Contents

January 2024

Regulars

Features

Reviews

30

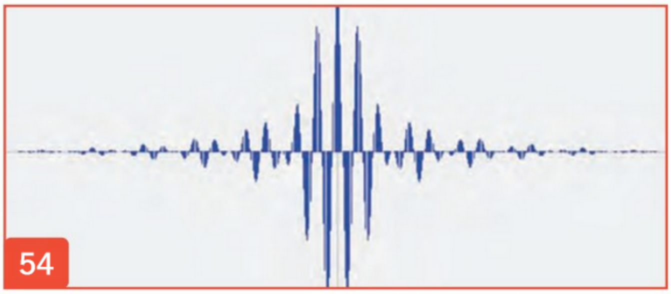

54

News and Reports

Technical Features

Cover design by Kevin Williams, M6CYB
Features covers from:
July 1925
June 1945
January 1955
August 1965
February 1974
June/July 1980
July 1993
April 2005
June 2020
August 2021

RadCom THE RADIO SOCIETY OF GREAT BRITAIN'S MEMBERS' MAGAZINE
Managing Editor: Edward O'Neill, M0TZX, edward.oneill@rsgb.org.uk
Technical Editor: Peter Duffett-Smith, GM3XJE
Layout and Design: Kevin Williams, M6CYB, kevin.williams@rsgb.org.uk

All contributions and correspondence concerning *RadCom* should be emailed to: radcom@rsgb.org.uk. Alternatively by post to *RadCom* Editor, 3 Abbey Court, Fraser Road, Priory Business Park, Bedford MK44 3WH Phone 01234 832 700.

RadCom is published by the Radio Society of Great Britain as its official journal and is sent free and post paid to all Members of the Society. **The February 2024 edition of *RadCom* is expected to arrive with most Members by 19 January 2024** although this can take up to a week longer in some cases; international deliveries can take longer still.

All material in *RadCom* is subject to editing for length, clarity, style, punctuation, grammar, legality & taste. Articles for *RadCom* are accepted on the strict understanding that they are previously unpublished and not currently on offer to any other publication. Unless otherwise indicated the RSGB has purchased all rights to published articles. No responsibility can be assumed for the return of unsolicited material. See www.rsgb.org/radcompix for info on taking photos for publication.

The online *RadCom* is at www.rsgb.org/radcom/
RadCom Plus is available to RSGB Members online at www.rsgb.org/radcom-plus
RadCom Basics for Members new to the hobby can be found at www.rsgb.org/radcom-basics/
Abbreviations and acronyms we use are listed at http://tinyurl.com/RC-acronyms

RADCOM (ISSN No: 1367-1499) is published monthly by the Radio Society of Great Britain and is distributed in the USA by RRD/Spatial, 1250 Valley Brook Ave, Lyndhurst NJ 07071. Periodicals postage pending paid at So Hackensack NJ. POSTMASTER: end address changes to RADCOM c/o RRD, 1250 Valley Brook Ave, Lyndhurst NJ 07071

www.facebook.com/theRSGB
Twitter @theRSGB

New Products

Rooster 40m CW transceiver from Kanga Products

The Rooster is a replacement for the old FOXX3 kit that Kanga produced for many years. Kanga still gets asked for it but decided to produce a new kit with better performance while keeping the cost low. Kanga wanted more power, better BCI rejection, some audio filtering and a more practical case that was supplied as standard with the kit. The result was the Rooster.

The kit is partly an SMD design but **all** the SMD parts are pre-installed for you leaving around 20 parts to fit. The Rooster is supplied with a strong cream-white aluminium case.

The Rooster will give 2W of RF at 13.8V (7.030MHz).

A 26-page colour instruction manual is available online.

Available from https://www.kanga-products.co.uk for £37.99

Bhi NEDSP1962-KBD amplified DSP noise cancelling module

The NEDSP1962-KBD amplified DSP noise cancelling module can be retrofitted inside many types of extension speaker or older style radios and transceivers. It incorporates unique digital signal processing (DSP) technology which can identify speech from within a noisy signal and can provide up to 40dB of noise reduction and 65dB of tone reduction. The module comes pre-wired and with an on-board power amplifier enabling it to be easily incorporated into existing equipment. All the module functions are controlled by the pre-wired keyboard switch assembly. The NEDSP1962-KBD uses a powerful programmable DSP audio processing chip with DSP noise cancelling technology inside. This high-performance chip delivers exceptional sound quality and incorporates a full audio signal chain with 16-bit A/D converters and digital interfaces utilising fully flexible digital processing architecture. The audio amplifier stage is provided by a Texas Instruments 7W Mono Class-D audio power amplifier chip which is 94% efficient, eliminating the need for heatsinks.

Available from https://bhi-ltd.com for £139.96

Geochron – new features

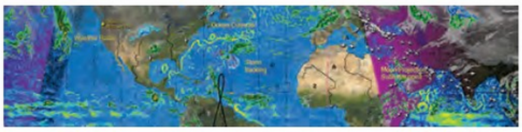

Geochron's new Premium Weather app has 30 live weather layers from AERIS, including animated radar.

The Geochron team is enhancing the regional zoom feature with CQ and ITU overlays for the Ham Radio Bundle, plus the Moon's sublunar point and visible projection.

New regional zoom layers will be released in the coming months. The first area to be released will be Australia / New Zealand.

Geochron's premium layers have been popular, and now you can bundle them all together for one low annual price.

Find out more at https://www.geochron.com/ham-radio-4k

Rigexpert Shackmaster power 500 (36A) power supply

The Shackmaster Power 500 is the pinnacle of power supply innovation, designed to redefine the way you power and interact with your devices. With an unwavering commitment to excellence, this power supply is engineered to provide you with unparalleled convenience, safety and flexibility, all while enhancing your workspace aesthetics.

Prioritising safety and security, it features an advanced Intelligent Protection System. So, your device is shielded from overcurrent, overvoltage and overheating, allowing you to work confidently, knowing your valuable equipment is well-guarded.

The ShackMaster Power 500 maintains a rock-solid 13.8V output voltage, ensuring a consistent and reliable power supply to your device. This level of stability is crucial for delicate electronics, ensuring peak performance and prolonged device lifespan.

The ShackMaster Power 500 is equipped with four Powerpole DC connectors, offering ample ports to effortlessly connect various devices without the need for adapters or complex configurations.

Featuring 2 USB Type-A ports and 2 USB Type-C ports, it empowers you to charge devices, connect peripherals, or power USB-compatible gadgets seamlessly, all while maintaining a clutter-free workspace.

The ShackMaster Power 500 includes a built-in USB interface that lets you connect it to your computer, enabling you to log and analyse consumption data, allowing for informed decisions about energy efficiency.

Features:

- Space-saving design with touch screen display
- Low acoustic noise
- Multifunctional touch display
- -30°C to +70°C operating temperature
- Intelligent Protection System

Available from https://moonrakeronline.com for £399.95

Antennas

Anyone living with a small garden or back yard at their QTH will no doubt empathise with the challenge this presents when seeking to use the HF amateur radio bands at home.

Similarly, those of us who indulge in portable operating may also have to consider not taking up too much horizontal space, to avoid tripping up other users of your outdoor location. Bearing all this in mind, the vertical antenna is often a preferred option. At my QTH, the ground space available is very limited, and ground radials are not a good option as other family members also use the garden.

Elevated verticals are therefore the preferred choice, as only a minimal number of radials (in my case two) need to be used to bring the efficiency close to that of an antenna using dozens of ground radials. The antenna which immediately springs to mind is the elevated quarter-wave vertical antenna. This antenna provides useful low-angle gain, and is quite simple to build. Aside from the 40m version (which can also be used on 15m as a three-quarter wave antenna), these antennas can usually only be used for one band. I decided, therefore, to examine the possibility of using the elevated vertical antenna, but with the idea of using shorter-length elevated radials, thus saving on space.

Shorter radials

Keen to explore possibilities, my online research saw me stumble across a very interesting series of articles **[1]** which addressed the possibility of shortening the usual quarter-wave long elevated vertical, and using two radials positioned at 180° from each other at a length appreciably shorter than a quarter-wave. This would also necessitate the lengthening of the vertical element so that the antenna can reach resonance at a chosen frequency. The radiator length is therefore longer than a quarter-wave and, at that length, before adjusting the quarter-wave-long radials, the measured impedance will be a large positive value. In order to aim for a resonant antenna, and therefore a low value of the VSWR, we need to decrease the length of the two elevated radials until the reactive component is reduced near to zero ohms, the resonant point.

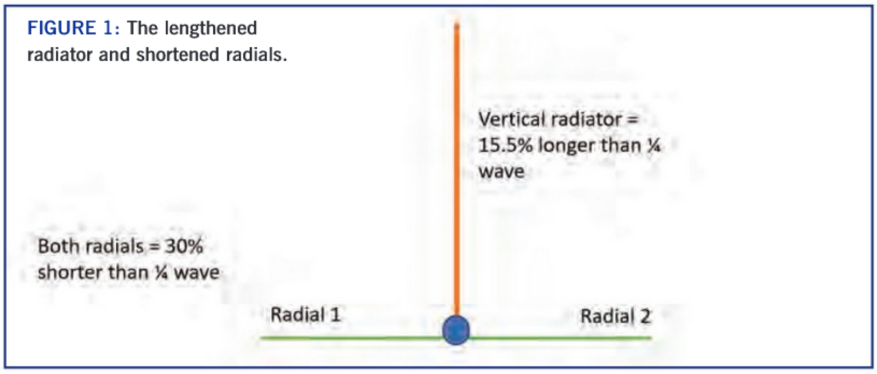

FIGURE 1: The lengthened radiator and shortened radials.

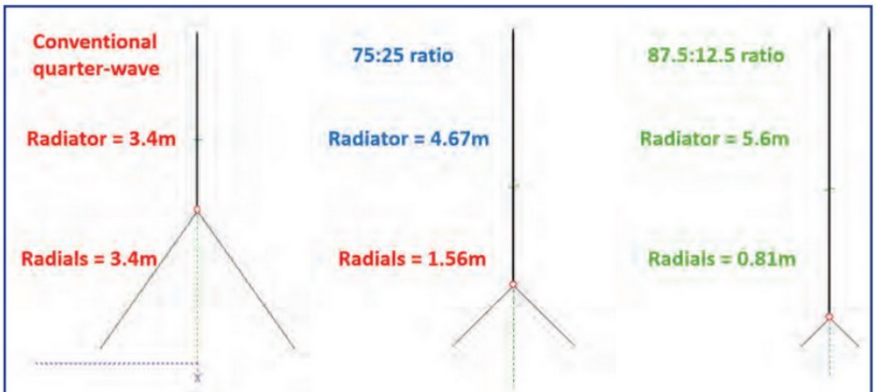

FIGURE 2: Three configurations of the antenna.

While the resistive component is largely determined by the length of the vertical *radiator*, decreasing the *radial* lengths almost entirely affects only the reactive component at the feed point (the jX part). The technique by which we increase the vertical radiator length to achieve 50Ω is known as 'over loading'. The article summarises that an increase in length of the radiator by 15.5%, and a reduction of both the radials by 30%, would produce resonance and a good VSWR match to 50Ω (see **Figure 1**).

Exploring options

My mind switched to considering how much we could possibly reduce the radial lengths by, so reducing the space required for using this antenna. I chose the 15m band, and the centre frequency of 21.250MHz, to assess these different configurations. Moreover, I modelled the top of the vertical element to be at a height of 6.9m above ground, to mimic the use of a 7m fibreglass pole to support the antenna. In each case, the radials ran at 45° from the vertical direction at the feed point, and were at 180° from each other.

To begin with, the orthodox quarter-wave vertical radiator, plus two quarter-wave radials each measuring 3.4m in length (the 50:50 arrangement), produced a modelled VSWR of 1.34:1 (using MMANA-GAL) at its resonant point (with the real part, R, being 67Ω). The next configuration which I modelled was a 62.5:37.5 arrangement, with the vertical radiator now measuring 4m and each radial 2.4m. Notice how this reduces the combined vertical-plus-single-radiator size from 6.8m to 6.4m, a reduction of 5.9%. What is also encouraging is that the VSWR at resonance is 1.49:1 (resistive part 74Ω and the reactive part zero, or $Z = 74\Omega + j0\Omega$), quite like the 50:50 configuration.

At a 66:33 ratio the vertical element increases in length to 4.21m, and each radial reduces further in length to 2.1m. Again, the combined vertical/single radiator length is reduced to 6.31m, with the radials now

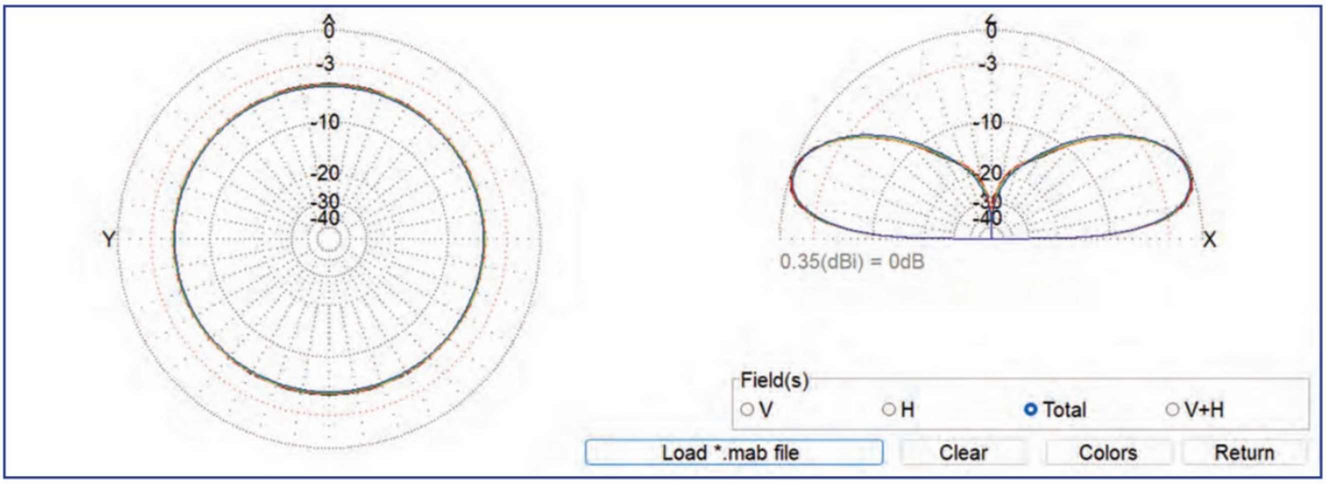

FIGURE 3: **The gain at 5° take-off angle is practically identical in all five configurations.**

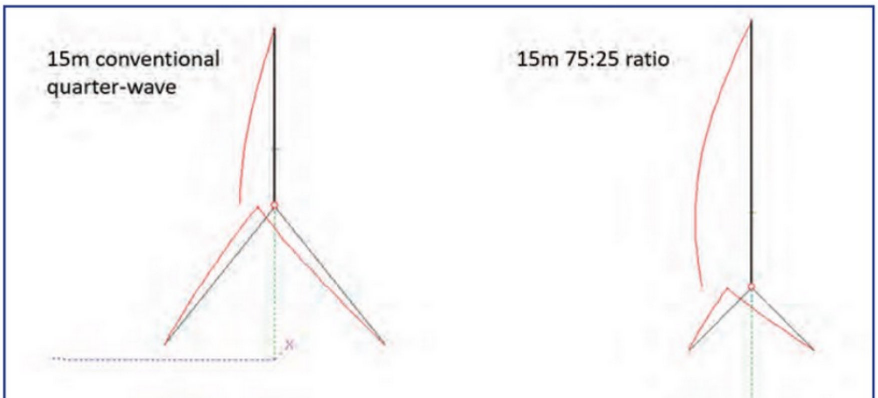

FIGURE 4: **The current distribution in a conventional quarter-wave antenna, and the antenna with a much-longer vertical element, and shorter radials.**

38% shorter and the vertical 24% longer than the conventional quarter-wave antenna. Encouragingly, VSWR at resonance is 1.56:1 ($Z = 75\Omega + j0\Omega$). Three configurations are illustrated in **Figure 2**.

How far can we take this trend in radial reduction? **Table 1** shows the data produced by modelling different configurations, from the conventional '50:50 vertical:radial quarter-wave' to a 7:1 ratio, whereby the vertical element is much longer than the elevated radials. The trend seems to suggest that we do get an increasing VSWR as we increase the vertical element and shorten our radials. However, in many cases this is not prohibitive, and a feed-point VSWR of 1.92:1, as found for the 75:25 ratio, would make little real-world difference compared with a 1:1 VSWR. It could be argued that the 2.43:1 VSWR, modelled with the 80:20 ratio, will probably be reduced at the operating end anyway (because of losses in the feeder), and a moderate run of coaxial cable carrying such a VSWR should not be discounted as an option. It isn't until we reach the 87.5:12.5 ratio that we see the VSWR begin to climb appreciably. As the vertical element begins to dominate the overall antenna configuration, we see the value of R begin to climb at the antenna's resonant point.

The overall length (measured as vertical plus one radial) also decreases as we make these adjustments. This reaches a maximum reduction of 8.3% at the 75:25 ratio. At this same ratio, each radial length is reduced to 1.56m, equating to a reduction of 54.1% compared with their initial quarter-wave lengths of 3.4m. At this same point, the vertical length has increased by 37% from the initial quarter-wavelength.

Antenna gain

It is also useful to gauge the impact these changes may make on the low-angle gain of the antenna. **Figure 3** shows a comparison of the gain in dBi as the vertical element of the antenna becomes progressively longer compared with the two radials. As we can see there is practically no difference between these different configurations in terms of low-angle (5°) gain. The modelled gain, in fact, reduces by a mere 0.3dBi between the best (the conventional 50:50 configuration) and the worst (the 87.5:12.5 configuration).

Scaling to lower bands

So far, we have looked at this antenna on the 15m band. Can it be scaled to accommodate the lower HF bands, which would, however, cause more issues in terms of squeezing the antenna into a small horizontal space?

On 20m and using 14.220MHz, the modelled results that we saw on 15m are mirrored. Not only does the conventional quarter-wave antenna show the same impedance and VSWR measurements as the 15m version, so too do the adjusted vertical- and radial-length configurations. This allows us to trim the radials for a 20m vertical to 2.32m in length as we extend the vertical element to a length of 7m. This 75:25 ratio gives a sub-2:1 VSWR with the top of the vertical element at 10m above ground, and both radials just under 2m above ground (importantly just over head height). Gain at the 5° take-off angle is identical to that of the 15m version across the various configurations.

On 40m, savings in terms of radial lengths become even more of a win. Taking the 75:25 configuration as an example, the radial lengths are reduced from 10m each to 4.59m. The trade-off here is that the vertical radiator is now extended to 13.79m, but if you have a tree or similar support for use and not much horizontal space, then this could

Tim Hier, G5TM
timhier@icloud.com

TABLE 1: Modelled data when adjusting vertical and radial lengths for 21.250MHz

% ratio of lengths (vertical: radial)	Vertical length (m)	Radial length (m)	Combined length of the vertical and one radial (m)	VSWR at resonance (X=0)	Feed-point impedance at resonance (W)
50 : 50	3.40	3.40	6.80	1.34:1	67 + j0
62.5 : 37.5	4.00	2.40	6.40	1.49:1	74 + j0
66.6 : 33.3	4.21	2.10	6.31	1.56:1	78 + j0
70 : 30	4.38	1.89	6.27	1.64:1	82 + j0
75 : 25	4.67	1.56	6.23	1.92:1	96 + j0
80 : 20	5.00	1.25	6.25	2.43:1	121 + j0
87.5 : 12.5	5.60	0.81	6.41	4.54:1	227 + j0

TABLE 2: VSWR range with the 75:25 configuration

Band (m)	Resonant frequency, VSWR – 1.92:1 (MHz)	VSWR at the bottom of the band	VSWR at the top of the band
15	21.250	1.94:1 (21.000MHz)	2.05:1 (21.450MHz)
20	14.220	2.01:1 (14.000MHz)	2.05:1 (14.350MHz)
40	7.150	2.01:1(7.000MHz)	2.16:1 (7.200MHz)

work for you. Bear in mind that if you are unable to increase the height of the feed point, then on 40m the radials, even when reduced, may need to be run away nearer to 90° than 45° from the vertical direction (this has been considered when undertaking the above modelling on 40m). One issue with the adjusted 40m vertical, is that the harmonic alignment we usually see with 15m using the quarter-wave is no longer available.

VSWR coverage and currents

As we based the modelling for each band on attaining resonance on a relatively central frequency, it would be necessary to check the VSWR curve we can expect to see. Using the 75:25 configuration as an example, and its 1.92:1 VSWR at resonance, it is useful to note that, for 15m, 20m and 40m, the VSWR curve is quite flat across each band (**Table 2**). Finally, it is worth noting that the current distribution in both the conventional elevated quarter-wave vertical antenna and, for example, the 75:25 configuration is also quite similar (see **Figure 4**).

Final thoughts

It was quite intriguing looking at how the conventional quarter-wave antenna could be adjusted and developed to cater for the occasions where shorter radial space is required. As ever, the idea could only be pushed so far, and the 75:25 configuration seems to be where we approach the limit between achieving less horizontal space and keeping the VSWR to just around 2:1. However, on 40m especially, the very-similar performance in terms of gain between the adjusted vertical antenna, and the conventional quarter-wave antenna, and the win in terms of reducing radial lengths, make this design something to consider. It is also worth considering the use of a choke at the feed point with this antenna. This is advisable whether this antenna is configured as a conventional elevated quarter-wave antenna, or if it is used as what would, in effect, be an off-centre fed vertical antenna.

References

1. https://labsdl.wordpress.com/2017/11/14/resizing-your-antenna-for-another-band/

Want to see *your* work published in *RadCom*?

RadCom, *RadCom Basics* and *RadCom Plus* are always looking for people with good, previously-unpublished projects and features. Many of the technical articles you see in these pages come from people with no previous publishing or writing experience. What they do have is enthusiasm and a good, original, unpublished idea.

Big or small, as long as it's something to do with amateur radio, we're interested. There is great demand for constructional projects, ranging from simple to complex. Antenna projects and any sort of buildable electronic widget are favourites. The most complex projects go into *RadCom Plus* but the majority will appear here in *RadCom*.

What to do

The first step is to get in touch with the editorial team via email to radcom@rsgb.org.uk and simply tell us what you have in mind.

It doesn't matter if you have never written before – we can provide as much or as little editorial support as you require. We'll need good quality photos to accompany your article **[1]** but line drawings (circuits and other diagrams) can be a simple sketch. We re-draw the majority of diagrams to give a consistent style across all articles. There is rarely any time pressure, so you can work at your own pace.

Acceptance procedure

Our acceptance procedure for technical articles **[2]** involves an independent review by a member of the Technical Panel, whose members are always willing to work with authors if there's something that needs more explanation or would benefit from clarification.

Feature articles aren't assessed in that way, but most of the information in [2] still applies: it's well worth reading that page before you send any articles to *RadCom*.

Although initial draft and discussion material can be accepted in most common formats, we will need to receive the final material in a suitable form for reproduction. Usually this means text in a Word file with the minimum possible of formatting, and photos as per [1].

[1] www.rsgb.org/radcompix
[2] https://tinyurl.com/RC-guidance

MUF and a November storm

Geomagnetic storms often lower the maximum usable HF frequency. In some cases a major storm can *raise* the frequency limit.

Carl Leutzelschwab, K9LA points out that the electron density in the F2 region of the ionosphere can be enhanced by a storm in a variety of ways depending on latitude **[1]**. This temporarily increases the maximum usable frequency (MUF).

For example, a geomagnetic storm that affected MUF around the globe began at 0900UTC on 5 November 2023. **Figure 1** shows the three-hour averaged planetary K index levels, Kp *vs* UTC time **[2]**. The peak at 1500UTC on 5 November was followed by fluctuating G1-class activity ending at 0600UTC on 7 November. The NOAA severity scale levels at the right side of the figure show that the peak level is Strong (G3). **Table 1** shows the NOAA Space Weather Scale for Geomagnetic Storms **[3]**.

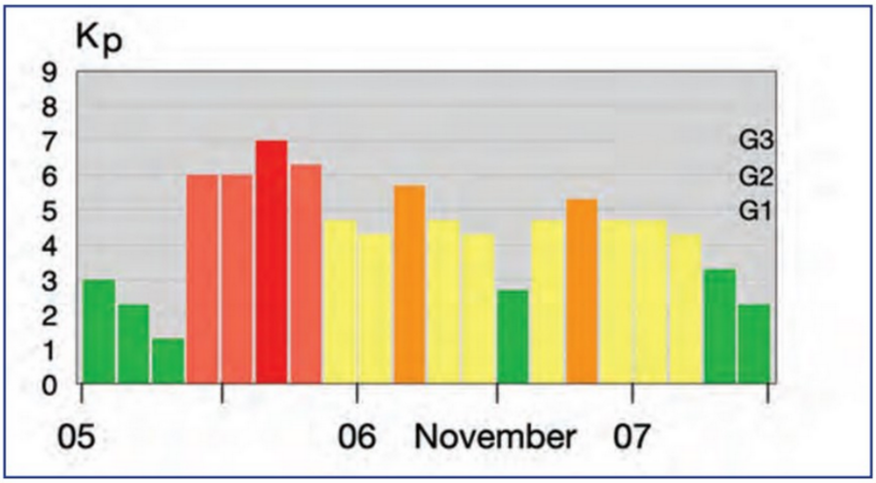

FIGURE 1: Three-hour averaged planetary K-index, Kp vs UTC time for 5-7 November, 2023, storm. Data from [2].

The global maps in **Figure 2** from a NOAA computer simulation show the greatest MUF effects, which are at 0000UTC on 6 November. Color contours in the top map show MUF is 50MHz or higher near the day/night terminator in the Central and North American sectors. The bottom map shows how much the MUF changed: an MUF anomaly (-25MHz to 25MHz) is the difference between MUF and the running ten-day average of MUF measured at the same time of day. Red regions of increased MUF and blue regions of decreased MUF are scattered around the globe.

The maps in Figure 2 are from the NOAA Whole Atmosphere Model (WAM-IPE), a global 3D model with altitude extending from 90km to 10,000km **[4]**. On the NOAA website you can view animated maps for the previous 24 hours as well as two-day forecasts. The MUF is for single-hop transmissions between stations 3000km apart, and the values are shown at the mid-points between stations. For comparison, MUF for 3000km is three times the critical frequency (foF2) measured with an ionosonde **[5]**.

The NOAA model results in Figure 2 show that this storm had no significant effect on MUF in the UK region. Measurements from the Chilton ionosonde confirm this. You can use the online PropQuest application by Jim Bacon, G3YLA, to compare measurements of foF2 during the storm with any previous day **[6]**. Long range transmissions to and from UK stations may have been affected if there was an MUF anomaly where the signal was refracted.

Storm-related changes in HF propagation depend on latitude, the season of the year, and the local time when a storm begins. A single storm can have both positive (higher MUF) and negative (lower MUF) phases. At middle latitudes, including the UK, storms that mainly lower MUF are more frequent than storms that raise MUF.

From K9LA: "So always keep an ear open during geomagnetic storms. If you're in the right place at the right time, you may be pleasantly surprised" [1].

TABLE 1: NOAA Space Weather Scale for Geomagnetic Storms [3].

Level		Kp
G5	Extreme	9
G4	Severe	8
G3	Strong	7
G2	Moderate	6
G1	Minor	5

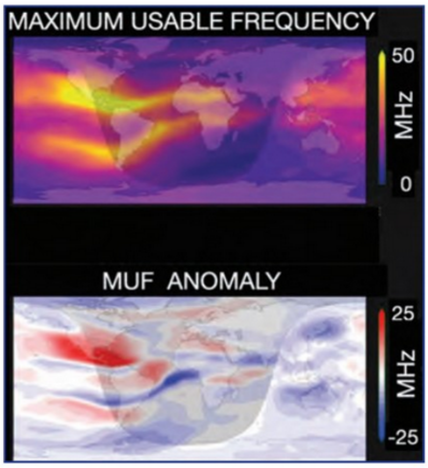

FIGURE 2: Top: Global map of simulated maximum usable frequencies (0-50MHz) at 0000UTC on 6 November. Bottom: Global map of simulated maximum usable frequency anomalies (-25 to +25 MHz) [4].

References

[1] Carl Leutzelschwab, K9LA, Electron density changes during geomagnetic storms. https://k9la.us/Jan18_Bonus_-_Electron_Density_Changes_During_Geomagnetic_Storms.pdf

[2] https://swpc.noaa.gov/products/planetary-k-index

[3] https://www.swpc.noaa.gov/noaa-scales-explanation

[4] https://www.swpc.noaa.gov/products/wam-ipe

[5] Steve Nichols, G0KYA, *Radio Propagation Explained*, RSGB.

[6] https://www.propquest.co.uk/graphs.php

Peter DeNeef, AE7PD
HamRadioAndVision@gmail.com

Icom ID-50E
handheld transceiver

Recently, I was lucky enough to have the chance to review the top-of-the-range Icom ID-52 2m/70cm D-Star and FM handheld transceiver.

At the time, I commented that Icom had just two VHF/UHF handheld transceivers for the amateur market, the IC-T10 which is a basic FM-only handheld transceiver for 2m and 70cm, priced at around £199, and the top-of-the-range ID-52, giving D-Star and FM on 2m/70cm at £550 – with nothing in the middle! The newly-announced ID-50E goes someway to plugging that gap, although priced at around £450, it's towards the higher end.

First impressions

Taking the rig out of the box, the first impression is of a solid and nicely-built unit (see **Figure 1**). There's no charger included in the box, because the rig can be charged using the USB-C port on the radio, and a cable is supplied for connection to a phone charger or a USB port on a computer. Getting the rig going was fairly straightforward, and I quickly found my way around the menu, with limited recourse to the instruction manual. A basic manual is included with the radio, but you can download an advanced manual from the Icom UK website at **[1]**.

The first test was to use my multi-mode digital voice modem (MMDVM) hotspot to connect to one of the D-Star reflectors, REF030C. I tuned the ID-50 to the hotspot's frequency, and very quickly I was hearing traffic from the reflector on the rig, as well as seeing callsign and position data associated with amateur transmissions. The ID-50 has a built-in GPS, so there's a very useful screen for when you are out an about which gives you your position in both latitude/longitude and maidenhead-square-locator format. If the ID-50 receives a D-Star transmission containing position data, it will automatically calculate the distance and bearing from you to that station, which can be quite interesting. The ID-50 can automatically announce the callsign of a station in audio, which can be useful and interesting if you are monitoring and can't see the screen of the rig. The display of the ID-50 is a good size but is monochrome (compared to the colour display on the ID-52). It's fairly easy to read, although suffers a little in some light conditions.

FIGURE 1: The ID-50E is a solid and nicely-built unit.

Having listened to the D-Star traffic through the hotspot, I wanted to transmit. Because it is a simplex hotspot, I found, as I have done before with Icom D-Star radios, that I needed to set the rig to duplex, but with a zero offset (so the rig is transmitting and receiving on the same frequency). Trying to use the rig on the simplex setting meant that I was not heard through the hotspot, but as soon as I set it to duplex, everything worked fine. I was able to make a couple of quick QSOs around the world through the hotspot to assure myself that everything was working well.

Like the ID-52, one of the nice things about the ID-50 is that it comes programmed with frequencies and information for both analogue and D-Star repeaters in the UK. Better still, because the GPS tells the rig where you are, it's able to offer a 'nearby repeater' function. This is very handy if you are in an unfamiliar area. The only disappointing thing, from my perspective, was that the repeater lists were solely UK-based. Living here on the coast in West Wales, some of my most accessible repeaters are across the water in Ireland. On the other hand, it's easy enough to add your own custom entries to overcome this issue. The repeater lists can be refreshed by downloading them from the Icom UK website **[2]** and transferring them onto the rig by means of the microSD card.

Of course, as well as D-Star, the ID-50 works on FM too and works well. Received audio was of good quality and had a nice full tone, especially considering the size of the loudspeaker in the rig. The receiver seemed quite sensitive, and I was easily able to pick up the 70cm repeater on Mount Leinster, Ireland, some 80 miles or so away, with the handheld connected to the supplied 'rubber duck antenna' in the shack, though this is nothing exceptional as there's not much in the way! With the repeater directory included on the rig, you won't generally have to do much setting up CTCSS tones and the like, but I found it quite straightforward to set up the ID-50 to access my Allstar analogue hotspot.

Receiver

The ID-50's receiver covers the FM broadcast band from 76MHz to 108MHz, and then from 108MHz to 174MHz (including AM on the civil air band), then from 375MHz to 479MHz. The ID-50's receiver incorporates a band scope, allowing you to see activity on nearby channels. This seemed to work quite well, and it was interesting to see signals pop up when monitoring. This might be useful on the air band as well. The band scope has three different modes: 'centre', 'scroll' and 'scroll (fix)'. The only issue I found here was that, when the backlight went off after a few seconds, it was easy to miss some activity being displayed on the band scope. However, you can easily configure the backlight in the 'Set/Display' menu to stay on permanently.

If you have a microSD card installed in the ID-50, you can record QSO audio. Recordings can either be played back on the ID-50 or, using the microSD card, you can move them over to your PC and listen to them

FIGURE 2: **Images may be sent and received on the ID-50E.**

there. You can also set up a voice-transmit file, to call CQ for example. I can't quite envisage when this would be useful on an FM/D-Star rig, but you never know.

Images

When I reviewed the ID-52, one of the things that I enjoyed, slightly to my surprise, was the ability to exchange pictures via D-Star. Essentially, it's a file-transfer utility (see **Figures 2** and **3**). With the ID-52's Bluetooth capability, it was a simple matter to install the appropriate application onto a Smartphone. The ID-50 does not have Bluetooth however, so you'll need to use a cable. If you have an Android phone with a USB-C port, you'll be able to install the ST-ID50A application which allows transfer of images from the phone to the ID-50. My Android device doesn't have a USB-C port, so I needed to use the Windows utility ST-ID50W **[3]**. This installed readily on a Windows-10 machine, and the USB cable which comes with the ID-50 can be used for the connection to the ID-50. Make sure you have a microSD card inserted into the ID-50, then plug the USB-C cable into the port on the radio with the other end of the USB lead going to a USB port on your computer. Windows will install an ID-50 driver, setting up a COM port, which you can then select in the ST-ID50W utility. You can use this utility to add text to a picture, rotate it, and upload it to the microSD card in the ID-50. Once the picture is on the SD card, you can transmit it, by selecting 'picture mode'. Received pictures are saved to the SD card and you can synchronise these with the PC using the ST-ID50W utility. There's a D-Star picture exchange net on REF055D on a Wednesday night at 1800UTC.

D-Star gateway function

Another interesting feature, included in Icom D-Star radios of late, is the D-Star gateway function. The idea behind this is that you don't need a hotspot to connect to the D-Star network, repeaters and reflectors. In the case of the ID-50, you'll need the USB cable supplied with the radio, and a Windows PC running the RS-MS3W software. This can be freely downloaded from the Icom Japan website **[4]**. If you have an Android device with a USB-C port, you can download the RS-MS3A app from the Google Play store. I used the Windows software and, unfortunately, the experience was not straightforward.

Software

Installing the software onto the PC was fine, as was installing the USB driver. According to the instructions, you then need to do some 'port forwarding' on your router, to allow UDP port 40000 to pass to the PC running the gateway software (although I was subsequently told this is not necessary). Then, you have to configure the gateway software, which requires a fair amount of research and experimentation. Although there are good instructions for the gateway software, they don't help with what values you need to enter in certain fields, or what servers might accept your connections. Even after considerable experimentation, I was still getting errors! I spoke to Icom UK who were helpful, suggesting I use their gateway server gb7ic.icomuk.co.uk to connect through. It seemed that I could use callsign routing to connect to repeaters, but I struggled with connecting to a reflector, although Icom UK thought it should work.

The gateway mode allows you to use the ID-50 as a speaker/microphone on the D-Star network, with the traffic all going through the PC or mobile phone. There's also an 'access point' mode, which allows you to use the ID-50 as a mini D-Star gateway/repeater, connected to the Internet. You'll need another D-Star radio to communicate with the ID-50. Given the expense involved, I suspect most people will use 'terminal' mode instead.

Although the idea behind the software is excellent, and should, in theory, allow users to connect to the D-Star network from almost anywhere there's an internet connection, setting it up seems tricky. I think most people would find it easier to use a hotspot. It's fair to say that this is true of the D-Star gateway feature across all the Icom D-Star radios and is not specific to the ID-50.

Battery life with the ID-50 seemed good. Admittedly, most of my operation was using 'super low power' and the hotspot and the battery lasted a very long time as you might expect. You can select between transmit power levels of 5W, 2.5W, 1W, and 0.1W. I was impressed with the USB-C charging; there was no need for yet another handheld charger under the bench.

The scanning speed on the ID-50 was excellent – one of the faster scanning speeds I've seen on a device that wasn't a dedicated scanner. You also have the ability to select a group of memories that you want to scan, so if you organise your memories, you can store, say, air-band frequencies in one group and amateur bands in another. That way, if all you want to do one day is to listen to aircraft, then so you can. Of course, you can skip channels and so on.

FIGURE 3: **Another example of a received image.**

Overall

Did I like the ID-50? I did. It has a lot of good functionality. Compared to the ID-52, I missed the colour screen and the Bluetooth capability, so if you are in the market for a D-Star handheld you will have to decide whether those features (and others) are worth around the £100 extra. From an ease-of-operation point of view, it's fairly easy to find your way around the rig, and the inbuilt repeater directory removes the need for a lot of programming. I don't use D-Star as often as perhaps I should, and am always interested in some of the conversations I hear on some of the reflectors, and I enjoy joining in. Having the ID-50 here prompted me to do this a little more.

The rig has IPX7 waterproofing, meaning that it could stand being submerged in a metre of water for 30 minutes (not tested by me!) so would suit operation in the outdoors. At around £450, the ID-50 is one of the more-expensive handhelds on the market. To get the best out of it, you'll probably need to be a committed D-Star user, either with your own hotspot or a D-Star repeater in your locality. My grateful thanks to the team at Icom UK for the loan of the ID-50 and for their willingness to answer my questions.

References

[1] https://icomuk.co.uk/files/icom/PDF/advancedManuals/ID-50A_50E_ENG_AM_0.pdf
[2] http://downloads.d-staruk.co.uk/
[3] https://www.icomjapan.com/support/firmware_driver/3918/
[4] https://www.icomjapan.com/lineup/options/RS-MS3W/?open=4#detail_content

Tim Kirby, GW4VXE
gw4vxe@icloud.com

Design Notes

Z80 / BBC BASIC computer reborn

Steve, G1KQH emailed to mention a new single-board computer that has been released. He states: "The AgonLight2 8-bit computer runs BBC BASIC and other stuff, and it's cheap too at around £50. If I remember back in the BBC days, there was a load of amateur radio software written for the BBC, so some stuff might be worth porting over to run on the Agonlight2. There is also a support group on Facebook." It appears the AgonLight2 is a development of an earlier single-board computer based on a Z80 processor that ran BBC BASIC. This is somewhat interesting as the BBC Computer itself was built around a 6502 processor. To summarise from the website **[1]**, the AgonLight2 is a redesign of the original AgonLight board designed by Bernardo Kastrup – an 8-bit retro Z80 computer running BBC BASIC! The original was a complete single-board computer with VGA display output, USB (PS2) keyboard input, and an SD card slot acting as an external disk, with no need for an external computer to set it up. This new version includes USB-C for power, USB-A for your keyboard, LiPo battery connector and charging, an improved voltage regulator, and more enhancements. The firmware of the original AgonLight, by Dean Belfield, reproduces the BBC BASIC for the Z80. The firmware is open-source, available on GitHub, and also works on AgonLight2 without any modifications. Judging from the comments and reviews on the website, the board appears to be reasonably straightforward to get going, and should appeal to those who were familiar with these first generations of home computers and would like to be able to use their old software and programming skills. Unlike the Arduino and Raspberry-Pi, it doesn't need any supporting PC or web access, just a monitor, keyboard, SD card and anything else you choose to add.

Mains waveform shape

In the November issue of *RadCom*, there was a design for a mains monitor measuring true-RMS voltage and current. Included in that article were a number of plots showing a flat-topped shape to the mains waveform rather than the nicely-rounded sinusoidal shape that might be expected. An oscilloscope trace of my own supply can be seen in **Figure 1** showing exactly the same flat-topped, truncated sine wave. Many others report similar waveform shapes. It would appear that such a flat-topped, or squashed, waveform appears to be universal, at least throughout the UK. After much discussion about the cause of this, on the various internet forums, no definitive answer was to be found, but several reasons were suggested. One possibility is the massive preponderance of switched-mode power-supply units (SMPSU) that directly rectify the mains. Conventional bridge rectifiers, driving a reservoir capacitor, draw current in short spikes at the peak of the waveform. Any residual impedance in the supply line will result in extra voltage drop during these peaks, and so serve to flatten the voltage waveform. But it is hard to believe that this is sufficient to affect the whole of the Grid! In addition, modern SMPSUs, rated at more than a few hundred watts, are required to include power-factor correction (PFC) on their inputs. This is a special arrangement, done in a single chip using DSP techniques and a boost converter, to deliberately force the input rectification process to appear resistive, and so remove this very-peaky current draw. PFC has been in use for a couple of decades now, so it looks like we can't really blame SMPSUs for the flat-topping. A second possibility is that it is deliberate, possibly caused by the Grid distribution transformers running almost at saturation. The local ones, from the few-kilowatt-rated transformers mounted up poles, to the 500kVA or higher ones in substations, transform the 11kV distribution voltage to domestic 400/230V supplies. The laminations in all mains-frequency transformers use a silicon steel that goes into saturation quite rapidly above a magnetic flux density of around 1.3T. To minimise cost, the amount of copper in the windings, and special iron in the core, both need to be reduced as much as possible, which is achieved by running the core as close to magnetic saturation as possible. What harm does a non-sinusoidal waveform actually do? Probably very little in practice. All SMPSUs, and anything that directly rectifies the mains, actually run better with a flat-topped waveform, as it removes that current spike as capacitors charge. With all modern electric lighting being LED type with bridge rectifiers, it's a real advantage. AC motors and other loads, that would normally require a proper sinusoidal drive, will see this flat top as third- and higher-order harmonics, probably running at slightly reduced efficiency. But these days motors in refrigerators and washing machines aren't directly-driven AC types; they are more likely to be MOSFET-controlled DC motors.The bigger commercial users usually get their power delivered at 11kV or higher, and they look after their own transformers, so if they need a good sine wave, they can

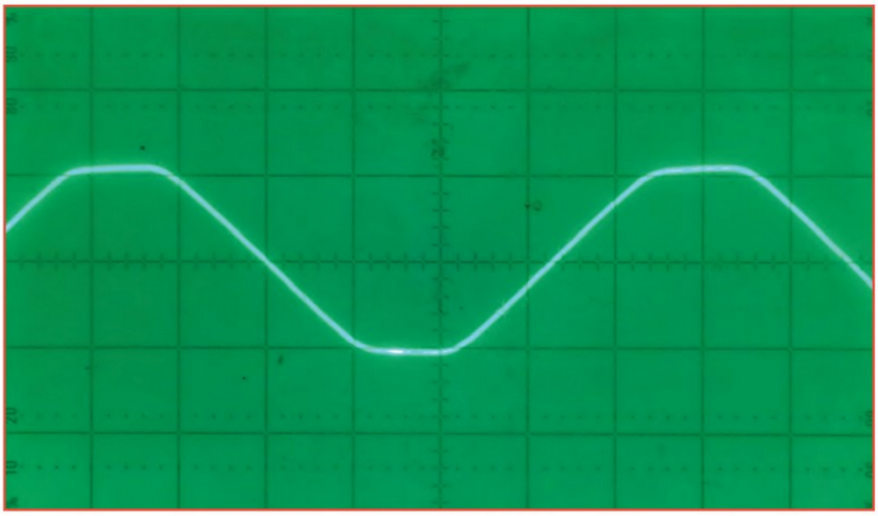

FIGURE 1: **The mains non-sinusoidal waveform here at 'JNT Labs. This appears to be typical throughout the UK.**

Andy Talbot, G4JNT
andy.g4jnt@gmail.com

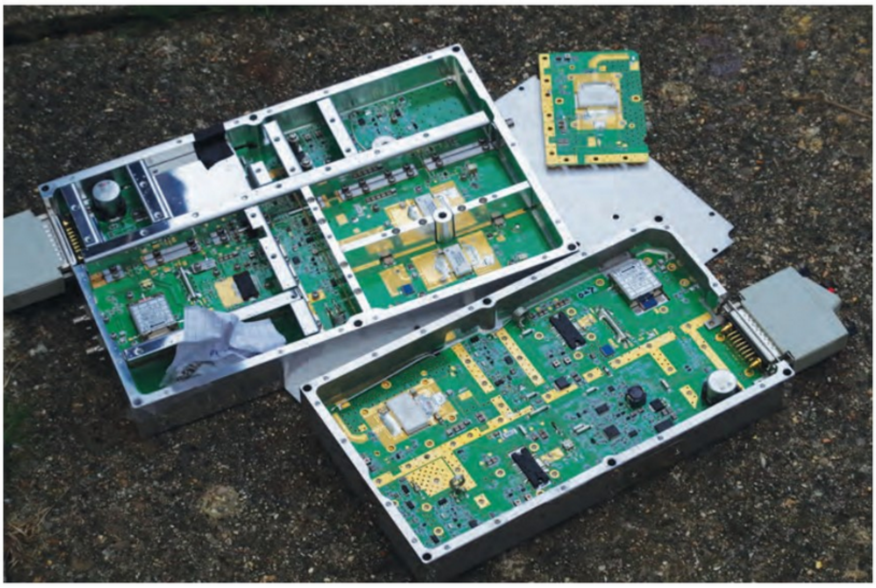

FIGURE 2: **Surplus modules from the mobile-phone market. These are worth dismantling to find RF devices for power amplifiers.**

ensure their transformer delivers it. It would be interesting to see what the waveform looks like on the 11kV network, whether it shows any sign of flat-topping. But it would, of course, be very foolish, and downright dangerous, and probably illegal, to probe the 11kV distribution network. So don't try to measure the high voltage waveform at home!

Transverters and crystals

The 2022 RSGB Convention included a talk by Sam, G4DDK on the benefit of using transverters, rather than dedicated transceivers, for the VHF and UHF bands, especially under contest and congested band conditions **[2]**. A transverter enables an HF radio to be used, mixing the frequency up for transmission and down for reception. This allows the better general ergonomics and strong-signal handling of most HF radios to be kept, instead of compromising with the more-basic functionality that most VHF-only transceivers can offer. Strong-signal handling and phase noise are the dominant technical issues that make this solution worthwhile. A good transverter uses a low-noise crystal for its local oscillator, mixing the VHF band often to 28MHz to be tuned on the HF driver. For the 2m band, a 116MHz crystal converts 144MHz-146MHz to 28MHz-30MHz. For 432MHz, a local oscillator frequency of 404MHz is needed for the same IF range, generated typically from a 101MHz crystal multiplied by four. But crystals for these obscure or specialist frequencies, indeed any crystals for made-to-order frequencies, are becoming difficult to find, and are getting very expensive. Simon, G8DMN in a post to one of the groups, mentioned this difficulty in obtaining high-specification transverter crystals. As a bit of lateral thinking, he suggested that an off-the-shelf crystal oscillator for 100MHz could be used instead, producing an IF at 44MHz-46MHz. Modern HF transceivers usually offer continuous coverage to 50MHz, and the problem of asking it to transmit out of band is probably not an issue if the low-level transverter drive port on these transceivers is used for the transmitter drive RF. It was also pointed out that there is one website that does offer TCXOs at various useful frequencies **[3]**.

Medium-power RF devices

Still on the subject of transverters, many kits and ready-built modules generate a transmit power of just a few tens or hundreds of mW. There is a need to amplify this up to a few W for general use, or to drive a high-power amplifier, the latter typically requiring around 4W-6W drive power. It is getting more and more difficult to find devices for these power levels for the VHF and UHF bands. The 'Mitsubishi Bricks', that were so popular as transmitter medium-power stages, are obsolete and getting harder to find; those that are available are attracting higher prices than ever. There are fewer and fewer discrete transistors for operation at this sort of power level. Here at 'JNT Labs, I acquired a stock of MRF134-MRF137 MOSFETs many years ago by dismantling surplus equipment. These are all rated in the 'several watt' region, run from a 28V supply, and are an absolute doddle to use as power amplifiers at HF, and up to at least 144MHz, and even up to 432MHz. These very same surplus devices have been used in the GB3ORK 5MHz beacon, a 50MHz beacon still in current use somewhere in the world, and most recently as driver stages in both the breadboard and operational GB3MBA 50MHz meteor-beacon power amplifiers. I now have no more left. At this power level, few more-recent devices can be found. There are many cases of special integrated packages and ICs for low-power transmitters that appear. Then, as soon as we try to buy them, production is ended and they are no longer obtainable. Such ICs go out of production faster than they appear in the catalogues, so it is no good for transverter designers or builders trying to specify or incorporate any nice-looking modern device appearing in the catalogues, if their design is to survive more than a few months.There are devices out there, bipolars and MOSFETs, discrete and ICs, that are designed for the high frequencies used in the mobile-phone bands. Common sense suggests that most of these will have capability at much lower frequencies for a similar power output, provided surrounding components, such as for coupling and decoupling, are designed to suit. A device rated for 10W at 1GHz ought to be quite capable of delivering 5W at VHF. The snag comes in getting it stable. A high-frequency amplifier device is quite likely to have a very high gain, and be difficult to match at a tenth of its design frequency. But, with a bit of feedback, resistive slugging, and some trial and error, it ought to be possible. And once a design is out there, it's probably going to be possible to drop in any surplus or current suitably-rated device, and only make a few minor tweaks around a standard layout. Once upon a time, radio amateurs pushed devices to beyond their limits, usually trying to make them work higher in frequency or power than they were designed for. Nowadays, why can't we do the same, only now pushing them lower in frequency than they were designed for? Next rally or boot sale, look for that scrap mobile-phone base-station module that no one wants, like the ones shown in **Figure 2**. Offer the seller a couple of pounds to save them the effort of their having to junk it, get inside, and extract the power devices. Try making a breadboard PA. Think of your predecessors using valves designed for medium-wave radios, taking off the bases, and wiring them in to get operational at HF, and do the modern equivalent. See what you can do, and let me know how you get on.

PCB vias and foil cutting

Peter, G8EZE wrote in: "I have long been a fan of the construction method using copper or brass foil as through-board vias on homebrew PCBs **[4]**, rather than pins, wire,

FIGURE 3: Versions of the PICkit-3 programmer.

or griplets, but have always found cutting 0.8mm-1.0mm wide strips difficult with scissors, usually winding up with a taper even over short lengths (I have exactly the same problem here at 'JNT labs. I just live with tapered strips).

"I have recently been trying out a miniature paper guillotine with a 16cm blade, and find that I can accurately cut long strips with it. I have used copper foil/tape with the adhesive stripped off in the past, but prefer 0.1mm brass shim which will take a right angle bend at the end and hang in place through the hole for the first side to be soldered without fluttering away. The foil can then be bent flat on the second side and trimmed with a craft knife before soldering.

"These guillotines seem to be readily available online for £15-£20. Obviously they are not as robust as the professional versions, but are good enough for occasional use and compact enough (210 x 130 x 40mm) to not need much storage space."

PICkit programmers and support

Design Notes for January 2022 described some issues with, and availability of, the PICkit-3 programmer for PIC microcontrollers. The original PICkit-3 marketed by Microchip has been obsolete for many years, having been replaced by the PICkit-4, PICkit-5, and other programmers. The advantage of the PICkit-3, that still makes it the preferred choice for many PIC users, is that it comes with a standalone driver package, meaning the over-bloated and complicated MPLAB programming-support suite doesn't have to be used. Many of us just like to write code using a simple text editor, a standalone assembler, and then blow it into the device; this is not possible using the later PICkit devices. The enduring nature of the PICkit-3 means there are now a number of clones out there, some better than others, as detailed in that earlier column. The fact that Microchip made the hardware design and the firmware openly available meant that nearly all the clones and copies were usually properly built and used the correct firmware, even if they did compromise a bit on protection of I/O lines, hence the longevity of this particular programmer, several versions of which can be seen in **Figure 3**. But now there is another problem with this setup. The standalone PICkit-3 driver software cannot cope with some of the latest-generation PIC devices. I first met this issue when blowing the Scriber software by G8HAJ for Hellschreiber into a 16F1615. The PICkit-3 software just couldn't cope with that device, and I was forced to use MPLAB, which is quite a challenge if you just want to blow in someone's already-compiled code! But there appears to be a solution. Colin, G3YHV stated in the RSGB-Workshop Group that the PICkit+ software is worth a look **[5]**. This is described as: "PICKITPLUS – a programming software that revitalises the outdated PICkit-2 and PICkit-3". It is a total replacement for Microchip's standalone driver packages, and can work with both original and clone programmers. It can cope with all old and new devices, and there is an assurance that it will be kept up to date. It can also do DSPics, another application that otherwise mandates MPLAB. There are different versions of PICkit+, for the PICkit-2 and PICkit-3 programmers, and also one (that works with both) that can be driven from the command line. See **[6]** for more details on purchasing. The price for UK buyers appears to be £23. It looks to be well worth the price if it can guarantee forward compatibility, and not have to keep changing the programmer firmware to swap between standalone or MPLAB software.

And on the subject of PICkits ...

Neil, G4DBN commented in the RSGB-Workshop Group: "I'm seeing some horror stories about Microchip cables for the PICkit-4 and PICkit-5 to adapt the new 8-pin socket to the old-style 6-pin connector. It seems they are often under-crimped and end up spreading the sockets on the PICkit and making them unusable on standard pin headers.

"I don't have a fancy crimp tool for these little connectors, so looking for a recommended non-Microchip interface cable to go from 8-pin male to fit the PICkit to a 6-pin female to connect to the enclosed male connector on the old PIC boards."

Amplitude keying the AD9850 DDS

Just as this column was being prepared, Greg Smith, ZS1IX, G3ZZI, ZL3IX, contacted me to say he is designing the driver and A1 keying circuitry for a 28MHz beacon in Cape Town. The source is an AD9850 DDS chip, and he was querying the wisdom of on-off keying it by programming it with zero Hz for the off state. He suspected key clicks would be an issue, and indeed they certainly would be. There would be no way to filter the resulting output, and spectral spreading would be horrendous, so that idea was a non-starter.

He then suggested controlling the I_{OUT} current from pin 12 of the AD9850 device, using a Gaussian or similar-shaped waveform to control the keyed output envelope. The current sourced from this pin to ground defines the RF peak output from the DDS, so it is a convenient way to control the RF level with a linear relationship between I_{OUT} and RF output. Normally, a resistor is placed from here to ground, so the voltage imposed across it, nominally 1.25V, defines the current through the resistor, but there is no reason why external circuitry cannot be used to force any particular value of sink current. That will need a bit of circuitry that works as a current sink, and does not attempt to 'push' or force the voltage on the pin; this would no doubt be possible with an op-amp and a bit of lateral thinking for the circuit design. Greg is looking at it and will report his findings.

An alternative is to use a diode, or double-balanced mixer, as described in Design Notes for June 2014. There it can be seen that a remarkably-linear attenuation vs voltage relationship is possible over at least a 40dB range.

References

[1] https://thepihut.com/products/agonlight2-z80-bbc-basic-retro-single-board-computer AgonLight2 BBC Computer lookalike

[2] https://www.youtube.com/watch?v=0-L4FBEhqgc G4DDK's presentation 'Transverters revisited', at the 2022 RSGB Convention

[3] https://fr.aliexpress.com/item/32794136550.html?gatewayAdapt=glo2fra%20TCXO

[4] Design Notes, *RadCom* October 2015

[5] https://www.electronics-lab.com/PICkitplus-software-enhances-outdated-PICkit2-PICkit3/

[6] https://PICkitplus.co.uk/Typesetter/

Technical

De-bounce

switches using software-defined Schmitt triggers

FIGURE 1: Oscilloscope traces showing the voltage before de-bouncing (magenta trace), and after de-bouncing (yellow trace).

```
bool Schmitt = false;
const int threshold_high = 600;  // in counts
const int threshold_low = 400;   // in counts
const int LED_PIN = 9;

void setup() {
  pinMode(LED_PIN, OUTPUT);
}

void loop() {
  // Better code suggested by Google Bard
  // Read input from switch (assumed to be connected to A7 pin)
  int A7_input = analogRead(A7);  // assuming 10-bit ADC on 5V reference
  // Check if the switch state has changed
  if ((A7_input > threshold_high) && !Schmitt) {
    Schmitt = true;
  } else if ((A7_input < threshold_low) && Schmitt) {
    Schmitt = false;
  }

  // Turn on or off LED based on boolean variable
  digitalWrite(LED_PIN, Schmitt);
```

FIGURE 3: Arduino code snippet.

Mechanical switches, sensors and rotary encoders do not normally produce clean and noise-free signals that can be used reliably by discrete logic, or microcontroller, circuits.

The state of a mechanical device can be incorrectly captured if software-driven polling is used as the state can be indeterminate at the instant of polling. Worse still, if the mechanical device is connected to a counter, either hardware or software, the noise produced can be interpreted as multiple events.

The magenta trace shown in **Figure 1** is the signal produced by a push-button switch connected to +5V via a 10KΩ resistor, as shown in **Figure 2**, and the noise generated by the metal contacts rubbing against each other is clear. A microcontroller unit (MCU), such as the Microchip ATMega328, is used on many commercial boards such as the Arduino UNO and Nano. Both digital and analogue input-output (I/O) pins are featured on these microcontrollers, and it is the digital inputs that are normally used to interface with mechanical switch-type devices. It is easy to find Arduino code on the internet that de-bounces switches using time delays to allow switch-induced noise to be ignored – just ask ChatGPT! But it can be difficult to configure the timing, and sometimes a trade-off has to be made between reliability and speed. If digital de-bounce software is either not adequate or reliable enough, the designer will often resort to a conventional analogue solution using passive RC networks and Schmitt trigger logic ICs such as a 74HC14 hex Schmitt device, using additional passive and active components and adding to the size and cost.

I wanted to create a software solution that would emulate an RC Schmitt circuit with the minimum of passive components.

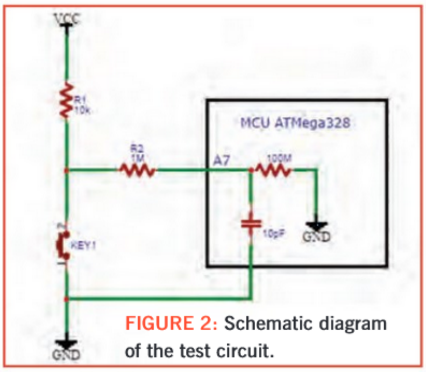

FIGURE 2: Schematic diagram of the test circuit.

My solution

First, I developed a software implementation of a Schmitt trigger with high- and low-voltage switching points, and therefore with an element of hysteresis. An analogue input on the MCU is used with the code shown in **Figure 3**.

A Boolean variable called Schmitt is defined in the program, which can have one of the two binary values TRUE (1) or FALSE (0). Also, two analogue thresholds are defined in terms of count values that can be compared with the number produced by the MCU's 10-bit analogue-to-digital convertor (ADC) with a range of 0 to 1023 equating to 0 and 5VDC respectively. The defined thresholds equate to approximately 2.9VDC and 1.9VDC, ie above and below the half-way voltage of 2.5VDC, and with 1VDC of hysteresis. The code listed in Figure 3 just drives an LED on digital-pin number 9 of the MCU, but could be used with the built-in timers and counters in the MCU, or various other software routines, eg to assess the steps and rotational direction of a rotary encoder.

Next came the problem of how to implement a low-pass RC filter to remove high-frequency switching noise. The obvious solution would be to use a resistor and a capacitor, but as a devout Yorkshire man I wondered if the cost of either could be eliminated? And yes – looking at the data sheet for the ATMega328, the analogue input pins are connected to a CMOS comparator with an input resistance of 100MΩ and capacitance of about 10pF. So the ADC input pin can be fed from a high-value resistor without significantly affecting the voltages from the mechanical switch ie 0V or 5V. Using a high-value resistor, eg 1MΩ, creates an effective low-pass filter just using the 10pF capacitance of the analogue input pin, with a time constant of 10pF x 1MΩ = 10µs or, in the frequency domain, a cut-off frequency of 15.9KHz.

The combination of a passive RC filter, together with a software-defined Schmitt trigger, is very effective at eliminating noise and de-bouncing switches as is shown by the yellow trace of Figure 1. If a longer time constant is needed, I have used a 15MΩ resistor on the analogue input pin, which gives a time constant of 150µs and a cut-off frequency of 1 KHz.

Space and cost are reduced in this solution as one resistor, one capacitor, and a logic IC have been replaced by a single resistor and a few lines of code. The saving in components, and the flexibility offered by a software Schmitt trigger, is even more apparent if interfacing to a rotary encoder with built-in push-button functions, as three such inputs would then be required. You can vary the value of the input resistor and the software-defined Schmitt thresholds to suit your particular application.

Web resource

Project materials are available for download at: https://rebrand.ly/sdst-m0yyt

Steven Bennett, M0YYT
m0yyt@outlook.com

Book Review

RadCom Team radcom@rsgb.org.uk

Best of 2023's books

RSGB Radio Communication Handbook, 15th edition

Edited by Ed Durrant, G8GLM/DD5LP

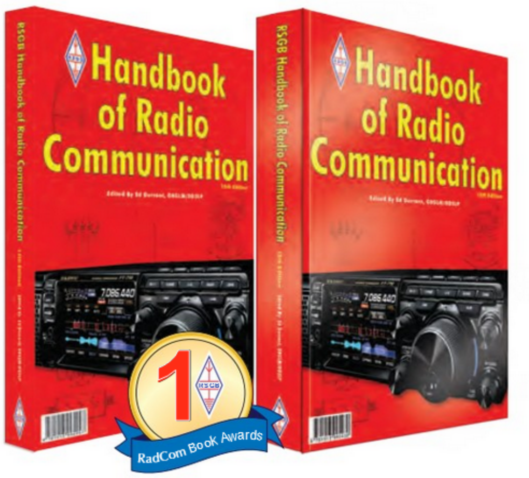

The RSGB *Handbook of Radio Communication* is regarded by many as being the holy grail of information on amateur radio. The 15th edition doesn't disappoint, with in-depth coverage over 25 chapters which can now also be purchased individually online if you are not seduced by the full, printed volume. Every chapter has been revisited and updated in line with current levels of subject knowledge and available equipment. The level of detail provided in each chapter is both graphically and textually enhanced and provides an exciting and informative read.

Full coverage of SDR radios is provided, including discussions on the merits of this technology and its various uses such as the use of SDR as test equipment, direct digital synthesis, use of panadapters and the factors affecting the dynamic range of SDR receivers. There is even a new section on using SDR for astronomy.

A very useful section on DMR and D-Star sets out the pros and cons associated with the use of digital voice modes. All the primary systems are covered including Wires-X and Yaesu System Fusion. The use of hotspot apps is also covered.

To accompany some of the projects in the book, a set of PCB templates are included which can be used to produce your own devices.

The book contains many more sections (too many to include here) including information on, for example, antennas, QRP operating and transmission line theory.

With over 800 pages, this is one of the best all-round technical books for the shack. It is recommended to all amateurs.

The Dirty Tricks Department

By John Lisle

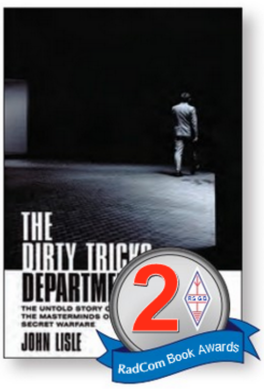

This entertaining book charts the story of the Office of Strategic Services (OSS) which later became the CIA and its joint working with the British Special Operations Executive (SOE). It is an insight into how the British and American governments decided to form intelligence and fighting units that could operate 'outside of the box' and create havoc amongst occupied territories.

Some of the ideas seem like they are straight out of a James Bond film script but many of the described devices were put to highly-effective use.

The book also deals with the moral dilemmas operatives faced when leaving the conventional rules of warfare behind them. As well as this, many unconventional techniques are included. These include psychological warfare, chemical warfare, truth drugs and exploding rats, to name a few!

This is an immensely entertaining book which shines a spotlight on a previously-closed subject.

The Radio Today Guide to the Xiegu X6100

By Andrew Barron, ZL3DW

Many readers enjoy Andrew Barron's books and report that they are often much easier to understand than the standard manuals that come with transceivers. This edition to the suite is clear and concise with well-presented pages of instruction on how to set up and program just about all aspects of the Xiegu X6100 radio. Among the comprehensive details are settings for linking the radio to a computer – information that will inevitably be valuable to the modern amateur.

The section on digital modes is interesting as it not only covers all the main modes but also provides a comparison on which versions are better to use in practice.

The book is an excellent read and is a useful companion in getting to know, and getting the most out of, your X6100.

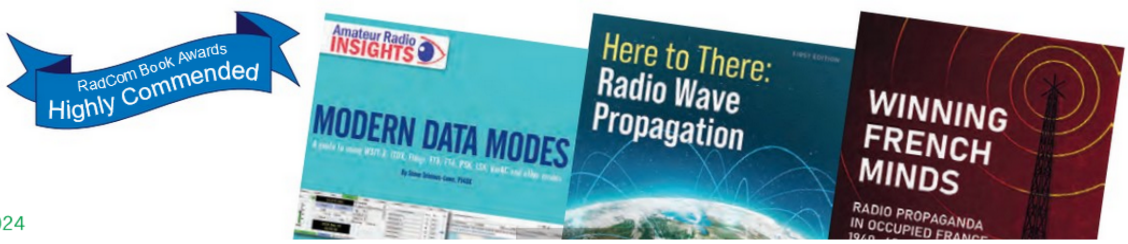

What actually goes on inside a digital filter?

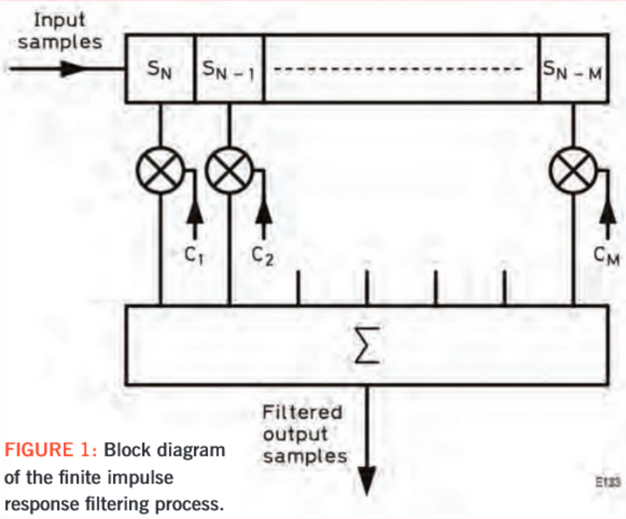

FIGURE 1: Block diagram of the finite impulse response filtering process.

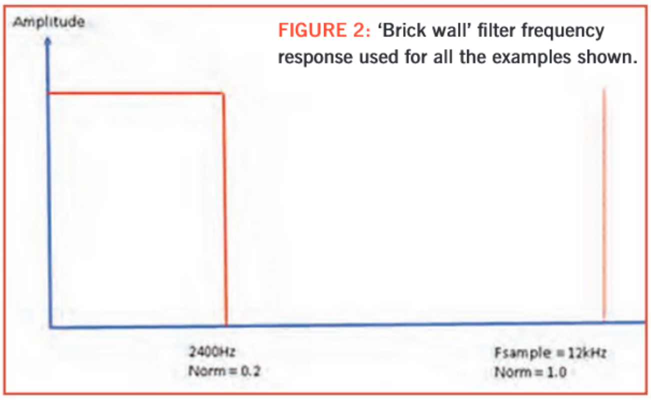

FIGURE 2: 'Brick wall' filter frequency response used for all the examples shown.

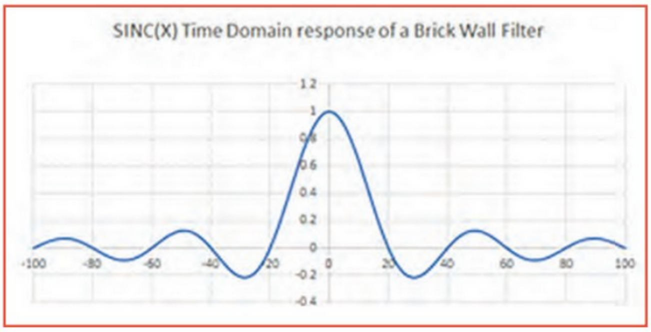

FIGURE 3: The normalised sinc(x) function, which is the time-domain equivalent of the filter in Figure 2.

All modern transceivers employ digital filtering of some form or another, but few radio amateurs really know how the digital filtering process works.

Here we will look at it from a hardware and what-it-does perspective, not worrying over-much about deep theory. 'Hardware' is perhaps not the best term to use since, as soon as the input signal is digitised in the input analogue to digital converter (ADC), it just becomes a series of numbers handled by a processor chip, stored, manipulated and shuffled around in registers or memory locations. The registers are real transistors in the silicon chip so perhaps hardware is as good a term as any.

Numbers

To get away from classic descriptions, we will try not to use too many algebraic representations here; real numbers will be quoted. But it does need to be said that the numbers used are just a representation of a particular case. They will change to whatever is needed, so '201' used below could just as easily be '65' or '511' or '8193'; 12kHz sampling could just as easily be 8kHz or 31.25kHz etc.

The hardware

Figure 1 shows the outline of a particular type of digital filter, the so-called finite impulse response (FIR) type. Infinite impulse types do exist, but they are used less frequently, and we don't need to study them here. The analogue signal arrives from the left-hand side, and is converted into a series of numbers by the ADC. The conversion is done at a sample rate that has to be at least twice the highest frequency in the input – the Nyquist criteria – and, for practical reasons, the sampling rate will often be perhaps four or more times the highest component in the waveform. For speech, a sampling rate of 12kHz is a firm favourite as it is comfortably above the 3kHz normally considered to be the maximum required. 48kHz is another common value, allowing audio up to beyond 14kHz to be worked on with a comfortable margin. In general, the higher the sampling rate the better, but this comes with the need for faster processing.

The series of numbers, the samples S, coming out from the ADC are pushed sequentially into a buffer, shown by the series of boxes in Figure 1. Each box is a memory location with a width, or number of bits, sufficient to store the ADC samples with adequate precision. In many audio digital signal processors (DSP), the values are often turned into floating point values in the processor so the range of numbers from a 16-bit soundcard from -32768 to +32767 may be scaled to floating point values in the range -1.00000 to +0.99997. Modern DSPs are designed to work with floating point values, so there is no processing speed disadvantage in doing this. But it doesn't have to be like this – low-cost simple DSP systems can stay with 'digital values', or integers, throughout.

As each digitised sample arrives from the ADC, it is put into the left-hand buffer cell with the previous samples (say 200 of them) moved one to the right and the oldest discarded as it is of no further value. The result is that at each sampling interval, 12000 times every second if we are sampling at 12kHz, we have a history of the last

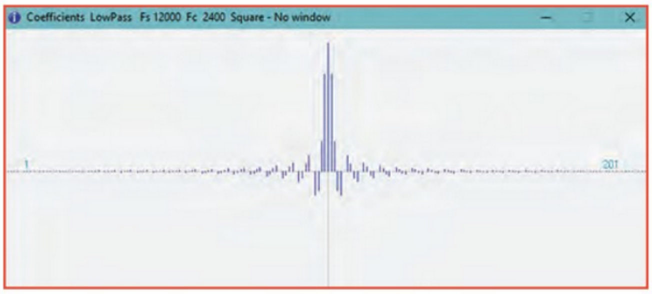

FIGURE 4: **The sampled version of Figure 3 with a uniform 'window' applied in the Fourier transform.**

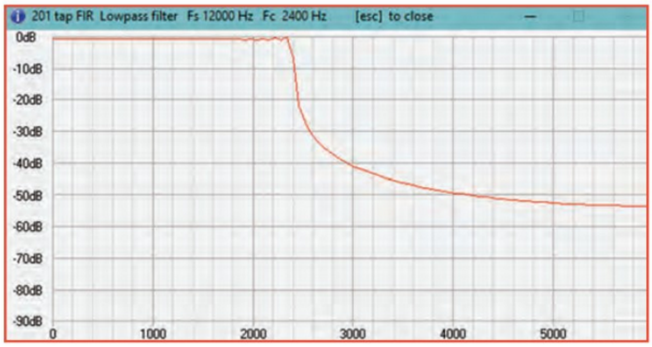

FIGURE 5: **The calculated frequency response of the digital filter with a uniform 'window' applied in the Fourier transform, with 201 coefficients.**

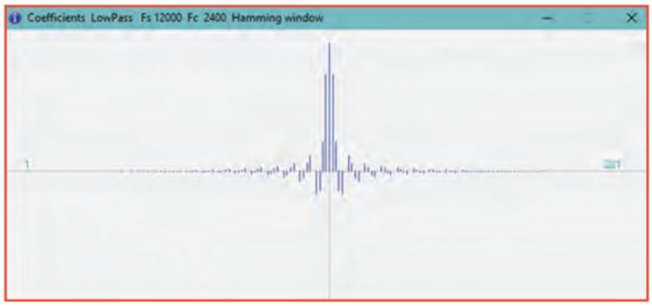

FIGURE 6: **The sampled version of Figure 3 with a 'Hamming window' applied in the Fourier transform.**

201 samples of the input waveform; this is about 16ms worth of the waveform.

Doing the filtering

As a new sample is pushed into the buffer, 12000 times per second, the contents of each of the 201 buffer cells is multiplied by a coefficient C and the results of each of the 201 multiplications are added. There are 201 separate coefficients made up from a set of fixed values that are predetermined for the filter type; how they are defined is covered later. The single number that is a result of the 201 multiply-and-accumulate operations, (called MAC in DSP devices), is the filtered output sample. It can be sent to a digital/analogue converter (DAC), an amplifier and a loudspeaker, and we have a DSP audio filter box.

The output value is updated 12000 times per second but, because is it made up of the last 201 samples of the input, the filtered output sample contains some part of the input signal going back 16 milliseconds. How much of each of the historical samples contribute is a direct function of the coefficients, and it is the design of those that determines exactly the shape of the filter.

How do we determine the coefficients?

We start off by defining the basic filter shape that we want, so for now we'll choose the 2.4kHz low-pass 'brick wall' filter shown in **Figure 2**. In DSP processing, it is standard practice to refer everything to the sampling rate, to normalise all frequencies, so the 12kHz sampling rate takes a value of 1. Our 2.4kHz cut off takes a value of 2400/12000 = 0.2. This is called the 'normalised bandwidth'.

Now we need to look at the relationship between signals in the time domain and their equivalents in the frequency domain. This may appear to be a digression, but bear with it; the result is an eye-opener.

Without going into the whys and wherefore, let's look at the spectrum of a square pulse with a duty cycle of 0.2 ie a signal that is high for 200µs and low for 800µs. The frequency response of such a pulse is shown in **Figure 3**, a sin (x)/x or sinc(x) shape, whose tails theoretically go out to plus/minus infinite frequencies. In reality, the levels out at the distant limits are so low, below thermal noise, that they can be ignored. Standard Fourier analysis is used to transform between time and frequency domains. This is just a number crunching exercise done inside a processor using a standard algorithm (the numbers used are actually complex numbers having two parts, a so-called 'real' part, and a so-called 'imaginary' part, or in another representation the 'amplitude' and the 'phase'. We need not be distracted by these complexities here. We confine our attention to the amplitudes only in order to understand how things work).

Now we need to go though a mental barrier, and appreciate that the transform from a time domain pulse into a frequency spectrum – which is quite easy to visualise – can be swapped. This isn't quite so easy to visualise. But a block of frequencies, or 'pulse' of frequencies, just happens to be exactly what our brick-wall low-pass filter is. The transform of this rectangular block of frequencies into the time domain would be that same sinc(x) shape, but now generating a smooth pulse, a real voltage waveform that starts at minus infinite time, goes through 'now' and runs to the end of time. Only its amplitude from a few milliseconds ago to a few milliseconds hence are, fortunately, of any practical use.

We clearly can't invoke time travel with a real signal occurring before it is generated, so for practical purposes we work with a delay such that the furthest-forward transformed time-domain voltage of any significance happening is 'now', and the earliest was the delay interval ago. Any outside this range are considered too low to be of value and are ignored. This is perhaps a bit surreal, but so long as we accept that the relationship between time and frequency representations of a waveform are reversible, ie the Fourier transform operates both ways round, everything works.

Now we transform our filter shape (a rectangular 'pulse' in frequency with 0.2 duty cycle) into the time domain, and end up with the same shape as in Figure 3, with the horizontal axis is labelled time.

Merging the two

If we take our time-domain values of the Fourier transform of the filter shape, quantise it into a block centred about zero truncated to 100 either side of the centre, we have 201 values with the middle, the 101th, being the highest and tailing off each side to nearly zero at each end. To do this, we know the frequency-domain shape we want is a

Filtrate
radioecm@googlemail.com

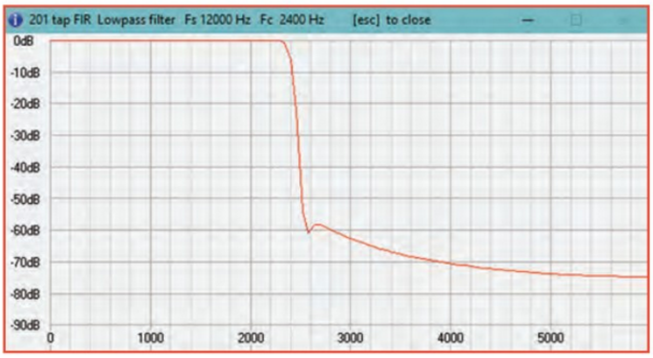

FIGURE 7: **The calculated frequency response of the digital filter with a 'Hamming window' applied in the Fourier transform, with 201 coefficients.**

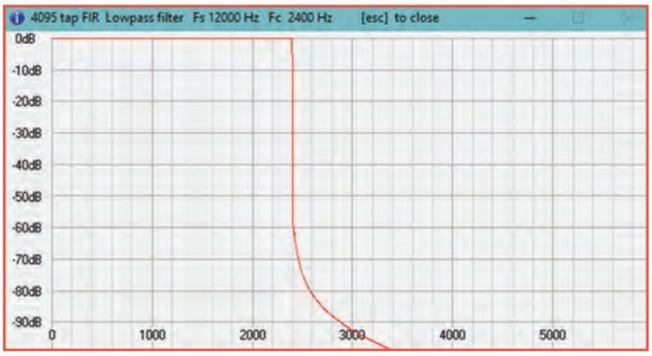

FIGURE 8: **The calculated frequency response of the digital filter with a 'uniform window' applied in the Fourier transform, with 4095 coefficients.**

brick wall, and that the Fourier transform of that is a sinc(x) response shape over time. From our normalised response of 0.2 we calculate a set of sinc(x) values over a range of 201 to give the coefficients for each of the multiplications in the hardware described above. Now the magic happens. The time representation of the wanted filter shape is multiplied by the latest and the previous 201 samples of the input waveform, the products are summed and sent to the output DAC. The result is a waveform that consists of the input spectrum after having passed through a 2.4kHz (normalised bandwidth of 0.2) low-pass filter.

Because we had to invoke the delay, our filtered output is 201 samples, or 16ms, behind the input, but that is the price to pay for FIR filtering. And that's all a digital filter is; digitise, store the last few samples in a buffer, multiply each of the buffer contents by its own coefficient and sum the result, then send to a DAC. Shift along one sample and repeat continuously.

Practical realisation

The plot in **Figure 4** shows the coefficients for a 2.4kHz low-pass filter at a sampling rate of 12kHz using a 201 long buffer. Note how the coefficients take on both positive and negative values and they are symmetrical about the middle 101th bit. **Figure 5** shows the calculated frequency response of such a filter made exactly as described above with coefficients calculated from sinc(x). Its stop-band performance isn't very good, is it? That is because of the truncation of the infinite time response to just 201 samples, which leads to a sharp step at each end where the coefficient might well be a small number but it, and any others beyond it, aren't actually exactly zero.

The solution is to multiply each calculated coefficient by a weighting function, or a window, so those at the tails are reduced to zero, or close to it, and all others are smoothed out with the centre one untouched. The resulting response has an improved stop-band performance but is compromised elsewhere such as filter roll-off rate. There is a wide choice of window types, and some windows allow user-defined settings. What we can say is that any particular window results in a trade-off between the sharpness of the roll off and the stop-band rejection. We've already seen one extreme, the rectangular or 'uniform window' case. **Figure 6** and **Figure 7** show the same filter response but with the coefficients modified by a so-called 'Hamming' window, usually considered to be a nice compromise. The stop-band rejection has improved, but the roll-off is not so sharp. A 'Blackman-Harris' window improves the stop-band rejection to more than -90dB, but has an even smoother roll off. The window type has to be chosen based on requirements; it is just part of the calculation done at the same time as the sinc(x) values are generated.

The choice of 201 samples was only a value to use for demonstration purposes; we can have any buffer length we like. A longer buffer takes

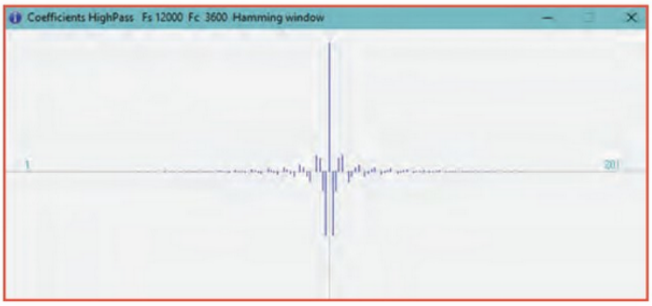

FIGURE 9: **The coefficients for a high-pass filter with a 'Hamming window' applied in the Fourier transform.**

up more memory, and requires more processing time, but can give a cleaner and sharper response. **Figure 8** shows the same filter (uniform window), but now implemented in a 4095 long buffer. The plot of its coefficients is too unwieldy to illustrate.

Why all these odd numbers, 201, 4095 etc? We want a symmetrical response so the centre functionally corresponds to zero, with coefficients numbered from -100 to +100. Negative numbers are not practically useable in counters, so the whole lot is shifted. To maintain zero as a real buffer number in the count we need an odd number of cells with an equal spread each side, which means the total has to be an odd number. Digital filters do work perfectly well with an even-numbered buffer length, but the calculations, especially when being done in DSP chips, becomes untidy with no proper centre reference to work from. Digital filter design software often mandates an odd number for the buffer length.

High-pass and band-pass filters

These filters are built in just the same way, but the calculation of the coefficients is more complex. A high-pass response is a mirror image of its equivalent low-pass. The low-pass coefficients are calculated backwards, mirrored around half the sampling rate and have alternate coefficients negated. **Figure 9** and **Figure 10** show the coefficients and response of a 2.4kHz high-pass digital filter using a Hamming window. Comparing the response with that shown in Figure 7, noting that the right-hand edge is at half the sampling rate, shows the mirror-image concept. Inverting alternate samples is, in effect, shifting the mirrored response back into the wanted region.

Calculating the coefficients for a band-pass filter is even more complex. We now have to bring in the concepts of normalised bandwidth and centre frequency, but we transform from a low-pass prototype in just the same way, modifying the coefficients based on centre frequency and bandwidth. To conclude, **Figure 11** and **Figure**

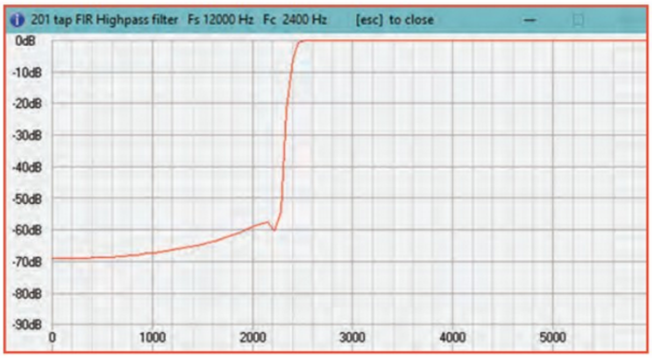

FIGURE 10: The calculated frequency response of the high-pass digital filter with a 'Hamming window' applied in the Fourier transform, with 201 coefficients.

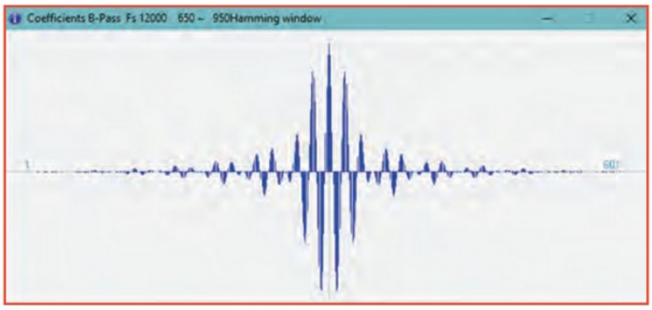

FIGURE 11: The coefficients for a band-pass filter with a 'Hamming window' applied in the Fourier transform, with a buffer length of 601.

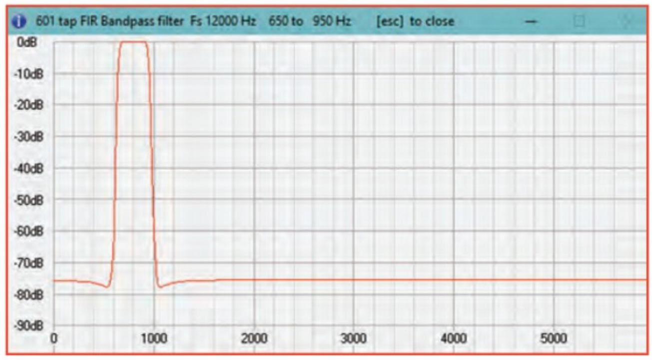

FIGURE 12: The calculated frequency response of the band-pass digital filter with a 'Hamming window' applied in the Fourier transform, with 601 coefficients.

12 show the coefficients and response of a band-pass filter centred on 800Hz with a width of 300Hz – the classic CW filter. To be effective, band-pass filters generally need longer buffer lengths than low-pass or high-pass ones, so a 601 long buffer was used for illustration. In practice, band-pass filters may more-easily be implemented in two stages, as a low-pass filter followed by a high-pass filter in cascade.

Design software and practical implementation

With so many variables to choose from, sampling rate, filter type and shape, buffer length, rejection vs roll off, delay etc, some type of design software is essential to allow optimisation. There are hundreds to choose from, and typing "*Digital Filter Design Software*" into a search engine is a good place to start.

All the plots shown here were generated from software written by me in order to get results into a wanted format, saving coefficients in a file for use in DSPic devices. They agree closely with coefficients and results from several other filter-design suites. The filter coefficients come from sinc(x) and window-function calculation, and the resulting frequency response is formed by taking a Fourier transform of those windowed coefficients.

In a real bit of hardware, with a filter that is user-adjustable, any time the filter is adjusted a new set of coefficients is generated. As we've seen, these are easily calculated on the fly, so the filter shape can be changed rapidly, and even changed automatically based on other criteria. This is one way that noise reduction and automatic notch rejection is implemented.

There are other ways to do digital filtering. For example, take a Fourier transform, shape the resulting spectrum and then do an inverse Fourier transform to get back to time domain again. That technique is very processor-intensive, but it does allow a weird and arbitrary filter shape to be formed.

Infinite impulse response (IIR) filters use a short buffer of very few past samples, but the coefficients are replaced by weighted versions of the output in a feedback arrangement. There is no long finite buffer length and the output, theoretically at least, is influenced by an infinite number of input samples – hence the name. IIR filters can be difficult to implement in practice and are used mainly when delay has to be minimised or processor hardware is limited.

Final twist: the circular buffer and the MAC command

We have seen that input samples are moved into a buffer, and all existing contents moved one place to the right. It would be very foolish to actually move, in software, the contents of each buffer cell as each new sample arrives. That would waste an awful lot of processor time. Instead, a circular buffer is maintained with two pointers. The head-pointer points to the next cell where a new ADC sample is stored then the pointer is incremented. When it reaches the maximum, or buffer size, the head-pointer count is reset to zero and just continues incrementing, in effect counting modulo the buffer length, ie 0, 1, 2 … 200, 0, 1 etc. So now the oldest sample is automatically discarded by virtue of its being overwritten by the new sample, with no physical reading and re-writing of the contents, just pointers manipulated. Every time the accumulated multiply and add operation is made, a tail-pointer starts at the current head-pointer position and moves around the buffer until it reaches the last or oldest location, just behind the current head pointer. Or it could go the other way round. That also counts modulo the buffer length so it automatically points to the right cell in turn crossing zero to do so and retrieving each historical sample in turn for the accumulation. Custom DSP chips like the DSPic family from Microchip all have a MAC command within their set of instructions. This is a very powerful instruction and the way DSP chips are designed can allow three things to happen at once, all in one processor clock cycle. First of all, the contents of two registers are multiplied, here the current buffer cell and its coefficient, with the product added into an accumulator. Simultaneously, two register pointers can be incremented modulo a value, one for the head-pointer, the other for the coefficient pointer. The modulo counting is automatic, being a function built into the DSP chip's hardware with the modulo value, the buffer length, just having to be set up in the processor at the start.

A few more processor clock cycles may be needed for setting up the routine and looping, but it does mean that a typical filter operation can be done with perhaps two or three clock cycles per buffer cell. With a DSPic chip running at, say 50MHz clock frequency, this means one entire filtering operation with this 201 long buffer, working at 12kHz can be done in around 600 clock cycles or 12μs. This gives plenty of time left to do other things with the data stream, and many DSP devices run an awful lot faster than that. Custom hardware, as used in direct-sampling DSP radios, may well be running at a clock speed of several GHz, meaning real-time filtering of a 120MHz sampled input waveform is possible. Other tricks are used there to speed things up even further, but this is not the place to go into them.

Technical

A simple 6m antenna

Many of us live in properties with small gardens, but we should not despair of putting up effective antennas, especially on the shorter wave bands. Here I describe a simple and effective antenna with low visual impact for use on the 6m band.

6m operating

For a number of years I have taken a mid-summer break in a caravan or lodge, and concentrated on 6m contacts using FT8: the propagation mode was almost exclusively Sporadic-E (Es) **[1]**. Sites near the coast of Kent and Sussex perform well, which can be attributed to a good take-off in directions from east round to south. This year, I have stayed at home (Hertfordshire), but I wanted to continue my collection of grid squares. I used the same equipment as for my caravan breaks: an Icom 7300 with a dipole antenna suspended from an extendable carbon-fibre pole (a Sotapole **[2]**), using FT8 in WSJT-X software running on an old Acer Netbook. When in the caravan, I attached the pole to the railings around the outside seating area, and tied off the lower end of the dipole to the railings. But when operating from home, I needed a different sort of antenna support. My garden is small, basically a lawn with shrubs round it, and is overlooked by other houses. I decided to use the rotary washing drier, which is roughly central in the lawn and therefore at maximum spacing from surrounding houses. I attached the pole to the vertical support of the washing drier using Velcro straps, and one end of the dipole to the top of the pole. By extending the lower end of the dipole with nylon cord, I tied it to a convenient shrub (see **Figure 1**). I brought the coaxial cable away from the dipole at roughly 90°. In Figure 1, you can see the black pole extending up from the centre of the drier, the dipole running down from the top, and the central 'T' insulator is seen just to the left of the chimney of the house in the background. The lower half of the dipole can also just be seen.

FIGURE 1: **The 6m dipole erected in my small garden.**

The details of the dipole are given in **Table 1**. The overall length of the dipole is 271cm. The construction is as straightforward as possible: one end of the dipole is attached directly to the top of the mast by the 'top antenna insulator', and a small loop in the wire is made at the other end with the nylon cord tied to it. I use a PL259 to PL259 feeder made of thin coaxial cable connected between the dipole centre and the transceiver. Items for improvement could be adjusting the length of the dipole for minimum VSWR, using thicker coaxial cable for lower loss, and using a proper insulator between the dipole and the antenna cord.

On the air

The results surprised me. Although the Es propagation was largely in the directions from east through to south, I was getting short periods of propagation to Norway, Sweden and occasionally Finland. These periods were short, about 60s, but just long enough to complete an FT8 contact if I concentrated hard! They were also limited in geographical area, and each involved only one station: in contrast, when the propagation was to, say, Italy, the decoding screen was full of I-prefix stations.

David Price, G4BIX
ddprice001@gmail.com

After thinking about this, I realised that the shrub I had tied the guy off to, was located roughly north of the clothes drier. I therefore had a sloping dipole antenna, with the low end to the north of the high end. Could this be producing directivity, allowing my system to explore weak propagation paths roughly to the north (Scandinavia)? This rang bells with me: I remembered seeing 'slopers', and arrays of slopers, mentioned in various antenna books over the years, but only for HF (typically 40m). Could this directional effect also work on 6m, on the boundary between HF and VHF? I untied the guy from the northern bush, and instead tied it off to a garden chair placed on the lawn roughly southeast of the clothes drier. This orientation gave me good propagation to the Balkans, and I was now able to work pretty much every station I could hear, which was a definite improvement.

Directionality with a sloper

It seems that a sloping half-wave 6m dipole, suspended from a carbon-fibre pole, can give directivity in the direction of downward slope. Changing the directivity is easily done; just move the lower tie-off position so that the antenna slopes down towards the desired direction. I realise of course that this is not a scientific study, but the results are sufficiently interesting to suggest that this is an inexpensive and simple method of obtaining directivity. Further practical improvements could be made: there are probably better ways of changing the directivity than by moving a garden chair around!

Propagation

What propagation mode was involved in the Scandinavian contacts? I'm pretty sure it wasn't meteor scatter. Listening to the signals as they were being received sounded nothing like meteor 'pings'. Also, they didn't sound like aircraft reflections; in my experience, these produce a very characteristic pulsing amplitude ('flutter') signal, with the pulsing getter slower then faster. But the distance was about right for an E-layer reflection, and they occurred during an Es opening, but away from the direction of the main reflective patch. I do wonder if the main patch extended in a very weak form between the UK and Scandinavia, and the high sensitivity of the FT8 mode allowed transient weak reflections to be heard. Alternatively, perhaps each 60-second propagation event could have been caused by a meteor delivering metal ions locally and forming a transient patch.

References

[1] D. D. Price, 'Using FT8 for summer sporadic E', *RadCom* Vol. 96, page 21, 2020

[2] https://www.sotabeams.co.uk/telescopic-masts-sotapoles-7m-and-10m/

TABLE 1: The details of the dipole.

Materials	Manufacturer	Part number
Dipole centre	Palstar	DP-239
Antenna wire--lightweight	Sotabeams	AWLW-103
Stealth antenna cord (nylon)	Sotabeams	STAC-050
Telescopic antenna mast	Sotabeams	TAM-101
Velcro antenna straps (2)	Sotabeams	VAS-100
Versatile top antenna insulator	Sotabeams	VTAI-100

Contest Calendar January 2024

Ian Pawson, G0FCT

RSGB HF Events

Date	Event	Times (UTC)	Mode(s)	Band(s)	Exchange
Sat 6 Jan	AFS 80m – 40m CW	1300-1700	CW	3.5, 7	RST + SN
Sun 14 Jan	AFS 80m – 40m Datamodes	1300-1700	PSK63, RTTY	3.5, 7	RST + SN
Sat 20 Jan	AFS 80m – 40m SSB	1300-1700	SSB	3.5, 7	RS + SN

RSGB VHF Events

Date	Event	Times (UTC)	Mode(s)	Band(s)	Exchange
Tue 2 Jan	144MHz FMAC	1900-1955	FM	144	RS + SN + Locator
Tue 2 Jan	144MHz UKAC	2000-2230	All	144	RS(T) + SN + Locator
Wed 3 Jan	144MHz FT8 AC (4 hour)	1700-2100	FT8	144	Report + 4-character Locator
Wed 3 Jan	144MHz FT8 AC (2 hour)	1900-2100	FT8	144	Report + 4-character Locator
Tue 9 Jan	432MHz FMAC	1900-1955	FM	432	RS + SN + Locator
Tue 9 Jan	432MHz UKAC	2000-2230	All	432	RS(T) + SN + Locator
Wed 10 Jan	432MHz FT8 AC (4 hour)	1700-2100	FT8	432	Report + 4-character Locator
Wed 10 Jan	432MHz FT8 AC (2 hour)	1900-2100	FT8	432	Report + 4-character Locator
Thu 11 Jan	50MHz UKAC	2000-2230	All	50	RS(T) + SN + Locator
Tue 16 Jan	1.3GHz UKAC	2000-2230	All	1.3G	RS(T) + SN + Locator
Thu 18 Jan	70MHz UKAC	2000-2230	All	70	RS(T) + SN + Locator
Tue 23 Jan	SHF UKAC	1930-2230	All	2.3G	RS(T) + SN + Locator

Best of the Rest Events

Date	Event	Times (UTC)	Mode(s)	Band(s)	Exchange (Info)
1 Dec 23 - 31 Jan 24	UKSMG Winter Marathon	All	All	50	RS(T) + Locator
26 Dec 23 - 6 Jan 24	WAB Christmas Party (not a contest)	All	All	All	WAB Book number
Wed 3 Jan	UKEICC 80m	2000-2100	SSB	3.5	6-character Locator
Mon 1 Jan	IRTS 80m Daytime Counties	1700-1800	SSB/CW	3.5	RS(T) + SN (EI & GI also send county)
Sat 6 - Sun 7 Jan	ARRL RTTY Roundup	1800-2359	Data	3.5-28	RST + SN (W send State, VE send Province)
Sat 6 - Sun 7 Jan	EUCW 160m CW Party	2000-2300, 0400-0700	CW	1.8	Club Member: RST + name + club + membership no. Other: RST + name + NM
Sat 13 Jan	WAB 1.8MHz Phone	1900-2300	SSB	1.8	RS + SN + WAB square
Fri 26 - Sun 28 Jan	CQ 160m DX	2200-2200	CW	1.8	RST + CW Zone (W send State, VE send Province)
Sat 27 - Sun 28 Jan	BARTG RTTY Sprint	1200-1200	RTTY	3.5-28	SN
Wed 31 Jan	UKEICC 80m	2000-2100	CW	3.5	6-character Locator

For all the latest RSGB contest information and results, visit www.rsgbcc.org

An improved, yet simple, direction-finding loop

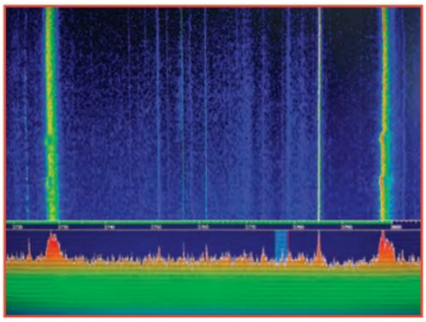

FIGURE 1: QRM on 80m.

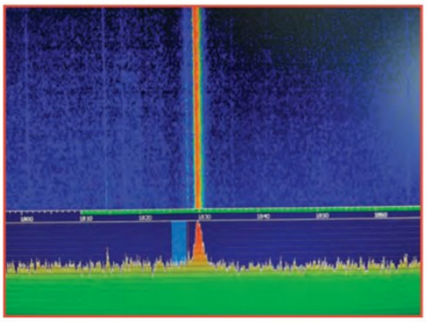

FIGURE 2: QRM on 160m

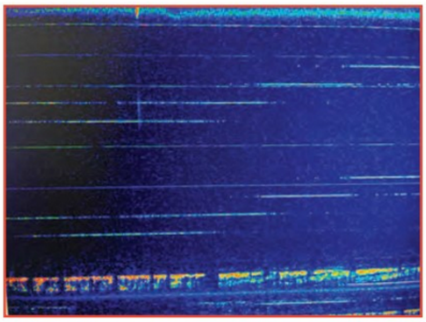

FIGURE 3: QRM on 40m.

It was during a chat on the 40m band in May that I realised my last HF QSO had been in March, and somehow I'd missed April completely! I had been active, but just on VHF and UHF, and mainly operating portable.

Later that afternoon, whilst mooching around the lower bands, I discovered to my horror a whole series of new repeating interferences across 40m, 80m and 160m. My heart sank…

Some history

RadCom readers may recall a previous article I wrote **[1]** on building a direction-finding loop and 'sniffer wand', which I used to track down and eliminate my local source of QRM. A neighbour's non-filtered car-battery charger was the culprit at that time (despite sporting CE markings), and it was promptly replaced at my own cost with a compliant UK-made model (and a neighbourly bottle of wine for their help). I also found a single dodgy LED bulb in my own home, which I replaced.

This new 80m and 160m interference looked suspiciously similar to last time on the waterfall display (**Figures** 1 and 2). The wide signals and associated sound, and the roughly 70KHz spacing, looked like the car battery charger-related QRM I'd experienced before. The 40m signal was something different though. There were four clustered narrow-band signals, repeating in groups across the band, which came and went in a variable burbling pattern (**Figure** 3).

Initial tests

Buoyed by my previous successes, I first did a mains power-off test on my own house to confirm I wasn't the source of my own QRM. A Yaesu FTDX10 with a small lead-acid battery was used to check the bands and, sure enough, the QRM persisted. The next step was a quick check with any of my neighbours who happened to be in at the time. No, no one could think of any new electrical items installed over the past few months. One neighbour had installed solar panels and a battery storage system, but I'd specifically checked the bands after that was commissioned earlier in the year, and all seemed okay. Not disheartened, it was time to dig the direction-finding loop (**Figure 4**) out of the garage and start a local search on foot.

Now the only 'portable' HF radio I have is the FTDX10, but it is not readily carried by a pedestrian. Nevertheless, carrying the rig in a backpack in my arms, powered from the lead-acid battery in a bum-bag around my waist and loop in hand, I set off to check my local small estate. I set the radio to the frequency of a strong interfering signal at the top-end of the 80m band. Yes, I did feel a bit awkward initially, but I soon got into the swing of things and forgot that I probably looked a bit weird!

Well, I heard nothing. Up and down the road I went and I couldn't detect a thing. I even went back into the shack to fire up my main radio and yes, the QRM was still there. All the home antennas were picking up the interference, but it just wasn't strong enough for the loop to detect. Perhaps a weaker source, or even further away than last time? Despite encouragement from the neighbours, I felt a little deflated. What if I couldn't trace the source of this QRM? I'd be stuck with it.

Direction-finding loop Mk2

I then remembered some feedback Don, WD8DSB had given to my QRM YouTube video **[2]** from a few years previously (which my original *RadCom* article had been based upon). He suggested a tweak to the original design which would increase the gain and sensitivity of the loop on the lower bands. Don believed that the original loop would perform well above 15MHz, but less so on the lower bands. With renewed optimism, into the garage I went looking for a section of RG213 and my toolbox.

It didn't take long to make the new version of the loop as it was very simple to make. It is just a loop of coaxial cable. I'd also remembered in the feedback that someone suggested a construction method without the various PL-259s and T-piece I'd used previously, and it was this simpler solder-only method that I used this time around. A diagram of the loop is shown in **Figure 5**. You build it as follows:

1. Using a 1m+ length of RG213, expose about 10mm of braid and core at one end, and solder the core to the braid to form a short, soldered stub.
2. Measure 440mm from the end of the outer jacket of this soldered stub and mark a 20mm section from this point ie 440-460mm. This will be the top-centre of the loop. Cut away the outer jacket for this 20mm section, and then also cut away the braid using side-cutters, but not the inner dielectric. Now from the other side of this 20mm gap, measure another 440mm and make a mark. This is the other half of the loop.
3. Roughly form the loop by bringing the soldered-stub end round to this mark, and cut away approximately an 8-10mm square of outer jacket to expose the braid. Do this on a flat surface to help ensure the loop lies reasonably flat and is not twisted.
4. Add some solder to this exposed square of braid and then solder the stub-end to it.
5. Tape up both the exposed top of the loop, and this soldered joint at the base. Also add a suitable plug to the tail of coaxial cable for connection to a patch lead and/or the radio.
6. Finally, support the structure of the loop and

FIGURE 4: The original loop antenna.

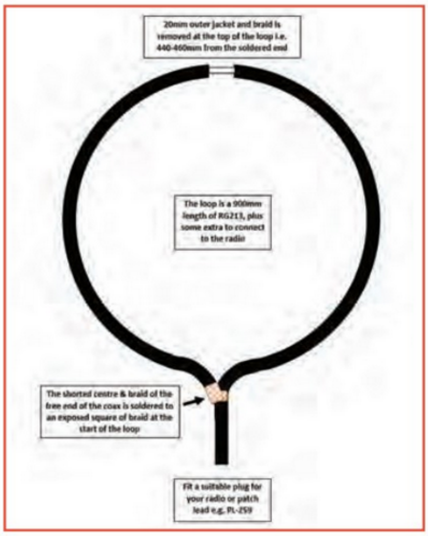

FIGURE 5: A diagram of the Mk2 loop.

FIGURE 6: The finished loop Mk2.

provide a handle using some wood, garden cane, or plastic conduit. I used cheap 25mm plastic electrical conduit which is readily available from DIY and electrical retailers. A cutaway channel at the top end of the conduit allows the 20mm section of the inner section of the coaxial cable to neatly sit in this recess and support the top of the loop. Whatever support you use, form a neat loop and then tape the top, the base and the coax tail to the handle. See **Figure 6**.

The only difference in this Mk2 version, compared to the original design, is that the end of the coax that returns to the base of the loop is shorted across its core and braid and then connected to the incoming braid at the start of the loop as described above. On the Mk1 design, the core at the end of the coax wasn't connected; only the braid on the end of the stub was connected to the braid at the start of the loop. This difference in design can been seen in **Figure 7**.

When I ventured back outside armed with the Mk2 loop, it was immediately obvious it was detecting the interference. I then started to search around my estate to try and establish its direction. It's worth pointing out at this point that a loop has a strong null when face on, ie when you're viewing the open circle. Its gain comes from the two aligned edges. So when you hear a peak in noise while slowly rotating the loop, it's worth remembering that, if the loop is edge-on when held in front of you, the source of the noise could be ahead of you or behind you. If you walk forwards and backwards a reasonable distance, then you should hear the noise increase or reduce and therefore establish which direction it's actually in. You therefore need to walk around the neighbourhood and start to establish a direction from multiple listening points using this triangulation technique. Where these imaginary direction lines cross gives a good indication of the location of the source of interference. By having the more-sensitive Mk2 loop, I was able to establish an initial direction on which to focus my search. These initial results can be seen as varying-strength yellow lines in **Figure 8** (the thicker lines show increased signal strength), with my 80m dipole location shown in red. As in my previous search in 2021, this new search took me beyond the boundaries of my own estate but, as the battery was running low, I retired at this point, at least happy I had a chance of tracking down the source.

Extending the search

After work the next evening, I set up the equipment and headed off with the Mk2 loop to the neighbouring estate which backed onto ours. By walking round the outer edge of the estate, I could confirm the source *was* lurking somewhere inside, but I found, as I closed in and walked up and down the various roads, that the signal was too strong, and it was difficult to pin down reasonably accurate directions in what was now around 40 houses. So back home I went and swapped to the Mk1 loop which I now knew was less sensitive on this lower band. I could also have inserted an attenuator between the loop and the radio (if I had one to spare). Sure enough, I could navigate inside the estate taking mental direction readings as I went. I'm sure I saw a few curtains twitching too as I entered various quiet cul-de-sacs. Anyway, the result of a lot of back-and-forth walking and aching arms, was homing in on two houses where I thought the noise was coming from. Again, suspecting my small battery was on its last legs, I retired. I didn't want to knock at a door, only to find my battery had died at the vital moment of my introductory pitch and radio demonstration!

The door knock

On the third day, I went back round and knocked

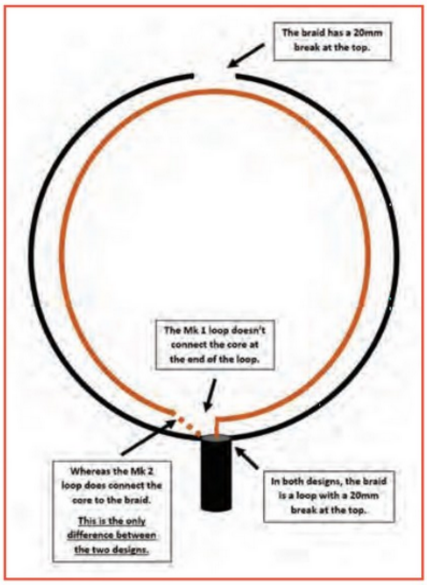

FIGURE 7: A diagram showing the differences between the Mk1 and Mk2 versions. The core is shown in a light brown colour, and the braid in black.

on the door of the main suspected house. I had a good chat with the nice lady who answered, and she was only too happy to help, but couldn't think of anything new and electrical that she might have in the house. When asked specifically though, she did confirm that she didn't have a car battery charger. However, while chatting on the doorstep and demonstrating the interference, I noticed that the signal had reduced a little and it was marginally stronger back on the road behind me.

Ian Miles, G0CNN
Ian.Miles@xentrall.org.uk

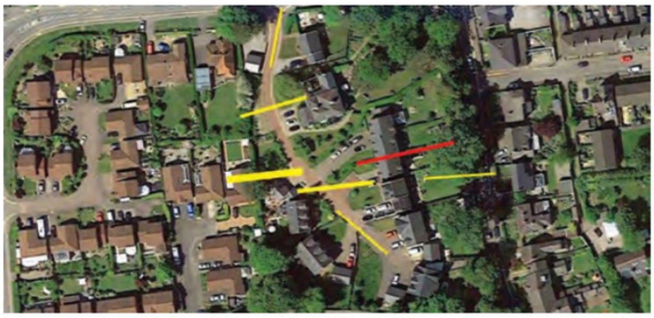

FIGURE 8: An initial search. The yellow lines indicate the directions to the source of interference, with the thicker lines indicating increased signal strength. My 80m dipole is shown in red.

Being restricted to the public road and paths meant that this was the first time I could get up close to the suspected source, and at this point I realised my mistake. I had forgotten the double-direction property of the loop and the source must have been behind me, not in front.

Having thanked her for her help, I popped over the road and knocked on another door. No one appeared to be in, but what I did notice was a mains lead coming from the bottom of the garage door and under the bonnet of a car parked on the drive. Turning the radio back on, it was obvious I was very close to the source as the noise was terrible. I felt certain that a trickle battery charger on the car was the source. Anyway, with nobody home there was little I could do. I gave the lady opposite an update, and she mentioned that she did think that the house had had new occupants for the last couple of months. This would tie in with me not seeing this QRM before. Arms aching, but up-beat, back home I went.

Success

I lugged the radio (**Figure 9**) back round to that house every evening that week, but still there was no one at home. By the second week I was beginning to get concerned. Perhaps they were away on holiday? Anyway, a bit more door knocking, and I found a neighbour who had a contact number for someone related to the house in question. I made the call, and the lady who answered was only too happy to help after I explained the situation. She then offered to make a special trip to come across and unplug the charger, but I said after two or so months of interference, a few more days wouldn't make much difference. I also offered to replace the trickle charger free of charge with a new one, if in fact that was the source of interference. I explained that I could come round and confirm this when she next planned to visit. That weekend, the QRM on 160m and 80m unexpectedly stopped. I sent her a message and she confirmed that it had been unplugged. I reaffirmed my offer to exchange the unit as it would be good to see the model and possibly send it off to the RSGB for analysis (as I'd done previously in 2021).

At this point in the story I'd be expecting to show you a photo of the offending item, but a few weeks on I haven't been contacted again, so I haven't seen it myself yet. Anyway, the offer of replacement still stands and I'm just glad the interference has stopped. I may get in contact with her again in a week or two.

A trilogy?

With this source of QRM gone, I've still got the QRM on 40m. Although now a different pattern, it has also started in the last few months, so I did hope that this could be extinguished too. It's not as strong as the QRM from the trickle-charger, but it's very annoying as the narrower spacing and clustering of four signals makes a mess of the narrower 200KHz allocation on 40m. So this is my next job. Wish me luck!

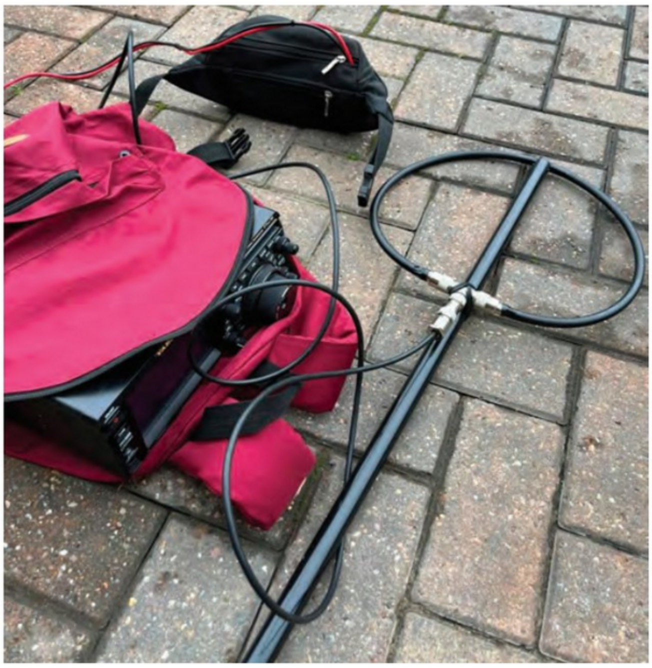

FIGURE 9: The 'portable' kit.

Takeaway points

Please don't take your QRM for granted. From the feedback I received from promoting the first loop, fellow radio amateurs have embraced the simple design and made one for themselves, and have had some really impressive successes. These loops are very easy to make and only require a short length of coaxial cable, a PL-259 connector, and some soldering. It could be a saviour in your hour of need. As in my case, with a small battery, even a 'non-portable' radio can be made 'portable' for this short-term searching exercise (just don't drop it!).

The Mk2 loop *is* a success, and has proven to be more sensitive on the lower bands. As I found, however, the Mk1 loop is still useful when you're closing in and your radio gets overloaded. I would have struggled to focus in on a few houses without it. Also, don't forget the 'sniffer wand' from the original article in 2021. It's very useful at close range.

In summary then:

- Use the Mk2 loop for distant work when the Mk1 loop hears nothing
- Use the Mk1 loop when closing in, and the Mk2 becomes too sensitive
- Use the sniffer wand when inside a property to pinpoint a noisy source
- Perhaps consider making a single loop with a switchable coaxial-core connection at the base. This will provide high- and low-sensitivity options
- The neighbours and strangers I've spoken to were genuinely interested, if not intrigued, by what I was up to. In comparison with the dark-sky campaigns in astronomy, you can help to explain this unseen 'pollution' issue in the radio world
- Remember your personal safety. You know your own neighbourhood
- Always bear in mind that you're looking for someone else's help and cooperation. You're not apportioning blame or challenging someone's past purchase decisions

Good luck if you do go out on your own QRM hunt!

References

[1] Original Loop, *RadCom* September 2021, Vol. 97, No. 9
[2] YouTube video of the first QRM hunt, https://youtu.be/56IUtdBK-7U
[3] Latest YouTube video on which this article is based: https://youtu.be/3EWRQYeRzdg