $\delta \to \pm\frac{1}{2}$. The estimator needs only a single snapshot, no target motion is required.

For a 4.88 m bin the theoretical accuracy of Eq. (15) is of order 0.35 m at 20 dB post-correlation SNR and 0.11 m at 30 dB. The metre-level errors reported in Sec. 6 are well above those figures, which indicates that the interpolator is not the limiting factor. A walking human is not a point scatterer, its reflection centroid moves as the limbs swing, and the corridor contributes multipath. This is the expected regime for an extended, moving target indoors.

The over-the-air system of Sec. 6 does *not* perform this coherent integration. There the stream is decimated ahead of matched filtering and MTI, and no range-Doppler map is formed. Radial velocity is instead obtained from a *pulse-pair* estimator: the phase advance of the complex correlation peak between consecutive pulses, averaged over three successive PRIs,

$$\hat{f}_d = \frac{1}{2\pi T_{PRI}} \arg\left\{ \sum_m \tilde{R}_{m+1}[\tau^\star]\, \tilde{R}_m^{*}[\tau^\star] \right\}, \qquad \hat{v} = \frac{\lambda \hat{f}_d}{2}. \tag{16}$$

Equation (16) is unambiguous over $|f_d| < 1/(2T_{PRI}) = 938\,\text{Hz}$, i.e. $|v| < 58.6\,\text{m/s}$, so a walking human is never aliased. As with range, resolution and accuracy differ sharply. Three pulses give a Doppler *resolution* of only $1/(3T_{PRI}) = 625\,\text{Hz}$ (39 m/s), so two targets cannot be separated in Doppler at all. But the *accuracy* for a single dominant target scales as $1/(2\pi T_{PRI}\sqrt{N \cdot \text{SNR}})$ and reaches a few Hz at realistic post-correlation SNR. The 0.26 m/s velocity RMSE measured in Sec. 6 corresponds to $\sigma_{f_d} \approx 4\,\text{Hz}$, consistent with roughly 32 dB SNR over three pulses.

Equation (14) and the M_{CPI}-dependent resolutions above therefore describe the simulation path of Sec. 5 only.

We quantify detection quality by the target-to-clutter ratio

$$\text{TCR} = 20\log_{10}\frac{|\text{RD}[\tau^\star, f_d^\star]|}{\text{mean}_{(\tau,f_d)\in\mathcal{N}}|\text{RD}[\tau, f_d]|}, \qquad (17)$$

where $(\tau^\star, f_d^\star)$ is the detected peak and $\mathcal{N}$ is a noise-reference region excluding a guard window around it.

4. Real-Time GNU Radio Implementation

4.1. Hardware

Fig. 3 shows the testbed. Two NI USRP-2901 devices (National Instruments) are used: one acts as the transmitter, radiating the continuous ZC or PSS pilot of Sec. 2.3, and the second acts as the two-channel sensing receiver. Both are locked to a common external rubidium reference, which keeps the transmit and receive local oscillators coherent. Without a shared reference, drift between the two devices appears as a spurious Doppler component and corrupts the pulse-to-pulse phase on which both the MTI of Sec. 3.4 and any velocity estimate depend. Directional horn antennas, separated by a 45 cm baseline, are aimed along the monitored corridor. Both channels are sampled at $f_s = 30.72\,\text{MS/s}$ and tuned to $f_c = 2.4\,\text{GHz}$.

4.2. Flowgraph and Real-Time Budget

Fig. 4 shows the GNU Radio Companion flowgraph. A `UHD USRP Source` acquires the reference and surveillance channels. A separate `UHD USRP Sink` drives the transmitter from a locally generated vector source. The `LMS Canceller` runs first, at the full sample rate. The stream is then decimated ahead of the `Fast Cross Correlator` and `ZC Real-Time MTI`, and the MTI output is streamed over UDP for peak detection.

No coherent integration is performed in the over-the-air chain. There is no Doppler FFT and no range-Doppler map; range comes from the interpolated correlation peak (Eq. 15) and radial velocity from the three-pulse pulse-pair estimator of Eq. (16), both reported per processed frame.

Figure 3. Bistatic testbed: NI USRP-2901 transmitter and two-channel sensing receiver sharing a rubidium reference, with TX and RX horn antennas on a 45 cm baseline aimed along the corridor.

The range-Doppler formulation of Eq. (14) applies to the MATLAB study of Sec. 5. This is a deliberate simplification: decimating before the correlator keeps the real-time chain well inside budget. It does mean the over-the-air system forgoes the coherent processing gain and the Doppler resolution that a full CPI provides in simulation, which is the most likely explanation for the weaker far-range performance discussed in Sec. 6.

The real-time requirement follows from the continuous pilot. At $T_{PRI} = 0.533\,\text{ms}$ an undecimated chain would have to sustain $1/T_{PRI} \approx 1876$ range profiles per second. The `Fast Cross Correlator`, which dominates the per-pulse cost, sustains more than 2000 pulses/s on the test laptop and so meets that bound on its own. Decimating ahead of the correlator lowers the requirement further, and the measured over-the-air detection rates of Sec. 6 correspond to a few frames per second.

All three blocks are Python out-of-tree blocks, which keeps them readable and easy to modify at the cost of some headroom. The measured margin above is modest, and a C++ implementation would be the natural step for higher sample rates or longer references.

5. Simulation Results

The chain is first characterised in MATLAB against a simulated UAV of RCS $\sigma = 0.05\,\text{m}^2$ moving at $v = 10\,\text{m/s}$

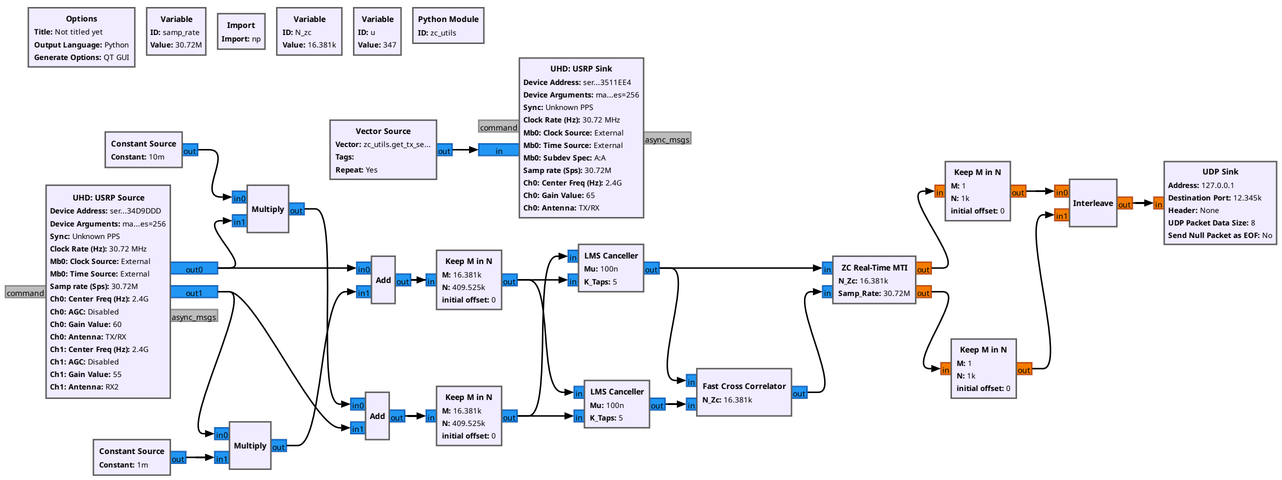

Figure 4. GNU Radio Companion flowgraph. The `LMS Canceller`, `Fast Cross Correlator` and `ZC Real-Time MTI` run as streaming Python out-of-tree blocks inside the GNU Radio scheduler at 30.72 MS/s. The over-the-air runs use $K = 5$ taps and $\mu = 10^{-7}$ in the canceller (Sec. 3.2).

in the $+y$ direction, at four bistatic excess ranges spanning the near and far extremes of interest. A drone target is used here because its $\approx 160\,\text{Hz}$ Doppler is well separated from zero and therefore isolates range-Doppler behaviour from MTI tuning. The slower human target is treated in Section 6.

Table 1 lists the parameters. Table 2 gives the four scenarios and their expected range, Doppler and Doppler bin.

5.1. Range Profiles

Fig. 5 follows the Mid-1 scenario ($\tau_s = 5$, $R_{bi} = 48.8\,\text{m}$) through the chain. The raw correlation profile is dominated by the direct path at zero delay, with the three modelled static clutter returns also visible and the target buried among them. Single-tap LMS cancellation removes the zero-delay component, leaving the clutter returns and the target. The MTI stage then removes the static returns, isolating the target at the expected range bin with a TCR of 52 dB by Eq. (17). Across all four scenarios the TCR exceeds 35 dB and the detected range and Doppler bins match Table 2.

5.2. Range-Doppler Maps

Fig. 6 shows the range-Doppler maps for the nearest and farthest scenarios. In both cases the target appears at the expected range and Doppler cell, and the residual clutter ridges sit at zero Doppler where the MTI has pushed them, confirming that the chain separates the moving target from static returns over an order of magnitude of range.

5.3. ZC Reference versus Full SSB

Fig. 7 compares range-Doppler maps formed with the extended ZC reference against the full SSB burst for the farthest scenario, with both waveforms normalised to equal mean transmit power so the comparison reflects waveform structure rather than energy. The ZC map isolates the target in a single cell against a flat floor. The SSB map recovers the same target but also shows a pronounced ridge extending across the entire range axis at the target Doppler, together with a broad low-Doppler pedestal spanning the first 150 m. Both artifacts follow from the SSB's OFDM structure: the SSS and PBCH symbols are not CAZAC, and the cyclic prefix introduces periodic self-similarity, so the autocorrelation is no longer an impulse. The simulated SSB places a ZC sequence in the PSS symbol, following the LTE construction rather than the NR m-sequence. We verified that this does not affect the comparison: rebuilding the SSB with the standards-correct NR PSS (length-127 m-sequence) and SSS (Gold sequence) leaves the worst autocorrelation sidelobe of the tiled SSB essentially unchanged at $-0.2\,\text{dB}$, against $-207\,\text{dB}$ for the extended ZC reference. The artifacts are set by the OFDM framing and by the unknown PBCH payload, not by which synchronization sequence occupies symbol 0.

In a multi-target scene such a ridge would mask weaker returns at other ranges and could be mistaken for detections in its own right. This is the basis for using the ZC-only reference in the processing chain.

Table 1. System Parameters

Parameter	Symbol	Value
Carrier frequency	f_c	2.4 GHz
Sample rate	f_s	30.72 MS/s
ZC length	N_{ZC}	16 381
ZC root	u	347
PRI	T_{PRI}	0.533 ms
Delay bin (path length)	c/B	9.76 m/bin
Range bin (target range)	$c/2B$	4.88 m/bin
Direct-path ratio	DPR	50 dB
Drone RCS	σ	0.05 m^2
Drone speed	v	10 m/s
TX-RX baseline	d	5 m
Simulation (drone, $f_d \approx 160$ Hz)		
Pulses per CPI	M_{CPI}	256
CPI	T_{CPI}	136.5 ms
Doppler resolution	Δf_d	7.33 Hz/bin
LMS taps	K	1
LMS step size	μ	10^{-3}
MTI forgetting factor	α	0.10
Over-the-air (human, $f_d \approx$ 16–24 Hz)		
Coherent integration	—	none
Doppler estimator	—	3-pulse pulse-pair
Unambig. velocity	—	±58.6 m/s
TX–RX baseline (OTA)	d	0.45 m
LMS taps	K	5
LMS step size	μ	10^{-7}
MTI forgetting factor	α	0.01

Table 2. Simulated Drone Scenarios

Scenario	τ_s	R_{bi}	f_d	f_d bin
Near	2	19.5 m	154.7 Hz	21.1
Mid-1	5	48.8 m	159.1 Hz	21.7
Mid-2	10	97.7 m	159.8 Hz	21.8
Far	20	195.3 m	159.9 Hz	21.9

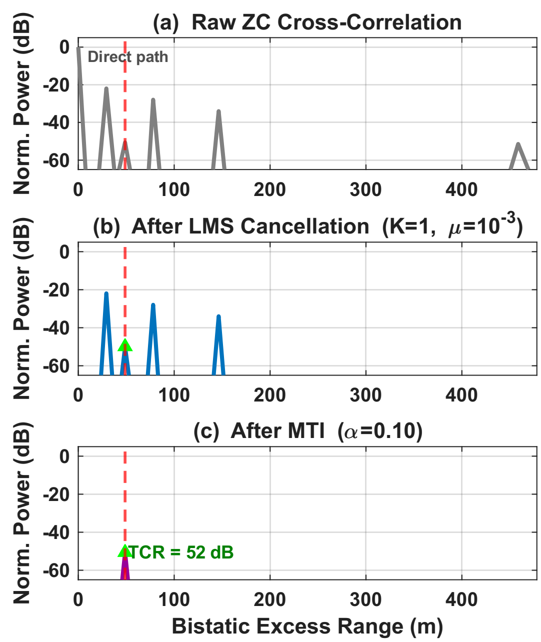

Figure 5. Simulated range profiles for the Mid-1 drone position ($\tau_s = 5$, $R_{bi} = 48.8$ m). (a) raw cross-correlation, (b) after single-tap LMS cancellation ($K = 1$, $\mu = 10^{-3}$), (c) after MTI ($\alpha = 0.10$). Green triangles mark the expected target location.

6. Over-the-Air Experimental Results

6.1. Experimental Setup

The over-the-air experiments use the testbed of Fig. 3 in a 38 m indoor corridor. One USRP radiates the continuous pilot at $f_c = 2.4$ GHz through a horn antenna. The second USRP provides the two sensing channels: the surveillance channel is fed by a horn antenna aimed along the corridor, capturing both target reflections and direct-path leakage, while the reference channel receives a wired copy of the transmitted waveform taken directly from the transmitting USRP. That copy is not passed through untouched: the multiply-and-add stage visible in Fig. 4 deliberately injects a scaled, delayed component so that the reference carries a controlled level of emulated impairment rather than being an ideal noiseless replica. This keeps the reference realistic enough to exercise the canceller while holding the impairment at a known, repeatable level, which an over-the-air reference antenna would not. A deployed passive receiver would need a second antenna steered at the illuminator, with uncontrolled multipath and noise on that path; that is a meaningfully harder problem and is left to future work.

A single human subject walked from the antennas to the far end of the corridor, waited there, and walked back. The run was repeated for two transmit waveforms: the extended ZC pilot of Sec. 2 ($N_{ZC} = 16381$, $u = 347$), and the 5G NR PSS itself, the length-127 m-sequence, repeated ten times per transmitted block, both radiated continuously. Note that this is a different pairing from the simulation of Sec. 5.3, which contrasts the extended ZC reference against the *full four-symbol SSB burst*. Over the air the comparison isolates sequence length and family at equal duty cycle, whereas in simulation it isolates the effect of the SSB's OFDM structure. The two are complementary rather than the same experiment. The receiver ran the flowgraph of Fig. 4 with $K = 5$, $\mu = 10^{-7}$ and $\alpha = 0.01$, reporting an estimated range and radial velocity for every processed frame.

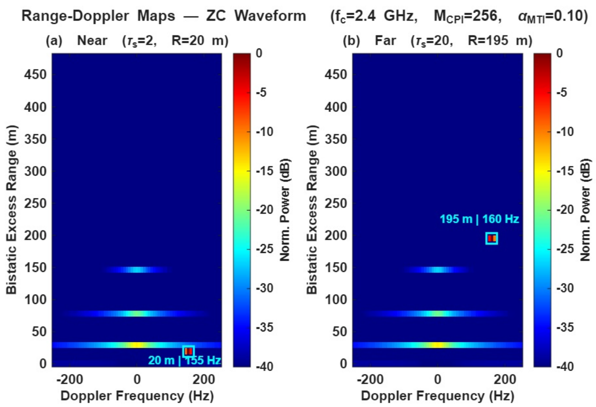

Figure 6. Range-Doppler maps for the near ($\tau_s = 2$) and far ($\tau_s = 20$) drone scenarios with the ZC waveform, $M_{CPI} = 256$, $\alpha = 0.10$. Residual static clutter remains confined to zero Doppler.

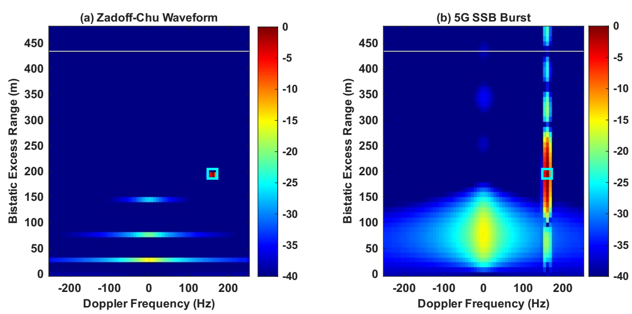

Figure 7. Range-Doppler maps for the far scenario ($\tau_s = 20$, $R_{bi} = 195.3$ m, $f_d \approx 160$ Hz), equal mean transmit power. (a) Extended ZC: single clean target cell. (b) Full SSB burst: the target is recovered but an OFDM-induced ridge spans all range bins at the target Doppler, with a broad low-Doppler pedestal below 150 m.

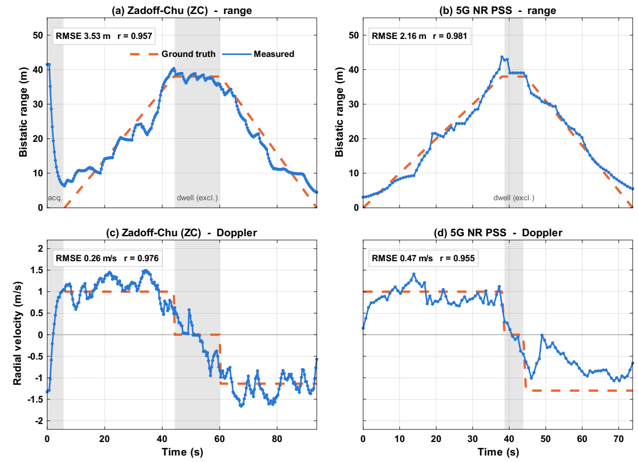

Figure 8. Over-the-air corridor results. (a),(c) extended ZC pilot ($N_{ZC} = 16381$); (b),(d) 5G NR PSS pilot (length-127 m-sequence, repeated). Measured range and radial velocity (solid) against the constant-speed reference trajectory (dashed). Shaded spans are the acquisition transient and the dwell at the far end, both excluded from the error statistics.

Table 3. Over-the-Air Corridor Results

Metric	**ZC**	**NR PSS**
Detections	265	102
Frames with no detection	0	23
Detection rate	100%	81.6%
Median detection magnitude	56.8	6.1
Range span tracked	4.5–40.4 m	3.1–43.7 m
Range RMSE	3.53 m	2.16 m
Range correlation r	0.957	0.981
Velocity RMSE	0.26 m/s	0.47 m/s
Velocity correlation r	0.976	0.955
Implied update rate	2.8 Hz	1.4 Hz

6.2. Tracking Results

Fig. 8 compares the reported range and radial velocity against the expected trajectory for both waveforms. The reference trajectory is a constant-speed walk over the 38 m corridor at a nominal 1 m/s. The interval during which the subject waited at the far end is shaded and excluded from the error statistics, as is the initial acquisition transient in the ZC run, during which the reported range collapses from 41.5 m to 6.3 m in 18 frames while the detection magnitude rises by a factor of 40, filter convergence rather than physical motion.

Timing caveat: The detector logs report range and velocity per frame but carry no timestamps, so the time axis is reconstructed by assuming a constant frame interval calibrated such that the outbound leg corresponds to the 38 m corridor at 1 m/s. This yields an implied update rate of 2.8 Hz for ZC and 1.4 Hz for the PSS pilot, consistent with the decimation applied ahead of the correlator. All time-referenced quantities below inherit this assumption, and the range errors in particular are conditioned on the assumed constant walking speed. Logging wall-clock timestamps would remove this dependence entirely and is the single most valuable change for future runs.

Table 3 summarises the comparison. Both waveforms track the walk over the full corridor, with range correlations above 0.95 against the reference trajectory, and neither loses the target during the turn at the far end.

6.3. Discussion

The measured comparison is more nuanced than the simulation alone would suggest, and it is worth stating plainly in both directions.

The extended ZC pilot is clearly superior on *detection robustness*. Its median detection magnitude is roughly $9\times$ that of the PSS pilot, which is what the $16381/1270$ ratio in correlation length predicts. A longer reference results in processing gain, and here that gain appears directly in the measurement. The ZC pilot produced a detection on every processed frame, whereas the PSS run failed to detect on 23 of 125 frames, an 81.6% detection rate with the dropouts concentrated at the far end of the corridor where the return is weakest. It also sustained twice the update rate. For a sensing system a track that never drops is worth a great deal, and the direction of this result is consistent with the clean correlation floor the ZC reference gives in simulation (Fig. 7).

The PSS pilot, however, gave the *better range track* in this experiment with 2.16 m RMSE against 3.53 m for ZC, with a slightly higher correlation. This does not follow from the simulation, and we do not claim it as a waveform property on the strength of a single run per waveform. Inspection of Fig. 8(a) shows the reason lies in the ZC track rather than the PSS one, between roughly 10 s and 40 s the reported ZC range advances much more slowly than a constant-speed walk while the reported ZC velocity remains near 1 m/s, so the range estimate stalls while the Doppler estimate does not. The most likely cause is that the correlation peak is held in a neighbouring range bin by a strong static corridor return that the $K = 5$ canceller has suppressed but not eliminated, with the interpolator of Eq. (15) then biased toward that residual.

Two further caveats bound what should be read from these numbers. The reference trajectory assumes a constant walking speed, so genuine pauses by the subject are charged as range error. And the nominal reference begins at 0 m whereas the detector never reports below 4.5 m (ZC) or 3.1 m (PSS), because the near-range bins are dominated by direct-path leakage and the LMS span. The reported range biases are nonetheless small (-0.58 m and $+0.02$ m), so the errors are spread rather than systematic.

7. Conclusion

We presented an open-source GNU Radio framework for bistatic sensing with 5G NR synchronization waveforms, combining adaptive LMS cancellation, FFT-based matched filtering and exponential-average MTI as three reusable streaming out-of-tree blocks running in real time at 30.72 MS/s on NI USRP-2901 hardware.

MATLAB simulation of a drone target across bistatic excess ranges of 19.5 - 195.3 m gave target-to-clutter ratios above 35 dB with correct range and Doppler recovery, and showed that an extended ZC reference yields a markedly cleaner range-Doppler map than the full SSB burst. Over the air, the chain tracked a walking human continuously across a 38 m corridor with both an extended ZC pilot and the 5G NR PSS. The ZC pilot detected on every frame with roughly $9\times$ the detection magnitude of the PSS, while the PSS gave a slightly better range track in this particular run, for reasons that appear to lie in a range-bin bias in the ZC track rather than in the waveform itself.

References

Abratkiewicz, Karol, Ksiezyk, Adam, Plotka, Marek, Samczynski, Piotr, Wszolek, Jacek, and Zielinski, Tomasz Piotr. SSB-Based Signal Processing for Passive Radar Using a 5G Network. *IEEE J. of Sel. Topics in Appl. Earth Observ. and Remote Sens.*, 16:3469–3484, 2023.

GNU Radio Project. GNU Radio: The Free & Open Source Radio Ecosystem. https://www.gnuradio.org. version 3.10, accessed July 2026.

Howland, Paul E, Maksimiuk, D, and Reitsma, Gjalt. Fm radio based bistatic radar. *IEE proceedings-radar, sonar and navigation*, 152(3):107–115, 2005.

Lin, Xingqin, Li, Jingya, Baldemair, Robert, Cheng, Jung-Fu Thomas, Parkvall, Stefan, Larsson, Daniel Chen, Koorapaty, Havish, Frenne, Mattias, Falahati, Sorour, Grovlen, Asbjorn, et al. 5g new radio: Unveiling the essentials of the next generation wireless access technology. *IEEE Communications Standards Magazine*, 3(3): 30–37, 2019.

Liu, Fan, Cui, Yuanhao, Masouros, Christos, Xu, Jie, Han, Tony Xiao, Eldar, Yonina C, and Buzzi, Stefano. Integrated sensing and communications: Toward dual-functional wireless networks for 6g and beyond. *IEEE journal on selected areas in communications*, 40(6): 1728–1767, 2022.

National Instruments. USRP-2901. https://www.ni.com/en/shop/hardware/software-defined-radios/model-usrp-2901. Accessed: September 19, 2026.

Widrow, Bernard and Stearns, Samuel D. Adaptive signal processing. *Englewood Cliffs*, 1985.

Willis, Nicholas J. *Bistatic Radar*. SciTech Publishing, 2 edition, 2005. ISBN 978-1-891121-45-6.