

Direct RF and Heterodyne Receivers for Congested and Contested Electromagnetic Environments

A. J. WILEMAN

Loughborough University, UK

Abstract—Recent advances in high-speed RF-sampling data converters and integrated digital channelization are enabling direct radio-frequency (Direct RF) digitization and instantaneous-wideband receiver architectures. These architectures offer the potential for improved probability of intercept against agile emitters, reduced sensing-to-response latency, and increased multi-mission flexibility in increasingly congested electromagnetic environments. At the same time, contested environments impose stringent constraints arising from front-end linearity, blocking, spurious responses, sampling-clock jitter, converter dynamic range, and real-time processing capacity. This paper presents a system-level comparison of Direct RF and heterodyne receiver architectures using a common Monte Carlo simulation framework in which both architectures are evaluated against identical realizations of the electromagnetic environment. Seven operationally motivated aspects are examined: probability of intercept, sensing-to-response latency, detection performance under emitter congestion, compute-limited scalability, blocking robustness, spectral exploitation efficiency, and multi-emitter deinterleaving performance. The analysis identifies the architectural trade space between persistent wideband observation and sequential spectrally selective reception, highlighting the roles of instantaneous bandwidth, finite digital-processing capacity, and analogue front-end protection. The resulting comparison motivates hybrid receiver architectures that combine selective analogue protection with wideband digitization and scalable digital processing for future electromagnetic operations.

I. INTRODUCTION

Future electromagnetic (EM) operations are expected to take place in environments that are simultaneously congested, contested, and highly dynamic [1], [2]. The

proliferation of frequency-agile emitters, low-probability-of-intercept waveforms, distributed sensing platforms, and increasingly software-defined RF systems places growing demands on receiver architectures in terms of instantaneous bandwidth, spectral awareness, reconfigurability, dynamic range, and response latency [1], [2], [6]. In such environments, receiver effectiveness depends not only on sensitivity to an individual signal, but also on how rapidly and concurrently the surrounding spectrum can be observed, interpreted, and acted upon.



Fig. 1. Congested, contested, and highly dynamic EM environment [2]

Traditional heterodyne receivers have provided the dominant RF-receiver architecture for more than a century [3]. By translating selected portions of spectrum to one or more intermediate frequencies prior to digitization, they enable analogue selectivity, controlled gain distribution, and protection of downstream stages. When wide spectral coverage is required, however, conventional architectures commonly rely on sequential tuning or scanning. This introduces an inherent trade-off between instantaneous bandwidth, revisit time, probability of intercept, and response latency, particularly for short-duration or frequency-agile signals.

Direct radio-frequency (Direct RF) receiver architectures move the analogue-to-digital boundary closer to the antenna by digitizing signals at, or near, their received carrier frequencies. This concept follows the broader evolution of software-defined radio, in which increasing functionality is transferred from fixed analogue hardware into programmable digital processing [4]. Its practical realization has been enabled by continued improvements in high-speed data-converter performance [5], and, more recently, RF-sampling converters and integrated digital channelization [8]. Such developments increasingly permit multi-gigahertz instantaneous bandwidths to be processed

without conventional analogue frequency translation, enabling persistent observation of a substantially larger portion of the spectrum.

This increased spectral visibility is accompanied by a different set of constraints. Wideband RF-sampling receivers expose analogue front ends and data converters to a greater aggregate signal environment and therefore place demanding requirements on front-end linearity, converter dynamic range, sampling-clock purity, spurious-free performance, data movement, and real-time processing capacity [5], [8]. Conversely, strong analogue selectivity in a heterodyne receiver can provide valuable protection in blocker-dominated environments. The architectural question is therefore not simply whether Direct RF can replace the heterodyne receiver, but how the two approaches redistribute limitations and resources across the analogue–digital boundary.

Much of the historical development toward Direct RF has concentrated on converter capability, RF-sampling feasibility, software-defined architectures, and component-level performance [4], [5], [7], [8]. The operational consequences of this architectural transition are less commonly examined using a common system-level framework. In particular, persistent wideband observation affects not only instantaneous spectral coverage, but also probability of intercept, sensing latency, processing demand, susceptibility to blocking, and the continuity of observations available for multi-emitter association.

This paper therefore presents a comparative system-level study of heterodyne and Direct RF receiver architectures using a common Monte Carlo simulation framework. Both architectures are evaluated against equivalent realizations of a stochastic electromagnetic environment so that observed differences arise from receiver architecture rather than from differing signal scenarios. Seven operationally motivated aspects are examined: (i) probability of intercept for frequency-agile emitters, (ii) sensing-to-response latency, (iii) detection performance under increasing emitter density, (iv) compute-limited scalability, (v) blocking robustness, (vi) spectral exploitation efficiency, and (vii) multi-emitter deinterleaving performance.

The objective is not to identify a universally superior receiver architecture or to reproduce a particular hardware implementation. Instead, the analysis seeks to establish the regimes in which persistent wideband observation, sequential spectral selectivity, analogue protection, and finite digital-processing capacity become the dominant architectural constraints. This system-level perspective also provides a basis for considering hybrid receiver concepts in which selective analogue front-end protection is combined with wideband digitization and scalable digital processing.

The remainder of this paper is organized as follows. Section II defines the receiver architectures and modelling assumptions and places Direct RF within the historical progression from heterodyne reception to software-defined and RF-sampling systems. Section III describes the common simulation framework and performance metrics. Subsequent sections examine the individual architectural trade-offs, followed by a discussion of their implications for future EM receiver design and emerging receiver concepts. The final section summarizes the principal conclusions.

II. RECEIVER ARCHITECTURES AND ASSUMPTIONS

Modern electromagnetic operations are shaped not only by waveform design, but by the underlying receiver architectures that determine how spectrum is observed, processed, and acted upon. This section defines the receiver architectures considered in this work and establishes the modelling assumptions used for their system-level comparison in congested and contested electromagnetic environments. The objective is not to reproduce specific hardware implementations, but to abstract the dominant architectural characteristics that influence spectral awareness, agility, processing concurrency, and robustness.

TABLE I
Timeline of milestones relevant to direct RF and wideband reception

Era	Milestone	Why it matters for today's congested/contested spectrum
Early 1900s–1920s	Superheterodyne development described by Armstrong (invented ~1918; details public by 1919; presented later). [10–11]	Establishes the “translate to IF, filter hard” paradigm that remains crucial for blocker handling and selectivity.
1990s	“Software radio” framing: place A/D as close to antenna as feasible; programmability of RF bands and modes as a hallmark [12–13].	Connects receiver architecture directly to agility, re-tasking, and multi-mission reconfiguration—core needs in contested EMS.
Late 1990s	ADC performance characterization and limitations (Walden survey): resolution vs sample rate trends and aperture jitter as a key limit at high sample rates [9],[13].	Provides foundational metrics (SNR/SFDR/ENOB) and frames jitter and speed/resolution trade-offs central to direct RF feasibility.
2000s–2010s	“Digital IF” and monitoring receivers	Bridges measurement-

	evolve, emphasis on preselection + AGC + recording + DDCs to improve probability-of-intercept (POI) in dynamic, uncontrolled environments [14-17].	instrument architectures and operational monitoring/ES needs; highlights how wideband capture changes detection and replay.
2017	RFSoc integration: RF sampling data converters with on-chip DDC/DUC and reduced footprint/power [17-18].	Marks practical integration of RF sampling + digital processing in a single device class, enabling multi-channel wideband architectures.
2020	ISSCC: 12b 18GS/s RF-sampling ADC with integrated wideband track-and-hold and background calibration. [19-20]	Demonstrates the circuit techniques (wideband T/H + calibration) required to push RF sampling into higher GHz with usable SFDR/SNDR.
2022	Distributed direct sampling in phased arrays (16-channel S-band platform): direct sampling ADCs in higher Nyquist zones + embedded DSP [21-22].	Shows system-level multi-channel considerations: synchronization, coherent combining, noise/spur behavior, and DSP configuration.
2023	ISSCC: 24GS/s 12b time-interleaved ADC for direct RF sampling [23-24]	Illustrates that “tens of GS/s with ~12b class” is not just research—now a reported state-of-the-art with explicit RF-sampling intent.
2025	ISSCC: 12.8GS/s sub-sampling ADC front end with 38 GHz input bandwidth [25-26]	Extends the RF-sampling front-end concept toward X/Ka-band, limiting reliance on mixers in certain mm Wave bands.
2025	Microcomb-based photonic-electronic ADC demonstration with extremely large acquisition bandwidth [27]	Shows longer-range paths to surpass electronic clock-jitter limits via photonic clocks/slicing, albeit with complexity and SWaP challenges.

A. Heterodyne Receiver Model

The heterodyne, or superheterodyne, receiver has formed the basis of RF reception for more than a century [10], [11]. In this architecture, received radio-frequency signals are translated through one or more intermediate frequencies using analogue mixers and local oscillators prior to digitization. Selectivity and dynamic-range management

are achieved through the distribution of gain and filtering across the RF and IF stages.

From a system perspective, a scanning heterodyne receiver is characterized by a finite instantaneous bandwidth and sequential spectral coverage. At any given time, only the tuned passband is observable, while coverage of a wider operational spectrum requires the receiver to retune or scan across frequency. This architecture enables strong analogue selectivity and mature protection against out-of-band interferers, but introduces trade-offs between instantaneous bandwidth, scan revisit time, tuning latency, and probability of intercept for short-duration or frequency-agile emissions.

In this work, the heterodyne architecture is therefore abstracted as a receiver with:

- a tuneable centre frequency;
- a finite instantaneous bandwidth;
- sequential frequency coverage across a wider operational band; and
- detection and downstream processing limited to the active tuned band at each instant.

This abstraction captures the architectural behaviour relevant to wide-area spectral surveillance without attempting to model implementation-specific mixer responses, filter characteristics, local-oscillator impairments, or proprietary signal-processing techniques. From an architectural standpoint, heterodyne receivers also exhibit well-known nonidealities that can influence performance in congested environments. Local-oscillator phase noise, mixer-generated spurious responses, and tuning transients may degrade detection performance, particularly for low-SNR or closely spaced signals [6], [8]. These effects are not modelled explicitly at the circuit level in this work; instead, the simulation captures their first-order architectural consequences through finite instantaneous bandwidth, sequential spectral access, and tuning-dependent latency




Fig. 2. Conceptual comparison of heterodyne and Direct RF receiver architectures.

B. Direct RF Receiver Model

Direct RF receiver architectures move the analogue-to-digital boundary closer to the antenna by digitizing RF signals directly at, or near, their received carrier frequencies. This development follows the broader software-radio concept of transferring increasing amounts of receiver functionality from fixed analogue hardware into programmable digital processing [4]. Continued improvements in high-speed data-converter performance [5], together with recent RF-sampling and integrated digital-channelization technologies [8], now enable multi-gigahertz instantaneous receiver bandwidths.




Fig. 3 Receiver example of a single-IF vs. a Direct RF mixing system.

At the architectural level, Direct RF receivers are distinguished by persistent wideband visibility rather than sequential spectral access. Within the limits of the RF front end, converter sampling bandwidth, and available digital-processing resources, multiple signals within the instantaneous capture band may be observed concurrently. Frequency translation, filtering, channel selection, and reconfiguration can then be performed digitally without conventional analogue retuning.

Direct RF does not, however, eliminate the analogue RF front end. Gain control, preselection, impedance matching, linearity management, and protection against strong signals remain essential. Furthermore, high-speed RF sampling introduces constraints associated with converter dynamic range, sampling-clock quality, spurious performance, data throughput, and real-time processing capacity [5], [8]. The principal architectural distinction is therefore not the complete removal of analogue circuitry, but the relocation of frequency agility and channelization across the analogue-digital boundary.

In this work, the Direct RF architecture is abstracted as a receiver with:

- a wide instantaneous RF bandwidth;
- concurrent observation of multiple sub-bands within that bandwidth; and
- digital channelization and detection limited primarily by available processing capacity rather than tuner dwell time.

This abstraction reflects the system-level consequences of RF digitization while remaining independent of a particular converter family, FPGA, or processing implementation.

C. Architectural Scope and Modeling Assumptions

To enable a fair and interpretable comparison, both architectures are evaluated under a common set of high-level assumptions. The analysis is intentionally system-level and comparative, focusing on architectural trends rather than absolute hardware sensitivity or implementation-specific performance.

The following assumptions apply throughout the study.

Architecture-level abstraction: The models capture first-order behaviours such as instantaneous bandwidth, tuning modality, spectral concurrency, analogue selectivity, and finite digital-processing capacity. Detailed circuit nonidealities, proprietary algorithms, and platform-specific implementations are not explicitly reproduced.

Platform-agnostic evaluation: No specific airborne, maritime, land, or space platform is assumed. Size, weight, power, and cost implications are therefore treated qualitatively rather than as fixed design constraints.

Representative congested and contested environments: Emitter density, signal agility, interference, and spectral occupancy are selected to provide representative comparative scenarios without implying classified threat models or operational parameters.

Comparative interpretation: The results are intended to identify architectural trends and trade-offs rather than best-case or worst-case achievable receiver performance. Both architectures are therefore assessed using common environmental and detection assumptions wherever possible.

Under these assumptions, the heterodyne and Direct RF models provide a controlled basis for examining how receiver architecture influences observation, processing, and response in dense and dynamic electromagnetic environments. The simulation framework used to quantify these effects is described in Section III.

III. SIMULATION FRAMEWORK AND PERFORMANCE METRICS

This section describes the Monte Carlo framework used to compare heterodyne and Direct RF receiver architectures under representative congested and contested electromagnetic-spectrum conditions. The analysis is intentionally architecture-level rather than hardware-specific. Its purpose is to isolate the consequences of spectral access, tuning modality, analogue selectivity, and finite digital-processing capacity under controlled and repeatable environmental conditions.

A key feature of the framework is that, within each Monte Carlo trial, the electromagnetic environment is generated once and subsequently presented to both receiver architectures. The two receivers therefore operate on identical emitter activity, carrier frequencies, signal bandwidths, received signal-to-noise ratios, and temporal events. This common-truth approach prevents differences between random emitter realizations from being interpreted as architectural performance differences.

A. Operational Scenario Abstraction

The simulated environment comprises multiple independent emitters distributed across a finite surveillance bandwidth. Depending on the experiment, emitters may be frequency-agile or fixed-frequency burst sources. Each emitter is characterized by an occupied bandwidth, received SNR, carrier frequency, and temporal activity pattern.

For congestion, compute-scaling, spectral-exploitation, and deinterleaving experiments, emissions occur as alternating finite ON and OFF intervals. Signal bandwidths and received SNRs are drawn independently from prescribed ranges, while emitter frequencies are distributed across the applicable mission bandwidth. Frequency-hopping behaviour is treated separately in the probability-of-intercept experiment so that hop dwell time can be controlled explicitly.

The models do not represent a particular threat library or classified emitter population. Instead, the parameter ranges are selected to provide controlled sparse-to-dense environments representative of electronic-support and wideband-surveillance problems considered in the open literature [28]–[31].

B. Receiver Architecture Models

The heterodyne receiver is represented as a sequentially scanning architecture with finite instantaneous bandwidth. The receiver dwells within one spectral sub-band before entering a finite retuning interval during which it is assumed to be unavailable for detection. It then advances to the next sub-band and repeats the scan across the surveillance bandwidth. A randomized initial scan phase is used for each Monte Carlo realization to prevent favourable synchronization between emitter activity and receiver tuning.

The Direct RF receiver is represented as an instantaneous-wideband architecture capable of observing a contiguous spectral region without analogue retuning. Signals lying within this capture bandwidth are immediately available to the digital receiver, subject to finite channelisation and processing resources. This abstraction is consistent with the

system-level Direct RF and concurrent multiband receiver concepts discussed in [36]–[38].

For experiments involving congestion, both architectures are subjected to a finite downstream processing capacity. When the number of simultaneously observable signals exceeds this capacity, only a limited number can be admitted for processing. This prevents wideband visibility from being treated as equivalent to unlimited digital-processing capability.

Representative architecture-level parameters are summarized in Table II.

TABLE II.

Receiver Architecture Parameters

Parameter	Heterodyne	Direct RF
Reference instantaneous bandwidth	200 MHz	2 GHz
Additional POI bandwidths	100 MHz	1, 2, 4 GHz
Dwell time	100 μ s	—
Retune time	100 μ s	—
Channelisation/update latency	—	10 μ s
Common acquisition time	20 μ s	20 μ s
Shared downstream processing limit, K	16 signals	16 signals
Direct-RF compute sweep	—	K = 2, 4, 8, 16, 32, 64
Classification latency	50 μ s	50 μ s
Decision latency	20 μ s	20 μ s
Action latency	50 μ s	50 μ s

C. Monte Carlo Procedure and Statistical Treatment

System behaviour is evaluated using a discrete-time Monte Carlo simulation with a nominal temporal resolution of 5 μ s. Unless otherwise stated, each reported operating point is estimated from 500 independent Monte Carlo trials. A fixed pseudorandom-number seed is used to permit reproducibility.

Within each trial, the environmental truth realization is generated before either receiver is evaluated. Both architectures subsequently operate on the same realization, ensuring that differences in the resulting metrics are attributable to receiver behaviour rather than to differing stochastic signal scenarios.

A signal is not declared detected merely because it overlaps the receiver bandwidth at a single time sample. Detection requires the signal to remain both observable and admitted to processing for a contiguous acquisition interval of 20 μ s. This prevents arbitrarily short spectral coincidences from being counted as successful interceptions.

Monte Carlo means are reported together with nominal 95% confidence intervals derived from the standard error of the trial ensemble. Increasing the number of trials therefore reduces statistical uncertainty without altering the underlying receiver models.

Representative environmental and simulation parameters are summarized in Table III.

TABLE III.

Environmental and Simulation Parameters

Parameter	Value
Total surveillance bandwidth	6 GHz
Time resolution	5 μ s
Monte Carlo trials	500
Emitter occupied bandwidth	1–20 MHz
Received SNR range	–10 to 20 dB
Detection threshold	–10 dB
Required acquisition time	20 μ s
Burst ON duration	40–160 μ s
Burst OFF duration	40–300 μ s
Emitter-count sweep	5, 10, 25, 50, 100
Frequency-hop dwell sweep	50–500 μ s
Congestion mission bandwidth	2 GHz

D. Detection, Processing, and Blocking Models

For the general detection experiments, an emission is eligible for detection when its received SNR exceeds the nominal detection threshold and its occupied bandwidth

overlaps the instantaneous receiver coverage. Successful detection additionally requires continuous visibility for the prescribed acquisition interval.

When more signals are simultaneously visible than the available processing capacity K , the model admits at most K signals for downstream processing. In the present implementation, candidate signals are prioritized according to received SNR. This provides a simple and reproducible abstraction of finite real-time processing resources without attempting to reproduce a particular scheduler or classifier. The importance of finite digital-processing capacity in wideband receiver architectures is consistent with the concurrent RF-sampling and channelisation systems considered in [36]–[38].

Blocking is examined separately from the general detection experiments. The wanted signal is placed within the tuned heterodyne passband, while a strong interfering signal is displaced by 400 MHz. The blocker therefore lies outside the reference 200-MHz heterodyne channel but remains inside the wider 2-GHz Direct RF capture bandwidth. Both receiver architectures are subjected to the same overload law; the architectural distinction arises from front-end exposure rather than from assigning intrinsically different desensitization laws to the two receiver types.

A representative 40-dB analogue rejection is applied to the out-of-channel blocker in the heterodyne case, while no equivalent rejection is applied when the blocker lies within the Direct RF capture bandwidth. The experiment is therefore intended to isolate the architectural effect of analogue preselection rather than reproduce a particular RF filter or ADC implementation. Strong-signal interference and blocking remain fundamental considerations in electronic-warfare receiver design [9].

E. Performance Metrics

Receiver behaviour is quantified using seven system-level metrics related to spectral access, response time, congestion, processing capacity, interference exposure, spectral exploitation, and emitter association.

Probability of Intercept: Probability of intercept has long been used in electronic-support and wideband-receiver analysis to characterize the likelihood of observing agile or intermittent emitters [28]–[31]. In the present frequency-hopping experiment, POI is evaluated on a per-hop basis. A hop is successfully intercepted when the signal lies within the receiver coverage and remains observable for at least the required acquisition interval. The metric is therefore defined as

$$P_{OI} = \frac{N_{\text{intercepted hops}}}{N_{\text{total hops}}}. \quad (1)$$

This definition directly captures the interaction between hop dwell time, receiver instantaneous bandwidth, and scan dynamics.

Sensing-to-Response Latency: Response latency is measured from emitter activation to completion of the modeled receiver-response chain. Cognitive radar and EW literature emphasize the importance of rapid sensing and decision cycles in dynamically changing environments [32], [33]. The modeled response time is

$$T_{\text{resp}} = T_{\text{detect}} + T_{\text{channel}} + T_{\text{classify}} + T_{\text{decide}} + T_{\text{act}} \quad (2)$$

where T_{channel} represents the additional Direct RF channelisation/update latency. Both the conditional response-time distribution for successful detections and the overall probability of response are retained.

Detection Performance Under Congestion: The ability to maintain detection performance as spectral occupancy and emitter population increase is relevant to spectrum-sensing and software-defined receiver operation [34], [35]. Detection probability is evaluated on a per-burst basis as

$$P_{\text{det}} = \frac{N_{\text{detected bursts}}}{N_{\text{eligible bursts}}} \quad (3)$$

This formulation ensures that finite signal duration and competition for receiver resources directly influence the measured performance.

Compute-Limited Scalability: Wide instantaneous bandwidth transfers a substantial portion of receiver complexity from analogue tuning to digital channelisation and processing [36]–[38]. Compute-limited scalability is therefore evaluated by varying the maximum number K of signals that can be processed concurrently. For a given processing capacity,

$$P_{\text{det}}(K) = \frac{N_{\text{bursts detected with capacity } K}}{N_{\text{eligible bursts}}} \quad (4)$$

This metric quantifies how effectively the potential advantage of wideband observation can be exploited as real-time processing resources are increased or saturated.

Blocking Robustness: Blocking robustness describes the ability to retain detection of a signal of interest in the presence of a much stronger interferer. It is expressed as the victim detection probability conditioned on interferer-to-signal ratio,

$$P_B(\text{ISR}) = P(\text{victim detected} \mid \text{ISR}) \quad (5)$$

The same overload response is applied to both receiver types, while architectural differences arise from whether analogue preselection removes the blocker before it reaches the vulnerable wideband receiver stages [9].

Spectral Exploitation Efficiency: To distinguish merely observing a wide bandwidth from successfully exploiting the signals contained within it, spectral exploitation efficiency is defined as the fraction of active information-bearing spectrum that is both visible to the receiver and admitted to processing:

$$\eta_{\text{SE}} = \frac{B_{\text{active, visible, processed}}}{B_{\text{active, total}}} \quad (6)$$

Overlapping occupied bandwidths are evaluated using their spectral union so that coincident emissions are not double counted. Unlike a conventional receiver-bandwidth occupancy ratio, therefore η_{SE} measures how much of the active electromagnetic environment the receiver can use at a given instant.

Multi-Emitter Deinterleaving Performance: Multi-target and emitter-association problems are commonly evaluated in terms of track purity, association completeness, and related multi-target tracking measures [39], [40]. In this work, deinterleaving performance is represented using an F1-type score obtained from weighted track purity P and mean emitter completeness C

$$F_1 = \frac{2PC}{P+C} \quad (7)$$

Unlike the previous implementation, completeness is evaluated against the full population of detectable emitters. An emitter that is never observed therefore contributes zero completeness rather than being omitted from the scoring population.

Each metric is evaluated using common environmental realizations wherever applicable so that the resulting differences reflect architectural behaviour rather than scenario bias. The purpose of the framework is not to establish absolute hardware-performance limits, but to quantify how heterodyne and Direct RF architectures redistribute constraints among spectral access, analogue selectivity, processing capacity, and interference exposure.

IV. PROBABILITY OF INTERCEPT FOR AGILE EMITTERS

Probability of intercept (POI) is a fundamental measure of receiver effectiveness against frequency-agile and intermittent emitters. In electronic-support and wideband-surveillance applications, an emitter may occupy a particular frequency for only a short interval before hopping elsewhere in the spectrum, making successful interception dependent on both temporal and spectral alignment between the emitter and receiver [28]–[31].

The present experiment evaluates the architectural effect of this alignment by varying the emitter hop dwell time from 50 to 500 μs . Heterodyne receiver bandwidths of 100 and 200 MHz are compared with Direct RF instantaneous bandwidths of 1, 2, and 4 GHz over the common 6-GHz surveillance band. A successful interception requires the signal to remain visible and available for processing for the acquisition interval defined in Section III. POI is evaluated on a per-hop basis according to (1), rather than as eventual detection of an emitter over the complete observation interval.

A. Architectural Dependence of POI

For the heterodyne receiver, interception requires coincidence between the emitter frequency and the receiver's currently tuned passband. The receiver therefore observes only a small fraction of the surveillance spectrum at any instant, and successful detection depends on the relationship between emitter dwell time, receiver dwell time, retuning interval, and scan position.

As the emitter hop dwell time increases, the probability that the scanning receiver encounters the emitter before its next frequency transition also increases. Wider heterodyne instantaneous bandwidth improves this probability by increasing the fraction of spectrum observed during each receiver dwell. This scan-dependent behaviour is consistent with the established limitations of narrowband electronic-support receivers operating against frequency-agile emitters [28]–[31].

Direct RF reception produces a different operating regime. Because a contiguous fraction of the surveillance bandwidth is observed persistently, interception does not require temporal coincidence with a scanning tuner. Instead, POI is governed primarily by whether the current emitter frequency lies within the instantaneous Direct RF capture bandwidth and whether sufficient processing resources are available.

B. Simulation Results

Fig. 4 shows the per-hop probability of intercept as a function of frequency-hop dwell time.




Fig. 4. Probability of intercept for a frequency-hopping emitter as a function of hop dwell time for heterodyne and Direct RF receiver architectures. Error bars indicate nominal 95% Monte Carlo confidence intervals.

The results show a clear separation between the two architectural approaches. For the 100-MHz heterodyne receiver, POI remains low over the complete dwell-time range, increasing from approximately 0.01 at the shortest dwell time to approximately 0.04–0.05 at 500 μs . Increasing the heterodyne bandwidth to 200 MHz approximately doubles the instantaneous spectral opportunity, with POI increasing to around 0.08 at the longest dwell times.

The Direct RF receivers achieve substantially greater POI because a much larger fraction of the surveillance band is continuously visible. Across most of the dwell-time range, the 1-GHz Direct RF receiver achieves a POI of approximately 0.15–0.17, increasing to approximately 0.31–0.34 for a 2-GHz capture bandwidth and approximately 0.60–0.68 for a 4-GHz capture bandwidth.

The approximate scaling of POI with instantaneous Direct RF bandwidth is physically consistent with the 6-GHz surveillance environment. Instantaneous bandwidths of 1, 2, and 4 GHz correspond to visibility of approximately one-sixth, one-third, and two-thirds of the total surveillance spectrum, respectively. The simulated intercept probabilities broadly follow these spectral-coverage fractions, demonstrating that persistent instantaneous coverage, rather than tuner revisit time, dominates the Direct RF result.

C. Effect of Hop Dwell Time

The heterodyne results exhibit an overall increase in POI with increasing hop dwell time. This behaviour follows directly from sequential spectral access: a longer-lived hop provides a larger temporal opportunity for the scanning receiver to enter the occupied frequency interval before the emitter moves elsewhere.

In contrast, the Direct RF curves are comparatively insensitive to hop dwell time once the emission duration exceeds the required acquisition interval. This is an important architectural distinction. Persistent wideband observation substantially decouples interception from receiver scan timing; increasing the emitter dwell time provides relatively little additional benefit when the occupied frequency is already continuously visible. Some point-to-point variation remains in the Direct RF results, particularly at the longer hop dwell times. Because the observation interval is finite, longer dwell times produce fewer hop events per Monte Carlo realization and therefore reduce the number of independent frequency opportunities represented within each trial. This finite-window effect should be distinguished from a fundamental dependence of Direct RF interception on hop dwell time.

D. Discussion

The results demonstrate that probability of intercept against a frequency-agile emitter is governed by two distinct mechanisms in the receiver architectures considered here. In a scanning heterodyne receiver, POI depends jointly on instantaneous bandwidth and temporal alignment between the receiver scan and emitter activity. Short dwell times are therefore particularly challenging because the emitter may leave a frequency before the receiver reaches it. For Direct RF reception, the dominant constraint instead becomes instantaneous spectral coverage. Provided that the emitter lies within the capture bandwidth and remains present for the required acquisition interval, no analogue retuning is required before detection can occur. Consequently, widening the Direct RF capture bandwidth produces an approximately proportional increase in the fraction of agile emissions available for interception. The result does not imply that arbitrarily wide Direct RF reception is without cost. Increasing instantaneous bandwidth exposes the receiver front end and data converter to a larger electromagnetic environment and increases the amount of data requiring channelisation and processing. These limitations become important in the congestion, compute-scaling, and blocking experiments considered in subsequent sections.

V. SENSING-TO-RESPONSE LATENCY

Sensing-to-response latency is an important consideration in electronic-support, cognitive sensing, and electromagnetic operations in which the value of an intercepted signal may diminish rapidly with time [32], [33]. Receiver architecture contributes directly to this latency because a signal cannot enter downstream classification and decision processes until it has first been made available by the spectral-access mechanism.

The latency experiment therefore evaluates the elapsed time between emitter activation and completion of the modelled response chain defined in (2). Both receiver architectures are presented with the same emitter realization. For the heterodyne receiver, detection is dependent on the instantaneous position of the scanning passband, whereas the Direct RF receiver maintains continuous access to signals falling within its capture bandwidth. The Direct RF model additionally includes the finite channelisation/update latency specified in Table II.

A. Architectural Sources of Latency

In the scanning heterodyne architecture, a detectable emitter may become active while the receiver is tuned elsewhere in the surveillance spectrum. Detection is therefore delayed until the scanning passband reaches the emitter frequency and the required acquisition interval has been satisfied. Because the initial scan phase is randomized in each Monte Carlo realization, this produces a distribution of detection delays rather than a fixed response time.

Once detection occurs, the common classification, decision, and action latencies are added to form the complete sensing-to-response interval. The resulting heterodyne response time can therefore be expressed conceptually as the sum of scan-dependent acquisition delay and the downstream response chain.

The Direct RF architecture removes the scan-alignment component for signals lying inside its instantaneous capture bandwidth. Following the required acquisition interval, the signal is passed through the modelled $10\mu\text{s}$ channelisation/update stage before entering the same classification, decision, and action processes used for the heterodyne receiver. Its response latency is consequently dominated by the fixed digital-processing chain rather than by spectral revisit time.

B. Simulation Results

Fig. 5 shows the cumulative distribution function (CDF) of sensing-to-response latency for the two architectures. Both receivers achieved a 100% response-success rate in this experiment, allowing the latency distributions to be compared without conditioning on different proportions of successful detections.




Fig. 5. Cumulative distribution of sensing-to-response latency for heterodyne and Direct RF receiver architectures. Both architectures achieved 100% response success for the simulated scenario.

The Direct RF result is concentrated at approximately 0.15 ms. This value is consistent with the modelled response chain: a $20\mu\text{s}$ acquisition interval, $10\mu\text{s}$ channelisation/update latency, $50\mu\text{s}$ classification interval, $20\mu\text{s}$ decision interval, and $50\mu\text{s}$ action interval give

$$20 + 10 + 50 + 20 + 50 = 150 \mu\text{s}$$

The narrow Direct RF distribution therefore reflects the absence of scan-dependent spectral-access delay in the modeled wideband receiver.

The heterodyne architecture exhibits a much broader latency distribution. Response times extend from values close to the minimum processing-chain delay to approximately 2 ms, with a median latency of roughly 1 ms. The distribution arises from the random relationship between emitter activation and the instantaneous receiver scan position. Emitters appearing shortly before the receiver reaches their frequency are detected relatively quickly, whereas those appearing after the relevant passband has been visited must wait for a subsequent scan opportunity.

This difference is visible across the complete CDF rather than only in the mean response time. The Direct RF architecture provides an almost deterministic response interval in the present scenario, while the heterodyne receiver exhibits stochastic latency associated with sequential spectral access.

C. Discussion

The latency result illustrates an architectural distinction that is separate from raw detection sensitivity. Both receiver types ultimately detect and respond to all emitters in the simulated experiment; the difference lies in when the signal becomes available to the downstream processing chain.

For the heterodyne receiver, sensing latency is coupled to the scan schedule. Increasing scanning speed can reduce this delay, but faster scanning generally reduces dwell time and therefore introduces a competing sensitivity and acquisition-time requirement. Sequential spectral access consequently creates an inherent trade-off between surveillance coverage, dwell time, and response latency. Direct RF reception largely removes this scan-dependent component for signals contained within the instantaneous capture bandwidth. Its latency is instead governed by acquisition, channelisation, and subsequent digital processing. In the present model, this produces a response time of approximately $150\mu\text{s}$, compared with a broad heterodyne distribution extending to the millisecond scale. This result should not be interpreted as implying that Direct RF processing is intrinsically instantaneous. Practical wideband receivers incur finite buffering, channelisation, detection, classification, and data-movement delays, which may increase with algorithm complexity and processing load. Rather, the architectural benefit is that these delays are not compounded by a sequential analogue tuning process.

The result therefore identifies a potentially important advantage of persistent wideband reception for time-sensitive applications: spectral access and downstream processing can begin without first waiting for a scanning receiver to revisit the occupied frequency. The following sections examine the corresponding cost of this persistent visibility as emitter density and digital-processing demand increase.

VI. DETECTION PERFORMANCE UNDER SPECTRAL CONGESTION

Congested electromagnetic environments place a different stress on receiver architectures from the frequency-agility and latency effects considered in Sections IV and V. As the number of simultaneously active emitters increases, receiver performance becomes dependent not only on spectral visibility, but also on the ability of downstream processing resources to channelize, detect, and maintain multiple signals concurrently.

The present experiment evaluates burst-detection probability as the emitter population is increased from 5 to 100. Detection is evaluated on a per-burst basis according to (3), so that each finite-duration emission represents an independent detection opportunity. Both architectures operate on identical environmental realizations and are subject to the common downstream processing limit $K = 16$.

A. Detection Probability Versus Emmitter Density

Fig. 6 shows burst-detection probability as a function of the number of emitters in the simulated environment.




Fig. 6. Burst-detection probability as a function of emitter population for heterodyne and Direct RF receiver architectures with a shared processing limit of $K = 16$. Error bars indicate nominal 95% Monte Carlo confidence intervals.

At low emitter densities, the Direct RF architecture achieves a burst-detection probability close to unity. Detection remains approximately 0.99 for populations up to $N = 25$, indicating that the combination of wide instantaneous visibility and the available processing capacity is sufficient to accommodate nearly all eligible emissions in this regime.

As emitter density increases, however, Direct RF performance begins to decline. At $N = 50$, the burst-detection probability falls to approximately 0.90, while at $N = 100$ it decreases to approximately 0.48. Because the receiver capture bandwidth is unchanged, this reduction does not arise from a loss of spectral visibility. Instead, increasing numbers of simultaneously active signals compete for the finite $K = 16$ downstream processing resource.

The heterodyne receiver exhibits a markedly different behaviour. Its burst-detection probability remains approximately 0.08–0.09 across the complete emitter-density range. This apparent insensitivity to increasing emitter population should not be interpreted as superior congestion robustness. Rather, the receiver is already strongly constrained by sequential spectral access: only signals falling within the currently tuned passband can become available for detection. Consequently, increasing the total number of emitters does not substantially alter the fraction of individual bursts intercepted by the scanning receiver.

The two architectures therefore occupy different operating regimes. The heterodyne receiver is primarily spectral-access limited, whereas the Direct RF receiver transitions from a spectral-visibility advantage at low density to a processing-capacity limitation as the instantaneous signal population increases.

B. Compute-Limited Scalability

To isolate this processing constraint, Fig. 7 evaluates Direct RF burst-detection probability at a fixed emitter population of $N = 75$ while varying the maximum number K of signals that may be processed concurrently. The compute-limited metric is defined by (4).




Fig. 7. Direct RF burst-detection probability as a function of maximum concurrent processing capacity K for $N = 75$ emitters. The dashed line indicates the heterodyne reference performance for the same environment. Error bars indicate nominal 95% Monte Carlo confidence intervals.

The dependence on processing capacity is pronounced. With only $K = 2$ concurrent processing channels, Direct RF detection probability is approximately 0.10, comparable with the heterodyne reference. Increasing the available capacity to $K = 4$ raises the probability to approximately 0.18, while $K = 8$ increases it to approximately 0.33.

A further increase to $K = 16$ produces a burst-detection probability of approximately 0.64. This is consistent with the degradation observed in Fig. 6 at high emitter densities and confirms that the reduction is associated with processing contention rather than a loss of wideband spectral access.

The most significant change occurs as processing capacity approaches the instantaneous signal demand. At $K = 32$, detection probability rises to approximately 0.99 and remains close to unity at $K = 64$. The plateau indicates that, for the simulated $N = 75$ environment, further increases in processing capacity provide little benefit once the number of concurrently active signals that require processing can be accommodated.

The dashed heterodyne reference remains near 0.09 because increasing Direct RF digital capacity does not alter the heterodyne architecture's sequential spectral-access constraint.

C. Architectural Interpretation

Taken together, Figs. 6 and 7 illustrate one of the central trade-offs of Direct RF reception. Wide instantaneous bandwidth provides access to a substantially larger fraction

of the electromagnetic environment, but this spectral visibility does not by itself guarantee that all observed signals can be exploited. As emitter density increases, the system bottleneck moves downstream from spectral access to digital processing capacity.

This distinction is important because it separates receiver bandwidth from usable receiver capacity. An architecture may digitize several gigahertz of spectrum while still being unable to channelize, classify, or otherwise process every signal present within that bandwidth in real time. The practical value of increasing instantaneous bandwidth is therefore coupled to the availability of corresponding FPGA, GPU, ASIC, or other real-time processing resources [36]–[38].

The compute sweep also demonstrates that this limitation is scalable rather than intrinsic. The degradation observed at $K = 16$ is progressively recovered as additional processing capacity is introduced, with near-complete burst detection achieved by approximately $K = 32$ in the present scenario. Direct RF therefore exchanges the fixed sequential-access limitation of the scanning heterodyne receiver for a resource-dependent digital-processing constraint that can, in principle, be addressed through system design.

This result also provides an important qualification to claims of wideband receiver superiority. Persistent spectral visibility offers a substantial advantage only when sufficient processing capacity exists to exploit the captured information. Conversely, the heterodyne architecture limits the instantaneous processing burden through analogue spectral selection, but at the cost of substantially reduced access to transient emissions distributed across the wider surveillance band.

The following section examines the complementary contested-environment trade-off: the same wideband exposure that improves spectral awareness can increase susceptibility to strong out-of-band or adjacent interferers before digital mitigation becomes available.

VII. ROBUSTNESS IN CONTESTED ENVIRONMENTS

Operation in contested electromagnetic environments places additional demands on receiver architectures beyond probability of intercept, response latency, and spectral congestion. In particular, strong interfering signals can expose limitations in front-end linearity and dynamic range before digital signal processing becomes available. At the same time, the ability to observe and exploit a large fraction of the active spectrum remains important when signals of interest are distributed across a wide operating band.

This section therefore considers two complementary aspects of architectural robustness: detection in the presence of a strong spectrally displaced blocker, and the

fraction of active signal bandwidth that can be simultaneously observed and processed.

A. Blocking and Interference Robustness

Blocking robustness is evaluated using the metric defined in (5). The wanted signal lies within the tuned heterodyne passband, while the interfering signal is displaced by 400 MHz. The blocker consequently falls outside the reference 200-MHz heterodyne channel but remains inside the 2-GHz Direct RF capture bandwidth.

The same overload relationship is used for both architectures. The distinction therefore arises from where the blocker is rejected. In the heterodyne case, representative analogue preselection provides 40 dB of attenuation before the blocker reaches the vulnerable downstream receiver stages. In the Direct RF case, the same blocker remains inside the wide instantaneous capture band and is therefore presented to the wideband front end and digitizer. This experiment is intended to isolate the architectural consequence of analogue selectivity rather than reproduce the response of a particular filter or ADC [9].

Fig. 8 shows victim-signal detection probability as a function of blocker-to-signal ratio.




Fig. 8. Victim detection probability as a function of blocker-to-signal ratio for a blocker offset by 400 MHz. The heterodyne receiver incorporates narrow analogue preselection, while the blocker remains within the Direct RF wideband capture region. Error bars indicate nominal 95% Monte Carlo confidence intervals.

At low and moderate blocker levels, the two receiver architectures exhibit similar detection probabilities of approximately 0.83–0.85. This indicates that, while the blocker remains below the effective overload region, the wider Direct RF capture bandwidth does not impose a significant detection penalty.

A clear architectural separation emerges as the blocker level increases. The Direct RF detection probability begins to fall rapidly at blocker-to-signal ratios above approximately 30–35 dB, falling to around 0.4 at 40 dB and approaching zero by approximately 50 dB. Because the

blocker lies within the Direct RF capture bandwidth, its full strength reaches the wideband receiver chain and progressively desensitizes the modelled victim channel.

The heterodyne receiver behaves differently. Analogue rejection of the spectrally displaced blocker substantially delays the onset of degradation, with victim detection remaining close to its nominal level across most of the simulated ISR range. Significant degradation occurs only at the highest blocker levels, when the residual interference remaining after preselection becomes sufficiently large to reach the modelled overload region.

The result therefore demonstrates an important consequence of wide instantaneous reception: signals that are digitally separable after conversion may still compete for finite analogue dynamic range before conversion. Digital filtering cannot recover information that has already been lost through front-end compression, desensitization, or converter overload.

B. Spectral Exploitation Efficiency

The complementary benefit of wide instantaneous reception is illustrated using the spectral exploitation efficiency defined in (6). Rather than measuring the fraction of receiver bandwidth that happens to be occupied, this metric measures the fraction of the active information-bearing spectrum in the environment that is both visible to the receiver and admitted to downstream processing.

Fig. 9 shows mean spectral exploitation efficiency as a function of emitter population for the shared processing limit $K = 16$.

At low emitter densities, the Direct RF architecture exploits almost all of the active signal bandwidth present within the simulated mission region. Spectral exploitation efficiency remains close to unity for populations up to approximately $N = 25$, indicating that the available processing resource is sufficient to accommodate the active signals visible within the wideband capture region.




Fig. 9. Mean spectral exploitation efficiency as a function of emitter population for heterodyne and Direct RF receiver architectures with $K = 16$. Error bars indicate nominal 95% Monte Carlo confidence intervals.

As emitter density increases, exploitation efficiency declines. At $N = 50$, approximately 0.87 of the active signal bandwidth remains visible and processed, while at $N = 100$ the value falls to approximately 0.48. This trend is consistent with the congestion results of Section VI: the receiver continues to observe the spectrum, but finite processing capacity prevents every simultaneously active signal from being exploited.

The heterodyne receiver exhibits an exploitation efficiency of approximately 0.05 across the simulated density range. This low value is primarily a consequence of sequential spectral access rather than insufficient downstream processing. At any given instant, only the narrow tuned region contributes to the active bandwidth that can be observed and processed, leaving much of the wider electromagnetic environment unavailable until the receiver subsequently scans across it.

The similarity between the Direct RF trends in Figs. 6 and 9 is expected but the two metrics describe different quantities. Fig. 6 measures the probability that individual signal bursts are successfully detected, whereas Fig. 9 is bandwidth-weighted and measures the fraction of the active spectral environment being exploited at a given time. Together, they show that increasing spectral density stresses both event-level detection and overall information capture.

C. Architectural Trade-Off

Figs. 8 and 9 expose a central tension in wideband receiver design. Increasing instantaneous bandwidth improves spectral awareness because a larger proportion of the active electromagnetic environment is continuously available for detection and processing. The same wider analogue aperture, however, exposes the receiver front end and data converter to a correspondingly larger aggregate signal environment.

The heterodyne architecture resolves this trade-off in favour of analogue selectivity. Signals outside the currently selected channel can be strongly attenuated before they reach sensitive downstream stages, improving tolerance to spectrally displaced high-power interferers. The cost is that most of the wider surveillance spectrum is not observed simultaneously.

Direct RF adopts the opposite emphasis. Wideband capture maximizes instantaneous spectral access and enables subsequent digital channelisation, but analogue components ahead of the ADC must tolerate the aggregate signal environment. Once the front end or converter is

driven into an overload regime, downstream digital filtering cannot restore the lost signal fidelity.

These results therefore support the broader interpretation that neither architecture removes receiver constraints; rather, each places them at a different point in the signal chain. Heterodyne reception accepts a spectral-access limitation in exchange for strong analogue protection, whereas Direct RF accepts greater front-end exposure in exchange for persistent wideband visibility.

This trade-off provides a strong motivation for hybrid receiver architectures in which selective or adaptive analogue protection is placed ahead of wideband digitization. Such an approach can potentially preserve much of the spectral-awareness advantage of Direct RF while preventing the strongest unwanted signals from consuming front-end and converter dynamic range.

VIII. MULTI-EMITTER DEINTERLEAVING PERFORMANCE

In congested electromagnetic environments, successful detection represents only the first stage of receiver exploitation. Individual observations must subsequently be associated with their originating emitters to support tracking, classification, geolocation, and situational awareness. Multi-emitter deinterleaving therefore depends not only on the quality of the association algorithm, but also on the continuity and completeness of the observations made available by the receiver architecture [39], [40].

The present experiment evaluates this architectural contribution using the population-aware F1 score defined in (7). The metric combines weighted track purity with mean emitter completeness. Importantly, completeness is evaluated against the full detectable emitter population, so an emitter that is never observed contributes zero completeness rather than being excluded from the analysis.

A. Architectural Influence on Deinterleaving

A scanning heterodyne receiver observes only the currently selected portion of spectrum. Emissions occurring outside the instantaneous tuned passband are therefore unavailable to the association process, and repeated observations of the same emitter may be separated by receiver scan intervals. This fragmented spectral access can reduce emitter completeness even when the observations that are obtained are correctly associated.

Direct RF reception provides a different input to the deinterleaving process. Multiple emitters falling within the instantaneous capture bandwidth can be observed concurrently, providing a more complete set of detections from which emitter-level tracks may be constructed. This does not eliminate the association problem, but it increases

the amount and temporal continuity of information available to the downstream process.

At high emitter densities, however, the same finite processing capacity examined in Section VI becomes important. When the number of simultaneously observable signals exceeds the available processing resource, some emissions are not admitted to downstream processing, reducing completeness and consequently degrading the population-aware F1 score.

B. Simulation Results

Fig. 10 shows deinterleaving performance as a function of emitter population.




Fig. 10. Population-aware deinterleaving F1 score as a function of emitter population for heterodyne and Direct RF receiver architectures. Error bars indicate nominal 95% Monte Carlo confidence intervals.

The heterodyne receiver produces an F1 score of approximately 0.11–0.12 across the simulated emitter-density range. The relatively low and nearly constant value reflects the limited fraction of the overall emitter population that is available to the association process through sequential spectral observation.

The Direct RF architecture provides substantially higher deinterleaving scores at low and moderate emitter densities. The F1 score is approximately 0.47 at $N = 5$, rises slightly to approximately 0.49 at $N = 25$, and remains around 0.44 at $N = 50$. This improvement is consistent with the greater completeness provided by concurrent wideband observation.

At the highest simulated density, $N = 100$, the Direct RF F1 score falls to approximately 0.27. This reduction corresponds to the same high-density regime in which Sections VI and VII showed increasing contention for the shared $K = 16$ processing resource. Although the receiver retains wideband spectral visibility, not every observable emission can be admitted to processing, thereby reducing the completeness of the emitter histories available for association.

C. Interpretation and Limitations

The results indicate that receiver architecture can influence deinterleaving performance before any sophisticated association algorithm is applied. Persistent wideband reception supplies a more complete and concurrent observation set, whereas sequential scanning inherently fragments the spectral record presented to downstream emitter-association processing.

The absolute F1 values should not, however, be interpreted as limits on achievable deinterleaving performance. The simulation intentionally employs a simplified association model and does not incorporate the full range of discriminating features available to practical electronic-support systems, such as pulse-repetition characteristics, angle of arrival, modulation features, waveform fingerprints, or learned emitter representations. The purpose of the experiment is therefore to isolate the architectural contribution to observation completeness, rather than to demonstrate an optimized deinterleaving algorithm.

The decline in Direct RF performance at high emitter density also reinforces the compute-scaling result of Section VI. Wideband observation can provide the raw information required for improved emitter association, but that advantage depends on sufficient processing capacity being available to preserve the relevant observations.

Taken together with the preceding results, the deinterleaving experiment completes the architectural comparison. Direct RF provides substantial benefits in persistent spectral access, response latency, spectral exploitation, and emitter observation completeness, while its principal limitations emerge through finite digital-processing capacity and increased exposure of the wideband front end to strong interferers. Heterodyne architectures provide much stronger analogue selectivity, but at the cost of fragmented and sequential access to the wider electromagnetic environment.

IX. DISCUSSION: ARCHITECTURAL IMPLICATIONS AND HYBRID RECEIVER CONCEPTS

The preceding results show that the transition from heterodyne to Direct RF reception does not simply remove receiver limitations; rather, it redistributes them across the signal chain. Sequential heterodyne reception limits instantaneous spectral access but provides strong analogue selectivity and controlled exposure of downstream stages. Direct RF reception substantially increases spectral visibility and reduces scan-dependent latency, but places greater demands on digital-processing capacity and exposes the wideband analogue front end and data converter to a larger aggregate signal environment.

The comparison therefore suggests that receiver architecture should be considered in terms of where selectivity, agility, and resource limitation are introduced, rather than as a binary choice between analogue and digital reception.

A. Architectural Trade Space

The probability-of-intercept and latency results demonstrate the principal benefit of persistent wideband observation. Against frequency-agile emitters, Direct RF substantially increases the fraction of emissions immediately available for detection, while the absence of sequential tuning removes much of the stochastic delay associated with scan alignment. These advantages are particularly important when emissions are short-lived or when subsequent classification and response are time critical.

The congestion and compute-scaling results expose the corresponding digital constraint. Once a large portion of spectrum is visible simultaneously, the limiting resource becomes the number of signals that can be channelized, detected, classified, and maintained in real time. Increasing instantaneous RF bandwidth therefore has diminishing operational value unless processing capability scales accordingly.

The blocking experiment illustrates the complementary analogue limitation. Wideband capture exposes the Direct RF receiver to strong signals that a more selective heterodyne front end may reject before they reach sensitive downstream stages. Consequently, digital flexibility does not remove the requirement for analogue dynamic range, preselection, gain management, and protection against strong interferers.

The deinterleaving results provide a further consequence of this trade. Persistent wideband access improves the continuity and completeness of observations available for emitter association, but this benefit again depends on sufficient processing resources being available to retain those observations as spectral density increases.

Taken together, the results indicate that the two architectures should not be viewed as universally competing solutions. Instead, they represent different allocations of selectivity and processing between the analogue and digital domains.

B. Hybrid Receiver Architectures

These observations provide a strong motivation for hybrid receiver architectures that combine selective analogue protection with wideband digital capture.

A practical hybrid architecture could retain a relatively broadband RF-sampling path while incorporating adaptive

or switchable analogue preselection ahead of the data converter. Depending on mission requirements, the analogue front end could employ switched filter banks, tunable preselectors, notch filtering, gain control, or other mechanisms to suppress dominant interferers before they consume front-end and converter dynamic range.

Such an arrangement would preserve much of the low-latency and wideband-awareness capability associated with Direct RF while selectively restoring one of the principal strengths of the heterodyne receiver: protection from strong unwanted signals before digitization.

The simulations suggest that the optimum degree of analogue selectivity is likely to depend on the electromagnetic environment. In relatively sparse or rapidly changing spectra, wide instantaneous capture may be prioritised to maximize probability of intercept and minimize sensing latency. In blocker-dominated environments, stronger analogue filtering may instead be beneficial even if this reduces instantaneous spectral visibility. A hybrid architecture could therefore allow the receiver operating point to move dynamically between these extremes.

Importantly, the results also indicate that receiver adaptation need not be confined to the RF front end. The compute-scaling experiment demonstrates that processing capacity itself represents a controllable architectural resource. Dynamic allocation of digital channels, prioritization of signals of interest, and selective reduction of processed bandwidth could therefore complement adaptive analogue filtering.

The resulting receiver is better regarded as a cross-domain architecture, in which analogue selectivity, instantaneous digitization bandwidth, and digital-processing resources are jointly managed according to the prevailing signal environment.

C. Future Perspective: Rydberg-Atom RF Receivers

A further development beyond conventional electronic Direct RF is the emergence of Rydberg-atom RF receivers. These systems use highly excited atomic states whose large electric dipole moments make them extremely sensitive to radio-frequency and microwave electric fields. The interaction between the incident RF field and the atomic energy structure is typically read optically using electromagnetically induced transparency or related techniques, allowing the atomic medium itself to perform part of the sensing and frequency-conversion process [41]. Rydberg receivers should not be considered Direct RF digitizers in the conventional ADC sense. Rather, they represent a parallel receiver paradigm in which RF-field sensing and, in some configurations, heterodyne mixing occur within the atomic medium before optical readout.

This distinction is important, because their potential system-level advantages and limitations differ substantially from those of either the conventional heterodyne or electronic Direct RF architectures considered in this work. Recent demonstrations indicate rapid progress toward broader-frequency atomic reception. In 2026, a Rydberg atomic superheterodyne receiver demonstrated continuous high-sensitivity microwave-field detection over 1–40 GHz by tuning atomic energy levels using the Zeeman effect [42].

A separate 2026 demonstration used a waveguide-coupled Rydberg receiver for frequency-hopping communication spanning 100 kHz to 20 GHz, with hopping rates up to 100 khop/s [43]. These results are particularly relevant to the present discussion because they demonstrate that frequency agility and wide tuning range need not necessarily be realised solely through conventional analogue mixers or broadband electronic data converters.

However, a wide tuning range should not be confused with unlimited instantaneous information bandwidth. Rydberg-receiver response bandwidth remains governed by atomic dynamics, optical readout, interaction time, and the particular sensing scheme employed; the recent technical literature treats bandwidth improvement and practical system integration as continuing research challenges.

Rydberg receivers therefore complement, rather than replace, the architectural trade space examined here. Conventional heterodyne receivers provide mature selectivity and dynamic range, Direct RF provides persistent electronic wideband access, hybrid receivers combine analogue protection with digital flexibility, while atomic receivers introduce the possibility of transferring part of the RF sensing and conversion process into a quantum medium.

From a system perspective, an especially interesting longer-term direction may be the integration of these technologies rather than their mutual replacement. Atomic sensing could potentially provide specialized wide-frequency or field-sensitive detection alongside conventional electronic receiver channels, while digital processing performs classification, association, and response. Such combinations remain speculative at the architecture level and are beyond the scope of the present simulation, but they illustrate how the analogue–digital boundary considered throughout this paper may ultimately evolve into an RF–atomic–optical–digital signal chain.

D. Overall Implications

The principal implication of this study is therefore that future receiver development is unlikely to be governed by instantaneous bandwidth alone. Effective operation in congested and contested electromagnetic environments

requires a balance among spectral visibility, analogue protection, digital-processing capacity, and response latency.

Direct RF moves a substantial part of this balance toward the digital domain and provides clear advantages when persistent spectral awareness and rapid response are required. Heterodyne techniques remain valuable where strong analogue selectivity and blocker rejection dominate the mission requirement. Hybrid architectures provide a means of combining these strengths, while emerging atomic receiver technologies suggest that still more fundamental rearrangements of the receiver signal chain may become possible.

The appropriate architecture will therefore depend not simply on how much spectrum can be digitized, but on how effectively the complete receiver chain can sense, protect, process, and exploit the electromagnetic environment.

X. CONCLUSION

This paper has presented a system-level comparison of heterodyne and Direct RF receiver architectures for congested and contested electromagnetic environments using a common Monte Carlo framework in which both receiver types are evaluated against identical environmental realizations. Seven operationally motivated aspects were examined: probability of intercept, sensing-to-response latency, detection under spectral congestion, compute-limited scalability, blocking robustness, spectral exploitation efficiency, and multi-emitter deinterleaving performance.

The results show that Direct RF reception provides its principal advantages through persistent wideband spectral access. Against frequency-agile emitters, increasing instantaneous capture bandwidth produces a corresponding increase in probability of intercept, while removal of sequential tuner alignment substantially reduces response latency. In the simulated latency experiment, the Direct RF response was concentrated around the approximately $150\mu\text{s}$ processing-chain delay, whereas the heterodyne response exhibited a much broader millisecond-scale distribution governed by scan position and revisit time.

The congestion and compute-scaling experiments demonstrate, however, that wideband visibility should not be equated with unlimited receiver capacity. As emitter density increases, finite downstream processing resources become the dominant limitation in the Direct RF architecture. Increasing the number of concurrently processed signals progressively restores detection performance, showing that the resulting limitation is resource-dependent and scalable rather than an inherent loss of spectral access. By comparison, the scanning

heterodyne receiver remains primarily constrained by the fraction of spectrum visible at any instant.

The blocking experiment identifies the complementary limitation of wideband reception. A strong spectrally displaced interferer that is rejected by heterodyne analogue preselection may remain within the Direct RF capture bandwidth and therefore consume front-end and converter dynamic range before digital filtering can be applied. Conversely, the spectral-exploitation results show the cost of this analogue protection: narrow sequential observation provides access to only a small fraction of the active electromagnetic environment at any one time.

The deinterleaving analysis further shows that persistent wideband observation can improve the completeness and continuity of the observations supplied to emitter-association processes. This advantage diminishes as processing resources become saturated, reinforcing the broader result that the usefulness of increased instantaneous bandwidth depends on the ability of the complete receiver chain to process the information that it captures.

Taken together, the results indicate that neither heterodyne nor Direct RF reception is universally preferable. Instead, the two architectures place their principal constraints at different points in the receiver chain. Heterodyne receivers trade instantaneous spectral awareness for analogue selectivity and strong-signal protection, whereas Direct RF receivers trade increased front-end exposure and digital-processing demand for persistent spectral visibility, rapid access, and processing flexibility.

These findings motivate hybrid receiver architectures in which adaptive analogue protection is combined with wideband digitization and dynamically allocated digital-processing resources. Such architectures offer a means of moving between wideband-awareness and blocker-protection regimes according to the prevailing electromagnetic environment, rather than fixing the receiver permanently at either architectural extreme.

Looking further ahead, emerging Rydberg-atom receivers suggest that the location of the receiver boundary itself may continue to evolve, with part of the sensing and frequency-conversion process transferred from conventional electronics into an atomic medium. While such systems remain outside the scope of the present comparison, they reinforce the broader conclusion that future receiver performance will depend not simply on how much spectrum can be digitized, but on how effectively sensing, protection, conversion, processing, and exploitation are integrated across the complete receiver architecture.

REFERENCES

- [1] U.S. Department of Defense, *Electromagnetic Spectrum Superiority Strategy*, Washington, DC, USA, 2020.
- [2] NATO, "Electromagnetic warfare," nato.int. <https://www.nato.int/en/what-we-do/deterrence-and-defence/electromagnetic-warfare> (accessed Mar. 6, 2026).
- [3] R. A. Poisel, *Electronic Warfare Target Location Methods*, 2nd ed. Artech House, 2012.
- [4] R. Akeela and B. Dezfouli, "Software-defined radios: Architecture, state-of-the-art, and challenges," *Computer Communications*, vol. 128, pp. 106–125, 2018. doi: 10.1016/j.comcom.2018.07.012.
- [5] A. De Maio, S. De Nicola, and Y. Huang, "Cognitive radar and EW systems: Spectrum sharing and adaptive sensing," *IEEE Signal Process. Mag.*, vol. 37, no. 2, pp. 85–98, Mar. 2020.
- [6] R. G. Wiley, *ELINT: The Interception and Analysis of Radar Signals*, Norwood, MA, USA: Artech House, 2006.
- [7] M. A. Richards, *Fundamentals of Radar Signal Processing*, 2nd ed. New York, NY, USA: McGraw-Hill, 2014.
- [8] M. Roberts, "The digitisation of direct RF," *Microwave Journal*, Jan. 2026.
- [9] R. Reddy and S. Sinha, "State-of-the-Art Review: Electronic Warfare Against Radar Systems," in *IEEE Access*, vol. 13, pp. 57530–57567, 2025, doi: 10.1109/ACCESS.2025.3555493.
- [10] E. H. Armstrong, "A method of reducing disturbances in radio signaling by a system of frequency selection," *Proc. IRE*, vol. 5, no. 5, pp. 429–438, Oct. 1917.
- [11] J. H. Collins and P. M. Grant, "A Review of Current and Future Components for Electronic Warfare Receivers," in *IEEE Transactions on Microwave Theory and Techniques*, vol. 29, no. 5, pp. 395–403, May 1981, doi: 10.1109/TMTT.1981.1130368.
- [12] J. Mitola III and G. Q. Maguire, "Cognitive radio: Making software radios more personal," *IEEE Pers. Commun.*, 1999.
- [13] R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 539–550, Apr. 1999, doi: 10.1109/49.761034.
- [14] B. Razavi, "Design considerations for direct-conversion receivers," *IEEE Trans. Circuits Syst. II*, 1997. vol. 44, no. 6, pp. 428–435, Jun. 1997, doi: 10.1109/82.592569.
- [15] J. Grajal, R. Blázquez, G. López-Risueno, J. M. Sanz, M. Burgos, and A. Asensio, "Analysis and characterization of a monobit receiver for electronic warfare," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 1, pp. 244–258, 2003, doi: 10.1109/TAES.2003.1188907.
- [16] J. Tsui, *Digital Techniques for Wideband Receivers*, 2nd ed., Artech House, 2001.
- [17] AMD Xilinx, *RFSoc Technology Overview*, White Paper, 2017.
- [18] S. Thomas et al., "RF-sampling ADCs and SoC integration for wideband systems," *IEEE Solid-State Circuits Mag.*, 2018.
- [19] Ali, Ahmed M. A. et al. "A 12-b 18-GS/s RF Sampling ADC With an Integrated Wideband Track-and-Hold Amplifier and Background Calibration." *IEEE Journal of Solid-State Circuits* 55 (2020): 3210-3224.
- [20] ISSCC 2020 Session 10 Highlights, *IEEE ISSCC Dig. Tech. Papers*, 2020
- [21] S. Henthorn, T. O'Farrell, M. R. Anbiyai and K. L. Ford, "Concurrent Multiband Direct RF Sampling Receivers," in *IEEE Transactions on Wireless Communications*, vol. 22, no. 1, pp. 550–562, Jan. 2023, doi: 10.1109/TWC.2022.3196279
- [22] S. Sun et al., "A 16-channel S-band direct-sampling phased-array receiver," *IEEE Trans. Microw. Theory Techn.*, 2022.
- [23] S. S. Kumar et al., "A 750mW 24GS/s 12b Time-Interleaved ADC for Direct RF Sampling in Modern Wireless Systems," *2023 IEEE International Solid-State Circuits Conference (ISSCC)*, San Francisco, CA, USA, 2023, pp. 1-3, doi: 10.1109/ISSCC42615.2023.10067793.
- [24] ISSCC 2023 Session 8 Overview, *IEEE ISSCC*, 2023.
- [25] J. Heel, H. S. Bindra, S. Louwsma, A. Dezzani and B. Nauta, "A 12.8GS/s Sub-Sampling ADC Front-End With 38GHz Input Bandwidth and >39dB SNDR for 1 to 32GHz in 22nm FDSOI," *2025 IEEE International Solid-State Circuits Conference (ISSCC)*, San Francisco, CA, USA, 2025, pp. 76-78, doi: 10.1109/ISSCC49661.2025.10904786
- [26] B. Murmann, "Trends in high-speed data converter design," *IEEE Solid-State Circuits Mag.*, 2024
- [27] X. Xue et al., "Microcomb-based photonic-electronic ADCs," *Nature Electronics*, 2023.
- [28] D. C. Schleher, *Electronic Warfare in the Information Age*. Norwood, MA, USA: Artech House, 1999.
- [29] D. L. Adamy, *EW 101: A First Course in Electronic Warfare*. Norwood, MA, USA: Artech House, 2001
- [30] A. Farina and F. A. Studer, *Radar Data Processing*, vol. 1. London, U.K.: Research Studies Press, 1985.
- [31] F. Neri, *Introduction to Electronic Defense Systems*, 3rd ed. Norwood, MA, USA: Artech House, 2018.
- [32] S. Haykin, "Cognitive radar: A way of the future," *IEEE Signal Process. Mag.*, vol. 23, no. 1, pp. 30–40, Jan. 2006.
- [33] A. Martone et al., "Cognitive EW systems for contested electromagnetic environments," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 33, no. 7, pp. 28–41, Jul. 2018.
- [34] J. Mitola III, "Cognitive radio architecture evolution," *Proc. IEEE*, vol. 97, no. 4, pp. 626–641, Apr. 2009.
- [35] E. Axell et al., "Spectrum sensing for cognitive radio: State-of-the-art and recent advances," *IEEE Signal Process. Mag.*, vol. 29, no. 3, pp. 101–116, May 2012.
- [36] M. A. Sánchez, M. Garrido, M. López-Vallejo, and J. Grajal, "Implementing FFT-based digital channelized receivers on FPGA platforms," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 44, no. 4, pp. 1567–1585, Oct. 2008, doi: 10.1109/TAES.2008.4667732
- [37] A. Blanchard et al., "Wideband channelization techniques for digital EW receivers," *IEEE Trans. Signal Process.*, vol. 65, no. 6, pp. 1554–1567, Mar. 2017.
- [38] Xilinx Inc., "RFSoc: Radio frequency system-on-chip for direct RF sampling," White Paper WP495, 2020.

- [39] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York, NY, USA: Wiley, 2001.
- [40] D. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. Autom. Control*, vol. 24, no. 6, pp. 843–854, Dec. 1979.
- [41] N. Schlossberger, N. Prajapati, S. Berweger *et al.*, "Rydberg states of alkali atoms in atomic vapour as SI-traceable field probes and communications receivers," *Nature Reviews Physics*, vol. 6, pp. 606–620, 2024, doi: 10.1038/s42254-024-00756-7.
- [42] J. Yao, Z. Sun, Y. Lin *et al.*, "Ultra-wideband continuous spectrum Rydberg atomic superheterodyne receiver with high sensitivity," *Communications Physics*, vol. 9, Art. no. 102, 2026, doi: 10.1038/s42005-026-02529-3.
- [43] J.-D. Nan, J.-R. Chen, B. Liu *et al.*, "Ultra-broadband anti-jamming communication via a Rydberg atomic receiver," *Nature Communications*, vol. 17, Art. no. 9643, 2026, doi: 10.1038/s41467-026-76592-3.



Andrew J. Wileman received the B.Eng. degree (Hons.) in electronic engineering from the University of Nottingham, U.K., and the M.Sc. degree in digital communication systems from Loughborough University, U.K. In 2012, he joined the Integrated Vehicle Health Management (IVHM) Centre, Cranfield University, as a Ph.D. Researcher in prognostics for electronics/defence systems, sponsored by Thales, U.K., and Meggitt. He worked with the U.K. Ministry of Defence and other engineering companies within the U.K. He has lectured in electronics/communication engineering with further and higher education institutions for a number of years. He is currently working as a Research Associate in RF & Microwave Systems in the Electromagnetic Environment Hub (EME Hub) at Loughborough University.