Perception Action Cycle Experimental Radar with Variable Interference and Adaptive Radar Waveforms

Rikki Masson[1], Robert Winter,
Colin Horne, Matthew Ritchie,
Department of Electronic and Electrical Engineering, University College London, UK,
{[1]rikki.masson.23}@ucl.ac.uk

***Abstract*—This paper presents both modelled and laboratory results of a radar system operating under dynamic interference conditions, evaluated with and without a Perception Action Cycle (PAC) adaptation mechanism to dynamically adjust the transmitted waveform in response to spectral interference. System performance is assessed using Signal-to-Noise Ratio (SNR) as a quantitative metric to determine the benefit of adaptive behaviour. Comparative analysis shows that applying PAC yields improvements of up to 17.8 dB with an average gain of 14.6 dB in the modelled environment, and a maximum improvement of 11.2 dB with an average gain of 8.5 dB on the ARESTOR hardware in a controlled laboratory setting. These findings are supported by a structured modelling approach and validated through physical measurements, reinforcing the potential of adaptive sensing for interference mitigation in contested electromagnetic environments.**

Keywords—Cognitive Radar, Electronic Surveillance, FMCW Radar, Adaptive Sensing

I. Introduction

Modern radar systems face increasing pressure to operate within complex, contested, and congested electromagnetic environments. Traditional non-adaptive radars, though effective in static scenarios, lack the agility and intelligence to respond to dynamically evolving threats or spectrum conditions. Their operation is often fixed and predetermined, with limited feedback or context-awareness, making them poorly suited to scenarios requiring real-time adaptability, spectrum coexistence, or collaborative sensing.

Cognitive Radar (CR), first conceptualised in [1], represents a biologically inspired evolution of conventional radar systems, drawing parallels with the perceptual capabilities of echolocating mammals such as bats. A CR system is characterised by its ability to learn from environmental interactions, employ a closed-loop PAC, and adapt transmission and processing strategies. Key elements include environmental learning, adaptive waveform control, and the integration of prior and sensed information. Further theoretical foundations and the evolution of CR are discussed in Section III. While the theoretical underpinnings of CR are well-developed, experimental validation remains limited. The gap between simulation and experimentation has slowed progress toward deployable systems. Integrating scheduling and resource optimisation into high-fidelity modelling to assess CR efficacy is a step forward [2]. Real-world evaluations remain essential to bridge this divide [3].

Before fully cognitive operation can be realised, radar systems must demonstrate adaptability within dynamic RF environments, including spectrum agility and interference mitigation. The growing number of RF devices and increased competition for spectrum access have made coexistence a fundamental challenge in both civilian and Defence applications. Future radar platforms must negotiate shared spectrum in real time and respond to external interference. This is reinforced by the UK Ministry of Defence's Future Operating Environment 2035 report [4], which anticipates a congested and contested electromagnetic spectrum and emphasises the importance of systems that are adaptable, autonomous, and resilient. It identifies real-time, spectrum-aware platforms as a strategic priority, aligning directly with the PAC-based adaptive radar behaviour demonstrated in this study.

Another key requirement is the integration of active radar and passive Electronic Surveillance (ES) sensing. In operational environments where situational awareness depends on diverse data sources, the fusion of passive and active information becomes critical. This calls for multifunction RF systems that can seamlessly switch between or combine sensing modes across a wide range of frequencies. Yet many current platforms are built around single-purpose sensors, limiting their ability to support integrated, adaptive operations.

While literature demonstrating degrees of CR behaviour experimentally is limited, notable contributions include the CREW radar system developed at The Ohio State University [5] and the FFAST framework, which demonstrated both interference avoidance and target-matched illumination [6]. These represent some of the few examples where adaptive and cognitive sensing have been validated in hardware. They underscore the importance of bridging theoretical frameworks with practical implementation, particularly when assessing system performance under real-world operational constraints.

Building on these foundations, ARESTOR [7] was developed by UCL to experimentally validate intelligent RF sensing concepts. Designed as a reconfigurable, multi-role RF sensor, it enables rapid experimentation with adaptive waveforms, mode switching, and PAC mechanisms. This paper presents results from simulated and experimental trials using ARESTOR, demonstrating adaptive waveform selection under variable interference to validate key principles required for CR development.

These investigations are part of a broader research trend advocating for fully adaptive signalling strategies to optimise radar performance in real time. Frameworks such as FFAST [6] show how dynamic waveform selection supports both interference mitigation and mission-aligned sensing, closely aligning with the PAC-based approach explored in this work. Meanwhile, increasing spectral congestion continues to drive the need for intelligent radar systems that can coexist with other spectrum users. The growing field of radar spectrum engineering highlights the urgent need for systems capable of real-time adaptation to evolving spectral conditions [8]. The work presented in this paper addresses this challenge by demonstrating how cognitive mechanisms enable spectral agility and mitigate interference in practical radar implementations.

This paper provides both modelled and experimental results of an FMCW radar mode operating under varying interference. The radar's response is analysed when interference is present in different parts of the sensing spectrum, comparing non-adaptive operation with a PAC-based approach that identifies interference and selects waveforms to avoid it. The trade-off between reduced sensing bandwidth and interference avoidance is evaluated.

The rest of the paper is structured as follows: Section II reviews the hardware used in the experiments; Section III provides an overview of CR theory; Section IV presents the results from modelling, Section V presents the results from experimentation, and Section VI concludes the work.

II. ARESTOR SYSTEM

The ARESTOR platform is a multi-role RF sensing system developed at University College London (UCL). It is built on the Xilinx ZCU111 evaluation board [9], ARESTOR leverages first-generation Radio Frequency System-on-Chip (RFSoC) devices to deliver a highly reconfigurable architecture capable of operating as an active radar, passive radar, and wideband ES receiver.

At its core, ARESTOR [7] is a tightly integrated RFSoC device, which combines eight high-speed Analog-to-Digital Converters (ADCs) and Digital-to-Analog Converters (DACs) with sampling up to 4 GS/s and 6.5 GS/s respectively, with Field Programmable Gate Array (FPGA) fabric and embedded processing cores. This enables significant control over signal generation and capture, as well as real-time processing and adaptivity. The flexibility of the system is extended through a modular development framework that allows rapid reconfiguration between roles and facilitates complex RF sensing experiments without the need for multiple disparate systems.

To enable scalable and robust RF experimentation, a bespoke FPGA framework has been developed. This framework defines generic transmitter and receiver hardware modules, which may be instantiated as needed to realise specific sensing configurations. The platform also supports extensive Digital Signal Processing (DSP) within the programmable logic, including decimation, digital mixing, deramping for Frequency Modulated Continuous Wave (FMCW) signals, and passive radar pre-processing. Efficient data transfer is supported via a Direct Memory Access driven architecture, enabling real-time access to high-volume radar data through the processing system and onward to external storage or analysis pipelines.

A key innovation within ARESTOR is its ability to operate multiple sensing modes simultaneously by dynamically allocating ADC/DAC resources, implementing mode-specific DSP chains in hardware, and managing real-time synchronisation across sensing threads.

The combination of high bandwidth, digital flexibility, and reconfigurable sensing modes positions ARESTOR as a capable experimental testbed for advancing research into the PAC in CR systems. Its integrated design enables dynamic selection of sensing modes and waveforms, in response to the perceived environment. This functionality makes ARESTOR ideally suited for experimental studies involving adaptive sensing in contested RF environments.

III. COGNITIVE RADAR THEORY

CR distinguishes itself from adaptive radar not solely through the presence of feedback or adaptability, but through the incorporation of memory, learning, and inference. A cognitive system must go beyond reactive control; it must possess the ability to improve performance over time by internalising experience [1].

The conceptual foundation for CR was established in early work that introduced PAC, memory, and learning as central to radar operation [1]. This theoretical vision positioned CR as a biologically inspired evolution of conventional radar systems, capable of dynamically adapting to its environment.

Later contributions further emphasised CR as a transformative technology, highlighting its potential for spectrum efficiency and adaptability in complex electromagnetic environments [10]. More recent overviews have tracked progress in adaptive waveform design, environment sensing, and real-time processing, demonstrating how CR is moving closer to operational deployment [11]. Collectively, these works underscore the importance of experimental systems that validate cognition under real-world conditions.

A critical differentiator between adaptive and cognitive radar is the role of memory. In CR, memory operates at multiple levels: short-term memory captures recent environmental observations and supports immediate adjustments, while long-term memory retains historical information to inform future strategies. This persistent knowledge base allows the system to detect temporal patterns, refine hypotheses, and make context-aware decisions. Without memory, even a system capable of adaptive behaviour cannot reason about past interactions or anticipate future outcomes and thus cannot be considered truly cognitive. This raises a key consideration: can a radar system be called cognitive if it selects waveforms based solely on current measurements? A system that ignores historical performance or evolving environmental conditions may exhibit agility, but it lacks introspection and the capacity to optimise behaviour over time.

Cognition implies strategic decision-making under uncertainty, guided by knowledge, experience, and feedback. This process, not just the presence of feedback, defines the CR paradigm. For a system to be considered cognitive, it must maintain an internal representation of the environment, adapt its behaviour based on past outcomes, apply reasoning to optimise sensing and transmission, and improve its performance through experience. These characteristics form the theoretical framework [3] for evaluating CR systems and establish the benchmarks against which experimental platforms such as ARESTOR can be assessed.

IV. MODELLED RESULTS

Advanced RF sensors of the future will require the capability to determine not just how to adapt, but when and why adaptation is necessary. For example, the presence of interference in a band currently used by the radar may not justify a change in behaviour. This decision depends on the relative levels of the radar signal and the interferer, the interferer's spectral location, and the availability of suitable alternate waveforms. To explore this, a modelling environment was developed to simulate radar operation and assess the consequence of waveform adaptation in the presence of interference.

All simulation testing was carried out entirely at baseband, spanning 0–100 MHz, and designed to reflect the signal generation and reception characteristics of the ARESTOR radar hardware. An FMCW radar signal was generated using a fixed-duration chirp and a linearly swept frequency profile. The waveform was constructed using a sample period of approximately 1.39 ns, calculated from the DAC sampling rate of 5.76 GHz and an interpolation factor of 8. These values were selected to accurately replicate the baseband characteristics of the ARESTOR system. The signal was sampled using a 4096-point array. This size was selected to provide sufficient frequency and range resolution while remaining compatible with practical hardware memory and processing constraints.

A library of noise interference waveforms was created in the form of Additive White Gaussian Noise (AWGN), either spanning the full operating band or confined to specific sub-band regions. The full spectrum was divided into 10 sub-bands and the noise library resulted in 55 different combinations of bandwidth and position ranging from single 10 MHz size interference to 100 MHz full bandwidth examples. This allowed the precise control of the spectral shape and bandwidth of the interferer to emulate a wide range of realistic operating conditions. The simulation environment supported the evaluation of radar performance with and without adaptive behaviour, enabling direct comparison across interference scenarios and waveform selection strategies. A representative example of the modelling process is shown in Fig. 1, which illustrates the waveform adaptation logic for the first 10 test configurations. In these scenarios, each interferer occupies a 10 MHz sub-band, incrementally stepped across the 0–100 MHz spectrum. All FMCW and noise transmission was programmed in this test configuration.

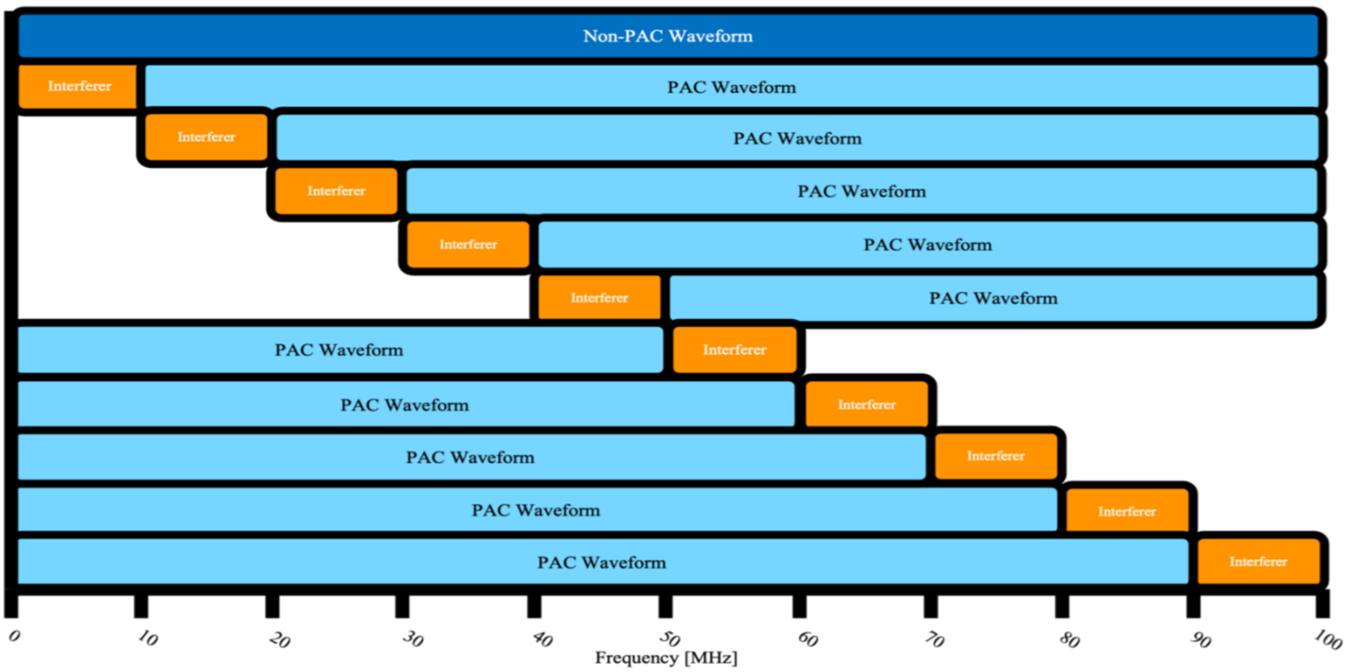

Fig. 1. Non-PAC, Interferer and PAC waveform bandwidths for the first 10 test configurations only.

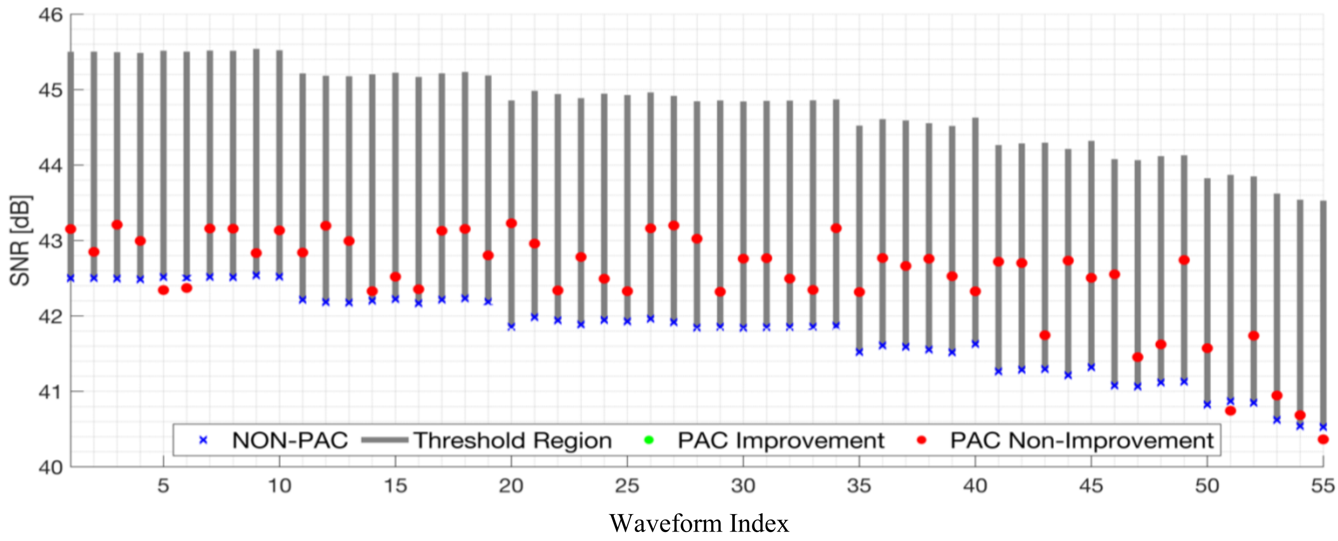

Fig. 2. Modelled SNR PAC vs Non-PAC – Lower Power Interferer.

Fig. 1 illustrates the non-PAC transmit waveform, the spectral footprint of the interferer, and the adapted PAC waveform for the first 10 test configurations. While the non-PAC system consistently transmits across the full 100 MHz band regardless of interference, the PAC system adaptively repositions or narrows the transmit waveform to avoid spectral overlap. This visualisation reinforces how PAC avoids interference where possible, preserving usable bandwidth and improving performance.

The PAC waveform is defined as the longest continuous section of bandwidth that doesn't overlap with the created noise interference defined above. This creates a trade-off: reducing bandwidth helps avoid overlap with interference but negatively affects the radar's resolution and reduces the SNR on the target.

To quantify performance across test cases, SNR was selected as the primary metric of success. It offers a clear and measurable indication of how interference affects the signal quality and the extent to which adaptive waveform selection improves system resilience. The signal component of the SNR was calculated by identifying the absolute global maximum within the FFT data. The noise component was derived by computing the absolute mean value of the FFT bins from 50 to 4000, a method chosen to avoid influence from the main lobe or target side-lobes, thereby capturing a realistic approximation of the background noise floor.

A comparative overview of SNR results across all 55 noise configurations is presented in Fig. 2 and Fig. 3, which illustrate outcomes under two different interference power levels. In the lower-power scenario (Fig. 2), the full-band noise interference was set approximately 22 dB below the average signal level. Under these conditions, PAC adaptation did not surpass the 3 dB threshold in any case and, in several instances, led to degraded performance due to the reduced FMCW bandwidth. PAC delivered an average SNR gain of just 0.8 dB. This outcome highlights the inherent resilience of the FMCW waveform and reinforces that adaptation decisions should not be based solely on the spectral position of an interferer but must also consider its relative strength.

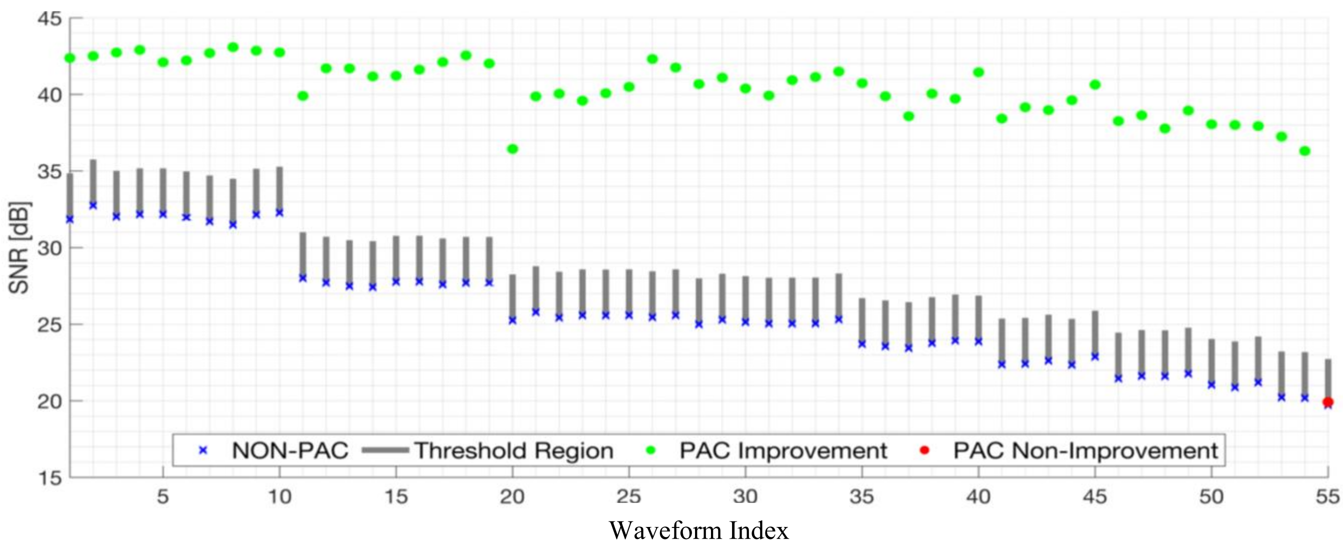

Fig. 3. Modelled SNR PAC vs Non-PAC – Higher Power Interferer

In contrast, Fig. 3 presents results from a higher-power interference scenario in which the noise power was increased by 16 dB, placing it approximately 6 dB below the average signal level. In this condition, PAC adaptation yielded a significant SNR improvement in 54 of the 55 configurations, exceeding the 3 dB threshold in all adaptation test cases. The only configuration without gain occurred when interference spanned the full 0 to 100 MHz band, resulting in the same waveform being selected in both modes. Across the remaining scenarios, PAC delivered a minimum SNR gain of 7.8 dB, a maximum gain of 17.8 dB, and an average improvement of 14.6 dB. These findings demonstrate the substantial advantage of employing adaptive waveform strategies in contested spectral environments, particularly when the interference strength reaches operationally significant levels.

V. Laboratory Results

To complement the modelled analysis, the same 55 test configurations were replicated in a controlled laboratory environment using the ARESTOR hardware [7]. The setup involved a direct cable loopback configuration, enabling consistent and repeatable signal capture without over-the-air variability. Interference was generated internally using a single DAC channel from the ARESTOR system, with the radar transmission and the interference combined prior to reception. A Mini-Circuits 2-way power splitter/combiner [12] was used to combine the two signals before they were received by the ADC, ensuring precise alignment and preserving signal integrity.

The interference waveforms used in the laboratory were based on the AWGN signals from the simulation environment and reformatted into binary files compatible with the ARESTOR system. The sample length of each binary file was explicitly designed to consist of 345600 samples to match the FMCW chirp period of 0.24 seconds, ensuring continuous interference throughout each FMCW chirp.

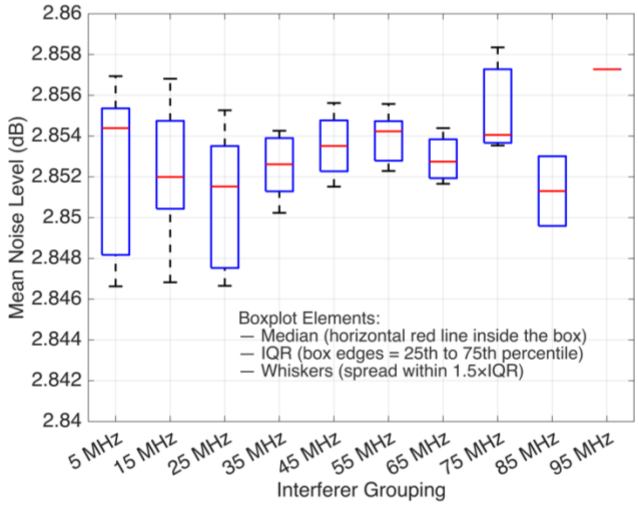

Fig. 4 Interferer Waveform Grouping Noise Comparison

This process produced 55 interference waveforms with comparable power spectral density (PSD). Fig. 4 illustrates the interference waveforms in their binary form prior to the PSD scaling applied by ARESTOR during FPGA loading. A scaling factor was applied to each waveform to ensure that, as the bandwidth increased, the PSD remained uniform across all test conditions.

The resulting binary files were assessed to ensure consistent interference power across all waveform bandwidths. Median PSD levels were evaluated by bandwidth group and found to be nearly identical, demonstrating uniformity across the 55 test configurations. Validation of the FMCW and interference waveforms was performed using two approaches: analysis of data captured in a short cable loopback configuration, and independent measurements using a Keysight FieldFox Microwave Analyser N9915A [13]. This dual method ensured that the transmitted waveforms adhered to the correct frequency profile, structure, and power levels for each test configuration.

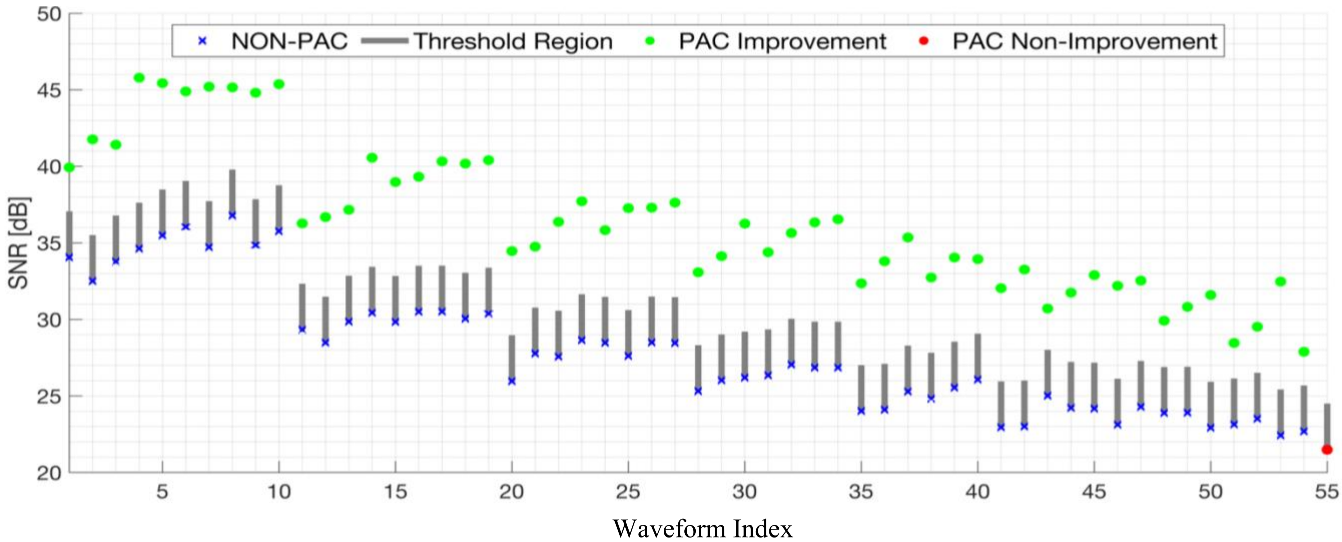

Fig. 5. Laboratory SNR PAC vs Non-PAC Results

Although ARESTOR supports ES and real-time waveform adaptation, a prescribed waveform schedule was implemented in these experiments to ensure repeatability and isolate the effect of interference.

Each test transmission and capture contained 50 chirps. To calculate the Noise level without artificially reducing the estimate through averaging, only the first chirp was analysed. Noise was computed as the mean magnitude across FFT bins 50 to 4000. The signal component was defined as the global maximum from the complete dataset of 50 chirps. This method ensured consistency with the simulation while enabling realistic evaluation of system performance under physical interference conditions

This hardware-based setup allowed for direct evaluation of how the PAC and non-PAC configurations perform when exposed to physically generated interference under realistic signal chain conditions.

A comparative overview of the SNR performance across all 55 test configurations was also evaluated using the laboratory measurements – in this configuration the transmitted waveforms were programmed. Fig. 5 shows the PAC and non-PAC SNR values were calculated for each test case using the same definitions as in the modelled environment. The PAC outcome was considered beneficial only if it exceeded the non-PAC SNR by a threshold of 3 dB, reflecting a cost associated with adaption. In all 54 test scenarios where adaptation occurred, the PAC waveform achieved a minimum SNR gain of 5.2 dB, a maximum gain of 11.2 dB, and an average improvement of 8.5 dB. These laboratory results reinforce the practical viability of PAC-based adaptation in real hardware and confirm its potential to mitigate interference and enhance signal quality under realistic operating conditions.

VI. Conclusions

Adaptive radar systems must operate in congested and contested RF environments, where spectral conditions are dynamic, and interference is prevalent. This paper has demonstrated a practical PAC-based waveform adaptation strategy that enables coexistence with other signals in-band. Compared to non-PAC operation, the approach delivered measurable SNR gains by avoiding spectral overlap.

In simulation, PAC adaptation outperformed non-adaptive operation in 54 of 55 configurations, achieving SNR gains from 13.5 dB to 35.2 dB, with an average of 24.1 dB. Hardware trials validated this behaviour, with gains ranging from 5.2 dB to 11.2 dB and an average of 8.5 dB. These results confirm the value of real-time adaptation under both modelled and practical conditions.

Performance improvements were driven by selectively reducing bandwidth to avoid interference, with the best results occurring when interference appeared at band edges. These findings highlight the potential of cognitive strategies to enhance sensing in dynamic spectral environments.

Future work will explore non-contiguous waveform designs that retain resolution while avoiding interference, over-the-air trials under realistic propagation conditions, and Reinforcement Learning approaches to enable intelligent, experience-based adaptation.

Acknowledgments

The authors would like to acknowledge the support of the Royal Air Force (RAF) which supports Rikki Masson’s PhD, AFRL (award FA8655-23-1-7081) which supports the authors research outputs in Intelligent Radio Frequency Sensing and finally Leonardo UK /Royal Academy of Engineering (RAEng) for their support of Dr. Matthew Ritchie’s Fellowship.

References

[1] S. Haykin, "Cognitive radar: a way of the future," in IEEE Signal Processing Magazine, vol. 23, no. 1, pp. 30-40, Jan. 2006, doi: 10.1109/MSP.2006.1593335.

[2] H. K. Nguyen, J. R. Guerci, S. K. Gogineni, J. S. Bergin, E. P. Blasch and M. Rangaswamy, "Cognitive Radar Scheduler Architecture," 2024 IEEE Radar Conference (RadarConf24), Denver, CO, USA, 2024, pp. 1-6, doi: 10.1109/RadarConf2458775.2024.10548605.

[3] C. P. Horne, An Experimental Investigation of Cognitive Radar, Ph.D. dissertation, Dept. Electronic and Electrical Eng., Univ. College London, London, United Kingdom, 2020.

[4] Ministry of Defence, Strategic Trends Programme: Future Operating Environment 2035, Aug. 2015. [Online]. Available: https://assets.publishing.service.gov.uk/media/6286575de90e071f69f22600/FOE.pdf

[5] A. M. Jones, C. P. Horne, H. D. Griffiths, G. Smith, A. Mitchell and P. John-Baptiste, "Experimentation of an adaptive and autonomous RF signalling strategy for detection," 2018 IEEE Radar Conference (RadarConf18), Oklahoma City, OK, USA, 2018, pp. 1213–1218, doi: 10.1109/RADAR.2018.8378735.

[6] C. P. Horne, A. M. Jones, G. E. Smith and H. D. Griffiths, "Fast Fully Adaptive Signalling for Target Matching," in IEEE Aerospace and Electronic Systems Magazine, vol. 35, no. 6, pp. 46-62, 1 June 2020, doi: 10.1109/MAES.2020.2970243.

[7] N. Peters, C. Horne and M. A. Ritchie, "ARESTOR: A Multi-role RF Sensor based on the Xilinx RFSoC," 2021 18th European Radar Conference (EuRAD), London, United Kingdom, 2022, pp. 102-105, doi: 10.23919/EuRAD50154.2022.9784466.

[8] H. Griffiths, L. Cohen, S. Watts, E. Mokole, C. Baker, M. Wicks, and S. Blunt, “Radar spectrum engineering and management: Technical and regulatory issues,” Proceedings of the IEEE, vol. 103, no. 1, pp. 85–102, Jan. 2015, doi: 10.1109/JPROC.2014.2360591.

[9] Xilinx, “Zynq UltraScale+ RFSoC.” [Online]. Available: https://www.xilinx.com/products/silicon-devices/soc/rfsoc.html

[10] J. R. Guerci, “Cognitive radar: The next radar wave?,” Microwave Journal, vol. 54, no. 1, pp. 22–36, Jan. 2011.

[11] M. S. Greco, F. Gini, P. Stinco, and K. Bell, “Cognitive radars: On the road to reality: Progress thus far and possibilities for the future,” IEEE Signal Processing Magazine, vol. 35, no. 4, pp. 112–125, Jul. 2018. doi: 10.1109/MSP.2018.2823780

[12] Mini-Circuits, “ZFRSC-183-S+ 2-Way Resistive Power Splitter/Combiner,” Mini-Circuits, [Online]. Available: https://www.minicircuits.com/pdfs/ZFRSC-183-S+.pdf

[13] Keysight Technologies, N9915A FieldFox Microwave Analyzer – Data Sheet, Keysight Technologies, Santa Rosa, CA, USA, 2023. [Online]. Available: https://www.keysight.com/gb/en/assets/7018-03314/data-sheets/5990-9783.pdf