Rathera

An X-band Pulse Compression Instrumentation Radar for RCS Characterisation

Prepared by: Alex Scarbro

Contributor: T.G. Pelham

Issue: 1

Issue	Date	Changes	Signed by
DRAFT	-	Initial draft	-
1	31-10-22	First release	A. Scarbro

Abstract

The design, development, and realisation of a system for Radar Cross Section (RCS) characterisation of an Uncrewed Surface Effect Vehicle (USEV). Consisting of an X-band Qorvo 4W PA/LNA, a near full-band up/down microwave block converter, and a wideband software-defined pulse processor, the system can also be used as a platform for the development of complex radar waveforms. An integrated electro-optics package provides wireless laser range finding and low-latency video for visual object recognition and tracking of cooperative and non-cooperative targets.

Figure 1 The radar in operation at Elvington Airfield, Yorkshire.

Background

Rathera were first tasked with developing a means of characterising the RCS of a small-scale model manufactured from the same materials as the full-sized craft. A Compact Range was created in the anechoic chamber at Unit 3 Compliance consisting of a network analyser (operating in Time Domain mode), a pair of TX/RX feed antennas, a 1m offset parabolic reflector, and a control computer to instruct the turn-table, trigger a measurement, and collect & process data.

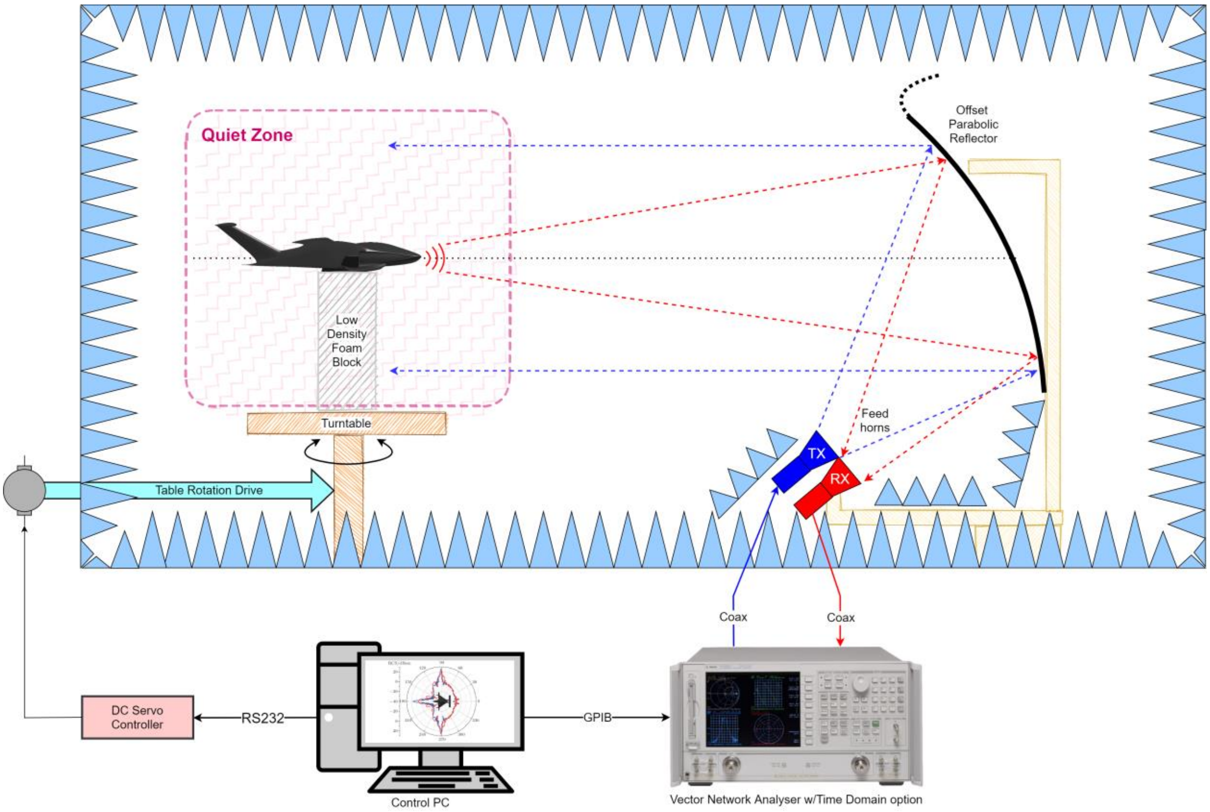

Figure 2 The Compact Range built at Unit 3 Compliance.

The feed antenna positions were first optimised by maximising S_{21} with the assistance of a third calibration antenna located at the rotational axis of the turntable. The network analyser was then configured to sweep in 1601 steps from approximately 1.5 – 11.9GHz to give a spatial response resolution of circa 56mm and an alias free range of 46m.

The scale model was placed upon a low-density polystyrene block and rotated through 360 degrees in azimuth to measure the effective RCS at all incident angles.

A full sweep with 5° steps could be completed in under 15 minutes allowing the placement and aerodynamic configuration of the USEV to be optimised within a reasonable time frame.

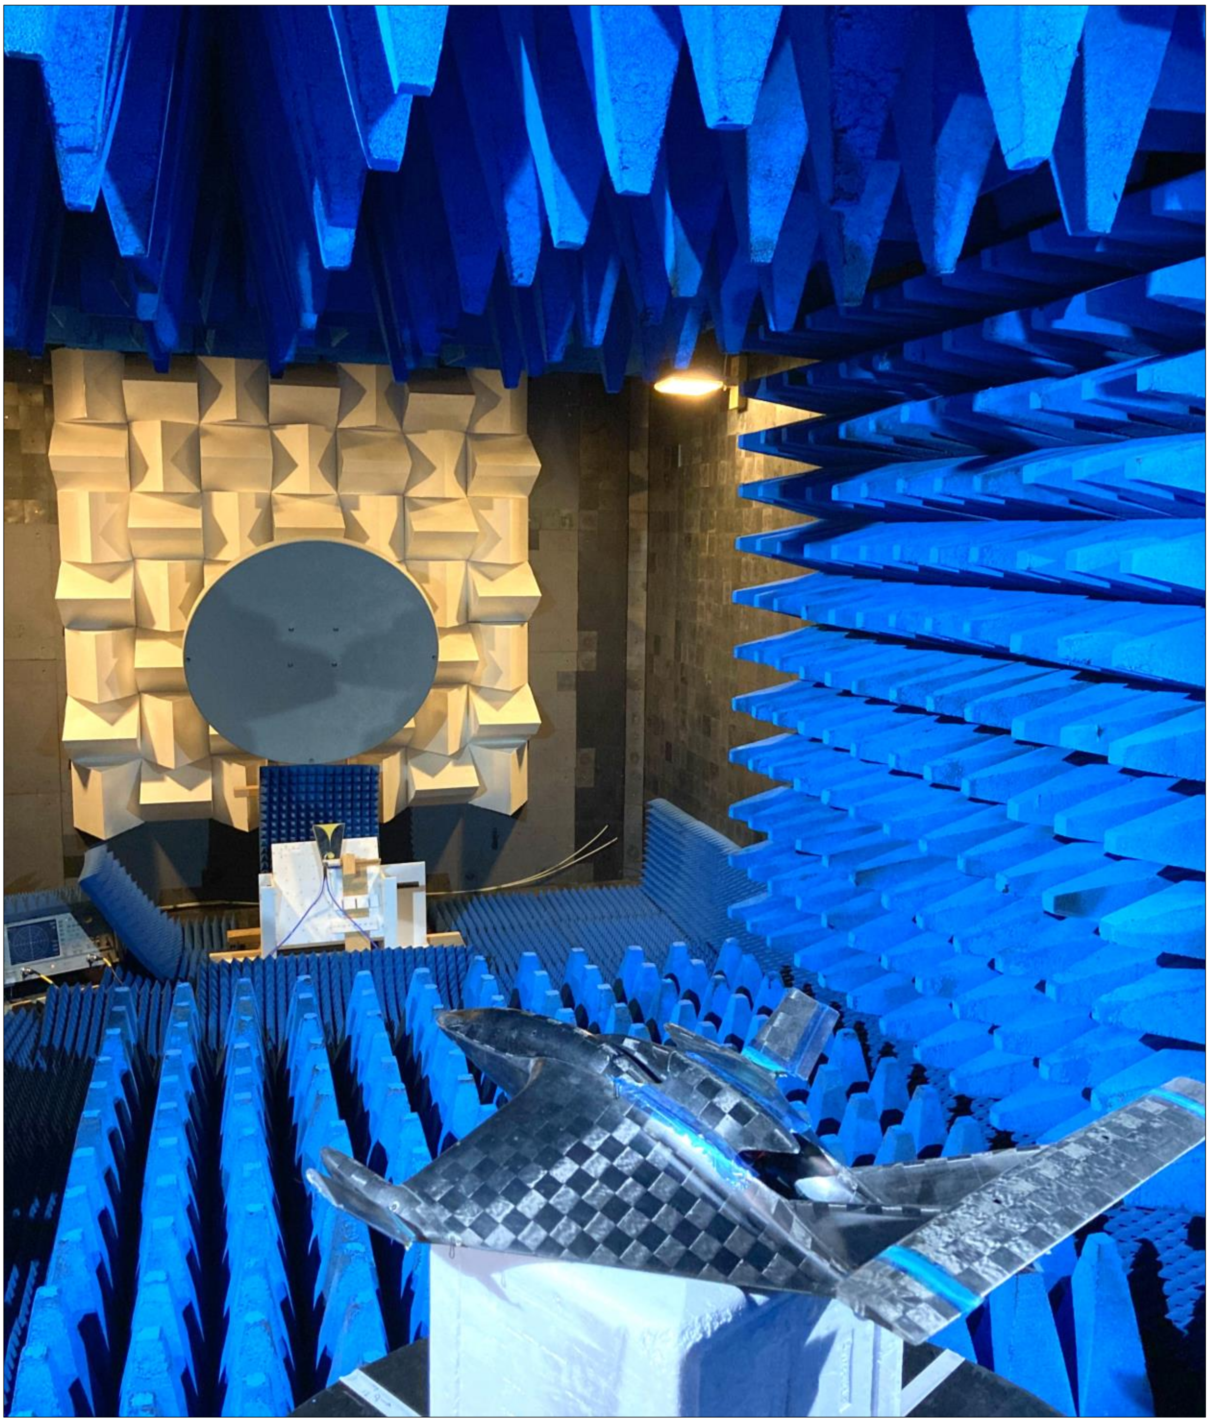

Figure 3 The scale model undergoing RCS characterisation.

Absolute calibration of the measured results was achieved through the use of a reference target of known RCS. A flat plate or corner reflector can be used, though a polished metallic sphere is preferred as it can provide an easily predictable frequency-independent RCS at all relative orientations. The only requirement is that the sphere radius must be at least 1.5 times the wavelength such that it is operating in or close to the "optical" region as shown in the graph below.

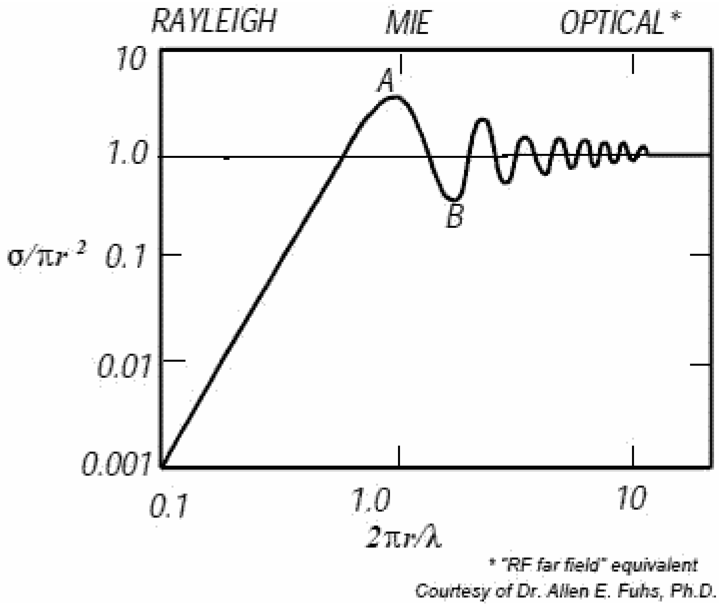

Figure 4 Rayleigh Scattering, Mie Scattering, and Optical regions.

Given that the above criteria are met, the RCS (σ) of a sphere can be calculated as followings:

$$\sigma = \pi r^2$$

Our 200mm diameter calibration sphere (shown on top of a low-density hollow polystyrene box) has a calculated RCS of 0.031m^2.

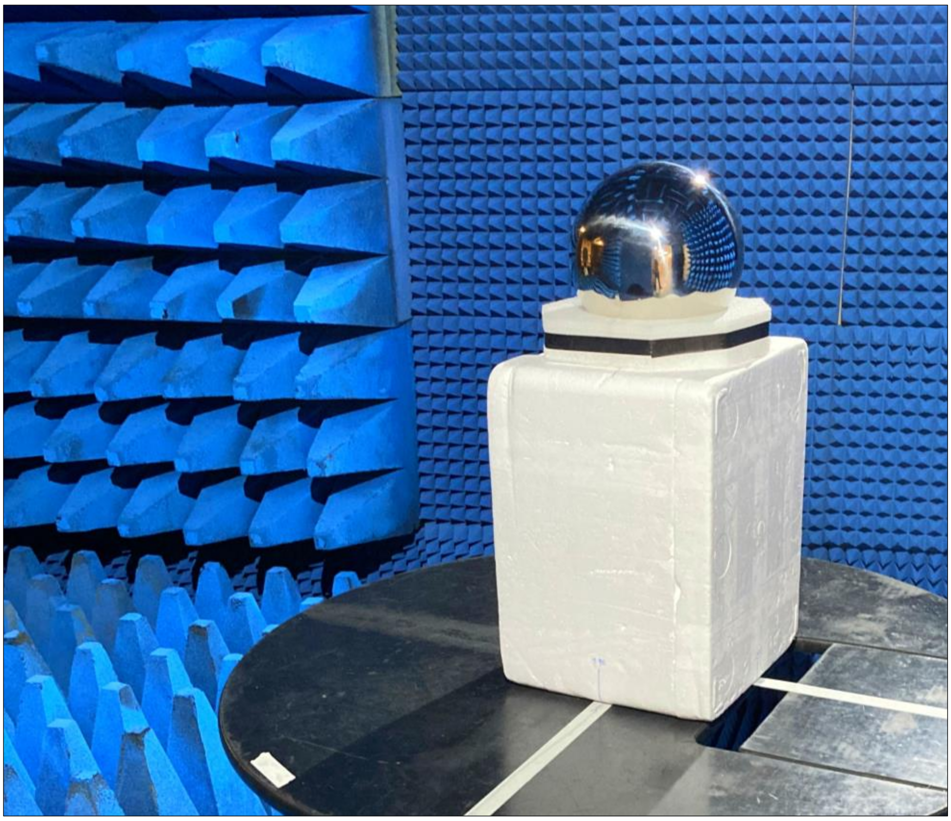
Figure 5 The RCS calibration sphere.

Processing of the result showed the Compact Range had excellent spatial resolution and provided the capability to identify specific features of the USEV's shape that degraded its otherwise low-RCS performance. A crude estimate of the RCS of a full-size craft could then be made.

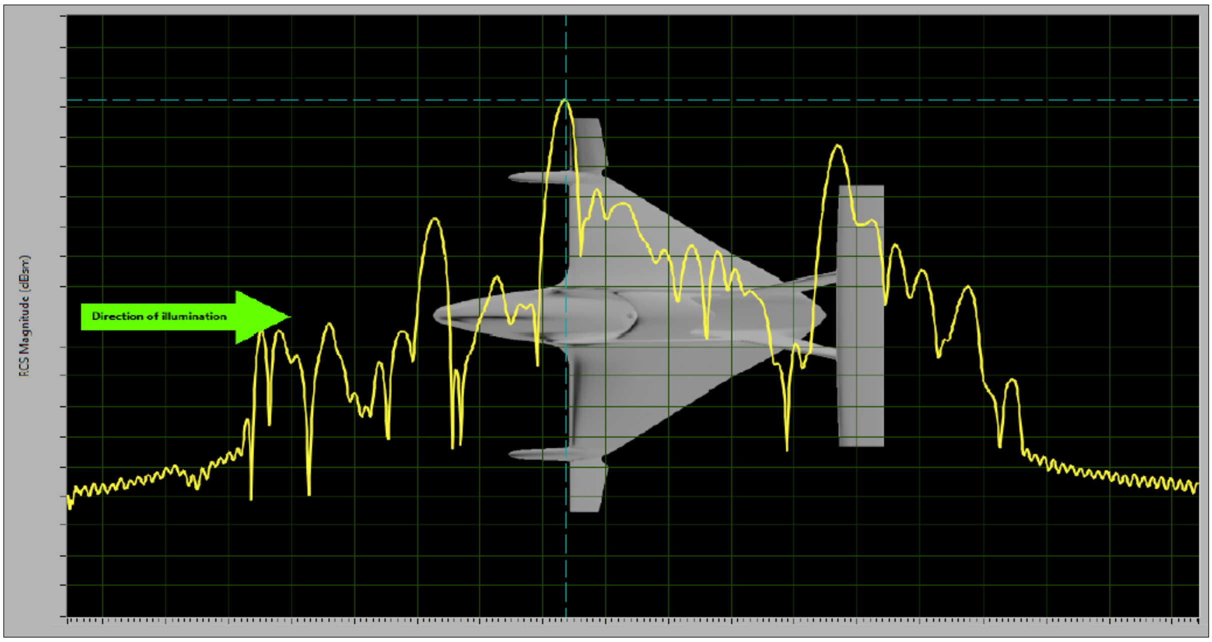

Figure 6 A plot showing RCS vs range and the dominant features degrading overall RCS.

Aims & Objectives

With chamber-based testing completed successfully, Rathera was then tasked with developing a means of performing real-world RCS measurements during outdoor flight tests of the first full-scale model.

Several aims & objectives were identified to meet the current and future needs of the measurement programme:

- Develop a self-contained instrumentation radar system operating in one of the common microwave bands (S, C or X band) with enough sensitivity to detect a $1m^2$ RCS target at >1000m and a minimum range of <75m.
- Target a range resolution less than or equal to the aircraft size (2-3m).
- Take a modular rather than integrated approach – we are only building one after all.
- Frequency plan around available connectorised filtering elements.
- Recycle existing designs and hardware wherever possible.
- Use our in-house PCB & electronic prototyping facilities.
- Minimise the number of bespoke metalwork designs (due to lead time and cost).
- Design the RF architecture such that the radar waveform can easily be developed and modified without the need to make hardware or electrical changes.
- Allow for future integration of different target acquisition and tracking technologies.
- Provide a means of absolute calibration for the measured results at the flight test site.
- Ensure the radar can operate from a single-phase AC $230V_{RMS}$ portable generator.
- Have the radar licensed and be considerate to other spectrum users – we don't need kilowatts of transmit power!

Back to Basics – Radar Range Equation

When setting out to design a radar we should start with the basic equation for range in a two-way (round-trip) monostatic radar:

$$Range \cong \sqrt[4]{\left(\frac{P_T G_T G_R \lambda^2 \sigma}{(4\pi)^3 S_{min}}\right)}$$

Where:

P_T is transmitted power.
G_T is the transmit antenna gain.
G_R is the receive antenna gain.
λ is the radar signal wavelength.
σ is the radar cross section (RCS).
S_{min} is the minimum detectable signal (e.g. thermal noise floor + minimum required SNR)

It should be noted that the above equation does not take in to account atmospheric absorption, however, and at X-band this is less than 0.02-0.03dB/km and so is inconsequential for our application.

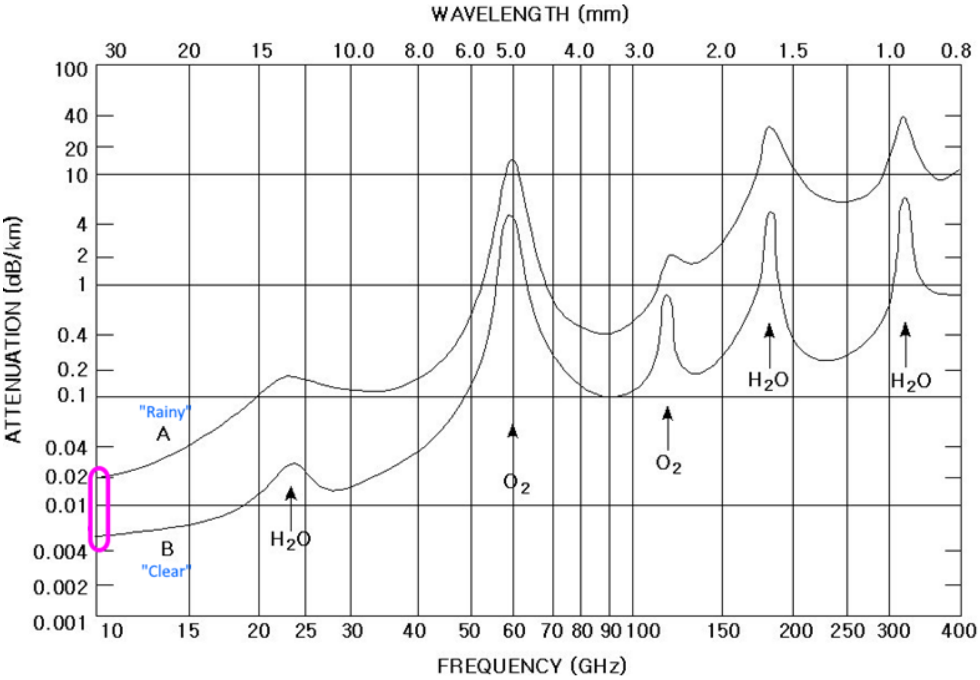

Figure 7 Atmospheric attenuation in "Rainy" and "Clear" weather vs frequency.

Having the largest gain for a given antenna size would be advantageous and would allow us to operate with a lower transmit power and so we chose to operate at the highest practical frequency (X-band, 8-12GHz).

The gain of a simple dish-type parabolic antenna can be calculated as:

$$G = 10\log_{10} k\left(\frac{\pi D}{\lambda}\right)^2$$

Where:

G is the gain over an isotropic source in dBi.
k is the efficiency factor which is generally around 0.5 to 0.6 (50-60%).
D is the diameter of the parabolic reflector in metres.
λ is the wavelength of the signal in metres.

A 0.5m dish, perfectly illuminated, operating at 9.25GHz would have a maximum gain of circa 30dB.

We can initially estimate the Minimum Detectable Signal (S_{min}) of our receiver as:

$$S_{min} \cong -174 + 10\log_{10}(BW) + NF + SNR_{min}$$

Where:

BW is the bandwidth of the receiver (matched to the signal BW).
NF is the Noise Figure (in dB).
SNR_{min} is the minimum required SNR for detection with a low false alarm rate.

Choosing a 10MHz as the signal BW, a NF of 2.7dB and an 8dB min SNR, we get a Minimum Detectable Signal of -93dBm.

Plugging these calculated values (along with a 1W peak transmit power) in to the basic radar range equation (with linear conversion of quantities) we get a maximum range of:

$$1.0km \cong \sqrt[4]{\left(\frac{1W \times 1000 \times 1000 \times (0.0324)^2 \times 1\,\mathrm{sq.m}}{(4\pi)^3 \times 4.67 \times 10^{-13}W}\right)}$$

The range resolution of a simple pulsed radar can be calculated as:

$$S_{range} \geq \frac{C_0\,\tau}{2}$$

Where:

C_0 is the speed of light.
τ is the pulse width.

Therefore, a 1us pulse, will have a range resolution of 150m which is clearly too large for our application. We could (in theory) decrease the pulse width to 20ns to achieve a range resolution of 3m. However, this causes two problems, when we decrease the pulse width, we decrease the energy that reaches the target, and we increase the required receiver bandwidth. Both of these negatively impact the SNR at the receiver.

Pulse Compression

As previously described, the range resolution of an unmodulated pulsed radar is solely dependent upon the pulse width (τ). Two targets located within the spatial extent of a single pulse are blurred in to one indistinguishable pulse envelope as shown below.

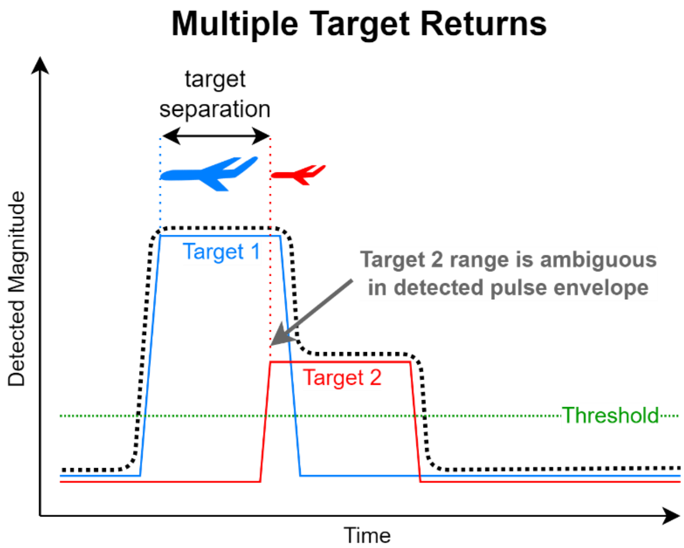

Figure 8 Multiple targets within the same pulse extent.

Essentially with a simple (non-modulated) pulse waveform, we can have either a long range or high resolution, but not both.

A long-standing solution to this challenge is Pulse Compression where the bandwidth of the transmitted pulse is increased (either by frequency or phase modulation) whilst a relatively long pulse width is still used to maximise the total energy impinging on the target.

A transmit waveform with Linear Frequency Modulation (LFM) is a popular choice that is easy to synthesize accurately either in complex baseband or at a low Intermediate Frequency (IF). In this case, the bandwidth of the pulse is the sweep stop frequency minus the start frequency ($F_2 – F_1$).

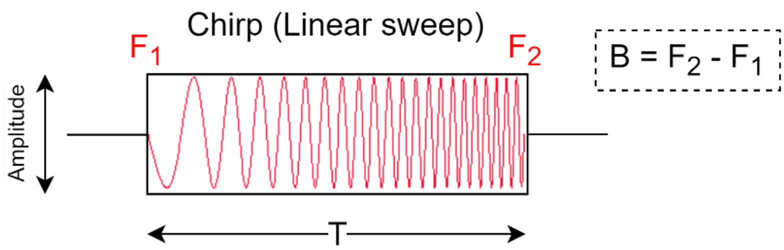

Figure 9 An intra-pulse Linear Frequency Modulation chirp.

Through compression of the received waveform(s) it becomes possible to localize multiple reflecting objects both within and without the pulse extent.

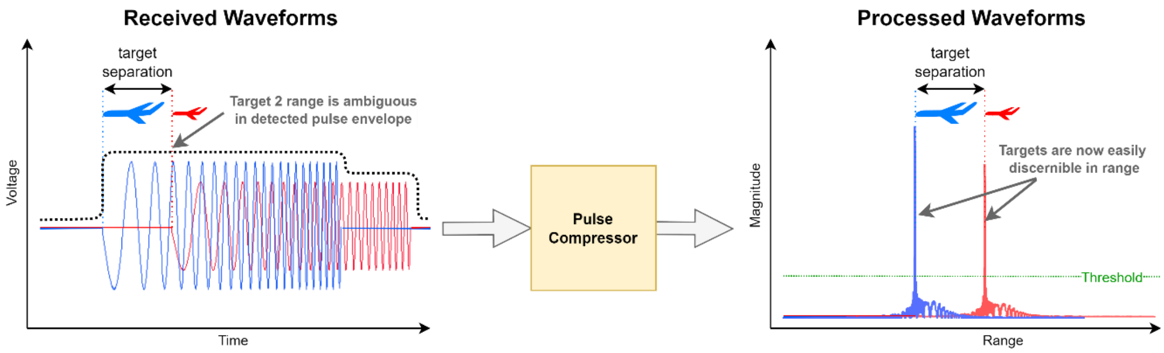

Figure 10 Increased range resolution through Pulse Compression.

Compression of the received waveform produces a sinc function in the time domain whose magnitude is given by:

$$\left|RX_{compressed}\right| \approx \sqrt{TB} \times \left(\frac{\sin(\pi Bt)}{(\pi Bt)}\right)$$

Where:

T is the uncompressed pulse width.
B is the bandwidth of the transmitted pulse. In the case of a linear frequency modulation, this value is the width of the frequency sweep ($f_2 - f_1$).
t is time.

The compressed waveform also experiences an amplification of $\sqrt{TB}$ as shown in the plot below.

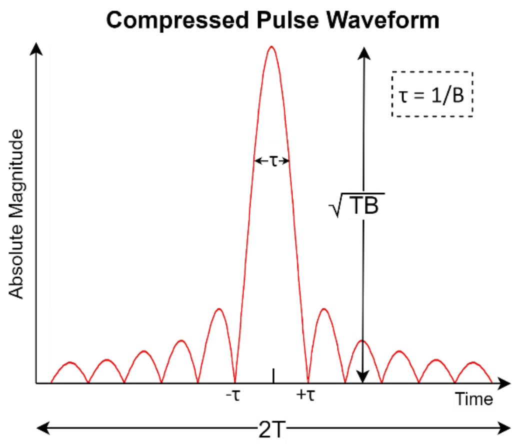

Figure 11 Pulse compressed waveform characteristics.

Since the waveform produced by compression is a sinc function, high sidelobes (in the time domain) will obscure smaller nearby target returns. Windowing of the received (or transmitted) waveform samples is a critical step and there has been much research into the optimisation of functions that maximise sensitivity and range resolution. A simple Hanning window is used for this radar.

The range resolution of a pulse compression radar can be calculated by:

$$S_{range} \geq \frac{C_0}{2B}$$

Where:

C_0 is the speed of light.

B is the bandwidth of the transmitted pulse. In the case of a linear frequency modulation, this value is the width of the frequency sweep ($f_2 - f_1$).

Consequently, a 100MHz pulse bandwidth can yield a range resolution as small as 1.5m.

Pulse compression, therefore, combines the high total energy of a long pulse with the high range resolution of a short pulse.

Whilst the compression process amplifies the main lobe (resulting from the correlated receive waveforms), the uncorrelated noise remains unaffected and so an improvement in SNR is observed that is proportional to the product of T x B.

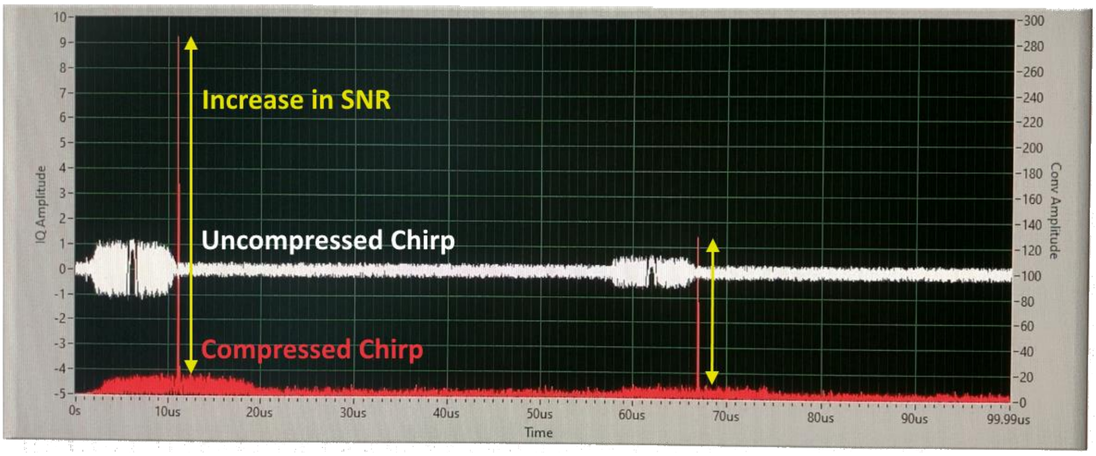

Figure 12 Pulse compression modelling with noise.

The illustration below (left plot) shows a low-amplitude chirp signal buried in the noise. After processing (right plot), the compressed pulse rises above the noise to a level where it would be easily detectable against a threshold.

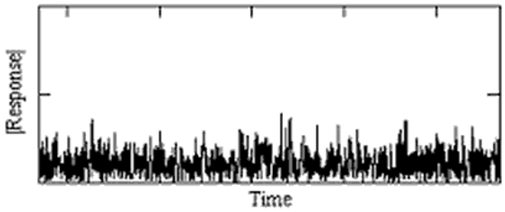

Figure 13 Noise sampled with embedded chirp.

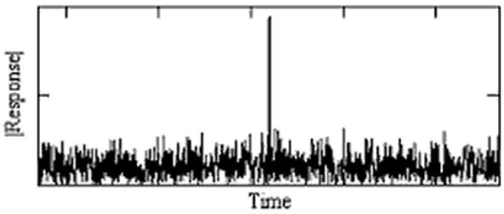

Figure 14 Chirp recovered through Pulse Compression.

Analogue Pulse Compression

Before the advent of modern Digital Signal Processing, pulse compression was achieved by analogue means. The radar transmitter could sweep the frequency of its local oscillator within the pulse duration (shown discretely as F_1 to F_4 in the diagram below). The receiver would then amplify any returned signals and pass them through a device (often a SAW filter) whose delay with respect to frequency matched the bandwidth and sweep characteristics of the transmitted pulse (hence the term "matched filter"). In the example below, the first frequency (F_1) is delayed more than the last frequency (F_4) such that the energy of the transmitted pulse (measured in watt-seconds) is compressed into a shorter time span. The total energy must be conserved and so the peak power increases along with a significant improvement in range resolution.

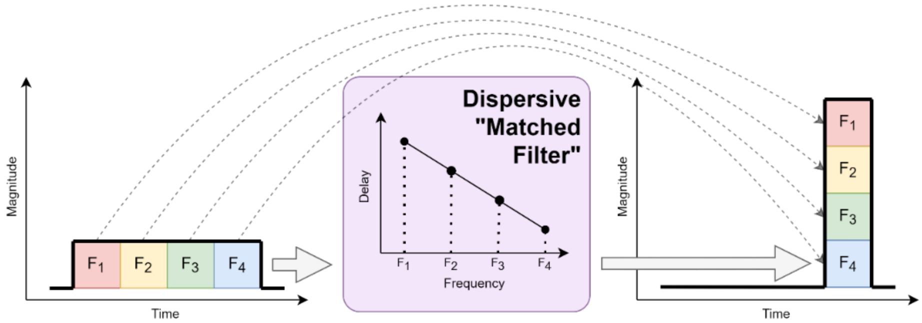

Figure 15 Analogue pulse compression used a matched filter.

Digital Pulse Compression

Pulse compression can now be implemented (with relative ease) using Digital Signal Processing (DSP) techniques. Whilst it is possible to design a perfect "matched filter" in the digital domain, there are other more direct methods that can achieve excellent results.

One solution uses the Fast Fourier Transform (FFT). In this implementation, the transmitted chirp waveform is passed through an FFT and stored in memory. Each block of received samples is passed through another FFT and multiplied by the stored TX chirp FFT. Lastly, the resultant samples are passed through an Inverse FFT to yield correlated Magnitude vs Range with a significant improvement in range resolution and SNR.

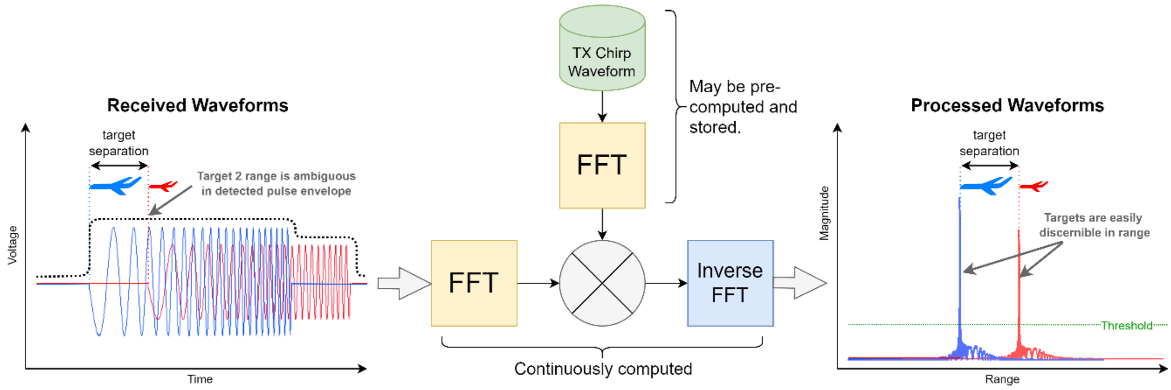

Figure 16 Digital pulse compression using FFT/IFFT.

A short trial of the digital Pulse Compression algorithm was undertaken at very low power in the ISM band. A fixed target was observed as a peak in the returned plot. Analysis using Google Earth identified the target as being a three-storey building wall at exactly 253m.

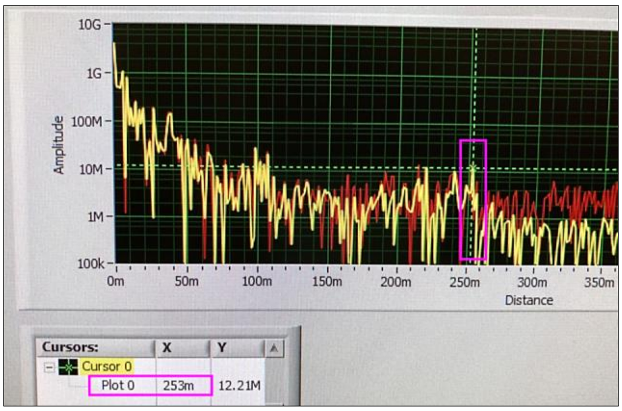

Figure 17 Amplitude vs range with an echo at 253m.

Figure 18 Distance measured using Google Maps.

Radar Range Equation with Pulse Compression

The basic radar range equation can be modified to take into account the compressed pulse width (τ_c vs τ_u) and the required SNR for reliable detection:

$$Range = \sqrt[4]{\left(\frac{P_t \tau_u G^2 \lambda^2 \sigma}{(4\pi)^3 \tau_c SNR_{min} k T_o B F L_{ges}}\right)}$$

Where:

P_t is the peak transmit power.
τ_u is the uncompressed pulse width.
G is the antenna gain (assumes the same antenna is used for both TX and RX).
λ is the signal wavelength.
σ is the radar cross section (RCS).
τ_c is the effective compressed pulse width ($\tau_c = \tau_u/PCR$).
SNR_{min} is the minimum required SNR for detection.
k is Boltzmann's constant.
T_o is the temperature in Kelvin.
B is the effective RX bandwidth (not exactly the LFM sweep BW due to filter mismatch).
F is the Noise Factor.
L_{ges} is the sum of the signal processing and system losses.

It should be noted that the compressed pulse width (τ_c) and the required RX bandwidth (B) are intrinsically related (decreasing τ_c requires that B be increased). However, the correlative receive processing has the ability to operate in a low (possibly even negative) SNR environment, pulling returns out of the noise floor in a way that a traditional pulsed radar cannot. The only caveat is that the gain before the ADC must be set high enough to allow the noise to be sampled with enough resolution to extract the chirp. This can limit the overall dynamic range of the system.

Whilst a pulse compression radar has excellent range resolution, its minimum range is still limited by the transmitted pulse width (i.e. the receiver is blind whilst the pulse is being transmitted). It is for this reason that we still vary the transmitted pulse width depending on the target range of interest.

The minimum (blind) range can be calculated as:

$$Range_{min} = C_0 \left(\frac{\tau_u}{2} + T_{turnaround}\right)$$

Where:

C_0 is the speed of light.
τ_u is the uncompressed pulse width.
$T_{turnaround}$ is the time taken to switch from TX to RX.

The $T_{turnaround}$ time is not to be underestimated. Key contributors include:

- Time to switch the transmit pulse modulator in to maximum isolation.
- Time to turn off the power amplifier.
- Time to switch the TX/RX switch to the LNA path.
- Time to switch the receive pulse modulator in to minimum isolation.
- Any settling times associated with the above.

Calculating RCS

Our radar signal processor will provide us with a magnitude value for each sample in the time domain.

We can reconfigure the Radar Range equation to give RCS as the result:

$$\sigma = \frac{(4\pi)^3 R^4 \tau_c SNR_{min} kT_o BFL_{ges}}{P_t \tau_u G^2 \lambda^2}$$

$$S_{rx} = SNR_{min} \times kT_o BF$$

$$\sigma = \frac{(4\pi)^3 R^4 S_{rx} L_{ges}}{P_t G^2 \lambda^2} \times \frac{\tau_c}{\tau_u}$$

Our digitiser returns uncalibrated I/Q or IF values (S_{adc}, relative to ADC Full Scale). We can use S_{adc} to compute an uncorrected RCS:

$$\sigma_{uncorr} = S_{adc} \times R^4 \times \left(\frac{\tau_c}{\tau_u}\right) \times \left(\frac{(4\pi)^3}{P_t G^2 \lambda^2}\right) \times L_{ges}$$

It would be necessary to scale S_{adc} by a correction factor to get the absolute received power S_{rx}. However, we can instead calibrate the whole measurement system (end to end) by placing a target of known RCS, at a known range and measuring the σ_{uncorr}. A calibration factor can then be calculated using the Reference Target RCS (σ_{ref}) value for each Pulse Compression Ratio utilised:

$$A_{PCR} = \frac{\sigma_{ref}}{\sigma_{uncorr}}$$

This then corrects the above σ_{uncorr} equation to give us absolute RCS:

$$\sigma_{corr} = S_{adc} \times R^4 \times \left[A_{PCR} \times \left(\frac{\tau_c}{\tau_u}\right) \times \left(\frac{(4\pi)^3}{P_t G^2 \lambda^2}\right) \times L_{ges}\right]$$

Note that everything to the right of R^4 (in square brackets) is essentially constant (for a given pulse compression ratio) and so we can quickly and easily compute corrected RCS at all ranges with minimum additional processing overhead.

Predicted Performance

The following table describes the predicted performance of the radar with some measured parameters assuming Linear Frequency Modulation (LFM) pulse compression:

	Units	Pulse Width / LFM Bandwidth		
Pulse width	**s**	**5.0E-6**	**1.0E-6**	**100.0E-9**
Max PRF (at 10% Duty Cycle)	**Hz**	20,000	100,000	1,000,000
TX to RX Turn Around Time	**s**	100.0E-9		
Linear FM Sweep BW	**Hz**	70.0E+6	80.0E+6	90.0E+6
Pulse Compress Ratio	**linear**	350	80	9
Effective Compressed Width	**s**	14.3E-9	12.5E-9	11.1E-9
TX Power (at the feed)	**W**	3		
Antenna Gain	**dBi**	25		
Antenna Gain	**linear**	316		
Total System Loss	**linear**	1.0		
Frequency	**Hz**	9.25E+09		
Wavelength	**m**	3.24E-02		
Target RCS	**m2**	1		
Noise Figure	**dB**	2.7		
Noise Factor	**linear**	1.862		
Temperature	**°K**	290		
Min Signal to Noise Ratio*	**dB**	8		
Min Signal to Noise Ratio	**linear**	6.31		
CW Pulse Range Resolution	**m**	750	150	15
LFM Pulse Range Resolution	**m**	2.1	1.9	1.7
Min (blind) Range	**m**	**780**	**180**	**45**
Max Range (at Min SNR)	**m**	**2,028**	**1,356**	**763**

Figure 19 Predicted performance with different pulse widths and intra-pulse modulation bandwidths.

As previously alluded to, it is possible to operate a pulse compression radar in an environment with a lower SNR than 8dB (chosen here as the typical Tangential Signal Sensitivity threshold for pulse detection). However, comparisons with a simple pulsed radar then become impractical, so we shall assume a reasonable (positive) SNR is required.

Three transmission pulse widths (100ns, 1us and 5us) were chosen to cater for targets at short, medium and long ranges, whilst the sweep bandwidth was adjusted to maintain an acceptable range resolution throughout.

System Topology

This instrumentation radar is very much a system-of-systems built around the scanner antenna. One chassis contains all the x-band converter electronics and another all of the motion control electronics (for both azimuth and elevation control). These are located on either side of the scanner support frame.

Atop the reflector sits and wireless visible light video camera with zoom lens, and a laser range finder.

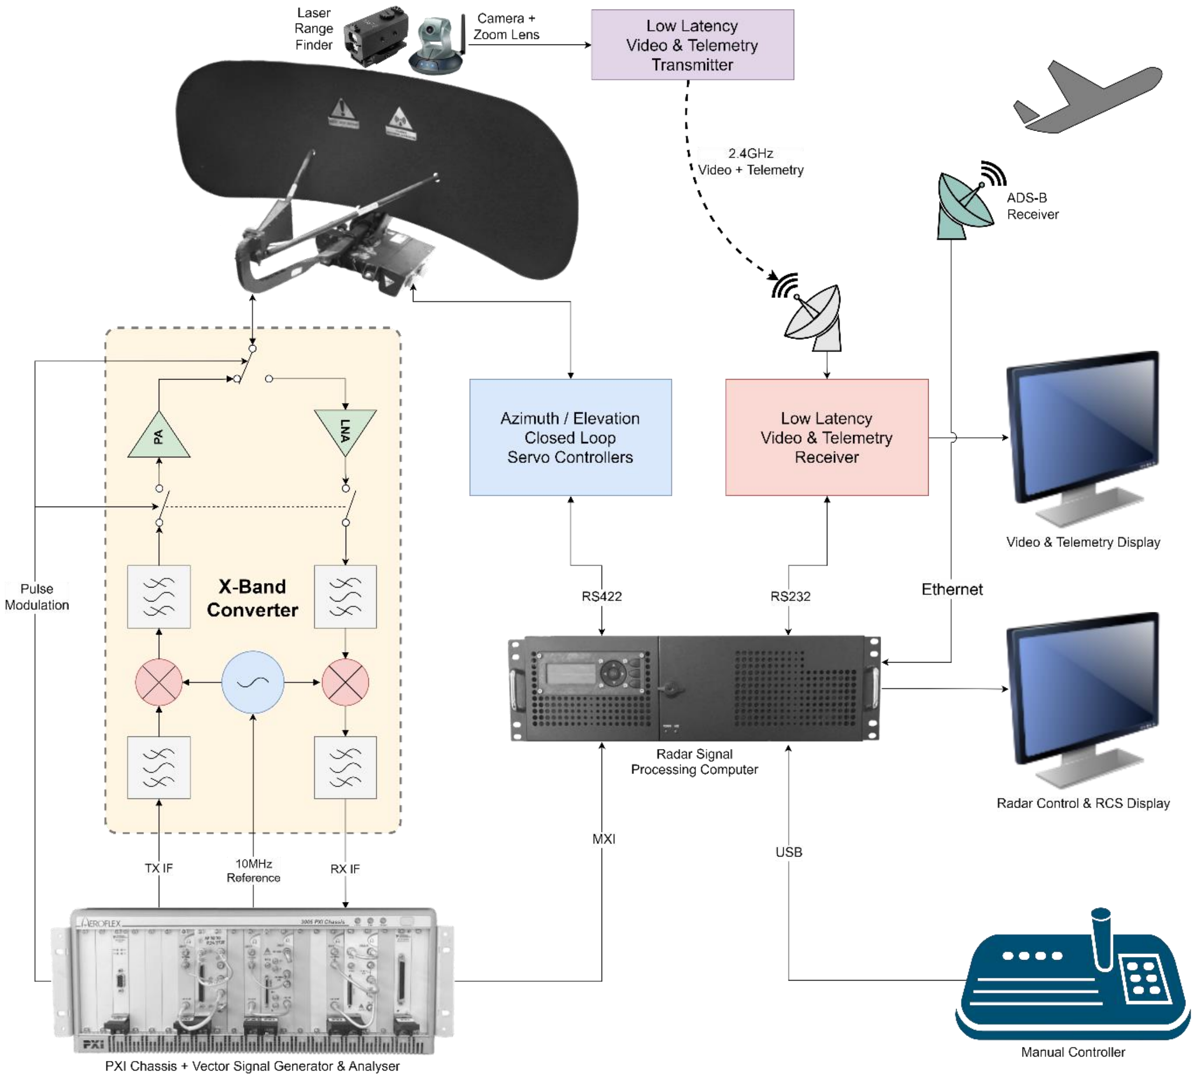

Figure 20 System topology.

An umbilical cord of cables connects the microwave electronics and motion control hardware to a 19" rack chassis containing the PXI transceiver and pulse processing PC. The PXI chassis connects to the processing PC's PCI-Express bus using an MXI link allowing high-bandwidth data transfer.

Two flat panel displays provide low latency video, telemetry and the main Graphical User Interface for the instrumentation radar (developed in LabVIEW).

An external ADS-B/AIS/Software Defined Radio receiver provides positioning data for airborne and marine targets.

Scanner

The scanner consists of a 1.1m x 0.4m reflector, horizontally polarised waveguide feed horn, rotary waveguide joint, and azimuth/elevation servos & resolvers.

Figure 21 MEL Super Searcher radar scanner.

It is a recycled antenna that once sat atop a Sea King Anti-Submarine Warfare helicopter (serial number ZA137, which last flew in 2015). This was part of the MEL Super Searcher radar system fitted to the HAS5 variant and would have been located in the dorsal radome behind the engines in the below photograph.

Figure 22 ZA137 Westland Sea king WS-61 HU5 at RNAS Culdrose. Picture Credit: Andrew Thomas

After servicing the antenna components, input match was measured through the waveguide, rotary joint and feed horn and found to be better than -12dB from 9.1GHz to 9.7GHz.

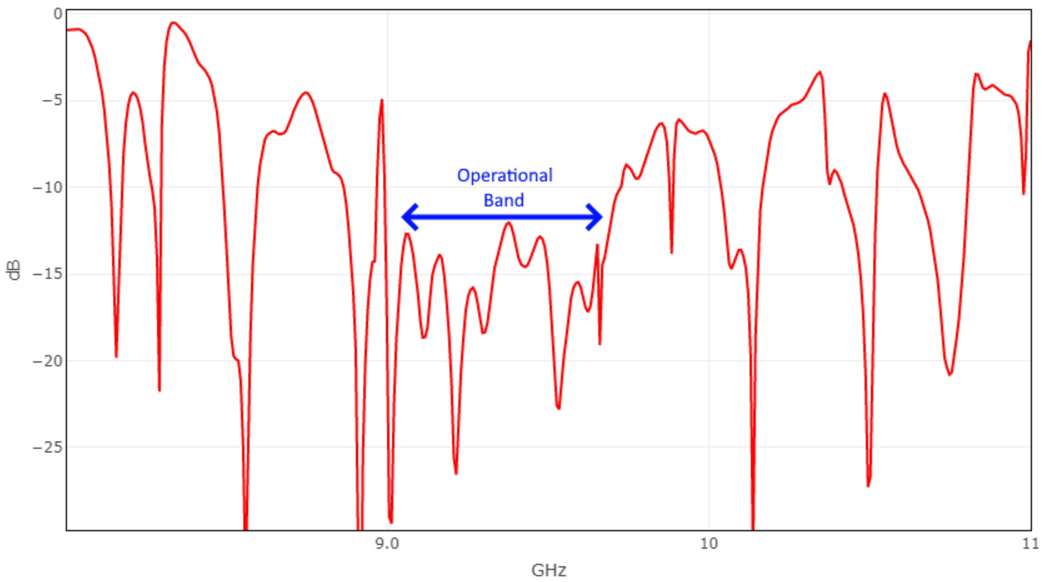

Figure 23 Measured S_{11} of scanner feed and reflector.

Timothy Pelham [1] (Senior Research Associate, Bristol University) has developed LyceanEM [2], a Python library for rapid virtual prototyping of antenna arrays and frequency & time domain channel modelling. LyceanEM allows the user to model a wide array of complex problems from antenna array architecture, to assessing beamforming algorithm performance and channel modelling. The model is built upon a ray tracing approach and allows for efficient modelling of large, low-density spaces. Timothy has produced an excellent paper demonstrating the computational performance advantage of LyceanEM versus a full electromagnetic solver (CST). CST had a peak memory usage of 13.4GB and took over 37 hours to complete, whereas LyceanEM consumed 5GB of memory and took just 109 seconds to complete. The comparative results plotted below use an experimental meshing algorithm and show impressive agreement given the size/complexity of the model and the order of magnitude faster computational time achieved with LyceanEM.

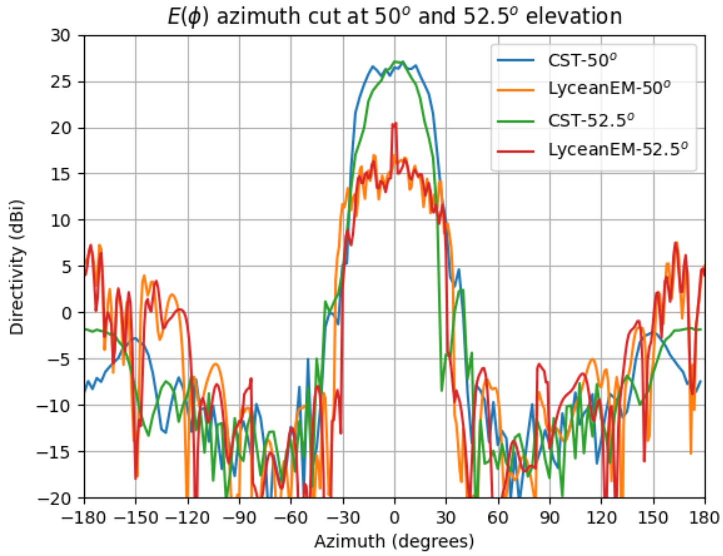

Figure 24 Comparison modelling of directivity using LyceanEM & CST. Picture Credit: Timothy Pelham.

A simplistic estimation of the gain of the reflector based on aperture alone is circa 38dB. The published literature on the Sea Searcher radar indicated an antenna gain of 34dB. The reduced gain (circa 27dBi) estimated in the LyceanEM & CST models could be attributed to many factors including non-optimal illumination of the reflector by the feed horn. Physical measurements were therefore required to qualify the above gain estimates and so with the help of James Pawson at Unit 3 Compliance, a brief measurement of calibrated gain was performed.

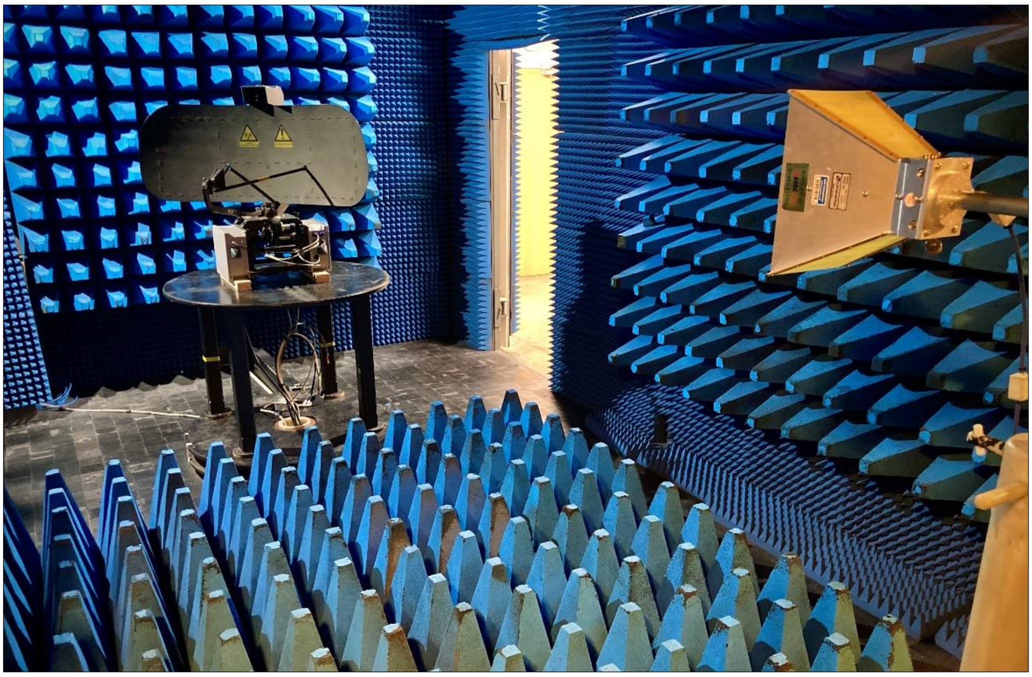

Figure 25 Physical gain measurements at Unit 3 Compliance.

Custom LabVIEW software was developed to characterise the scanner. The intensity plot shown below was taken with the azimuth servo varied ±20° and the elevation servo varied ±10°. The peak gain was measured to be circa 27dBi depending upon frequency and azimuth/elevation angle.

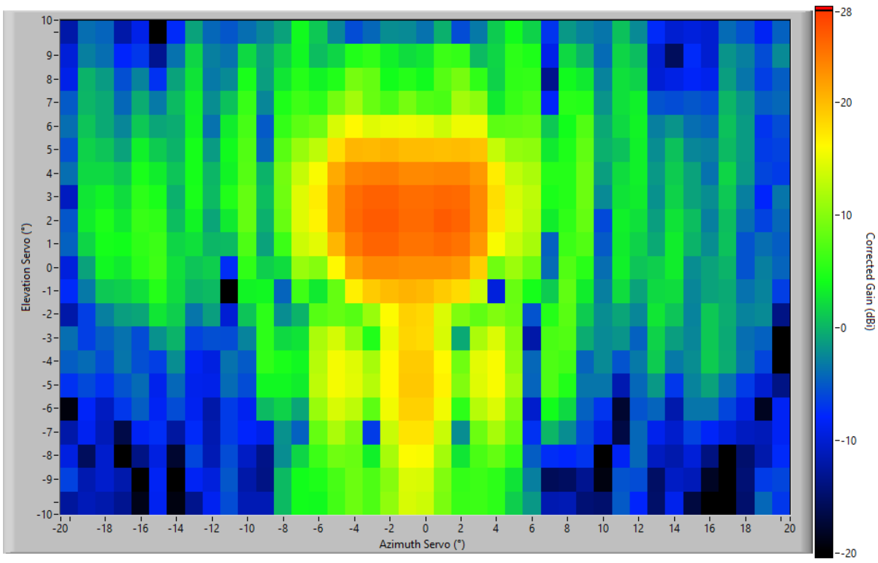

Figure 26 Measured gain vs azimuth & elevation servo angle.

It was noted that the main lobe had split in to a “doublet” resulting in a gain ripple of ±1dB and so this will be an area for future investigation and optimisation.

X-Band Converter

All of the microwave hardware is located in a single chassis close to the waveguide interface on the scanner to minimise cable losses (2.2dB/m at 9GHz). From an RF perspective, it has SMA IF inputs and outputs, a 10MHz reference input, and an X-band RF input/output. An additional blanking input is provided via a BNC connector. This is controlled by the adjacent servo driver system which can be configured to disable transmission over a specified azimuth angle range.

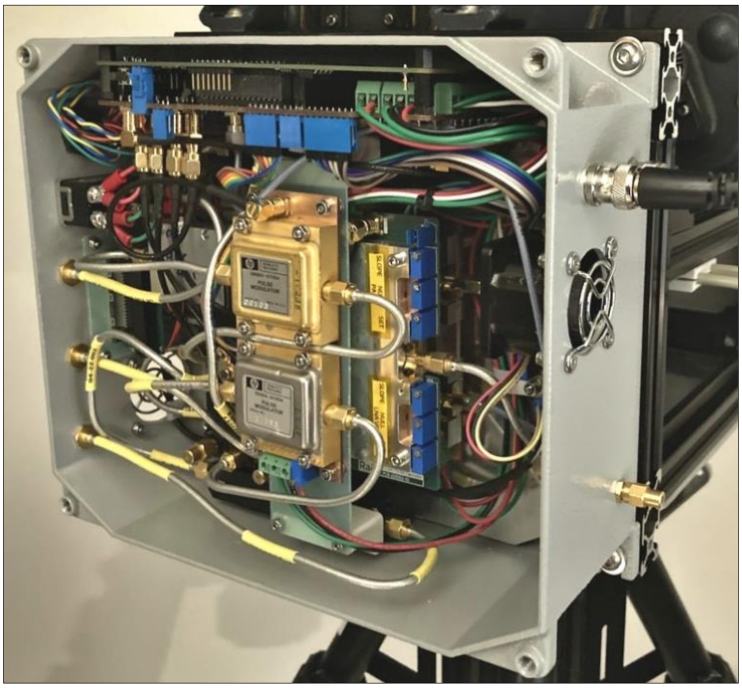

Figure 27 X-band converter chassis.

RF Front End

After a review of the available semiconductor solutions, we settled upon Qorvo's QPM1002 which combines a 4W GaN Power Amplifier, Low Noise Amplifier with 2.2dB noise figure and T/R switch in to a single 5mm x 5mm package. A very impressive level of integration.

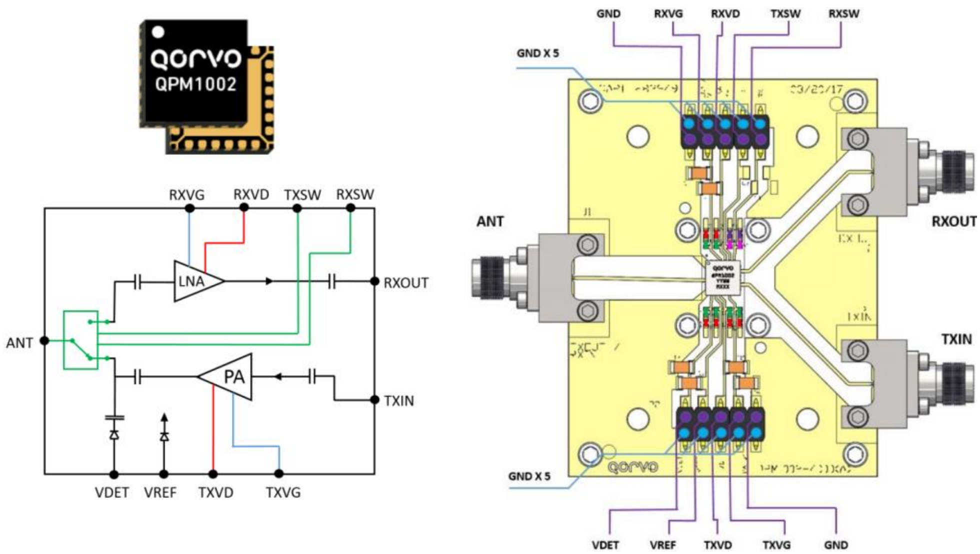

Figure 28 Qorvo QPM1002 3W GaN PA/LNA/Switch.

A custom PCB was designed using the Qorvo evaluation board as a reference. To ensure adequate isolation from the TX to RX path, an additional GaAs switch was included with T_{ON} - T_{OFF} time of 6ns.

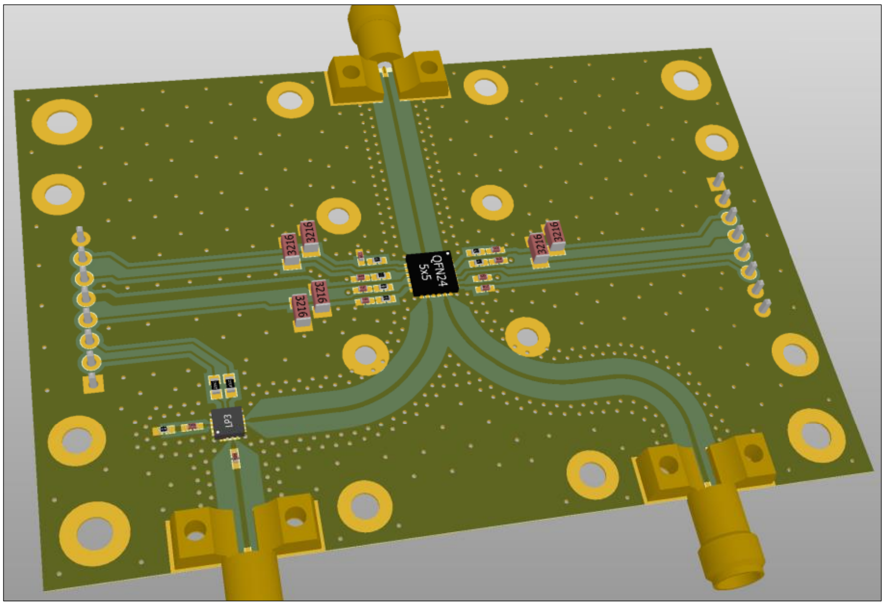

Figure 29 RF front end PCB assembly.

In keeping with our project's aims, we standardised the surrounding metalwork across the various microwave elements of the RF front end. A generic Aluminium base plate and cover were developed with support for up to six edge launch SMA connectors, adequate fixings for thin RF laminates, and a provision for an FR4 DC/control PCB on the rear of the assembly. Templates were produced in the PCB design software allowing painless customisation for each application. A generic DC breakout board was produced to allow quick "bring-up" of RF PCBAs without the need to initially also produce custom DC/Control PCBAs (shown below right).

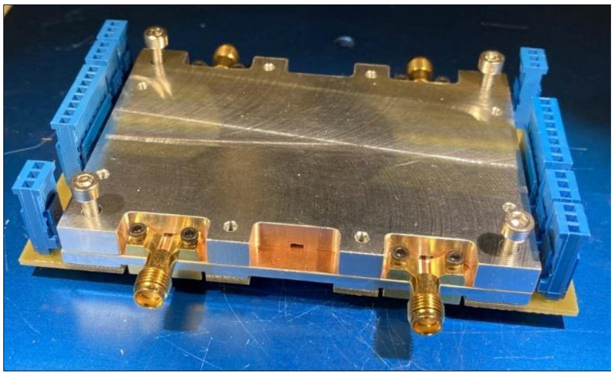

Figure 30 Generic RF assembly.

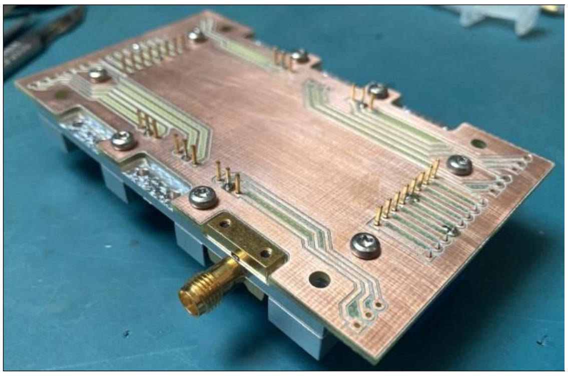

Figure 31 Generic DC break-out PCB assembly.

All the RF PCBAs needed for this project were produced in-house from bare Rogers 4350B 10 thou (0.254mm) RF laminate on an LPKF PCB milling machine. Holes and vias were activated and then electroplated using a bright copper process (See Appendix A for more details on the process used).

The lack of a solder mask to control wicking requires some design considerations and compromises. We hope to develop a solder mask process to add to our repertoire.

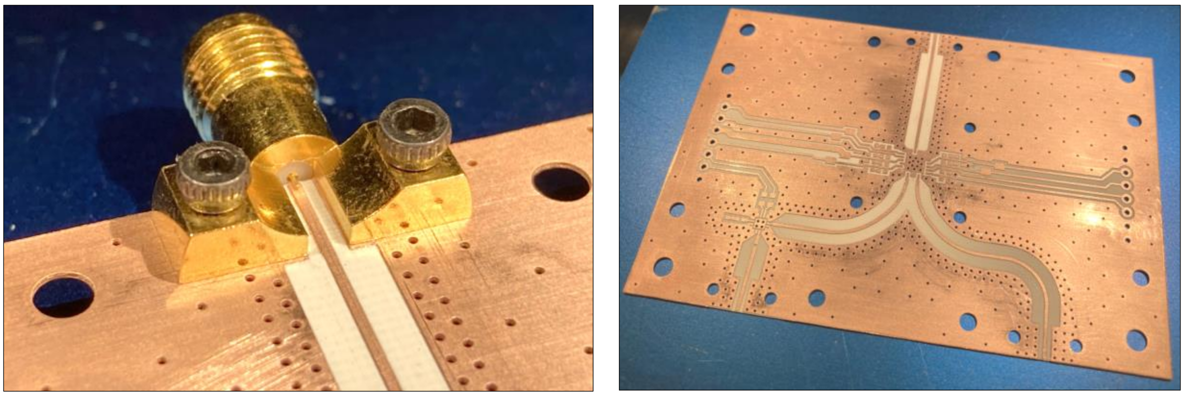

Figure 32 X-band PA/LNA/Switch PCB after machining and plating prior to population.

The RF front-end PCBA was initially tested with the generic DC/Breakout PCBA attached and bench power supplies generating all the bias and control voltages.

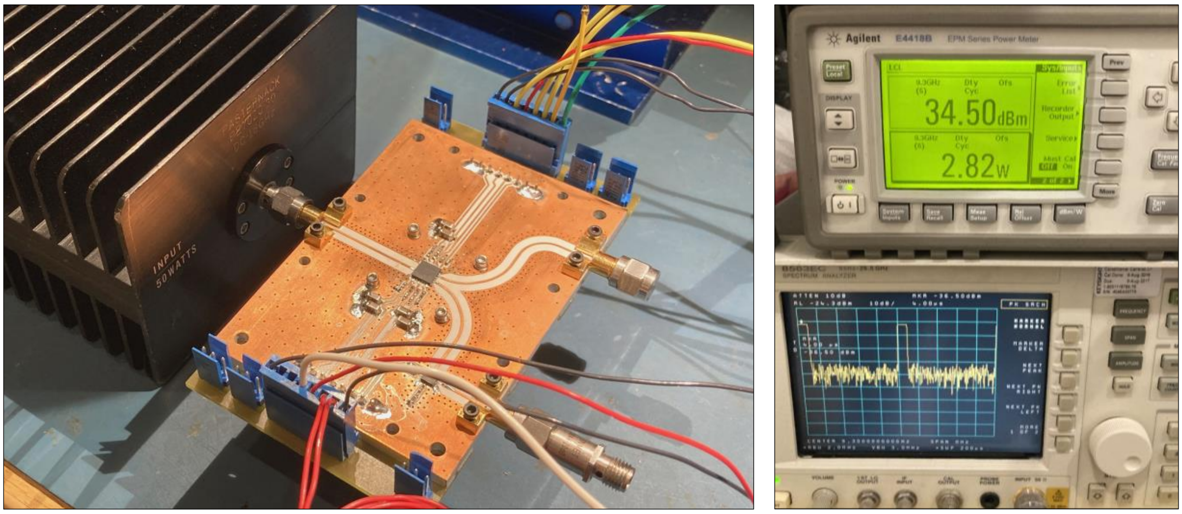

Figure 33 X-band PA undergoing pulsed RF power testing.

The QMP1002 datasheet notes the requirement to turn off the PA drain voltage supply when in receive mode. The TX/RX switch is the reflective type and so the PA has a tendency to oscillate when presented with an open-circuit load. This introduces an interesting challenge as we must turn the drain voltage off as quickly as possible after completing transmission of a pulse before switching to the receive path, otherwise, we could be blinded by our own oscillations and face EMC issues.

At first inspection this might seem straightforward, however, switching a +25V well-decoupled supply in circa 150ns yields a slew rate of nearly 170V/us. In practice, the PA gain also falls with the drain voltage and so we do not have to reach 0V before the close-loop gain falls below the level at which oscillation is supported.

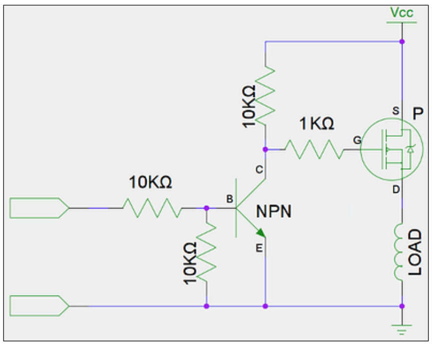

Figure 34 An overly simplified drain switch.

The current handling capability of a MOSFET is proportional to its semiconductor geometry/area. Therefore, high current devices have a larger gate-source capacitance (C_{GS}). Simply placing a P-channel MOSFET in the drain rail with an NPN driver would not be able to charge/discharge C_{GS} fast enough to achieve the required T_{ON} & T_{OFF} times.

With care, a discrete push-pull driver could be developed to serve this function, however, this is not without time-consuming challenges. Qorvo provides several useful presentations on the considerations of PA, LNA & Switch biasing and drain pulsing. One proposal makes use of an integrated high-side MOSFET gate driver that can sink/source 100-200mA whilst having T_{rise}/T_{fall} times less than 100ns. The gate driver also includes a bootstrap circuit to generate the high gate bias voltage needed to fully turn on an N-channel Power MOSFET operating on the high-side.

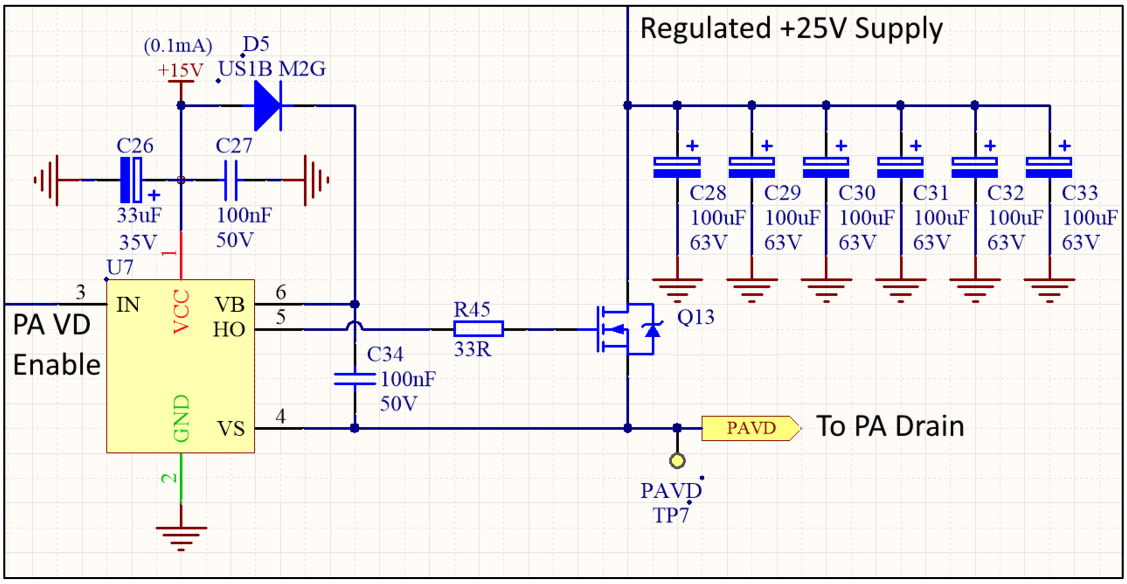

Figure 35 Drain pulse modulator.

Initial experiments with the circuit driving the PA demonstrated less than ideal drain voltage pulse fidelity (left plot, green trace). Further investigation revealed the decaying oscillation on the drain supply was due to the large low-frequency decoupling capacitor on the PA side of the drain switch. This was removed since there was already enough local decoupling present on the nearby drain voltage fast transient linear regulator. This change resulted in a fast-switching drain voltage with little to no ringing and a clean decay to 0V (right plot, green trace).

Figure 36 Drain voltage (green) with too much decoupling.

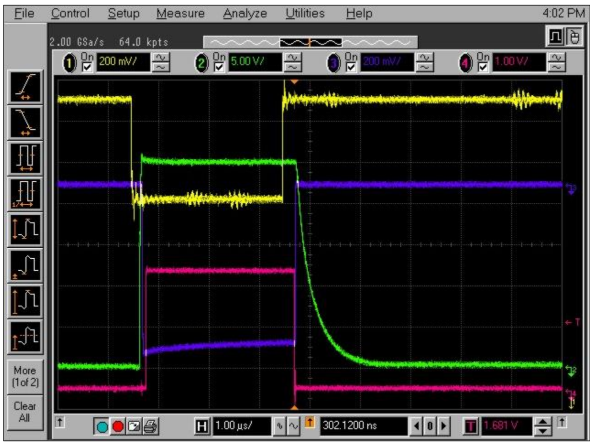

Figure 37 Drain voltage (green) with optimal decoupling.

It must be noted that High-Side Gate Drivers, whilst having very fast rise and fall times, have relatively large propagation delays (200-300ns is not uncommon) and so it is necessary to switch the driver well in advance of when the transition is actually required. How this is achieved in our system will be discussed later.

Gate bias generation, temperature compensation, and drain bias interlocking are important considerations for depletion mode GaN devices. The QPM1002 is no exception with the bias voltage needing to be driven more negative as the temperature increases to maintain a constant drain current. Again, Qorvo provides a starting point for circuit design with a solution proposed that can sink/source gate current and provide adjustment of offset & slope. The circuit shown below uses the linear temperature dependence of a transistor's PN junction to provide compensation. The first Op-Amp acts as a summing amplifier, whilst the second amplifier provides a high current buffer with a discrete output stage.

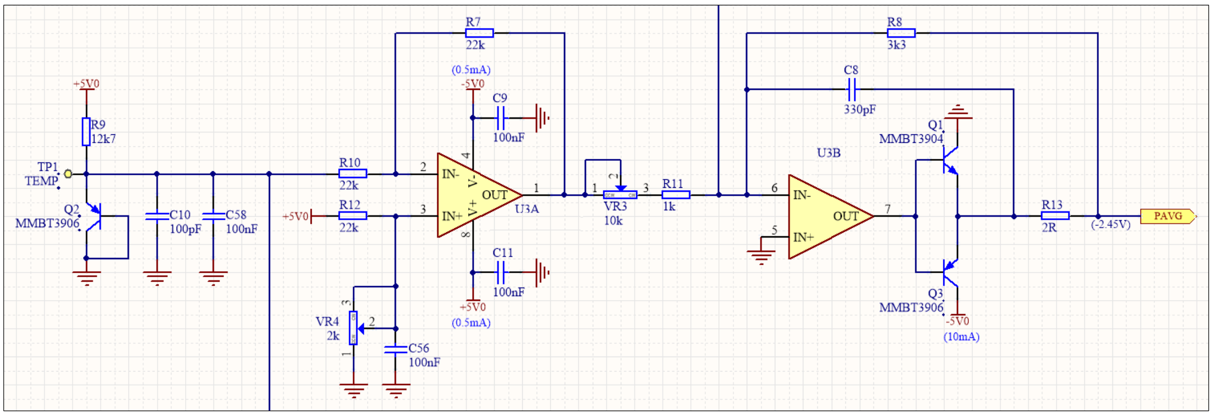

Figure 38 Gate bias generation and linear temperature compensation.

The QPM1002 integrated TX/RX switch is a welcome addition that significantly reduces the microwave circuit development time. However, its driver circuit must be capable of switching quickly between 0V and -28V to minimise radar "blind" time. The circuit developed (shown below) trades power dissipation for simplicity. Cascaded logic gates provide complementary signals whilst the Zener diodes provide the necessary voltage offset to drive the P-channel MOSFETs. The 1k ohm 1W resistors provide a pull-down to the negative rail.

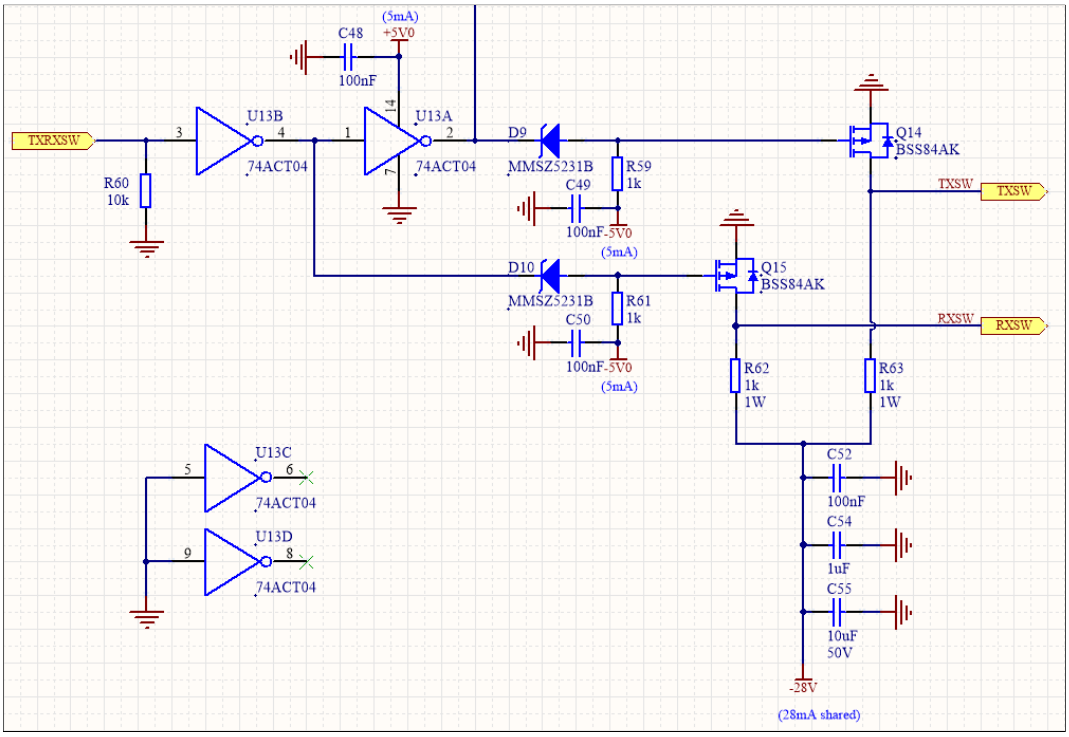

Figure 39 GaN TX/RX high-voltage switch driver.

The compact size and high level of integration of the QPM1002 are very attractive, though the real estate requirements of the accompanying bias and control circuits should not be underestimated! The developed bias & control PCBA demonstrates this point exactly (shown below).

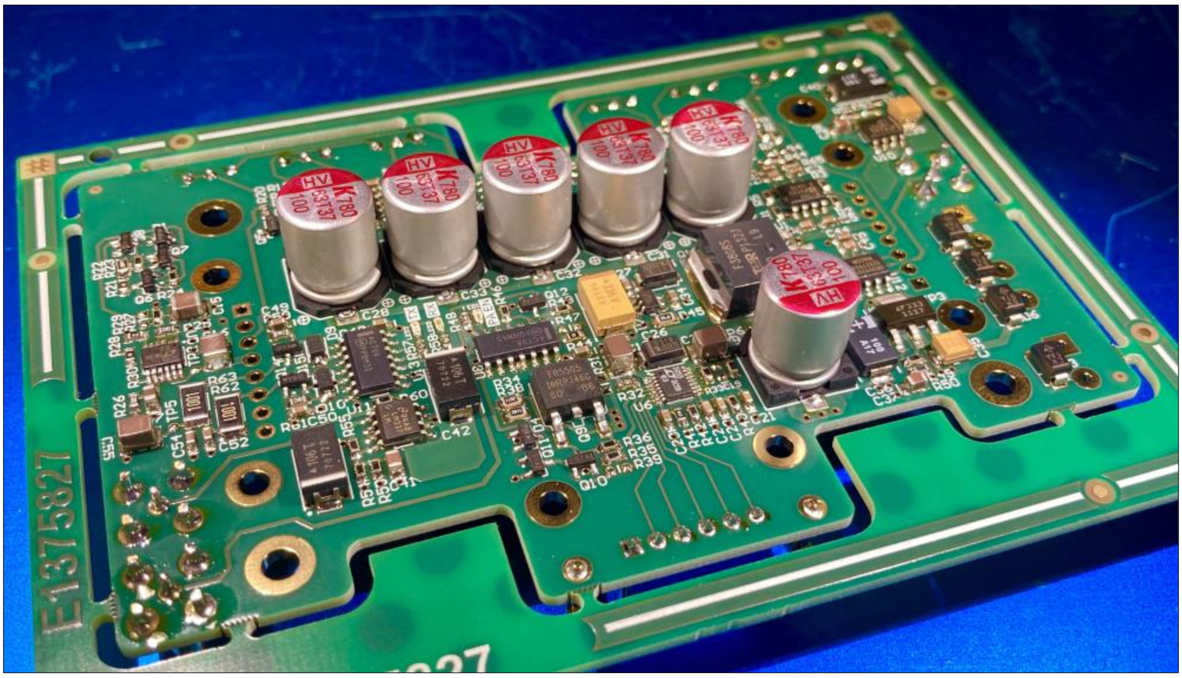

Figure 40 PA/LNA/Switch bias & control PCB assembly.

At the time of designing the PA drain switching circuits, the required pulse width and duty cycle had not yet been determined and so it was decided that, when it comes to local bulk charge storage, “if some is good, more is better”.

Dual Output LO Synthesizer

To allow coherent processing of the transmitted and received pulses, a common Local Oscillator (LO) frequency synthesizer was designed around an MMIC VCO and an integrated Phase Locked Loop.

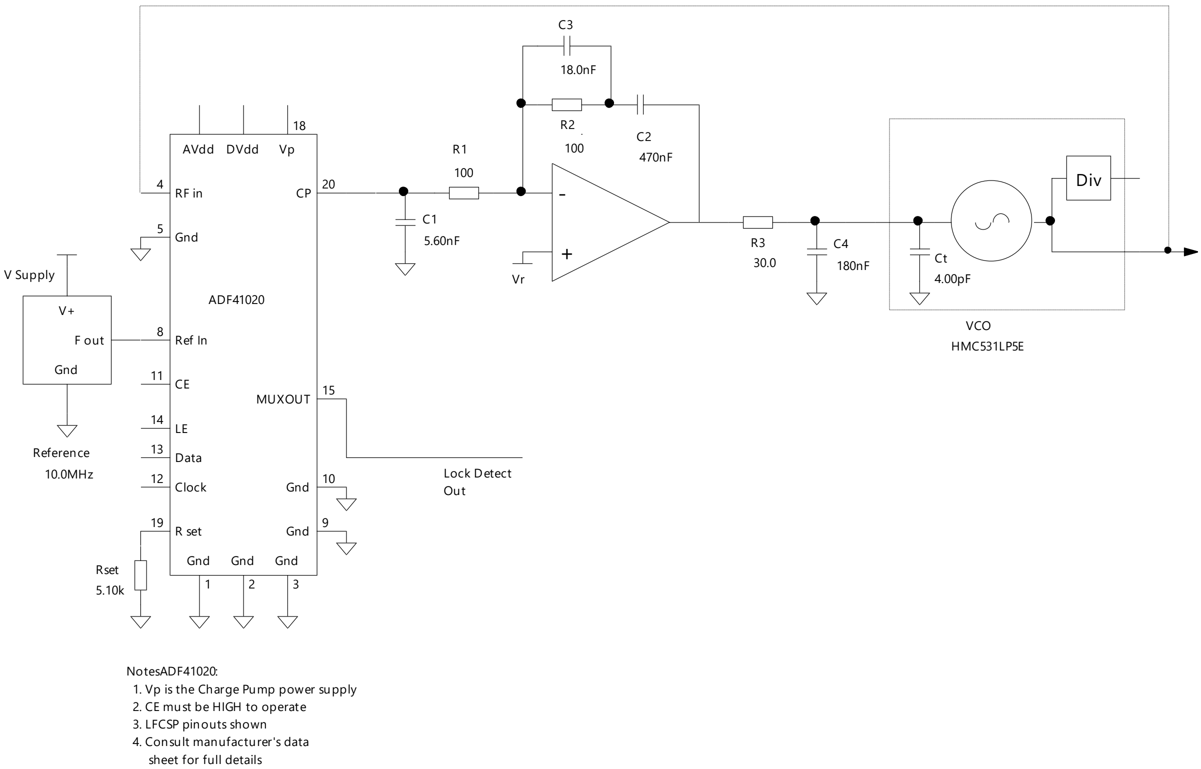

Figure 41 LO synthesizer modelling using ADiSimPLL.

The loop is closed around the doubled VCO output (circa 14.3GHz) whilst the regular VCO output (circa 7.15GHz) drives a splitter and two LO driver amplifiers operating near saturation. The divided VCO output is switchable and is routed out to an external connector for test purposes.

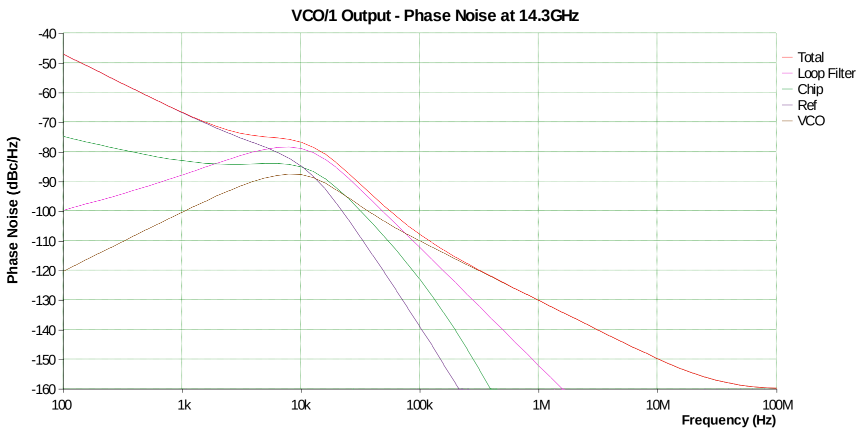

Figure 42 Estimated phase noise profile at 2xLO frequency.

The phase noise of the utilised output is at half the synthesizer frequency and so experiences a $6\mathrm{dB} = 20\log_{10}(2)$ improvement and so we can expect to achieve around -81dBc/Hz at 10kHz offset from carrier. Likewise, the divided output at a quarter of the frequency should achieve a figure of -87dBc/Hz. A basic measurement of phase noise at this test output using a Spectrum Analyser verifies this assertion.

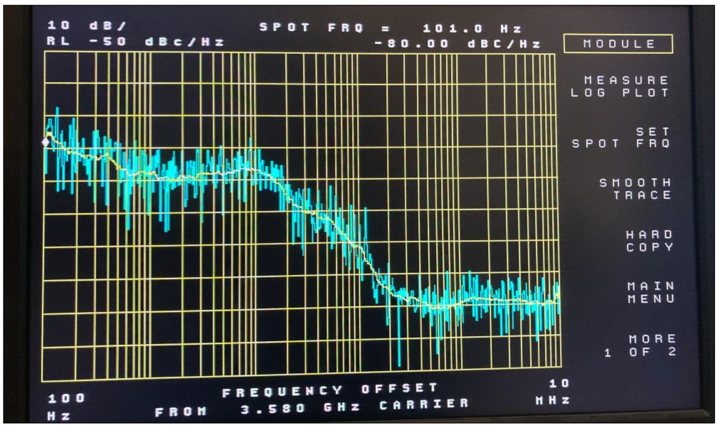

Figure 43 Phase noise measured at the divided output (LO/2).

Once again, the LO synthesizer and its associated DC/Control PCB were manufactured in-house using the developed templates significantly accelerating realisation.

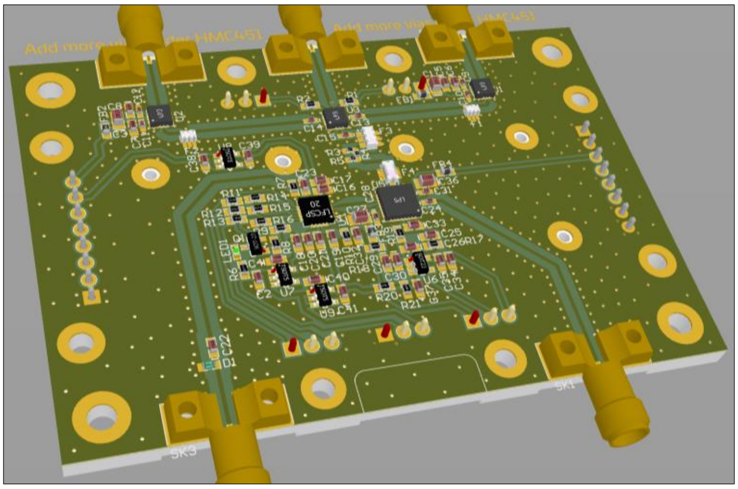

Figure 44 Dual output LO synthesizer PCB assembly model.

Initially it was assumed we would need to switch the LO between the two LO driver amplifiers to maximise TX to RX path isolation. However, the two LO driver amplifiers tended to oscillate at around 16GHz when not selected as the LO path (despite the SPDT switch terminating the unused LO path). Further investigation determined the cause to be too few ground vias under the LO amplifiers.

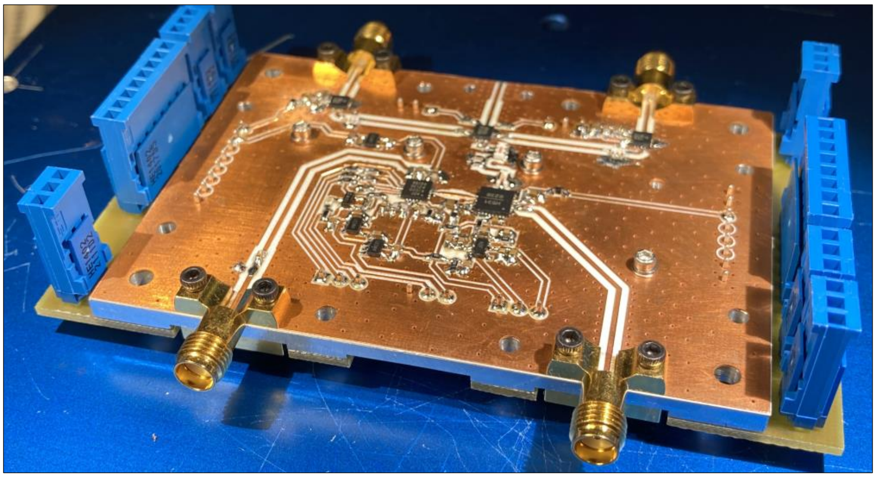

Figure 45 Dual output LO synthesizer assembled in the generic housing.

After undertaking system level testing, it was found that LO switching was unnecessary and so a Knowles DLI ceramic power splitter was fitted in place of the SPDT switch. This along with the judicious use of Radiation Absorbent Material (RAM) calmed the amplifiers' propensity to oscillate.

RF Pulse Modulator

In the past, it was possible to purchase single chip RF pulse modulators, but such parts don't appear to be readily available nowadays. Not wanting to commit additional time towards designing a discrete circuit, we instead recycled a pair of pulse modulators used in the Hewlett Packard 8684B RF Signal Generator (of 1982 vintage).

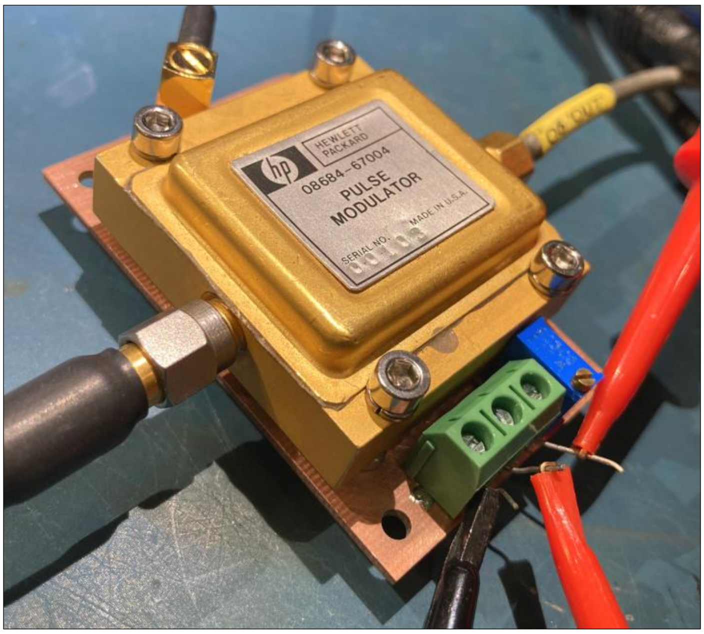

Figure 46 Hewlett Packard PIN pulse modulator.

The RF modulator is PIN based, has isolation greater than 80dB and requires a bipolar voltage swing to switch between the on and off states. The original discrete driver schematics were sourced from the test equipment service manual. These allowed us to determine the required on and off state voltages and to estimate of the required sink/source currents (typically 25-50mA).

An integrated driver solution was preferred and the MA-COM MADR series of PIN modulator drivers are hard to beat in terms of price, real estate, and level of integration. They have a single TTL input, can generate complementary bipolar outputs, and have a short propagation delay and fast slew rate needed to generate high fidelity RF pulses.

Frequency Conversion & Filtering

Identical connectorised triple-balanced mixers are used for up and down conversion. A triplexing filter on the RF port of the mixer performs both the function of a pass band defining filter (from 8-12GHz) and a good broadband match for out of band mixing products (from 4-8GHz and 8-20GHz). This is achieved by terminating the unused ports with 50 Ohm loads. The same frequency plan is used for both and up and down conversion.

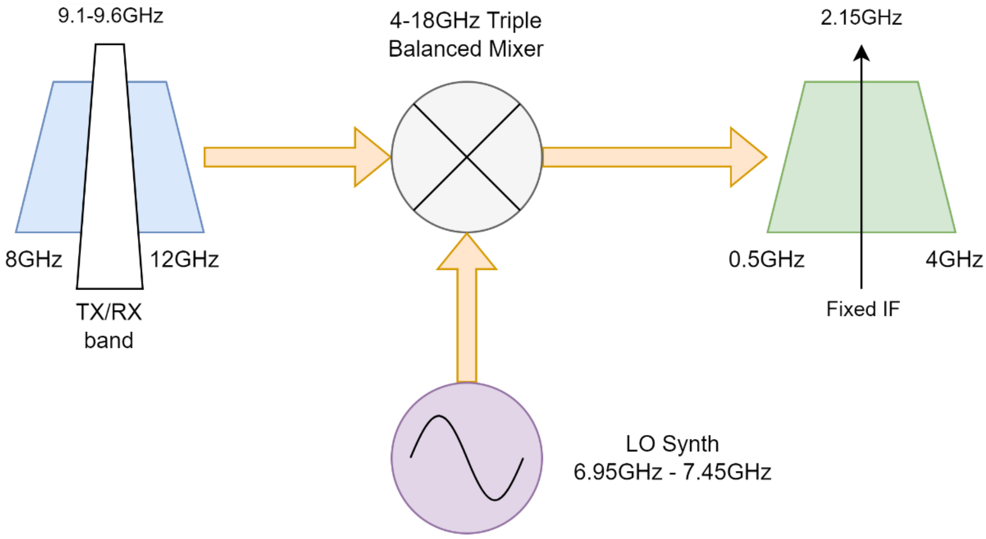

Figure 47 Frequency conversion plan.

The IF spectrum is defined by a connectorised 0.5GHz to 4GHz band pass filter. Combined with the triplexer's 4GHz wide RF passband and programmable LO frequency, the chain provides a great deal of flexibility for other applications of the transceiver. If out of band interference becomes a problem, an external connectorised narrow-band IF filter can be fitted between the X-band converter and the PXI RF digitiser.

The filters utilised in this system were also recycled parts coming from the RACAL MIR-2 "Orange Crop" ESM Receiver previously fitted to a Westland Lynx HAS2 (XZ255). The filters were originally produced at Filtronic Components in Shipley (Now Teledyne Defence & Space).

Figure 48 XZ255 Westland Lynx HAS2 Fleet Air Arm Royal Navy. Picture Credit: Crown/Royal Navy.

Power Supply

The system supply voltage is 28V (to match the Servo Drive enclosure). This is fed to several DC/DC converters to produce the required positive and negative rails.

Each active RF element is powered by a local low-noise linear regulator. The PA has a particularly fast transient response linear regulator with a programmable current limit and monitoring feature.

Control & Monitoring

A Microchip PIC based microcontroller is used to provide control of the LO Synthesizer frequency and monitoring of the phase-lock state, alarms, consumed currents and analogue voltages.

Communication with the system controller is via a multi-drop RS422 bus that is shared with the Servo Driver & Feedback module. This is both more resilient to noise and faster when compared to regular RS232. A Rathera proprietary protocol is used across all devices with each one uniquely identified with an 8-bit address. The protocol can be efficiently implemented in low-power microcontrollers and uses a 16-bit CRC to ensure packet integrity.

Timing

The radar transceiver has a number of key components that must be accurately switched and sequenced to generate a high-fidelity chirped pulse as short as 100ns:

- TX PIN Modulator
- TX RF Isolation Switch
- PA VDRAIN Switch
- TX/RX Switch
- RX PIN Modulator

Each of these components has different rise times, fall times, and in some cases rather lengthy propagation delays. These control signals must also maintain synchronicity with the LFM chirp generated in the PXI chassis. Creating custom control logic in the microwave enclosure for this purpose would be a significant challenge that would also limit the flexibility of the transceiver. Such logic would likely have to be implemented in an FPGA operating at up to 100MHz and be synchronised to the IQ generation clock.

Fortunately, there is a feature in the Aeroflex 3025 PXI RF Signal Generator that can be brought to bear on this problem. The IQ memory in the 3025 is 32 bits wide and mega samples deep. The I & Q waveforms used to the modulate the transmit IF are 14 bits wide each (28 bits in total). The remaining 4 bits are allocated as waveform synchronous “markers” that can be routed out of the LVDS front panel connector. We transfer these signals to the microwave enclosure, convert them from LVDS to single-ended logic levels, and then use them to drive the different system elements.

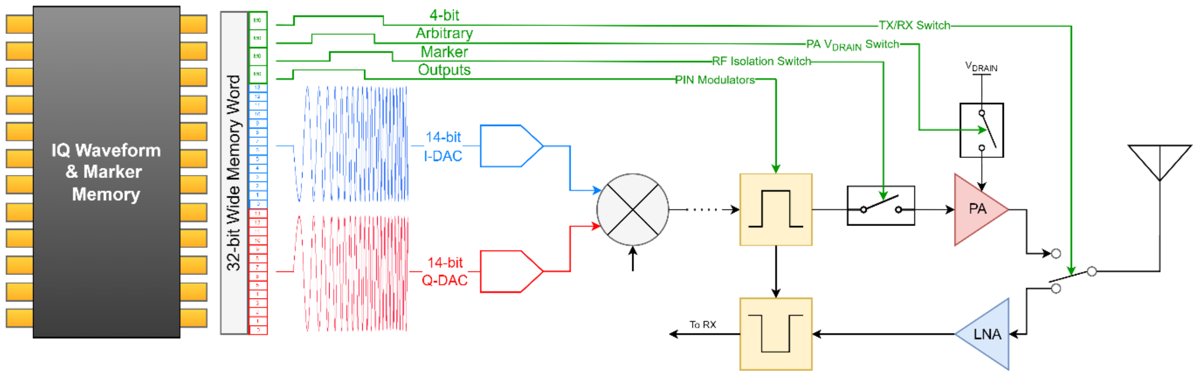

Figure 49 IQ synthesis and control signal mapping to LVDS markers.

Aeroflex provides a Dynamic Linked Library for waveform generation that allows markers to be specified alongside the I & Q waveforms. Custom LabVIEW software running on the control PC synthesizes the IQ waveforms along with the marker positions. Through careful inspection of the transmitted waveform envelope and adjustment of pre- and post-transmission delays, a high-fidelity pulse with a perfectly aligned chirp can be generated.

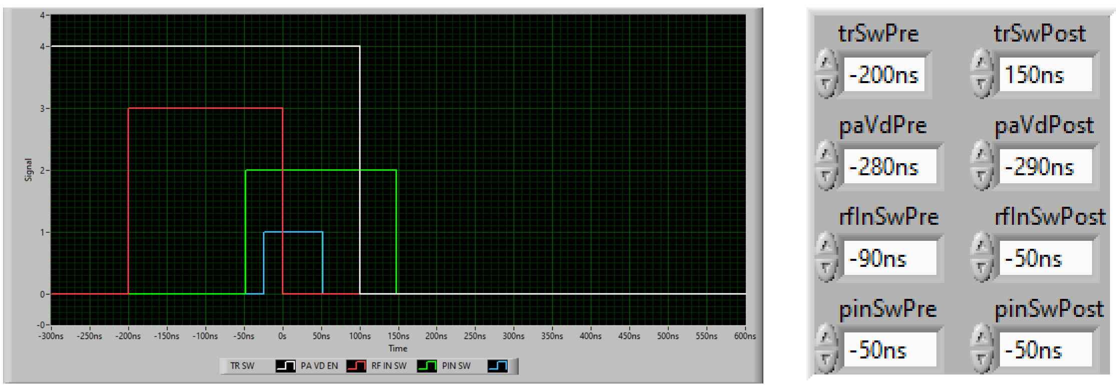

Figure 50 pre- and post-transmission delays of switch & modulator control signals.

The generated IF was recorded using an Aeroflex 3035C RF analyser and processed in both the time and frequency domain to verify the transmit pulse fidelity and chirp alignment.

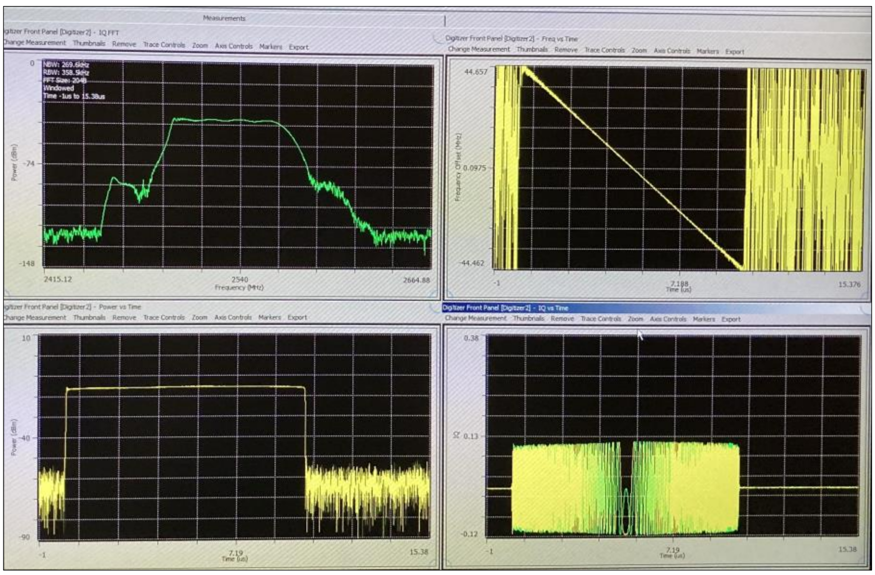

Figure 51 Testing chirp alignment to the pulse envelope.

Motion Control System

All of the motion control electronics are located in a single chassis close to the AC servos & resolvers, adjacent to the microwave enclosure. A temperature controlled 40mm fan provides cooling.

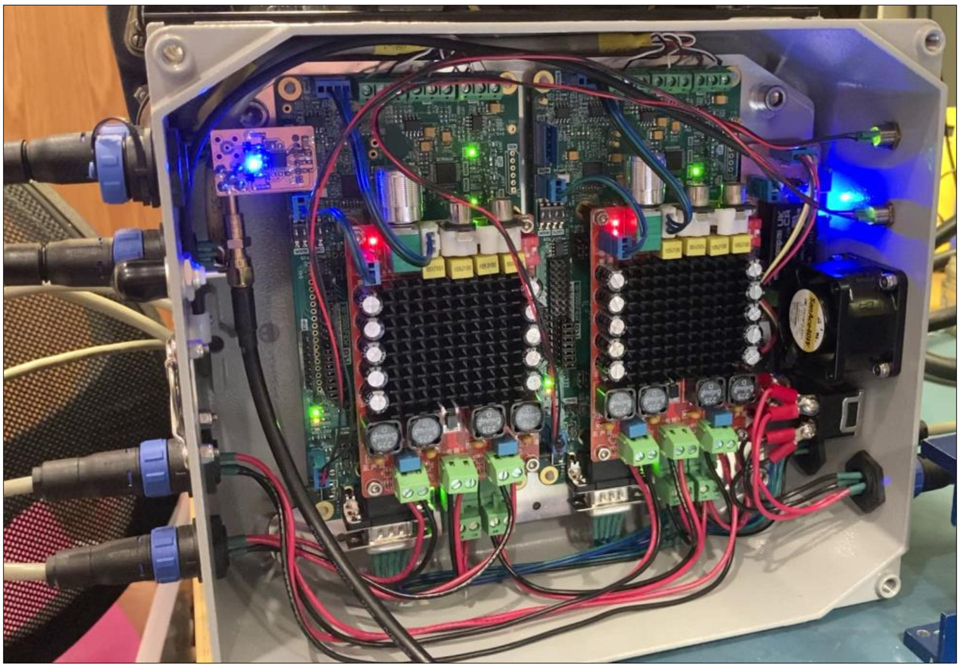

Figure 52 Motion control system chassis.

Servo Drivers

The azimuth & elevation axes of the scanner are driven by two-phase AC servos through a gear train. No information could be found on the electrical or mechanical drive characteristics of these parts (likely custom-made for MEL) and so significant experimentation was required to understand their behaviour.

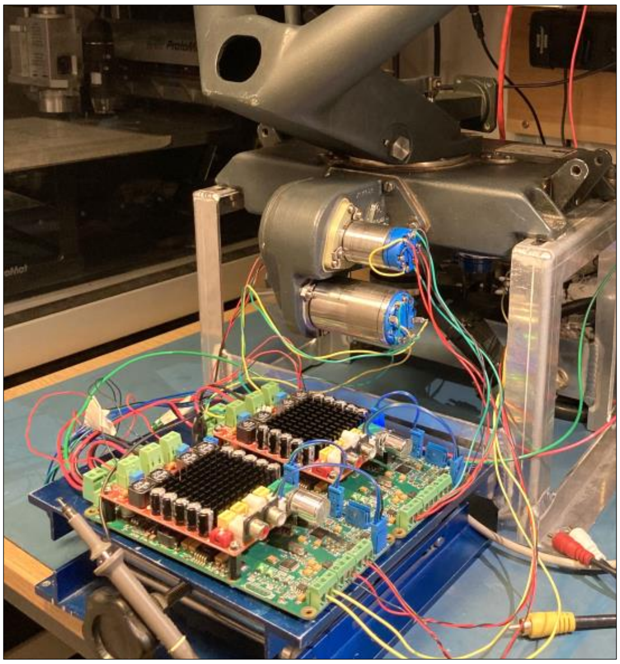

Figure 53 Testing the motion control system with the azimuth & elevation axes.

It was found that the servos can operate from two sinusoids of equal frequency (circa 10Hz - 80Hz). Beyond 80Hz, the maximum available torque decreases significantly. The relative phase of the drive sinusoids determines the direction of rotation and the torque applied. At 28Vpeak, the servos responded as required with minimal self-heating. Given the established drive characteristics and the need for two sinusoids, we decided to utilise an off-the-shelf stereo Class-D audio amplifier module modified to extend its lower frequency limit down to 10Hz and to increase roll-off above 1kHz.

Rather than developing a custom dual sinusoid generator using a DSP microcontroller, we determined it would be expedient to utilise a two-channel Direct Digital Synthesizer (DDS), but drive it with a much lower frequency clock than typical. The DDS can be programmed with amplitude, frequency, and phase, which it will continue to generate without external intervention, relieving the microcontroller from this real-time task.

Feedback

The scanner uses two Muirhead Aerospace resolvers to provide continuous analogue azimuth and elevation angle feedback. The resolvers deliver accuracy in the hundredths of a degree (or better) but are fabulously expensive.

Figure 54 Muirhead resolvers.

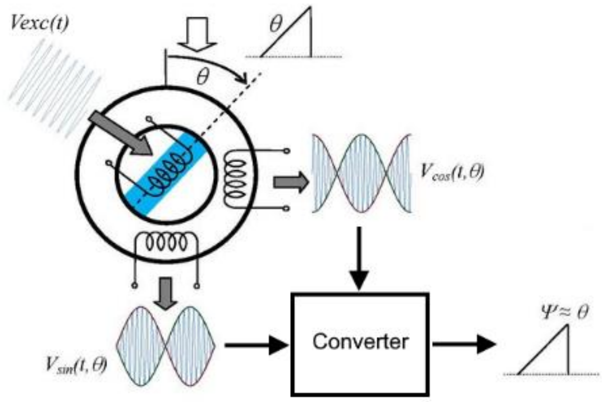

Figure 55 Resolver excitation and output signals.

A resolver is excited with a fixed-frequency sinusoid from which it generates $V_{\sin}(t,\theta)$ and $V_{\cos}(t,\theta)$ outputs whose relative amplitude and phase are translated in to a continuous angle measurement.

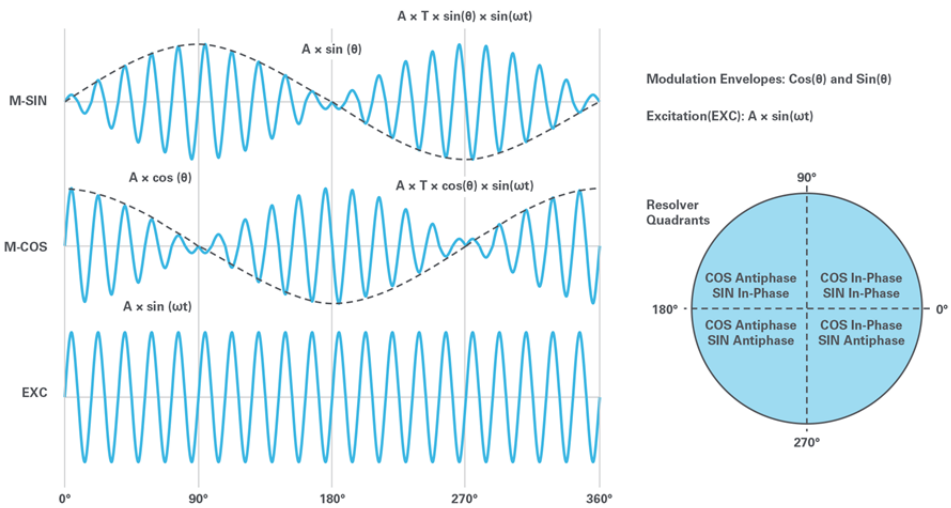

Figure 56 Derivation of resolver quadrant from signal phase.

Relatively complex analogue processing circuits are required to interpret the analogue output of resolvers and provide angular position and velocity feedback. Thankfully digitisers now exist specifically to interface resolvers with more modern digital systems. These provide continuous measurement of angle, velocity, and system health, all accessible over a fast SPI bus.

Servo Controller

A single-axis controller was developed with a fast microcontroller, DDS, resolver digitiser and buffering to drive an attached class-D amplifier daughterboard.

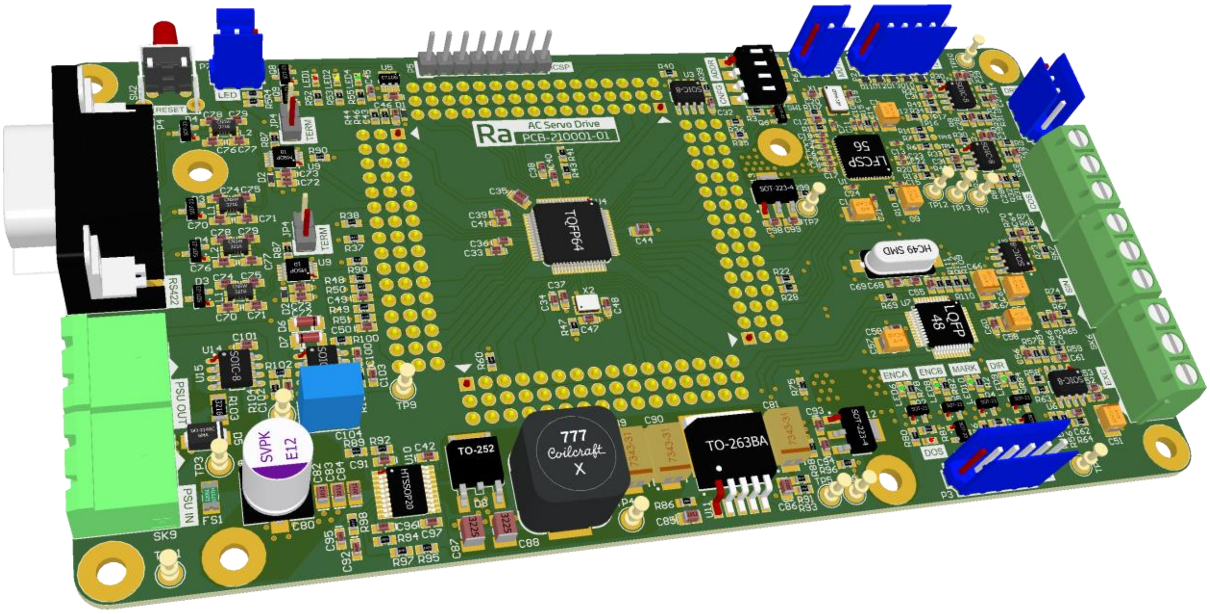

Figure 57 Servo controller PCB assembly.

A common RS422 multi-drop bus connects between the radar processor PC and the two axis controllers with each having a different address on the bus.

Enclosures & Core Frame

The enclosures chosen for the microwave and motion control sub-systems were custom-modified ILME APV-19 [3] parts. These are die-cast Aluminium, very strong and IP66/67 protected.

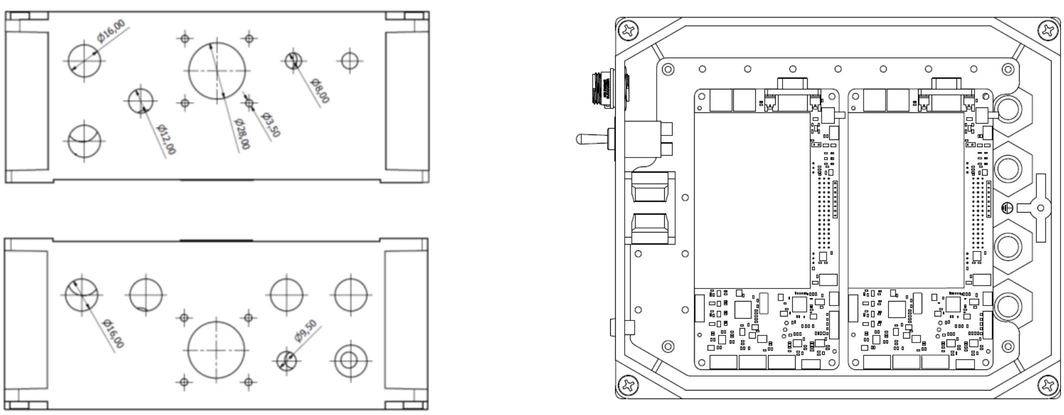

Figure 58 CAD modelling of APV-19 chassis & components for modification.

The core frame to which the scanner, two enclosures, and tripod attach is assembled from 20mm extrusion.

PXI System

Conversion between the RF/IF and digital domains is performed by an Aeroflex 3025C Digital RF Signal Generator and a 3035C Wideband RF Digitiser installed in a PXI chassis.

Figure 59 PXI chassis and pulse processor computer in 19" rack.

PXI (PCI eXtensions for Instrumentation) is a standard which combines the PCI bus with a trigger and synchronisation infrastructure across a backplane. The standard includes electrical, mechanical, and thermal characteristics for compliant modules which can occupy one or more “slots” in the chassis. Control of the of the modules is either performed by a local PXI controller (a compact rugged computer) or an external computer via a high-speed serialised PCI bridge. Each PXI module appears as a device on the computer’s PCI bus. It is therefore important to ensure that the computer motherboard and BIOS can support the significant number of PXI devices which may be attached. Since we may require significant computational resources (including GPU acceleration), an external rack-mounted computer is used with an MXI-Express interface to bridge between the PCI & PXI bus.

Figure 60 Aeroflex 3025C Digital RF Signal Generator.

Figure 61 Aeroflex 3035C Wideband RF Digitiser.

Aeroflex 3025C Digital RF Signal Generator

The RF signal generator operates from 1MHz to 6GHz and includes an integrated dual-channel arbitrary waveform generator with 2GByte of waveform memory. It provides output power control from –120 dBm to +6 dBm and a modulation bandwidth of up to 90MHz. A front panel port provides optional LVDS I/Q data input and marker outputs. The RF signal generator requires an accompanying LO module operating from 1.5 – 3.0GHz. A 10MHz reference can be supplied either by the LO module or daisy chained across several PXI modules for the coherent operation of several signal generators and digitisers. Commencement of generation can be triggered either by software, a dedicated front panel input or via the PXI backplane.

Aeroflex 3035C Wideband RF Digitiser

The RF digitiser operates from 250kHz to 6GHz with an IF bandwidth of up to 90MHz. The digitiser down-converts the bandwidth of interest and samples it with a 250MSps 14bit ADC in to a local 512Mbyte deep memory. Commencement of sampling can be triggered either by software, a dedicated front panel input or via the PXI backplane. As with the RF signal generator, an external LO module is required and there are various options for sourcing the requisite 10MHz reference.

Radar Processor

The radar processor is a rack-mounted Intel 64-bit computer of very modest horsepower. It serves as the central hub for all the connected sub-systems and performs a number of keys roles including:

- Configuring all the PXI hardware, reference clocks, samples clocks, local oscillators, and trigger busses between the RF signal generator and RF digitiser.
- Monitoring of the health of all sub-systems.
- Computing & synthesizing the chirp waveform and transferring it to the RF signal generator sample memory.
- Triggering the transmission of trains of IF pulses and sample acquisition by the receiver.
- Performing processing of the returned samples through the Pulse Compression algorithm along with correction for absolute RCS determination.
- Tracking of the motion control system on to targets both manually and automatically using the ADS-B data feeds.
- Processing of telemetry data from the electro-optics sub-system including target differentiation using the laser range finder.
- Presentation of the received target echoes on cartesian and polar plots.
- Recording of large amounts of sample, telemetry, and target data for After Action Review.

All of the “code” executing on the computer has been developed in LabVIEW. This was chosen for its wealth of connectivity options, ease of integration, and the simplicity with which multi-threaded high-performance applications can be developed.

Health Monitoring

The Motion Control and X-Band Converter modules (both attached to the RS422 multi-drop bus) have built in analogue voltage, current, temperature & status monitoring.

Figure 62 X-band converter health monitoring & status tab.

Motion Control

The Motion Control System can operate in Manual, Joystick, PID or Automatic mode.

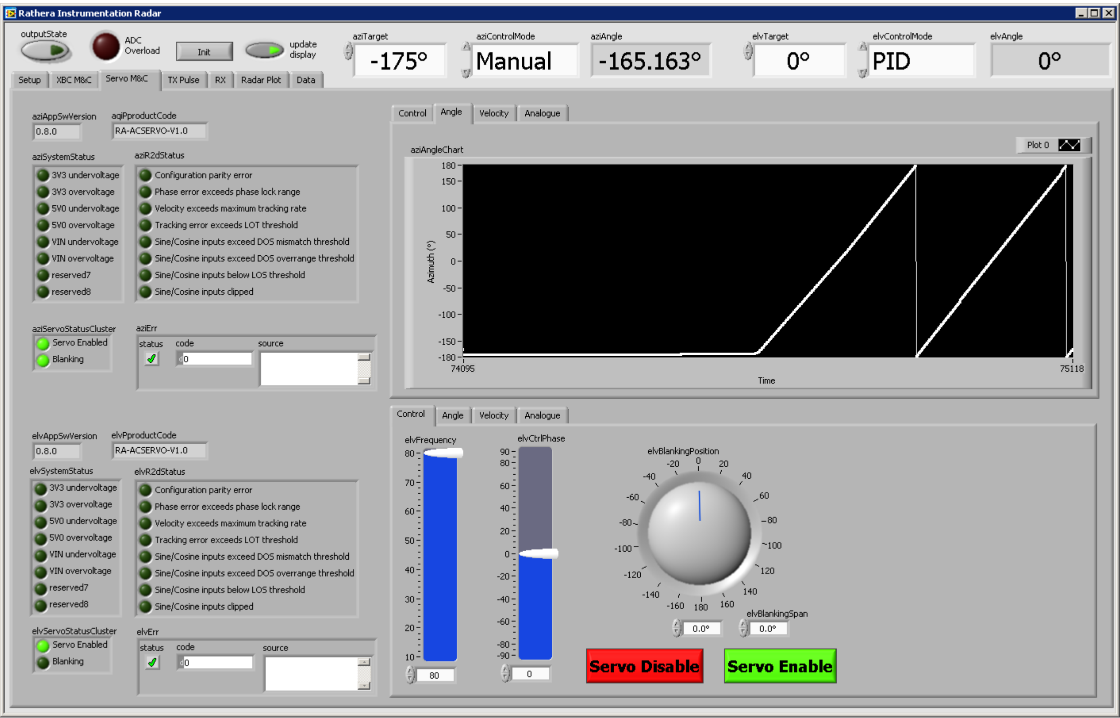

Figure 63 Servo monitoring & control tab.

Transmit Waveform Synthesis

The transmit pulse waveform is synthesized along with the markers used for TX/RX Switch, TX/RX PIN Pulse Modulator, RF Isolation Switch and PA VD Enable control. The Linear Frequency Modulation (LFM) chirp bandwidth is specified as a fraction of the overall RF Signal Generator bandwidth (whilst observing Nyquist). Windowing is applied to the transmit waveform before later processing in the receive pulse compression algorithm.

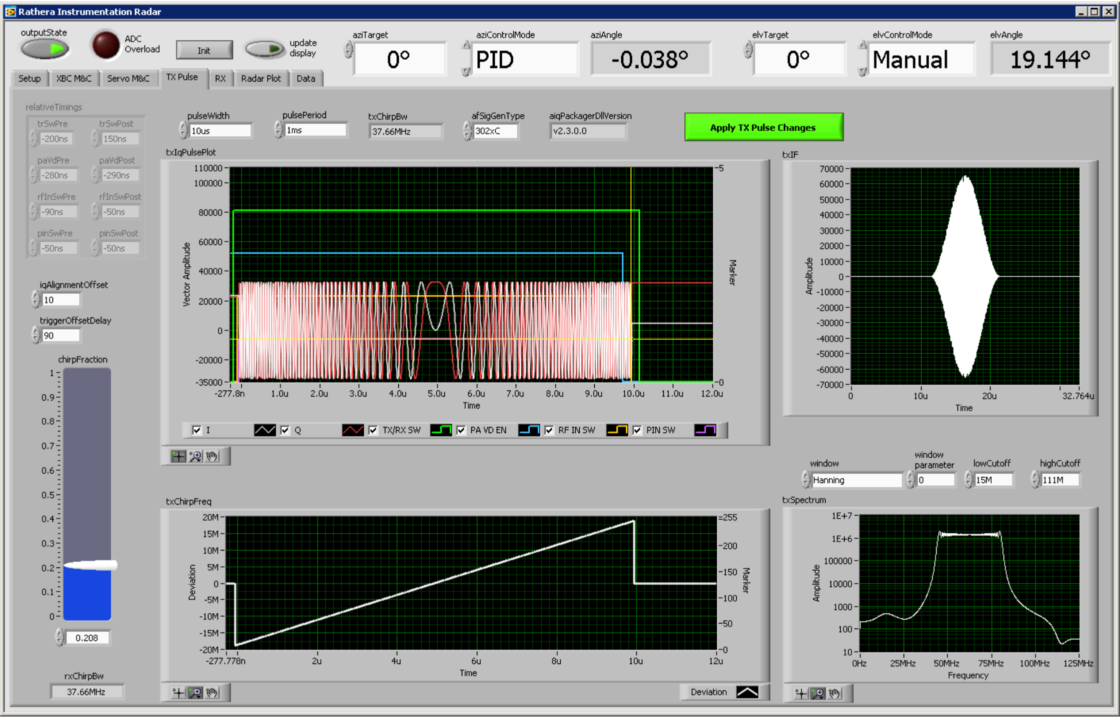

Figure 64 TX waveform synthesis and marker generation.

Controls are provided for fine adjustment of the chirp timing and receiver sampling trigger with respect to the transmitted pulse envelope.

The complex waveform and timing markers are "packed" using functions provided in an external Aeroflex Dynamic Linked Library (DLL) specific to the Digital RF Signal Generator model utilised. The packed waveform is automatically uploaded in to the Digital RF Signal Generator memory for continuous playback, leaving all the PXI backplane bandwidth available for transferring the received samples.

Receive Processing

The received sample blocks are displayed as linear/logarithmic power vs time (top plot) along with being passed through the pulse compression algorithm for improved range resolution (bottom plot).

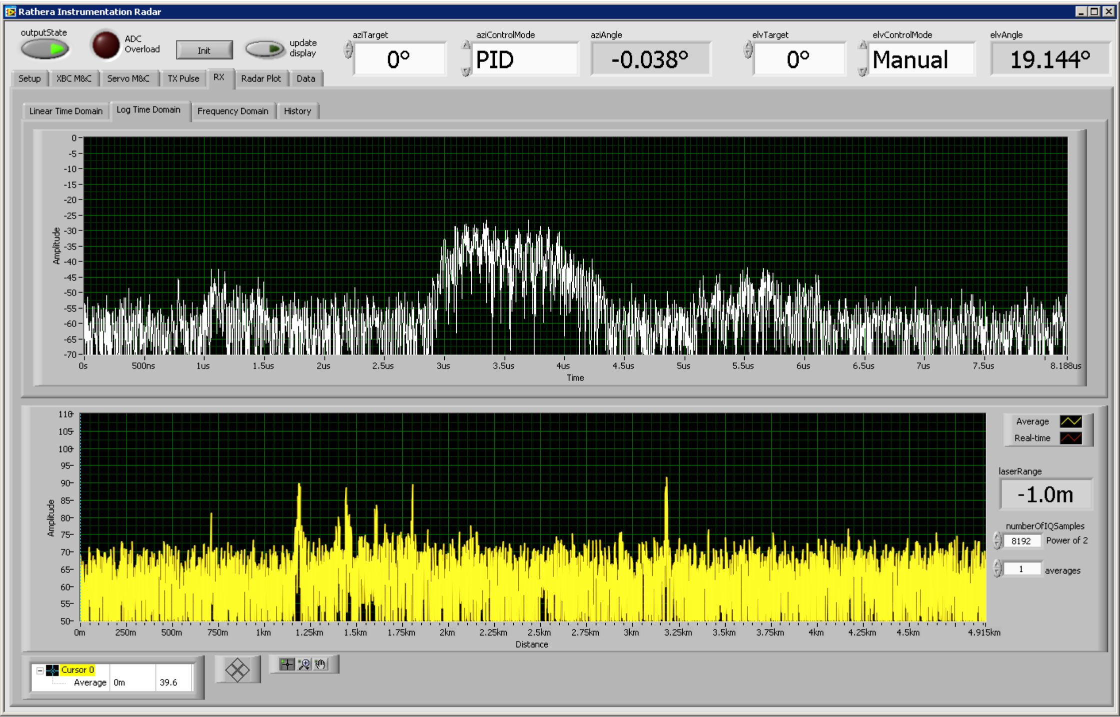

Figure 65 Received signal processing tab.

Cartesian Radar Plot

A 2D intensity plot shows the echo amplitude vs azimuth angle and range.

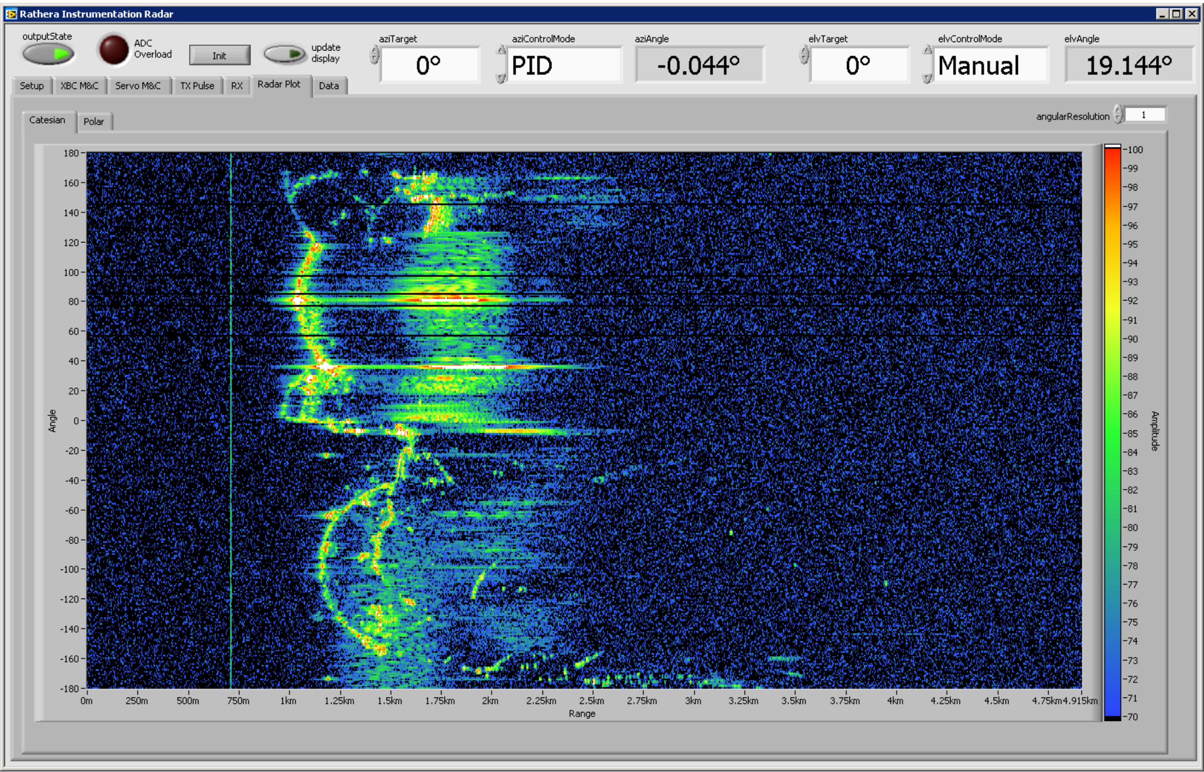

Figure 66 Cartesian radar intensity plot.

Polar Radar Plot

The polar plot overlays the pulse compressed echoes over a pre-loaded map image of the test location.

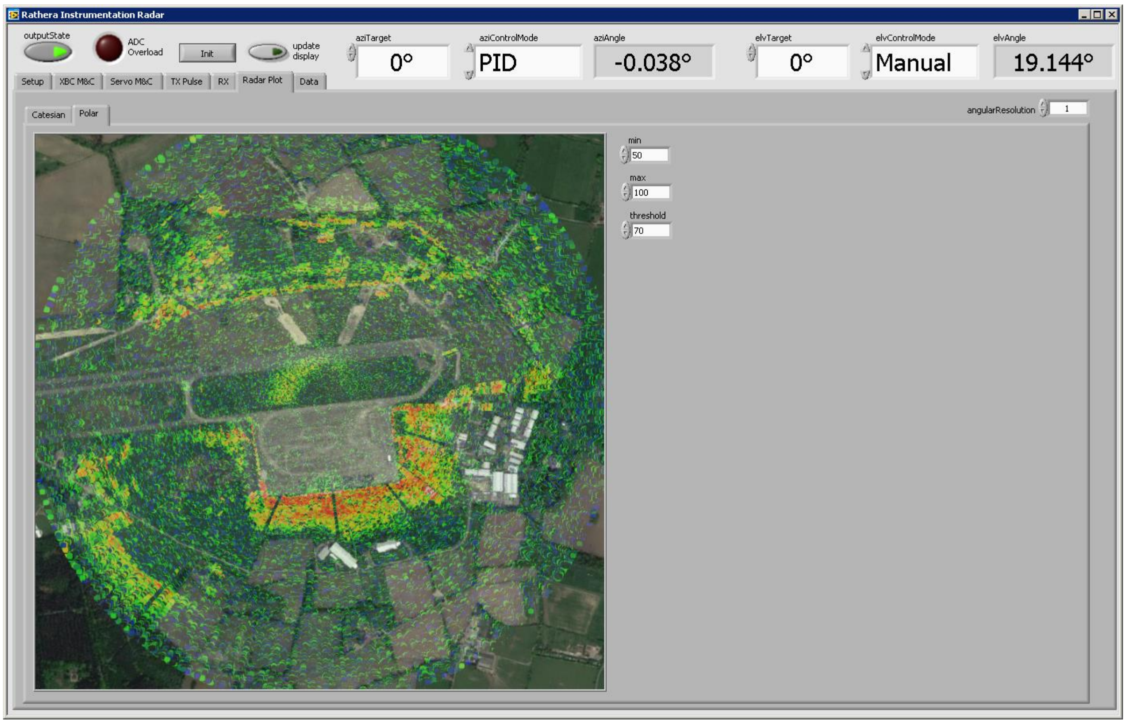

Figure 67 Polar radar plot.

Electro-Optical Sub-system

The electro-optics sub-system consists of a 12M megapixel high-quality camera with C-mount 50mm zoom lens (manually adjusted to near infinity), a 700m laser range finder with red laser dot, a digital compass/accelerometer, a GPS receiver, and a Raspberry Pi 4 running OpenHD software (an Open-Source software project developed for use on small drones).

The package is mounted on the top edge of the dish on the azimuth axis of rotation and aligned such that it is parallel with the main lobe of radiation. A lightweight 3D printed cover provides shade from the sun and weather.

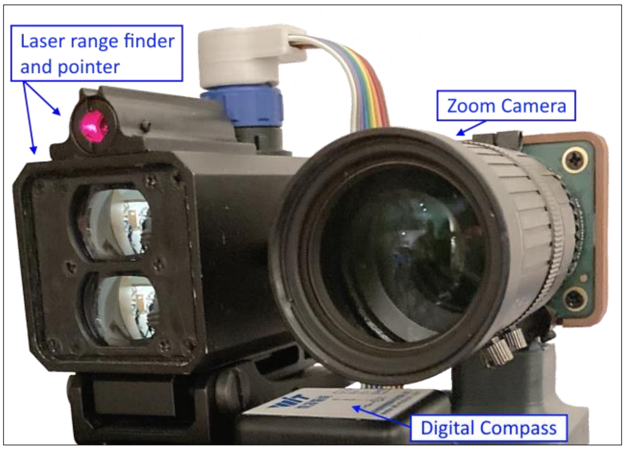

Figure 68 Electro-optical sub-system components.

Power is provided by two LG 18650 Lithium batteries installed in a Waveshare UPS HAT [6] which can operate the complete package wirelessly for approximately 3 hours.

An Alpha Networks AR9271 Wi-Fi Adapter provides a link to the OpenHD Ground Station with the transmit output power configured to its lowest setting to minimise interference with other spectrum users and maximise battery life.

OpenHD Software

The OpenHD [4] software running on the Raspberry Pi provides wireless high-definition, low-latency video, and bidirectional telemetry/control all over a single 2.4GHz or 5GHz radio link. The same software operates in Ground Station mode to take the received wireless data and convert it to an HDMI video signal for real-time display with a customisable overlay.

Figure 69 OpenHD Ground Station display.

OpenHD normally expects to interface with a Flight Controller, however, we did not wish to add further hardware to the system to support our range, position, and attitude sensors. Instead, custom Linux services ("daemons" in Linux parlance) were developed to take data from sensors and inject it in to Lightweight TeleMetry (LTM) protocol [5] frames for transmission to the Ground Station.

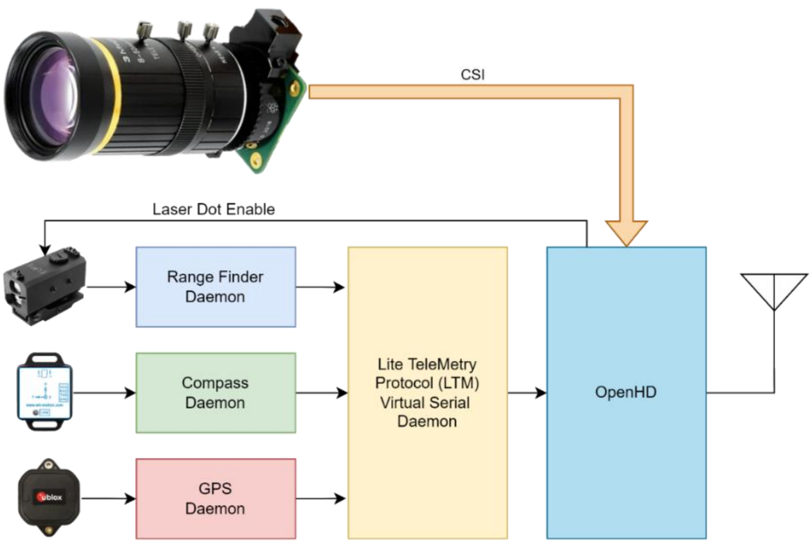

Figure 70 Additional software components added to OpenHD.

The telemetry is received by the Ground Station, shown on the OSD and forwarded via a serial interface to the Radar Processor.

Airborne ADS-B/Marine AIS Receiver

Airborne and marine craft, above a certain class, are required to broadcast their identification, position, and heading. It was deemed a useful feature to be able to track such craft using this information and so a Kinetic SBS-3 ADS-B & AIS SDR receiver was integrated in to the system.

The receiver connects to the processing computer via Ethernet and uses an external co-linear antenna to cover the ASD-B 1090MHz and VHF FM marine & AM air bands. Support for the newer Universal Access Transceiver band at 978MHz is also being considered.

Custom LabVIEW code receives and parses the data sent in ADS-B and AIS frames in to continuously updating latitude, longitude, and altitude values for each unique craft identifier. A set of equations, originally derived for calculating bearings between mountain peaks [9], is used to work out the azimuth & elevation the radar must steer to track the craft. Filtering the data by the supported extremes of elevation and any obscured azimuth angles along with sorting by range allows the prospective targets to be prioritised.

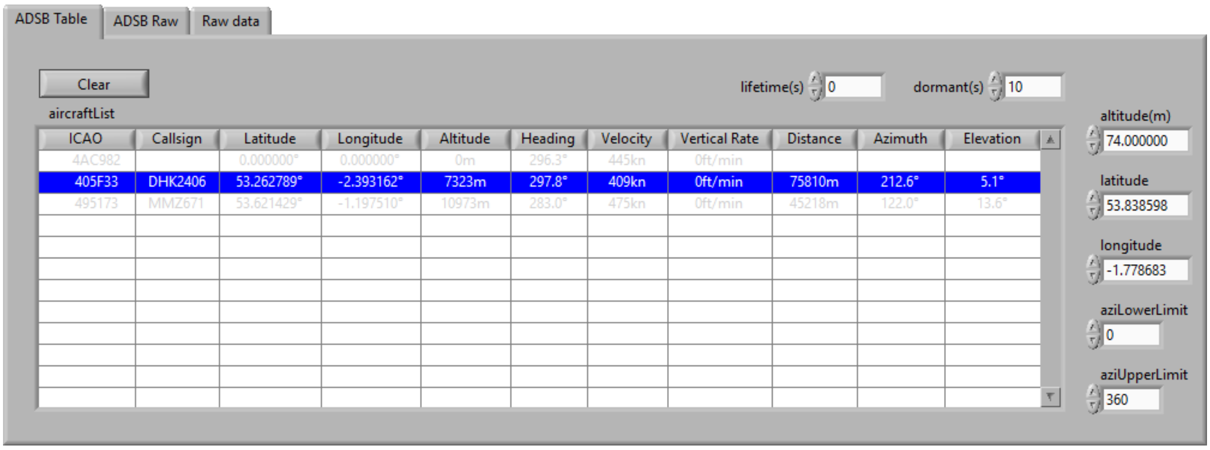

Figure 71 ADS-B data parsed and converted for airborne target tracking.

Radio Licensing

An OfCom Innovation & Trials License [8] was applied for in advance of deployment. This ensured that we had permission for our radar to emit in the specified band and provides a means for them to contact us in the unlikely event that unwanted interference is caused to other spectrum users.

The Innovation and Trial License allows the operation of the radar at fixed locations with the specified antenna, power output, modulation, duty cycle, and frequency band for a one-year period. This non-operational license offers quick, inexpensive access to radio spectrum for wireless tests and costs from as little as £50 per location per year. Our application took approximately 10-12 weeks to be granted.

Safety

The radar operates with a mean frequency of approximately 9.3GHz and a peak RF power of 3W. With the maximum estimated antenna gain (27dBi) and a 10% maximum duty cycle, the expected field strength will be circa 22V/m at a distance of 3m. This falls well below the ICNIRP/ARPANSA Public and Occupational maximum limits of 27.4V/m & 61.4V/m respectively.

The antenna is highly directional and a blanking capability means that transmission can be automatically suspended when the antenna is pointing towards the operator or any bystanders. Nonetheless, a visible cordon should be placed around the radar when operating so as to maintain a safe distance for on-lookers.

Flight Path Planning

The original flight trials were planned to be over a large body of water. The minimum and maximum operating range of the radar placed some requirements on the location, specifically that the closest the target could get to the radar and still record meaningful RCS was at least 50m and the furthest would be no more than 1-2km.

An example facility is shown below with a useable area of approximately 400m x 400m. Three flight paths were proposed to measure the RCS from six views (front/rear, left/right and diagonals):

- Cross Range Path
- Down Range Path
- Diagonal Path

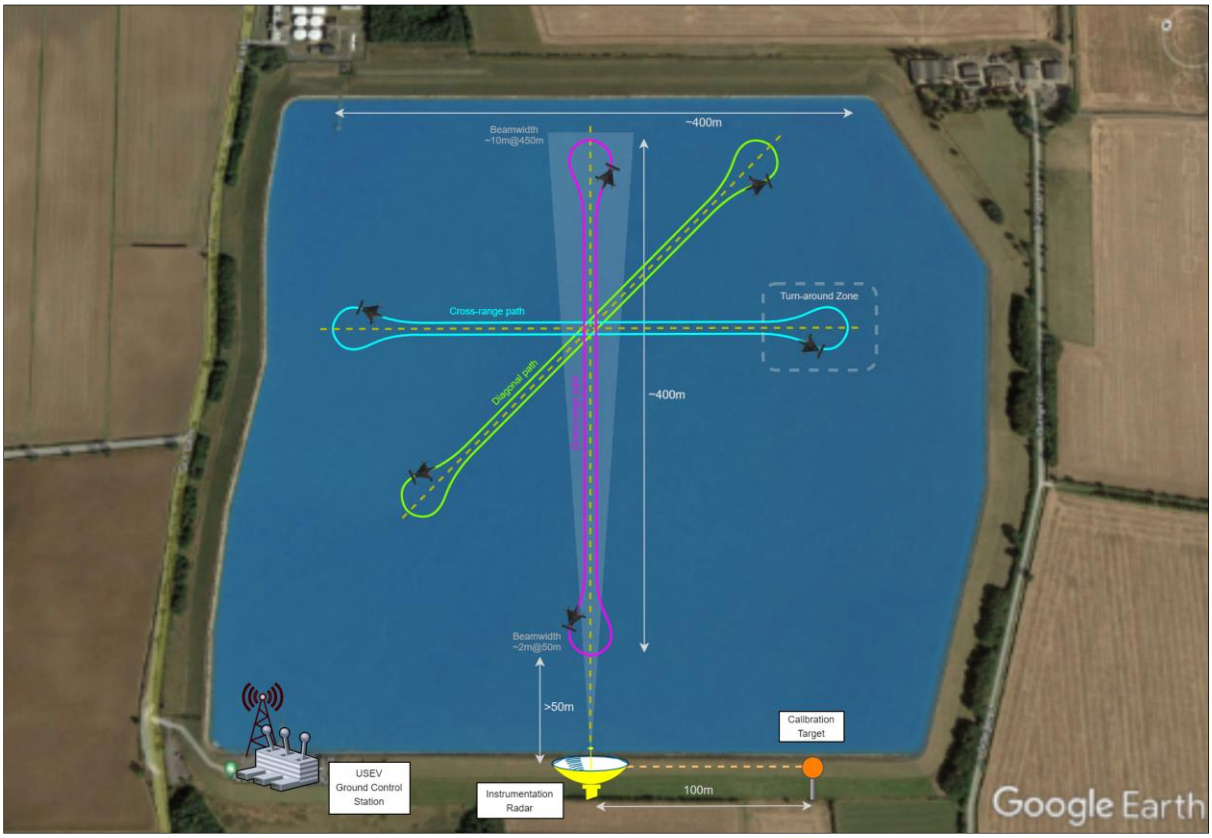

Figure 72 Initial over water flight test plans.

The instrumentation radar would be installed mid-way along one edge of the body of water ideally approximately 50m from the nearest designated turn-around zone. It was noted that we have limited capacity to mitigate multi-path interference and so some judgement would be required to analyse the recorded RCS data for each pass.

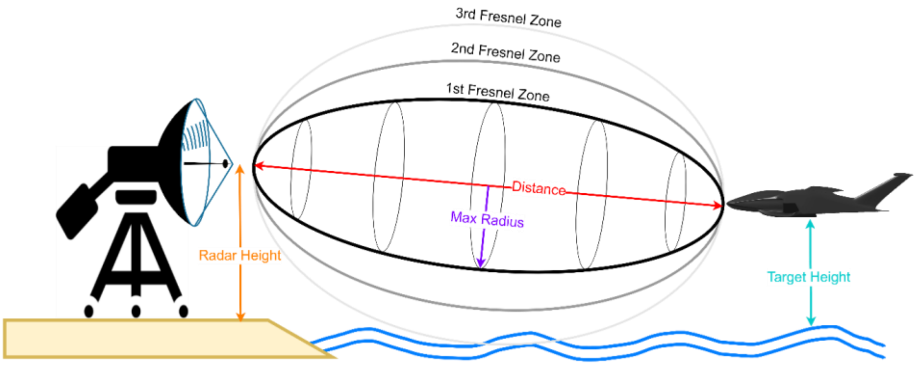

Figure 73 Fresnel zone estimation model.

It was expected for the USEV to travel at least 2-3m above the water's surface. It was calculated that the 3rd Fresnel zone [10] (at a 400m range between the radar and target) would have a maximum radius of 3.1m. It's therefore advantageous to elevate the radar's height on the shoreline if the USEV travels at a lower height above the water's surface.

It was noted that it may not be possible to extract useful RCS measurements in the turn-around zones if the angular velocity is too high or where local clutter becomes dominant.

Trials

The flight trials were rescheduled to be performed on the concrete hardstand at Elvington Airfield, Yorkshire. The runways were enlarged and the hardstand was added in 1952 for use by the USAF B36 nuclear deterrent bombers. The solid surface has the advantage of providing less drag during take-off (compared with water), but poses a higher risk of "unplanned disassembly" during landing.

Figure 74 The concrete hardstand at Elvington Airfield, Yorkshire.

Figure 75 A “chaser car” and quadcopter drones provided vital video footage for the After-Action Review.

Unfortunately, the first flight of the Uncrewed Surface Effect Vehicle was rather short but very distinguished! As a result, no further airborne testing could be performed that day, so our validation of the radar was limited to the surrounding terrain. The cartesian below shows a 360° sweep in azimuth covering the concrete apron and runways.

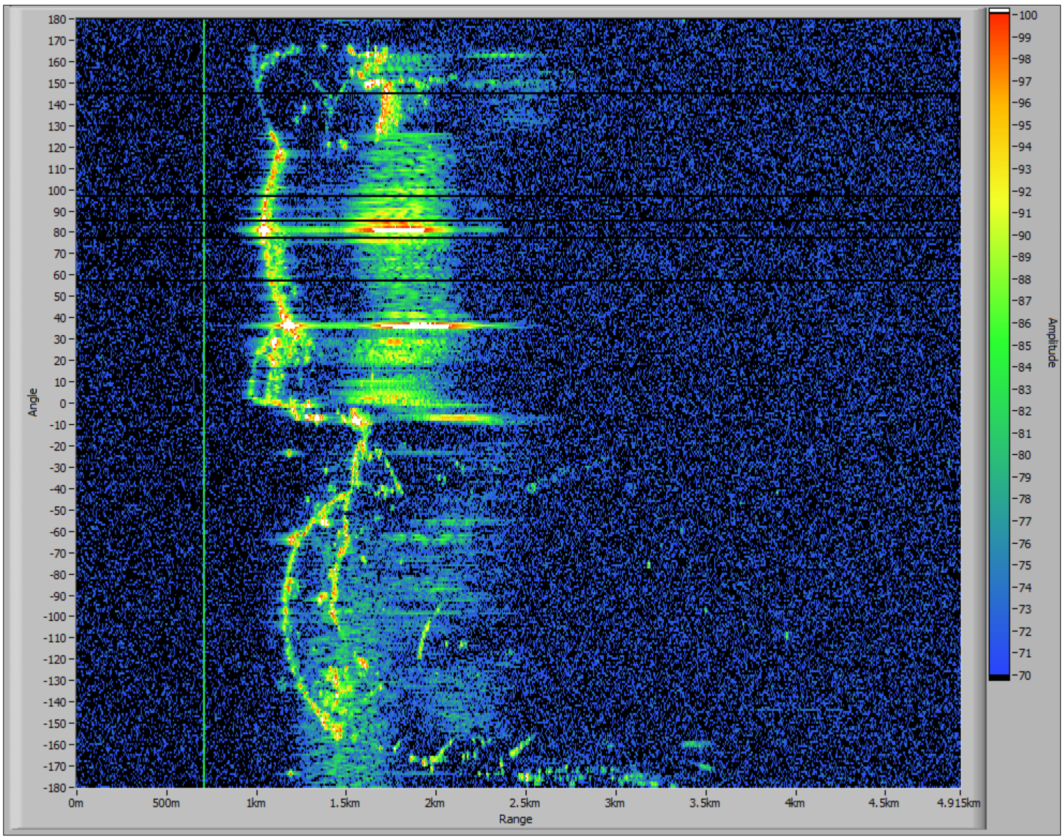

Figure 76 An intensity plot of range vs azimuth angle showing fine resolution.

The same data was transformed into a polar coordinate system and overlaid on to a Google Maps presentation of the surrounding area. The blind range of the radar can be seen with no targets detectable within the cyan circle. With the configured pulse width, targets were visible at ranges well over 1km. Tall trees to the south and east of the hardstand shadowed the surrounding buildings and warehouses.

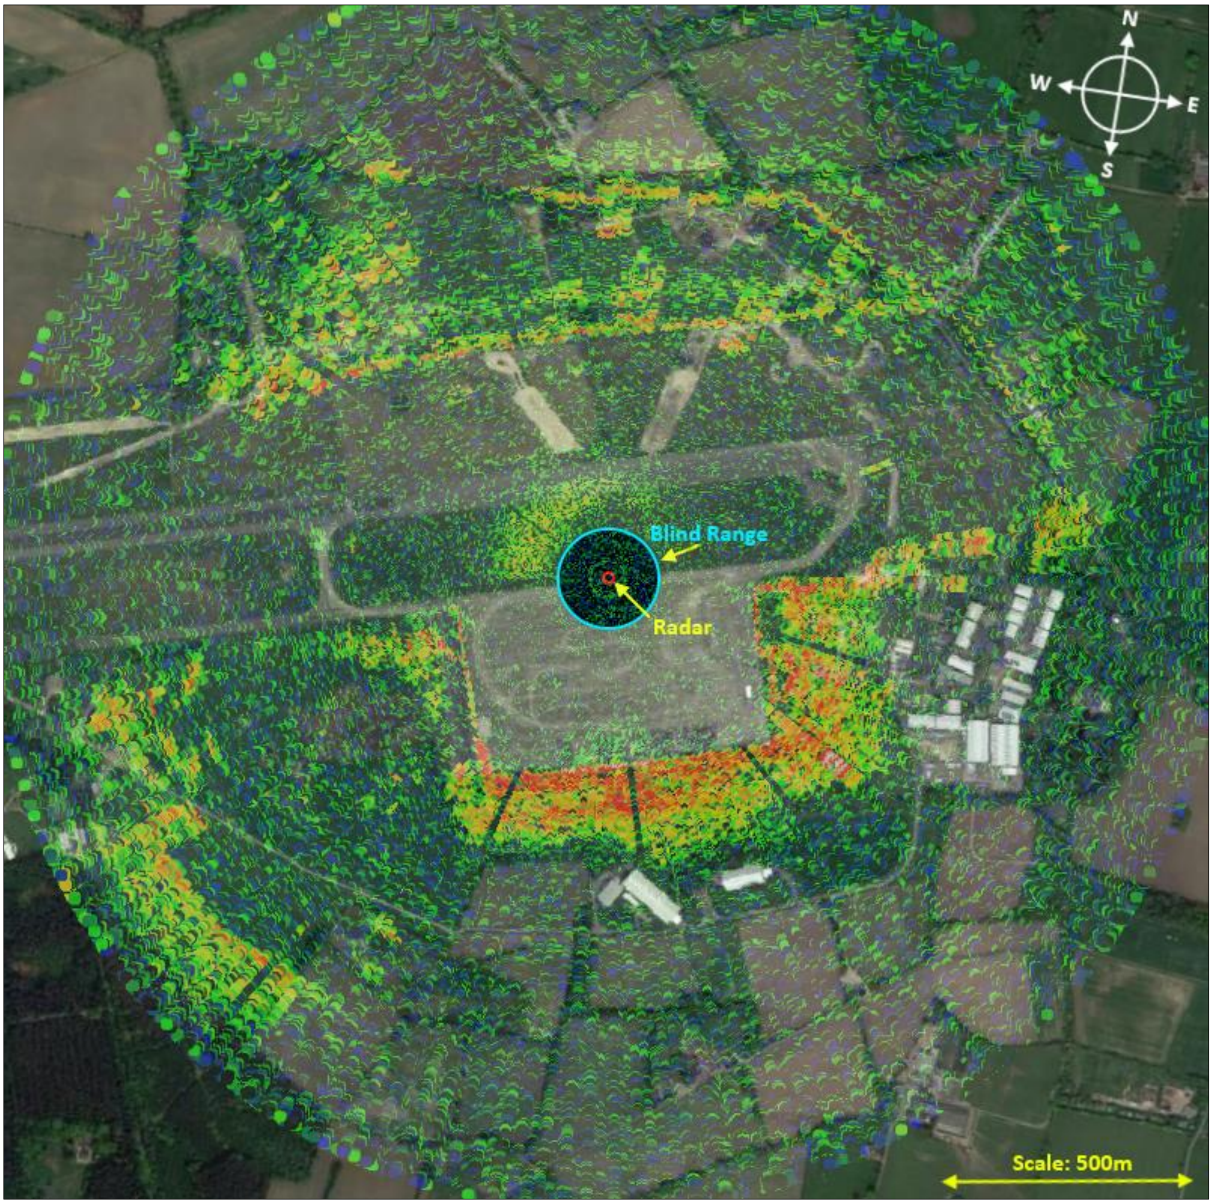

Figure 77 Processed data transformed on to a polar plot and overlaid on a map of the flight test area.

We were very happy with the sensitivity and range resolution of the instrumentation radar given its very low transmit power of just 3W and we hope to return to Elvington Airfield soon to complete RCS characterisation of the USEV.

Future Work

With many of the original design aims met, we can consider additional areas of development such that this system can become a more general platform for radar experimentation.

- Upgrade azimuth servo to direct-drive to reduce backlash and increase torque. This will help combat strong winds pulling the scanner slightly off target.
- Migrate FFT & IFFT CPU processing to NVIDIA GPU as previously performed in Rathera's Wideband Receiver.
- Develop visual non-cooperative target tracking using OpenCV methods in a co-executing Python thread.
- Add Kalman filter to ADS-B aircraft trajectories for smoother target tracking and forward prediction.
- Investigate alternative intra-pulse modulation waveforms including:
    - Hyperbolic FM for immunity to target doppler.
    - Non-linear FM for lower side lobes, but less tolerance to doppler.
    - Phase modulation (Barker codes).
- Interleave short, medium, and long pulses with composite processing in to a single continuous range extent.
- Create an API for remote control and streaming of raw and processed data over the internet.
- Recompute the LyceanEM scanner model using a parabolic reflector meshing algorithm.

Credits

James Pawson – Unit 3 Compliance. Thank you for your support throughout this project and access to your excellent facilities & biscuits.

T.G. Pelham, UK Intelligence Community Postdoctoral Research Fellow, University of Bristol. Thank you for your extremely valuable insights and for performing the LyceanEM scanner characterisation & modelling.

References

[1] Pelham, T. Rapid Antenna and Array Analysis for Virtual Prototyping, Radar 2022, 26/10/2022

[2] https://lyceanem-python.readthedocs.io/en/latest/index.html

[3] https://www.ilme.com/en/view-product/?code=APV+19&country=uk&origin=uk

[4] https://openhd.gitbook.io/open-hd/

[5] https://github.com/iNavFlight/inav/wiki/Lightweight-Telemetry-(LTM)

[6] https://www.waveshare.com/wiki/UPS_HAT_(B)

[7] https://www.instructables.com/Inexpensive-method-of-industrial-level-quality-PCB/

[8] https://www.ofcom.org.uk/manage-your-licence/radiocommunication-licences/non-operational-licences

[9] http://tchester.org/sgm/analysis/peaks/how_to_get_view_params.html

[10] https://www.omnicalculator.com/physics/fresnel-zone

Appendix A – Prototype RF PCB Manufacture

All of the RF PCBs used in this project were manufactured in-house on our LPKF milling machine. This machine provides us with the ability to iterate RF designs very quickly, also being ideal for prototyping filters and distributed structures.

Figure 78 LPKF PCB Milling Machine.

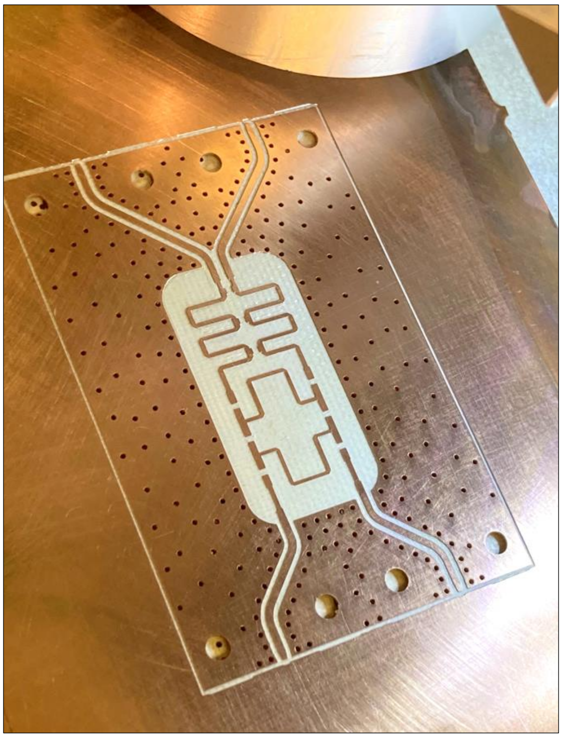

Figure 79 A distributed filter made on the LPKF.

The thinnest RF PCB laminate we routinely machine is 0.254mm (e.g. RO4350B). The smallest isolation channel between tracks that we can typically machine is 0.1mm to 0.2mm and the smallest drill we use is 0.3mm (e.g. for under QFN package exposed ground pads). There are practical limits to the maximum density of vias that can be machined whilst maintaining the mechanical integrity of the dielectric and laminated copper.

A small board can be drilled and isolation routed within an hour. However, there is one significant step not performed by the machine: Through Hole Plating.

The processes used by the commercial PCB industry for Through Hole Plating are not that practical for small companies and using tinned copper wire or ribbon is very labour-intensive.

One popular process used by hobbyists involves squeegeeing a conductive acrylic ink across a drilled (but not isolated) bare PCB. Once cured, the ink acts as a conductive surface on which copper can be electroplated. However, the ink is fairly viscous and whilst the via walls must be coated, we do not want the holes to be blocked with the ink as this would prevent deposition of the electroplated copper. Extraction of the residual ink can be performed with compressed air or a vacuum, though again this is quite labour-intensive and may not be repeatable.

We therefore sought an alternative process.

Prior to the fall of the USSR, a group of Russian chemists [7] developed a process that utilised the ability of some compounds to decompose at high temperatures in to a very thin layer of metallic copper nano-particles. This conductive layer was then used as the seed to electroplate additional copper using the conventional process. This technique has been found to be very reliable on low aspect ratio holes (I.e. holes in thin materials such as 0.254mm RO4350).

Figure 80 Drilled PCB activiation.

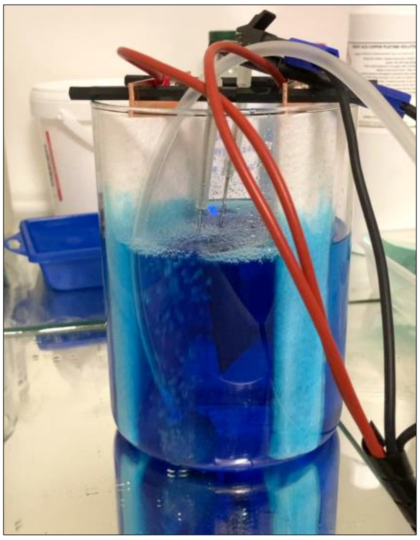

Figure 81 Copper electroplating.

The process first involves attaching the bare PCB laminate to the milling machine using locating pegs. The machine then drills the PCB to define the holes, vias, and slots that are to be plated. The PCB is briefly dipped in to the activation solution taking care to ensure that the liquid flows in to all holes. It's important to make sure that the bare surface of the PCB is not touched (even when wearing gloves) as this can degrade the performance of the process. The dipped PCB is then placed in an oven and heated using a defined temperature profile.

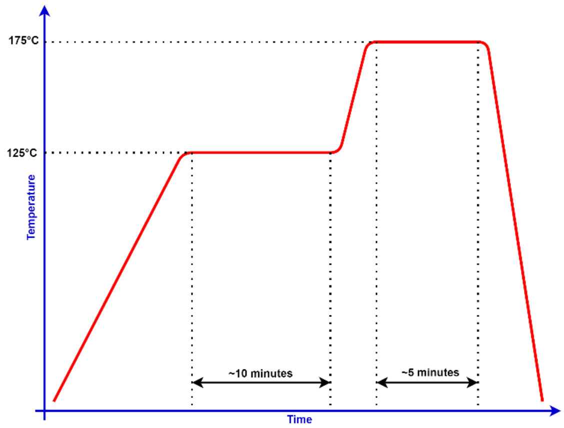

Figure 82 Activation temperature profile.

The activation solution decomposes at high temperatures, depositing nanometres of copper on to all the exposed surfaces (including the via walls). The PCB is then cleaned by hand with soap & water to remove the by-products and inspected to ensure all holes are clear of residue. The activated PCB is placed in a conventional copper electroplating bath with an anode plate on either side at equal spacing. With an appropriate current setting, 17-35um of copper can be deposited on to both sides of the PCB in approximately 30 minutes.

The PCB is washed to remove the acidic plating solution from the holes, dried, and then placed back on the milling machine (using the locating pegs for reference) to perform the single- or double-sided isolation routing steps.

This entire plating process adds approximately 1 hour to the PCB production time. It provides a reliable electrical and thermal connection that can easily survive the temperature cycling required to solder and desolder packaged parts (e.g. QFNs).

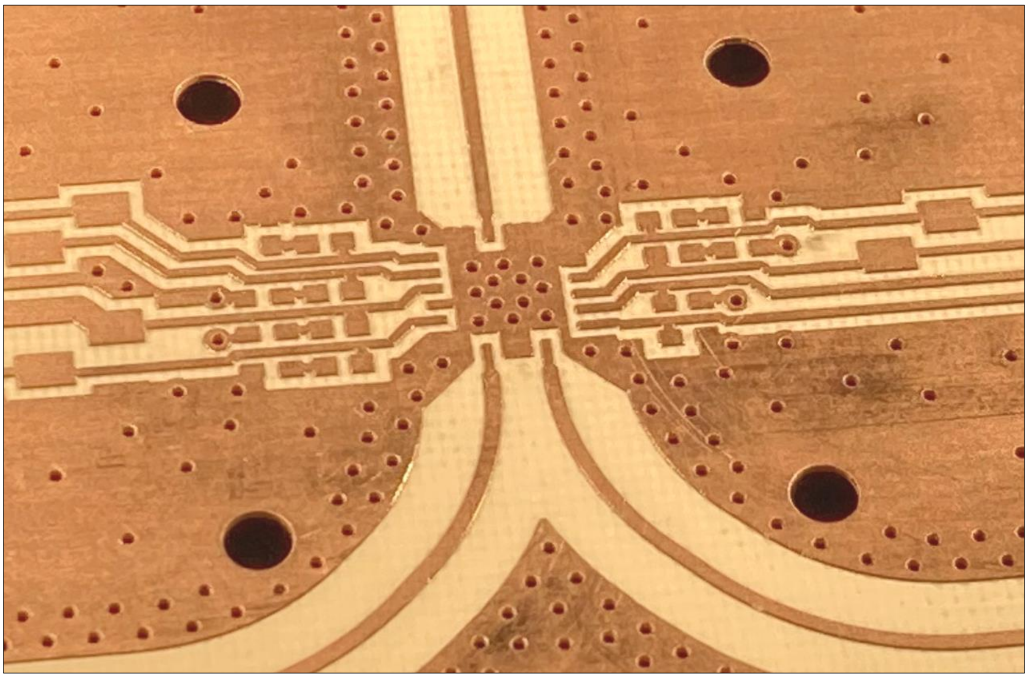

Figure 83 A PCB with Though-Plated Holes using the process described.

It is important to recognise that we can essentially have as many vias as we need for a design, without incurring significant additional assembly time/cost. This process works very well with Grounded Coplanar Waveguide (CPWG) transmission lines as the bare minimum of copper must be machined away and the many vias needed are essentially free.

The activation solution consists of:

- Distilled Water
- Copper Sulphate
- Calcium Hypophosphite
- Ammonium Hydroxide
- Liquid Soap

All of the chemicals used are readily available with the exception of the Calcium hypophosphate. This has unfortunately gained notoriety in the production of illicit drugs and so must be procured from a limited pool of chemical distributors. It is possible to use Sodium Hypophosphate in conjunction with Calcium carbonate and some additional steps, but it is believed that this solution will be a little less effective.

The volume of solution used for each PCB activation is extremely small with the run-off being returned to the container. Whilst the investment in chemicals & glassware is not inconsiderable, 500ml would likely last a lifetime!