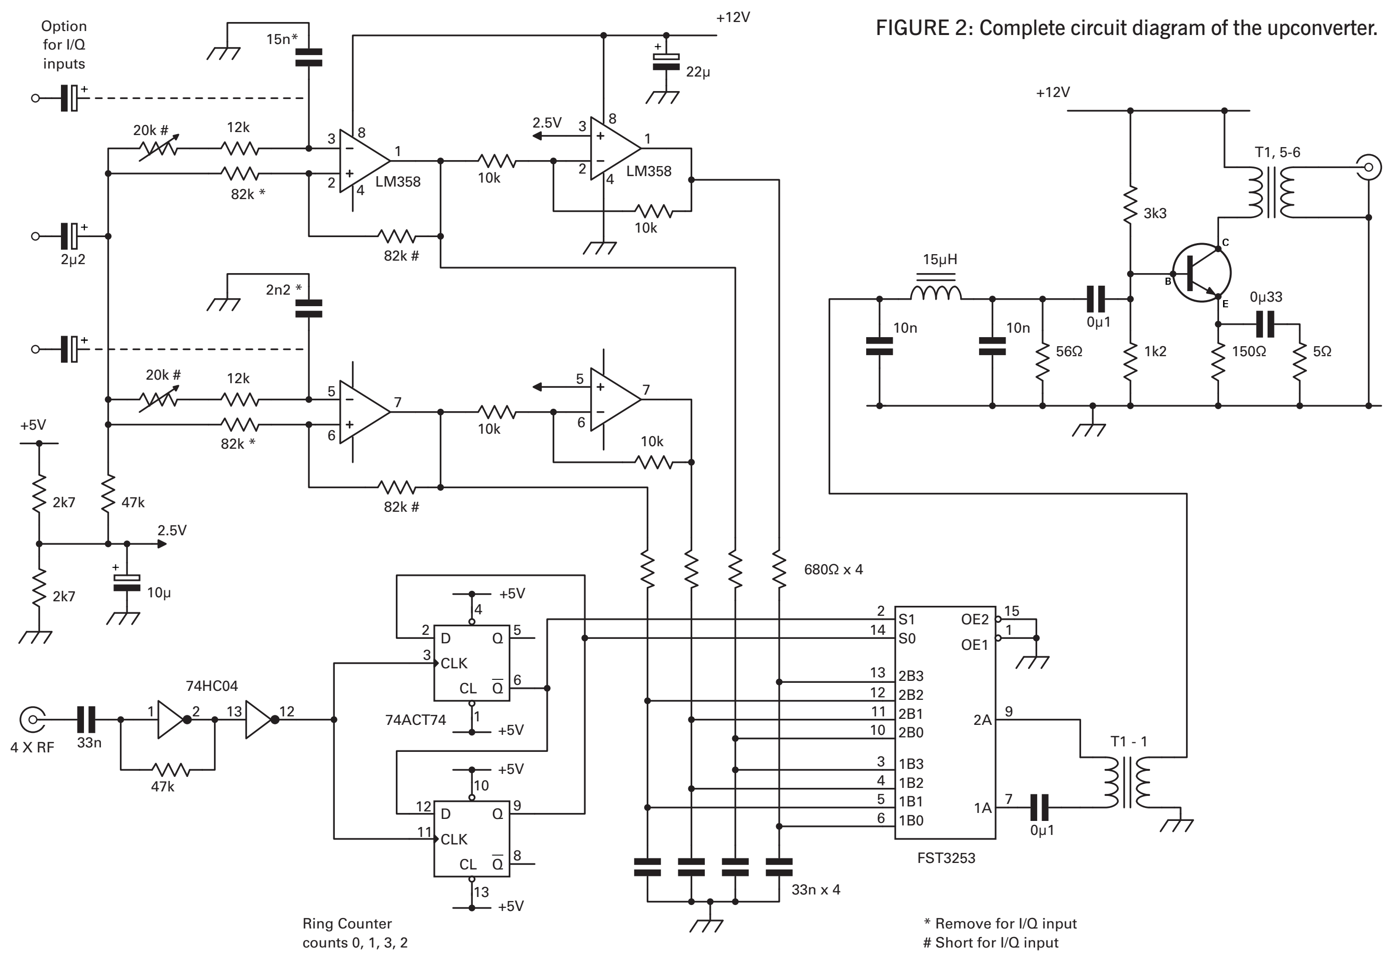

FIGURE 2: Complete circuit diagram of the upconverter.

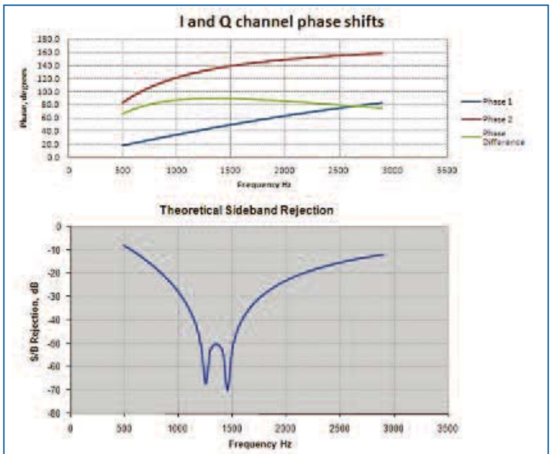

FIGURE 3: I and Q channel phase shift with its resulting sideband rejection.

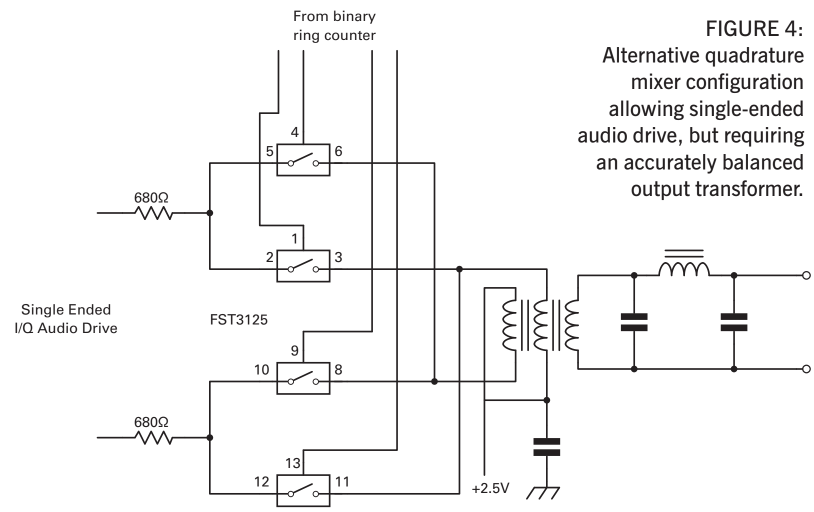

FIGURE 4: Alternative quadrature mixer configuration allowing single-ended audio drive, but requiring an accurately balanced output transformer.

See [1] for an all-pass network design tool.

Only single I and Q channel signals are supplied and we need differential drive, so a second pair of opamps are used as unity gain buffers to generate minus-I and minus-Q.

If I and Q signal can be generated by software instead of having to be synthesised by analogue networks, the connections shown by the dotted lines can be used to configure the first two pair of opamps as buffers rather than all-pass networks. Several components need to be shorted out or removed to do this. They are marked on the diagram.

ALTERNATIVE SWITCH CONFIGURATION. There is an alternative configuration that may be preferable in some cases. Instead of using a dual pole requiring four audio drive signals, four single switches with a single ended I and Q drive is possible. A balanced centre tapped transformer on the output as shown in **Figure 4** then becomes necessary. The accuracy of this transformer's centre tapping and balance dictates the level of carrier rejection possible. The arrangement is simpler than the first design, but does rely on the output transformer so carrier rejection and possibly sideband rejection may be worse.

SETTING UP. Alignment is no more than adjusting the I/Q networks for optimum sideband rejection and checking / trimming carrier rejection. As an absolute minimum you will need a receiver that can be tuned either side of the output band of interest to look at sideband rejection and a suitable clean sinusoidal drive signal. Ideally the receiver should have a spectral display several kHz wide – any soundcard based SDR will do admirably.

For the source, it is helpful if the drive can be mixed with white noise to show sideband rejection over the whole output in **Table 1**. All the devices were run with -5dBm input to the attenuator, giving their recommended -15dBm drive level. Audio amplitude was set to the maximum specified.

TABLE 1: Measured results from the three upconverters. * = see text.

Device	Frequency (MHz)	Carrier Rejection (dB)	Sideband Suppression (dB)
U2793	30	35	24
	50	43	49
	145	42	47
	250	55	42
	320	42	45
	435	36	41
AD8345	145	47	38
	432	41	41
	600	40	40
	900	28	40
	1000	25	39
AD8346	800	38	47
	1000	38	44
	1300	32*	45
	2320	27*	43
	2400	28*	43

The low carrier rejection seen on the AD8346 above 1GHz is almost certainly due in part to the imperfect isolation between input and output connectors that were mounted directly on the PCB and not fitted into a screened box. Another rats' nest wired construction using this chip as part a breadboard 1296MHz beacon that *was* properly screened showed around 40dB of carrier rejection at this frequency. However, the lower frequency device also showed a reducing carrier leakage at the upper limit of its specification so this could in part be a characteristic of the family.

The U2793 did not perform terribly well at the lower end of its specified frequency band in the 30MHz region, which was a little surprising. However, this was made up for by an excellent showing above 50MHz, which extended well up above its stated upper limit – turning in a useful carrier and sideband rejection at 435MHz. (The companion downconverter, the U2794 using a similar I/Q splitter did not perform at all well at the lower end of its frequency range either – but more about that next time).

IF INPUTS. Up to now we've assumed the inputs are audio or low frequency baseband signals for direct upconversion. Both manufacturers' products are specified for operation with IF inputs up into the low-VHF range: 75MHz for the AD devices and 50MHz for the U2793. This makes them particularly attractive for use as the Tx mixer in transverters. Transmit upconverters have particularly stringent filtering requirements to prevent the mixer image appearing as an out of band spurious transmitted signal. Typically, for high power operation we ought to be looking at the image (and any other non harmonic spurii) being at least 50dB and preferably nearer 70dB down. Such a filtering specification is quite arduous to maintain with a low IF like 28MHz and an output at 432 or 1296MHz – even more so at 2320MHz.

By splitting the 28MHz drive signal into a pair of I/Q components using an external network such as the high/low pass PI network shown in **Figure 3** or a MiniCircuits 90° hybrid splitter **[4]**, at least 30dB of image rejection is immediately found. The subsequent filtering task now becomes a lot easier as only an additional 20–40dB of attenuation has to be provided by the filter, rather than the 50–70dB originally needed. The simulated response of the dual PI network circuit is shown in **Figure 4**.

GENERATING DATAMODES AT BASEBAND.

Conventional audio speech drive needs a specially built I/Q network or a DSP engine to form the precise 90° I/Q split over the entire audio band. The requirement for such an additional complex module makes I/Q converters slightly less attractive for SSB rig replacement. The 30–40dB of opposite sideband rejection may also prove troublesome in critical applications.

But the real advantage comes about when datamodes are wanted. DSP software can generate I/Q signals almost as easily as single channel audio and the ability to generate signals with a centre frequency equal to the carrier, allowing the spectrum of the audio drive to wrap round on top of itself, means the leakage from the suppressed sideband falls in band, on top of the main signal. Being at least 30dB lower in amplitude, this will go unnoticed in the subsequent demodulation and, as it is on the same actual transmitted frequency, cannot cause QRM in adjacent channels.

A necessary requirement for any practical implementation of upconversion from baseband is that the modulation must not rely on being able to transfer a DC component. This is because the frequency response of soundcards – and even the input to the modulator – are AC coupled. So true baseband generation, without adding any deliberate frequency offset, of any modulation that maintains a carrier like AM or FM is not feasible without some ingenuity. Peter, G3PLX, came up with an interesting solution, though, in his *SDR transmitter* software. He adds a low frequency sub-audible frequency modulated tone at a level where the modulation index (deviation divided by modulating frequency) is at exactly the value where the carrier component falls to zero. This tone is added to the speech for both AM and FM so that now, during pauses in speech, there are sidebands that keep the baseband drive alive but with no DC component and at a level below audibility. Try Peter's software to see this in operation.

Narrowband datamodes can be formed at baseband using simple direct generation circuitry and applied directly to the modulator's inputs. Two modulation types used by amateurs particularly lend themselves to being generated in a small microcontroller like a PIC and digitised with low cost D/A converters. PSK31 has a basic envelope and modulation that is only 31Hz wide. The half-sinewave envelope can be stored in a lookup table and accessed at a sampling rate at a quite modest 8 to 10kHz. At this frequency, which is many times that of the data modulation, sidebands due to the sampling rate step and / or low D/A resolution quantisation can be made arbitrarily small. For example, PSK31 generated at 8kHz sampling rate will have its first set of sidebands 31Hz away from the modulation at a level, to a rough approximation, of 20LOG(31/8000)= –48dB. In a PIC running at a moderate clock speed, 8kHz sampling generated in a timer-interrupt routine allows plenty of clock cycles for character generation, two table look-ups (for 0° and 90° amplitudes) and even for sending the data on an SPI type interface to a serial D/A converter. The low cost TLV5625 dual 8 bit D/A converter is suited to this task as it allows simultaneous update of both analogue outputs. Any delay here would affect I/Q balance, to the detriment of sideband suppression.

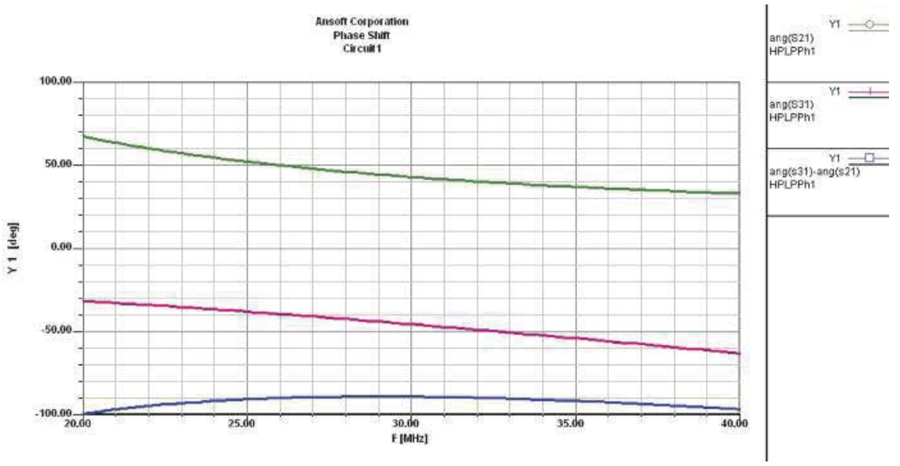

FIGURE 4: Simulated results for the circuit of Figure 3. The red and green traces are the individual phase shifts of each path. The blue trace shows the difference between the two.

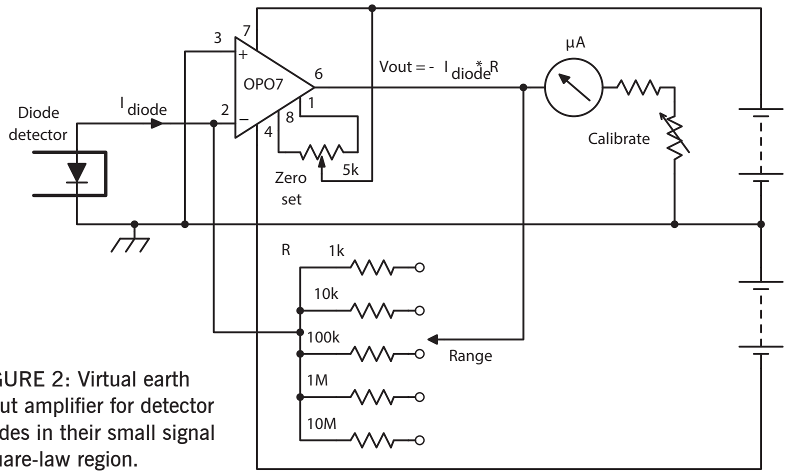

FIGURE 2: Virtual earth input amplifier for detector diodes in their small signal square-law region.

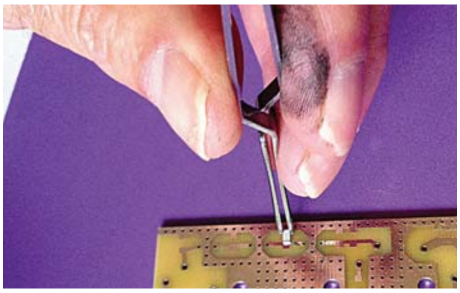

PHOTO 3: Press-to-open tweezers used by G3SBI for holding SMT components during circuit board assembly.

its melting point, after which the removal process becomes easier.

Figure 3 shows the connections for a standard USB B connector, with the pins that would normally drop into holes in a PCB uppermost. Two chokes filter the 0V and +5V supply leads. I used 2.2μH ones as I have a large stock of them, but any value in the microhenry rage will do provided they are rated to carry 100mA or more. It is arguable whether there is any advantage to be had in adding a choke in the ground connection as this will be in parallel with, and shorted by, the cable shield. The screen should be firmly connected to the metal case and to the PC casing at the other end. It always feels 'wrong' to me to have two such parallel paths, one inside a screened cable and the other the shield itself. There is the possibility of ground loops and the resultant pickup of magnetically coupled interference or chassis currents. The additional choke can't do any harm. The hi-fi audio world often uses low value resistors in a similar position in signal leads, to break ground loops that may be set up due to safety earthing on separate equipment.

Don't even think of putting separate chokes in the two USB signal leads! The USB signal is a balanced transmission line carrying signals at hundreds of MHz and chokes would kill it. Instead, construct a common mode choke from a twisted pair of two pieces of enamelled wire wound around a ferrite bead or toroid. The balanced transmission line action is not affected by the ferrite material as its effects exactly cancel on each conductor, but any common mode interference imposed against the screen or power supply leads running alongside will be attenuated.

Depending on where you buy the module, and to a lesser extent on how much you pay, some of the dongles are supplied with an adapter lead that matches the MCX coax connector on the module, with an SMA or BNC on the other end. If you are lucky this may even be a bulkhead mounted connector – many of the SMA ones are – so the adapter lead can be used unmodified by mounting that in a hole in the metal box. While I could have done this, my supplied adapter lead was too long to fit inside the small diecast box properly. Both the SMA bulkhead and MCX connectors were crimped so neither could be disassembled to shorten the lead. So instead I kept the crimped SMA bulkhead and removed the MCX socket completely from the PCB, soldering the shorter adapter cable as shown in Photo 2.

LOW NOISE AMPLIFIER CHIP. Colin Horrabin, G3SBI wrote in to say "A recent development from Mini-Circuits is the PHA-1+ MMIC and its dual matched version, the PHA-11+. What took my eye whilst looking at the data online was a quoted noise figure (NF) of 1.8dB, gain of 18dB and an output IP3 of +42dBm. The chip is really aimed at microwave applications well above 100MHz, so I expected its NF to rise sharply and its IP3 to drop at lower frequencies. Much to my surprise this was not the case.

I obtained a sample of the PHA-1 and put it down on a larger than necessary PCB to keep the chip temperature down and, along with George Fare, G3OGQ, we paid a visit to Dave Roberts, G8KBB to measure its characteristics on an N2PK vector network analyser. At 45MHz its input impedance was 80Ω in parallel with 25pF. NF was 2.2dB and gain was 18dB. When we used an L match from 50 to 80Ω so the chip was correctly matched, the NF was raised to 2.7dB. The PHA-11 (which has gain at 8GHz) was completely stable in my application at 45MHz "

SURFACE MOUNT TWEEZERS. Colin then went on to design a state of the art 45MHz front end using the PHA-11. He writes: "It had to be a commercial PCB so that thermal vias could be used to keep the chip temperature down. The circuitry was inside a tin box and with the lid removed you see all the surface mount components. I had never previously done any serious surface mount assembly so to construct this I bought a set of tweezers from Hobbycraft. However when I came to populate the board I picked up an SM part with some tweezers and before I got it anywhere near the board there was a 'ping' and the part disappeared. I realised I was going to be in serious trouble until I remembered that one of the pairs of tweezers opened when you pressed it. So you could clip it onto one end of a SM part and then hold it on the PCB to solder the other end. Using this technique was easier to fit than through hole parts. See **Photo 3**.

"Maplin Electronics offer a two-tweezer set that works the same way, one straight and one angled, for about £5. The wrong-way-round type tweezers make it easy. Not only are SM caps cheaper, they are technically superior to the wire ended part. I would really make the point that the use of the type of tweezers I used makes it easy even for someone with the shakes".

And to conclude, Colin also recommends "... the magic L and C measuring box made in the USA by Almost All Digital Electronics. [1] I used it to wind T37-10 toroids in the range 180 to 500nH and it was spot on. The same was true of capacitors in the 2 to 6pF region when I came to measure the values set by some trimmer capacitors. I think you can buy it ready made for about £100 or build it as a kit."

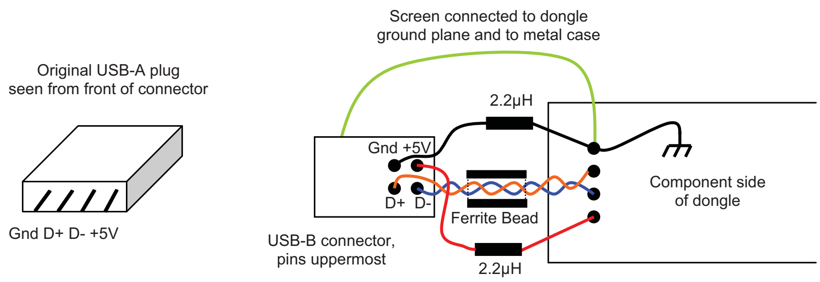

FIGURE 3: Connections for a separate USB-B socket with DC filtering and common mode choke on the USB signal leads.

REFERENCE

[1] Almost All Digital Electronics – www.aade.com

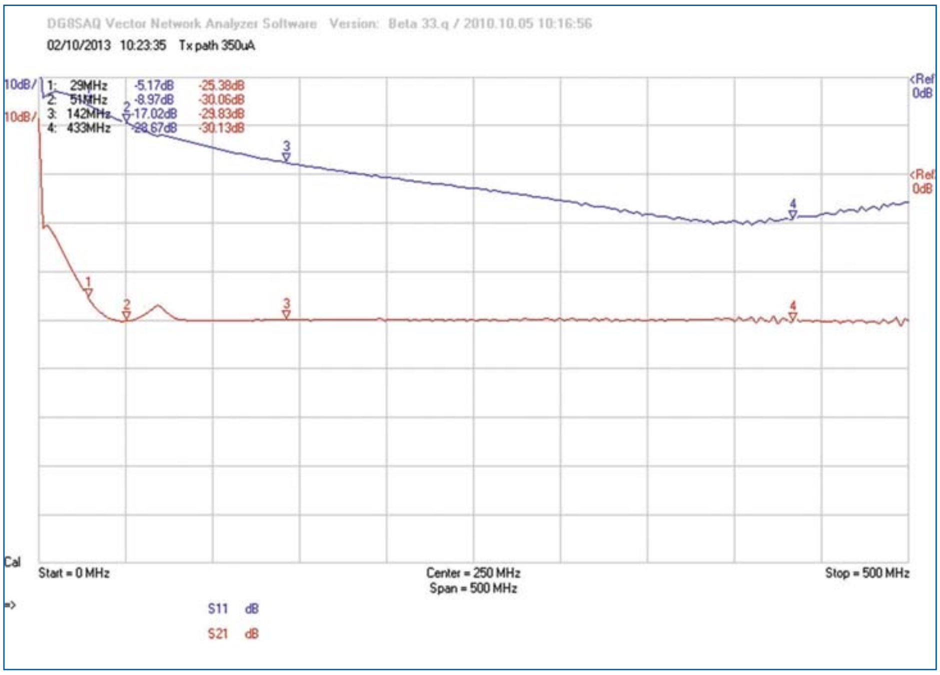

FIGURE 4: Network analyser plot of the transverter interface path on transmit (the red S21 plot). Note the broad constant attenuation above 100MHz. This can be varied by adjusting the diode forward current. S11 shows the return loss at the transmitter input. A value here of 17dB corresponds to a VSWR of about 1.3:1.

the diode drop) superimposed across it and defines the current that passes though the forward conducting diode. With the values given this is around 7mA and results in a diode resistance of around 3Ω for the BAP64 devices used. The unpowered line (RXP when in transmit) is maintained at zero volts with a pull-down resistor and, since the common of the diodes is at +V, means the one in the RF path not in use is reverse biassed to minimise junction capacitance and give the best possible isolation.

On transmit the RF passes first into a 6dB attenuator stage to dissipate most of the power. I build everything using SMT components, so several 0805 sized 0.1W resistors were used in parallel-series combinations to give the required attenuator values and stay within their rating. A BAP70-03 attenuator diode follows, the 4k7 resistor feeding this allowing a control voltage of 1 to 5V to give a resistance variation of around 200Ω to 30Ω This results in additional attenuation that can be varied over approximately 5 to 16dB. Intermediate resistors either side of this were selected to give an acceptable match on the Tx input for all diode current settings and to supply the attenuation range needed. A final fixed attenuator of 10dB completes the chain to make up the total attenuation for the path. Another switching diode transfers the attenuated RF signal to the mixer port. On receive, the other pair of PIN diodes conduct, generating an almost lossless path from mixer to transceiver Rx input. Note that in neither case is the second PIN diode, on the right hand side of the diagram, run with reverse bias. In practice there is sufficient isolation from the first one to make this unnecessary.

MEASURED RESULTS. A DG8SAQ network analyser was used to measure performance at 144 and 432MHz. The important parameters are the loss in the receive chain due to non-zero ON resistance of the switching diodes, the leakage of Tx power into the Rx path to the mixer that would interfere with the adjustable attenuation (Rx and Tx mixer ports were separated for the measurement), and the variable attenuation range.

Figure 3 shows a plot of transmit attenuation versus control current and **Figure 4** shows a network analyser plot of the Tx path for 350μA over the frequency range 0 to 500MHz. Performance drops off below 144MHz due to the value of coupling capacitors and bias inductors chosen. Operation at 50 or 28MHz would be possible while still using the same choice of PIN diode by increasing these components to 10nF and 10μH respectively.

TABLE 1: Measured results of the transverter interface.

		144MHz	432MHz
Rx path loss		1.7dB	1.8dB
Rx path isolation on Tx		67dB	48dB
Tx attenuation	I_{FWD} = 100μA (~1.3V)	37.8dB	38.5dB
Tx attenuation	I_{FWD} = 350μA	29.8dB	30.1dB
Tx attenuation	I_{FWD} = 1mA (5V)	25.4dB	25.6dB

Photo 1 shows the complete transverter interface module built on a 53 x 35mm PCB with a ground plane on the underside. The two PIN diode types used in this project are available from RS Components **[1]** for a few pence each (subject to buying a minimum quantity of 20).

GETTING HOLD OF SMT COMPONENTS.

For years I have always built everything using surface mount components, as I find them very much easier and convenient to work with than wire ended ones, allow simpler and better PCB layouts and behave better at high frequencies. They are also easier to get hold of than wire ended components and generally cheaper. You do often have to buy a seemingly large number at a time – but the ridiculously low price of most routine items ends up still reasonable.

However, not everyone has access to the suppliers of these and for a useful selection of SMT resistors, capacitors inductors and several semiconductor types there is an (almost) free alternative for supplies of small quantities of anything on offer.

The UK Microwave Group (UKuG) has over the years been the recipient of several gifts of surplus stock and has built up quite a collection of SMT components – the Chip Bank **[2]**. These are available *completely free* to any member of the UKuWG who needs a supply of SMT bits, and UKuWG will even pay the postage as a benefit to members. (Apparently, it is simpler that way). The only proviso is that you do have to actually be a member of the Group and anyone requesting unnaturally large quantities to subsequently sell on eBay will be spotted and "dealt with"! UKuG represents those amateurs who are into higher frequency operation where homebrewing and construction is the norm. Membership costs £6 per annum, and for that you get an electronic newsletter delivered about 10 times per year that covers microwave activity and operation and construction and design techniques for the bands. Considering the £6 annual membership is less than most manufacturers' delivery charge for non account holders, membership for the Chip Bank alone is worth it.

Browse the catalogue on the UKuWG website to see the range of devices available. The 1μH inductors shown in Photo 1 came from here. The list gets added to regularly as new donations of stuff come in – SMT components are just so cheap now when supplied in bulk that manufacturers usually just throw their surplus stocks away.

REFERENCES

[1] RS Components: http://uk.rs-online.com/web/ (type 'PIN Diodes' into the search box)
[2] UK Microwave Group: www.microwavers.org/

Andy Talbot, G4JNT • e-mail: ac.talbot@btinternet.com

Design Notes

Taking a good look at filters

FILTERS. Most items of RF hardware need a filter in them somewhere, whether to remove harmonics on the output of a transmitter, kill image responses in a receiver mixer, or for some more esoteric task like amplitude shaping a keyer waveform. There is quite an array of software tools out there offering plug-and-play filter design, so many in fact that constructors can become bewildered by the choice of offerings and where to start. So, to help in designing your custom filter, we're going to cover the basics of filters and the theory, going back to first principles to make the starting point a bit clearer.

The simplest filter type, both in theory and visualisation is the low pass ladder L-C filter such that in **Figure 1**. All low pass filters pass signals with a certain minimum attenuation up to a corner frequency, Fc, then progressively attenuate at frequencies above Fc. For some filter types, when frequency is plotted on a logarithmic scale and amplitude in dB, the roll off with frequency is close to a straight line.

A PI network is shown with a shunt input capacitor as the first element, but a Tee network with input inductor is equally possible. Nearly all filter design techniques start from a low pass prototype then, when this has been defined, elements are transformed to high pass or band pass structures.

A high pass design swaps the positions of capacitors and inductors. Band pass designs are more complex and introduce the concept of bandwidth – we'll cover these later.

The first thing to consider is the order of the filter. For simple ladder designs this is simply the number of components in shunt and series. Figure 1 is therefore a fifth order design.

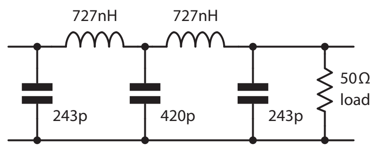

FIGURE 1: Typical lumped element 5th order low pass filter. This is a 0.1dB Chebyshev design with a cutoff of 15MHz. Typical of that used for harmonic filtering on a 14MHz transmitter.

NORMALISED FREQUENCY AND IMPEDANCE. To simplify discussion and filter basics, it is usual to talk in terms of normalised frequency. Here, the cutoff, or corner frequency Fc is given a value of exactly one. Then all other (actual RF) frequencies are converted to normalised values by dividing by Fc. So, for Figure 1, the response in the cutoff, or stopband region at 30MHz would correspond to a normalised frequency, F' = 2. Conversely, the response in the passband at 7MHz would correspond to approximately F' = 0.47

Using normalised frequency in this way allows standard, or 'normalised' component values to be determined for any particular response without having to worry about the actual design frequency or design impedance. When calculating final L and C values the maths come out such that F' = 1 is equivalent to one radian per second, or 2π Hz. The component values for a particular frequency are then scaled from the normalised values by multiplying by Fc / 2π.

Many filters operate in a 50Ω system, but not always. In the same way as it is done for frequency, a normalised impedance of 1Ω is assumed. To get the final component values, the frequency corrected values then scaled appropriately for the (usually) higher Zo of the final version. Inductors need to be multiplied by Zo and capacitor values divided by Zo.

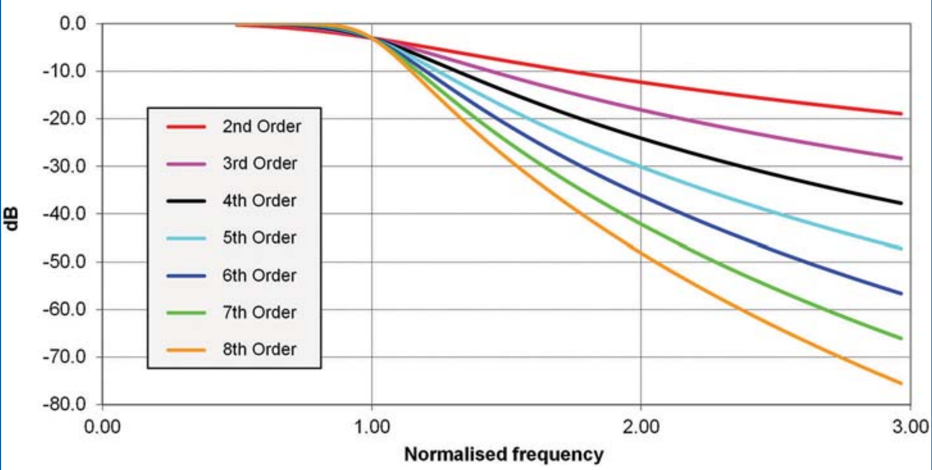

FIGURE 2: Attenuation versus Normalised frequency for Butterworth filters of different orders.

RESPONSE SHAPE AND MATCHING. The input and output impedance of a filter varies over frequency. Over the passband, we want as good a match as possible, so that a filter terminated with a resistor of Zo will look very close to Zo looking into the input terminals. How close to Zo it is, depends on the final application. A harmonic filter on a transmitter requires a good match at all frequencies of operation. An image filter on a receiver can often accept a quite high degree of mismatch if it only has to go between amplifier or attenuator stages. Impedance match and filter response are inextricably linked together.

A huge variety of filter responses is possible for any design of a given order, but for most cases we only really need to cover three, or possibly four families of response types. The simplest is the Butterworth, or maximally flat response as shown in **Figure 2** for up to 8th order. The passband is essentially flat, with no ripple, until it approaches Fc. At Fc the response is exactly 3dB down then falls linearly (on a log frequency / dB scale) at a rate of 20.N dB per decade, or 6.N dB / octave. The match is perfect at the bottom end, and progressively worsens as Fc is approached. In fact at Fc with 3dB attenuation the return loss is also 3dB, corresponding to a VSWR of 5.85:1. So, if a Butterworth design were to be used as an harmonic filter on the output of a transmitter, Fc would have to be made appreciably higher than the actual highest frequency of use. In fact, to guarantee a return loss of better than 20dB (VSWR = 1.22) on the filter input for a fifth order deign, F' needs to be below 0.64. For a 7th order design, this increases to F' = 0.73 for 20dB return loss. So to meet the required attenuation at harmonics, a higher order filter may be needed. For high power transmitters, the components can be quite expensive, so a low order filter is preferred. Butterworth filters have their uses, but these are mainly where the flatness of the passband is more important than complexity

The next family of designs to consider is the Chebyshev (sometimes spelt Tchebychev). These have the same circuit topology as Butterworth designs, but use different component values to deliberately introduce a ripple and corresponding mismatch into the passband. Ripple and mismatch are directly related: a return loss of 20dB (one hundredth) means 99% of the power is transmitted corresponding to a passband loss, or ripple, of 10.LOG(0.99) = 0.043dB. For this family, Fc = 1 is sometimes defined as being at the point on the roll off curve where attenuation reaches the ripple value, However, other references define cutoff as being at the 3dB

Andy Talbot, G4JNT • e-mail: ac.talbot@btinternet.com

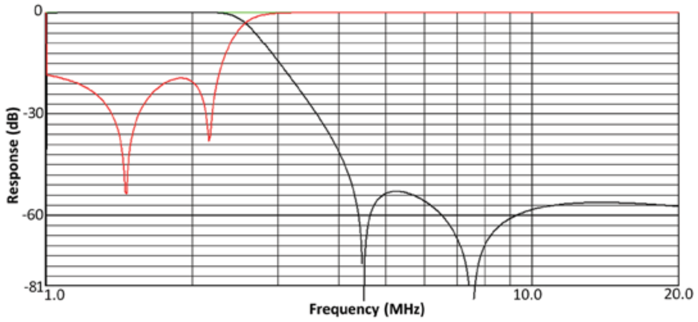

FIGURE 2: Output low pass filter frequency response.

that never resulted in a PA current of more than 2A. A separate 12V supply was used for the input circuitry to ensure the MOSFET always had an adequate switching waveform.

Figure 3 shows the measured value of PA efficiency (RF power out divided by DC power in) and **Figure 4** the actual RF output power, both with a PA supply voltage of 8.5V. The odd shape to the curve in Figure 3 is almost certainly due to measurement inaccuracies and it's not actually that shape – see the small bump that coincides on the P_{OUT} curve. It is notoriously difficult to measure RF efficiency with high levels of accuracy, and I didn't try to smooth the plot or frig anything. However, the general trend can be seen.

Optimum efficiency lies just slightly lower than the design frequency, at 1.87MHz, but the actual point is hard to determine; efficiency does not change by more than a couple of percent over the range 1.83 to 1.92MHz. However, output power does change considerably, from 11W down to 8W over the same range. So we have a range of operating frequencies over which the efficiency is more or less constant, but the output power is changing by 1.5dB. The best operating point is not obvious.

INDUCTORS AND CAPACITORS. There are a number of wound components in this design, and a few capacitors where voltage ratings are important. The tank inductor needs to be a good quality air wound one. It consists of 20 turns of 0.8mm wire close wound on a 15mm diameter former. The value was roughly determined from Rayner's formula **[2]** then fine tweaked using a G4HUP inductance meter. The DC feed choke consists of 32 turns of 0.5mm wire on a T68-2 iron dust core. The two filter inductors of 4.5 and 3.8μH are 28 and 26 turns respectively, again on a T68-2 core.

The 1:3 matching transformer is wound on a ferrite core, similar to FT50-43. The core type is not that critical, provided it is of a type designed for RF transformers. Matching transformers used in this way are usually transmission line types, made by twisting together sets of turns and winding the twisted set together around the core, then connecting individual windings as shown. Here, a trifilar set would normally be used for a 1:3 winding. However, I already had several metres of 5 core wire already twisted (ie quintifilar wound) and ready to use. So three strands out of the five were used in parallel for the 5.5Ω primary to give a thicker conductor on this side to reduce resistive losses. The remaining two strands were connected as normal to give the high impedance secondary. Six trifilar turns were used on the core (resulting in 18 in total across the 50Ω load).

The capacitor making up the tank circuit, shown as having a value of 2.5nF, was made from 1.5nF in parallel with 1nF. With a Q of 6, and assuming a peak output of 50W if a conventional AM high level modulator is used, they should have a voltage rating at least $6 * \sqrt{(50 * 5.7)} = 100$V. I had a batch of 100V polyester ones to hand. For a 14V supply, these are adequately rated, but an increase of input voltage on modulation peaks could cause problems. The capacitor in shunt with the MOSFET, with a $4*V_{DD}$ voltage peak expected, needs at least 100V rating.

Photo 1 shows the finished PA. Ideas for a suitable modulator will be covered next time.

HIGH SIDE CURRENT MONITOR CHIP. The July 2009 Short Circuits column described how current in a positive supply rail can be monitored. The small voltage drop across a series sense resistor is transferred via an opamp and transistor to a voltage that can be measured with respect to ground. This is often needed if power supply current is to be monitored together with other analogue parameters via a multi-input A/D converter that needs a ground reference input to function.

Several dedicated ICs are around to make such a high side current monitor a practicality. The INA169 available from **[4]** is one such device, suitable for supply rails up to 60V. This high voltage capability is particularly useful as 60V is well above the capability of most opamps that could be used in a discrete design. The 5 pin device is very flexible with regard to power supply and will work over an input voltage range of 2.7 to 60V. It generates a current output of 1mA per volt proportional to the voltage drop across the sense resistor. This output current is then driven though an external resistor to develop a voltage to measure. If, for example, a sense resistor of 0.02Ω were used in the positive rail, then a 10A load gives 200mV drop. The resulting output from the INA169 device is a 200μA current that, if driven though a 20k resistor to ground, gives 4V. This is a convenient value for a PIC-based ADC running from a 5V supply. Download the INA169 datasheet from [4] for more details

REFERENCES

[1] Calculating component values for class-E transmitters: http://g3nyk.ham-radio-op.net
[2] Rayner's formula for inductance: L (μH) = (N * D)^2 / (0.46.D + 1.02. l,) where N = number of turns, D = coil diameter and l = winding length
[3] *AADE Filter Designer*: www.aade.com/filter.htm
[4] RS Components; INA169 is part number 461-6586: https://uk.rs-online.com

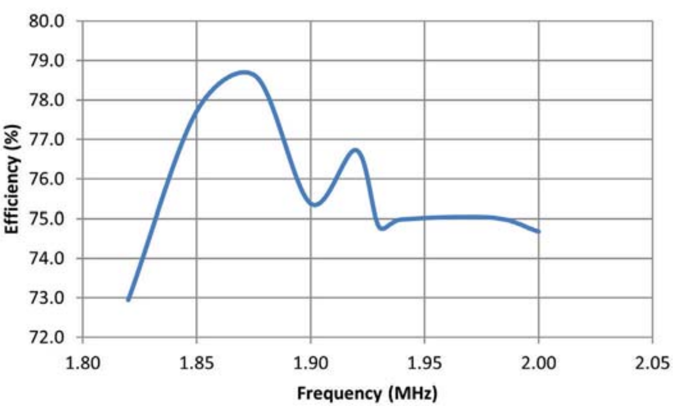

FIGURE 3: Measured value of PA efficiency (RF power out divided by DC power in).

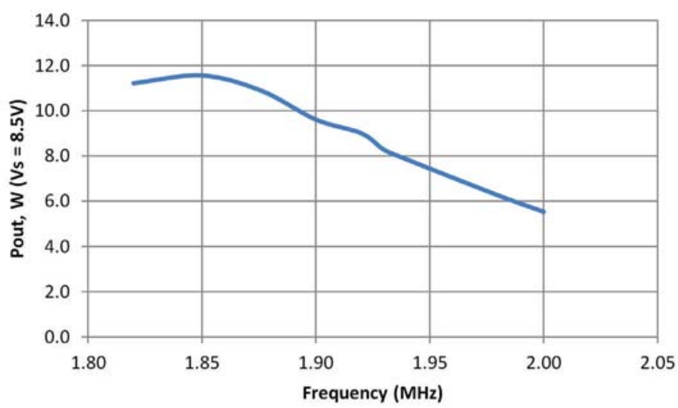

FIGURE 4: Actual RF output power with a PA supply voltage of 8.5V.

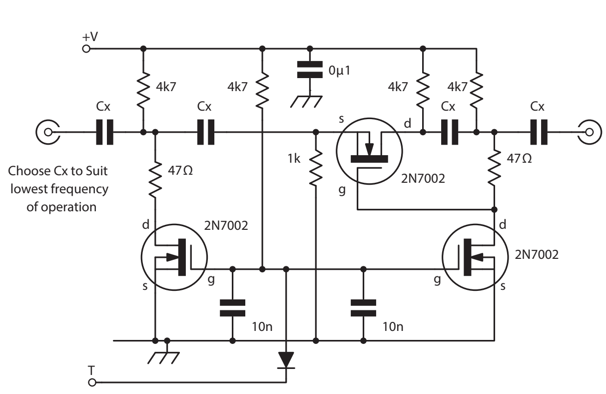

FIGURE 2: Circuit diagram of the LF switch, which provides over 60dB of isolation and a matched termination in the off state.

then the diode would become non-linear, approaching operation of a normal diode. While PIN diodes are made that function down to HF, they do not go down to the sub-MHz region.

A diode ring mixer would work as an RF switch, as detailed in this column a while ago, but its loss and power handling were inadequate. My Class E power amplifiers require a volt or two of input signal to function reliably. The solution is a FET switch. A single FET either in series with, or shunting the 50Ω drive signal would work perfectly well, and I had used a series FET in the dim-and-distant past for this very job. But such a switch is *just not nice*. When it is off, it causes a severe mismatch for both input and output ports. An open circuit in the case of a series switch element; a short circuit for a shunt switch. For my application, this would do, but thoughts idly turned to ways in which a FET switch could be made that presented a good match on both ports when off. The circuit in **Figure 2** was the result.

The control input is a ground to Tx signal applied to the terminal labelled 'T'. The diode is just to take care of any obscure DC level that may occur on the input control. Both series and shunt FET switches are employed. A series 2N7002 MOSFET connects the two ports for the 'on' state. When this is switched off, each port is shunted through two more devices to a terminating resistor. The series and shunt FETs need to be driven with opposite switching polarity. So rather than add a fourth device purely as a voltage inverter, I used one of the existing shunt devices to perform this function, inverting the switching polarity of the drive signal to control the series switch as well as doing its own RF switching. The RF switching results were excellent. **Figure 3** shows a network analyser plot of the switch in its on state, showing minimal loss and good through match. **Figure 4** shows a plot for the off state where 68dB of attenuation is achieved at 1.9MHz, with over 20dB return loss. The attenuation is still holding up at 14MHz, with 48dB.

As there are four DC blocking capacitors in series in the 50Ω path when the switch is on, they must have a higher value than would otherwise be used for a single DC block. 0.1µF capacitors proved too small a value to allow a decent through match at 137kHz, although these were more than adequate at HF and up. I had a considerable number of 4.7µF 10V ceramic capacitors from the UK Microwave Group Chipbank **[2]** (free to UKuG members) and used these. The series capacitors need to be non-polarised as the DC potential across some of them swaps as the switch is operated. The performance of the LF switch is serious overkill for its intended operation driving my LF PAs with the *Opera* RF waveform, but is a useful standby to have on hand as a piece of test equipment or a low loss matched switch for any future task. **Photo 1** shows the finished unit.

OLD PCS. I recently had to scrap a 2007 vintage desktop PC due to a failure on the motherboard. The hard disk drive, DVD and PSU were all recovered 'just in case'. I noticed that this PSU has two 12V supplies, each rated at around 17A (though not at the same time), as well as the normal +5 and +3.3V supplies. It seems modern PCs have their own switch mode regulators supplying the very high current and low voltage required by the modern processors. A bit of web searching suggested that the latest chips can require up to 100A at 0.8 to 2V – somewhat mind boggling! The finned heatsink cube and fan attached to the top of the CPU were more reminiscent of high power RF amplifiers than a computer! Nowadays, the bulk of the DC power is supplied to the motherboard at 12V to be reduced adaptively *in-situ*.

The PSUs are optimised to deliver the bulk of their output capability at 12V instead of on the 5V rail as before. This may be good news for us. Back in the February 2013 edition of *RadCom*, Mark Atherton, ZL3JVX wrote about the use of server type 12V PSUs for amateur applications, with units available at 23A or higher rating. But these are special PSUs that have to be purchased individually. Normal desktop PC power supplies ought to be far more available – as scrap from your nearest tip perhaps? I tried paralleling the two 12V outputs and nothing untoward happened. I found a twiddle pot inside that allowed the voltage to be raised to 12.7V, although the 5V one rose with it, increasing to 5.4V. As I wanted to keep this PSU in its original state for powering computer equipment, no further investigations were made and the lid was screwed back on. However, I see no reason why it couldn't be easily modified to deliver the 13.8V supply many of us want.

READING HARD DISKS EXTERNALLY. That failure occurred after I had transferred most of the data from the old PC to backup, but there were still quite a lot of obscure files in odd places that hadn't been considered worth backing up. (The synthesiser phase noise plots on page 24 were amongst them). A bit of research revealed that it was possible to connect the high speed serial (SATA) interface used on later disc drives to USB via a low cost interface. Looking on eBay for 'USB SATA adapter' threw up many for sale, and I purchased two for the grand sum of £3.78 each. Using the old PC PSU to then power the drive, the adapter gave full access to my old files.

Two lessons here: 1) always back up everything, even if you think certain odd files are not worth doing. You'll want them later if you don't! 2), if you haven't backed them up, all is not necessarily lost when things do fail.

REFERENCES

[1] Ideal bridge rectifier controller chip – www.linear.com/product/LT4320

[2] UK Microwave Group Chipbank, with free SMT components on offer to members – www.microwavers.org/

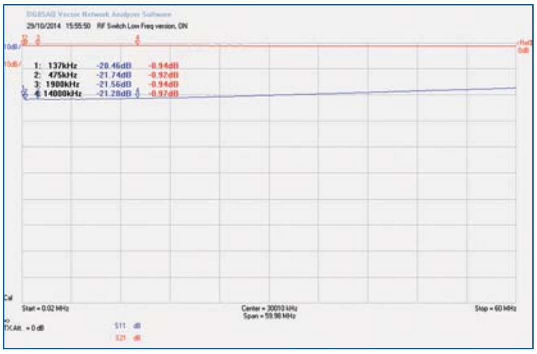

FIGURE 3: Match and input return loss of the FET switch in the ON state.

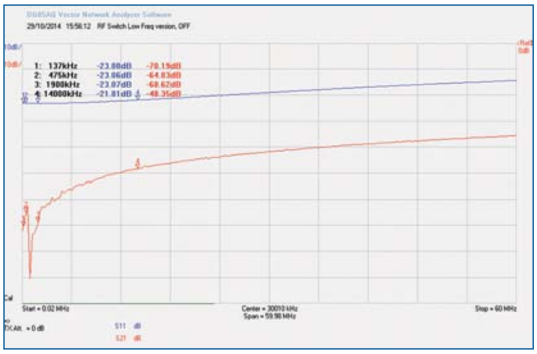

FIGURE 4: Match and input return loss of the FET switch in the OFF state.

Andy Talbot, G4JNT • e-mail: ac.talbot@btinternet.com

Design Notes

What has our experimenter been up to this month?

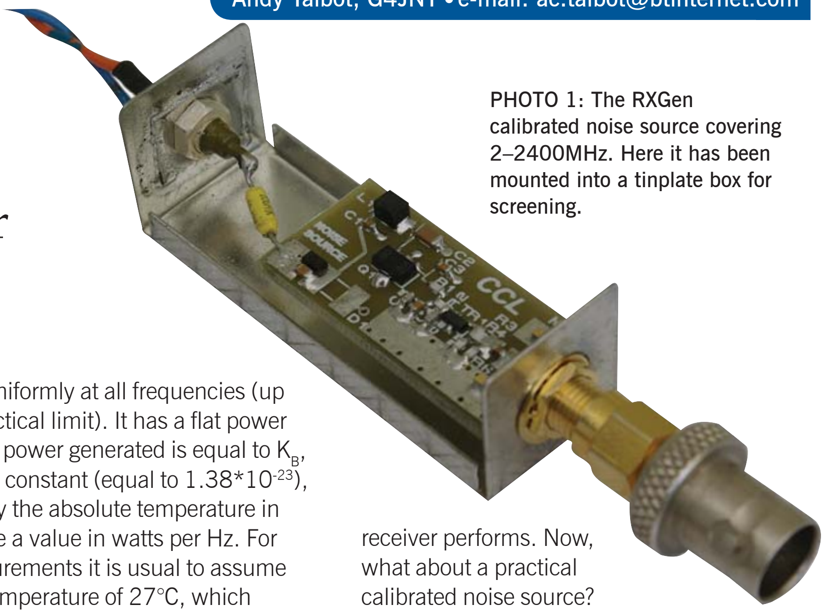

PHOTO 1: The RXGen calibrated noise source covering 2–2400MHz. Here it has been mounted into a tinplate box for screening.

HIGH POWER PIN DIODE SWITCH. John Pink, G8MM wrote in to mention this useful component he came across. "MACOM have introduced an interesting new product – a PIN switch capable of handling 100W that they claim will operate over the range 30MHz to 3GHz. Take a look at the MA-Com site [1]. It looks particularly interesting for VHF and microwave applications; perhaps even for preamp switching. The datasheet shows it as having 60dB isolation."

NOISE FIGURE AND RECEIVER SENSITIVITY. The ultimate weak signal capability of any receiver is governed by noise. At LF through to HF bands, and more these days up into VHF too, this noise is usually generated outside the equipment and received via the antenna. In quieter areas, and certainly at frequencies in the upper VHF through to microwave region, it is the receiver's own thermal noise that dominates and it is up here that we often strive for the best sensitivity, particularly for space working such as moonbounce or satellite communications. The usual way to refer to the sensitivity of a receiver is by its noise figure (NF), but to understand what this means we need to go back to the basics of what noise actually is.

The electrons in a conductive medium oscillate randomly at any temperature above absolute zero (0K, about -273°C) – faster and more violently as temperature rises. The effect is to generate small random currents that result in power being delivered to a load. As the variations are random and their speed or rate of vibration can be anything from just above DC to 'very fast', you can probably visualise that the power will be generated uniformly at all frequencies (up to some practical limit). It has a flat power density. The power generated is equal to K_B, Boltzmann's constant (equal to $1.38*10^{-23}$), multiplied by the absolute temperature in kelvin to give a value in watts per Hz. For noise measurements it is usual to assume a working temperature of 27°C, which equates to 290 kelvin. At this temperature the noise power that can be delivered to a load then becomes $4.0*10^{-21}$ watts/Hz, more commonly expressed as -174dBm/Hz. This defines an absolute lower limit to the sensitivity of a radio receiver when operating at normal room temperature and when the antenna is aimed at terrestrial sources.

In any practical receiver, imperfections and losses mean the sensitivity is reduced and the effective noise level is raised. Signals now have to be that much stronger to overcome the higher internally generated noise. Noise figure is defined as being the degradation in signal to noise ratio from that present at the input of the receiver to that measured at the output. But, although we can physically measure output S/N ratio, how do we know what is at the input in the first place? There is also the added complication that bandwidth is often unknown, so we need a method that is independent of signal type or bandwidth. So we make use of a calibrated noise source that generates a precisely known level of noise that is flat across the frequency spectrum at an amplitude that is greater than thermal noise by a precisely known amount. Noise sources are usually specified by their Excess Noise Ratio (ENR). Usually expressed in dB, ENR is the ratio of noise power when the source is switched on to the power of thermal noise when it is off. Since thermal noise power depends on temperature, a standard source-off temperature of 290K (27°C) is always assumed and ideally maintained at the test bench. By switching the source on and off, we generate a test input with a precisely known S/N ratio at the receiver input. We now measure the resulting output S/N ratio and, by comparing the two, work out how much it is reduced, and hence how our receiver performs. Now, what about a practical calibrated noise source?

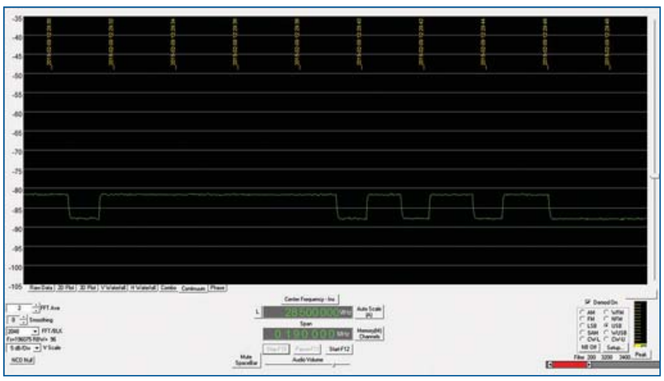

FIGURE 1: Screen capture of *Spectravue* in continuum mode used for noise level measurement. Noise source driving a Microwave Modules converter feeding the SDR-IQ receiver. Y factor = 6.5dB.

NOISE SOURCES. The most accurate and absolutely calibrated test noise source is simply to heat a resistor. We know that thermal noise = K_B.T watts per Hz, so by maintaining one resistor in liquid nitrogen at 77K and one in boiling water at 393K, by switching between them we know the change in level is exactly 393 / 77 = 7.08dB. This is OK for calibration labs, but totally impractical for the test bench – where by the definition of noise figure we ideally need the 'cold' one to be at 290K anyway. So, enter the noise diode.

Most semiconductor diodes operated with avalanche or Zener breakdown generate appreciable noise. We can make our own using a reverse biassed base-emitter junction of a transistor. A BC108 type device breaks down at around 9V. Try passing a 10-40μA backwards though the B-E junction, couple a reasonably sensitive receiver onto the device and listen for the added noise.

Proper noise figure measurements are made using custom noise diodes that are designed to generate a uniformly flat spectrum from low MHz to several GHz. They are initially calibrated against hot/cold sources, or 'better' noise diodes that have already been calibrated, then made available with their own calibration data. Until recently, the only noise sources available to amateurs have been either surplus professional ones – usually quite expensive if they have kept their calibration – or homebrew ones that need calibration from scratch. Remember, if the ENR isn't accurately known to start with, it is impossible to know how much the S/N has been degraded after passing through the receiver.

RXGEN CALIBRATED SOURCE. A calibrated diode-based noise source, suitable for the majority of amateur applications, is now available from Continental Compliance Ltd [2]. The RXGen consists of a noise diode, voltage regulator and decoupling components for a stand-alone noise source. The supplied PCB module is shown in **Photo 1**, where it has been installed in a tinplate screening

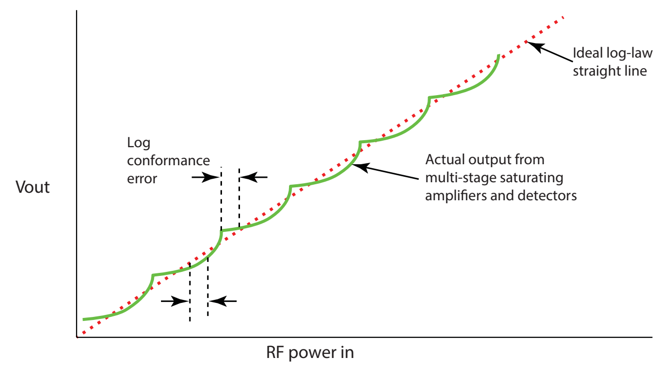

FIGURE 3: The distortion or log conformance error inherent in the multi-stage approximation to logamp detectors.

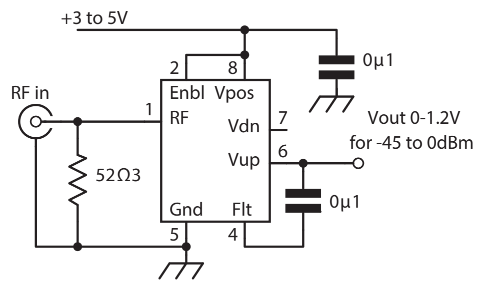

FIGURE 4: Basic connections for using the AD8314 100MHz to 2.5GHz power detector chip

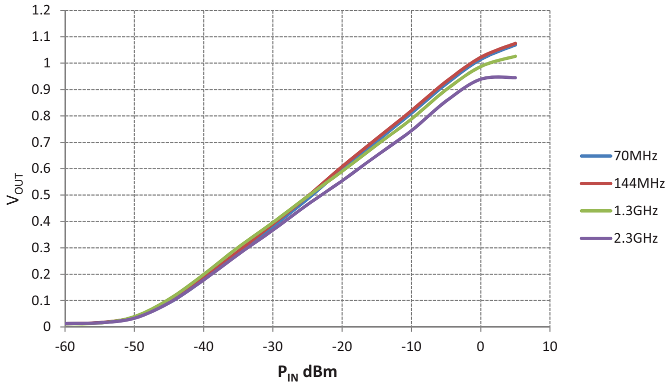

FIGURE 5: Measured results on the breadboard AD8314 at four frequencies.

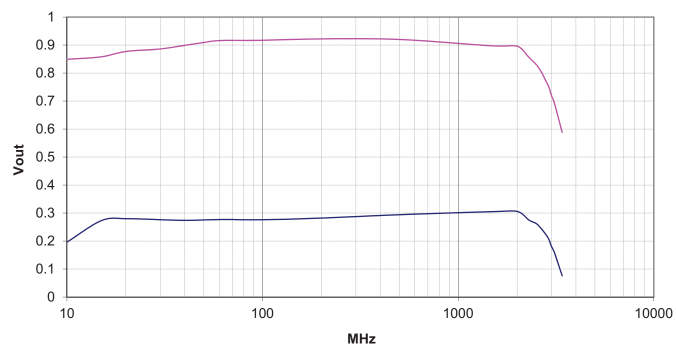

FIGURE 6: Variation of output voltage with frequency for fixed -35dBm and -5dBm input power levels.

AD8314 UHF TO MICROWAVE POWER METER. Two devices look appealing as they are so cheap. Their dynamic range is a fair bit lower but for the majority of purposes, do we actually need 90dB measurement range in one go? And perhaps, having fewer internal amplifier stages, the log conformance (accuracy) may be improved. Certainly, resolution on an output meter will be greater for a lower dynamic range The data sheets for the AD8312 and 14 devices don't actually quote a single figure for log conformance; the values shown in brackets in Table 1 are taken from a family of curves, so the same criteria may not have been applied to getting an overall figure of merit.

Some AD8314 chips were purchased and one was put through its paces. **Figure 4** shows the simple circuit configuration for the test bed of **Photo 1**. The chip was mounted on a small homebrew PCB with two parallel input terminating resistors to give the required value, and just a couple of 100nF capacitors for supply decoupling and output voltage filtering. See the datasheet for full details of how this chip can be used, but the basic connections shown are enough to start with.

Figure 5 shows a plot of output voltage versus RF input power at four frequencies, showing that there is a definite roll off at the upper end of its frequency range. This does not manifest itself just as an RF gain reduction, which would be a shift of the curve to the right, but as a change in the slope of the graph, measured in mV/dB. However, for frequencies up to above 1GHz, there is almost imperceptible change in performance from at least 70MHz right up to 1300MHz. For completeness, **Figure 6** shows a plot of frequency response at two fixed input levels of -35dBm and -5dBm.

NOISE MEASUREMENT. Last month we looked at measuring noise figure and noted the need for a high resolution readout, especially for low Y factors associated with higher noise figure equipments. The random nature of noise means that the log conformance errors will be averaged out when the random amplitude spikes of white noise are applied to the chip. However, the wide dynamic range available is way more than is needed for noise figure testing, so accuracy will be reduced simply by trying to make out the possibly sub-millivolt change for a small increase in input noise. Here the lower dynamic range log detector chips are likely to show an advantage over the wider range '07 and '10 devices.

SWR METERS. In conjunction with a directional coupler, a pair of such chips can form a versatile automatic SWR meter and power monitor. With one detector on the forward and one on the reflected ports of an in-line directional coupler, to calculate the match or return loss (and hence SWR), we divide the return power reading by the forward one to obtain a ratio independent of the actual level. As we are generating a voltage corresponding to the log of the values, the log of the ratio is given by simply subtracting the reverse reading from that of the forward. So we can generate a direct reading of return loss in dB just by placing a voltmeter between the two DC outputs from the two detectors. Another voltmeter on the forward detector can be used as a power monitor. A number of such projects can be found on the web.

TABLE 1: Summary of Analog Devices log detector ICs.

Type	Frequency range	Dynamic range	Linearity	Approx cost
AD8307	DC – 500MHz	92dB	±1dB	£12
AD8310	DC – 440MHz	95dB	±0.4dB	£10
AD8312	50 – 3500MHz	45dB	(±0.4dB)	£2.40 (6 pad BGA)
AD8313	100 – 2500MHz	70dB	±1dB	£16
AD8314	100 – 2700MHz	45dB	(±0.6dB)	£1.60

REFERENCES

[1] For measuring power levels from a few tens of milliwatts to a few watts, the most accurate solution for amateurs is probably just a diode across a 50Ω resistor, into a high impedance voltmeter like a DVM. The diode gives the peak rectified voltage so, after allowing for the diode forward voltage drop, V_D, the power in W is given by $(V_{MEAS} - V_D)^2 / 100$. Assume a V_D of 0.2V for a small Schottky diode, or 0.6V for IN914 types, so 100mW will result in around 2.8V and 5W will give 22V.

[2] Farnell Electronic Components, http://uk.farnell.com/ – search for device part number to download datasheets.

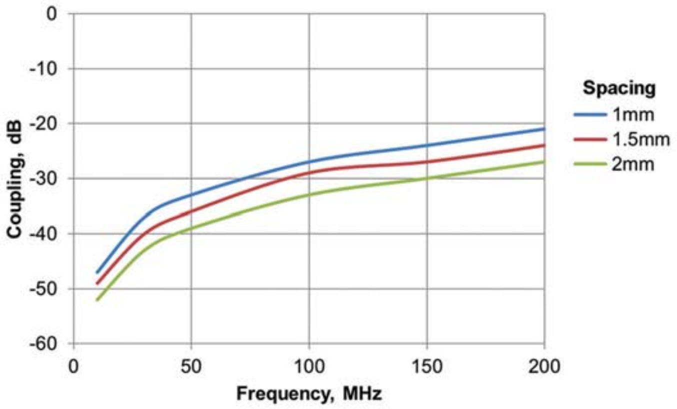

FIGURE 2: Measured performance of the three 66m long FR4-based test couplers over different frequencies with different line spacings.

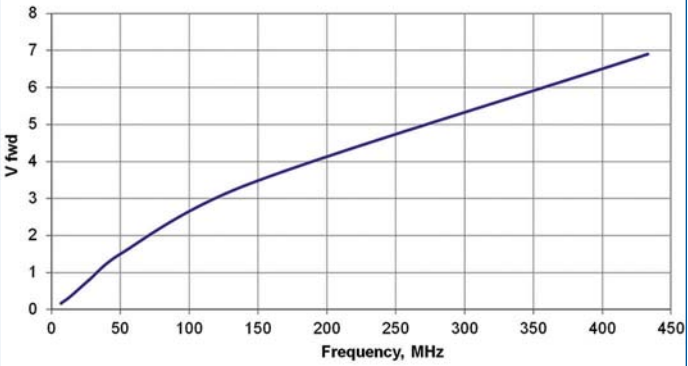

FIGURE 3: Frequency response of the 1.5mm spacing 66mm microstrip forward detector diode detector using with 5W of RF.

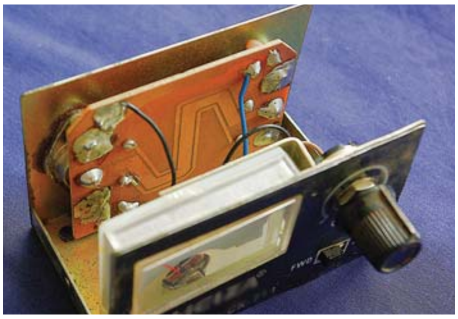

PHOTO 2: Low cost ex-CB SWR meter innards. Note the simple open-air printed coupling lines on the single sided PCB.

Microstrip is usually made by etching tracks onto a printed circuit board with a continuous ground plane on the reverse. The technology is so convenient that is has been analysed to death, with dozens of online design tools available to users of the medium. Two such online tools are given at **[1]** and **[2]**. You specify the dielectric constant Er, its height and the thickness of the copper. The software generates Zo for any specified width so to get a specific value of Zo, various values of W need to be tested. The version at [2] includes a synthesis function that allows Zo to be specified to generate a value for W. This calculation is often less accurate than analysis based on width input, so the result should always be checked 'backwards'.

PRACTICAL MICROSTRIP COUPLER TESTS – AND SOME SURPRISES. Three test couplers were built on standard FR4 fibreglass PCB as shown in **Photo 1**. Dielectric thickness was 1.59mm with Er = 5.0 **[3]**, copper thickness = 36μm. Both online calculators gave a width of 2.71mm for 50Ω line. An arbitrary value of 75Ω was selected for Zo of the coupled line, which required a 1.22mm line width. Three test couplers were built, with the spacing between main and coupled lines of 1.0, 1.5 and 2.0mm.

On test, I expected to have to use 75Ω resistors on each of the coupled lines but using this value resulted in very poor directivity. Experimenting with different values of resistor to find the optimum gave slightly different results for each of the three line spacings. The values required (to the nearest E24 resistor value) were 110Ω, 120Ω, and 160Ω for 1.0, 1.5 and 2mm spacing respectively; appreciably higher than the expected values. The influence of the adjacent line does influence Zo of each line, but for low values of coupling it was surprising to see such variation. The fact the line is electrically short also contributes to the discrepancy, and suggests, perhaps, the whole thing is acting more as a lumped element bridge rather than a true directional coupler. But it works and directivities in the range 20 – 30dB could be achieved with these load resistors. Improved directivity is no doubt possible with further trimming. Turn the bridge round to optimise each line, adjusting for minimum reflected power in each case.

With electrically short lines, coupling is dependent on frequency and **Figure 2** shows the values over the frequency bands of interest showing how coupling falls rapidly as frequency is lowered. **Figure 3** shows the output voltage from a diode detector on the forward line of the 1.5mm spacing unit with 5W of RF.

A simple design like this is so easy to replicate that its use can extend outside actual measurement. Used with a voltage threshold detector it could sound an alarm if the antenna becomes disconnected. Or build one into a PA for local power and poor match warning.

INSIDE A CHEAP COMMERCIAL BRIDGE. **Photo 2** shows the innards of a low cost CB type SWR bridge purchased at a rally. The main and two coupled lines are clearly visible. The termination resistors and diode detectors are on the rear of the PCB. What is really surprising is that there is no ground plane! This is not microstrip; the PCB dielectric has little effect and the transmission lines are, to all intents and purposes, in air. It could probably be considered a sort of stripline using the rear face of the housing as a groundplane, but if so the estimated Zo works out in the region of 250Ω. So the manufacturers made no attempt to match the line but as it is electrically so short at 27MHz doesn't matter too much.

The big surprise was how good its directionality was when a good matched 50Ω load was applied. Using a 100W transmitter to make sure the meter was getting plenty of rectified volts, the measured SWR (shown on the meter) of the load varied from less than 1.1 at 2MHz to around 1.3 at 50MHz. Even at 144MHz it was suggesting a SWR better than 1.5:1. Those values correspond to directivity values of 26dB, 18dB and 14dB respectively. Using a 150Ω load showed a SWR not far off 3:1, suggesting the FWD/REV balance isn't too far out, ie the forward and return paths are matched. All in all, not too bad for a £4 rally purchase. This does show just how uncritical design of SWR bridge heads can be for 'adequate' performance.

One even bigger surprise was to see that the termination resistors weren't tweakable or presets. They were just two fixed value wire ended resistors. At some point they must have been selected, but presumably the manufacturing process is stable enough to remove any need for final calibration.

REFERENCES

[1] Microstrip calculator – www.cepd.com/calculators/microstrip.htm

[2] Microstrip calculator with synthesis function – http://www1.sphere.ne.jp/i-lab/ilab/tool/ms_line_e.htm

[3] The dielectric constant Er of standard PCB material is often not specified and can vary across manufacturers and batches. Sometimes the value is supplied, but if not, it is an easy parameter to measure if you have a capacitance meter like the G4HUP design. Take a square of double sided PCB, measure its dimensions accurately and calculate the area *A* in m^2. Determine the thickness, *d*. Measure the capacitance *C* between the two sides. For a parallel plate capacitor, C (pF) = 8.854 * Er * *A* / *d*. So, knowing A and d, Er can be calculated from the measured capacitance.

So to get best accuracy, I carefully removed the insulation (using an absolute minimum of force to avoid squashing it) then, with a tiny amount of solder and liquid flux, tinned the wire so solder flowed into the strands but contributed nothing to overall diameter. Using a pair of vernier callipers the two dimensions of my cable were D = 2.18mm, d = 1.1mm.

The oft-quoted equation for twin line is Zo = 276 * vf * LOG(2D/d) but this is only accurate for wide spaced line where D >> d. Here the more accurate equation Zo = 120 * vf * $COSH^{-1}$(D/d) has to be used. Now we need to know the velocity factor to get Zo. So again, guesswork or looking at tables is the only way.

For my twin feeder, knowing vf = 0.7: Zo = 120 x 0.7 x $COSH^{-1}$(2.18mm / 1.1mm) = 105Ω (which for my mechanical engineering / measurement skills is pretty impressive!)

($COSH^{-1}$, otherwise known as ARC-COSH, is available on many scientific calculators by pressing something like INV-HYP-COS, or ARC-HYP-COS or $HYP\text{-}COS^{-1}$. In the Microsoft Excel spreadsheet, it is the obtained using the function 'ACOSH').

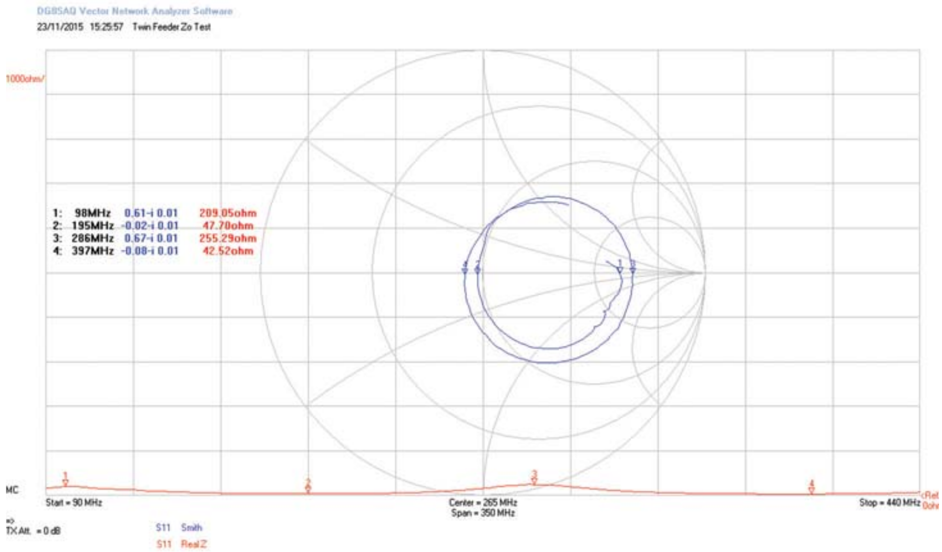

FIGURE 2: Network analyser plot of a 513mm test piece of twin feeder of unknown characteristic impedance terminated in a 50Ω load.

So the twin feed is nominally 100Ω as sold by many suppliers. Not what I wanted, but useful to have on the shelf.

REFERENCE

[1] Clicklock, locking a soundcard to 1 PPS GPS signals: www.qsl.net/zl1bpu/SOFT/click.htm

Design Notes

Last month we looked at two low cost transceiver modules made by DORJI. These modules have no harmonic filtering on the output and if used straight into an antenna could, and probably would, radiate on non-amateur frequencies, making their use illegal if left as-is. No mention of this lack of filtering, or indeed any indication of harmonic levels is indicated in the data sheet.

PHOTO 1: The two breadboard filters, constructed using SMT capacitors on a generic filter PCB.

Transmitter output filtering

So… we need external low pass output filtering to attenuate these harmonics. But what is an acceptable level? The CEPT requirements for commercial equipment define spurious outputs (not just harmonics) to some extent based on the output power and the frequency band. The specification is that all unwanted spurious outputs outside the necessary bandwidth are below the carrier (ie dBc) by 43 + 10.LOG(P) dB, where P is the power in watts.

The equation holds up to a maximum of -50dBc for frequency products below 30MHz and -70dBc for all products above 30MHz. The lower value of -50dBc only really holds for HF equipment operated at the bottom end of the band, so we will only consider the -70dBc case for frequencies over 30MHz.

In practice this means that a low power VHF transmitter of 1 watt needs to have its harmonics attenuated by at least 43dB, but a 100 watt VHF transmitter requires 63dB. The limiting case is at 500 watts, where suppression fixes at -70dBc Of course, as amateurs we are not *obliged* to meet these requirements in any home built kit – but there is no excuse whatsoever for not doing so; and if it subsequently does cause interference – enough said.

The harmonic levels I measured on the VHF and UHF DORJI modules are shown in Table 1. From this we can see the main problem with the VHF unit is going to be attenuating its second harmonic by at least 33dB to get it from -10dBc to below -43dBc.

TABLE 1: Measured harmonic levels on the DORJI modules at 1W output.

Module	2nd	3rd	4th harmonic
VHF	-10dBc	-22dBc	-28dBc
UHF	-19dBc	-57dBc	(all others <-70dBc)

Low pass filter design

Design of simple inductor capacitor low pass filters is, nowadays, almost a plug-and-play process using filter design software such as *AADE Filter Designer* [1]. All designs presented here use the diagrams from *AADE* but many other filter design packages are available. Search on the web for 'RF filter design software' to find them.

Filters for the output from a transmitter come with one or two special requirements, and some relaxations that can make for neat short cuts. All low pass filters allow the designer to specify a certain amount of passband ripple. Higher ripple generally means a sharper cut off but also means that the input impedance, at the points of peak ripple, may be quite poor. To show what this means, consider two theoretical designs for a 5^{th} order, low pass filter with a cutoff of 150MHz. Figure 1 shows the ideal Chebyshev values and response for a 0.05dB ripple design. Figure 2 shows the changes that happen when a ripple of 0.5dB is specified. In each case the red curve shows the input return loss [2] and you can see that for the latter case is goes to a rather poor 10dB at the points of maximum ripple (VSWR = 2:1). BUT, see how at other frequency values its return loss is still acceptable. Contrast with the 0.05dB ripple case where return loss is acceptable at better than -20dB over the whole of the passband – but a slower roll-off, giving less attenuation of the second harmonic. High ripple filters are often used in low level stages where match is not critical; transmitter output filters call for low ripple designs, or careful placing of the point of maximum ripple.

Neither of the two designs is suitable as it stands as the output filter for the DORJII module. In the 0.05dB ripple version, the attenuation at 290MHz is just 28dB and not sufficient for our needs. In the other, the match at 145MHz is too poor at just 9dB return loss. Also, the component values are the 'ideal' ones generated by the design software. We will have tolerances and unknown inductors to contend with, any departure from the ideal will alter the response. Fortunately the *AADE* software allows users to change components at will and see what happens to the response as things are varied (as do all the other packages). This works to our advantage as it allows 'odd' and custom designs to be modelled and shows how they can be adjusted for a final version.

Other filter responses could be used, but offer little advantage in this application. The Butterworth, or zero ripple, response has a slower roll off and offers nothing over a low ripple Chebyshev. An

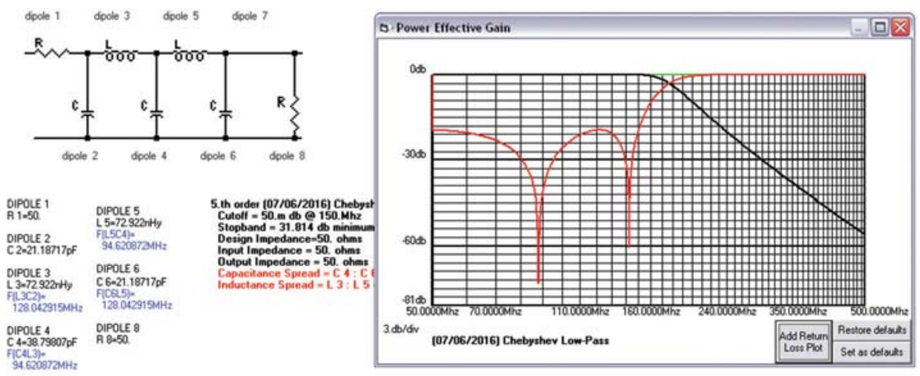

FIGURE 1: Idealised *AADE* filter for 150MHz LPF with 0.05dB passband ripple, with a good match over all the passband.

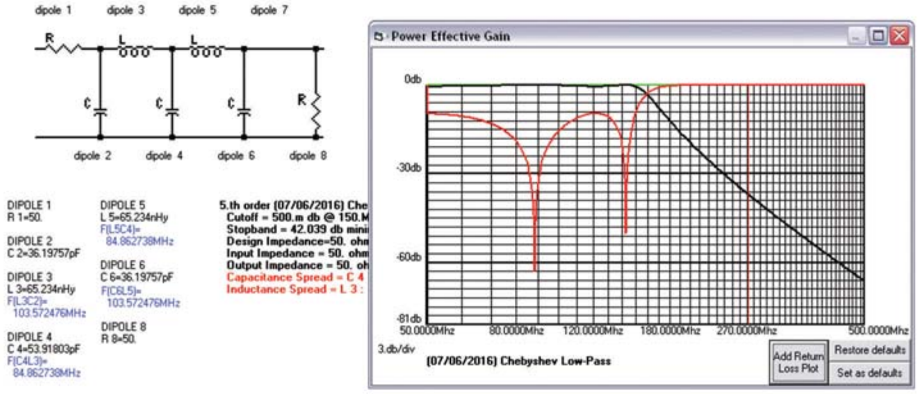

FIGURE 2: Idealised *AADE* filter for 150MHz LPF with 0.5dB passband ripple, showing the poor match at certain frequencies in the passband.

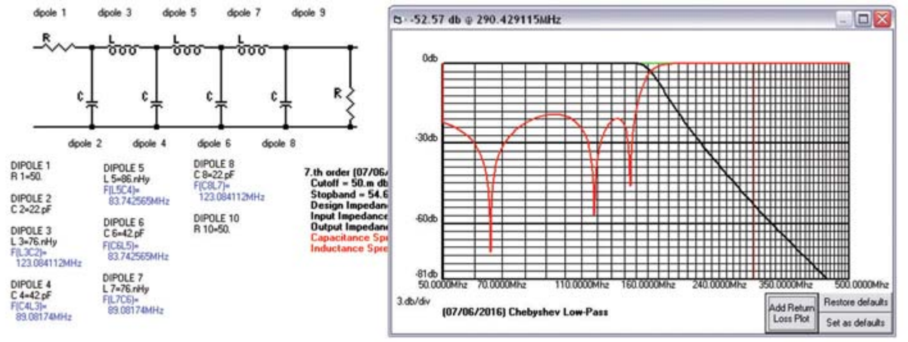

FIGURE 3: *AADE* design for 7^{th} order low ripple filter with components adjusted to 'convenient' values.

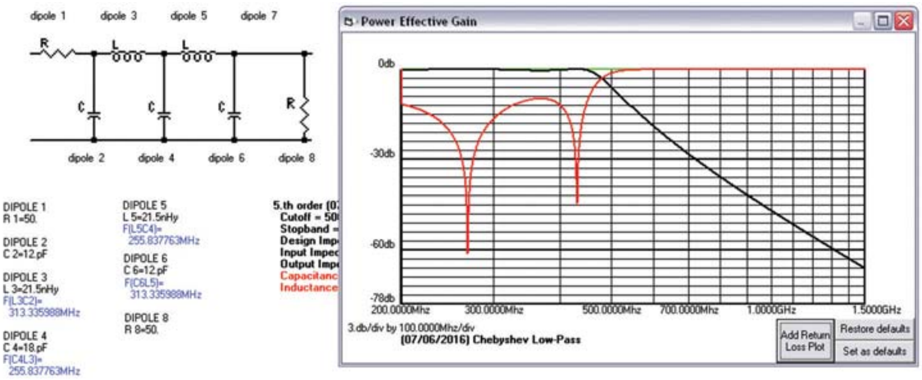

FIGURE 4: 450MHz high ripple filter adjusted to place an optimum match region at the frequency of interest.

elliptic design, using capacitors across the top Tee inductors to give notches of near infinite attenuation at certain stopband frequencies can be very helpful in cases where a wider passband and sharp cutoff is desired. Tuning the inductors by setting the correct notch frequency can also make for very easy adjustment. But in this application modelling showed no advantage to an elliptic over the 7^{th} order Chebyshev in terms of the total number of components.

Some practical designs

The easy way is just to increase the filter complexity and use a 7^{th} order low ripple design. **Figure 3** shows such a filter, that more than meets the requirements for passband return loss of 20dB and second harmonic attenuation of 52dB. The component values have been adjusted to the nearest pF and nH and this filter can be built and tuned as described later.

Now, what about the UHF module? The second harmonic level is already at -19dBc so we only need an extra 24dB to meet the requirements. If the design of **Figure 1** were scaled in frequency it would just about do the job with a few dB in hand, but let's have a play and see what can be done with a higher ripple 5^{th} order design, adjusting component values to optimise the position of the ripple minimum.

If we start out with a 450MHz low pass Chebyshev design specifying 0.5B ripple, worst case return loss is around 10dB (as before). Now look back at **Figure 2**, which is just a frequency scaled version of this one. See how the second passband dip lies at 140MHz. On the 0.5dB UHF design, this dip scales to 429MHz – tantalisingly close to where we want it! Can we play with component values and cutoff frequency to move this dip to 433MHz? One way is to scale the design cutoff frequency, to move it higher in the ratio of 433/428 so the ripple peak moves in sympathy. Modelling for a cutoff of 454MHz, the result of **Figure 4** is generated. Again the ideal component values have been adjusted to the nearest pF and (this time) 0.5nH resolution. The return loss at 433MHz is a satisfactory 20dB and, more importantly, the attenuation at the second harmonic is 40dB. Perfect.

Building them

L/C filters really lend themselves to just about any construction technique suited to RF circuitry. Rats nest layout over a copper ground plane works well, but for a more robust solution a PCB is preferable. The physical layout can follow the electrical schematic remarkably closely, and for designs up to UHF a single sided PCB can suffice if size is kept down. I already had some small PCBs made up for filters and use SMT components for everything. **Photo 1** shows my breadboard for the two filters – the 7^{th} order VHF low ripple design and the optimised 5^{th} order UHF one.

To get the calculated response, component values really need to be close to the design values – getting to within 5% should be good enough, assuming you haven't asked for too critical a filter cutoff or tried to be too clever! Capacitors of this accuracy are readily available: my stock of SMT devices are all 5% tolerance, and devices are available to this tolerance down to less than 1pF. The required values for the filter are made up of several in parallel. Doing this also helps reduce stray inductance and shares out current flow and heating effects at high power. Try to get really close to the calculated values of capacitance, as that makes the next stage easier.

Inductors

These can be a real pain at VHF and UHF. One solution is to use ready moulded ones with adjustable cores, but unless you happen to have good junk box stocks these are a bit expensive nowadays. Air wound ones work well, and can be adjusted by squeezing the turns, but getting the right value from the start is near impossible.

Andy Talbot, G4JNT
andy.g4jnt@gmail.com

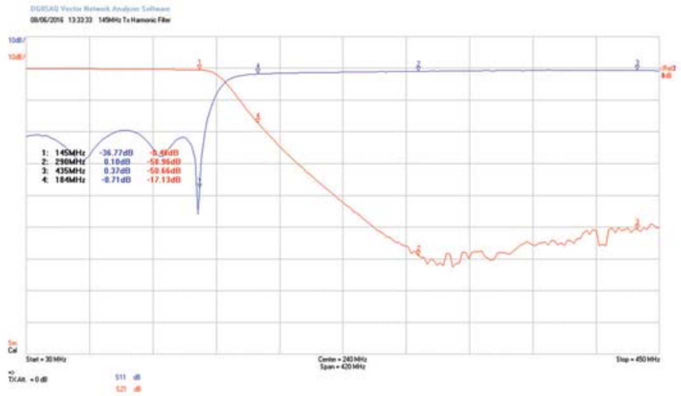

FIGURE 5: Network analyser plot of first breadboard version of Figure 3.

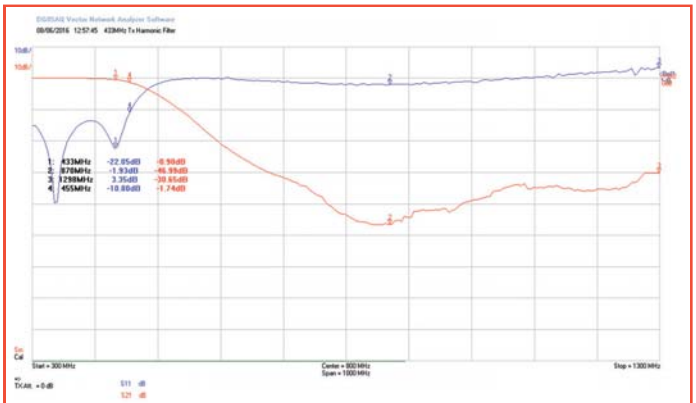

FIGURE 6: Network analyser plot of the first breadboard version of Figure 4.

Wheeler's equation [3] provides a starting point; expect something like 20% accuracy and, for two turns coils at UHF, even worse. But all is not lost: the coils can be adjusted at test. Even if not spot on, the fact that all the capacitors are correct means that over half the components are already at the right value. The VHF coils in Photo 1 were wound using 0.5mm enamelled wire on a 3.1mm drill and the UHF coils with 0.6mm silver plated wire on a 4mm drill.

Tuning up and testing

Rule one about tuning up filters: never, ever, just try to do it based on 'maximum smoke', otherwise known as just tuning for a peak. You'll go down an infinity of wrong ways and never get the right response. Instead, adjust while looking at the return loss of the filter with the output port terminated in a good load. This means only a single port measurement, meaning lower cost, simpler antenna analysers or even a SWR bridge and tuneable transmitter can be used. Of course, a two port network analyser means you can simultaneously monitor transmission loss, but this really doesn't show much extra value. Return loss tells you everything!

Look at Figures 1 to 4 again; see the red return loss curves. Those sharp dips depend on the component values. Since we already know the capacitors are correct, all that remains is to adjust the inductors to get the sharp dips more or less into the right place.

Results

Figures 5 shows the measured results for my breadboard 150MHz 7th order Chebyshev 0.05dB ripple filter. Compare with **Figure 3**. The dips are smoothed out a bit as I used averaging on the network analyser, but they are not too far off the predicted ones after a bit of coil squeezing. The attenuation at 290MHz measures 58dB, remarkably close to the predicted value of ~56dB.

Similarly, **Figure 6** shows the measured results of the 5th order Chebyshev with optimised passband. The plot doesn't look so nice because my DG8SAQ network analyser is reaching the limit of its capability (it is only properly useful up to 500MHz, with degraded performance above that). But it is good enough to tune up with. Since here the main aim is to place the second return loss dip at the band of interest, the first one is further away from predicted than expected. But notice the second harmonic attenuation, at 47dB (it may even be better than that as the network analyser is starting to show its limits) it is several dB better than the predicted value. I have encountered this before and realised it is probably due to unwanted coupling between the coils. Mounting at alternate right angles minimises coupling but, particularly for the UHF design, some will remain. While unwanted coupling is never desirable, in cases like this it *can* work to our advantage. I have rarely seen significantly less passband attenuation is such a filter unless obvious bad layout has compromised it. Good circuit simulation software should even be able to model coupling.

Conclusions & some second thoughts

The filters were designed for 150MHz and 454MHz cutoff, which in each case is only just above the wanted passband. That was too close and means that component losses, which were not modelled, cause a small insertion loss that increases as we approach cutoff – ie just where we want to use it. The VHF filter shows 0.7dB insertion loss and the UHF one 0.9dB. Had they been designed for a slightly higher cutoff, the stopband attenuation at the second harmonic would be degraded, but insertion loss would also be reduced. Some more time should have been spent trying designs for, say 160 or 165MHz cutoff and 470 to 480MHz.

For a low power transmitter such as these, an extra 0.5dB is hardly important, but at the 400 watt level, every fraction of a dB matters. Which brings us onto the final point…

Voltage and current ratings

All filters show slight resonances internal to their operation, which multiples applied voltage. To get an idea of capacitor ratings, first calculate the applied peak voltage at maximum power. $V_{PEAK} = \sqrt{(2 \times \text{power (W)} \times 50\Omega)}$, which is 200V for 400W. In normal usage, a safety factor of 1.5 or 2 would be used, so 400V rated devices are sufficient. But with the possibility of voltage multiplication, a factor of 4 or 5 is safer. So look for 1kV rated components in Tx filters at the 400W level. ATC do a range of porcelain capacitors suited to this application with associated high current ratings. For 1W only about a 20V rating is necessary and the currents will be fairly low – but not negligible.

Make inductors as low loss as possible, using wire capable of carrying the currents expected. Silver plated wire for high powers is good, but enamelled is often adequate.

Errata – correction to July Figure 1

G8ACA spotted an error: there shouldn't have been a short where ringed in the excerpt below.

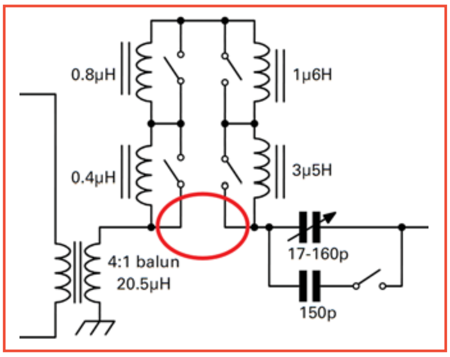

References

[1] *AADE Filter Design* Download – www.dxzone.com/dx17873/aade-filter-design.html

[2] Return loss is simply the reflection coefficient expressed in positive dB terms, and is related to VSWR (V) by the equation RL (dB) = 20.LOG[(V – 1) / (V + 1)] Filter ripple is also related via the equation Ripple (dB) = 10.LOG [2 * (V2 + 1) / (V + 1)2]

[3] Wheeler's equation for inductance: L (µH) = (D.N)2 / (0.46 * D + G), where D = diameter of the coil, G = length (both in mm) & N = number of turns. For large coils with many turns, expect around 10% accuracy. For small coils of only a few turns treat the result as just a starting point and expect an accuracy of 20 – 50% from the calculation.

Design Notes

Novel homebrew CW transceiver

At a recent club meeting Richard Harris, G3OTK gave a presentation on a three band CW (Morse) transceiver he designed and built from scratch. The transceiver has more than a few novel features and Richard put a lot of effort and thought into the design to get something that was easy to use, worked well with strong signals on crowded bands and was optimised for CW working in contests. So it is worth looking at each of his critical design points to see what was done and why. Hopefully the ideas will provide food for thought for others who want to build their own customised radios and other projects.

Transceiver summary

The requirement was for a CW transceiver that covered the 80m, 40m, and 20m bands for CW operation only. For low power portable use 3W output from an efficient PA stage with low power consumption was needed and for ease of operation it had to include an automatic ATU (antenna tuning unit). An electronic keyer was essential and the radio had to be useable with computer control, integrating properly with logging programs like *N1MM+* or *SD*. A block diagram of the transceiver is shown in **Figure 1**.

Frequency control

The radio frequency (RF) source is an AD9850 direct digital synthesiser (DDS); one of the modules popularised through eBay. It is frequency agile, tuning in steps of just over 10Hz. The frequency is changed to act as either a direct RF source on transmit or as the local oscillator (LO) 5MHz away on receive. A 400 steps per revolution optical incremental encoder (obtained from eBay) drives a PICAXE microcontroller to set the DDS registers for the correct frequency output. The frequency controller has a serial computer assisted transceiver (CAT) input port for remote operation that, for ease of integration with logging and transceiver control software, was set to emulate a Kenwood TS-2000.

Another pair of PICAXE controllers drive a 2 line x 16 character organic LED (OLED) display to show frequency and the received signal level via an analogue to digital converter (ADC) from the receiver intermediate frequency (IF).

Receiver

A three pole inductor-capacitor (LC) bandpass filter is used at the antenna input, with one of three units switched in for each band. The Rx is a conventional single conversion superhet with the IF at 5MHz. Two bandpass crystal filters are used with a design bandwidth of 310Hz. A six crystal linear-phase IF filter sits at the input to the IF amplifier. **Figure 2** shows the predicted and achieved filter response. A second, three crystal, linear-phase post-IF filter sits at the output to clean up noise introduced by the amplification and to add extra channel selectivity.

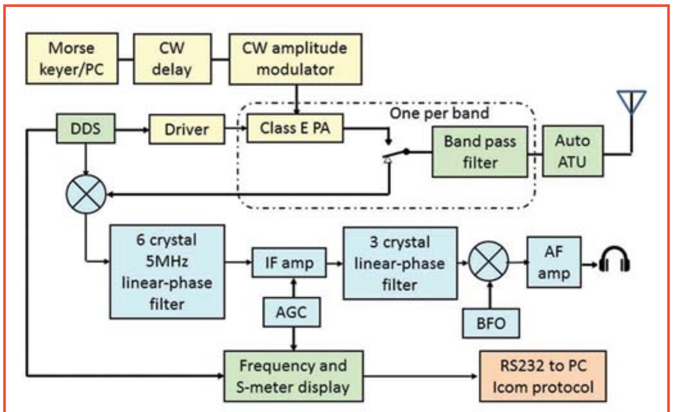

FIGURE 1: Block diagram of the G3OTK CW transceiver.

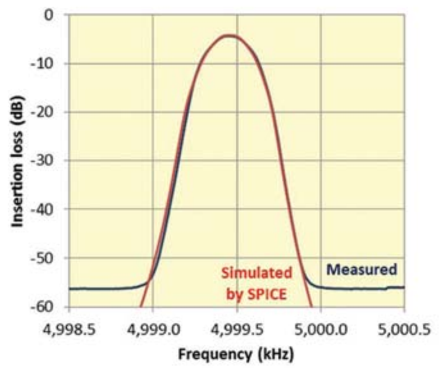

FIGURE 2: 6 crystal predistorted 0.5° linear phase IF filter with Fc = 4,999500Hz and Fbw = 275Hz, as simulated and measured.

The receiver is only for CW reception so ladder crystal filters were designed and optimised for this mode. With pulsed signal waveforms a sharp edged 'brick wall' filter is a poor choice of frequency response. This is because the time delay though the filter is longer at the edges of the frequency response than it is at the centre. The result is ringing and an unpleasant sound. The linear phase filter response is optimised for pulse waveforms and commonly used. Normalised values for this response are tabulated in sets of filter tables, and linear phase is an option in filter design software such as *AADE* and other packages.

The ladder crystal filters were designed from scratch using principles detailed in *RadCom* and other publications over the years in conjunction with Zverev **[1]**. They were all constructed from a bulk purchase of 100 low cost 5MHz crystals from Farnell, with individual units selected using a crystal measuring test set **[2]** to be as close to each other as possible.

Power amplifier

For a long operating life when running from batteries, high efficiency 3 watt Class-E switching power amplifier stages were employed. Class-E operation is a particular type of switch mode circuit configuration where the tank tuning is optimised to ensure the output device switches at the zero current crossing point – thus minimising dissipation during the on-off transitions. Very high efficiency PA stages are possible, approaching 90% if designed and built with care. A couple of Class-E stages for 1.8MHz and 475kHz operation were described in the October 2010 and December 2014 Design Notes respectively. More information on the Class-E amplifiers can be found at **[3]**.

As Class-E amplifiers are inherently narrowband with a medium to high Q tank resonator being essential for their operation, Richard's transceiver used three separate amplifiers, one for each band. **Figure 3** shows the circuit diagram of the 40m one. At just a few pence for the switching MOSFET devices, this proved a cost-effective solution. The output filtering was done by the same filters as used for the receiver input.

CW break-in

Semi-automatic break-in CW operation was considered essential but based on experience some additional thought had to go into the break-in process actually employed on the transceiver. Conventionally, semi break-in, which is similar to voice operated switching (VOX) on SSB, works like this:

- Go to Tx as soon as the Morse key is pressed
- Stay on transmit while keying, with the keying resetting a timer
- 0.4s after the last dot/dash, switch to receive.

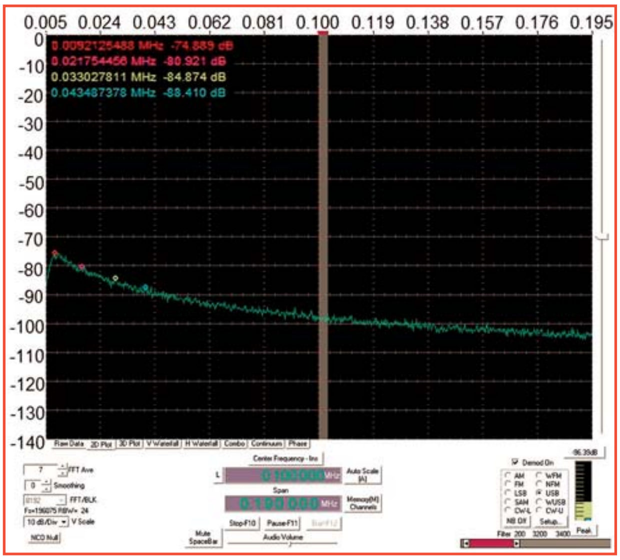

FIGURE 2: Output noise from the LM1117 LDO regulator when decoupled with 100μF on the output and 10μF on the input. Input and output decoupling capacitors have to be increased together to control noise from this device.

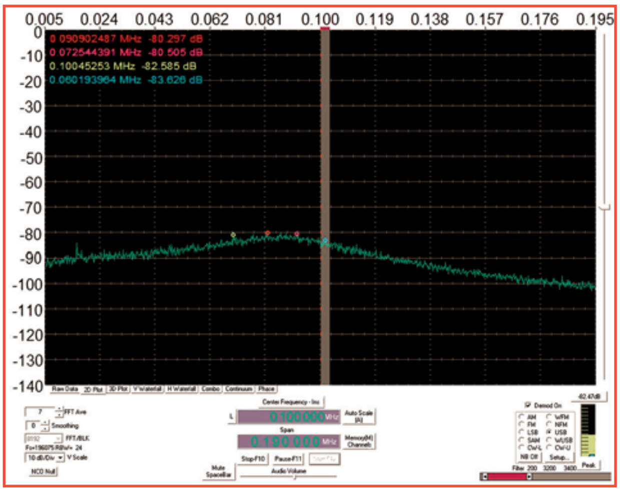

FIGURE 3: LM317 adjustable regulator with the minimum specified values of 1μF on the input and 680nF on the output. The noise peak at 84kHz reduces in frequency as the output decoupling C is increased, but values up to 1000μF or more are needed to get it right down to insignificant levels.

-120dBm at 1.5MHz. 100kHz noise improves by 10dB if the output capacitance is raised to 100μF, but LF noise needs even more to reduce it. *2200μF* was needed to reduce it by a further 10dB in the 10kHz region.

Input noise was measured as -100dBm at 100kHz and could be reduced by increasing the input capacitor to 100μF, as expected. BUT, this noise on the *input* terminal was made worse by any increase in the *output* decoupling capacitor. So both capacitors have to be varied together. **Figure 2** shows the result for 100μF on the output and 10μF on the input. A further 2200μF had to be placed on the input to reduce noise to close to the receiver's measurement limit.

The device oscillated in the LF region with anything less than 1μF on the output and 470nF on the input.

LM317

The datasheet specification says "1μF on the input if spaced far from the PSU filter; 0.1μF on the output is optional to improve transient response with typically 1μF to 1000μF as needed". 0.68μF was used here.

Noise is worst at low frequencies, with a definite peak at 84kHz as shown in **Figure 3**. But the noise does roll off faster with this device, dropping to -110dBm at 300kHz, -120dBm at 500kHz and down to the measurement limit at 1.2MHz. Adding extra output C shifts the LF peak lower: a value of 100μF drops the noise at 100kHz to -110dBm, while 2200μF takes it nearly to the measurement limit.

Although not normally advised, some applications bypass the adjustment pin to ground to reduce noise. Testing this with an extra 0.68μF in that position flattened the low frequency response of Figure 3 at a level of -100dBm over a wide band. Proper noise reduction and decoupling in the LM317 can become quite complex, but one thing that must *never* be done is just to connect a capacitor between output and the adjustment pin. A capacitor here will *definitely* make it go unstable **[3]**. Apart from this scenario, the LM317 could otherwise never be made to go unstable and oscillate.

Conclusions

All three voltage regulators tested generate quite appreciable levels of noise at LF to MF frequencies when using the minimum decoupling specified in the manufacturer's datasheets. It is easy to see why they give their greatest problems with sensitive PLL circuitry and are also regarded badly in the Hi-Fi audio world. Of the two fixed voltage ones, the low dropout LM1117 fared quite a bit worse at LF but did roll off slightly faster as frequency rose into the HF region. The adjustable LM317 with minimum decoupling fared badly at lower frequencies, but rolled off faster above 300kHz. It also cleaned up the best of the lot with extra C added on the output – exactly as the datasheet suggested would happen.

To properly clean up the resulting supplies requires large values of additional capacitors on the output. Some extra series resistance and/or a choke and a large C to form a low pass filter are going to be needed in critical applications.

Noise appearing on the regulator input pin is seen in some devices but should not cause a problem unless the input supply line goes elsewhere to sensitive circuitry, where conducted noise can leak through. This was the issue seen during the loop amplifier tests where the DC input to the LDO was provided 'up the coax'. The noise got into the RF feed to the receiver. A bigger input C and optionally a series inductor improved matters.

Paul, M1CNK, who was one of the experimenters who found the noise on the loop amplifier, commented that "…the main issue, apparently, is with noise coming from the voltage reference. Therefore the better [low noise] regulators have a separate pin for decoupling the voltage reference. That way the capacitance used for decoupling is not inside the control loop. One problem with just adding larger C to the input and output is that these become part of the control loop. Hence by adding too much, you slow down the response of the regulator to noise spikes – so they just pass through".

References

[1] All noise levels were measured in dBm; the power delivered into the 50Ω input of the SDR-IQ in a measurement (FFT) bandwidth of 24Hz. The SDR-IQ has a minimum sensitivity of around -128dBm at this setting but a lot of local computer interference is picked up at frequencies below 100kHz so this sensitivity could only be achieved above 500kHz. Fortunately (or not) the voltage regulators under test exceeded even this pickup level when using the minimum recommended decoupling. The QRM was also spiky, so 'real' noise could be discerned between the spikes.

To convert a measured dBm reading to a voltage, first normalise from the 24Hz measurement bandwidth by subtracting 10.LOG(24) = 14dB from the dBm value. Then calculate the voltage needed across a 50Ω resistor to give this power. -100dBm measured = -114dBm/Hz or 4×10^{-15} watts per Hz. In 50Ω this corresponds to 450nV per root-Hz.

[2] The tantalum capacitor was hand-held in my fingers while it was placed across the regulator terminal pins for testing. I got it the wrong way round. It is at times like that you really appreciate what happens when a tantalum bead capacitor is connected with the wrong polarity!

[3] It is easy to fall into this trap unintentionally, as a capacitor directly across these two pins is often seen with 7805 type devices. So it can be tempting to just place one when designing and laying out PCBs. The story may be apocryphal, but a certain spacecraft manufacturer on the south coast apparently fell into this trap many years ago. It was only when a PSU badly failed EMC testing (fortunately while still at the breadboard stage) that the mistake was found.

Andy Talbot, G4JNT
andy.g4jnt@gmail.com

this is easy to arrange. The core and number of turns is designed based on the lowest frequency of operation and, as frequency rises, at any given voltage input, B falls in direct proportion. Winding length and copper and ferrite losses usually define the upper frequency performance of a ferrite transformer.

Inductors

This is where it gets more interesting. A pure inductor of L henries has a voltage across it defined by the *rate of change* of current. Alternatively, the volt-seconds across it relate to the static current such that V • t = I • L or V = L x (amps per second). It is this property of volts across a coil (for a short period of time) and current though it that allows an inductor to store energy. In a transformer, with its two windings, all energy is taken out from the secondary as it enters the primary so a transformer is not storing energy, only transferring it.

We've already seen earlier that H, the field strength, is related to amps and B to volt-seconds. So the relationship between B and H is the inductance (leaving out all unchanging things like number of turns, area and path length). It's hopefully clear that in an air wound coil B = μ_0 • H and, practical issues aside, any coil wound in air will work whatever the voltage, frequency or current. But not so for cored components. As soon as saturation is approached, that B/H relationship breaks down and inductance falls off. Add to that the fact that μ_r is often temperature dependent as well, and hopefully you can see the futility of attempting to wind inductors on a continuous core. So how do we make a fixed inductor with a value much higher than an air wound one can offer?

Real inductors with a core

We use the core to concentrate the magnetic flux generated from current in a winding so it is contained completely. If the core is continuous, saturation rapidly approaches so we may well get a real inductor with a very high value of L, possibly many henries, but it will be unpredictable, prone to saturation and incapable of carrying DC – totally useless.

If we introduce an air gap into the magnetic circuit so the core is no longer continuous, everything changes. Now all those amp-turns are poured into the small air gap – which may only be a fraction of a millimetre over the area of the core. The resulting B caused by the H is lower and as it is now defined by the gap and not by the ferrite, the linear relationship between B and H will hold. However B must still be below saturation as the ferrite still has to carry it – so the volt turn equation still applies.

Now we revisit the exact meaning of H expressed as amps per metre. What it actually means for this core is that if all the flux is concentrated in the gap, the amp-turns divided by the width of the gap gives H in amps per metre *inside the gap*. This is pure geometry based on the size of the gap and the number of turns on the core with the current though them. The core property is irrelevant, so long as it can be assumed to be carrying all the flux, ie it has very high μ_r so no flux is leaking from other parts of it than at the gap, and it is below saturation. So now we can define our H exactly in a medium (air) that allows it to relate directly to B. If B and H are controlled, so is the inductance. By just adding a small air gap to a magnetic core we've now generated an inductor of a constant value we have control over. By adjusting core area, gap length and number of turns we can select our inductance to be any value we need. And *the ferrite properties have no direct influence on it*.

By a bit of arithmetic manipulation [1], combining the equations for B and H and knowing that inductance, L is given by V • t / I (volts seconds per amp) we eventually end up with the equation for an inductor made from a small gap of width d in a ferrite core of cross sectional area A:

L = μ_0 • N^2A / d

For practical cores μ_0, A and d are lumped together into a value of specific inductance often specified in nanohenries per turn squared.

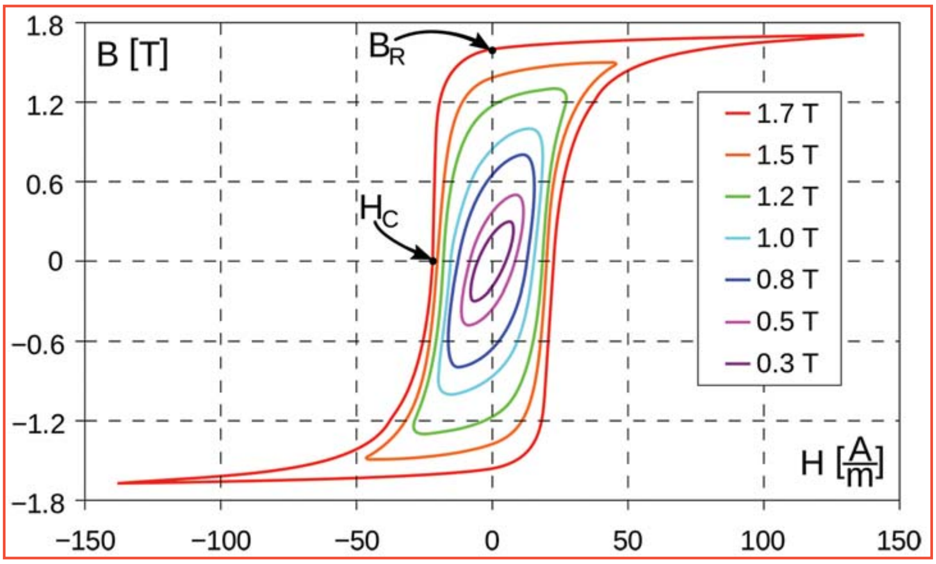

FIGURE 1: A family of B-H curves for grain orientated steel at 50Hz, showing how B does not retrace when H is varied. This results in residual magnetism stored in the material, and energy loss per cycle. *Image: CC BY-SA 3.0, courtesy of 'Zureks', Wikimedia Commons.*

Practical core materials

Magnetic materials, ferrite in particular, are fickle things and many types exist. The main problem is that the B-H relationship is not only non-linear, but does not even retrace itself. So if H is driven upwards then back down, the resulting B does not fall back to the same value. The result is that if H is continuously varied, as in an AC waveform, a plot of the resulting B-H relationship (the curve often seen in data sheets on ferrites), shows as a loop. The area inside this loop corresponds to lost energy per cycle, which is dissipated as heat. So at high frequencies and high B, ferrite materials begin to fail. In transformers and inductors this is mitigated by keeping B as low as possible – consistent with not having too many turns that add copper loss, and choosing a ferrite material that minimises the B-H loss at the frequency range of interest. This is why there are so many different ferrite materials specified for different frequency ranges and purposes. Consider a ferrite material deliberately designed to have a large area inside the B-H curve. It will absorb energy from AC current flowing in its windings and dissipate it as heat. This is exactly what is needed for suppression and EMC protection, and explains why ferrite cores designed for EMC (and often available cheaply at rallies) are virtually useless for RF transformers, baluns and whatever else we'd like to use them for.

Iron dust cores have built in gaps between the particles so they have an 'average' gap that is impossible to define. Which is why they, as well as most gapped pot cores, are specified in terms of inductance per turn-squared rather than by gap dimension.

References

[1] For the mathematically inclined:
Using L = V • t / I (the definition of inductance),
V • t = N • A • B (Faraday's law), B = μ_0 • H
and H = N • I / d
For a core with gap
L = V • t / I = N • A • B / (H • d / N)
= N^2 • A • μ_0 • H / (H • d)
L = N^2 • A • μ_0 / d (turns, gap area and gap spacing only)

Andy Talbot, G4JNT
andy.g4jnt@gmail.com

Design Notes

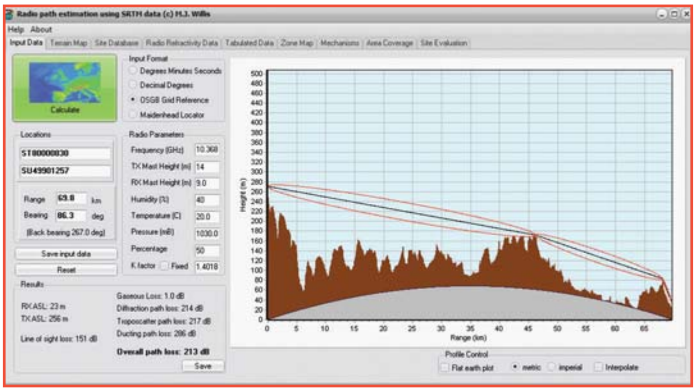

FIGURE 1: A typical plot from G0MJW's Path Profile software. Here showing the path from my home QTH to the GB3SCX Beacon in Dorset on 10GHz. Path loss is calculated using standard RF propagation algorithms for troposcatter and diffraction, taking atmospheric and meteorological effects into account.

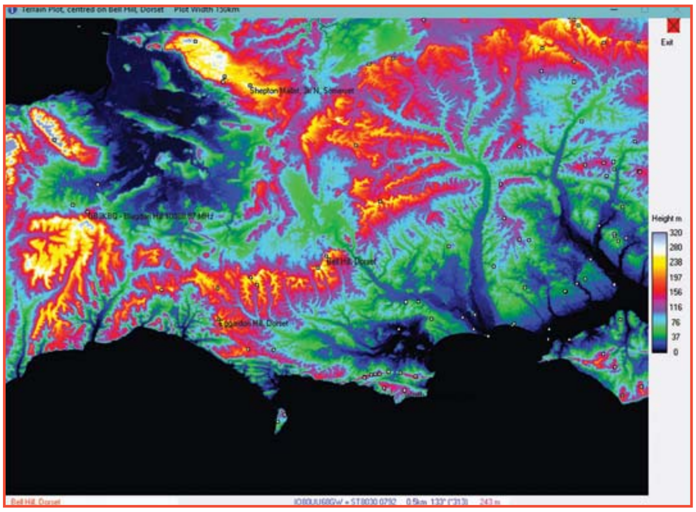

FIGURE 2: A terrain map of the Dorset area centred on the Bell Hill beacon site at IO80UU59NR. Ground height is represented by colour. NGR and Locator corresponding to the cursor position can be read off and the white dots correspond to sites stored in the site database. Double-clicking the mouse names the nearest site.

Where are we, and where are we going?

This column has now been running for a little over ten years, starting off originally as the bi-monthly *Short Circuits* where it alternated with *Data*. When I retired early from full time slavery, otherwise known as employment, back in 2010, there was more time to devote to actually experimenting and playing with projects and writing them up. Often this was with the sole intention of designing and building just for publication as I enjoyed playing with new ideas and hardware and components just for the sake of it.

I had no issues with buying components on a whim, using them, building something up, even if it was a small project that was just of interest at the time and one that I, personally, may have little use for it; someone else would surely be interested. The fun was in the design and build, not the use. Many such PCBs, modules and even boxed items of equipment lie abandoned, or waiting on a shelf, or in cupboards for that opportunity to burst into real usefulness. Or to be passed on to someone who could use them.

But, after ten years, *JNT Labs* is slowing down. There's a limit to the amount of stuff one can just pull out of the junk box, find at rallies, buy from new, play with, test out and write about that is applicable to the amateur radio community and of interest to you, the Members.

My own interests are also moving on, even more than they used to be, towards SDR, DSP and software generally. Can you see where this may be heading? I'm writing this month's column with only a few days to go, and a near complete block on anything of an idea technical or construction-wise to write about. The only hardware project in the current pipeline is the 12V to 42V charger for the lithium battery described last month – an item that will actually be used from time to time, but not essential enough to rush it to completion. That hasn't progressed far enough to be worth writing about, but may well form part of next month's column, unless other inspiration strikes.

So, your input please

Talking to amateurs at events such as round tables and rallies and monitoring several technical internet groups, there is still a fair bit of interesting home construction and design going on. So, to keep this column going as a primarily higher-end technical construction and home brewing forum, could I have your stuff please. What are you building, or doing? Modifying unusual kit, ripping it apart to get useful bits? Using modern components, either as they should be, or in an unconventional manner? Share it, however wacky the ideas may be. Please send anything you want to be included to the address at the bottom. If you want, just send notes or comments and I will put them into a readable form; or if you can, send a full block of text for copy and paste with little editing. It's the ideas and the interesting bits that count.

And now for something completely different

Back in the 1980s, even in the late 1970s, *RadCom* published quite a few articles on software for calculating distance and bearings for contest scoring, pointing antennas and the like. The new IARU world-wide locator system (abbreviated to *Loc*) had just arrived; an alphanumerical sequence such as IO90IV defining any QTH in the World to an accuracy of a few kilometres. These replaced the old QRA, codes like ZK04g that were only applicable to western Europe.

RadCom published quite a lot of software back then as home computers were new and exciting, and many radio amateurs were putting them to use in many areas of the hobby. There was considerable interest in using computers for contest scoring and algorithms were published for distance and bearings between two locators. There was even a very crude approximate conversion between Lat/Long and the National Grid Reference (NGR) as used for UK mapping and defining locations; it wasn't the exact Ordnance Survey conversion, but accuracy was good enough for the purpose at that time.

Moving on to the 1990s, I became interested in the microwave bands where most operating was done from hilltops with highly directional antennas that needed bearings to be calculated accurately. Several amateurs had put-together a database of sites used for microwave operating, good hill tops, some useable coastal spots and a few home stations. The database

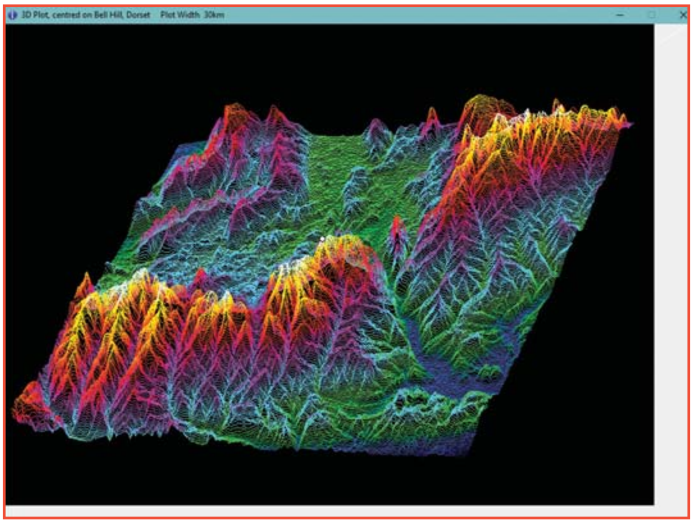

FIGURE 3: 3D representation of ground terrain showing a more localised view around the same site as Figure 2. The pseudo 3D image allows the effects of local obstructions to be visualised.

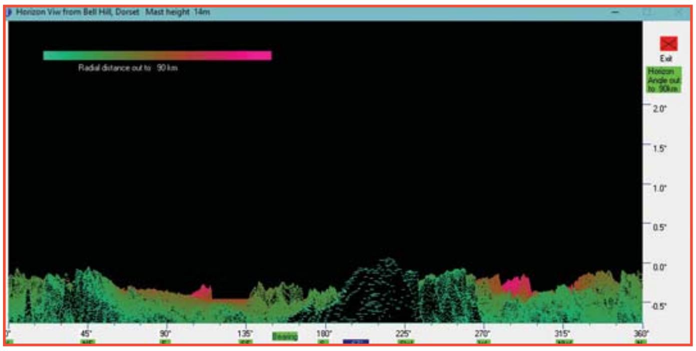

FIGURE 5: This plot shows the horizon angle in a linear form, as if you were up the mast on the reference site and scanned around through 360°. Cursor readout gives the bearing, showing-up narrow slots where DX line of site paths may exist.

FIGURE 4: Plot of visible terrain from the same site. The colour corresponds to horizon angle, with black sections not visible. The cursor position reads out NGR, Locator, and the nearest site, allowing various high spots and line of sight paths to be discovered.

was, back then, used with BBC computers or other ones of the era, or hand calculators for the bearing calculations. I took the original site database and integrated it with locator conversion software to and from Lat/Long. Then added in NGR after finding the exact Ordnance Survey equations for the conversion to and from Lat/Long. That became a programme that printed out a list of distances and bearings from any user defined site, or any values typed in, as either NGR or Locator, or as a site name in the database. It was taken up widely by the microwave operating community. On the back of that I was given a copy of a height database of the UK that gave spot heights for every interval on a 500 metre grid. That was integrated into several programmes that plotted path profiles, terrain maps and even showed a three-dimensional plot of the local terrain. All the code was all written in the 16 bit programming language PowerBasic for DOS, which later-on ran perfectly well in a command screen within the early dialects of Windows. That original suite, *GEOG2*, worked happily up to Windows XP with several operators still using it in a DOS emulator after Windows 7 and then Windows 10 appeared.

Meanwhile, Mike Willis, G0MJW, had taken the idea and developed a full-blown path profile analysis tool for Windows with RF propagation calculations built in [1]. He used a higher resolution height database generated by the Shuttle Radar Terrain Mapper (SRTM). This is available for anyone to download for free and covers the whole world [2]. Individual 'tiles' of one degree latitude and one degree longitude are stored with file names such as N50W002.HGT (50° North, 2° West). There are an awful lot of them covering the World's landmass, so download only those needed. The spatial resolution is 3 arc-seconds, or 1/1200 of a degree which at mid latitudes corresponds to a squaroid [3] of about 93 metres – six times better than the original UK-only height database. A typical plot from Mike's software is that shown in **Figure 1**, which is the path profile from my house to Bell Hill in Dorset, the site of the South Coast Microwave beacons. Parameters are set for 10GHz propagation.

Meanwhile there had been a global shift in location accuracy available to the masses that changed everything. GPS (now GNSS) had come along and anyone could get their Lat/Long anywhere on Earth to a few metres accuracy. Any modern smartphone can do it, and the display apps are mostly free. One effect of GPS taking over was that a global terrain reference frame, or spheroid, referred to as WGS84 had to be used for defining Lat/Long instead of local mappings for each country or region – our original one was OSGB36. This meant the old UK Lat/Long, locked to the national grid from surveys dating back to the eighteenth century, was now in error more than hundred metres when compared with GPS measurements. It wasn't 'wrong', just a different, global, reference. The OS developed conversion algorithms and a WGS84 / NGR conversion accurate to about three metres with stand-alone software was included by Mike in his *Profile* package. The full high accuracy OS conversion is accurate to millimetres, but involves a lookup correction table based on 1km squares – unnecessary for our purposes. The OSGB36 / WGS84 error can be seen by comparing old 1980s generation OS maps in the 1:50000 series with their modern equivalents. The blue crosses indicating 5' intervals that are shown where they do not obstruct other detail have moved by up to a couple of millimetres.

After an idle period earlier this year I decided it really was time to update the earlier software that gave the nice plots that Mike's didn't go in for. Using a 32 bit version of PowerBasic, several of the earlier *GEOG2* suite were converted to use the SRTM data and run properly in a Windows environment. And, of course, use the WGS84 to NGR conversion. Full details can be found at [4] including more on the various programmes, and on mapping and conversions. Some of the plots typical of those generated can be seen in **Figures 1** to **5**. I also discovered the somewhat geeky joy of, while out hill walking, placing my iPhone on a Trig point, leaving it for several minutes to acquire and fix properly then comparing the displayed 1m NGR (with its 'declared' 5m accuracy) against the 'true' value for that trig point. A list of all 60000 of the old OS survey points, many of them now unfortunately lost or destroyed can be downloaded from the OS website; the NGR of each is given to one-metre accuracy. The iPhone GPS readout of NGR has never been more than a few metres out, with a couple showing, at least at the time of measurement, the exact value listed in the OS database. Very satisfying.

Andy Talbot, G4JNT
andy.g4jnt@gmail.com

Design Notes

Gwynn Griffiths, G3ZIL has designed an adaptive noise cancelling system for HF reception, based around a Raspberry Pi. He sent in the following write up.

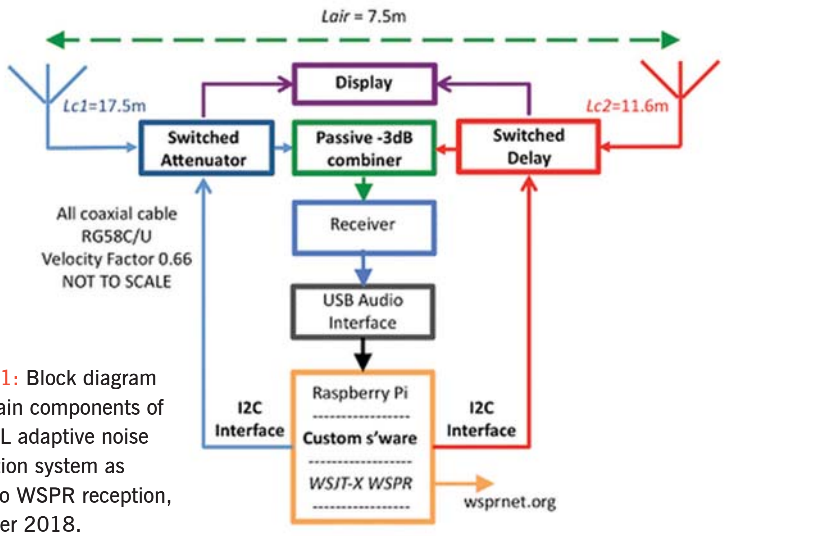

FIGURE 1: Block diagram of the main components of the G3ZIL adaptive noise cancellation system as applied to WSPR reception, September 2018.

Adaptive noise canceller

Noise cancellation prior to the antenna socket of a receiver is a well-established technique with several commercial devices available including the MFJ-1026, the NCC-2 from DX Engineering and the ANC-4 from Timewave. This method needs an auxiliary antenna, a way to match the noise amplitudes from the auxiliary and main antennas and alter the phase so that at a combiner the two noise components are 180° out of phase. I've described a manual passive noise canceller using a switched attenuator and tapped delay line [1]. Results were encouraging, with noise reduction of 12dB possible on 7MHz. However, the optimum settings changed throughout the day, and from day to day. Those changes provided the motivation for this adaptive version currently configured to be compatible with WSPR reception. The key features are:

- Fully adaptive using closed-loop feedback control of delay & attenuator settings based on measured noise in a defined audio band
- Usable from 80–10m (40–15m in prototype)
- Delay and attenuator step size to give over 20dB noise reduction for a single source of small spatial extent
- Front panel display of delay and attenuator settings
- Compact aluminium case, approximately 160x210x60mm
- User interface with 3D display of noise with delay and attenuation, time-stamped .CSV log file for diagnostics including noise levels, noise notch depth and width.

System overview

Figure 1 shows the main components of the adaptive noise cancellation system. Both antennas are shallow inverted V dipoles.

The motivation is to improve SNR for WSPR reception, hence the output from the receiver in [2] is fed to a USB audio interface connected to a Raspberry Pi 3. The Pi runs standard WSJT-X software and a custom set of Python and Bash shell scripts to implement adaptive noise cancellation.

The relay-switched attenuator and delay lines are configured with powers-of-two steps: 0–7dB in 1dB steps, and 0–77.5ns in 2.5ns steps. Delay and attenuation values are set over the I^2C bus from the Pi via I^2C to 8 bit I/O expanders. Front panel LED bargraph displays show the current settings in binary.

WSPR transmissions last 110.6s, starting 1s after an even minute within a two-minute frame, making available a 9.4s gap with no transmissions (in theory) within which noise measurements can be made. Currently, it takes the gaps over three two-minute WSPR frames to step through all the delay and attenuator settings.

Each measurement cycle comprises setting the relays, waiting 6ms for contacts to settle, acquiring 200ms of audio, an FFT analysis to find the noise amplitude within the 1400–1600Hz band and then discard of the top 20% amplitudes (a transmission may have over-run or started early, or there may be a non-WSPR signal present). Once a set of noise levels at the complete set of delays and attenuator settings has been gathered (currently every six minutes), the filtered noise profiles are updated, the optimum delay and attenuator settings found, ready to be applied for the next three acquisition cycles.

The current filtered delay and attenuation profiles are stored as a .CSV file and a Python program can plot a 3D representation, from which the noise level and the depth of cancellation can be seen as shown in **Figure 2**. With experience of the local noise environment, one can infer direction from the delay setting, and the spatial extent of the noise source(s) from the width and depth of the cancellation notch.

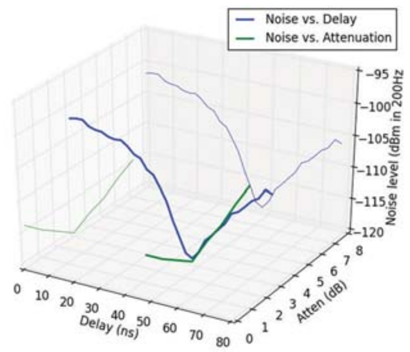

FIGURE 2: 3D representation of the noise levels with delay and attenuator settings. Here there was a single, spatially narrow noise source.

While there is some interaction between the delay and attenuation settings, eg with coaxial cable delay lines longer delays mean greater attenuation in that path, the intersection of the minima for the delay and attenuator settings in Figure 2 is a very good combined minimum.

A notch width of 10ns or less between 3dB points is often seen at G3ZIL's location on 40m, implying a delay step size of 5ns (about 12° in phase) or less is needed to deliver the 20dB noise cancellation that is possible. On this basis, a step size of 1.25ns should be fine for operation to 10m.

Daily noise environment and noise canceller data

Filtered noise level profiles are appended to a .CSV file suitable for post-processing (eg in Excel, as used for the examples here). **Figure 3** shows the wealth of useful information from the data logged. Interpretation comes with experience of using the system, and each location may well be different, but the caption shows what can be gleaned from a day of data.

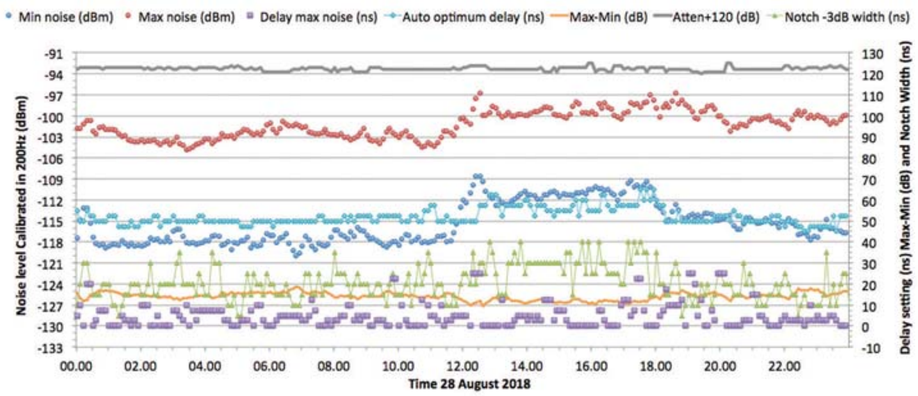

FIGURE 3: Example of a very good day's noise record: until ~1130 noise level after cancellation (blue) is as low as has been achieved so far at around -118dBm in 200Hz, the noise direction is steady, inferred from the delay setting (cyan), and the noise cancellation (orange) averages 14.9dB. After 1130 the pre-cancellation noise level is higher (red), as the noise level rises from 1130-1230 the notch width (green) increases, implying a greater angular extent, and the notch depth (orange) decreases, hence the post-cancellation noise (cyan) increases by more than the noise level has increased. While the noise level (red) is much the same from 1730–2000, there is a narrowing of the angular distribution after 1730 (green) resulting in a deeper notch (orange) and a ~6dB reduction in post-cancellation noise.

The longer-term noise environment

Figure 4 shows another form of presentation, a form of contour plot of noise level with time and delay setting over a nine-day period. This shows a regular, approximately daily, pattern to the minimum noise obtainable and periods of high noise that have a relatively sudden onset and end. The wider spectrum available on an RSP1 SDR shows many of these high noise events to be slowly drifting interference with bandwidths of tens of kHz. Only on few occasions does the noise canceller fail to reduce the noise to less than -110dBm in 200Hz. It is intriguing that the most extensive period of low noise in this record, 26–27 August, was during and immediately following a G3 geomagnetic storm. It is possible that band noise was lower during that period.

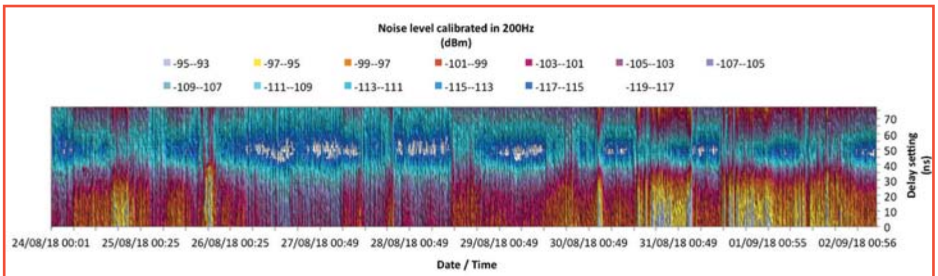

FIGURE 4: Nine days of noise level variation with delay setting shown as a contour chart. There is a regular, approximately daily, pattern to the noise minima around 50ns and the maximum around 0-10ns. Note that there was a G3 geomagnetic storm on 26 August, with Kp at 7 between 0300–0900.

Using the system to drive reductions in local noise level

With the consistent measurement capability of this system, **Figure 5** shows the improvements in noise cancellation and average position in the 40m WSPR Challenge **[3]** for the following changes:

- Until 27 July (blue), manual noise cancellation with fixed delay and attenuator settings, 11m separation between the two antennas, first antenna 4m from the house.
- 29 July–3 August (red), same antenna arrangement, but now with adaptive noise cancellation.
- 4–23 August, adaptive noise cancellation but antenna spacing 7.5m, first antenna 7.5m from the house. Increased noise reduction probably due to smaller angular extent as first antenna further from house.
- 24 August–2 September, as previous but with my 60m dipole and its feeder into the house taken down. Improved noise reduction due to less noise originating from the house being coupled to the outside.

Coax cable at low frequencies

A recent post on the LF Group highlighted a problem with using low cost feeder on the LF and MF bands. One poster used low cost satellite TV coax to feed LF loop antenna, and encountered much higher loss than expected. It turned out that the coax centre conductor was copper plated steel instead of pure copper. Steel offers the advantage of greater strength and lower cost, and possibly works better when forming the centre pin of F-type connectors.

Over the 950MHz to 2GHz frequency band the cable is designed for, the skin depth of the copper plating, which may only be a few microns in depth, is enough to contain all the RF power. But at LF, where the skin depth extends to hundreds of microns, most of the current is flowing in the lossy steel part of the centre conductor.

Skin Depth

Skin depth is defined as the distance into a conductor at which the current has fallen to 1/e, or about 37% of that at the surface. It can be calculated from:

$$d = 503\sqrt{\frac{\rho}{F \times \mu_r}}$$

where **d** is the resistivity ($1.7 \times 10^{-8}\Omega$m for copper, typically 1 to 10×10^{-7} for steel), F is in Hz and μ_r is the relative permeability (1 for copper and several hundred for magnetic materials like iron or steel).

For copper, **d** at 137kHz is 0.18mm (180 microns, μm); at 950MHz it works out to be just 2.2μm so it is clear that a few microns of copper plating is sufficient for the intended purpose. However, at LF very nearly all the current is in the steel part of the centre conductor.

Satellite TV feeder is not the only cable to use a plated centre. RG174, the thin stuff we often use for interconnects and as

Andy Talbot, G4JNT
andy.g4jnt@gmail.com

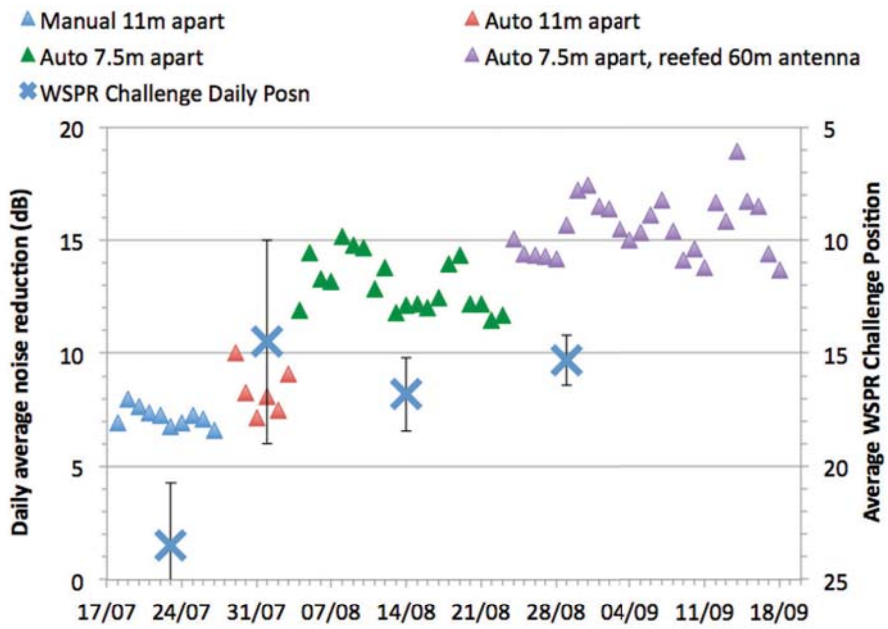

FIGURE 5: Daily average noise reduction for four configurations with average WSPR challenge positions for each configuration with standard error bars. Dates are in 2018.

flexible BNC and SMA leads is often made this way. High spec microwave cables with PTFE insulation can use silver plated steel. These are all things to watch out for if using these cables outside their intended frequency range.

As an aside, if you calculate the skin depth for the 50Hz mains frequency in copper or aluminium (d = 2.7x10^{-8}), you can demonstrate why you never see power cables with a conductor more than about 18mm in diameter. Then do the same exercise for 50Hz in steel (with a magnetic permeability of, say, 1000) and think of the issues supplying electric trains. Some rail systems use 16⅔Hz supplies for this reason.

References

[1] Griffiths, G , 2 018. H F b and p assive noise cancellation, *Practical Wireless*, 94(9): 12-15.
[2] Griffiths, G, 2016. A Direct Conversion WSPR Receiver for 30, 40 or 60 metres, *Practical Wireless* 92(4): 30-35.
[3] WSPR Challenge: http://wspr.pe1itr.com/

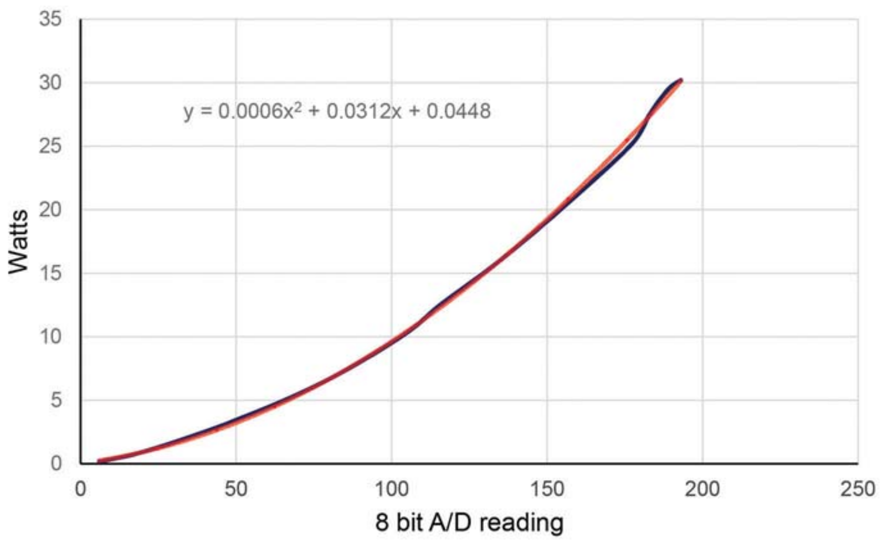

FIGURE 2: The black line is a plot of a diode detector response showing 8 bit A/D value (with a 5V reference) corresponding to detector voltage output versus RF power. The red line is a 2nd order trendline whose coefficients are used to calculate precise intermediate values from these measured results.

For any A/D output value, X, the resulting RF power is given by plugging that value into the equation.

There are now two routes to using the results in a processor. If you are using an Arduino or Raspberry Pi or something with support for floating point arithmetic, it may be easiest to do the calculation in the controller directly. Even a PIC or small controller with integer maths would work, so long as the fractional coefficient values are pre-multiplied by some constant to allow the integer maths to offer sufficient accuracy.

In a PIC, lookup tables are simplest to work with, so I just used the polynomial coefficients in an equation in the spreadsheet to generate formatted values of P_{OUT} for each possible A/D reading; about 200 table entries. Since a lookup table can contain any output format of choice, I let Excel generate the BCD equivalent ready for direct display.

Now all the PIC in the transmitter controller has to do is read the voltage from the A/D, use this as an input to the lookup table and display the resulting BCD value on the LCD. It was complicated slightly by having two formats in the table, ie decimal watts up to 9.9W then units from 10 to 30W, but is still simpler and faster than direct calculation and BCD conversion. Anyone that wants to study the process in more detail can download the spreadsheet from **[3]**.

An alternative to copying the coefficients from the chart is to use the LINEST function that does the same as the trendline but with numbers in cells. I did do an alternative spreadsheet using this method, but it's not the easiest function to remember how to drive; just copying the coefficients is easier when low-order polynomials are used.

Modifying microwave coaxial relays

John, G0API sent in his findings about modifying a microwave relay.

"I have just completed work on a mod to a standard 24V DC, SMA-ported relay, shown in **Photo 3** and **Photo 4**. This model requires DC to be maintained in the operated state, reverting to the rest state when the supply is removed.

"The relay was recovered from some ex-test gear and had probably never carried anything but low level RF. It was lossy on the N/O port when tested at 3.4GHz, so in standard amateur fashion I clamped the SMA ports block into a vice and gave the case a few taps with a large chunk of wood, to ease removal. The case seemed to be held with a film of superglue and the shock did the trick.

"Slowly withdrawing the cover exposed the workings. Studying the arrangement, it was clear that four captive screws needed to be removed to open it all up – there are some small bits in there but as long as you note where they came from it can all be put back again. I found some crud on the N/O and N/C contact faces and used a soft plastic wedge to remove it back to the original face metallisation. Looking at Photo 3, you can see the see-saw contact bar and just above that a central black block and two white-wrapped coils. The switch action works by use of the black block, which is a permanent magnet and magnetic pole assemblies energised by the coils. On closer inspection it turns out that the coils are wired in SERIES, with a combined static resistance of approximately 350Ω.

"So, I thought, if I cut the link wire between the coils and took new lead extensions to the appropriate plus and minus terminal wires, would it work at 12V DC?

"And it does! But there is more. Notice in Photo 4 that there are two holes either side of the centre line and at the upper ends of the coils – these are adjustable magnetic flux cores, which can be set to project below the coil structures. By careful adjustment at nominal 12V DC, it is possible to make the switch action into a latching state type. The switch is activated with supply and the supply reversed to unlatch it.

"Use an adjustable, current settable PSU and make sure the relay works reliably at 12V. Check the minimum reliable pull-in voltage; if just at 12V, adjust the right hand core screw very slowly. On my sample, pull-in can be achieved at 9.3V and the current draw at 12V is 0.14A – I ran it for 30 minutes energised and it was mildly warm but functioned fine."

Master index

A brief index and summary of all editions of this Design Notes column, and the earlier Short Circuits, can be found at **[4]**. The equivalent for the Data column is at **[5]**.

References

[1] ww2.minicircuits.com/

[2] Image rejection as a function of phase error (departure from 90° phase shift) is given by A = 20.LOG(TAN(phase error)). So 2° error results in 20.LOG(TAN(2)) = -29dB. Image rejection due to amplitude error is given by A = 20.LOG((G + 1) / (G - 1)), where G is the linear amplitude error between I and Q channels, equal to $10^{(Ampl / 20)}$. So 0.5dB amplitude imbalance gives G = 1.059 and A = 31dB.
A discrete network can incorporate adjustments to optimise rejection; a commercial package needs to be selected and there is little scope for trimming.

[3] www.g4jnt.com/MikomPA_PowerMon.xls – spreadsheet illustrating the use of Excel's 'Trendline' curve fitting to generate a lookup table for an accurate power monitor facility from a diode detector.

[4] http://g4jnt.com/DesignNotesIndex.pdf

[5] http://g4jnt.com/DataColIndex.pdf