

Fast computational methods for panoramic spectrum analysis with hardware artifacts compensation for SDR-based direction finding applications

Fedir Katushonok^{1,*}, Olga Shcherbyna^{1,†}, Oleksandr Zadorozhnyi^{1,†} and Anatolii Simchuk^{1,†}

¹National University "Kyiv Aviation Institute", Liubomyra Huzara Ave., 1, Kyiv, 03058, Ukraine

Abstract

This article contributes to the advancement of both the theory and practice of radio emission source (RES) direction finding (DF) systems and panoramic spectrum analysis methods. The practical significance of this work is in demonstrating how different panoramic spectrum analysis methods can be implemented on popular SDR platforms, and how hardware artifacts can be compensated for to improve the accuracy and performance of radio emission source detection. Radio emission source direction finding is a critical technology in various applications, including wireless communication, navigation, and surveillance. With the advent of software-defined radio (SDR) platforms, radio emission source direction finding has become more accessible and cost-effective. However, SDR-based DF systems often suffer from hardware artifacts that can degrade the accuracy of direction estimation, as well as limited bandwidth, which makes it difficult to determine the angular position of a radio emission source with unknown frequency. In this paper, fast computational methods for panoramic spectrum analysis with hardware artifacts compensation for SDRs are being proposed. The article compares the performance, efficiency and accuracy of different panoramic spectrum analysis methods for 2 of the most popular SDR platforms with wide range on the market: HackRF One and bladeRF micro 2.0.

Keywords

radio emission source, direction finding, software-defined radio, panoramic spectrum analysis

1. Introduction

Radio emission source (RES) direction finding (DF) remains a cornerstone technology in modern radio engineering [1, 2], with applications ranging from civil aviation navigation aids and search-and-rescue operations to spectrum monitoring, electronic warfare and counter-unmanned-aerial-system (C-UAS) defense [3, 4, 5]. Traditional DF installations have historically relied on purpose-built, narrow-band receivers tuned to a known frequency of interest. Such systems offer high sensitivity and angular accuracy, but they are expensive, inflexible and ill-suited to scenarios in which the carrier frequency of the target emission is unknown a priori or changes rapidly, as is increasingly common with agile communication waveforms, frequency-hopping radios and consumer-grade drone control links.

The rapid maturation of software-defined radio (SDR) over the last decade has substantially lowered the entry barrier for both research and operational DF systems [6, 7, 8, 9, 10, 11]. General-purpose SDR front-ends such as the HackRF One and the bladeRF micro 2.0 provide hundreds of megahertz of tuning range, configurable sample rates and an open software stack at a price point that makes multi-channel and distributed DF architectures economically viable. At the same time, these affordable platforms inherit the limitations of their consumer-oriented analogue front-ends and low-cost mixers: the instantaneous bandwidth is bounded (a few tens of MHz for the HackRF, up to several tens of

CMiGIN'26: Fifth International Workshop on Conflict Management in Global Information Networks, June 26–27, 2026, Krakow, Poland

*Corresponding author.

†These authors contributed equally.

✉ 5307190@stud.kai.edu.ua (F. Katushonok); olha.shcherbyna@npp.kai.edu.ua (O. Shcherbyna);

oleksandr.zadorozhnyi@npp.kai.edu.ua (O. Zadorozhnyi); 8863303@stud.kai.edu.ua (A. Simchuk)

🌐 <https://github.com/CeSiumUA> (F. Katushonok)

🆔 0009-0008-3483-9871 (F. Katushonok); 0000-0002-6058-2749 (O. Shcherbyna); 0000-0002-2660-8829 (O. Zadorozhnyi);

0009-0007-8041-2724 (A. Simchuk)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

MHz for the bladeRF), and the I/Q signal path is contaminated by hardware artifacts - in particular the DC offset and LO leakage spike at band centre, I/Q imbalance images, spurious tones and frequency-dependent gain ripple [12]. When such a receiver is repeatedly retuned across a wide search range in order to locate an unknown emitter, these impairments translate into false detections, missed signals and biased bearing estimates, undermining the very advantage that the wide tuning range was meant to provide. Panoramic spectrum analysis - the rapid sweep of a wide frequency range followed by detection and parameter estimation of the strongest emissions - is therefore a prerequisite step for any practical SDR-based DF system that has to operate against emitters of unknown frequency [13]. The quality of the panoramic scan directly determines the upper bound on the DF system's reaction time, probability of detection and the fidelity of the bearing estimates produced in subsequent processing stages. Two competing requirements drive the design space: on the one hand, the sweep must be fast enough to capture short transmissions and to keep up with frequency-agile emitters; on the other hand, it must be clean enough - free of hardware artifacts and tuning transients - to avoid corrupting the downstream DF estimator [14]. Existing tools such as `hackrf_sweep` and `rtl_power` address only the first requirement, while published artifact-compensation techniques are typically demonstrated on laboratory-grade equipment and rarely benchmarked across heterogeneous low-cost SDR platforms.

The objective of this work is to formulate, implement and experimentally compare a family of high-speed panoramic spectrum analysis methods tailored to the constraints of low-cost SDR receivers used in DF applications, with explicit compensation of hardware artifacts. Specifically, the contribution of the paper is threefold. First, several panoramic scanning strategies are formalized, including the naive contiguous retune-and-capture scheme, an oversampled variant that trades sample rate for fewer retune events, and a side-FFT approach that discards the corrupted central bins around the LO and stitches the remaining usable spectrum into a seamless panorama [15]. Second, all methods are implemented uniformly on two representative SDR platforms covering the bulk of the affordable market - the HackRF One (8-bit, half-duplex, up to 20 MS/s) and the bladeRF micro 2.0 (12-bit, full-duplex, up to 61.44 MS/s) - so that platform-specific behavior can be disentangled from algorithmic behavior. Third, the methods are evaluated through a series of progressively more demanding experiments (smoke tests, settling-time sweeps, side-FFT validation, quick scans and oversampling studies) using objective metrics of sweep rate, achievable revisit time, residual artifact level and spectral fidelity, allowing concrete guidance to be derived for designers of SDR-based DF systems.

The remainder of the paper is organized as follows. Section 2 describes the mathematical formulation of the panoramic scanning algorithm, its performance metrics, and scanning artifacts that can be present during different scan techniques. Section 3 describes the measurement set-up, including the two SDR receivers, the 5.8 GHz video transmitter used as a controlled target emitter, and the receive antennas characterized with a vector network analyzer. Section 4 reports the experimental results obtained for each method on both platforms and compares them against the formulated performance criteria. Section 5 summarizes the findings and outlines directions for future work.

2. Problem statement

2.1. Problem formulation and performance metrics

In direction finding systems, sometimes the exact frequency (channel) of a radio-emission source is unknown. Thus a direction-finding system based on an SDR receiver of instantaneous bandwidth $B \ll B_{\text{tot}}$ must cover a search range $[f_{\text{start}}, f_{\text{stop}}]$ of total width $B_{\text{tot}} = f_{\text{stop}} - f_{\text{start}}$ in order to detect a target emission source whose carrier frequency $f_c \in [f_{\text{start}}, f_{\text{stop}}]$ is not known in a moment. Since $B_{\text{tot}} \gg B$, the resulting spectral panorama of the search range needs to be composed from multiple scanning tiles, each tile being received by tuning the SDR to some fixed frequency.

Let f_k represents the centre frequency of the k -th tile, with $k = 0, 1, \dots, K - 1$, K being the total number of tiles required to cover the search range, then the set of tile centres $\{f_k\}$

$$f_k = f_{\text{start}} + \frac{B}{2} + k \Delta f, \quad K = \left\lceil \frac{B_{\text{tot}}}{\Delta f} \right\rceil, \quad (1)$$

where $\Delta f \leq B$ is the *tile step* - the frequency increment between two consecutive LO settings. Choosing $\Delta f < B$ introduces an overlap of $B - \Delta f$ between adjacent tiles, which is later exploited to discard frequency bins corrupted by hardware artifacts (Section 2.2).

For each tile centre f_k , the SDR is tuned, allowed to settle for a time T_{settle} , and then a complex baseband segment $x_k[n]$, $n = 0, \dots, N_{\text{cap}} - 1$, is captured at sample rate $f_s = B$. The tile being captured is then processed by an overlapping, Hann-windowed short-time FFT of size N_{FFT} with overlap factor $\alpha \in [0, 1)$, producing the resulting averaged power spectral density (PSD)

$$P_k[m] = \frac{1}{M_k} \sum_{i=0}^{M_k-1} \left| \sum_{n=0}^{N_{\text{FFT}}-1} w[n] x_k[n + iH] e^{-j2\pi mn/N_{\text{FFT}}} \right|^2, \quad m = 0, \dots, N_{\text{FFT}} - 1, \quad (2)$$

where $w[n]$ is the Hann window, $H = \lfloor N_{\text{FFT}}(1-\alpha) \rfloor$ is the frame hop and $M_k = \lfloor (N_{\text{cap}} - N_{\text{FFT}})/H \rfloor + 1$ is the number of frames being averaged. The resulting bin frequencies, expressed in the RF domain after `fftshift`, are

$$f_k[m] = f_k + \left(m - \frac{N_{\text{FFT}}}{2}\right) \frac{f_s}{N_{\text{FFT}}}. \quad (3)$$

The final panoramic PSD $P(f)$ over $[f_{\text{start}}, f_{\text{stop}}]$ is obtained by selecting, in each tile, only the bins that lie in a *usable* central band of width $\Delta f_u \leq \Delta f$ and concatenating them:

$$P(f) = \bigcup_{k=0}^{K-1} \left\{ (f_k[m], P_k[m]) \mid |f_k[m] - f_k| \leq \frac{\Delta f_u}{2} \right\}. \quad (4)$$

The two parameters Δf (tile step) and Δf_u (kept band) define the main degrees of freedom that distinguish the scanning methods: the naive method takes $\Delta f_u = \Delta f \approx B$, the side-FFT method takes $\Delta f_u < \Delta f$ to discard the corrupted central band around the LO, and the oversampled method increases f_s above the nominal B so that more bins per tile fall outside the artifact region.

The proposed methods are compared using four objective performance metrics, defined as follows. *Scan time*. The total wall-clock time required to cover the search range,

$$T_{\text{scan}} = \sum_{k=0}^{K-1} \left(T_{\text{tune}}^{(k)} + T_{\text{settle}} + T_{\text{cap}}^{(k)} + T_{\text{fft}}^{(k)} \right), \quad (5)$$

where $T_{\text{tune}}^{(k)}$ is the LO retune latency, T_{settle} is the post-retune settling delay, $T_{\text{cap}}^{(k)} = N_{\text{cap}}/f_s$ is the capture duration, and $T_{\text{fft}}^{(k)}$ is the time spent computing Equation (2). In this work T_{scan} is reported as the median of $n = 3$ independent runs together with their inter-quartile range.

Sweep rate. The effective speed of the scanner, expressed as the search bandwidth covered per unit wall-clock time,

$$R_{\text{sweep}} = \frac{B_{\text{tot}}}{T_{\text{scan}}} \quad [\text{Hz/s}]. \quad (6)$$

A higher R_{sweep} implies a shorter revisit time $T_{\text{rev}} = T_{\text{scan}}$ between two consecutive observations of the same frequency, which directly bounds the probability of intercepting short, bursty emissions.

Residual artifact level. Even after stitching, the panoramic PSD $P(f)$ may exhibit residual spurs caused by DC/LO leakage and I/Q imbalance images at the boundary between adjacent tiles, and the ratio between the peak amplitude of the strongest noise spike in a signal-free portion of a spectrum and the median noise floor of the same portion represents the level of residual artifact:

$$A_{\text{res}} = 10 \log_{10} \left(\frac{\max_{f \in \mathcal{F}_{\text{noise}}} P(f)}{\text{median}_{f \in \mathcal{F}_{\text{noise}}} P(f)} \right) \quad [\text{dB}], \quad (7)$$

where $\mathcal{F}_{\text{noise}}$ is a frequency subset known a priori to be free of intentional emissions in the laboratory environment. Lower A_{res} is better.

Spectral fidelity at the target. This value shows how accurate the known signal is being reproduced. The PSD value at the known VTX carrier f_{VTX} is compared to the reference value from single-tile capture on f_{VTX} :

$$\varepsilon_{\text{fid}} = |P(f_{\text{VTX}}) - P_{\text{ref}}(f_{\text{VTX}})| \quad [\text{dB}]. \quad (8)$$

This metric represents the summing effect of windowing, averaging, artifact-induced bias and any frequency offset caused by the stitching scheme being selected.

2.2. Hardware artifacts of direct-conversion SDR front-ends

Both HackRF One and bladeRF SDRs are direct-conversion transceivers: the incoming RF signal is mixed with a local oscillator tuned exactly to the frequency being tuned to f_k , and the resulting complex baseband is digitized by an ADCs. This architecture has a benefit of avoiding an explicit IF stage, but it introduces some set of artifacts at and/or symmetric to LO frequency in Equation (2). These artifacts can be classified to 4 different categories.

DC offset and LO leakage. Mismatches in the I and Q signal paths and finite isolation between the LO and the RF port produce a non-zero average value at baseband, which appears as a spike at the bin $m = N_{\text{FFT}}/2$ of every tile. The amplitude of this spike depends on several factors, the most significant are the receiver gain, the temperature of the front-end and the tuning frequency.

I/Q imbalance images. Due to even a small amplitude or phase mismatch between the I and Q channels, each true tone at $f_k + \Delta f$ can have an additional mirror image at $f_k - \Delta f$. After stitching tiles together, a resulting panorama can contain these emissions located symmetrically around each tile centre and can be wrongly classified as a target signal by a peak-based detector.

Spurious tones. Some narrow-band tones in baseband frequency grid can be present due to various factors, like internal clock harmonics, USB-related switching noise and PLL reference spurs. After retuning, these tones appear at the same baseband offset in every tile, which means they draw a regular pattern across the stitched panorama at frequencies $f_k + \Delta f_{\text{spur}}$ for the same Δf_{spur} in each tile.

Frequency-dependent gain ripple. The measured gain of a received frame can vary across different frequencies not only because of different noise or target signal levels, but also due to the unevenness of the analog front-end amplitude response. Usually, the gain rolls off on the edges of the baseband filter and can contain tiny ripples across the usable band. Despite the deterministic behavior of this artifact for a given gain and frequency, detected power peak can be biased and some discontinuities could appear between the adjacent tiles while stitching.


Combination of these 4 artifact types motivate the development of the scanning methods. The naive scan is the simplest one, being heavily impacted by all these 4 artifacts. Side-FFT method is able to suppress DC offset and LO leakage by discarding a central portion of each tile. The oversampled scan reduces the relative weight of the corrupted central bins by enlarging the usable fraction of each tile.

3. Measurement setup

All measurements reported in this paper were collected in a controlled indoor laboratory environment. Both the SDRs were tested sequentially in the same room, same target signal emitter and same receive antenna. A drone video transmitter (VTX) operating on a known, constant frequency acts as the reference emitter, while the two SDR receivers under investigation - a HackRF One and a bladeRF micro 2.0 - observe its signal from a short, fixed distance through the same antenna. Both SDRs are driven from the same single-board computer over independent USB 3.0 links. Fig. 1 shows the receive bench with both SDRs and their antennas; the transmit side is shown separately in Fig. 2 and detailed in Section 3.2. The transmitter and the receive bench are placed approximately 2 m apart in line-of-sight, with no intentional attenuator inserted between the VTX and its antenna. This setup ensures a strong direct-path signal at the receivers, significantly stronger than the noise floor of both SDRs across the 4.9-6.0 GHz scan range.



(a)



(b)

Figure 1: Receivers used in all experiments. (a) HackRF One with the log-periodic antenna. (b) bladeRF micro 2.0 xA4 with the same log-periodic antenna.

3.1. SDR receivers

Two SDR platforms selected are those that are mostly used among engineers, military and radio amateurs and are accessible and affordable on mass-market: the Great Scott Gadgets HackRF One and the Nuand bladeRF micro 2.0 xA4. Both devices are direct-conversion (zero-IF) transceivers with a tuning range that comfortably covers the 5.8 GHz video band of interest, but they differ significantly in ADC resolution, maximum sample rate and USB throughput, which makes them complementary for benchmarking panoramic scanning methods. The key parameters used in the experiments are summarized in Table 1.

Table 1

Key parameters of the SDR receivers used in this study.

Parameter	HackRF One	bladeRF micro 2.0 xA4
Frequency range	1 MHz - 6 GHz	47 MHz - 6 GHz
Architecture	Direct conversion	Direct conversion
ADC resolution	8 bit	8 bit, 12 bit
Max. sample rate	20 MS/s	120 MS/s (8 bit), 61.44 MS/s (12 bit)
Sample rate used	20 MS/s	120 MS/s (oversample, 8 bit), 60 MS/s (12 bit)
Duplex	Half	Full
Host interface	USB 2.0	USB 3.0
RF front-end gain stages	AMP / LNA / VGA	Unified (AGC off)
Gain settings used	LNA 40, VGA 50, AMP off	40 dB manual

Both receivers are connected to the same Raspberry Pi 5 host through independent USB 3.0 ports. All driver- and userspace-side software is identical across runs: the scanning code is implemented in Python on top of the `pyhackrf2` and `bladerf` bindings, runs in the same Jupyter kernel, and is exercised on each device in turn through a common high-level scan loop, so platform-specific behavior can be cleanly separated from algorithmic behavior.

3.2. Target emitter: 5.8 GHz video transmitter

A commercial 5.8 GHz analogue FPV video transmitter (Fig. 2) is used as the reference emitter throughout this work. The unit covers the standard 48-channel 5.8 GHz video band and is configured to operate on channel 22, which corresponds to a nominal carrier frequency of $f_{\text{VTX}} = 5752$ MHz and a declared RF output power of 600 mW. No video source is connected to the modulator input, so the transmitter encodes a constant black frame; in the spectral domain this produces an essentially stationary, narrow-band CW-like emission centred on f_{VTX} , which is ideal for benchmarking panoramic scanners because the target feature does not move, fade or modulate during a sweep. The transmitter is mounted on a small

tripod approximately 1 m above floor level and equipped with the manufacturer-supplied rubber-duck antenna, oriented vertically.



Figure 2: Reference emitter: 5.8 GHz, 48-channel analogue video transmitter configured on channel 22 ($f_{\text{VTX}} = 5752 \text{ MHz}$, $P_{\text{out}} = 600 \text{ mW}$).

3.3. Receive antenna

To keep the cross-platform comparison fair, both SDRs are equipped with the same directional log-periodic antenna [16, 17] with SMA connector (Fig. 3). The antenna characteristics were measured with a LiteVNA portable vector network analyzer over a 4.8 GHz-6 GHz sweep range in order to verify the voltage standing wave ratio (VSWR) in the band of interest around the VTX channel. Fig. 4 shows the resulting VSWR trace, with the marker M1 placed at the closest available sweep point to the VTX carrier (5750 MHz, limited by the analyzer's frequency-step resolution).




Figure 3: Receiver log-periodic antenna used in all the measurements




Figure 4: VSWR trace of a receive antenna measured with a LiteVNA over 1 GHz-6 GHz. The marker M1 at 5750 MHz lies near the VTX channel 22 carrier used throughout the experiments.

4. Experimental results

4.1. Smoke tests and baseline capture

In the first experimental phase, which is smoke test, both the receivers are tuned to the VTX carrier frequency, $f_c = f_{\text{VTX}} = 5752$ MHz, and a single block of $N_{\text{cap}} = 60\,000$ complex samples is received from each SDR device. Both SDRs are configured to their maximum native parameters: 12-bit mode with $f_s = 60$ MS/s samplerate for bladeRF, and $f_s = 20$ MS/s samplerate for HackRF. This capture is an idealized scenario that represents the best possible operating conditions for both SDRs, as there is not frequency retuning, thus LO settling and tile-stitching are absent. For each capture the averaged PSD defined in Equation (2) is calculated with $N_{\text{FFT}} = 4096$, a Hann window and 50 % overlap, yielding $M = 14$ averaged frames for the 60 000-sample blocks. The resulting baseband bin resolution is $\Delta f_{\text{bin}} = f_s/N_{\text{FFT}} \approx 14.65$ kHz for the bladeRF and ≈ 4.88 kHz for the HackRF. Fig. 5 shows the resulting spectrum, normalized to the median noise floor of each receiver.




Figure 5: Single-tile spectrum at $f_c = f_{\text{VTX}} = 5752$ MHz, normalized to the median noise floor of each receiver. Left: bladeRF micro 2.0 at 60 MS/s. Right: HackRF One at 20 MS/s.

4.2. Settling-time sweep

The SDR retuning requires some settling time to be waited (or some number of samples to be drained) in order to stabilize after a frequency change. The optimal settling time is ideally short enough to keep scanning rate high, the same time long enough to ensure that LO became stable on a new frequency, producing artifact-less results. For this purpose, some of the settling time variations were tested, measuring four most significant indicators: median scan time (Equation (5)), mean and standard deviation of the detected peak power at f_{VTX} , and the standard deviation of the detected peak frequency σ_f . All the results are summarized in Table 2.

Table 2

bladeRF settling-time sweep over $[4.9, 6.0]$ GHz. N_{drain} is the number of post-retune samples discarded at $f_s = 60$ MS/s; the corresponding settling time is $T_{\text{settle}} = N_{\text{drain}}/f_s$.

N_{drain}	T_{settle}	T_{scan} [s]	$\bar{P}_{\text{peak}}$ [dB]	σ_f
1 024	17 μ s	0.598	$+11.0 \pm 8.5$	2.51 MHz
131 072	2.18 ms	0.607	-12.3 ± 0.6	4.15 MHz
262 144	4.37 ms	0.607	0.3 ± 4.9	3.32 MHz
1 048 576	17.48 ms	1.101	$+27.8 \pm 0.6$	7.18 kHz

Fig. 6 and Fig. 7 illustrate the same data graphically: the former plots T_{scan} , σ_P and σ_f as a function of the drain-buffer size, while the latter shows the resulting stitched panoramas, in which insufficient settling time causes the artifacts in the lower-frequency tiles with retune transients.




Figure 6: Settling-time sweep metrics on the bladeRF: scan time, VTX peak-power standard deviation and VTX peak-frequency standard deviation versus drain-buffer size (5 repetitions per configuration).




Figure 7: Stitched bladeRF panoramas across the 4.9–6.0 GHz search range for four drain-buffer sizes from 17 μs to 17.5 ms. The step-shaped artifacts in the upper subplots result from the LO still slewing during the capture window and gradually disappear as the settling delay increases; only at $N_{\text{drain}} = 2^{20}$ does the VTX carrier emerge as an isolated peak at 5.752 GHz.

Thus, some preliminary conclusions can be made. Firstly, for any T_{settle} below about 5 ms the detected carrier frequency drifts does not match the expected value, and it drifts by several megahertz between tile retunes (σ_f on the order of 10^6 Hz), which means that during the capture window being received the LO synthesizer is still unstable, causing artifacts in a measurement. Thus, only at $N_{\text{drain}} = 2^{20}$ samples (≈ 17.5 ms) σ_f collapses to about 7 kHz, which corresponds to a single FFT bin, which is the state that can be considered as stable. Secondly, the detected peak power, as well as detected carrier frequency, is unreliable at short delays: only at 17.5 ms it stabilizes to a +27.8 dB, indicating a transition to a stable working condition. Third, the settling time variation has a direct scanning speed impact: increasing T_{settle} from 17 μs to 17.5 ms nearly doubles T_{scan} (from 0.60 s to 1.10 s), which degrades the direction finding performance in cost of a measurement stability.

According to an experiment conducted, $T_{\text{settle}} \approx 17$ ms is the minimum delay that produces a frequency-accurate, amplitude-stable detection on the bladeRF, this value is being selected as a default for all the subsequent experiments.

4.3. Side-FFT validation against the VTX target

The side-FFT method is validated on the HackRF, where the central DC/LO spike characterized in Section 4.1 is most severe. Two panoramic scans of the same 4.9-6.0 GHz range are compared back-to-back: a naive scan (phase 2) and a side-FFT scan (phase 3) using identical tile geometry but with the ± 100 central bins around each LO replaced by the corresponding bins of an overlapping neighbor tile. Fig. 8 overlays the two resulting panoramas.




Figure 8: HackRF panoramic scans of the 4.9-6.0 GHz range. Top: naive retune-and-capture (phase 2, $T_{\text{scan}} = 2.73$ s). Bottom: side-FFT with central-bin rejection (phase 3, $T_{\text{scan}} = 4.05$ s).

Comparing the results of a naive and side-FFT methods, it can be clearly seen that the massive DC spikes spaced by the tile step Δf present in a naive method are absent in the side-FFT method, while the overall wideband structure of the panorama is preserved, confirming the correctness of the obtained results. However, this improvement is a trade-off between scanning speed and spectral purity. The VTX frequency is being detected with the same ≈ 130 kHz error, the same as with a naive method. This confirms that the side-FFT method does not degrade a frequency detecting accuracy.

4.4. Quick scans and oversampling study

A noticeable amount of time that takes SDR to change its operating frequency is a host-to device communication and LO setting. While trying to maximize the scanning speed, this can lead to poor

performance result. For this case, most of the modern SDRs, as well as bladeRF and HackRF being studied, have a firmware-assisted quick-tune mechanism. This mechanism allows to save a significant amount of time by pre-calculating LO parameters for all the frequencies that SDR will be tuned to during the scan. Since the range being scanned remains the same, instead of tuning to a new frequency each time, the SDR firmware just takes already pre-calculated LO settings from the memory. Fig. 9 illustrates results being obtained by comparing bladeRF and HackRF (in both naive and side-FFT modes) scanning modes.

Another method that allows to improve scanning speed, is an oversampling mode. From the 2 SDRs being taken in this research, only bladeRF supports oversampling, since it's ADC has a native 12-bit resolution, and the oversampling is achieved by changing resolution from 12 to 8-bit, dramatically increasing maximal sample rate allowed from 60 MS/s to 120 MS/s. Fig. 10 depicts how oversample relates to a normal sample rate in both naive and quick tune mode.




Figure 9: Quick-tune compared to naive panoramas on both platforms.

Quick-tune method provides a $5.3\text{--}6.6\times$ speed-up on either platform without introducing any additional noise or artifacts, because the overhead it removes is independent of the bandwidth covered. Doubling the sample rate adds another $\sim 1.5\times$ on the bladeRF by halving the tile count, at the cost decreased resolution, which is not critical for scanning. The VTX carrier remains clearly detectable at ≥ 30 dB above the floor. Combining these 2 methods builds up a fastest possible configuration of this study: 0.093 s for the full 1.1 GHz sweep on the bladeRF, i.e. a sweep rate of $R_{\text{sweep}} \approx 11.8$ GHz/s.




Figure 10: bladeRF 120 MS/s SC8_Q7 oversampled panoramas (naive and quick-tuned).

4.5. Methods comparison and cross-platform summary

The panoramas obtained with every method on both platforms are assembled side-by-side in Fig. 11 and the associated metrics are summarized in Table 3.




Figure 11: All seven panoramic configurations evaluated in this study on the common 4.9–6.0 GHz search range with the VTX active at $f_{\text{VTX}} = 5752$ MHz. Top row: bladeRF (60 MS/s SC16_Q11 naive and quick-tune, 120 MS/s SC8_Q7 naive and quick-tune). Bottom row: HackRF (naive 20 MS/s, side-FFT 20 MS/s, firmware-sweep quick-tune).

Table 3

Cross-platform summary of the panoramic scanning methods on the 4.9-6.0 GHz range.

Device	Method	T_{scan} [s]	Peak err. [kHz]	Noise [dB]	Speed-up
bladeRF	Naive 60 MS/s SC16_Q11	0.921	41.0	−17.4	×1.00
bladeRF	Quick-tune 60 MS/s SC16_Q11	0.173	41.0	−17.4	×5.34
bladeRF	Naive 120 MS/s SC8_Q7	0.632	31.2	−0.4	×1.46
bladeRF	Quick-tune 120 MS/s SC8_Q7	0.093	41.0	0.0	×9.94
HackRF	Naive 20 MS/s	2.728	124.0	19.0	×1.00
HackRF	Side-FFT 20 MS/s	4.046	133.8	18.6	×0.67
HackRF	Firmware-sweep quick-tune	0.414	133.8	19.5	×6.60

5. Conclusions

This work benchmarked three panoramic scanning methods - naive retuning , side-FFT with central segment rejection and SDR firmware-assisted quick-tune - on two affordable mass-market SDRs, the HackRF One and the bladeRF micro 2.0, using a 5.8 GHz FPV video transmitter as a known fixed central frequency signal over a 4.9-6.0 GHz search range.

Three main conclusions can be made. First, a sufficient settling delay (about 17 ms on the bladeRF) is a prerequisite for any meaningful comparison: below this threshold both the detected carrier frequency and its power are dominated by retune transients rather than by the scanning algorithm. Second, firmware-assisted quick-tune is the single most effective optimization, yielding a 5.3-6.6× speed-up on both platforms without measurable degradation of the panorama. Third, bandwidth-oriented techniques are platform-specific: 8-bit oversampling at 120 MS/s adds a further $\sim 1.5\times$ on the bladeRF, while on the HackRF the side-FFT method trades a 0.67× scan-time penalty for the suppression of DC/LO spikes. Combining quick-tune with oversampling on the bladeRF reaches a sweep rate of $R_{\text{sweep}} \approx 11.8$ GHz/s for the full 1.1 GHz range, while keeping the VTX carrier at ≥ 30 dB above the noise floor.

The calculated peak-frequency error is caused by the FFT bin resolution of the chosen configuration rather than by the scanning method itself, which confirms that the proposed methods do not bias the target carrier and that any further refinement of bearing estimation in the direction-finding chain must come from longer per-tile captures rather than from the panoramic scanner. Future work can include integrating advanced direction finding algorithms (like MUSIC, MVDR, ESPRIT) into the scanning loop, and testing the methods in more realistic scenarios with multiple emitters and frequency-hopping signals.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

- [1] S. Zhyla, E. Tserne, Y. Volkov, S. Shevchuk, O. Gribsky, D. Vlasenko, V. Kosharskyi, D. Kovalchuk, Statistical synthesis and analysis of functionally deterministic signal processing techniques for multi-antenna direction finder operation, *Computation* 12 (2024) 170. doi:10.3390/computation12090170.
- [2] A. Martian, C. Paleacu, I. M. Marcu, C. Vladeanu, Direction-finding for unmanned aerial vehicles using radio frequency methods, *Measurement* 235 (2024) 114883. doi:10.1016/j.measurement.2024.114883.
- [3] A. Rutkowski, A. Kawalec, Some of problems of direction finding of ground-based radars using monopulse location system installed on unmanned aerial vehicle, *Sensors* 20 (2020) 5186. doi:10.3390/s20185186.

- [4] H. A. Cook, K. M. T. E. Kahn, V. Balyan, Radio direction-finding techniques for an unmanned aerial vehicle, *Micro-Electronics and Telecommunication Engineering, Lecture Notes in Networks and Systems* 106 (2020) 1–10. doi:10.1007/978-981-15-2329-8_1.
- [5] N. Ruzhentsev, et al., Radio-heat contrasts of uavs and their weather variability at 12 ghz, 20 ghz, 34 ghz, and 94 ghz frequencies, *ECTI Transactions on Electrical Engineering, Electronics, and Communications* 20(2) (2022) 163–173. doi:10.37936/ecti-eec.2022202.246878.
- [6] I. Bae, J. Hong, Survey on the developments of unmanned marine vehicles: Intelligence and cooperation, *Sensors* 23 (2023) 4643. doi:10.3390/s23104643.
- [7] Z. Hussein, S. Francois, B. Oumaya, V. Thierry, Syloin: Measuring angle of arrival of lora signals using software defined radio, in: *Proceedings of International Conference on Indoor Positioning and Indoor Navigation, IEEE, Lloret de Mar Spain, 2021*, pp. 1–8. doi:10.1109/IPIN51156.2021.9662518.
- [8] Z. Dai, Y. He, V. Tran, N. Trigoni, A. Markham, Deepaoanet: Learning angle of arrival from software defined radios with deep neural networks, *IEEE Access* 10 (2022) 3164–3176. doi:10.1109/ACCESS.2021.3140146.
- [9] H. B. Tuğrel, H. Alakoca, G. K. Kurt, C. Ayyıldız, Angle of arrival (aoa) estimation by using software defined radios, in: *Proceedings of 24th Signal Processing and Communication Application Conference, IEEE, Zonguldak Turkey, 2016*, pp. 1429–1432. doi:10.1109/SIU.2016.7496018.
- [10] F. Katushonok, M. Zaliskyi, O. Shcherbyna, O. Zadorozhnyi, Reliability analysis of the angle of arrival system based on software defined radio applications, in: *2026 IEEE 18th International Conference on Advanced Trends in Radioelectronics, Telecommunications and Computer Engineering (TCSET), 2026*, pp. 1–5. doi:10.1109/TCSET65181.2026.11461095.
- [11] S. Zhyla, et al., Practical imaging algorithms in ultra-wideband radar systems using active aperture synthesis and stochastic probing signals, *Radioelectronic and Computer Systems* 1(105) (2023) 55–76. doi:10.32620/reks.2023.1.05.
- [12] P. Flak, Hardware-accelerated real-time spectrum analyzer with a broadband fast sweep feature based on the cost-effective sdr platform, *IEEE Access* 10 (2022) 110934–110946. doi:10.1109/ACCESS.2022.3215800.
- [13] M. Șorecău, E. Șorecău, A. Sârbu, P. Bechet, Real-time statistical measurement of wideband signals based on software defined radio technology, *Electronics* 12 (2023). URL: <https://www.mdpi.com/2079-9292/12/13/2920>. doi:10.3390/electronics12132920.
- [14] N. Bello, K. Ogbeide, Designing a real-time swept spectrum analyser with usrp b210, *Nigerian Journal of Environmental Sciences and Technology* 5 (2021) 329–339. doi:10.36263/nijest.2021.02.0275.
- [15] S. Seven, P. Rothana, S. Channareth, S. Sreypov, T. Brosopov, T. Tevy, U. Sokoeun, R. Vannak, N. Sros, Low-cost sdr spectrum analyzer and radio receiver using usrp n210, gnu radio, and raspberry pi 5, in: *2024 International Conference on Electrical Engineering and Computer Science (ICECOS), 2024*, pp. 7–12. doi:10.1109/ICECOS63900.2024.10791167.
- [16] O. Shcherbyna, O. Zadorozhnyi, F. Katushonok, Antennas Combination for Ground and Airborne Segments of UAV Communication Systems, 2025, pp. 231–250. doi:10.1007/978-3-031-91992-3_15.
- [17] O. Shcherbyna, F. Katushonok, O. Zadorozhnyi, M. Riabyi, Selection of antenna elements for a radio emission source detection and tracking system, in: *CEUR Workshop Proceedings, volume 4024, 2025*, pp. 144–156. URL: <https://ceur-ws.org/Vol-4024/paper12.pdf>.