RADIO SOCIETY OF GREAT BRITAIN
ADVANCING AMATEUR RADIO SINCE 1913

RadCom plus

SPRING 2024
VOLUME 09 ◆ Number 01 ◆ £4.95

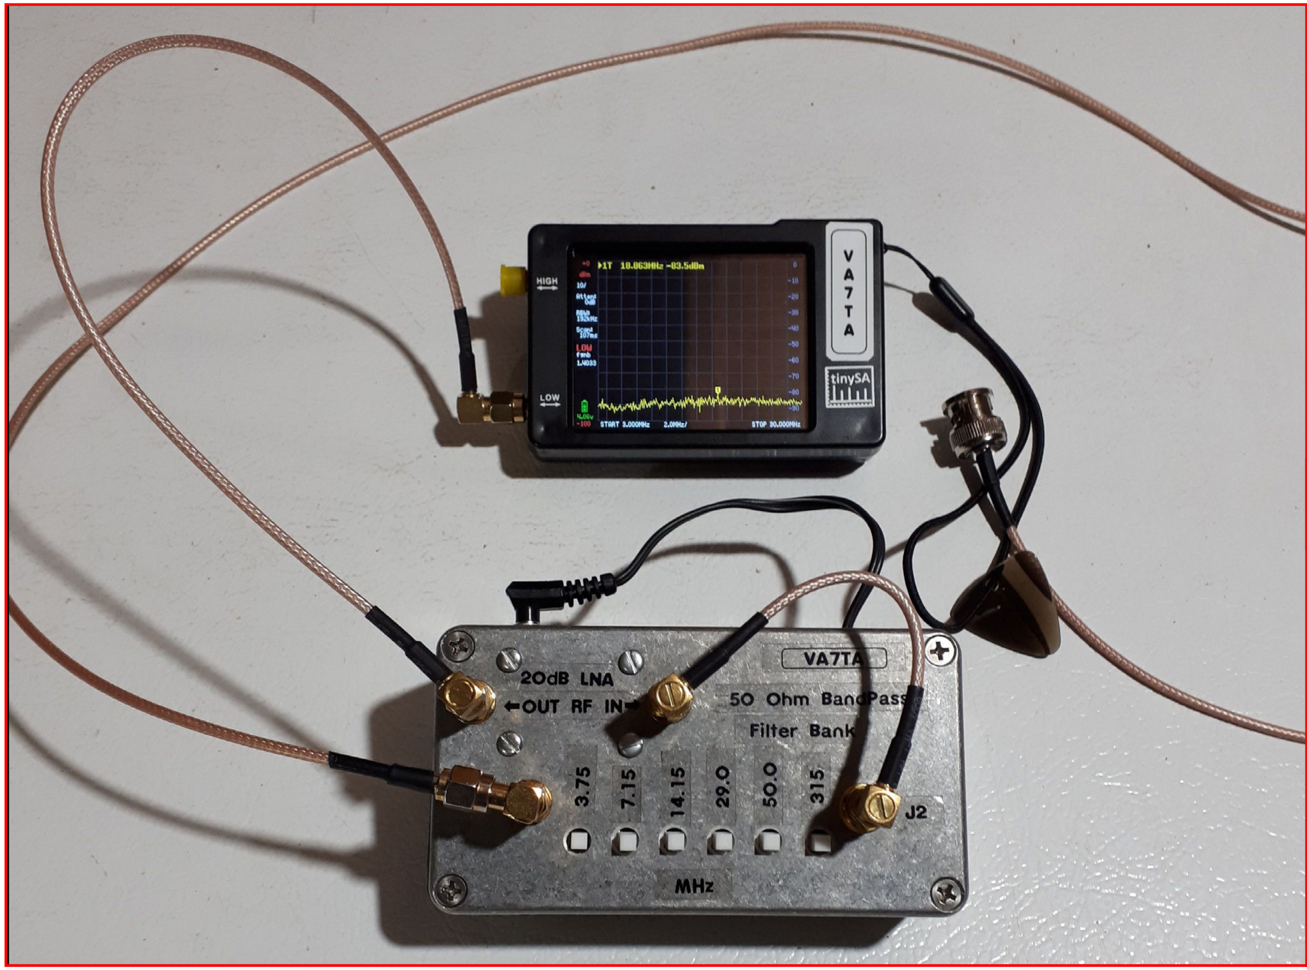

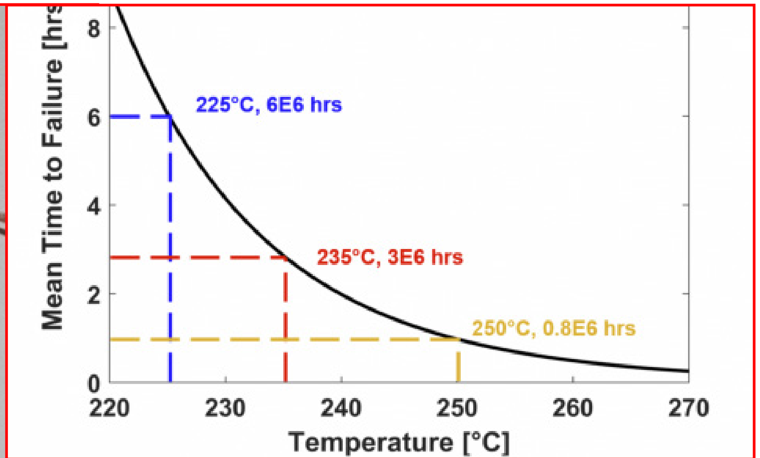

RadCom plus

THE RADIO SOCIETY OF GREAT BRITAIN'S MEMBERS' MAGAZINE

EDITOR:
Peter Duffett-Smith, GM3XJE
E-mail pjds@mrao.cam.ac.uk

MANAGING EDITOR:
Edward O'Neill, M0TZX
Email: Edward.ONeill@rsgb.org.uk

All contributions and correspondence concerning the content of RadCom Plus should be sent to: The Editor, RadCom, 3 Abbey Court, Fraser Road, Priory Business Park, Bedford MK44 3WH Telephone 01234 832 700 Facsimile 01234 831 496, e-mail radcom@rsgb.org.uk

RadCom Plus is published by the Radio Society of Great Britain as an addition to its official journal, *RadCom*, and is avaiable free to all Members of the Society via the RSGB website at www.rsgb.org

All material in *RadCom Plus* is subject to editing for length, clarity, style, punctuation, grammar, legality and taste. No responsibility can be assumed for the return of unsolicited material (if in doubt, call us first!)

Welcome to the Spring 2024 edition of *RadCom Plus*, the online journal for the technically-minded radio amateur.

We have four articles for you in this edition. The first article, by Tom Alldread, VA7TA, is about how to make a protection filter for your expensive and sensitive SDR or VNA. The trouble is that you can't be sure that there are no very strong signals present within the broad bandwidth of the input circuits which could drive them into overload or, worse still, damage them. Tom's filters prevent this happening. Next we have a contribution from Sheldon Hutchison, N6JJA, who discusses in great detail how to make a really good antenna noise canceller. You may recall the article by Andy, G0FVI, in October's RadCom which discussed the principles of antenna noise cancellation. Sheldon investigates exactly how the canceller should work, and his investigations have resulted in his superior design which you can construct for yourself. Our third article, by Michael J. Toia, K3MT, is all about understanding the relationship between a train of square pulses and the harmonic content of its frequency spectrum, not with complicated mathematics, but by using an entirely diagrammatic approach. Those less familiar with mathematics might especially find this article helpful. Finally Andy Nehan, G4HUE, discusses in great detail the design of analogue power supplies. You will certainly want to think again about using those old electrolytic capacitors in the bottom of your junk box after reading what Andy has to say!

As always, I am on the lookout for articles to go in the next edition of RadCom Plus, due to appear during the summer of 2024. Would you like to contribute? Your article need not be very long, but the online format does allow for more extensive contributions than would be possible in the printed pages of RadCom, with a lot more pictures and diagrams. If you want to discuss this, why not drop me a line at the email address shown at the top?

Peter Duffett-Smith GM3XJE

Technical Features

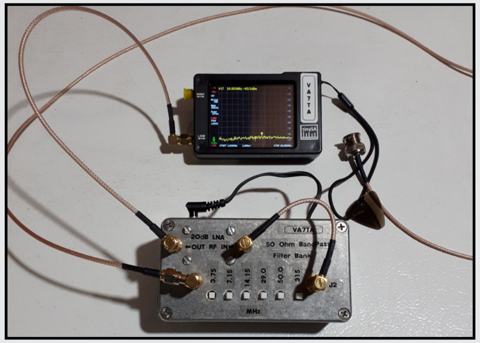

COVER IMAGE 1: Excessive RF level protection filters

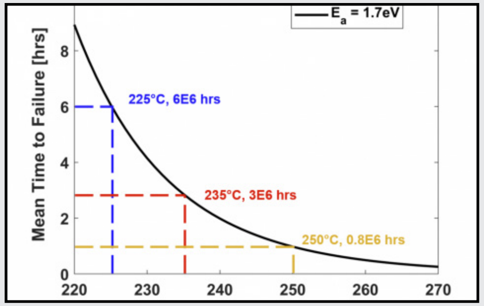

COVER IMAGE 2: The design of linear power supplies

COVER IMAGE 3: Building and testing antenna noise cancellers

RSGB BOOKSHOP

Always the best Amateur Radio books

Introduction to RF Circuit Design for Communication Systems

By Roger C. Palmer, VE7AP

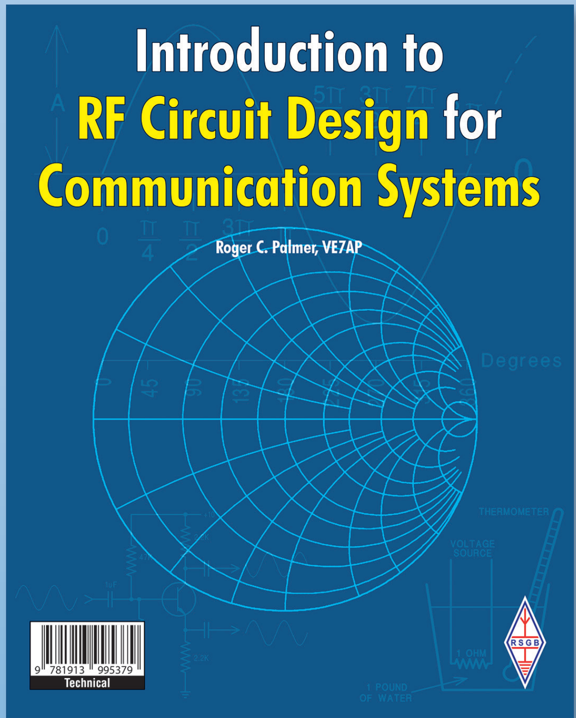

If you have ever wondered how your transceiver actually works, then his book provides a real insight into exactly that. It lays out in an easy-to-understand way the techniques that are commonly used in the design of modern RF communications equipment.

Introduction to RF Circuit Design for Communication Systems is not intended to be a formal textbook with rigorous explanations, derivations, and difficult mathematics. What it is though, is a book for a reader has a general understanding of basic electronic concepts and is looking to understand more of how circuits work, with just enough detail so that designs can be analysed in a basic manner. Where appropriate, approximations and 'rules of thumb' will be disclosed that can often simplify the design process. The book includes several design examples. Although the emphasis is on equipment or circuits that are part of communication systems, you will find information is also provided on a variety of general electronic design topics. Working through from 'plug to antenna' this book is a mine of revealing information that you may never thought about until reading it.

Introduction to RF Circuit Design for Communication Systems is a great read that starts at the level of reader familiar with Ohms Law and reactance. The information included is invaluable to anyone wanting more or a simply great reference book.

Size 203x254mm, 208 pages,
ISBN: 9781 9139 9537 9
Non Member's Price £19.99
RSGB Member's Price £16.99

RSGB Handbook of Radio Communication - 15th Edition

Edited By Ed Durrant, G8GLM/DD5LP

Biggest Update in Years!

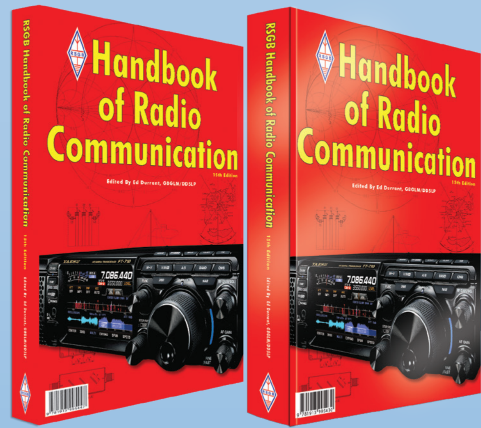

The RSGB Handbook of Radio Communication has now undergone its biggest update in years with this new 15th edition. New editor Ed Durrant, G8GLM/DD5LP has had every one of the 25 chapters revised to bring it bang up to date – which makes the new *RSGB Handbook of Radio Communication* the absolute must have for every shack.

All the important data is still included in the *RSGB Handbook of Radio Communication* but you will also find many new additions including the brand new "Islander-3" receiver building instructions. Some chapters such as 'The Great Outdoors' have been completely rewritten and now provides a concise, up-to-date chapter covering current options and helping those new to 'Portable Ops' get a great start in this expanding part of the hobby. Whilst chapters like 'Transmission Lines' now goes much deeper into the theory than ever before.

New options

As a new alternative, if you are only interested in specific areas of information contained in the *RSGB Handbook of Radio Communication* you can now chose to download individual chapters from the RSGB online shop as 'mini books'. With 25 chapters to choose from you are sure to find just what you are looking for. Simply visit www.rsgbshop.org for these copies.

With 752 pages the *RSGB Handbook of Radio Communication* remains the ultimate reference manual to be read and re-read as we meet challenges in construction and understanding of technologies, old and new.

The RSGB Handbook of Radio Communication is available in both a durable hardback format and a more affordable paperback version – can you afford not to own your own copy?

Paperback
210x297mm, 752 pages,
ISBN: 9781 9139 9544 7
Non Members' Price: £39.99
RSGB Members' Price: £33.99

Hardback
210x297mm, 752 pages,
ISBN: 9781 9139 9543 0
Non Members' Price: £49.99
RSGB Members' Price: £39.99

Don't forget RSGB Members always get a discount

Radio Society of Great Britain **www.rsgbshop.org**
3 Abbey Court, Priory Business Park, Bedford, MK44 3WH. Tel: 01234 832 700

FROM FREE P&P on orders over £30. SeeT&Cs

Excessive RF level protection filters

Tom Alldread, VA7TA
tma.7ta@gmail.com

INTRODUCTION

A piece of RF equipment with sensitive multi-octave broad-bandwidth input circuits, that is directly connected to an antenna, is exposed to the possibility of harmful interference from strong signals that are present within the local RF environment. Spectrum analyzers, and also those wide-spectrum-coverage receivers designed without effective RF input band-pass filters, are particularly vulnerable. In many situations, band-pass input filtering is needed to limit the RF spectrum bandwidth in order to block interference from strong signals such as those from local AM/FM/TV broadcast stations. Such high-level signals can drive sensitive input circuits into regions of non-linear response which, in turn, results in the internal generation of intermodulation noise that can flood broad segments of the frequency spectrum. The associated rise in the RF noise floor can, in many cases, mask reception, or prevent the accurate measurement, of weaker signals of interest across the HF spectrum. In the worst case, a strong signal from a nearby transmitter can be sufficiently high in level to damage the input circuit.

Highly-sensitive portable handheld RF equipment is particularly vulnerable to high-level signals. During typical mobile or portable situations, the local RF environment is often unknown and changing. High-level signals at frequencies beyond the currently-displayed frequency spectrum, and thus beyond the operator's viewing range, can lead to unfortunate results. As an example, consider a handheld spectrum analyzer configured to display the noise-floor level within a portion of an HF amateur radio band. The presence of strong, but out-of-view, signals at frequencies beyond the displayed spectrum edges can easily go unnoticed, yet can cause an erroneous measurement or, in the worst case, result in damage to the input circuits. The keying up of a VHF/UHF handheld transceiver, in close proximity to an HF EMI sense antenna, could over-drive, or possibly even damage, the sensitive RF input circuit of a broad-bandwidth spectrum analyzer rated for a maximum input level of just a few mW.

My recently-acquired TinySA Ultra handheld spectrum analyzer, shown at the top of **Figure 1**, has the potential to be a very useful instrument for tracking down EMI noise sources. It has a broad input bandwidth range of 100kHz to 6GHz, and is very sensitive. With its built-in low-noise preamplifier (LNA) activated, it is capable of displaying signals at levels below 10nV. As with other highly-sensitive instruments, it has a maximum non-destructive aggregate signal input power-level specification of just a few mW, specifically in this case +6dBm (4mW).

Considering my TinySA Ultra's vulnerability to overload-level signals that may be unexpectedly encountered during field measurements, I decided it would be prudent to have the means to limit the input bandwidth to the HF frequency spectrum. I therefore designed the small-form-factor protection filters shown in Figure 1 and described below. The multiple filter stages can be individually selected and, if desired, can all be bypassed to provide unrestricted access to the RF-spectrum. For the support of electrical power-related EMI noise-source tracing activities, a 315MHz surface-acoustic-wave (SAW) band-pass filter section was also provided that could be switched into the RF path to limit the pass band to roughly 312MHz to 318MHz. The relatively easy-to-construct highpass/lowpass (HP/LP) filter version shown in **Figure 2** and described below, includes a 10dB flat-response resistive-attenuator section that can also be individually selected. It provides a convenient means for shifting signal levels to check for the possible internal generation of intermodulation products. DC blocking and static discharge protection are included in all the designs. An

FIGURE 1: My TinySA Ultra handheld spectrum analyzer (top) and three overload-protection filter alternatives: passive high or low pass, band pass, and with or without an LNA for a sensitivity boost.

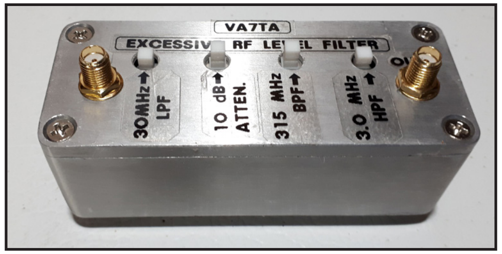

FIGURE 2: The easier-to-build HP/LP filter version.

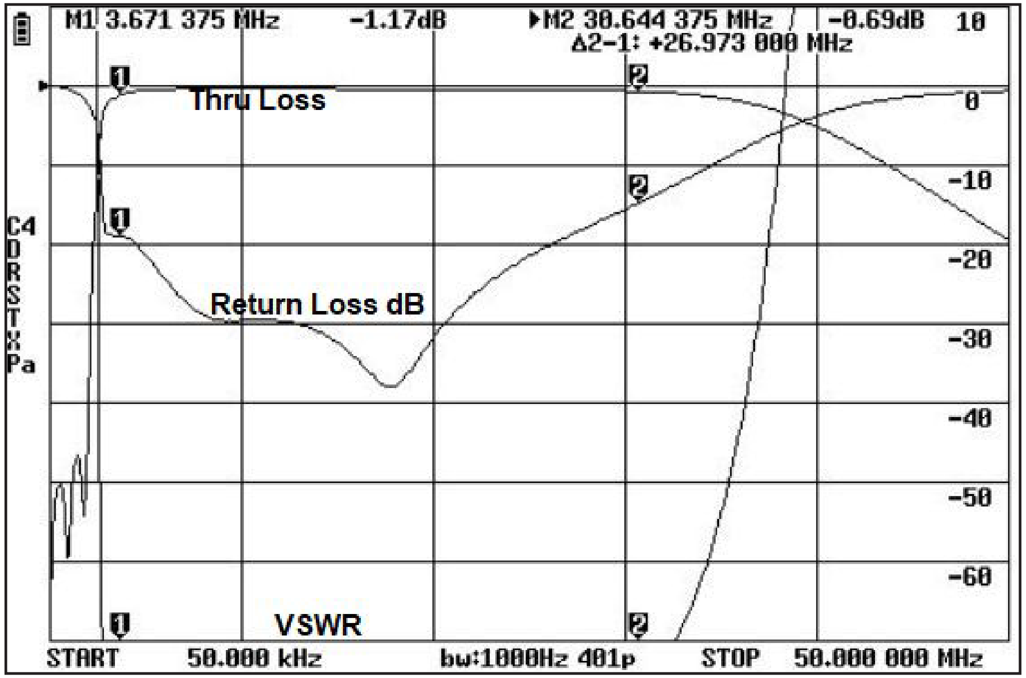

FIGURE 3: 30MHz pass-band swept-frequency response (10MHz/div). Note that the VSWR trace at the bottom is 1.25:1.

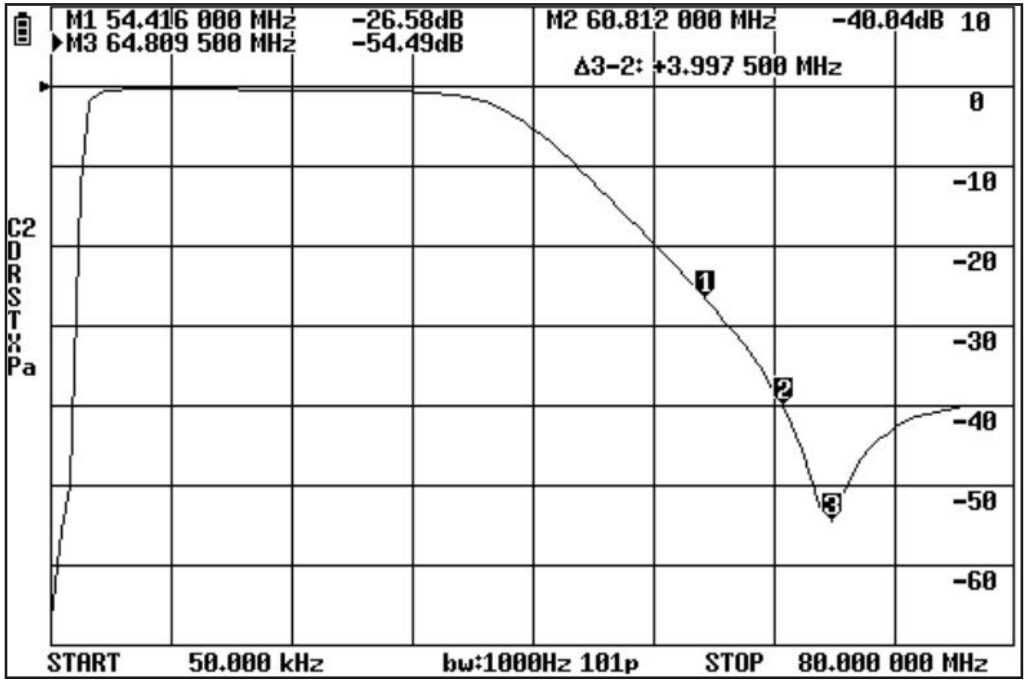

FIGURE 4: The 30MHz LPF upper-edge response, shown here in this NanoVNA screen capture, 10MHz/div.

optional battery-powered LNA for extending the usefulness of less-sensitive instruments (the filter in the larger enclosure shown in Figure 1) is also described.

SWITCHABLE LOW-PASS AND HIGH-PASS HF FILTERS

The project goal here was to built a switchable HF band-pass filter in a small enclosure that would become a portable accessory for my TinySA Ultra. I decided to use the same rather simplistic enclosure design concept that I had previously used for a small step attenuator that has been found easy to replicate **[1]**. The filters without the LNA employ the Hammond Manufacturing 1590A die-cast aluminium box (or equivalent) for the enclosure. The required enclosure preparation work simply involves drilling holes in the front panel in the exact locations specified by a template that is available for download **[2]**. The printed-circuit board (PCB) is mounted by the input/output SMA coaxial connectors, and so there is no need for additional mounting hardware. As shown in Figure 2, the enclosure could be vaguely described as a handheld rectangular box, about 9cm (3.5in) long, with the removable lid used as the front panel. I employed the same plunger-actuated latching slide switches for selecting the filter sections that were previously used in my step-attenuator design. These particular switches, manufactured by E-Switch, have silver-plated contacts. As they have proven to be reliable within my step attenuator for low-level RF signal switching, I concluded they should also be a good fit for this project.

Initially, I chose a design concept based on five switchable band-pass filter sections centred on the 80m, 40m, 20m, 10m and 6m bands, plus roughly 6MHz bandwidth at 315MHz provided by the SAW band-pass filter mentioned above. The 5-pole electrical design for each L/C filter section was generated using the ELSIE CAD freeware by Tonne Software **[3]**. The band-pass filter design concept was chosen initially as it could achieve individual amateur-band selectivity using readily-available surface mount (SMD) components that would fit within the limited PCB space. However, as the bandwidth-to-centre-frequency ratios involved were only a few percentage points, this concept turned out to be very sensitive to component tolerances, and was tedious to tune. By using a pair of nanoVNA network analyzers (one equipped with NanoVNA tweezers **[4]**), and swapping and stacking SMD capacitors, I managed to achieve good results. However, the tuning procedure was a tedious time-consuming process that made me consider a different design.

I decided to adopt a new approach and moved on to the switchable LPF/HPF HF band-pass design alternative that is described in detail below. It is much less affected by component tolerances, has practically no pass-band insertion loss, and does not require any tuning. However, it does not offer the advantage of being able to select individual HF amateur-bands. Although the HP/LP design is susceptible to high-level signals anywhere within the entire 3MHz to 30MHz HF spectrum, in most cases all that is needed is the blocking of high-power AM broadcast signals and local VHF/UHF signals including FM and TV broadcast. To allow for use in an area that happens also to be host to a high-power short-wave broadcast station, an unused set of PCB solder pads have been included for adding a custom-designed trap. An added bonus is that the transparency for the full HF spectrum makes the HP/LP design a good fit for general-coverage SDR short-wave listening applications.

To support those that might prefer the more-complex multi-band pass filter design alternative, the work-in-progress circuit diagram, template, and PCB manufacturing files, are available for download. Frequency response measurements and other information are also available **[5]**.

For the LP/HP design, I chose to provide individually-selectable, 7-pole 3MHz high-pass filter (HPF), and 7-pole 30MHz low-pass (LPF) L/C filter sections which, when both switched-in, form a 3-30MHz band pass filter response as shown by the sweep in **Figure 3**. In comparison to the 5-pole single amateur-band designs used previously, the 7-pole filters provide steeper HF band edges. The ELSIE CAD freeware was once again employed for designing the filter sections. The Cauer (sometimes called elliptical) filter type was used for the 3MHZ HPF. It has a very sharp cut-off characteristic for signals below 3MHz, and provides deep rejection of the AM broadcast band. In the interest of achieving optimum VHF/UHF rejection, I chose a Chebyshev design for the 30MHz LPF. This provides good rejection of the heavily-used VHF and lower-half of the UHF spectrum band allocations.

As shown in **Figure 4**, the 30MHz LPF provides more than 20dB rejection of signals

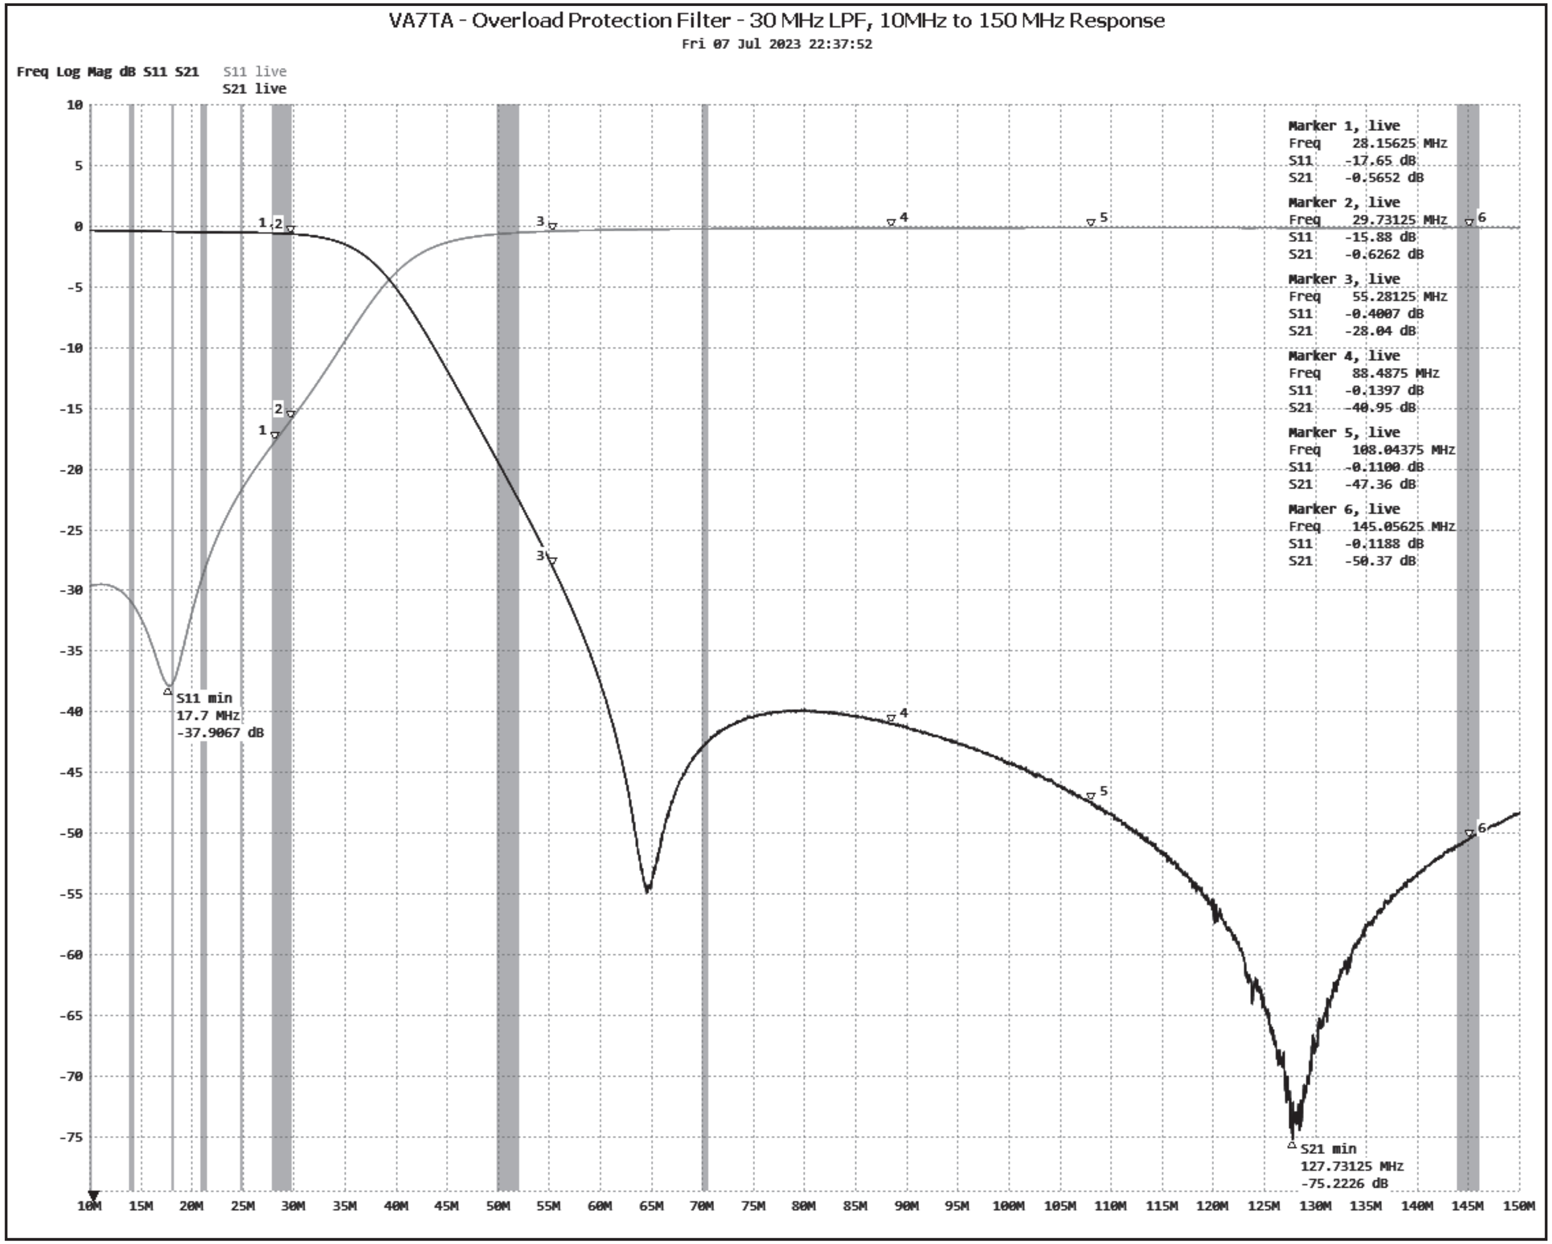

FIGURE 5: 30MHz LPF response. The shaded vertical strips represent amateur bands.

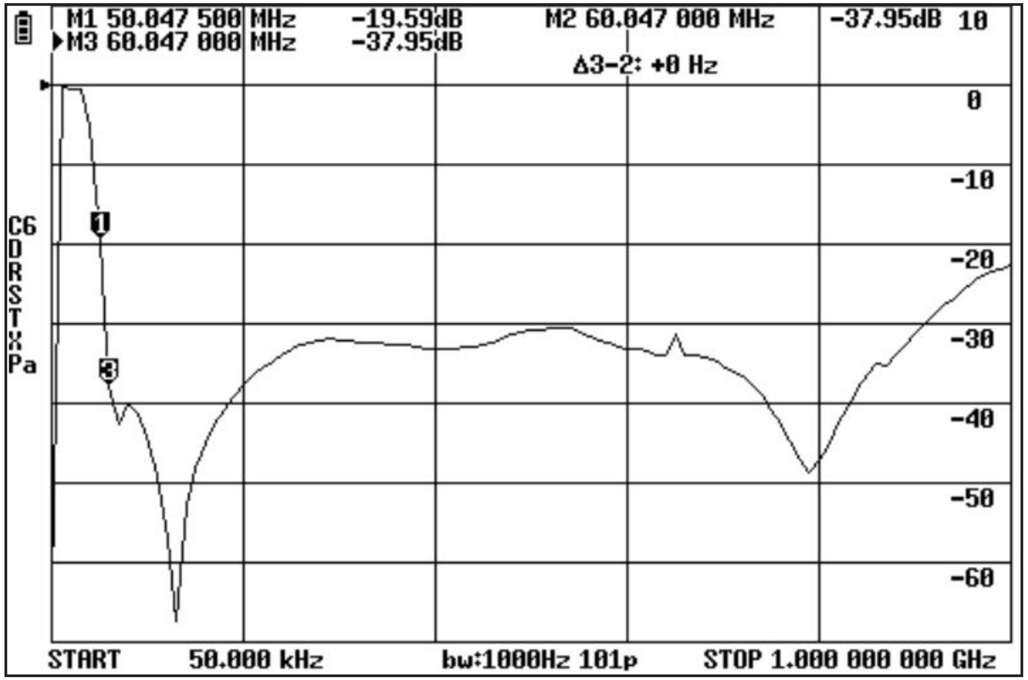

FIGURE 6: VHF/UHF spectrum rejection shown in this NanoVNA screen capture, 200MHz/div.

FIGURE 7: (a) 3MHz HPF Monte Carlo plot C±1%, L±2%;

above 50MHz, and more than 40dB for signals above 60MHz. The rejection in this part of the spectrum is of significance considering the possibility of powerful broadcast signals within the television VHF low band. Although 20dB does not sound like much rejection, it should be kept in mind that the level of intermodulation products generated within test equipment caused by signal overload would be reduced in absolute level by 2 or 3 times the level of reduction of any high-level signals located within the stop band. For every dB an overload signal is reduced, the absolute level of the internally-generated intermodulation products decreases by 2dB or 3dB, dependent upon on the order of the intermodulation products. Thus for all but the most extreme cases, this rejection depth should suffice. Also, as shown by **Figure 5**, the rejection for the higher TV broadcast channels, FM broadcast, air-to-ground, land mobile, and the 2m amateur bands, increases to 40dB or more which should be more than sufficient.

As shown by **Figure 6**, the rejection of the UHF band remains reasonably good up to about 1GHz. Although the rejection falls off beyond 1GHz, it is unlikely to matter in practice as the UHF signals would probably be significantly attenuated by the lossy antenna response.

A very useful feature provided by the ELSIE filter design application is the Monte Carlo function, which generates graphical-plots of the effect of component-value tolerance variations on a filter. This feature calculates and plots the filter response variations for all combinations of component-value maximum deviations because of manufacturing tolerances. This requires many repetitive frequency-sweep calculations and plot cycles, perhaps several hundred depending on the number of components in the circuit. The multiple sweep cycles, and corresponding plotted response curves, paint an area that graphically illustrates the expected range of filter responses. The plot shown in **Figure 7(a)** shows the expected response range obtainable with higher-cost 1% or 2% tolerance components whilst the **Figure 7(b)** illustrates the wider range of responses that can occur using commoner lower-cost 5% tolerance components. This demonstrates that, for this application, the LPF/HPF cascade design concept used here is not particularly sensitive to small component value variations. As the primary purpose of this filter is to reject out-of-band signals at frequencies that are, in most cases, far from the 3-30MHz band edges, lower-cost 5% tolerance components should suffice.

ELECTRIC POWER DISTRIBUTION NETWORK EMI TRACKING SUPPORT

The use of a compact 315MHz UHF directional yagi antenna, with a bandwidth of 6MHz, for triangulating the location of interference from loose or defective arcing hardware on power distribution poles (known to generate broad-band 50Hz or 60Hz buzzing EMI), has previously been documented in an ARRL QEX Magazine article **[6]**. As previously mentioned, the 315MHz SAW band-pass filter stage is provided to support the tracking down of electrical power EMI sources. When desired, the 315MHz band-pass filter is individually selected, with the other filters

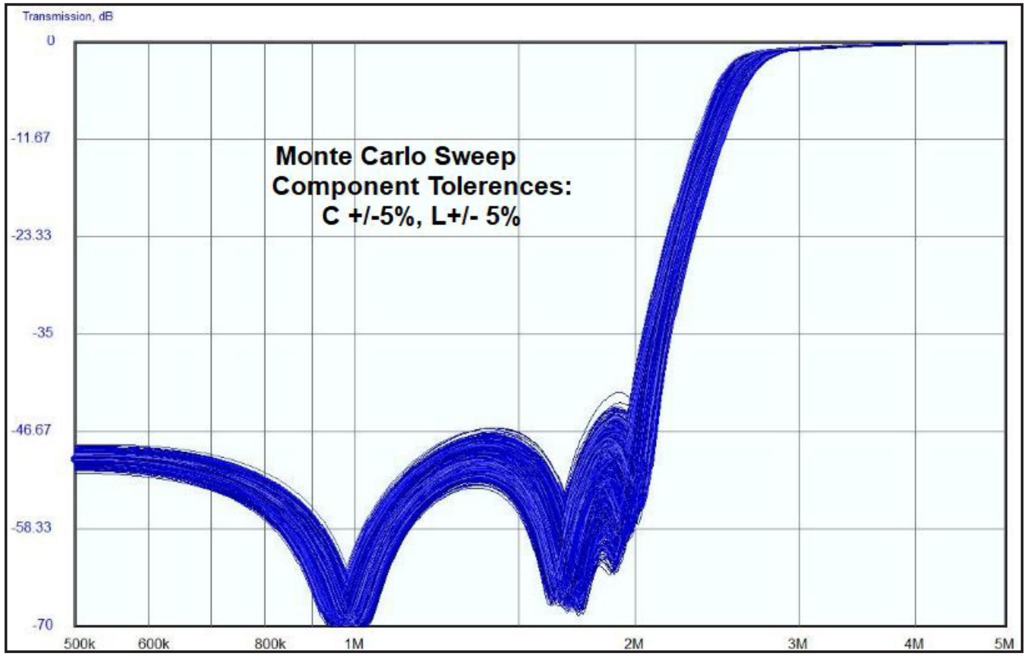

FIGURE 7: (b) a 3 MHz HPF Monte Carlo Sweep C±5%, L±5%.

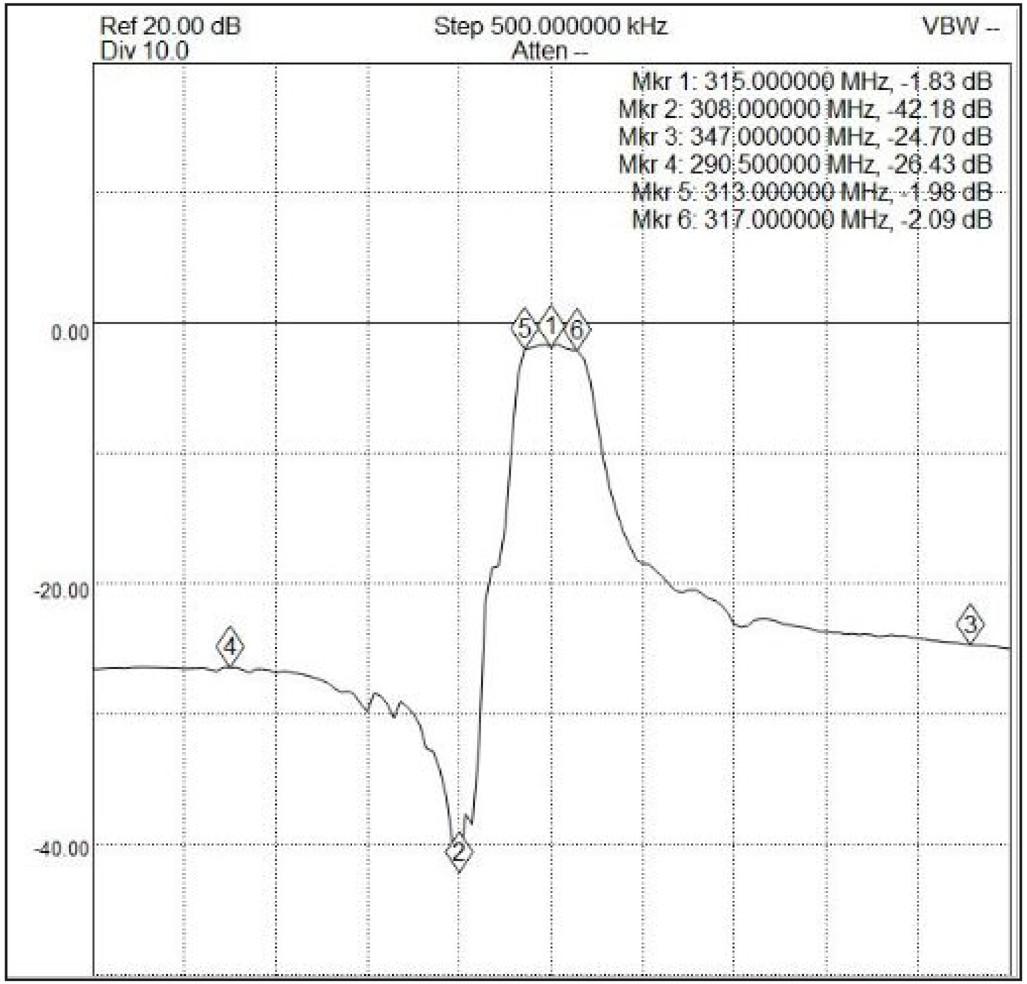

FIGURE 8: 315MHz SAW band-pass filter frequency response.

by-passed. The SAW filter is used to provide overload protection from strong signals that may be present outside of its pass band.

In North America, the 315MHz spectrum slice is allocated to low-power short-range wireless remote-control applications such as the control of garage doors. The band-pass filters needed for the common remote-control applications has resulted in the large-scale manufacturing of tiny 315MHz SAW filter chips, which perform well and are readily available at a very low cost. As shown in **Figure 8**, this filter stage has a low insertion loss of only 1dB or 2dB, and has a bandwidth of about 5MHz. It should permit sensitive measurements within the 313MHz to 318 MHz range. For EMI noise-source hunting, a frequency within the 5MHz wide filter pass band, that is free of local signals, needs to be selected. The 312-315MHz segment in the UK is assigned to fixed/mobile, earth-to-satellite services. Considering the use of typical SSB bandwidth detection devices, such as the TinySA or SDR receivers, a vacant 2KHz narrow slice within this relatively-broad 5MHz filter bandwidth should be easy to find when setting up in the field.

STATIC DISCHARGE PROTECTION

Transient voltage suppression (TVS) diodes are included within the design to protect against equipment failures caused by high-voltage static-electricity discharges. Without protection, static discharges can destroy sensitive RF semiconductors. Internally, the 315MHz SAW band-pass filter chip is a voltage-sensitive component that also needs static discharge protection. The TVS diodes employed here will shunt voltages greater than about 5V. They present a very-high impedance to voltages less than 3V, and have a shunt capacitance of 0.5pF or less. Thus the presence of the pair of TVS diodes is transparent to the desired low-level HF/VHF spectrum of interest, which will therefore pass through without noticeable attenuation.

LOW-NOISE AMPLIFIER OPTION

One of many improvements, introduced in the recent TinySA Ultra spectrum analyzer design, is the built-in low-noise pre-amplifier (LNA), which can be switched in to increase the RF sensitivity. When selected, it offers greater RF sensitivity than that obtainable from numerous high-cost spectrum analyzers, many of which were not

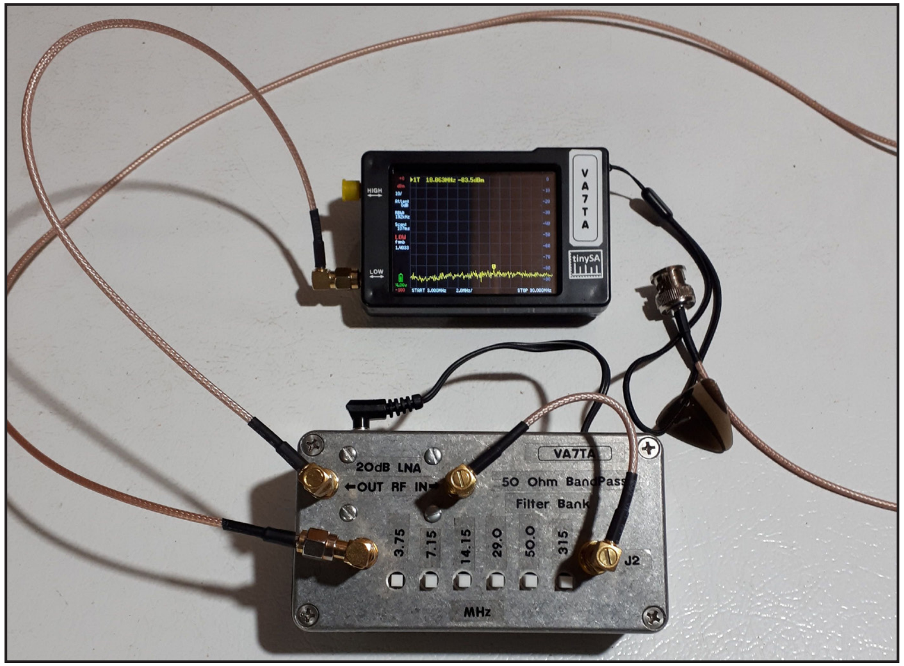

FIGURE 9: An earlier version of the LNA band-pass protection filter, mounted in the larger Hammond 1590B size enclosure, shown connected to the TinySA spectrum analyzer.

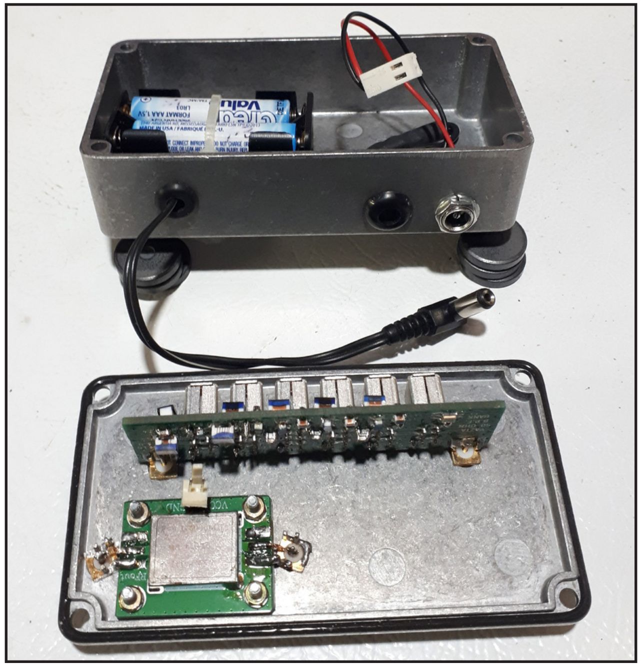

FIGURE 10: The protection filter and 5-600MHz 20dB LNA, purchased online from AliExpress, installed on a Hammond 1590B enclosure lid. The 4 AAA cells and the plug/jack switch assembly is also shown.

designed with built in LNAs, and consequently have higher noise figures. As better sensitivity makes locating low-level noise sources easier to pinpoint, I decided that it would be desirable to provide an alternative design option that included an LNA for supporting the use of less-sensitive instruments. As shown in **Figure 9** and **Figure 10**, this option requires a larger enclosure to provide the additional space needed for the 20dB LNA module, along with the associated AAA-cell battery power source. The LNA can be easily powered up and patched into the signal path with a short SMA coaxial jumper cable to provide a sensitivity boost when needed.

CIRCUIT DESCRIPTION

The circuit diagram is shown in **Figure 11**.

The filter bank itself can pass signals in either direction but, for the purposes of explanation, we will consider SMA coaxial connector J1 as the input, and the signal flow will be from left to right and out via J2. The input signal is coupled by DC blocking capacitor C1 to slider contact pin 5 of DPDT latching slide switch SW1 pole 'B'. Shown latched in the by-pass position, the signal path passes on to the next filter section via pin 2 of the SW1 pole 'A' slider contact, and then on to the next section via the SW2B slider contact. Similarly, when all stages are bypassed as shown in Figure 11, the signal path bypasses all the filter sections and exits via DC blocking capacitor C9 to J2. The transient-suppression diodes D1 and D2 are connected between the SW1B and SW6A slider contacts and ground respectively.

The 3MHz HPF section consists of C2 to C8, and L1 to L3 inclusive. As previously mentioned, it was designed with ELSIE CAD, in this case using the Cauer capacitor input filter type selection options. When SW1 is pressed and latched, signals above 3MHz are coupled through to the next stage via series capacitors C2 to C5 inclusive. For decreasing frequencies below 3MHz, the shunt loading, presented by L/C circuits C6/L1, C7/L2 and C8/L3, increases significantly because of the cancellation of the opposing phases of the L and C reactance values. Additionally, the coupling capability of capacitors C2 to C5 decreases because of the increased capacitive reactance injected in series with the path. As shown in Figure 7(a), this combination blocks the passage of lower-frequency broadcast-band signals.

The second filter in the cascade, controlled by switch SW2, is the 315MHz SAW band-pass filter. It simply consists of a single integrated circuit with pin 2 as input and pin 5 as output. Conveniently, it has an impedance of 50Ω, thus does not require any external matching components.

The third stage, controlled by SW3, consists of the 10dB resistive pi attenuator, provided for level-shift testing for identification of internal intermodulation product generation as described below.

The fourth stage, based on the 30MHz Chebyshev LPF design, is controlled by SW4, and consists inclusively of L4 to L6 and C12 to C15. Although the Chebyshev design does not offer as sharp a cut-off characteristic as the Cauer type, it was chosen for its better rejection characteristics of the UHF band. The signal-path parallel-resonant circuit capacitors within a Cauer design tend to couple the UHF band, whereas the stand-alone series inductors of the Chebyshev design tend to have the opposite effect, resulting in better UHF rejection.

PCB pads for two series-resonant RF traps have been added to the 30MHz LPF circuit. The RF pads provide for the installation of C10, L7 and R4 and C11, L8 and R5. The C10, L7 and R4 pads are spares to permit the installation of an HF trap should there be a need to attenuate a local HF broadcast signal (eg a continuous standard-time HF broadcast signal such as WWV). Note that, to avoid de-tuning the 30MHz LPF, C10 must not be larger than 2pF. The value

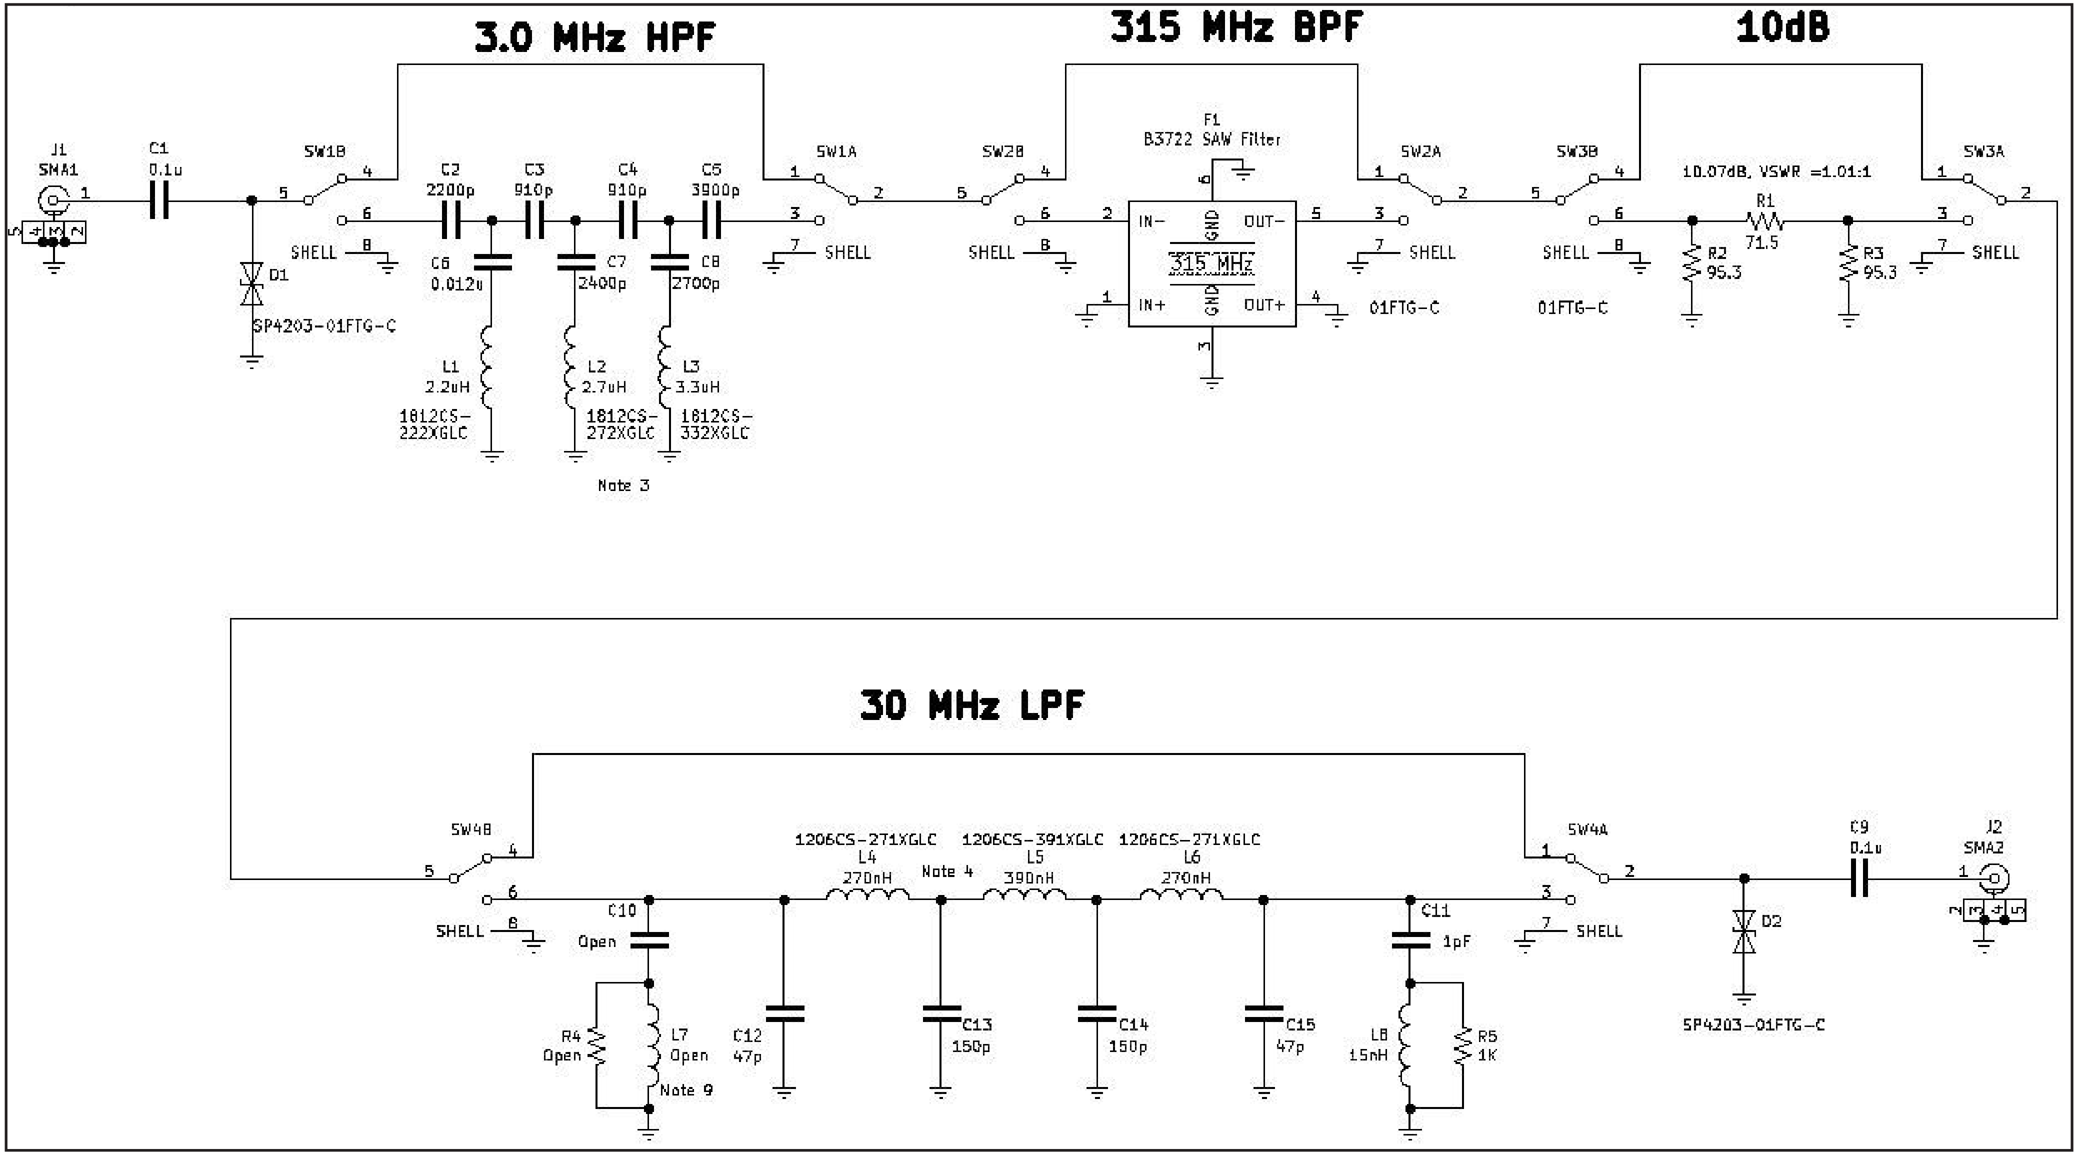

FIGURE 11: The circuit diagram of the high-pass, low-pass, band-width limit protection filter.

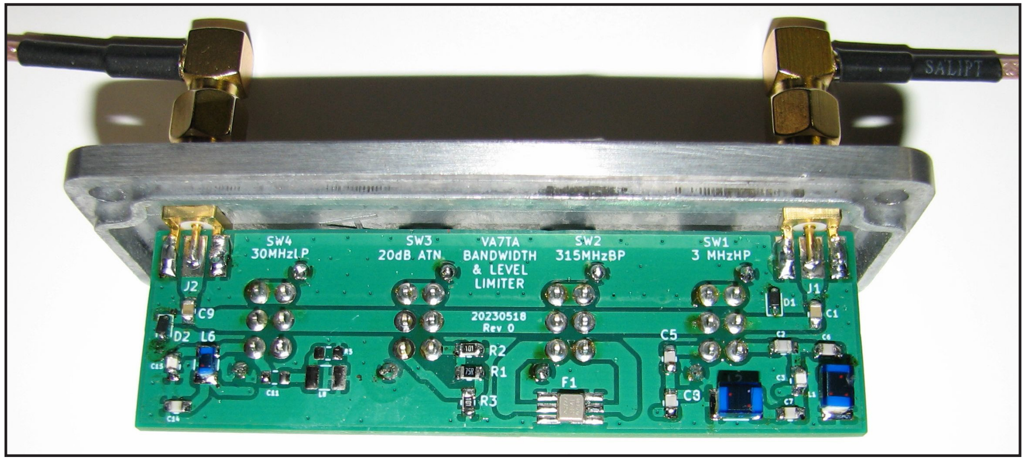

FIGURE 12: Populated PC board: the signal path and component side.

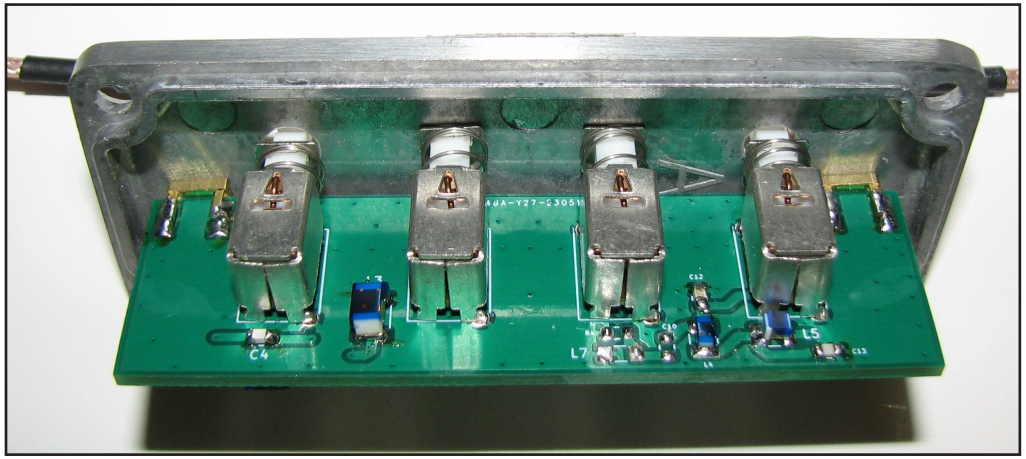

FIGURE 13: The switch side of the populated PCB.

of L7 must be selected to resonate at the desired trap frequency based on the C10 value of 2pF or less. The C11, L8 and R5 circuit is populated with components to improve the UHF band rejection around 1.3GHz. The resistors R4 and R5 provide a means to reduce the Q-value of the traps if it is found necessary to broaden the nulls.

CIRCUIT LAYOUT

Figure 12 shows the underside of the PCB where the majority of the SMD components are mounted. J1 is shown on the right. The large blue inductors on the bottom right are L1 and L2, part of the 3MHz HPF. The inductors used here do not have magnetic shielding. Note that L1 and L2 are mounted at right angles to each other in order to minimize the magnetic coupling between them. Inductor L3, also part of the 3MHz HPF circuit, is mounted on the opposite side of the PCB at right angles to L2. The shielding from the PCB ground plane and switch body help to decrease the coupling between the inductors. Also note that the inductors for the 30MHz LPF at the opposite end of the PCB are mounted in a similar fashion. The components for the 3MHz HP and 30MHz LP filters are mounted at opposite ends of the PCB to minimize coupling between the filter HPF/LPF stages.

The 315MHz SAW filter, F1, is below the SW2 footprint. Wide traces, of roughly 50Ω characteristic impedance, connect the SW2 poles to the input/output pins of F1.

The 10dB resistive pad, consisting of resistors R1, R2 and R3, are shown to the left of F1, connected to the SW3 poles.

The tiny D1 and D2 static-discharge protection TVS diodes are shown mounted close to the J1 and J2 connectors, as are the C1 and C9 DC-blocking input/output capacitors.

Figure 13 shows the switch side of the PCB where the remainder of the SMD components are mounted. Note that the switch bodies themselves provide good shielding between stages which helps minimize inter-stage coupling.

PARTS PROCUREMENT

Figure 14, intended only as a guide, lists the details for the particular components I used. Note that many of the parts, such as the capacitors and resistors, are generic and alternative equivalents will function just as well. Digikey is shown as just one possible supplier. The capacitor voltage ratings are minimums. Some of the capacitors listed might have higher voltage ratings than those shown with 'in Figure 14.

Inductors have relatively-complex specifications. I prefer CoilCraft inductors for documentation and quality reasons. They may be purchased directly from CoilCraft or via one of their distributors **[7]** . There may be similar inductors available from other manufacturers. Besides the inductance values, check and compare the SMD size, the Q value, the self-resonant frequency (SRF), and the tolerance specifications. Although I used 2% tolerance inductors, 5% tolerance should be sufficient for this application as it is not particularly critical.

The J1 and J2 SMA connectors, and the die-cast aluminium enclosures, can be obtained from many sources. I used an Asian supplier as the prices were lower, and the quality was good enough for this application.

The Gerber files, which define the PCB, are contained within a ZIP archive available for download **[8]**. This archive when uploaded to PCB manufactures, such as JLCPCB, provides the design details needed for the manufacturing process.

PRINTED CIRCUIT BOARD COMPONENT INSTALLATION

The RF-performance and small-size objectives of this project could not have been met without the use of surface mount (SMD) components. There are various methods used to solder SMD components onto PCB pads, according to personal choice. For small-scale one-off projects like this, I like to use my soldering station, Teflon-tipped tweezers, a magnification lamp, a flux pen, and very-fine 22-gauge solder. My soldering station has a sharp pencil tip, and is temperature controlled. I also use a temperature-controlled hot-air tool.

It is important, particularly for this project, to

Intermod Protection Filter Parts List						Rev. 1	2023-07-24
Part ID	**Value**	**Tol. +/-%**	**SMD Size**	**Type**	**Mfr.**	**Mfr. Part #**	**DigiKey Part #**
C1,C9	0.1uF/100V	20	0805	X7R	Kemet	C0805C104M5RAC7800	399-C0805C104M5RAC7800TR-ND
C2	2200pF/50V	5	0805	COG/NP0	Murata	GRT2165C2A222JA02D	490-GRT2165C2A222JA02DTR-ND
C3,C4	910pF/50V	5	0805	COG/NP0	Murata	GRM2165C1H911JA01D	490-1621-2-ND
C5	3900pF/50V	5	0805	COG/NP0	Murata	GRM2165C1H392JA01D	490-3045-2-ND
C6	0.012uF/50V	5	0805	COG/NP0	Murata	GRM2195C1H123JA01D	490-3320-2-ND
C7	2400pF/50V	5	0805	COG/NP0	Murata	GRM2165C2A242JA01D	490-14872-2-ND
C8	2700pF/50V	5	0805	COG/NP0	Murata	GRM2165C1H272JA01D	490-1630-2-ND
C10	OPEN	N/A	0805	COG/NP0	N/A	N/A	N/A
C11	1pF/50V	0.25pF	0805	COG/NP0	Wurth	885012007094	732-12215-2-ND
C12,C15	47pF/50V	5	0805	COG/NP0	Kemet	C0805C470J5GAC7210	399-C0805C470J5GAC7210TR-ND
C13,C14	150pF/50V	5	0805	COG/NP0	Wurth	885012007083	732-12193-2-ND
D1,D2	3V TVS	N/A	SOD323	TVS	Littlefuse	SP4203-01FTG-C	F11893TR-ND
J1,J2	SMA PCB Edge	N/A	N/A	PCB Edge	RF JKM	SIMILAR – SEE TEXT	343-CONSMA003.042-L-G-ND
L1	2.2uH	2	1812	WireWound	Coil Craft	1812CS-222XGLC	https://www.coilcraft.com/en-us/
L2	2.7uH	2	1812	WireWound	Coil Craft	1812CS-272XGLC	https://www.coilcraft.com/en-us/
L3	3.3uH	2	1812	WireWound	Coil Craft	1812CS-332XGLC	https://www.coilcraft.com/en-us/
L4,L6	270nH	2	1206	WireWound	Coil Craft	1206CS-271XGLC	https://www.coilcraft.com/en-us/
L5	390nH	2	1206	WireWound	Coil Craft	1206CS-391XGLC	https://www.coilcraft.com/en-us/
L7	OPEN	N/A	0805	N/A	N/A	N/A	N/A
L8	15nH	2	0805	WireWound	Murata	LQW2BAS15NG00L	490-16073-2-ND
R1	71.5 Ohm	1	0805	ThickFilm	Yageo	RC0805FR-0771R5L	311-71.5CRTR-ND
R2,R3	95.3 Ohm	1	0805	ThickFilm	Yageo	RC0805FR-0795R3L	311-95.3CRTR-ND
R4	Open	N/A	0805	N/A	N/A	N/A	N/A
R5	1K Ohm	1	0805	ThickFilm	Stackpole	RNCP0805FTD1K00	RNCP0805FTD1K00TR-ND
SW1,2,3,4	DPDT switch	N/A	N/A	Slide	E. Switch	LC2255EESP	EG5891-ND
Enclosure	DieCast Alum Box	N/A	N/A	DieCast	Hammond	1590A Hammond – see text	HM150-ND
PCB					JLCPCB	See Text	See Text
PCB	MfrFiles:	https://drive.google.com/file/d/11ZFWwXtmFQU3JtYCQ2sEO9shtjcf-64N/view?usp=sharing					

FIGURE 14: A table of components.

FIGURE 15: The front panel labelled with transparent label tape.

install the smaller components first so that the larger parts do not block access for positioning the smaller parts. I started with the tiny TVS diodes and the SAW filter. Then I installed all the 0805-size capacitors and resistors, starting with the input/output coupling capacitors C1 and C9. Next I installed the inductors, and finally the switches. It is important NOT to solder connectors J1 and J2 until the board is fitted to the front panel to ensure that the PCB can be shifted a little to optimize the switch-plunger alignments with the corresponding panel holes. I use a slightly larger soldering pencil iron, that has more heat transfer mass in the tip, for soldering the switches and connectors. This provides a faster temperature rise which helps ensure proper time-efficient soldering, whilst minimizing the build up of heat in the surrounding area.

As shown in Figure 12, the bottom side of the printed circuit board provides mounting space for the majority of the SMD components. With amateur construction limitations in mind, larger-sized SMD components were chosen where possible. On the SMD size scale, the capacitors are the relatively-large 0805 size. The inductors sizes are 1206 and 1812 which, as the numbers suggest, are even larger. I believe that most amateur builders find these sizes sufficiently large to install with just the help of a magnification lamp or hood for a more detailed view. An exception here is the tiny size of the over-voltage protection diodes which were selected for their optimum electrical characteristics and availability. Although the diodes are in quite a tiny package, the terminal spacing happens to match nicely the 0805 SMD footprint pad spacing. The much larger 0805 footprint pads, which are easier to solder manually, were adopted here for the TVS diodes to make installation of these minuscule packages easier. The SAW filter is also packaged in a chip that is a bit challenging to install. It is only available in a 6-pin package with lead-less contacts exposed only on the bottom of the chip; I give more details below about how to install this item.

To install the common two-terminal SMD chips, I wet one of the pads first with a thin layer of solder then slide the chip into place with my tweezers. Once the chip is stuck in place, I solder the opposite side and then return to the first pad

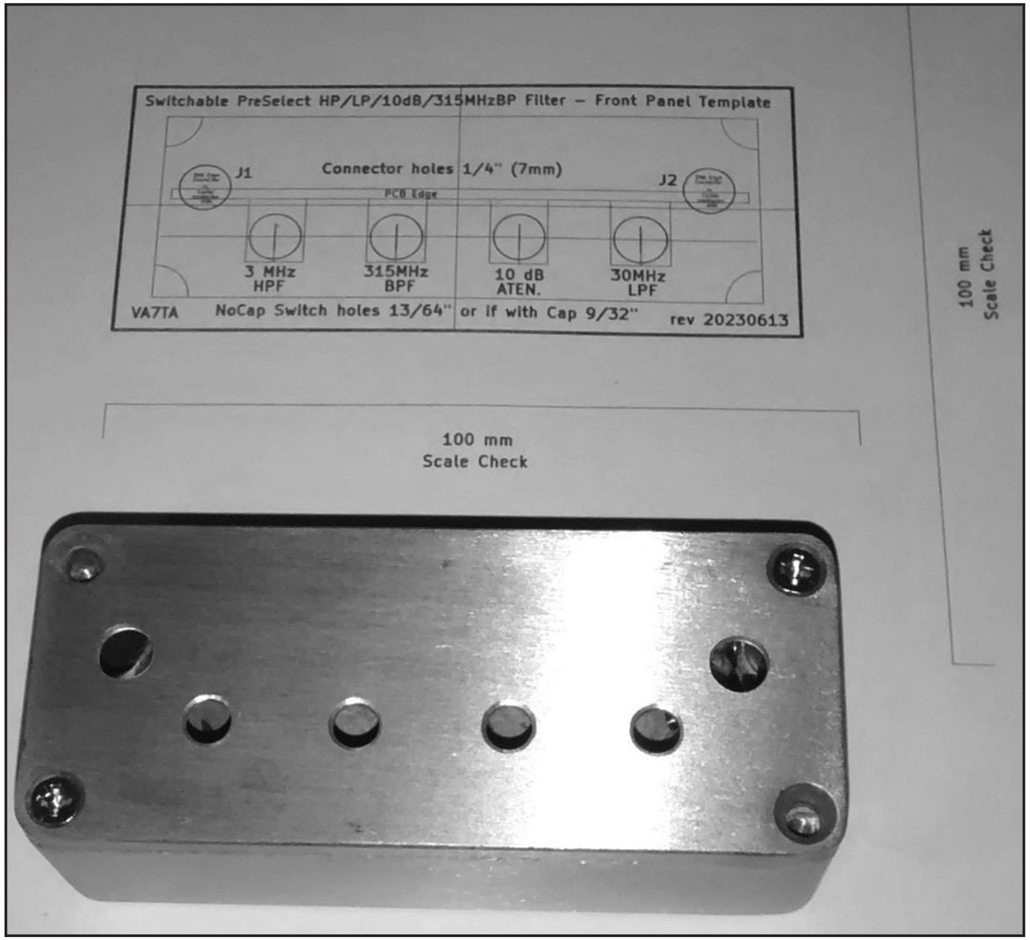

FIGURE 16: A view of template printout with scale-check lines, and the enclosure showing the drilled front panel. Switch-plunger holes have been chamfered to provide smooth edges. The SMA coaxial connector holes are not chamfered.

to solder it properly. Should I ever need to remove a two-terminal passive chip, I use a pair of pencil-tip irons and simultaneously heat each pad. In the vast majority of occasions, the chip will pop off the pads and stick to one of the irons until knocked off onto the work bench surface. The use of a solder flux pen to coat the pads in advance helps the soldering process significantly.

In the case of the SAW filter, I first tin each of the 6 footprint pads with a very-thin layer of solder then, whilst holding the chip in place with my Teflon-tipped tweezers, I heat the chip and the immediate surrounding area with the hot-air tool until the solder melts on all 6 pads. Once the chip is soldered down, I reheat the two middle input/output transmission-path pads with my iron,

FIGURE 17: A close-up view of a switch-plunger front-panel hole with chamfered edges, taken with a low-power microscope.

whilst adding a tiny bit more solder, to ensure good connections to the centre pins.

ENCLOSURE

A close-up view of the front panel is shown in **Figure 15**. The labelling was simply done with a Brother P-Touch label maker using black-on-clear 6mm wide and 12mm wide label tape as appropriate. It is a somewhat crude method of labelling, but effective and quick.

Figure 16 and **Figure 17** illustrate the switch-plunger holes with the smoothed, chamfered edges. The holes must not be left with sharp edges as otherwise the nylon switch plungers, which are a bit sloppy, can get caught by a sharp circumference edge and stick part-way in or out. The hole edges must be smooth chamfered on both the outside and the inside of the front panel. I very cautiously spin a 5/16in drill bit in my drill press to shave the sharp hole-shoulder edges off on both the inside and outside of each switch plunger hole. If doing this manually with a hand drill, one must be very careful to ensure that the larger drill bit does not grab, dig-in and inadvertently pass all the way through; the resulting larger hole would ruin the front panel. Note, do not chamfer the SMA connector holes, as the connector barb washers work best with a clean sharp hole edge.

Figure 17 was taken with a low-power microscope, and it illustrates good switch alignment. The trimmed hole edge that was shaved off is also shown. The plunger is shown in the out (stage bypass) position. The spring-loaded washer, which is part of the switch plunger assembly, can be seen on the far side of the hole flat up against the inside of the panel. Note that the square white plastic switch plunger is well centred in the hole. This well-centred alignment, and chamfered hole edges, will ensure unrestricted operation of the switch.

For those that so wish, there are more-decorative push-button caps available for these switches that will snap onto the plunger shafts. As noted on the template, use of the caps that I tried a few years ago required larger 9/32in plunger holes in the front panel to provide sufficient operational clearance. As the size of available plunger caps probably varies, a different hole size may now be needed to match. Personally I find the bare plunger shafts are functional without the caps, and that it is easier to tell the position of the switch without caps. However, the button design does improve the appearance of the front panel.

Figure 18 is a detailed view of the drilling template that can be downloaded in PDF format from the author's Google Drive account [8]. The template, when accurately printed to scale, can be used for precise front-panel hole centre-punch marking. Once printed out, the scale accuracy must be confirmed. This can be done by measuring both 100mm X- and Y-axis scale-check lines, which are printed out adjacent to the template, as shown in Figure 16. It is very important to check the scale as printer drivers often shift the scale according to option settings. In my case, using my Brother laser printer, I need to select 100% and the 'none' setting to defeat the automatic fit-to-sheet-margins scaling that is commonly used for printing documents. Once I selected the correct settings, my 100mm scale check lines each measured 100mm within 0.5mm, which is about as accurately as I could measure with a ruler.

Once cut out of the sheet using paper scissors, the confirmed-to-scale template can be carefully centred on the bare enclosure front panel cover, and then firmly attached using transparent tape. The hole positions can then be carefully marked/indented for drilling with a centre punch. I use a small 1/16in drill bit to make the pilot holes. Then I gradually increase the hole sizes in small increments until I reach the hole size specified on the template drawing.

Once the front panel has been drilled, and the switch-plunger hole chamfering has been completed on both sides, the SMA PCB edge-mount connectors can be loosely installed. As shown in **Figure 19**, some SMA connectors are supplied with two types of lock washers. The split-ring type washer is not used for this project. The inside barb washer is installed on the outside of the panel directly under the nut. Do not use a washer on the inside of the front panel as this would affect the switch-to-panel spacing.

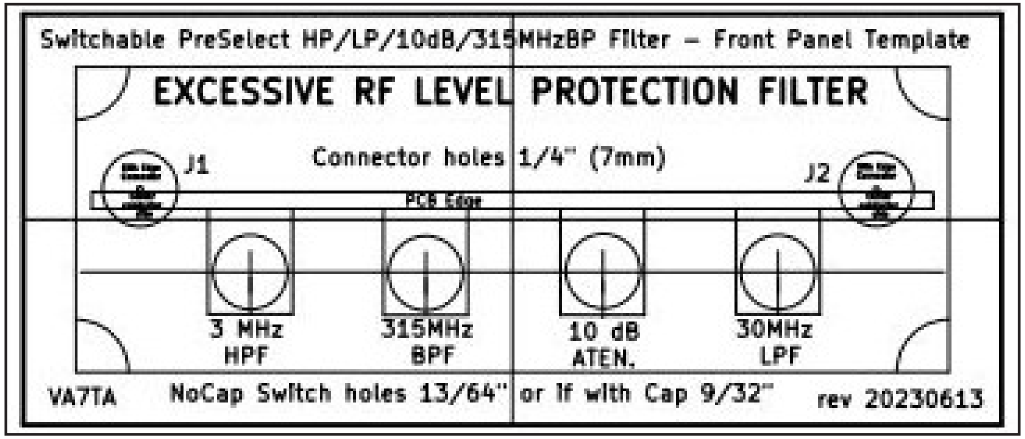

FIGURE 18: Front-panel drilling template (not to scale).

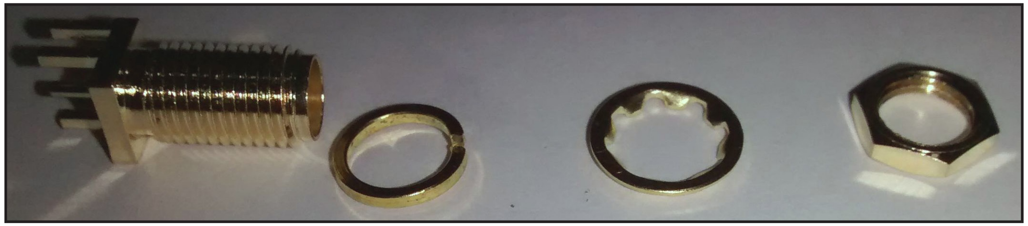

FIGURE 19: An SMA connector supplied with two types of lock washers.

Once the connectors are loosely installed, the PCB can be slid into the edge slots of the connectors as shown in Figure 12 and Figure 13. Once firmly seated in the slots, shift the PCB position laterally+/- for optimum switch-plunger hole alignment as shown in Figure 17. Once all four switch plungers are optimally aligned with their corresponding holes, then the connector nuts can be snugly tightened. It is important to ensure that the switch alignment remains optimum, and that the PCB remains fully seated within the connector slots. The connectors can then be soldered to the PCB pads with use of a larger-tipped soldering iron. Tighten the connector nuts to complete the PCB installation.

The final construction steps are to mount the front panel into the enclosure base and, if desired, install stick-on rubber feet on the bottom of the enclosure, a carrying strap loop (**Figure 20**), and a label on the front panel similar to that shown in Figure 15 with your call sign or name.

Figure 9 shows the earlier version 6-stage band-pass protection filter with the LNA option installed. It is shown connected to an earlier model TinySA spectrum analyzer that does not have a built in LNA. A larger enclosure (Hammond 1590B) is used to provide sufficient room to mount the LNA on the inside of the front panel alongside the filter PCB. The larger enclosure also gives enough room for a 4-cell AAA battery holder. I installed AAA alkaline cells for the power source as shown in Figure 10. The newer 4-switch version HPF/LPF filter described in detail here could also be mounted into a 1590B enclosure, along with an LNA and battery, in the same manner as shown in Figure 9 and Figure 10. The same drilling template described above would need to be used by attaching it to one side of the front panel in order to align the switches offset from the centre towards one side to provide room for the LNA.

Figure 10 also shows a jack and plug for external power, a combination which doubles as the power switch. A polarity protection diode (a 1N4001), which is not shown, is in series with the output of the power jack. In addition to polarity protection, it drops the 6V battery voltage down to about 5.3V, which is sufficiently close to the 5V level specified for this LNA. The LNA only draws about 20mA of current, so the AAA cells should last a good while.

The plug and power-jack switch approach allows powering the amplifier from an optional wall-wart supply, should the LNA happen to be needed on the test bench etc. The rubber grommet in the side of the enclosure provides a place to stow the power plug when the LNA is not in use. On the inside, a 2-pin connector is used to connect power to the LNA. This makes it easy to free the front panel from the enclosure base when it is necessary to replace the batteries. The battery holder is held in place with a sacrificial nylon tie wrap looped through two small holes on the bottom of the enclosure.

FIGURE 20: A carrying strap can free up hands for viewing a TinySA whilst pointing a directional antenna.

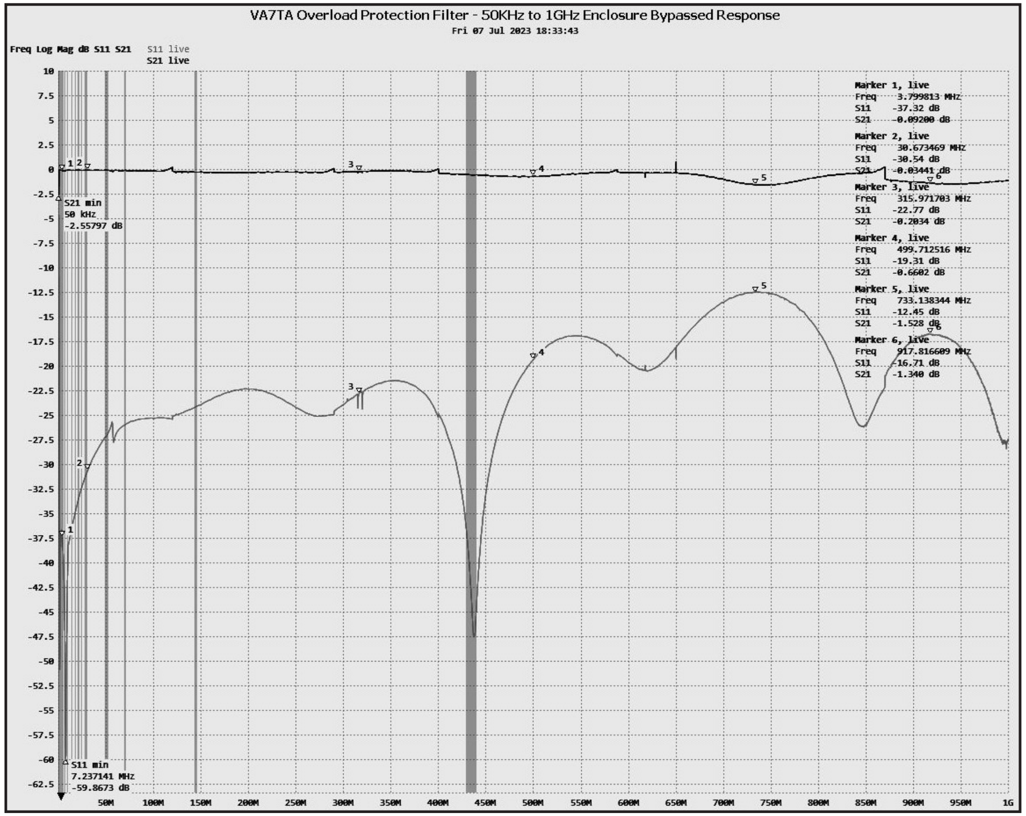

FIGURE 21: The enclosure response with all filter stages bypassed.

The shielded 20dB gain LNA module used here is designed for 5V and has a bandwidth of 5MHz to 600MHz. The LNA PCB was used directly as a drilling template to mark the holes. As shown, the SMA connectors were removed from the LNA PCB and mounted at right angles to the board through the front panel adjacent to the LNA PCB connector footprints. By mounting the SMA connectors at the shown rotated angle, one of the ground pins can be soldered directly to the PCB ground pads. Also as shown, a very short length of copper braid was used to re-connect the connector centre pins to the SMA pads on the PCB. Note that this LNA has shielding over the electronic components which is desirable for this application to minimize possible coupling to the filter coils.

TESTING

Testing can best be carried out with an RF sweep transmission test-set, such as a vector network analyzer (VNA); the sweep should begin at a frequency well below the AM broadcast band and extend to at least 500MHz. I used my genuine NanoVNA H4, calibrated for a frequency sweep from 50kHz to 1GHz. This wide sweep range was used to confirm the wide-bandwidth low-level signal transparency of the PCB and switches when in bypass mode (all 4 stages bypassed). The test checks for excessive insertion loss and/or for any significant impedance mismatch that could be caused by a bad solder joint, defective switch, or a PCB issue. The objective is to have an excellent match across the 3MHz to 30MHz HF spectrum, and a reasonably-good match well into the UHF 300MHz to 3000 MHz band.

Figure 3, Figure 4, Figure 5, **Figure 21**, and **Figure 22** illustrate measured frequency responses for all of the variations of by-pass switch positions. Any significant deviation from

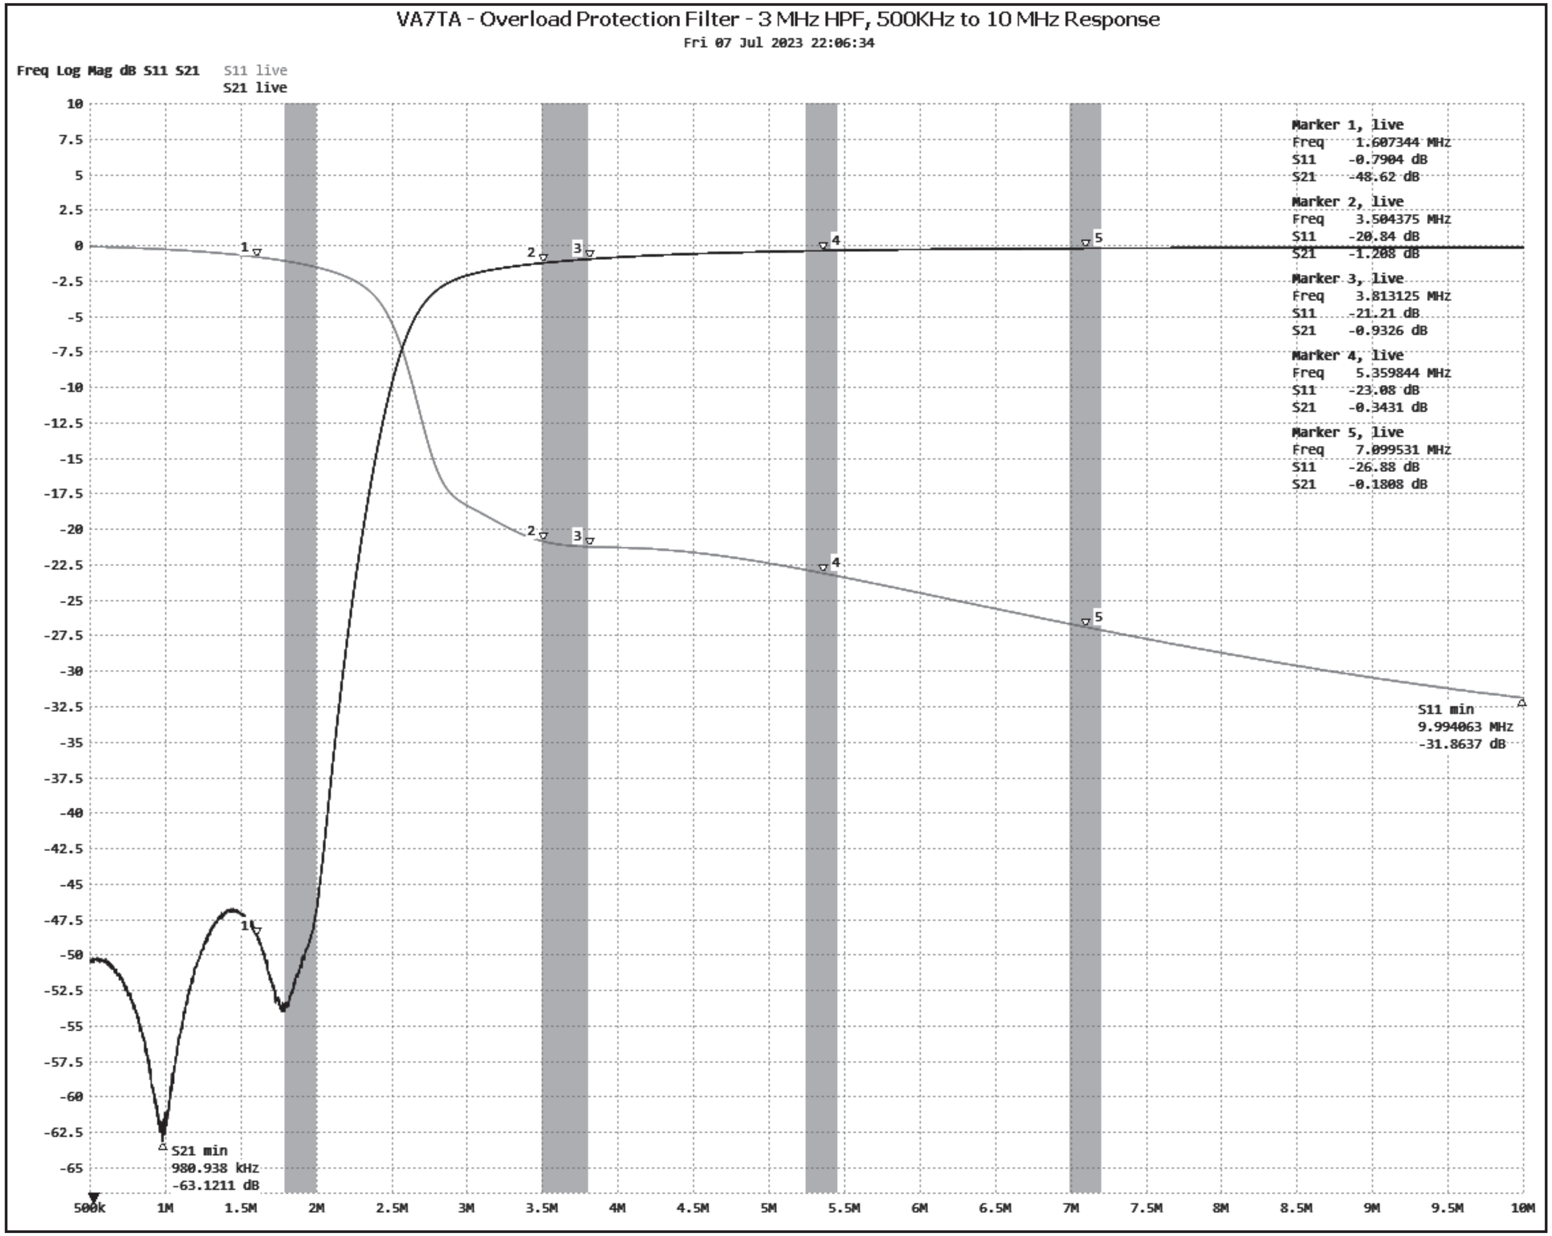

FIGURE 22: The 3MHz HPF response. Shaded vertical strips represent amateur bands 160, 80, 60 and 40m.

these test results would suggest there is an issue that needs to be identified and corrected. Figures 5, 21 and 22 are high-resolution (3201 measurement steps) screen captures from the NanoVNA-App PC application, which I use to control my nanoVNA. The vertical grey broad shaded stripes in the charts represent amateur radio allocated bands.

The 10dB attenuator can be tested by switching it in and out of the transmission path. The whole sweep width level should correspondingly shift by 10dB (within half a dB) at HF, although it may be less accurate at higher VHF/UHF frequencies. All the swept-frequency measurements of the filters shown here were obtained with the 10dB attenuator switch set to by-pass.

Figure 21 displays the measurement response from a wide 50kHz to 1GHz sweep of the transmission path straight through the connectors, PCB, and switches, with all filter stages switched out. The upper trace shows the insertion loss and the lower trace the impedance-mismatch reflection coefficient expressed in terms of return loss **[9]**. As shown, the HF-band insertion loss is less than 0.1dB, almost transparent. The loss at 316MHz, which is of interest for the SAW-filter pass band, is less than 0.3 dB. In general, the loss up to 650MHz is less than 1dB (less than 1/4 of a standard S unit). There is a dip in the response around 750MHz where the loss increases to about 1.5dB (1/4 of an S unit). These UHF-band losses are insignificant in this case since they are outside the usual pass bands of interest to the radio amateur.

Figure 22 shows detailed 500kHz to 10MHz swept-frequency response results for my 3MHz HPF. As shown, the loss at the bottom end of the 80m band is about 1.2dB, and it decreases to less than 1dB at 4MHz. A loss of 1.2dB is only about 1/5th of an S unit. VSWR is better than 1.2:1 (RL of 21dB or better) for all the bands from 80m to 30m.

Also as shown, the poorest rejection of the AM broadcast band is better than 46dB around 1.4MHz, and the deepest rejection at around 1MHz, the middle of the AM broadcast band, is greater than 60dB. In terms of power ratio, 46dB represents a reduction of about 40,000:1 and 60dB is 1,000,000:1.

Figure 5 illustrates the measured response of my 30 MHz LPF. This 10MHz to 150MHz measurement is focused on the pass band through loss up to 30MHz, and the rejection of the lower VHF television, FM broadcast, air-to-ground, land mobile and amateur 2m bands. As shown, the worst case through loss for the bands from 30m to the top of the 10m band at 29.7 MHz was measured to be 0.6dB equivalent to about 1/10th of an S unit. The return loss for all the bands up to, and including, 12m was measured better than 20dB (VSWR <= 1.2:1). The worst case pass-band VSWR at 29.7MHz was measured as VSWR <= 1.4:1 (RL = 16dB).

TIPS AND PITFALLS

Note that since 160m is very close to the AM broadcast band, it is rejected by the 3MHz HPF; similarly the 30MHz LPF rejects the VHF/UHF amateur bands. Of course, if measurements need to be made below or above the 3MHz to 30 MHz HF spectrum segment, the corresponding 3MHz HPF or 30MHz LPF respectively must be switched out.

The 10dB attenuator was included to provide a means to determine if an interfering noise, or unwanted signal, is being internally generated through intermodulation within the test equipment (or receiver), or whether it comes from an external source. When both LPF and HPF filters have been switched in to pass only the HF band, the attenuator can provide a quick check for the presence of any unnoticed high-level overload signals which are at HF frequencies but are beyond the current analyzer display range. For example, a high-level signal from a local HF transmitter would pass straight through the HPF/LPF filter pair without attenuation, and could overdrive an RF input circuit, resulting in intermodulation. If a noise spur, or extraneous signal, drops in level by 20dB or 30dB when the attenuator is switched into the transmission path, then that would be an indicator of overload. On the other hand, if a noise spur drops by 10dB, the same loss as the attenuator, then it is highly likely that the interference is originating from an external source.

It should be understood that the use of this filter will only reduce the noise floor level during periods when there are, in fact, high-level out-of-band signals present. If there are not any high-level out-of-band signals present, switching the filters in and out will not make any difference. My QTH is in a rural area where there are not any near-by broadcast transmitters, nor any VHF/UHF dispatch-type services. The strongest continuous signal at my QTH comes from an FM broadcast station about 15km away, thus my RF receiving environment is almost always free of signals of overload strength. For a case-in-point, I recently did a test of my filter by connecting it between my TinySA Ultra RF input and my 20m beam feed-line. The propagation K index at the time was high at 3, thus the band was pretty quiet. The strongest signals were only about S9 (50mV or -73dBm). Switching the filters in and out didn't make any noticeable difference to the 3kHz resolution band-width noise floor displayed on my TinySA. In contrast, mobile radio operators, who have experienced occasional VHF-rig squelch openings spewing out mixed audio from a multitude of signals whilst navigating within a city, will attest, from first hand experience, the potential intermittent nature of intermodulation noise.

CONCLUSION

The spectrum exposure limiting that is obtainable with the use of this filter can provide a means for protection against high-level signals outside of the 3MHz to 30MHz HF band. It can reduce the measurement noise floor level in cases when there are high level out-of-band signals present. This filter is an accessory that can protect instrumentation inputs, and possibly save the measurement day in cases where overload protection is needed because of an unexpected exposure to a high-level signal.

HF EMI often limits the reception of anything but very strong signals, which is frustrating for radio amateurs and shortwave listeners. In some cases, nearby powerful broadcast and commercial transmitters can overload sensitive receiving equipment, causing the generation of intermodulation products that can also raise the reception noise floor. With the seemingly ever-increasing sources of HF EMI emanating from the aging electrical power grid, modern-day microcontroller-controlled household appliances, and/or digital-based consumer electronics, there is an increasing need for portable tools that support the identification of the causes. Instrumentation tools that support the identification of noise sources can be helpful. The recent introduction of the handheld portable TinySA spectrum analyzer test-set is a big step in the right direction, in that it provides an affordable type of instrumentation that previously was only available at industry-level costs. It is hoped that the construction details provided here for this bandwidth- and level-limiting overload-protection accessory, will be found useful for those wishing to extend the capability of their sensitive instrumentation and/or receiving equipment.

REFERENCES

1. QEX July/August 2021, Miniature SMA RF Step Attenuator by VA7TA
2. Drilling template download link: https://drive.google.com/file/d/1KSKeyo87Se-FhW-g8SAANBB9IhmsgfjFt/view?usp=sharing
3. http://tonnesoftware.com/elsie.html
4. QEX Nov/Dec 2021, NanoVNA SMD Tweezers by VA7TA
5. Supporting information for multi-band-pass intermodulation protection filter work-in-progress design: https://drive.google.com/drive/folders/1OEd-jqyekZ_O0fw2Q5UY7oTRPdixpKP4I?usp=sharing
6. QEX Nov/Dec 2003, "The EMI Finder" by K0DK https://www.arrl.org/files/file/Technology/tis/info/pdf/031112qex048.pdf
7. https://mouser.co.uk
8. PCB Manufacturing Gerber files and other items: https://drive.google.com/file/d/11ZFW-wXtmFQU3JtYCQ2sEO9shtjcf-64N/view?usp=sharing
9. The term return loss (RL) it is a alternative term commonly used for expressing impedance mismatch similar to VSWR. RL is the ratio of the forward power to the reflected power expressed in dB, and a larger value is better. VSWR to RL conversion calculators can be found on various RF tools websites.

Building and testing antenna noise cancellers

Sheldon Hutchison N6JJA
N6JJA@arrl.net

From time to time, a few intrepid experimenters have published articles describing their contributions to antenna noise-canceller designs, and a nice introductory article has recently been published in RadCom **[1]**. But for me, a lot of information has been missing in all of this work, and so I set out to try to gain a better understanding of how things actually work. That led to about 2 solid years of R&D in my lab.

INTRODUCTION

As I begin, I should note that the project name for all of this is the 'GCR' project, not the name of another of our cats, but an homage to Giles C. Read, G1MFG (now SK), who was the *RadCom* technical editor when all this began, and who became a good friend. My conversations with Giles were what caused all this work to happen.

The Editor of *RadCom Plus* asked me to make this a 'construction' article; that is, offering something new that others might like to build, and so helping others to expand and improve what I report here. Such has been the history of amateur radio, and so it continues today. However, I'm choosing to offer the story of this antenna noise canceller first, starting, as it were, at the end, before introducing the technical background. What you will find below is the result of those months and years of work. The story about what I learned along the way is not entirely necessary in order to build a comparable antenna noise canceller, and is very long. I may publish the long story in a future edition of *RadCom Plus*. Meanwhile, if you have questions regarding the design published here, please send them along.

As I begin this article, however, I'm going to introduce an extra circuit that has actually proven to be beneficial in the final prototype designs. After that, it's on to the family of designs that I feel are worth reporting on, along with the test results.

MY PREAMPLIFIERS

Throughout the development of this work, there was one idea that stuck with me. Simply, for users who can't build multiple antennas that can be tuned to specific frequencies to do the work of supplying enough signal for an antenna noise canceller to work properly, perhaps just adding antenna amplification on both inputs could improve matters. If nothing else, it could help with weak signals, or at least balance the 'Main' and 'Aux' inputs better. I actually spent considerable time designing, building, and testing a small family of low-noise (2dB noise figure or lower) preamplifiers that can be easily reproduced by anyone. Of those, however, the best overall fit for noise-canceller work uses a dual-gate MOSFET, the 3SK293, in a reasonably-simple circuit. It has a bandwidth of at least 500MHz, and gain can be adjusted from -5dB to over 20dB with a single control. In addition, unlike a lot of low-noise RF components, this device is still currently produced. The bad news is that it's an SOT-243 package which I find fiddly to use, but the good news is that there are lots of sources (especially on eBay) for very small adapter PCBs to allow that tiny dot to become part of a prototype PCB circuit. **Figure 1** shows the completed 'adapted' configuration, soldered to 0.1-inch-spaced header pins. Noise measurements of a completed amplifier showed the expected noise figure of roughly 1.6dB. Using this 'adapted' transistor, I've never noticed any problems or instabilities, even at 100MHz.

The broadband version of the 3SK293 preamplifier at the inputs of a noise canceller worked okay, but there were background noise signals that overpowered the preamplifier at frequencies I don't care about, lowering available gain in the frequencies I do care about. That led me to add a group of tuneable (by amateur band) high-pass L-C impedance-transforming filters at the preamplifier inputs. I chose a high-pass design since the general amplitude distribution for random noise signals always means that there's more noise at frequencies below the tuned-in frequency than above it. So having less noise to compete with the desired signals in the preamplifier means that a high-pass filter is desirable. I started with a rotary switch to select filters, finally settling on some small relays (like the Panasonic TQ2-12V and many others) to do the switching with fewer tangled wires. Also, the filters supply impedance matching that works better over narrower bandwidths, causing me to reconfigure this circuit for my favourite six amateur-radio bands as indicated in **Table 1**.

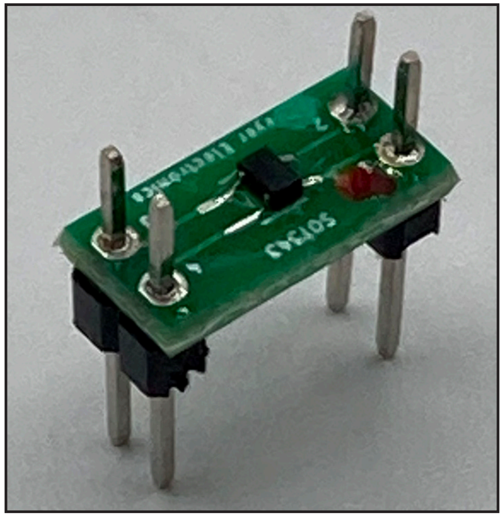

FIGURE 1: A 3SK293 mounted on an adapter PCB and pins.

The circuit diagram of my preamplifier is shown in **Figure 2(a)** and **Figure 2(b)**. For all of the circuit diagrams that you'll find in this article, non-electrolytic capacitors are in microfarads unless otherwise noted, and are of 50V ceramic type. Resistors, unless otherwise marked, are rated at ¼W, 5%. For all of the relays shown in the circuit diagrams, they're in their un-energised or 'normally closed' positions. Transformer T1 consists of 10 turns on the primary winding, 2 turns on the secondary winding, each of #28 AWG enamelled wire on an FT50-43 core. The tuning control (R1), marked 'frequency', can actually supply a tuning voltage to multiple preamplifiers in parallel. This is a much better solution than using multiple-ganged units. All the rotary potentiometers are low-power linear-scaled units. Choose RF chokes (1mH) that have a series resistance under 10Ω if you can. You

Band	Frequency range (MHz)	L (μH)	Winding data, #24 AWG enamelled	Capacitors
160 m	1.27-3.41	26.6	L1 = 22 turns on FT37-61	C1 = 0.1μF
80 m	2.01-4.45	12.4	L2 = 15 turns on FT37-61	C2 = 1200nF C0G type
40 m	6.60-8.60	4.5	L3 = 9 turns on FT37-61	C3A = 100pF C0G type in parallel with C3B = 120pF trimmer
20 m	11.95-14.95	1.4	L4 = 5 turns on FT37-61	C4 = 120pF trimmer
15 m	19.75-22.5	0.8	L5 = 13 turns T50-10	C5 = 10-55pF trimmer
10 m	29.2-30.2	0.53	L6 = 11 turns on T50-10	C6 = 10-55pF trimmer

TABLE 1: Band-selection components for the preamplifiers in Figure 2(b).

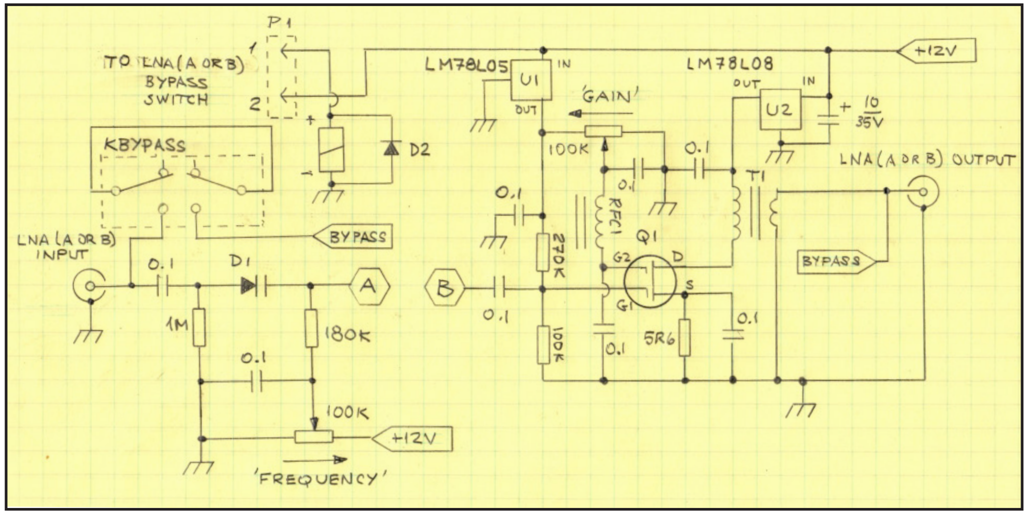

FIGURE 2(A): The circuit diagram of the preamplifier. Q1 is a 3SK293 MOSFET, U1 a 78L05 regulator, U2 a 78L08 regulator. D1 is a Toshiba 1SV149 varactor diode, D2 a 1N4001 rectifier. KBYPASS is a Panasonic TQ2-12V or equivalent DPDT 12-volt unit. The high-pass filters shown in Figure 2(b) are inserted between points 'A' and 'B'.

can use whatever connectors you have to hand. I tend to use SMA types that fit onto the edges of PC boards, but you can always wire directly if convenient. I tend to take things apart several times during a build, but I could also do that with a soldering iron. The pin-type connectors are the inexpensive 'Dupont' style available on eBay.

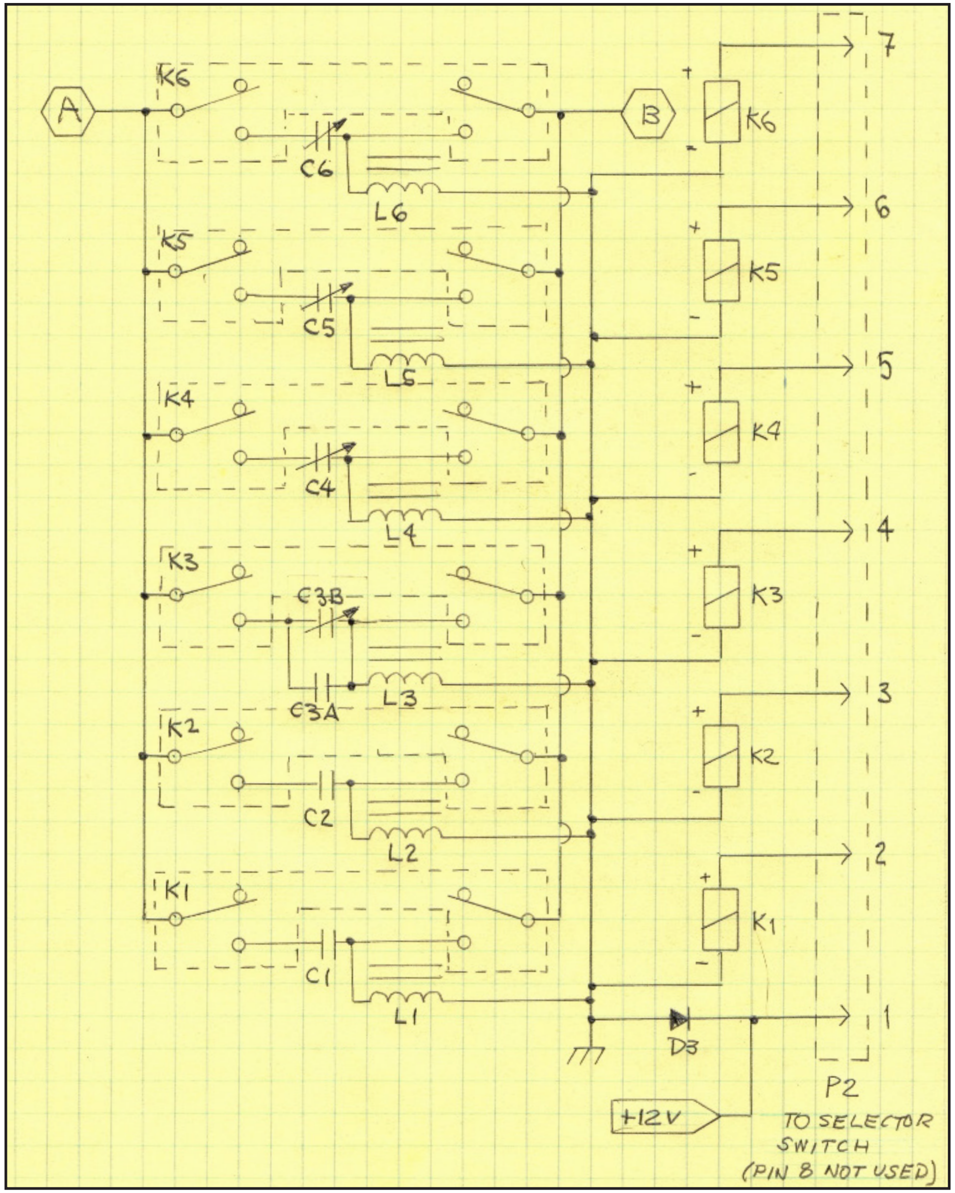

FIGURE 2(B): The LNA filter unit. See Table 1 for values. C1-C6 are typically Murata-style units. K1-K6 are Panasonic TQ2-12V relays or equivalent DPDT. D1 is a 1N4001 or equivalent. P2 is an 8-pin 'Dupont' style header, not using pin 8. Wires can be used instead of the header, and these go to a band-selector rotary switch.

As I've noted, I use six bands to cover my most-used amateur HF spectrum and, as I've done many times over the years, I use electrical tuning for this circuit. Coil and capacitor data are shown in Table 1. The varactor diode I use here is the Toshiba 1SV149. It's no longer in production, but Toshiba apparently made tons of them, and they're inexpensive and readily available on eBay. Note, if you're going to buy some, buy a lot. I use them in so many designs that I found it useful to buy 100. And why not? They have a range from 20pF to over 500pF using a 12V supply, a good Q value, and are robust and reliable. Having said that, mechanically-adjustable capacitors should work very well too, but they are physically much larger and more difficult to isolate. If you need to stick with mechanically-adjustable capacitors, consider also switching from high-pass to low-pass filters. These are not quite as effective, but much easier to construct electrically. For the 3SK293, I use a DC bias voltage on gate 1 of about 1.3V, and on gate 2 (for gain control) a DC bias voltage of about 1V to 5V. About 8V on the drain gives the best noise performance, with an input impedance of 1.5kΩ-2kΩ, and an output impedance of about 1.2kΩ. The drain current is 12-13mA with the components shown. The 1mH choke should be chosen for as low a series resistance as possible (<10Ω), and a test frequency as high as possible. I use two simple and inexpensive regulators to provide the bias voltages shown and a single +12V supply. The first regulator is an LM78L08 supplying +8V at the drain of the transistor. The second regulator is in parallel with that one, an LM78L05, providing the gate bias voltages. When building them as a pair, I install the regulators on one of the preamplifier boards, and then run wires to share the bias (and tuning) voltages with the second preamplifier.

I carried out extensive testing of the preamplifier design and, indeed, I built a pair of them, one for each antenna input. The trimmer capacitors (I use Murata-style units) allow me to overlap (to a reasonable extent) the tuning of both preamplifiers on each band, since they need to work in tandem. I use a single tuning control for both preamplifiers, but left the gain controls independent. Designing the L-C high-pass circuits is actually very easy **[2]**, and can be done to set up a user's device for optimum performance for different frequencies or bands as desired. I have the luxury of having a lot of equipment here, including multiple VNAs and spectrum analysers, but it is possible to tune the preamplifiers using a broadband source and an SDR, or one of the inexpensive NanoVNAs that other experimenters have written about.

NOISE-CANCELLER CONFIGURATIONS TESTED

The goal of all this work, naturally, was to build something that demonstrates one or more 'most useful' overall designs to use in an amateur radio station. I hope to explore some of the novel approaches that I tested and set aside in future articles. The results of my 2 years' worth of exploration and testing produced the following list of design variations, all based on the 'WIMO' or 'VK5TM' basic R-C design **[3]**. The in-depth work over this period provided enough data to convince me that this was the best choice, and the simplest. So the list is:

- a basic JFET-based design to use as a baseline;
- a copy of that design that also replaces the J310 JFETs with the 3SK293 MOSFETs, and thus also allows gain adjustment for each part of the adding circuit;
- the use of Toshiba 1SV149 varactor diodes to replace the fixed shifter capacitors with tuneable capacitors for a broader range of frequencies;
- try two cancellers in series (suggested as useful through an accident in my lab), either JFET-based circuits using independent adjustments for each stage, or MOSFET-based circuits with ganged adjustments; and
- add two preamplifiers (as described above) with independent gain adjustments, but common tuning, to insert at the antenna inputs of a unit under test.

The 'accident in the lab' I refer to occurred when I inadvertently tried to test a design on my IC-7300 but forgot to switch off my commercially-built WIMO box. Noise dropped a bit and suggested that, with proper phasing between stages, it might be worth trying. I did try it out by deliberately building 2 circuits in series (with some phasing adjustments), but the initial results were disappointing. My current experience with these circuits is that regaining simultaneous control in 2 stages of both phase and gain is going to require more work than I've allotted to this project, but I plan to get around to it next year. I report the current body of measured results in this article.

The resulting test bed is shown in **Figure 3**. The circuit diagram of the basic JFET canceller is available in many places, and reference 3 will take you to the VK5TM website to review it. The basic MOSFET design is shown in **Figure 4**. By now, the input phasing circuit should be familiar to anyone who's investigated these designs. In Figure 4, the primary winding of T1 consists of 4 turns, and the secondary of 16 turns, of #28 AWG enamelled wire on an FT50-43 core. It's best to centre the primary winding over the middle of the secondary winding. The primary winding of T2 has 10 turns, and the secondary 2 turns, of #28 AWG enamelled wire on an FT50-43 core. R1 and R2 adjust the phase (denoted 'Phase 1' and 'Phase 2' respectively) and R3 the gain (denoted 'Gain'). All three have linear scaling.

FIGURE 3: The GCR testbed with multiple circuits for testing (named after Giles C. Read).

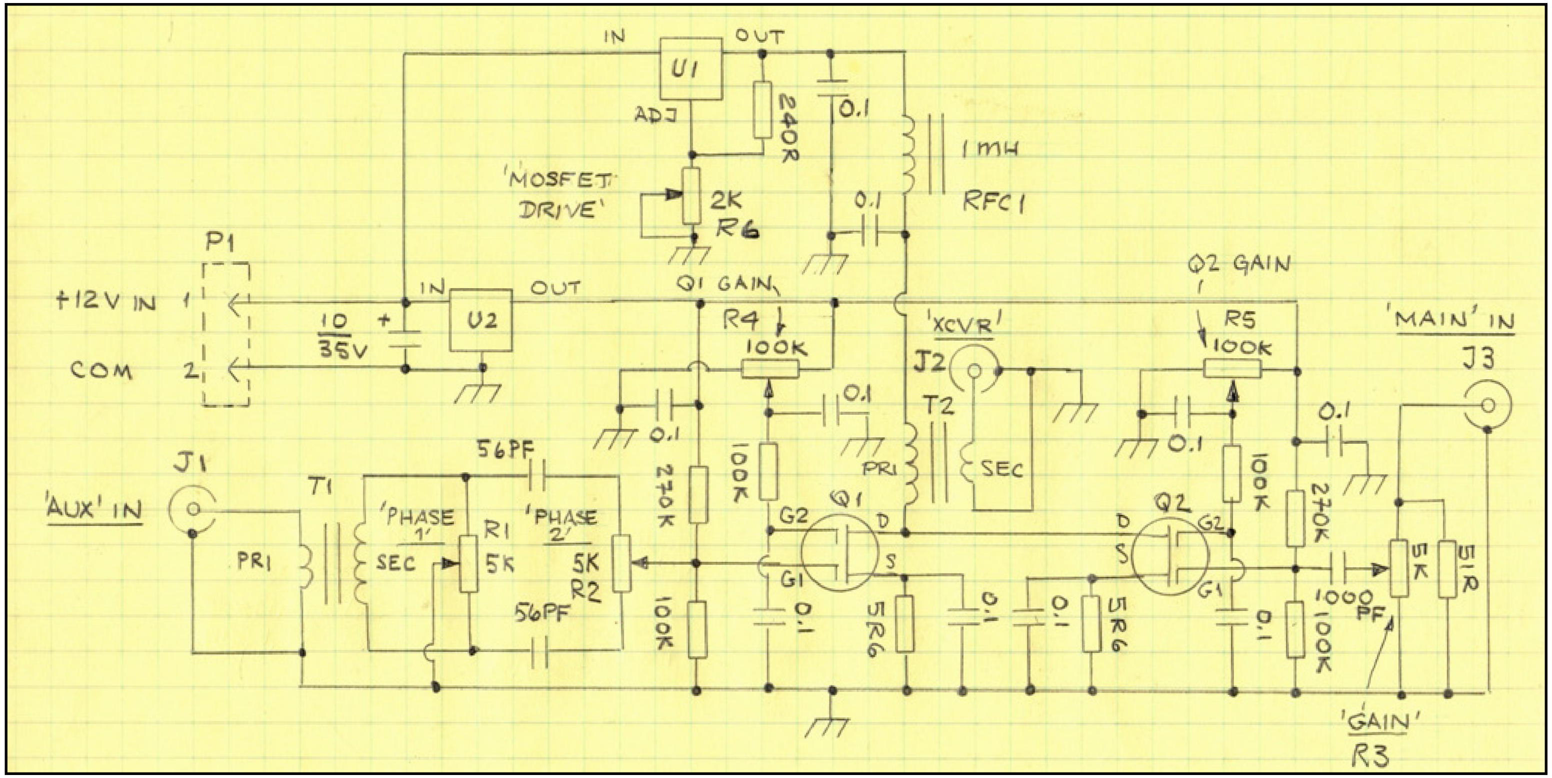

FIGURE 4: The basic dual-gate MOSFET canceller design, as tested. Q1 and Q2 are 3SK293 MOSFETs.

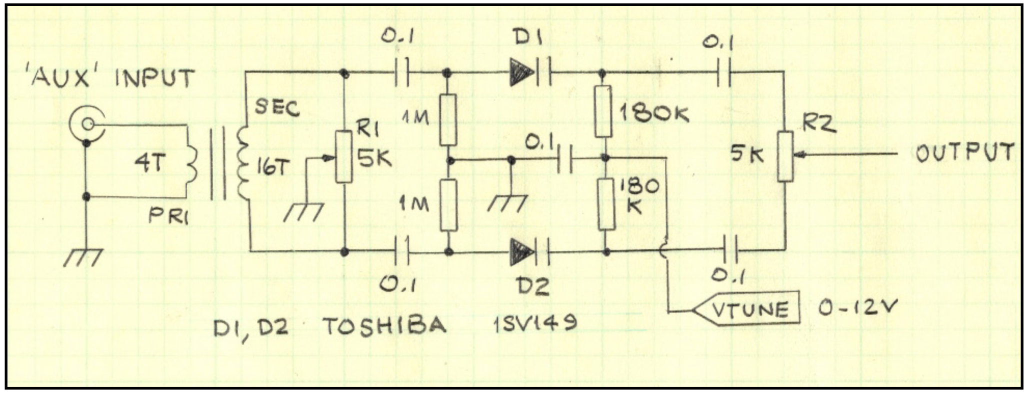

FIGURE 5: The modified phasing circuit. D1 and D2 are Toshiba 1SV149 units .

R4, R5 and R6 are trim-pots. While it's possible to use resistor dividers to get the adjustable drain and gain biases, I find it easier and more reliable to use simple low-power LM317L regulators, U1 and U2, as shown. I tried working without the output transformer, T2, but the circuit worked much better with it, suggesting how important the impedance at that point is when compared to the JFET design. I generally set the gate 2, or gain voltage, on each stage initially at 2V, and then make very small adjustments once I connect the antennas. R1, R2, and R3 then become the now-familiar 'Phase 1,' 'Phase 2' and 'Gain' controls.

One of the first honestly-useful modifications I made (even before the preamplifiers) was to provide variable phasing capacitance for the 'Aux' input. **Figure 5** shows the modified phasing circuit that I've been using, successfully, so far. In the figure, R1 is the 'Phase 1' control, and R2 is the 'Phase 2' control. There are 4 turns on the primary of T1, 16 turns on the secondary, all of #28 AWG enamelled wire on an FT50-43 core. I follow the winding procedure described in a previous publication **[4]**. The potentiometers are linear and can be of low wattage. This phasing unit again incorporates 1SV149 varactors diodes (D1 and D2) to allow the phasing capacitance to be varied, something that I found very useful to extend the frequency range of the overall design. In all my R&D work, I used prototype PCBs with copper foil tape as a ground plane and PC-pin potentiometers that fit the 0.10-inch spacing between holes. The PCBs and potentiometers were purchased from Amazon. Trimmer capacitors for the preamplifiers were found on eBay, as were the J310 JFETs.
A note about those transistors: since they're no longer in production, there are a lot of 'not quite up to scratch' J310s around, so it's a good idea to buy a few from various sources and test them. I use the PEAK Atlas DCA75 to do the testing **[5]**. It helps sort out units that might cause problems by allowing me to compare measured results of tests and the component specifications. But, having said that, I've even tried the 'duds' in the basic JFET unit, and they do actually work, just not quite as well. For those who aren't quite so adventurous, the surface-mount version, MMBFJ310LT1G, is still in production, and is available (but fairly inexpensive), and can always be mounted on an adapter PCB just as I've done with the MOSFET. I've also tested J211 (still in stock in some locations), BF245A, and 2N5485 transistors I've discovered on eBay. The 3SK293

and 1SV149 units I've found on eBay are proper production units, in my experience, and I haven't encountered any 'duds' of those types.

THE RESULTS OF TESTING

I first tested noise extinction, in several parts. First, the 'Annie' noise generator [4] was used with the low-impedance output split with a tee-piece to go into both 'Aux' and 'Main' inputs on a unit. I then made adjustments, first for maximum throughput (as measured on my Rigol DSA-815TG spectrum analyser), and then readjusted for best extinction at 4MHz, 14MHz and 28MHz. For some configurations, I then repeated these measurements using the preamplifiers described above. The units, built as 1- or 2-stage devices, were tested first as 1-stage, then 2-stage, and then again as 1-stage units with the preamplifiers. **Table 2** shows the results of the measurements.

Next, a more-complex set of tests was performed to examine the range of extinction with the same inputs as above. This was also performed by measuring the gain of the 'Aux' and 'Main' channels by using the tracking generator (TG) on the DSA-815TG as input, and measuring the output, one channel at a time. This test verified the capability of a design to match both phase and amplitude, and to uncover insertion losses and measure gain. I consider this test a 'health check' for a particular design. This turned out to be an interesting set of data. In particular, looking at the post-adjustment (for 'best extinction') curves in **Figure 6** and **Figure 7**, it becomes clear how closely the amplitudes of the two channels need to be matched for good extinction. Also, comparing those curves with the 'TG input' curves when everything is adjusted for 'Max throughput' for each channel one can see how much signal amplitude can get sacrificed along the way. Remember, these curves use the TG on only one channel at a time. The curves marked 'Annie input' split the Annie noise generator output and apply it to both 'Main' and 'Aux' inputs and these curves show the output at the 'XCVR' connector.

In my final test, I used the multi-component or 'MC' noise test to check response and extinction when the phasing and gain of non-preamplified units were tuned to a single frequency. On many lab-type signal generators, there are usually several 'pseudo-noise' waveforms that the generator 'constructs' based on user inputs. I found this useful because it produces an output that has a wide bandwidth, like random noise, but has closely spaced peaks that allow the experimenter to look at what's happening between the peaks. Spacing was set at 250kHz, amplitude under 100mVRMS, and AM modulation at 100% at 1000Hz. Extinction was adjusted while observing the output on the spectrum analyser, after which the 'Annie' noise generator was again substituted (into both inputs, without any additional adjustments) to show the response of the extinction curves to more random, broadband noise.

In Table 2, the simple 'noise extinction' test was performed, again observing output on my spectrum analyser. To achieve the extinctions values shown, I used the following adjustment

Configuration	Input conditions	Best extinction (dB) 4MHz	14MHz	28MHz
JFET, simple design	Raw	-42.5	-31.4	-20.3
	Preamplified	-60.3	-49.8	-38.1
MOSFET, simple design	Raw	-26.8	-29.9	-14.9
	Preamplified	-60.8	-50.3	-40.6
JFET, 1 of 2 stages	Raw	-32.9	-18.7	-13.0
	Preamplified	-55.9	-47.9	-36.3
JFET, 2 stages	Raw	-32.9	-17.5	-11.1
MOSFET, 1 of 2 stages	Raw	-26.7	-33.9	-33.0
	Preamplified	-50.0	-46.8	-44,1
MOSFET, 2 stages	Raw	-28.4	-23.4	-22.5

TABLE 2: Results from using the 'Annie' noise generator, measured by my spectrum analyser.

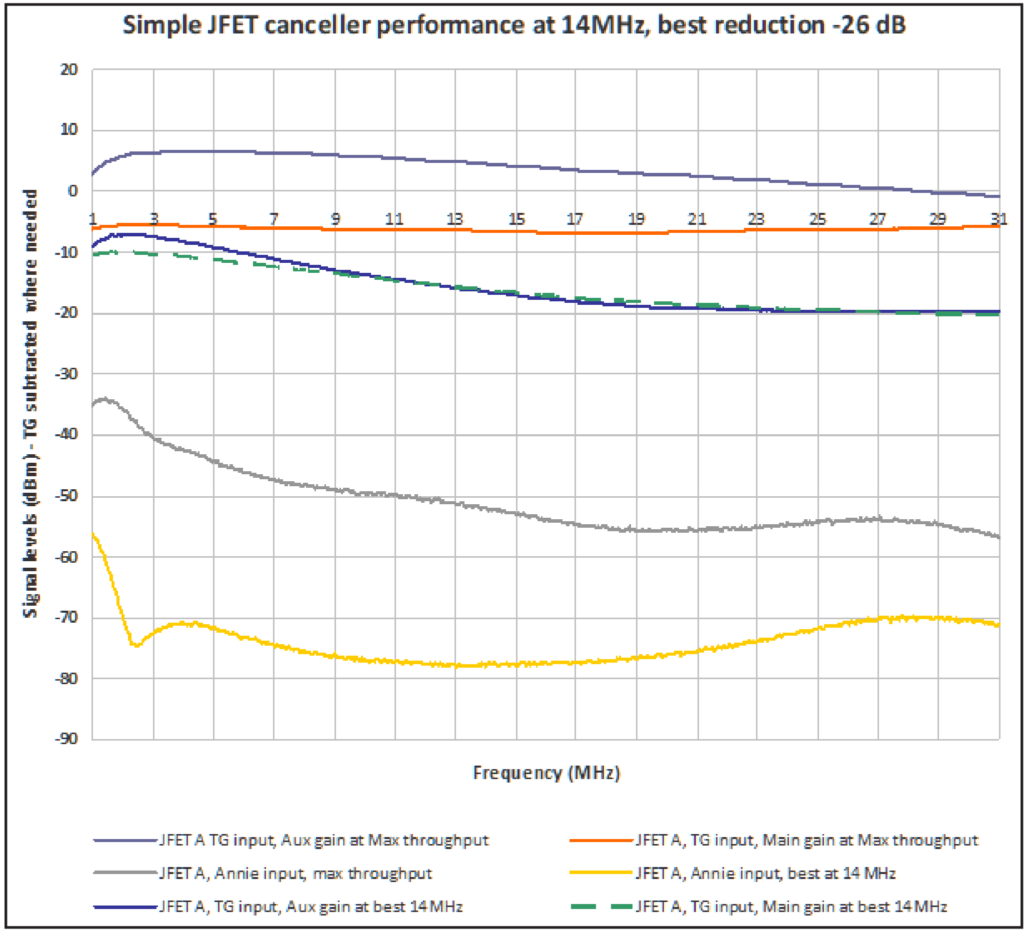

FIGURE 6: Circuit adjusted for optimum rejection at 14MHz: the overall performance of a simple JFET unit.

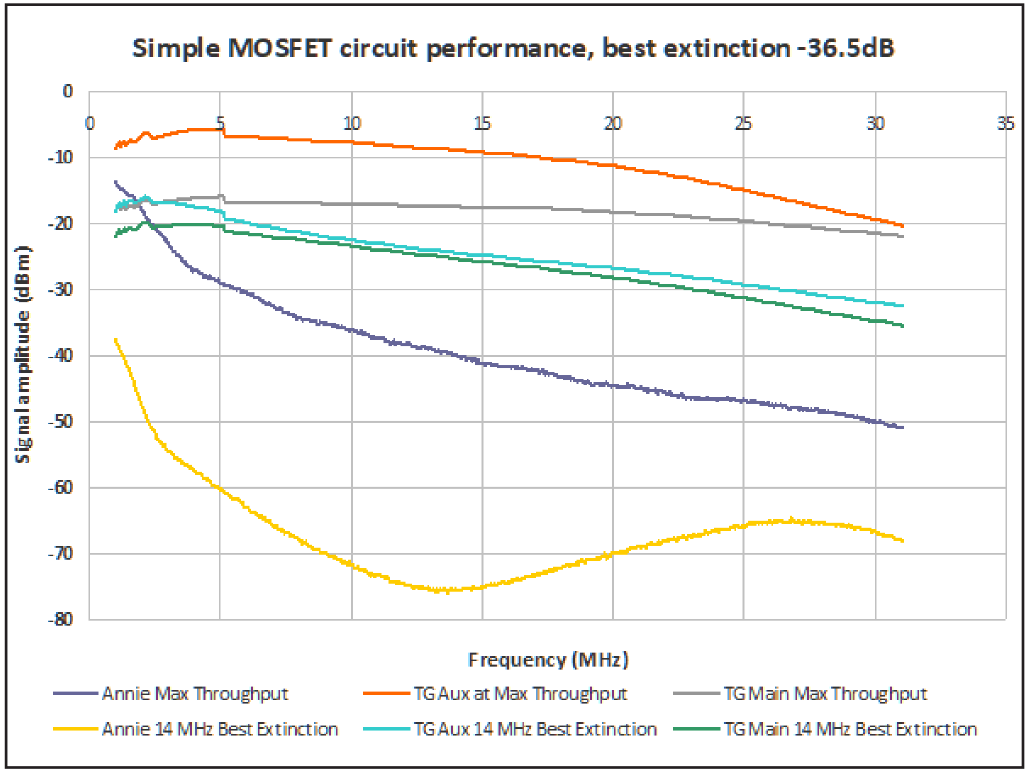

FIGURE 7: Circuit adjusted for optimum rejection at 14MHz: the overall performance of a simple MOSFET unit.

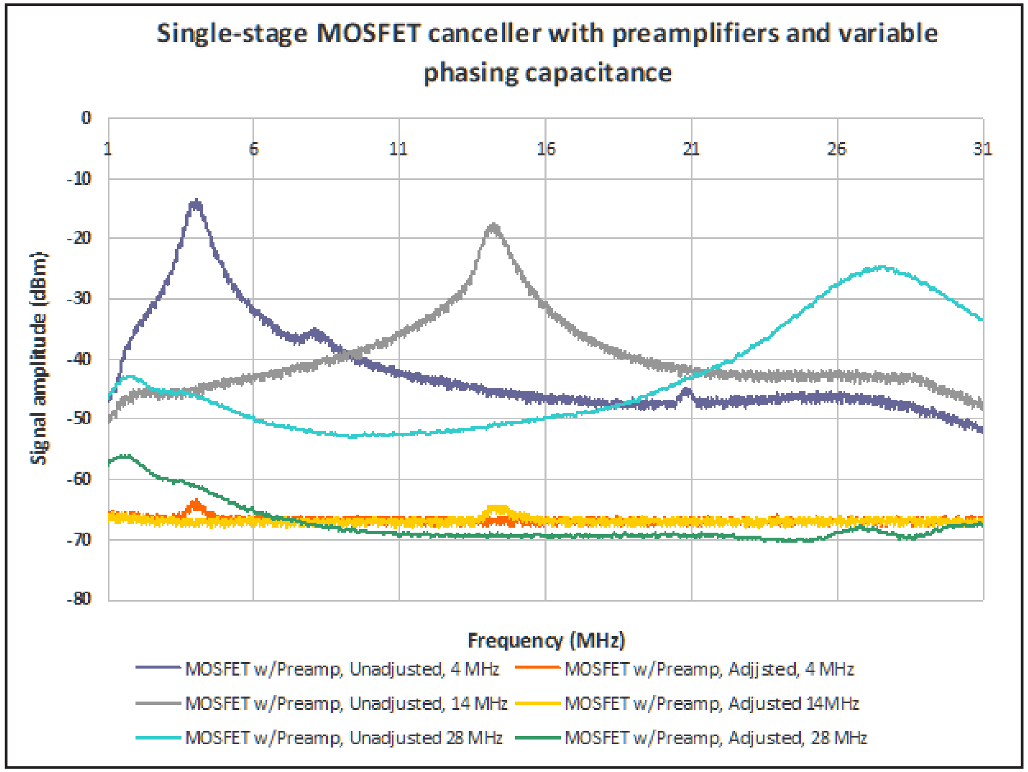

FIGURE 8: Extinction-test data for the MOSFET design (Figure 4) with preamplification.

procedure for each test:

- set all controls at 50% at the start and adjust ONE control at a time;
- adjust 'Phase 1' for the best overall null;
- adjust 'Phase 2' for improvement in the null at the frequency of interest;
- adjust the 'Gain' control for the lowest noise signal output at frequency of interest;
- repeat these steps until no further improvement is seen.

I've found that this is a reasonable way to adjust R-C phase shifters in practice, and have been using it for years. It's sufficient for the normal S3-S5 background noise here at N6JJA but, when it's hot and the Sun is intense, the background noise (probably because of local solar-panel inverters) blasts out at about S9. The results in Table 2 were obtained using this procedure and are for a 'Raw' configuration (equal 'Annie' noise applied to both ports) or 'Preamplified' (same noise, but preamplifiers added in front of the canceller).

The results indicate a number of things. First, the simple JFET design does overall very well in its 'Raw' configuration. Like most of the designs tested, its performance at 28MHz isn't as good as it is elsewhere, and that may not be important for most man-made noise which tends to decrease in amplitude with frequency.

Next, having a second stage did not work as well as hoped, perhaps because of additional phasing mismatches that can creep in when you have multiple stages, or requiring both stages to share the same 'Aux' input. I spent considerable time looking for other weaknesses, and I found a few, but nothing that added a lot to its performance. Nevertheless, the second stage also added complexity that may not be worthwhile, but it will take time to sort that out.

That being said, the addition of the tuneable phasing capacitance makes a big difference, especially at lower frequencies. Often, a null at 4MHz (or a better null at 28MHz) could be made, but results seemed disappointing until the phasing capacitance was adjusted. I'd add here that, as a direct observation, the losses in the R-C phasing circuitry can be reduced drastically by choosing the best phasing capacitance.

The MOSFET design in Figure 4 needed further iterative tweaking of the gain controls of each transistor to provide better results than tabulated here. This was learned in subsequent tests. But it was definitely the winner when the preamplifiers were added.

For all of the preamplified results, performance at 80m is definitely improved, and the depth of the nulls come with considerable bandwidth. **Figure 8** shows a typical set of data taken with the preamplifier pair.

LESSONS LEARNED FROM THESE TESTS

The simplest building blocks are the best and easiest to implement and adjust, but adding the variable phasing capacitance helps. Fixed capacitors can always be used with a selector switch if the varactor-diode solution isn't available. The lowest capacitance I use is about 25pF, and the highest probably around 300-400pF.

I won't catalogue all that I learned from the 2-stage designs here; it's still a work in progress, and I am still learning. Please send me an email if you would like more information. The preamplification turns out to help a lot. Additional phase shifting caused by the preamplifier input circuits arrives equally (more or less) at both inputs and thus shouldn't do a lot of harm. The data presented in Figure 8 show that the preamplifier actually prevents a lot of noise from even reaching the canceller circuits. Since most man-made noise has higher amplitude at lower frequencies, it makes sense to use a high-pass input to the preamplifier as I've done here. A better filter design at 28MHz might help further, but I doubt that the extinction will be much different. Also, please remember that under each peak is both signal and noise! That's just a general warning when interpreting test results; it's typically possible to get rid of all the noise when lowering the strength of the desired signal too much. And do note how flat the amplitude lines are, away from the peaks, when adjusted. And, just to be clear, all the controls required adjustments for best performance. Having a gain control in the 'Main' input's preamplifier isn't enough to do everything required. More on that in the next section.

The 'Adjusted' curves in Figure 8 are also close to the local noise floor on the spectrum analyser. In practice, it is necessary to adjust the controls so as not also to eliminate too much of the desired signal, as I've noted above. In the 'Preamplified' configuration, this can happen more easily than in the 'Raw' configurations.

CIRCUIT PERFORMANCE TESTS DURING CANCELLATION

Figures 6 and 7 are used for this discussion. They represent testing on the 'simple' JFET and MOSFET circuits, and reveal an interesting difference.

The data shown in Figure 7 show something important that I discovered during these tests. First, it's evident that, at the frequency of interest, the signal amplitudes, here the gain from the TG data, are nearly identical. Focus on the curves in which gain is measured using the TG, first for 'Maximum throughput' (highest gain), then for 'best extinction'. It is clear that, without further adjustments, the gains in the 'Main' and 'Aux' channels are quite different. But the pair for 'best extinction' are very close together at the frequency of interest, a requirement of the extinction processes, ie they are of roughly equal amplitude. Note that the same close alignment of the TG-measured gain is seen in Figure 6. The best cancellation occurs when phase and amplitude are matched closely, a point that will be expanded upon in a future article. As a brief preview of that analysis, it turns out that random processes like noise are tough to cancel very well except as an aggregate of unwanted RF amplitude/phase 'signals.' But, in Figure 7 you see a 'typical' broad null, not far different from that in Figure 6 but deeper in the middle. Curves more like Figure 6 were something I saw in most of the data, and the curves that resulted in the extinction data shown for this circuit in Table 2 are based on that. Table 2 data for the simple MOSFET circuit is -29.9dB, but in Figure 7 it is -36.5dB. The phasing is better than it was in the data taken for Table 2. This seemed a puzzle, so I turned to the last test in the series and found an answer. However, data for all units shows similar characteristics, as they should, and this is again confirmation of the 'mathematics of cancellation.' I hope to address that topic in a future article. This test really only addresses a basic question that should be asked about any new design I try. If I can't get a null that deals with matching both phase and amplitude, I need to try something different. All the R-C circuits passed this test.

NARROW-BAND ADJUSTMENTS AND NOISE EXTINCTION

There's no test already in use that I'm aware of that allows one to compare noise canceller circuits, so I decided to develop my own. This is a test that uses my 2-channel signal generator (Siglent SDG 2122X) set to produce a pseudo-arbitrary waveform called 'MC noise' at a base frequency of 250 kHz, AM 100% at 1kHz, and an amplitude under 100mVRMS. **Figure 9** shows an average of 100 scans for the MC noise output from one channel of my signal generator. The analyser settings for these and subsequent data use a resolution band width (RBW) of 10kHz, 10dB of input attenuation, the internal

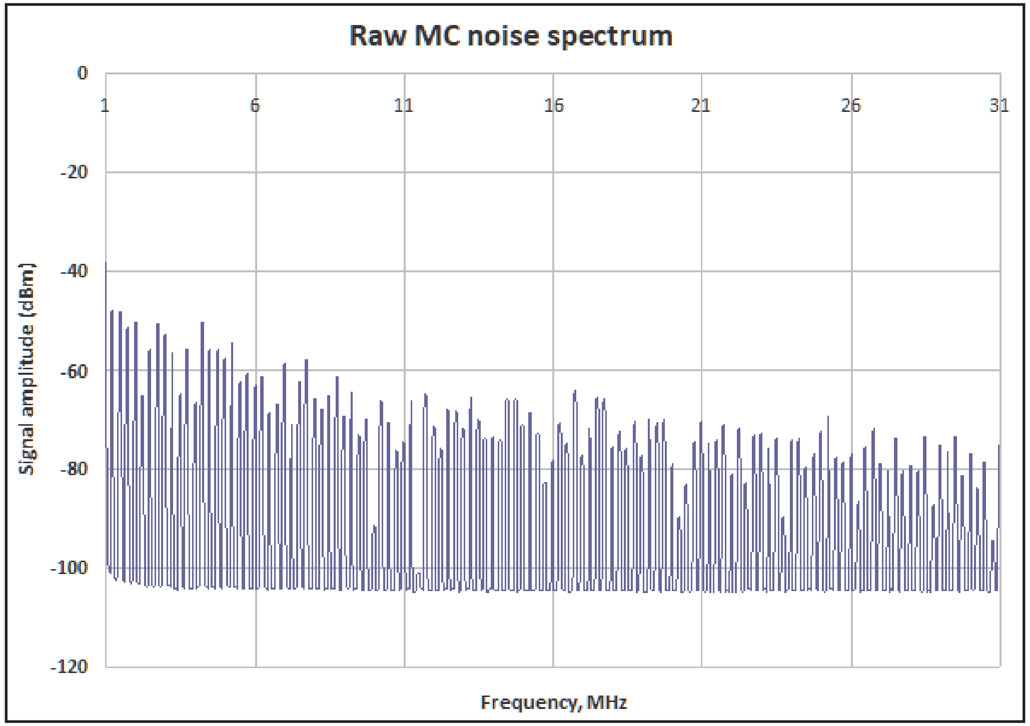

FIGURE 9: The amplitude envelope for an average of 100 scans of MC noise.

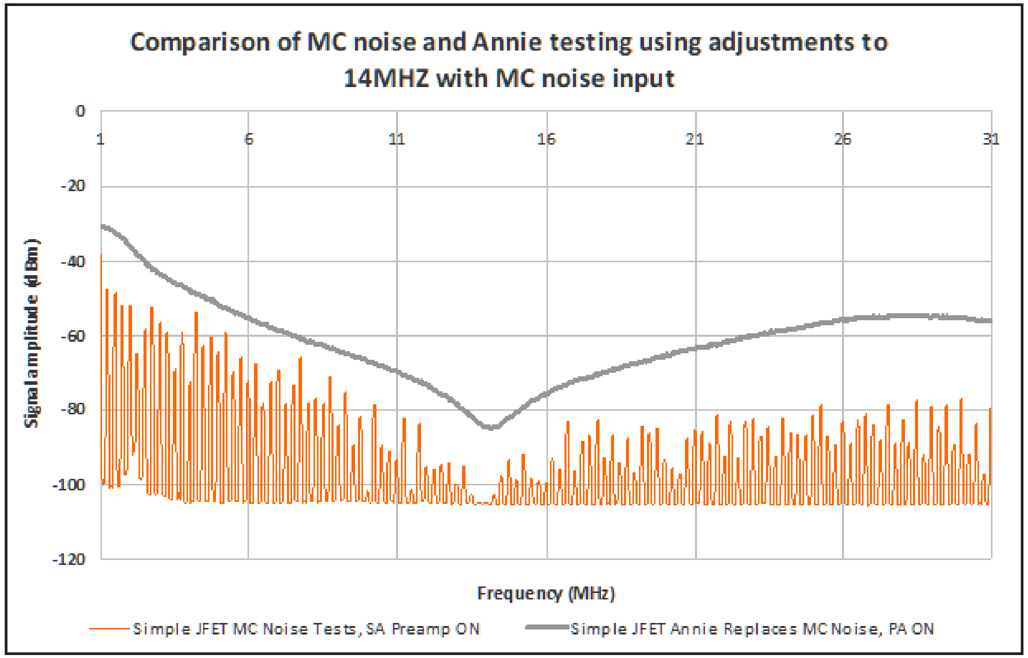

FIGURE 10: MC and Annie noise nulling at 14MHz for the simple JFET circuit.

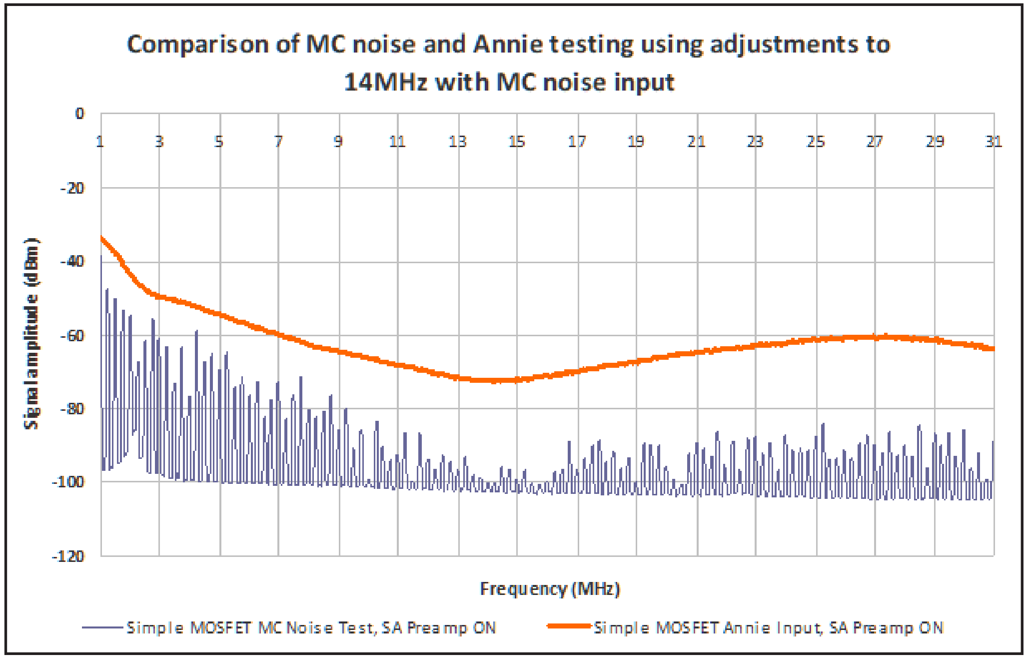

FIGURE 11: MC and Annie noise nulling at 14MHz for the simple MOSFET circuit.

preamplifier switched on, and a power average of 100 scans.

The thing that I learned to appreciate about this waveform is the mixture of random and distinct (and persistent/repetitive) signals. Note that the low points at about 10MHz and 12MHz are persistent artefacts of the noise generated by the arbitrary waveform generator in my SDG-2122X. They can be made closer together by lowering the base frequency in the signal generator. I then used the usual adjustment procedure I listed above to attempt to null the available signal in the scan shown in Figure 9 at 14MHz. Having used this 'MC noise' to obtain a null, I then connected the 'Annie' noise source and rescanned (without further adjustment) to see what the noise extinction looked like. **Figure 10** and **Figure 11** show the results.

Here's where I learned one of the most valuable lessons I have encountered in how to adjust these devices. In Figure 10 and Figure 11 there's something troubling, to me, at least. While I, and I suppose many others, use the adjustment procedure I listed above, the data here suggest a problem: simply, it's possible to adjust a canceller to remove much more desired signal than you'd like. The 'MC noise' signals are both discrete peaks while covering a broad spectral range. It's in the way these nulls overlay, between the discrete 'MC noise' and the broadband 'Annie' noise. Making any of those discrete peaks (that might also contain the signal from some desirable DX station) go away at the same place the noise is nulled, might not be avoidable, but it isn't what I want.

So back to work, first exploring what each knob actually does.

THE KNOB ROSETTA STONE

After doing the hundreds of painstaking adjustments for measurements as shown in the figures above, one begins to get a feel for which knob does what. Watching the results roll by on the spectrum analyser instilled in me a real appreciation for what these circuits actually can do. So let me cut to the chase. I developed my own way of making the necessary adjustments and this led me to the data shown in **Figure 12**.

Comparing this 14MHz data with the extinction levels listed in Table 2 for that frequency, it's clear that something has changed. And these data are without preamplifiers and their narrowed bandwidth! The curves aren't as broad as one can obtain with the old 1-2-3 adjustment procedures, but a lot more is happening. Without the extra pages of data and discussion, let me boil it down thus:

THE NEW ADJUSTMENT PROCEDURE

What the knobs actually control:

(a) 'Phase 1' primarily controls the frequency at which a null can occur.

(b) 'Phase 2' is indeed a phase control, actually adjusting the depth of the null.

(c) 'Gain' is the amplitude matching control, also moving the frequency of a null around a bit as 'Gain' is adjusted. The DC-blocking capacitor at the transistor input (see **Figure 13**) on the main side of the adder along with the 'Gain' control represent the other and necessary phase shift network.

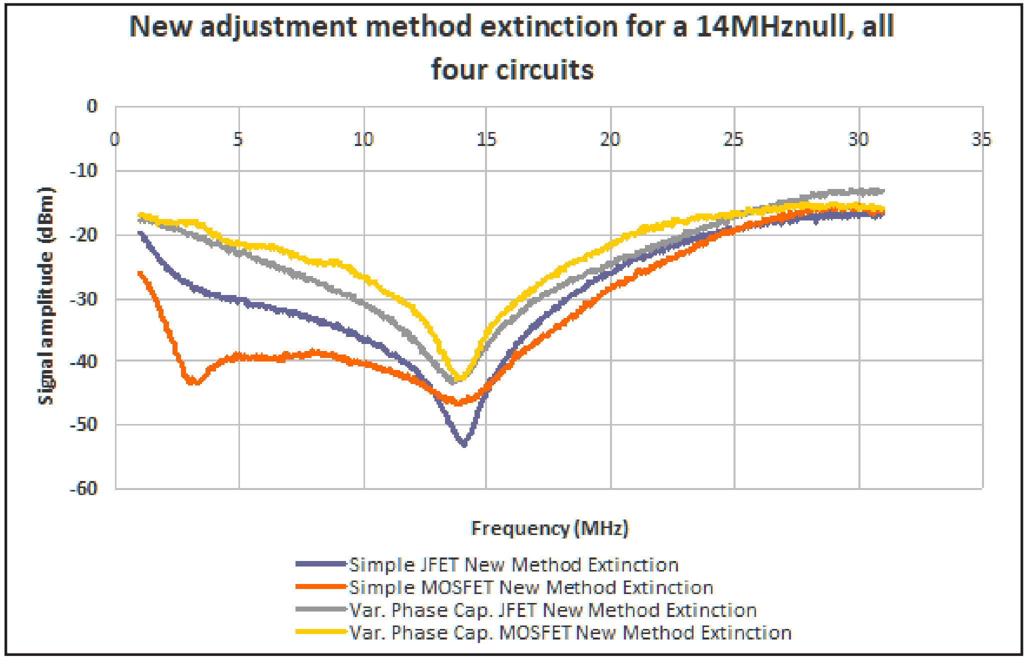

FIGURE 12. 'New method' extinction curves.

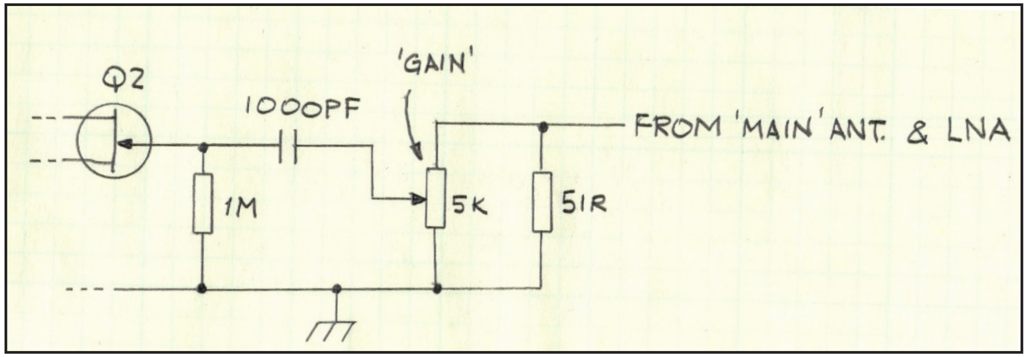

FIGURE 13: 'Optional' input circuit from the 'Main' antenna. Adding the 1000 pF capacitor allows better adjustment of the phase in the 'Main' channel of the canceller.

How to adjust to get curves in Figure 12:

(a) Start, as always, with all controls at their 50% positions, and adjust ONLY one control at a time!

(b) First adjust 'Phase 2', NOT 'Phase 1'! Some reduction in noise will occur, often with a null somewhere in the frequency range of interest. You may or may not see it or hear it, but it's there, waiting to be found. A small improvement is all you need at this point.

(c) If you have a variable phasing capacitance control, use it here. Higher capacitance for lower frequencies, etc. Some additional reduction may be seen, and it might only be small, but don't lose patience.

(d) Next, adjust 'Phase 1' to move the frequency of the null to the frequency of interest, evident when more reduction takes place.

(e) Then adjust the 'Gain' knob to reduce noise further. Remember, adjustments to this knob can show a null, beyond which noise rises again. (As noted, there's an extra component in my circuits that helps produce that null.)

(f) Use 'Phase 1' and 'Gain' to bring the null to the right frequency while adjusting the amplitude matching.

(g) Return to 'Phase 2' to then attack phase discrepancies.

(h) Repeat c-g until no more reduction can be seen.

(i) Remember, always adjust 'Phase 1', at least a bit, every time you adjust 'Gain.'

It's not more complicated than the procedure I outlined earlier, it's just more careful. It works in the lab with my test equipment and on the air. As I write this we're in an S9 noise day and my WIMO box and its own 'Aux' antenna can't cut it. I used my 40-foot 'Aux' antenna and could kill the noise (easily to S0) with all 4 circuit types using this procedure. Granted the waterfall display on my IC-7300 helps a lot, but your ears are the final judge. I was somewhat surprised (though I shouldn't have been) how much phase/frequency influence the 'Gain' control has in these R-C configurations. But look at Figure 13, drawn for the JFET-based design I am using. Most circuits use a similar input signal level control on the 'Main' side of the adder, but mostly without the 1nF capacitor. But it's pretty clear that the 'Gain' control along with the combined capacitances that follow (1nF and JFET input capacitance) create one more frequency-dependent phase control. This is the extra phase adjustment that makes this circuit work as it does. I added a switch to that 1nF capacitor on the 'standard' JFET canceller used in the tests described here to observe the difference in performance directly using the TG setup to measure extinction. With the 1nF capacitor, nulls were more easily found and were about 10dB deeper. I didn't do extensive testing for this observation, but it confirmed my collected data showing that it's a very useful piece of the whole canceller puzzle.

And as I did the on-the-air tests with the new procedure, I saw less degradation in the desired signals. Why is that possible? In Figure 10 you see what happens using the old method of adjustment. But aligning a circuit to null the 'Annie' noise input using the new method, then comparing it with the resulting 'MC noise' null, results in the data shown in **Figure 14**.

The discrete signal signature of the 'MC noise' gets through while the broadband noise nulls out. This is essentially the kind of result I found when doing the on-the-air tests. And in those tests, the preamplifier really made things work so much better. I found that, once the frequency was tuned in (signals get louder), starting at moderately low gain on both channels allowed the procedure above to just use the preamplifier gain controls (one at a time!) to make the waterfall display go completely dark except for the stations transmitting. Keep the gain as low as possible, by the way. Adding or reducing gain on the 'Aux' antenna input acts to lengthen or shorten that antenna to help in the cancellation process.

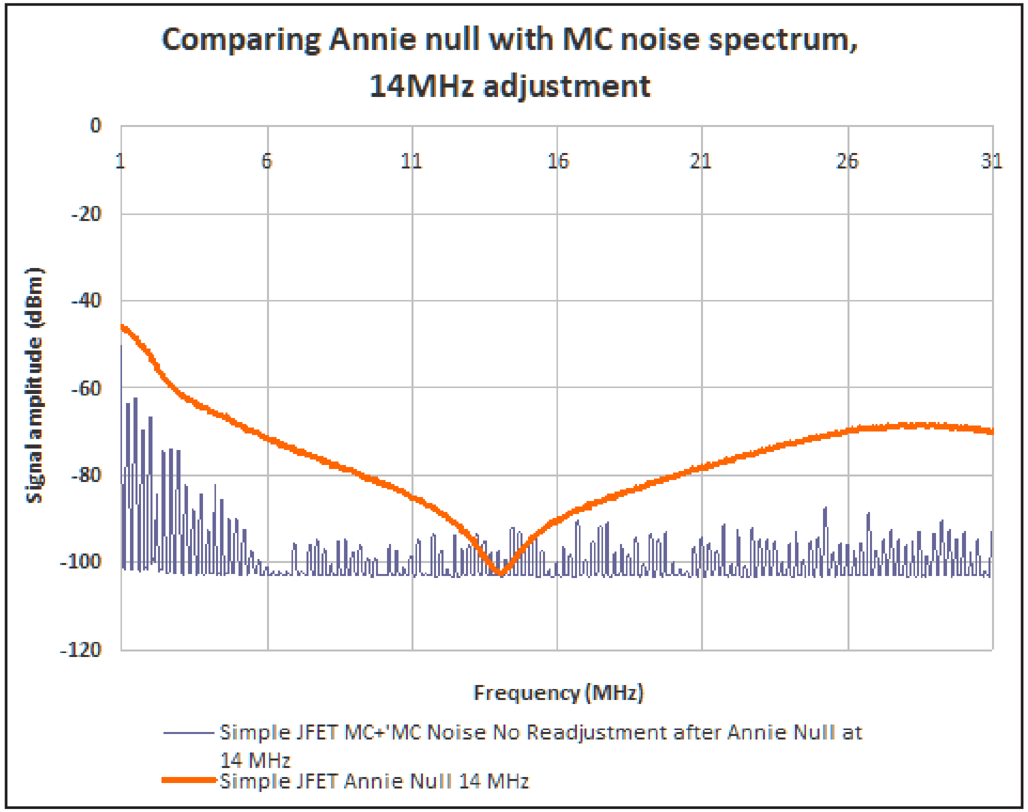

FIGURE 14: The new adjustment method and a new result.

commercial designs that have been, for so long, 'almost there.' This has been a long project for a retired guy, but it is far from over.

WHAT I FINALLY BUILT

For personal reasons, based on my experiences of doing this project, I built my noise canceller as a JFET type, along with a pair of the preamplifiers per Figure 2. **Figure 15** and **Figure 16** show their circuit diagrams. **Table 3** and **Table 4** list their parts. All relay positions shown are for the un-energized (no switching voltage applied) condition. In the world of relays, 'NC' refers to the 'normally closed' or un-energized position, while 'NO' refers to the 'normally open' position that only gets closed when voltage is applied to activate the relay.

My goal was to build, test, and use this unit to continue to tweak the circuitry, then report my results again in a *RadCom* article update. The final product is shown in **Figure 17**. What remains of the circuitry to be described is the bypass and signal relays that I've added to handle the 100+W of RF from my transceiver and allow me to experiment with how I use the preamplifiers. Relay K1 in Figure 15(b) won't be enough for the transceiver output, but is part of the canceller circuit here to protect its input and output during transmissions. A more robust relay switches the main RF power.

ON-AIR TESTING

For me this was anticlimactic. It turned out that I have built and tested this circuitry so many times over the past years that about all I can say is that 'it works as advertised.' For any engineer, experimenter or enthusiast, that's about as good as it gets. I noted above that I can bring S9 noise back to about S0 with this setup, and still hear signals, and that's indeed what I find in the on-air work. One thing about signal reports. I find that my S-meter underreports received signals by as much as 2 S-units, depending on the amount of background noise that needs to be removed. I try to inform contacts that my readings are estimates, but that's what happens when you're

RECOMMENDATIONS

First, try the newer adjustment procedure, if you own any type of R-C style noise canceller. Be patient and adjust carefully. In my experience, that may work well enough to make all the difference you need. I hope to show, in a future article, how to address issues to help you improve your 'Aux' antenna. If you already own an R-C canceller, consider building and adding the preamplifiers, if you're up to that small challenge. But, if you're up for a greater challenge, I'd build the JFET circuit with the added variable phasing capacitance and the preamplifiers, knowing that your success depends on your 'Aux' antenna. For me, I'm going to tinker with the two circuits, JFET and MOSFET, both with preamplifiers, perhaps using both in my shack as the mood strikes me. I've learned enough about them to understand their strengths and weaknesses.

At some point, if one learns anything at all in a project, the time comes to report what you have learned up to that point. I learned a great deal from this work and look forward to seeing it reflected in the designs that will bring credit to many others, perhaps as well as improving the

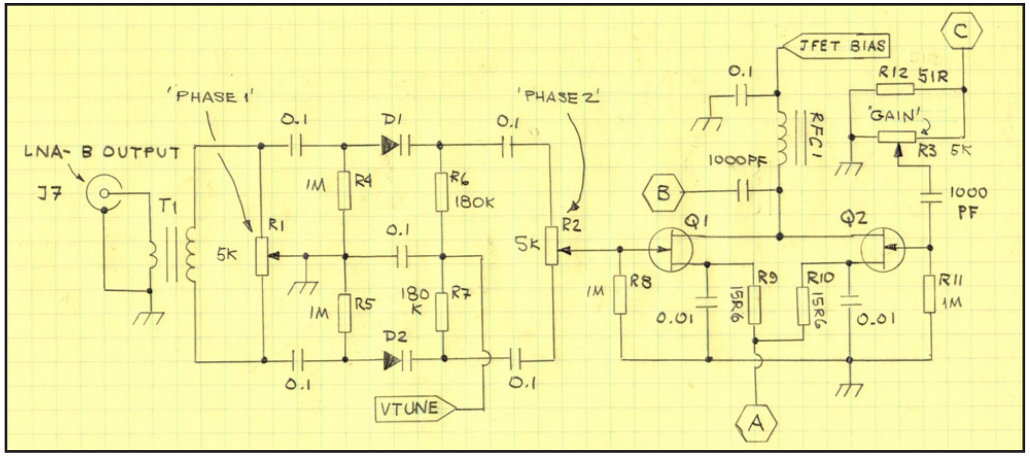

FIGURE 15(A): The noise canceller circuit of the as-built unit at N6JJA. Q1 and Q2 are MMBFJ310 JFETs. Points 'A', 'B', and 'C' refer to connections in Figure 15(b).

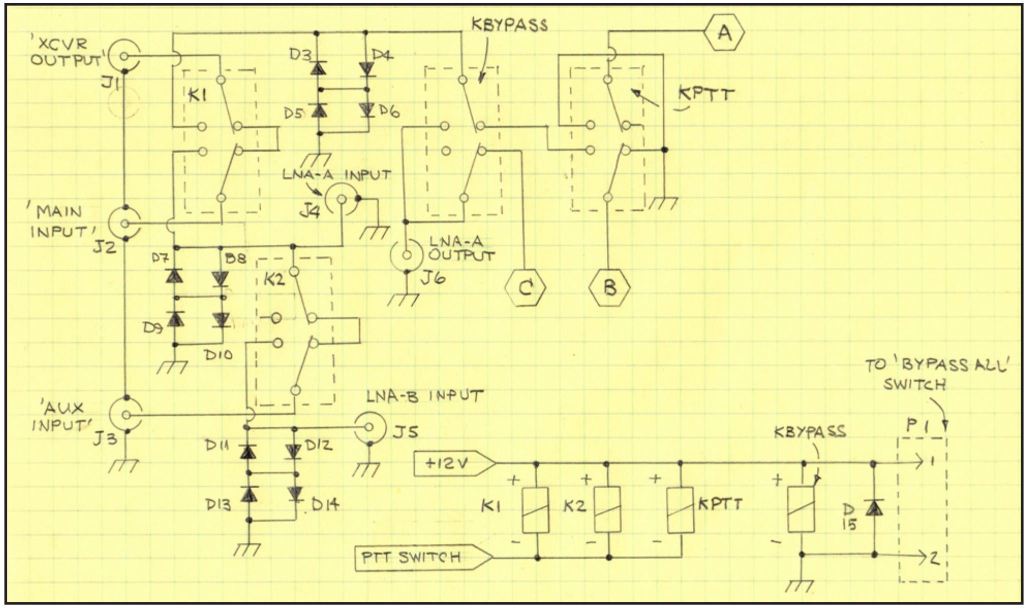

FIGURE 15(B): Routing and control circuits for the new JFET canceller. K1 is an Axicom V23105A5003A201 DPDT 3A 12-volt relay. All other relays are Panasonic TQ2-12V units. D3-D14 are 1N4148 or equivalent. J1-J3 are SO-239 UHF connectors. P1 connects to a panel switch that allows me to bypass the canceller.

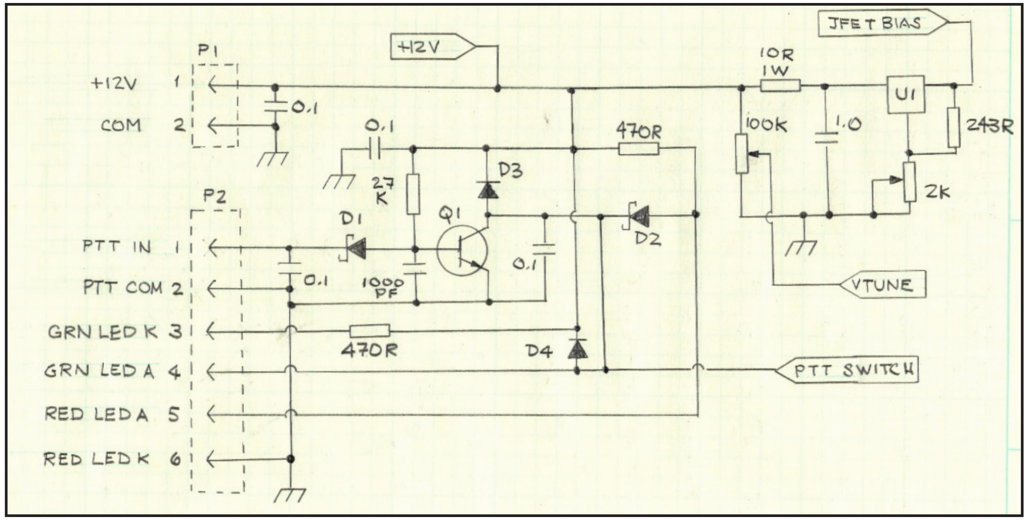

FIGURE 16: The noise-canceller 'Push-to-Talk' (PTT) circuits. Q1 is a BC337-40 transistor. D1 and D2 are BAT85S Schottky diodes. D3 and D4 are 1N4001 rectifiers. U1 is an LM317L regulator, and the 2K trim-pot should be adjusted to deliver approximately 8 volts to the 'JFET BIAS' line. The two 470W resistors operate the LEDs on the front panel that have approximately 2.1 volts of forward drop. The green LED indicates that the receiver is active, the red LED lights during transmission.

fighting noise. I'll also repeat that the waterfall display on my Icom 7300 is a great addition to how well I'm able to zero in on the kind of noise reduction I report here. I'd suggest that if you're working with a rig that lacks that kind of display, and you'd like something more than just your ears to guide adjustments, having a small SDR 'dongle' that can be switched in instead of a transceiver can help with adjustments. Another caveat I'd add is that, when you employ additional preamplification to the input of a receiver, you're increasing the likelihood of reducing the dynamic range of the receiver. I've found that I use only modest amounts of gain for the 'Main' preamplifier. I always find that the 'Aux' gain control winds up being set significantly higher than the 'Main.' I guess I could increase the length of the 'Aux' antenna, but for the time being I'll just turn the knob.

I'm not going to offer any more operational descriptions of this circuitry here. Why? First, because, like all the other designs out there, every installation is different, and each pair of ears can hear things in different ways. In other words, noise cancellation is always subjective, to some degree. Second, I feel more comfortable answering questions about specific detail of interest to a reader than go on at great length here. My goal in this and future articles is to guide our endeavours toward more constructive and objective results. As I say, I'll be interested to hear from others, too. We all have something to offer to this important niche of our hobby.

REFERENCES

1. Andy Gilfillian G0FVI, 'An introduction to antenna noise cancellers', RadCom October 2023, pp 24-26
2. I've used the simple calculator at www.toroids.info for years. That page not only helps you wind a toroid, but also allows you to calculate, for input and output impedances, the L and C values you need.
3. https://vk5tm.com/homebrew/noisecancel/noisecancel.php
4. Sheldon Hutchison, 'Annie, when what you want it noise,' RadCom, November 2022, pp 41-47
5. https://www.peakelec.co.uk/acatalog/dca75-dca-pro-semiconductor-analyser.html

Part	Value	Notes
Fixed resistors		Typically 1/4W, 1% or 5%
Fixed capacitors		Typically 50V ceramic, X5R or X7R
Potentiometers		Linear, 0.1W to 0.25W rating
D1, D2	Toshiba 1SV149	Source: eBay
K1	Axicom V23105A5003A201 or equivalent	3A signal relay, DPDT, non-latching
Other relays	Panasonic TQ2-12V or equivalent	Signal relays, DPDT, non-latching, 2A or less
Q1, Q2	MMBFJ310 or J310	See text about caveats on leaded parts
RFC1	1mH moulded choke	Chose a unit with a low DC resistance

TABLE 3: Parts list for the JFET canceller.

Part	Value	Notes
Fixed resistors		Typically 1/4W, 1% or 5%
Fixed capacitors		Typically 50V ceramic, X5R or X7R
Trimpot	2k, 0.25W	Single turn
'Tuning' Potentiometer	100k	Linear, 0.1 to 0.25 W rating
D1, D3	BAT85S	Schottky
BC337	BC337-40	Note pin reversal
LEDs	Red and Green	Typically 2.1-2.5V forward drop for series resistors shown

TABLE 4: Controller parts list.

FIGURE 17: The final product.

Fourier transforms in digital communications: a geometric approach

Michael J. Toia, K3MT
k3mt@arrl.net

Much of today's communication is effected digitally with binary waveforms grouped into eight-bit bytes. Their messages exist as serial bit streams, and circuits transmitting these must provide the necessary bandwidth. We can calculate the bandwidth using a Fourier transform, but unless you have an advanced understanding of mathematics, the nature of a Fourier transform may be elusive to you. Here I provide an alternative approach which does not rely on equations, but rather uses pictures: a geometric approach.

A FRAMEWORK FOR ANALYSIS

For discussion purposes, let us assume that we have a bit rate 8kbit/s. The waveform in **Figure 1** shows the binary pattern 11110000, hexadecimal F0, with the time axis divided into eight intervals of 125μs. I call these 'ticks'. Sent repeatedly on a circuit, it is a 1kHz square wave, and then its spectrum is as shown in **Figure 2**. This is familiar as a fundamental frequency of 1kHz, and its odd harmonics at intervals of 1kHz. I call the fundamental frequency the 'first harmonic'. The amplitude of any harmonic is inversely proportional to its harmonic number. For example, the amplitude of the third harmonic is one-third that of the first harmonic. Figure 2 shows harmonics only as far as the seventh. Although not shown, all the others exist as well.

Each of the 256 possible bit combinations in a byte has a similar relationship between its time waveform and frequency spectrum. Here I calculate those relationships in a gallery of simple sketches, using only plane geometry.

The square-wave pulse may also be defined by a pair of impulses (see **Figure 3**). The first, positive-going, at tick 0, is a very short current event depositing, say, a fixed charge on a capacitor (or a perfect battery). The second, at tick 4, drains all charge, returning the device to its previous state. A study of these paired impulses leads to our goal of understanding the frequency spectrum of an arbitrary bit stream. They are always interleaved, of the same strength, each positive-going pulse followed some time later by a negative-going pulse. Their amplitudes in my sketches are proportional to the charge they dump onto, or remove from, the capacitor (or perfect battery).

Now let's consider the summation of a 1kHz cosine wave and all of its harmonics. All harmonics have the same amplitude and all start at tick 0 with the same simultaneous maximum, as shown in **Figure 4**. The (infinite) sum has a surprisingly simple feature: it is zero everywhere except at eight-tick intervals. There it suddenly becomes a large, positive-going impulse! A repeating train of these at 1ms intervals has a spectrum of components at 1kHz, 2kHz, 3kHz etc., all of the same amplitude. Figure 4 shows the first five harmonics over a span of 2ms.

Figure 5 presents the waveform as 'picket fences', in both time and frequency. Either of these, along with phase information, describe the waveform of the impulses.

We now introduce phasor diagrams (**Figure 6**). In this particular form, the 1kHz cosine wave is represented as a vector spinning clockwise at 1000rpm. The vertical distance between its arrowhead and the horizontal axis is the cosine's amplitude, and the angle of the arrowhead round from the vertical axis is a phase delay. Spinning at twice that rate, it generates the second harmonic, and at thrice the rate it generates the third harmonic, and so forth.

SINGLE-PULSE SPECTRA

I now encapsulate the foregoing discussion in a new type of drawing, **Figure 7**. The cosine waves of positive impulses are represented by a set of eight phasors at the bottom, depicting the relative phasing of harmonics 1 to 8. As they're all at zero phase at tick zero, all point vertically as shown. Now consider a negative impulse at tick 1. This generates negative cosine waves, whose phasor arrows point in the reverse direction or, equivalently, they are rotated a half cycle, ie by four ticks in the phasor diagram. The first harmonic impulse position, one tick later, adds one more tick delay. Its phasor is therefore rotated clockwise by a total of five ticks. **Figure 8** shows how it goes. The phasor is shown at the bottom left, labeled 1. The second harmonic phasor likewise begins four ticks later because the impulse is negative. Additionally, it spins twice as

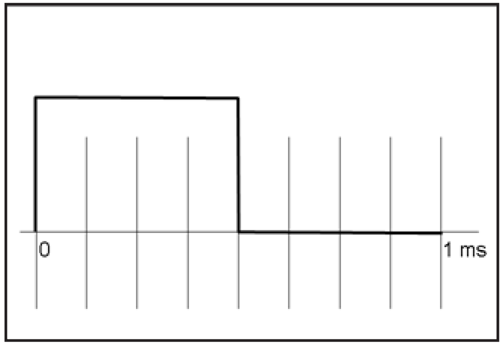

FIGURE 1: A 1kHz square-wave voltage pulse.

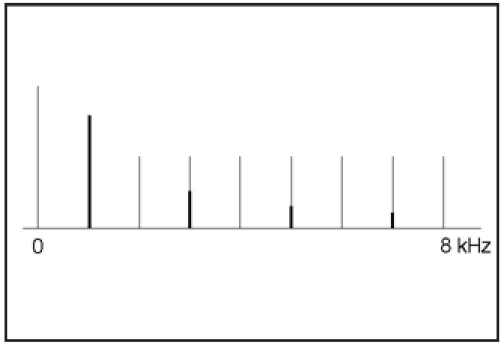

FIGURE 2: The square-wave amplitude spectrum (without any phase information).

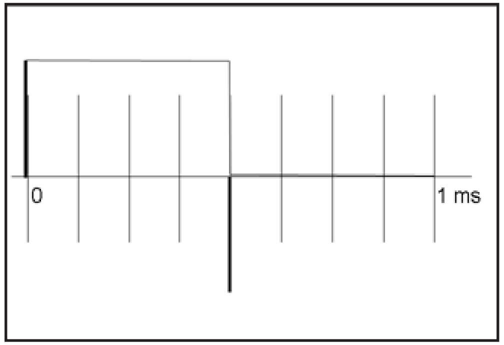

FIGURE 3: A pair of impulses creates the square wave.

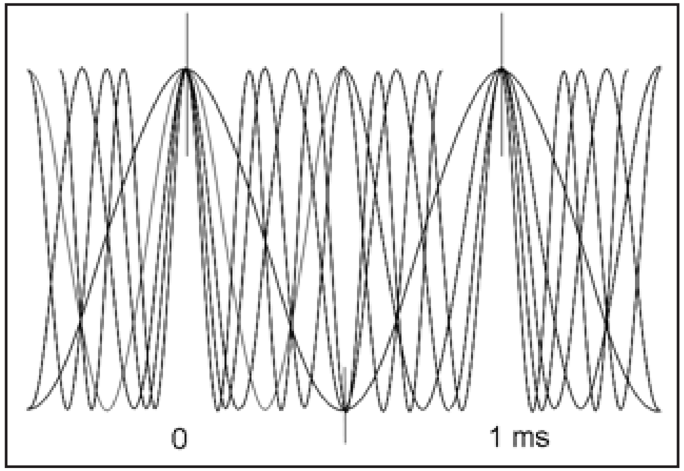

FIGURE 4: The spectral components of an impulse.

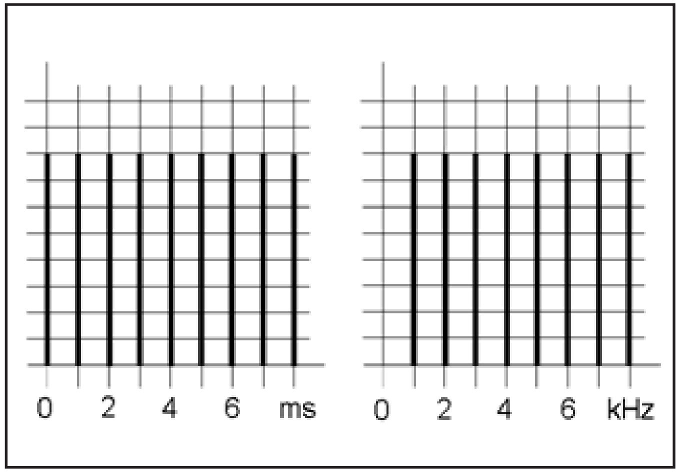

FIGURE 5: The time and frequency components of an impulse train at a rate of 1kHz. Impulses exist for all time, and at all 1kHz intervals without limit.

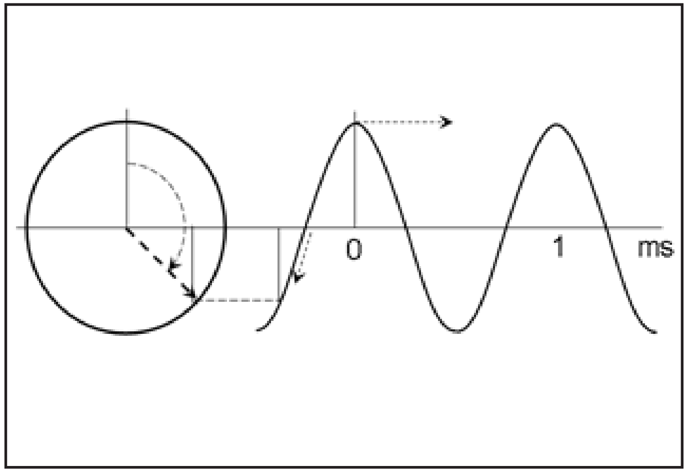

FIGURE 6: Rotating phasor representation of a cosine wave.

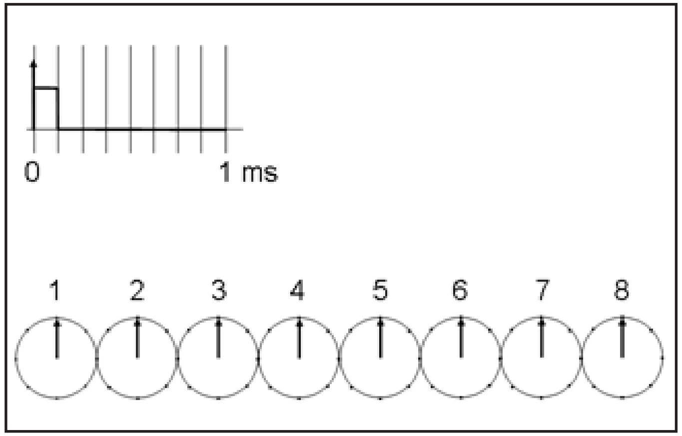

FIGURE 7: Phasors of the first eight harmonics of a single positive impulse. The height of the impulse is proportional to the charge it deposits on the capacitor.

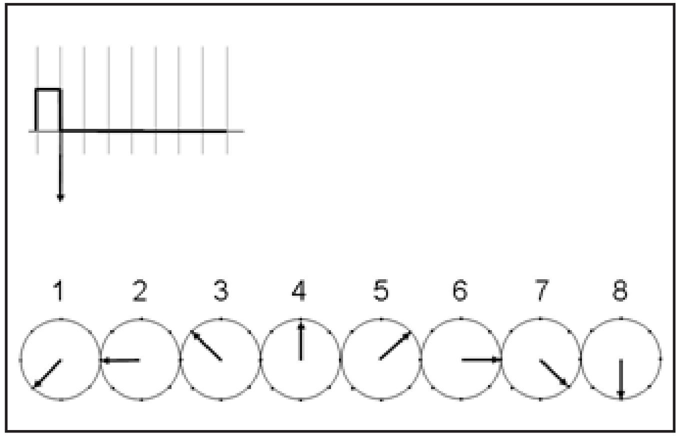

FIGURE 8: Phasors of the negative impulse at tick 1.

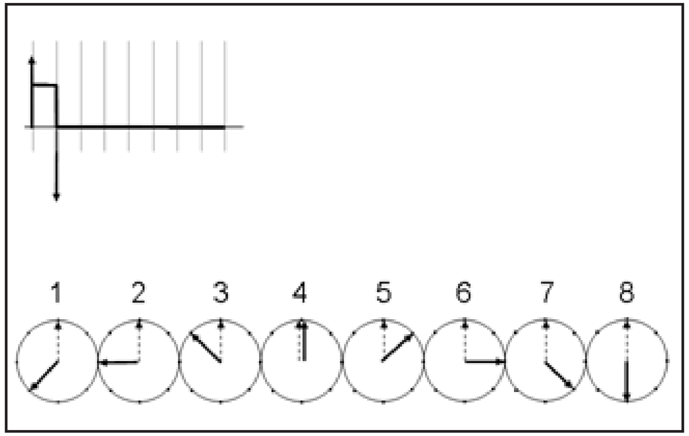

FIGURE 9: Phasors of the impulse pair.

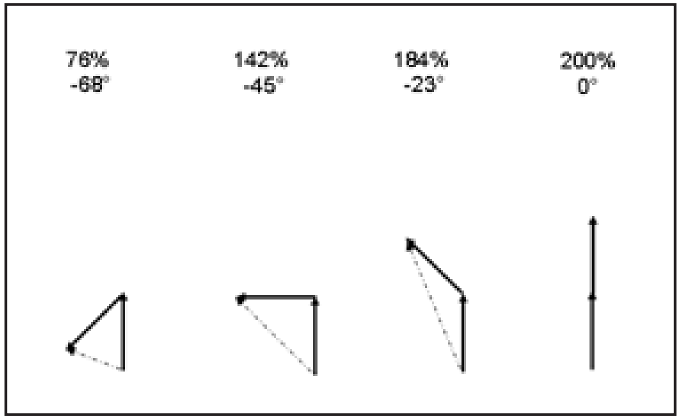

FIGURE 10: The total phasor amplitude and phase calculation. Amplitude is expressed as a percentage of that of a single cosine wave.

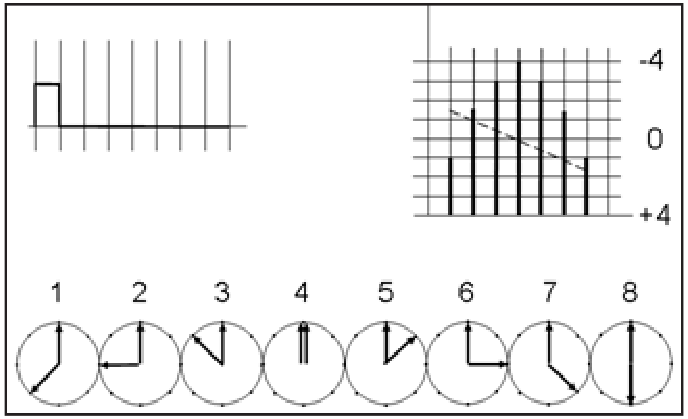

FIGURE 11: Amplitude and phase information calculated from Figure 10. Phase delay is in ticks relative to tick zero.

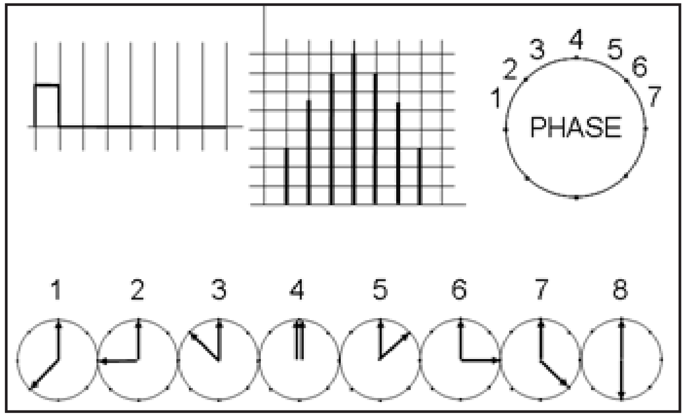

FIGURE 12: An alternative way of presenting the spectral content in amplitude and phase delay (by ticks). Tick 4 is now the reference phase (zero degrees).

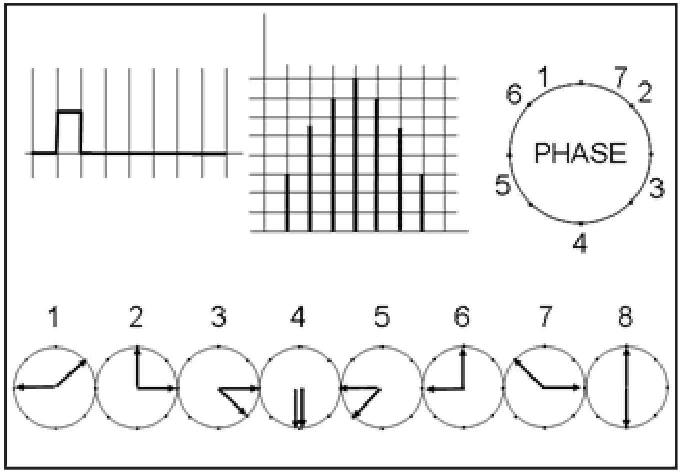

FIGURE 13: The pulse is delayed by one tick. The amplitude spectrum remains unchanged, while the phases move about in a predictable pattern.

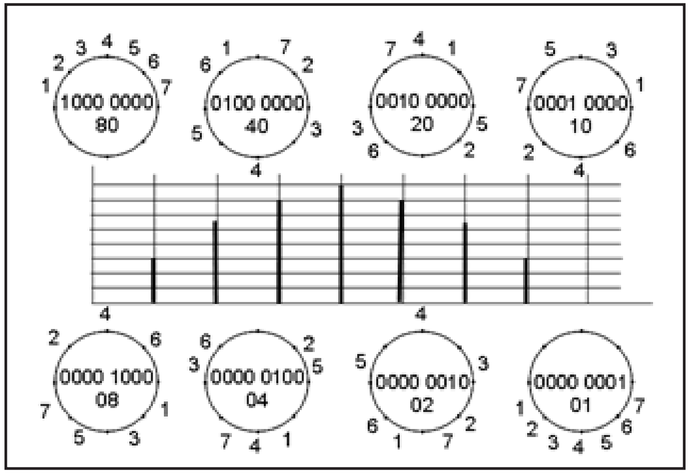

FIGURE 14: Spectra of the eight possible single one-tick wide pulses.

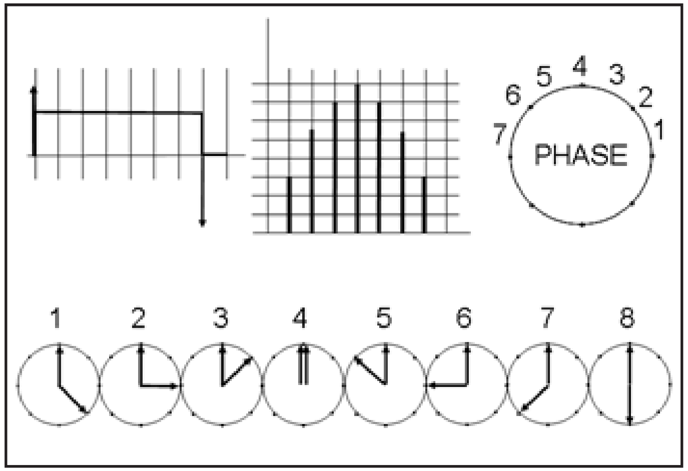

FIGURE 15: Spectra of the eight, single, 7-tick wide pulses.

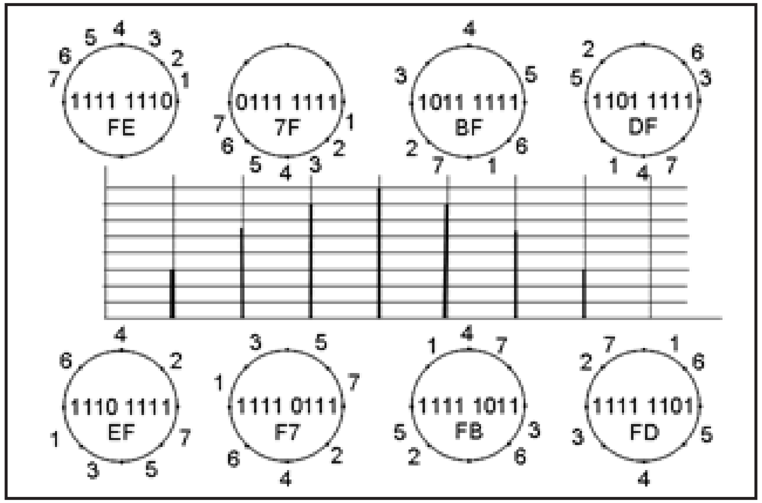

FIGURE 16: Spectra of the single, seven-tick wide pulses.

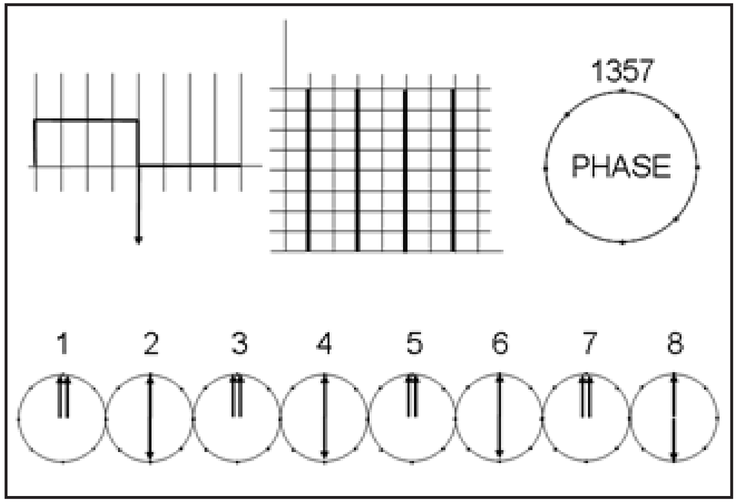

FIGURE 17: Analysis of a square wave.

fast as the first, adding two more tick's delay, a total of six as shown, labeled 2. In like manner, the remaining harmonics' phasors change by their harmonic number times one tick, plus four ticks from the negative impulse. The eight phasor diagrams in Figure 8 show this.

Combining Figure 7 and Figure 8, the drawing at **Figure 9** is obtained. The first phasor set (Figure 7) is drawn with a dotted line, and the second (Figure 8) is drawn with a solid line. You can see that the phasors assume one of four possible orientations with respect to each other, as in **Figure 10**, illustrating geometric vector addition. On the right-hand side of the diagram, they are in phase, and add to 200% of that of a single phasor. This is the maximum amplitude possible in this impulse pair's spectrum.

Percentages and phase delays for the other alignments are shown, relative to a single phasor. Using this, the amplitudes of the harmonics are found and added in **Figure 11**, revealing the spectrum of the impulses generating 10000000, hexadecimal 80.

There is now the issue of the phases of the harmonics. Figure 10 shows their phase delays. These are applied and plotted in Figure 11, showing the phase from tick zero. The first harmonic is advanced one and a half ticks. The 4th is at the reference, tick zero (zero degrees), and the 7th is delayed a tick and a half. The dotted line on the spectral display indicates these phase shifts, as is commonly done. Interestingly, the 4th harmonic is always at the reference phase, and the 8th vanishes. Any pair of equal-intensity impulses, regardless of their timing, does the same! Play with the phasor diagrams and convince yourself that this true.

On continuing to higher harmonics, the pattern repeats itself over and over. Furthermore, it's also seen that the 7th harmonic repeats the first, with a negative phase shift, and likewise for harmonic pairs 2-6, and 3-5. This phenomenon is known as aliasing.

The dashed-line phase-change becomes more and more contorted on these diagrams, so we really need an alternative way of displaying phase shifts. In **Figure 12**, the circle with the word 'phase' in the middle, shows tick marks and harmonic numbers. On it, the 4th harmonic, at zero phase shift, has a vertical phasor. Harmonic 1 peaks a tick and a half earlier, and the others are as shown. If the original pulse occurs one tick later, it encodes binary 01000000, hex 40, as in **Figure 13**. Compare its eight phasor diagrams with those in Figure 12. The fundamental rotates a tick clockwise (cw), and its mate, the 7th, a tick counter-clockwise (ccw). They remain as mirror images. The second harmonic rotates two ticks cw, with its mate, the 6th, moves two ticks ccw. The 3rd and 5th do the same, moving correspondingly three ticks, and the 4th, four ticks. The pairs remain horizontally across from each other on the PHASE diagram. But the amplitude spectrum remains the same! Only phasing changes, and its rotation is simple to see.

Continue moving the pulse one tick to the right, generating patterns 00100000 to 00000001, hex 20 to 01. With each step, the eight phasor sets continue their simple progressions, pair 1 turning one tick cw, and its mate, 7, ccw the same amount. Pair 2 turns two ticks cw, with 6 turning the same ccw. Likewise for pair 3 and 5, while pair 4 turns four ticks, a half-turn. These eight combinations are summarized in **Figure 14**, all having the same amplitude spectrum. Only the relative phases change.

Consider now the inverses of these waveforms, a pulse "ON" for seven ticks and "OFF" for one. **Figure 15** shows the pattern 11111110, hex FE. A positive impulse occurs at tick 0. For harmonic 1, there is a negative at tick 3, the opposite of tick 7. The eight phasors at the bottom show those for the positive impulse as vertical. The negatives occur at seven ticks (negative one tick) around the circle, times their harmonic number, plus four ticks to accommodate the negative impulse. Compare this with Figure 12. All eight phasors along the bottom are simply mirror images of each other. Thus the amplitude spectrum remains unchanged. Only the phase data change. The progression around the phase circle reverses as shown.

Shifting the pulse along the time axis, one tick at a time, will develop **Figure 16**, showing the phase changes for each configuration. In each case, the amplitude spectrum remains unchanged.

One may now continue with every possible waveform involving a single pair of impulses, such as 11000000, 11100000, 11110000, 11111000, 11111100, with each pattern shifted along the time axis a tick at a time, the last bit reappearing as the first in shift-register style. I leave it up to you to construct the phase diagrams.

Finally in this section, let's examine a square wave (**Figure 17**). First we consider the pattern 11110000, hex F0. As earlier stated, only

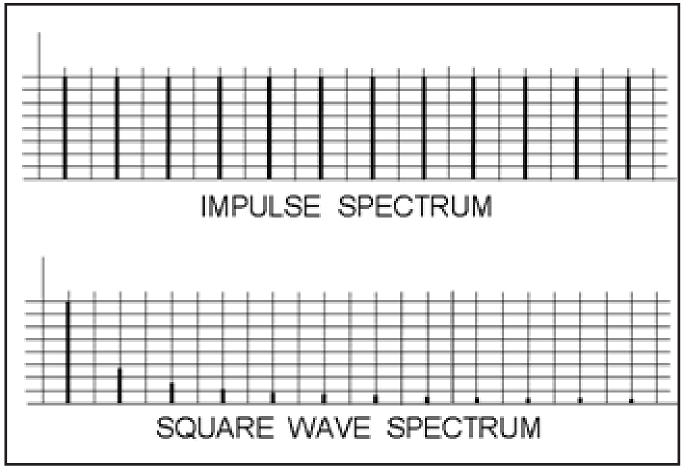

FIGURE 18: The amplitude spectra of a square wave (bottom) and its impulse pair (top).

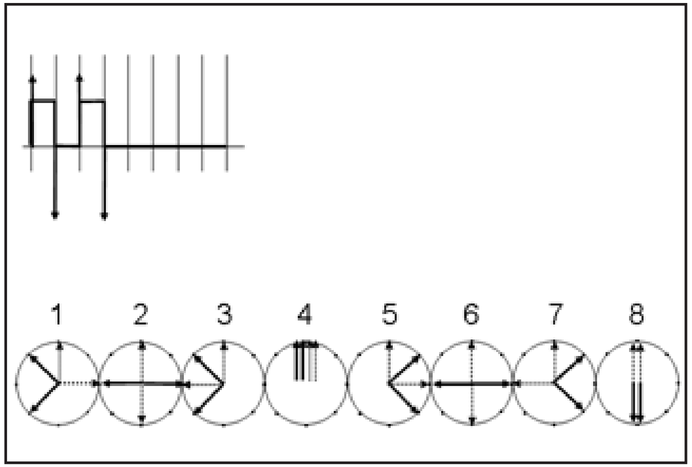

FIGURE 20: Analysis of a pair of pulses.

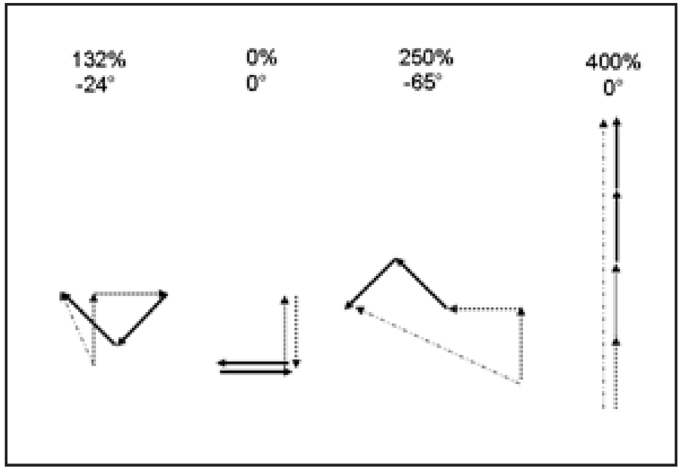

FIGURE 21: Amplitude and phase calculations for Figure 20.

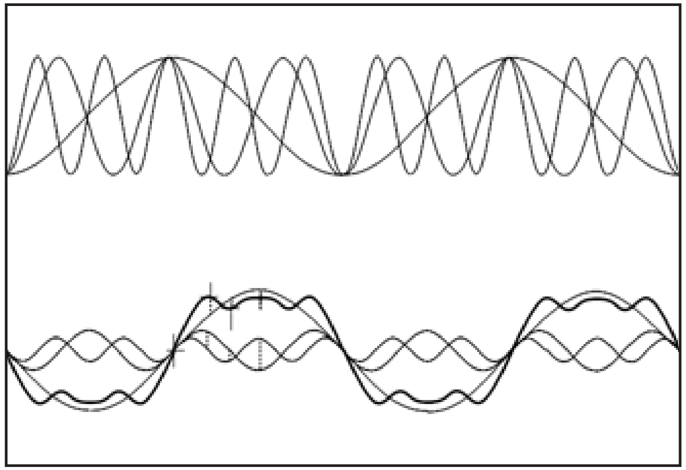

FIGURE 19: Construction of an approximate square wave from its first, third, and fifth harmonic.

the odd harmonics appear, all of the same amplitude. This is the spectrum of the impulses. To obtain that of the waveform, here and in all cases, two additional changes are required: divide each amplitude by its harmonic number, and convert each cosine wave to a sine wave. The first is depicted in **Figure 18**, showing the relative amplitudes of the impulse and square wave pulse spectra. From Figure 17, observe that all have the same phase. The higher-order harmonics drop off rapidly. One can get a fair approximation of the pulse using only the first, third, and fifth harmonics. This addition, done simply by geometric constructs, is shown in **Figure 19**.

A PAIR OF PULSES

The above arguments dispatch every possible combination of a single pulse using a pair of impulses. Now turn attention to waveforms with two pairs of impulses. Begin with 10100000, hex A0, As shown in **Figure 20**. The impulse at tick zero establishes the phasors oriented vertically on phasor diagrams 1 to 8 as the dotted, vertical phasors. The positive impulse at tick two establishes the dotted phasor oriented to the right on diagram 1. Its second harmonic spins twice as fast, and establishes the dotted phasor oriented “down” on diagram 2. Likewise its phasors on diagrams 3 to 8 advance two ticks per step, as shown.

The negative impulse at tick one establishes a vector opposite tick 1 on diagram 1, the black phasor toward tick position 5. In a similar fashion, it advances one tick position per harmonic, as shown. Finally, the negative impulse at tick three establishes a phasor opposite tick 3 on diagram 1, the other black phasor. It advances three ticks on successive diagrams 2 to 8. The eight phasor diagrams maintain a mirror symmetry configuration, 1 and 7 being mirrored, as are 2 and 6, and 3 and 5. Thus aliasing remains.

Phasor pairs for harmonics 2, 6, and 8 represent cancelling sets of cosine waves. Thus these harmonics are absent in this spectrum, while harmonic 4 has maximum amplitude, four times that of a single cosine wave. There remains the task of obtaining the amplitudes of harmonics 1 and 3, and their phase angles. This is addressed as before in **Figure 21**, accomplished with but a ruler and protractor. These results are added to the total spectrum in **Figure 22**.

As we did earlier, shift the impulses a tick at a time to the right, through its seven additional positions, from its start as hex A0 through hex 50, 28, 14, 0A, 05, 82, and 41. The amplitude spectrum remains unchanged, and the phases shift about as shown in **Figure 23**. Similarly, one may obtain the spectra of the eight inverse waveforms, binary 01011111 through to 10101111.

In like manner, spectra for waveforms based on 10010000 and its inverse 01101111 may be obtained. But observe carefully the pattern 10001000. It has but 4 variations, not eight. And patterns based on 10000100 and 10000010 have previously been discussed.

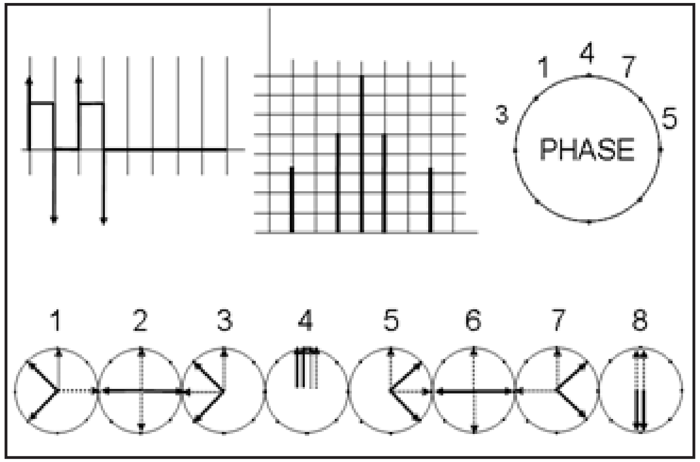

FIGURE 22: Time and frequency of the dual-pulse spectrum shown.

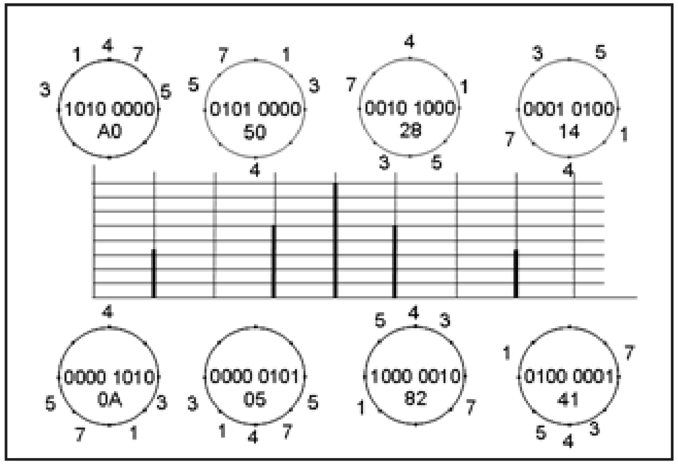

FIGURE 23: Analysis results for two-pulse waveforms of shown hex values.

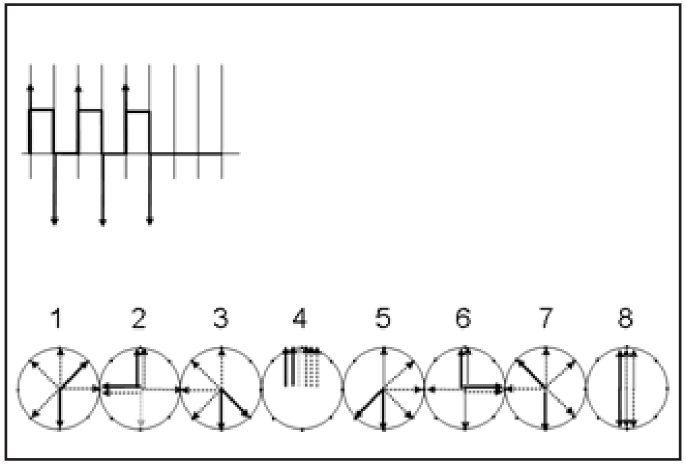

FIGURE 24: Analysis set-up for the trio of pulses shown.

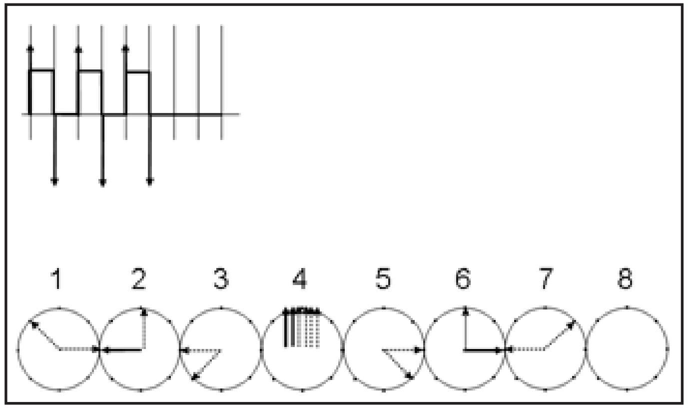

FIGURE 25: Removal of opposing (cancelling) phasors.

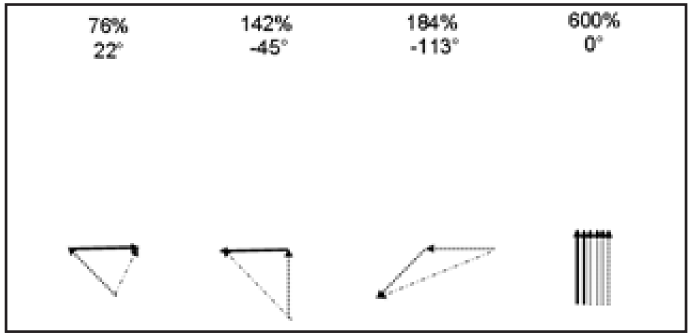

FIGURE 26: Amplitude and phase calculations.

A TRIO OF PULSES

Continue with spectra for three pairs of impulses, binary 10101000, using the same technique (**Figures 24-28**). These are shown graphically: **Figure 24** is just Figure 20 with two additional impulses added, a positive at tick 4 and a negative at tick 5. Their phases are added in the eight phasor diagrams.

FOUR-PULSE WAVEFORMS

There are two waveforms with four pairs of impulses. Their analysis is presented in **Figures 29-31**, following the methods outlined previously. The spectrum contains only odd multiples of the 4th harmonic.

Two additional waveforms are possible: hex 00 and FF. Neither has impulses. Their spectra do not have any harmonic switching from one to the other. **Figure 32** shows the situation. Since there is no paired negative impulse, harmonic 8 now appears. All harmonics are of the same amplitude and zero phase. Its inverse is identical, with phasors yet all vertical, but the phase shifts four ticks around the circle.

I AND Q

Each harmonic has an amplitude and a phase. Its phasor is the sum of two components, one vertical at zero phase shift, and one horizontal at 90° phase shift. These are the in-phase and quadrature components respectively, the I and Q phasors of common parlance, as in **Figure 33**.

CONCLUSION

I hope that you can now visualize the simple extensions of these teachings to sampled, non-binary waveforms, and to systems of differing bits per sample. The principles are the same, and the geometric approach may bring you insights that are, perhaps, not readily apparent using straight mathematical formulations.

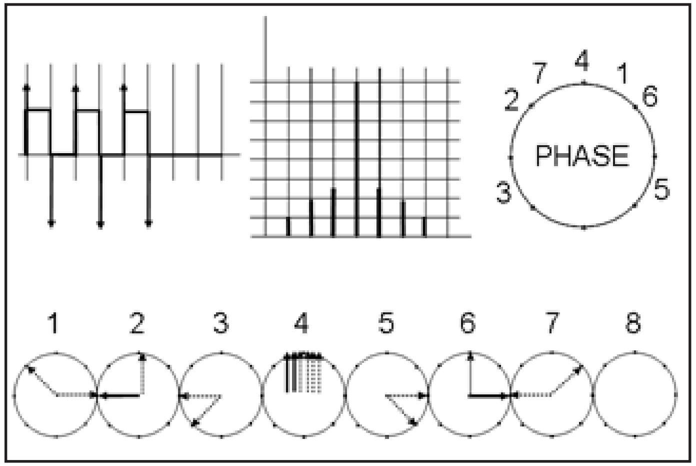

FIGURE 27: The spectrum of 10101000.

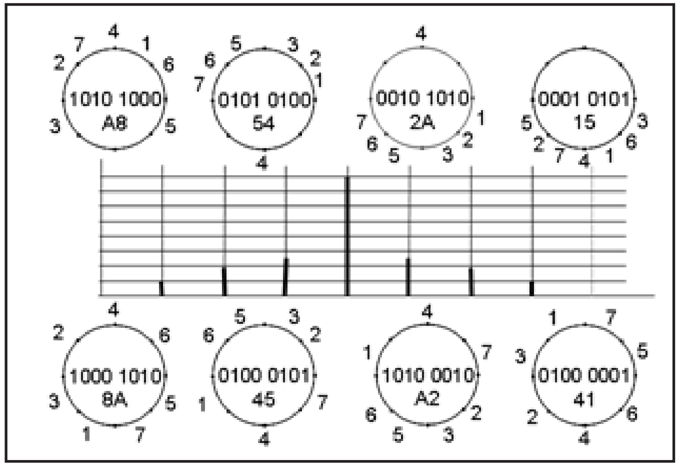

FIGURE 28: The complete spectrum of the eight variations of 10101000.

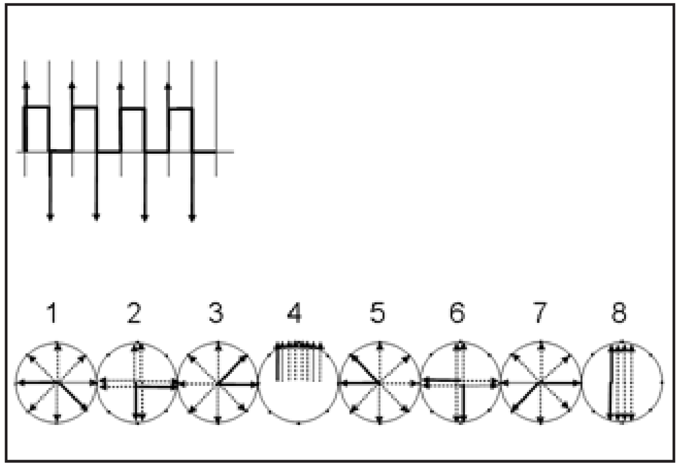

FIGURE 29: Analysis set-up for 10101010.

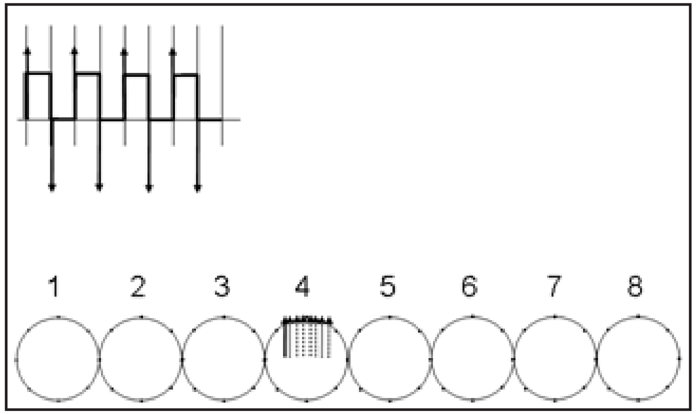

FIGURE 30: Removal of opposing phasors. Only the 4th harmonic remains.

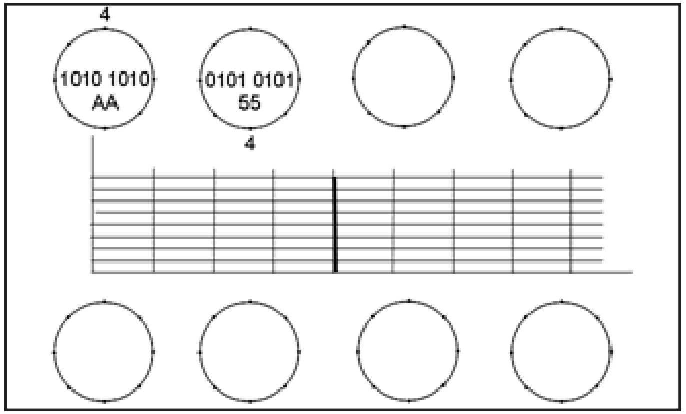

FIGURE 31: The complete spectrum of the two possible patterns.

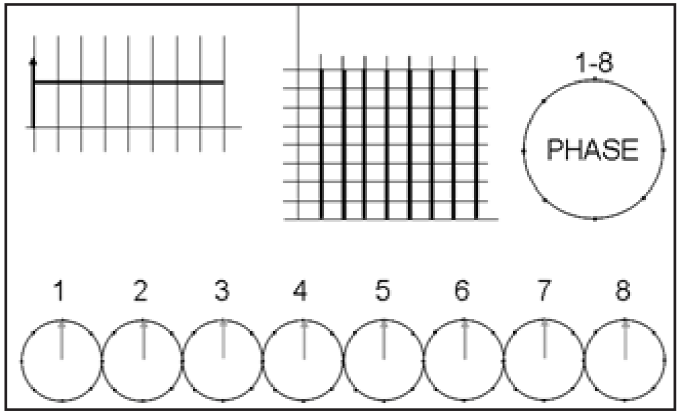

FIGURE 32: The spectrum of a positive impulse at tick zero.

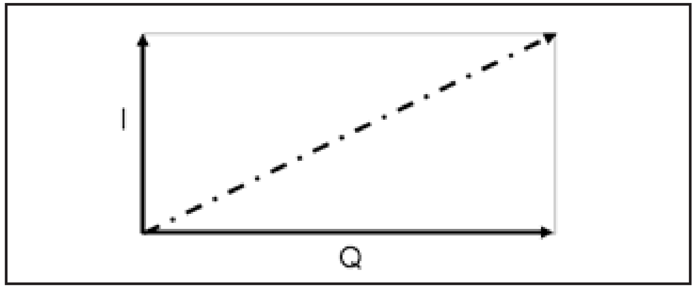

FIGURE 33: Calculating I and Q components of any phasor (any harmonic).

The design of linear power supplies

Andy Nehan, G4HUE
andy.nehan@btinternet.com

This is an article about linear power supplies. I have many years of experience designing and building linear and stabilized power supplies, but have no expertise in switch-mode power supplies; this is something I will leave to others.

TOOLS

There is one tool which I regard as *de rigueur* for the design of any linear supply, namely Ben Duncan's PSUDII **[1]**. I would only consider the design of a linear power supply after having first run the numbers through this tool. It works with both high- and low-voltage supplies, for both high- and low-current delivery. Let's take an example or two to see how it may be employed.

EXAMPLE 1

We need a regulated 13.8V power supply which can deliver 3A, using Schottky diode rectifiers. These are named after William Schottky (1886 to 1976) who investigated metal/semiconductor interfaces. These diodes exhibit a forward voltage drop of between 150mV and 450mV. Choosing, for example, 400mV and a current of 3A, the power dissipated in the diode is 1.2W. If the diode is housed in a TO220 package, it will probably require a heat sink! A Schottky diode usually has an inverse maximum voltage below 100V, although those produced by CREE, amongst others, can withstand higher voltages. 1kV types are available, but one needs to search for them. **Figure 1** shows the basic circuit diagram, and the information produced by Ben's tool. Note that the left-hand vertical scale indicates the voltage across the capacitor and the right-hand scale the current through the transformer. As can be seen in the left panel it is possible to select, for plotting, current and voltage waveforms throughout the circuit.

There is a lot of information is displayed in Figure 1. As selected, the voltage across the load is shown in green and the current through the transformer in red. The left-hand side of the plot begins after a delay of 1s from switching on, and shows 25ms worth of the voltage and current waveforms. Note that the voltage scale is offset from zero. The mouse may be used to select different configurations. It may be useful to know the peak current flowing through the diodes, transformer, and capacitors, for example, and such information is relatively difficult to obtain by other means. You can select the transformer, a diode, a capacitor, a load type (including a stepped load where the load changes at a specified time), capacitor input, choke input and stages of smoothing. You will need to know the series resistance (ESR) of the smoothing capacitors, the regulation of the transformer, and may need to specify diodes of other types than the ones you want to use (the tool may issue a warning but will still give good results). I suggest you use regulation figures of 5% for large transformers (4in each side and upwards), 10% for small transformers, and 20% for very small transformers of a few cubic inches.

The impact of making a poor assumption regarding the series resistance of the capacitor is quite significant (which alone shows just how important a 'good' capacitor is, as also is its provenance). That old electrolytic capacitor found in the junk box or mobile rally may be well be past its sell-by date! This is illustrated in **Figure 2**, in which you can see that the capacitor does a very poor job of providing an output with minimum ripple. The ripple using a capacitor with an ESR of 30mΩ was approximately 0.9V (Figure 1), but with a 1Ω ESR the ripple is 6.6V. This brings home the importance of good capacitors.

EXAMPLE 2

We need an unregulated 400V power supply supplying 100mA. We can use a standard semiconductor diode. These became widespread in the 1960s, although at that time they were of limited current capability and peak inverse voltage rating. Today, even an inexpensive 1N400X series diode will support a high voltage power supply. Forward voltage drop is typically in the region of 0.7V, which in Example 1 would imply a dissipation of 2.1W, certain to require a heat sink if a TO220 package is used, but only 0.07W in this example. **Figure 3** displays this example in Ben's tool. The red line shows the ripple voltage, which is 7V peak to peak. One must also make an allowance for mains fluctuations (see below).

BRIDGE VERSUS BI-PHASE RECTIFICATION

Bridge and bi-phase rectification have some similarities, but equally important they have differences, and these two rectification schemes are not interchangeable. The RSGB manual 3rd Edition dated 1963 and the 4th edition dated 1968 had a most interesting table, shown here as **Figure 4** (not present in editions 1 or 2 and slightly expanded and modified in editions 5 onwards). The figure shows that the peak inverse voltage rating (PIV) required of the rectifier configured as a full-wave bridge is half that required of a bi-phase configuration. This may be important if the PIV requirements are severe, but be wary of 'cutting it fine' when selecting a diode as spikes on the mains may destroy it. The figure also shows that the rms current per diode is greater for the bridge, although this is not usually a problem. However, Figure 4 does not show that there are two diode voltage drops with a bridge, rather than the one in a bi-phase configuration. This may be important for low-voltage supplies. It is also worth noting that twice as much copper wire is used for the secondary winding in a bi-phase configuration, and therefore the cost of the wire, the size of required bobbin, and the size of the transformer core may all be greater. If you are using valve rectifiers, the bridge configuration may be difficult as two of the cathodes will be at different potentials and are likely to require different heater windings.

Additional information can be found at **[2]** and other locations on the Internet.

Wherever possible use standard, off-the-shelf, transformers which, for low-voltage types, often have two outputs. In this situation, a bridge rectifier is often the best approach. If a transformer is rated at 40VA and has two windings of 0-20V each at 1A, then 2A may be drawn if the windings are paralleled (ensuring that the 0V connects to the 0V and the 20V to the 20V). If the bi-phase configuration is used, only 1A should be drawn as otherwise the wire rating of 1A will be exceeded.

Selecting amongst differing types of silicon rectifier diodes, I would suggest using Schottky diodes wherever possible, especially as high-voltage versions are now available. For a fuller exploration of the differences between diode types, see **[3]**. In summary, Schottky diodes are usually to be preferred as they

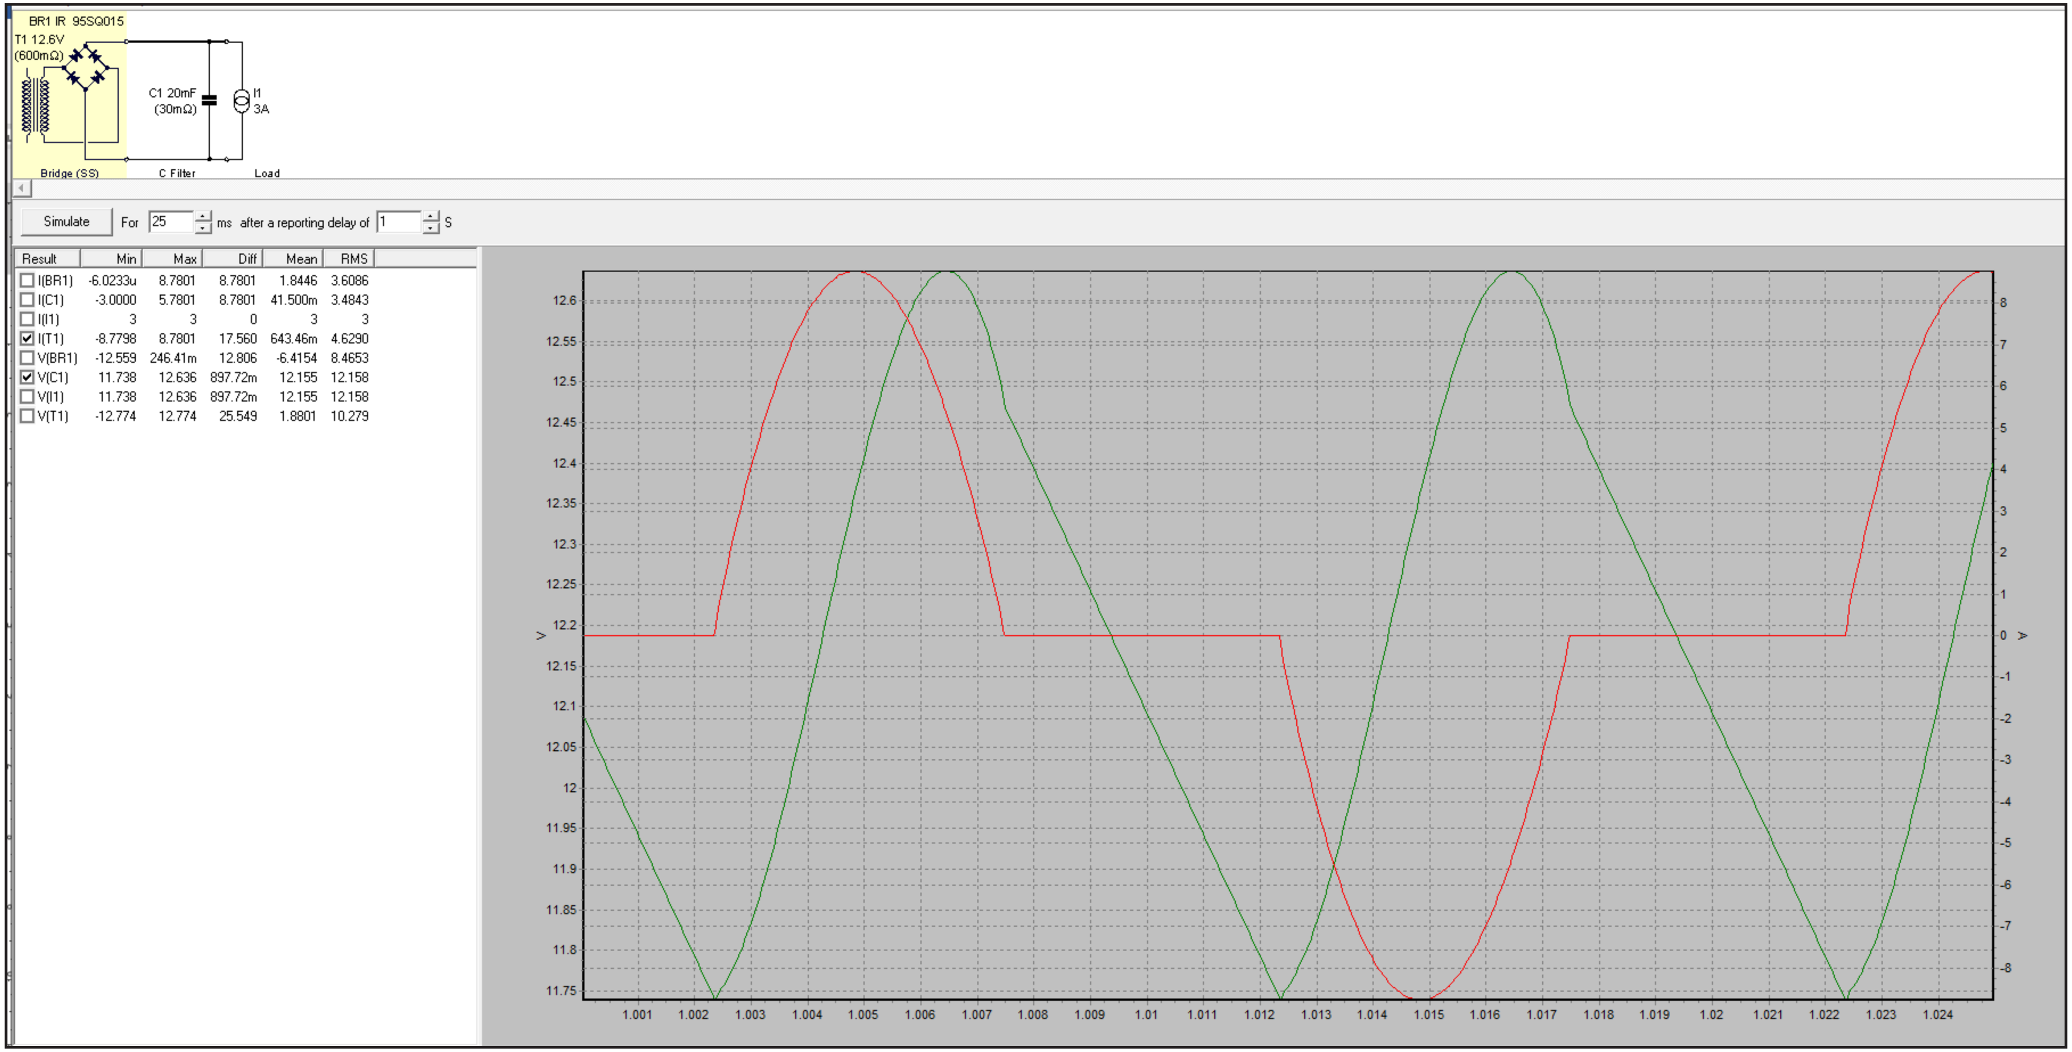

FIGURE 1: A 13.8V 3A supply.

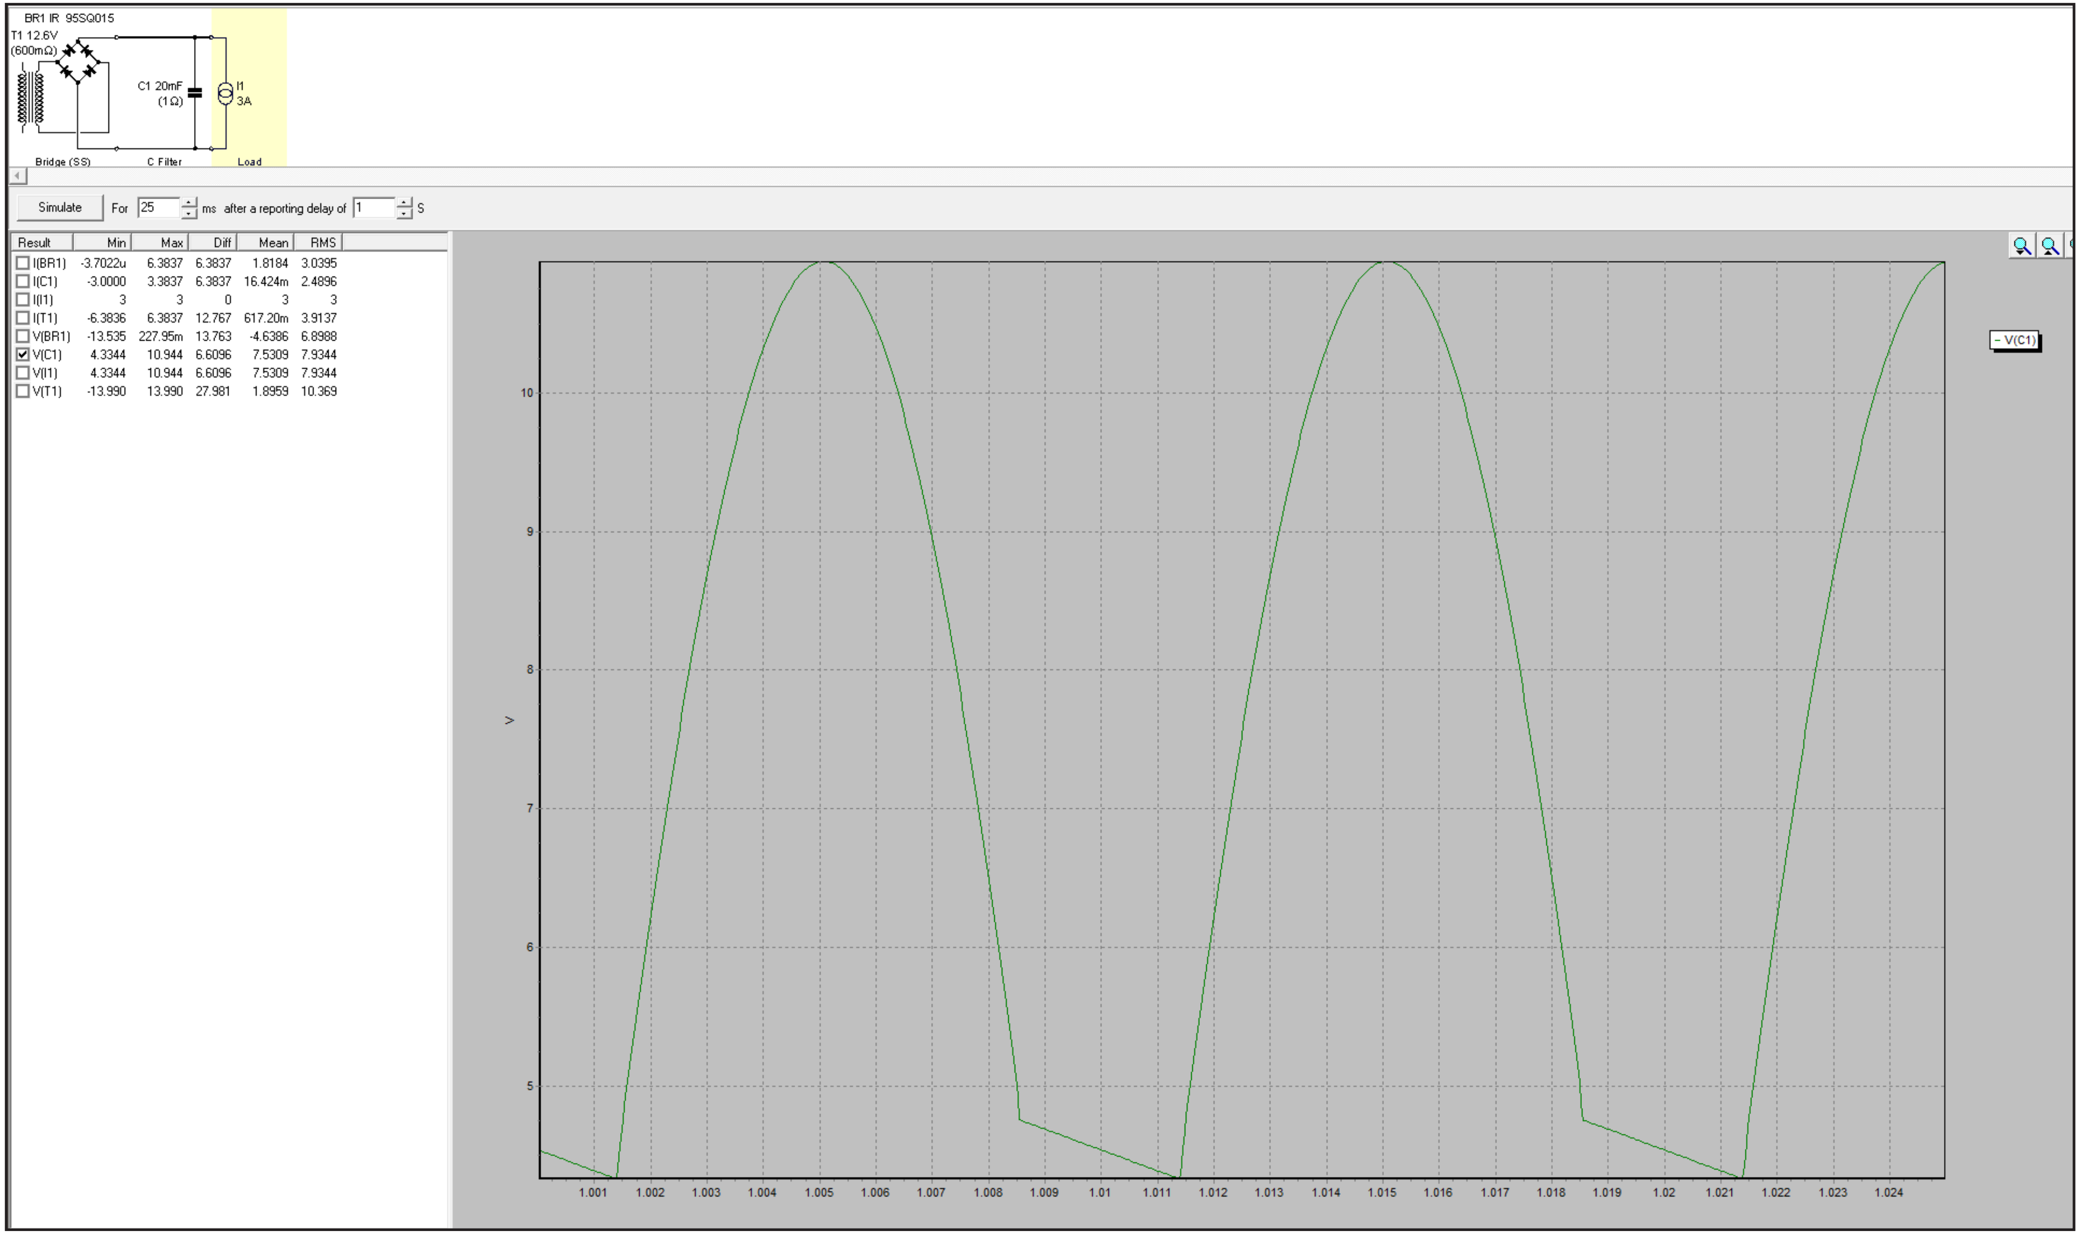

FIGURE 2: Using a poor-quality reservoir capacitor.

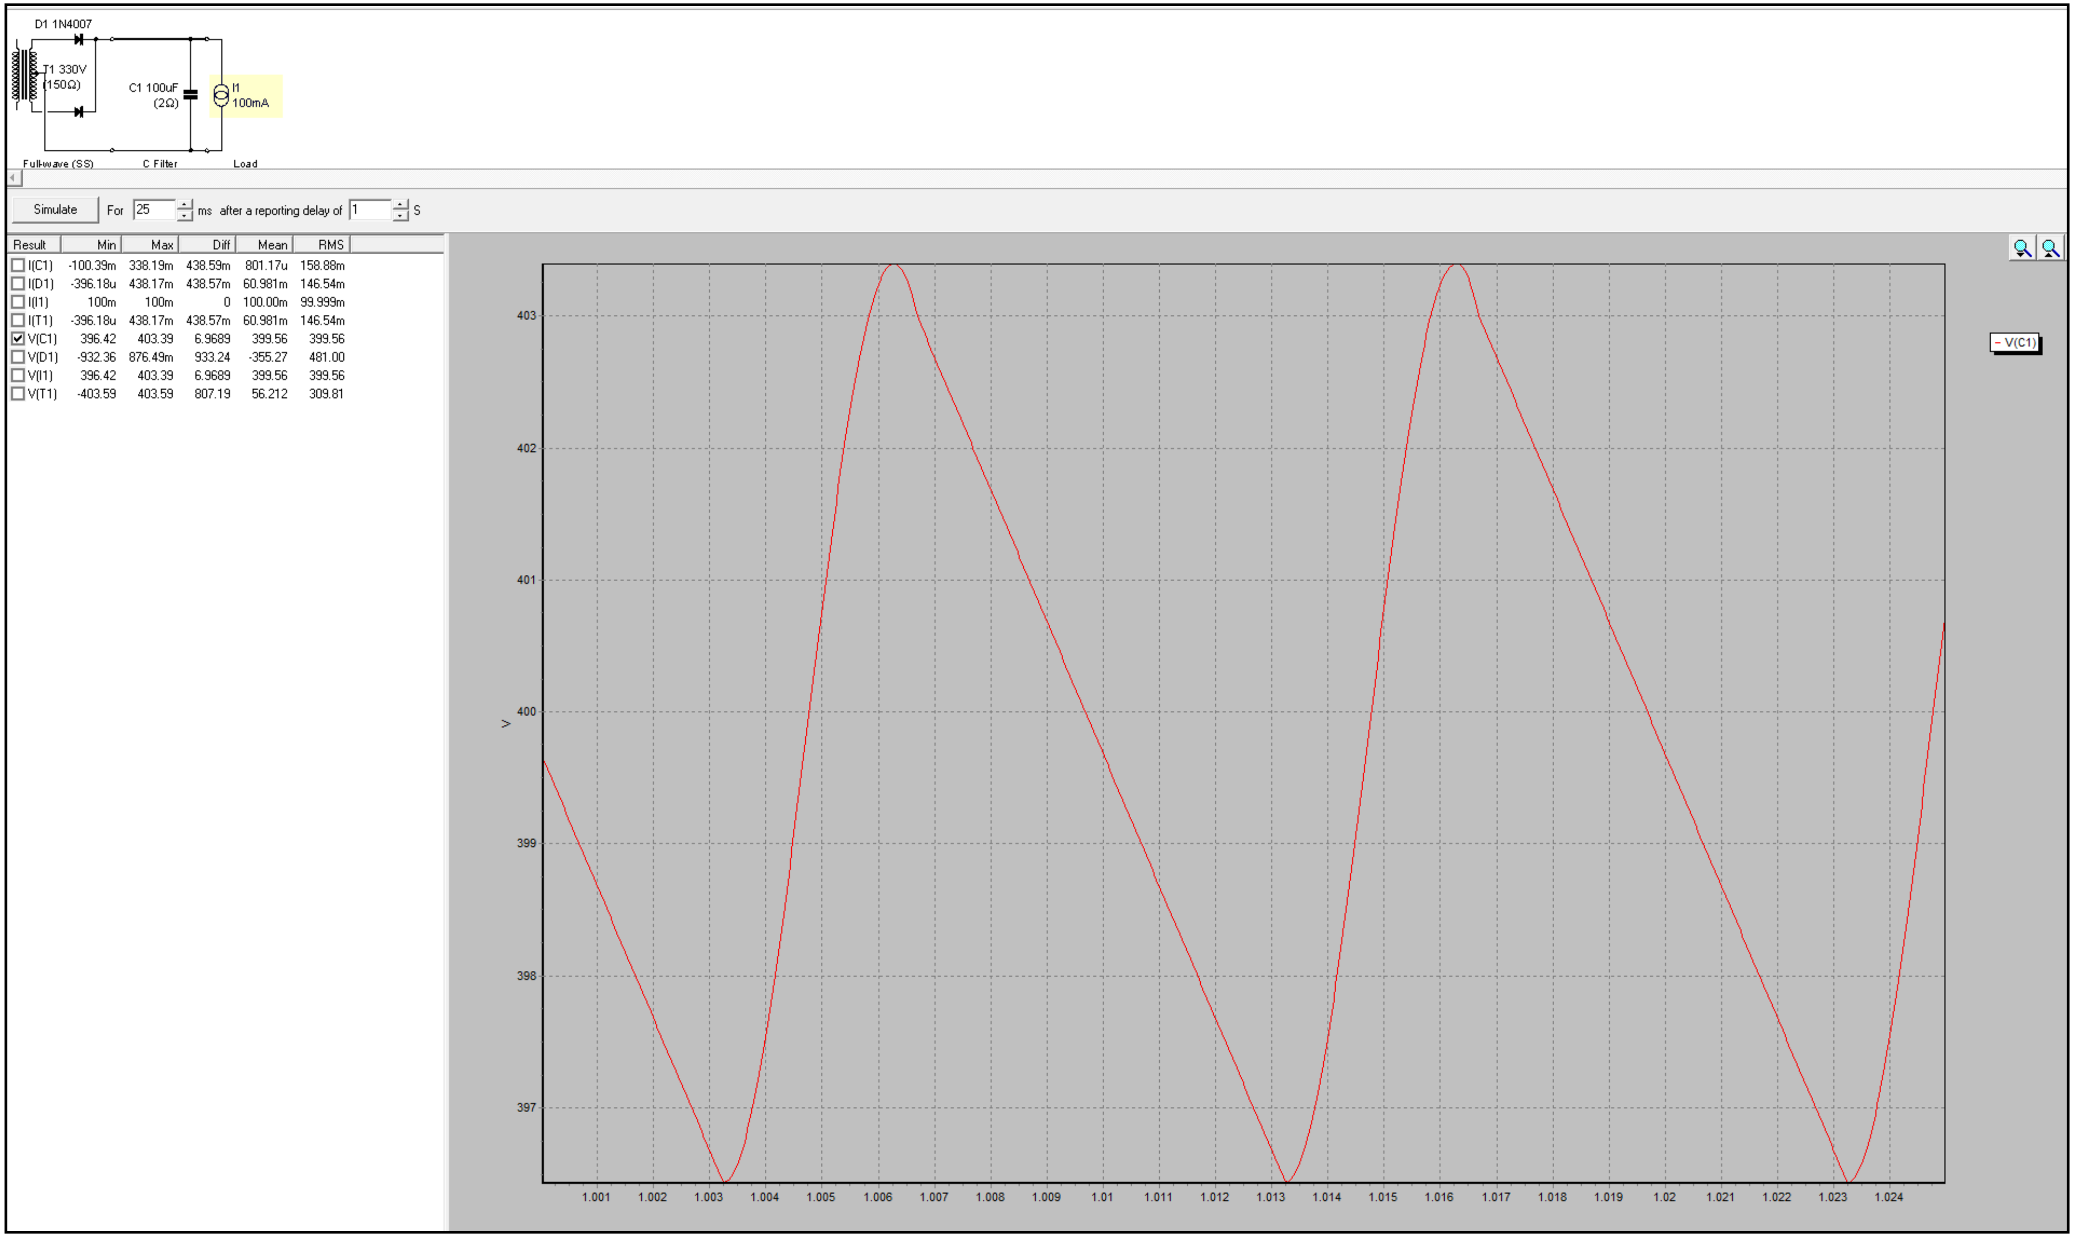

FIGURE 3: A 400V 100mA supply.

Single Phase Rectifiers

	Half-wave	Full-wave or biphase half-wave	Full-wave bridge
Circuit	Fig. 1 (a)	Fig. 1 (b)	Fig. 1 (c)
Number of rectifiers	1	2	4
Current per rectifier —mean	1·0 I_{dc}	0·5 I_{dc}	0·5 I_{dc}
—peak (inductive load)	—	1·0 I_{dc}	1·0 I_{dc}
—peak (resistive load)	3·14 I_{dc}	1·57 I_{dc}	1·57 I_{dc}
—r.m.s.	1·57 I_{dc}	0·725 I_{dc}	1·11 I_{dc}
Peak inverse voltage across rectifier A	3·14 V_{dc}	3·14 V_{dc}	1·57 V_{dc}
B	2·83 V_{ac}	2·83 V_{ac}	1·41 V_{ac}
R.m.s. secondary voltage	2·26 V_{dc}	1·13 V_{dc} *	1·13 V_{dc}

V_{dc} = output voltage I_{dc} = output current V_{ac} = input a.c. voltage
* Each half of the h.t. winding.
A = choke input B = capacitor input

FIGURE 4: Rectifier classes and their circuit impact (from the RSGB's Communications Handbook 4th Edition, 1968).

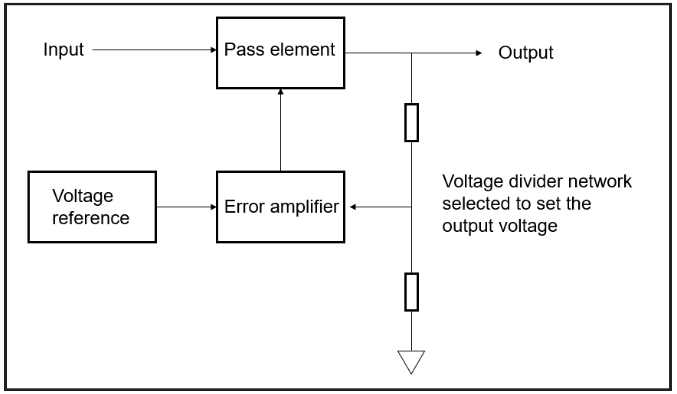

FIGURE 5: Simplified voltage regulator block diagram.

possess none of the unwanted recovery characteristics of other silicon diodes. When a non-Schottky silicon diode switches off, it does not do so instantaneously but rather it has a recovery phase which may have an abrupt transition. Such a transition will introduce unwanted high-frequency components into the circuit, and may also encourage 'ringing' in the transformer, something we wish to avoid

REGULATORS

There are multiple voltage regulators available, and it is not the intention of this article to recommend one particular type. But we can look at the generic design characteristics of regulators, whether they be IC-based using devices such as the aged and venerable 78xx series, valve-based, or using discreet components. A simplified view is shown at **Figure 5**. It will be observed that, provided that the feedback amplifier has a finite gain-bandwidth product, then the following statements are true:

The output impedance is affected by the limited gain-bandwidth product such that it increases from a low value at a turn-over frequency, which is often a few kHz. Effectively, the device represents an inductive output impedance (that may be modified by adding a capacitor across the output terminals), and this is usually of little or no concern. The characteristics of the 7805, for example, are shown in **Figure 6**, although these may vary slightly between manufacturers. Its turn-over frequency is a little more than 1kHz.

The circuit is not immune to high-frequency pass-through as it has significant input-to-output capacitance that will vary according to the type of the pass element. This is illustrated in **Figure 7** for the 78xx series.

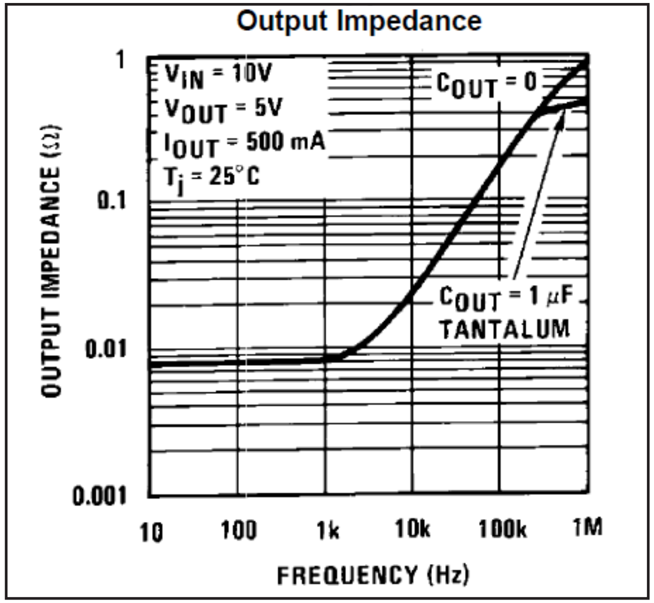

FIGURE 6: The output impedance of a typical 78xx series regulator.

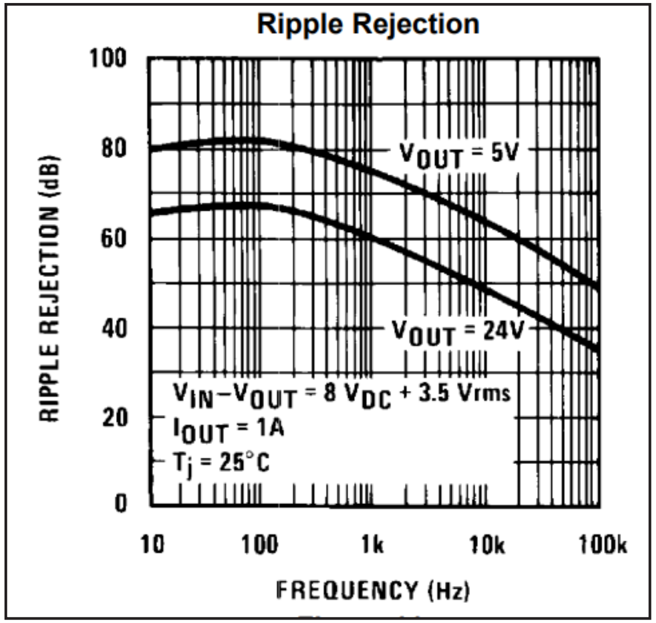

FIGURE 7: Ripple rejection for a series 78 regulator.

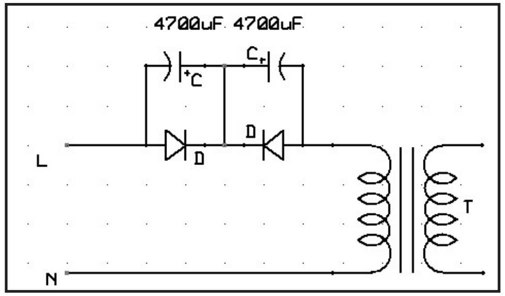

FIGURE 8: A DC blocker.

FIGURE 9: The variation of the mains voltage as observed in late December 2022.

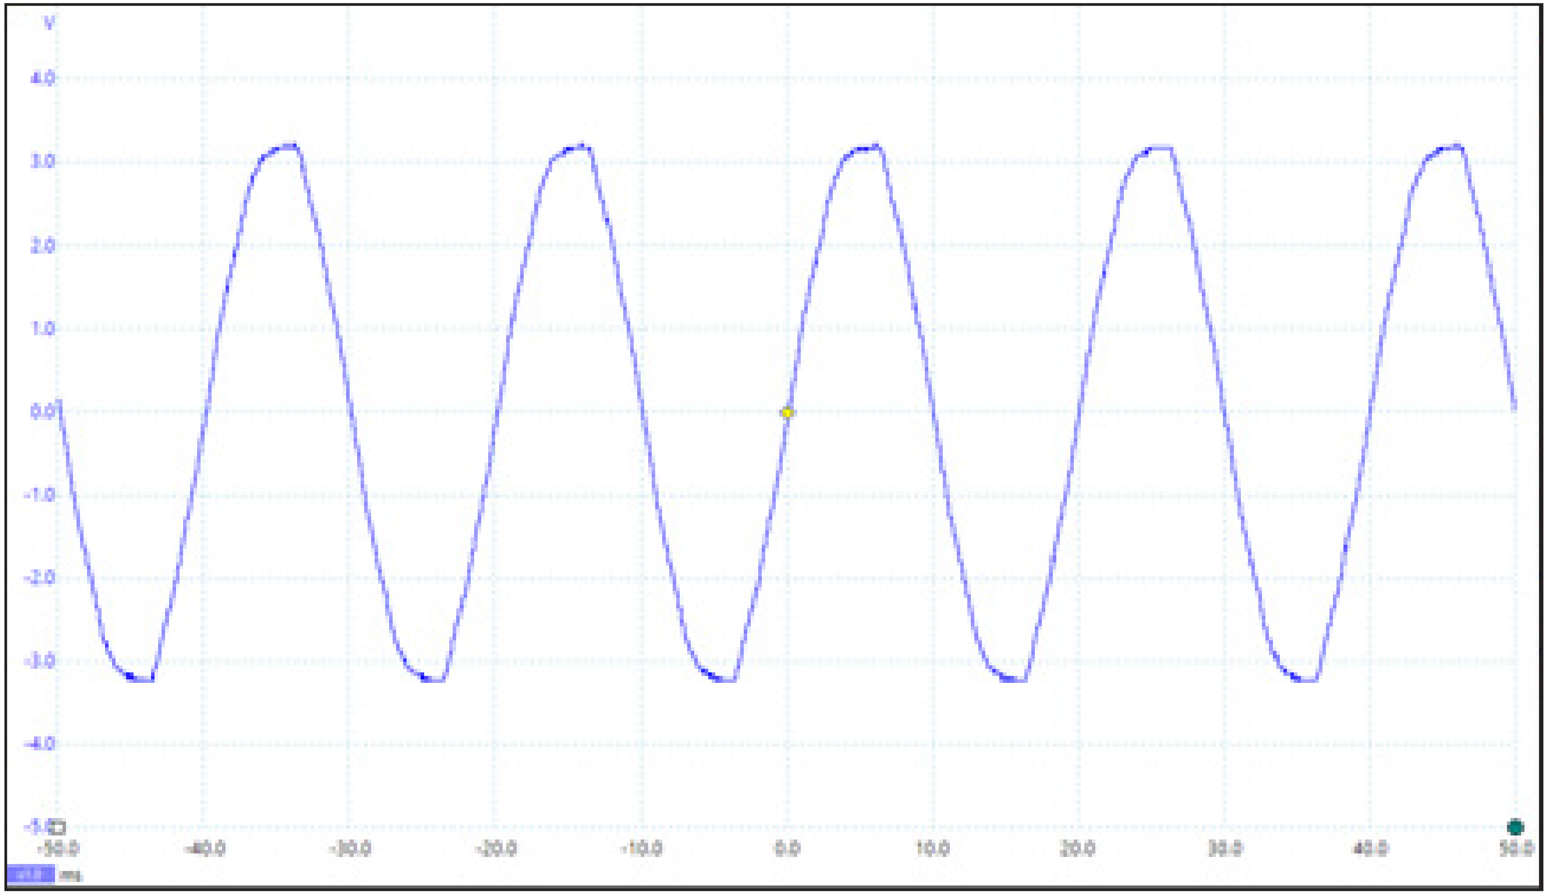

FIGURE 10: The waveform of the mains.

The device is not immune to variations in input voltage, even with a constant output load. This is typically only a few mV. For some highly-sensitive applications, a very stable and clean supply may be important. In this case series-connected regulator stages can often fix the problem.

THE MAINS

In an ideal world, in which the mains supply was a waveform of equal average power on both the upper and lower half cycles, there would be no DC offset. Unfortunately, the use of switch-mode power supplies, of sometimes dubious conformance to regulations, has meant that both halves of the mains waveform are not necessarily of equal power. However, this is usually only a problem when using toroidal mains transformers. They can hum quite badly as they are particularly susceptible to DC that can cause saturation of the core. Toroidal transformers are not 'gapped', as in so-called 'EI' transformers (composed of layers of E and I shapes), and this makes them susceptible to saturation. A gap is introduced by the nature of joining the E- and I-shaped laminations, so that a continuous metal structure is not produced. It would appear, from my experience, that the larger the transformer the greater the problem. A

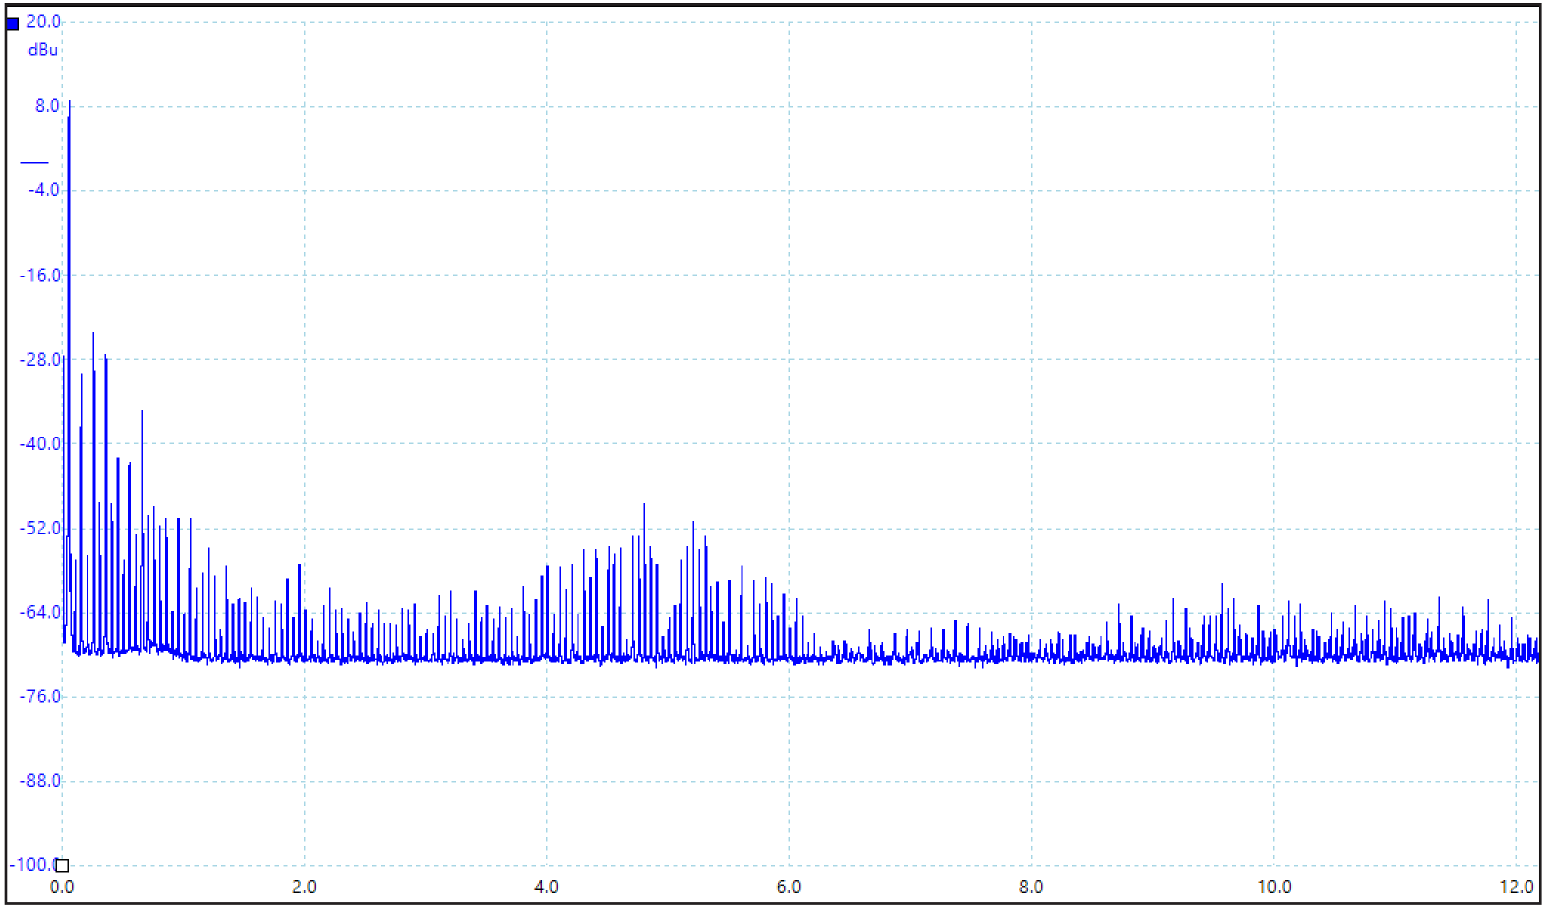

FIGURE 11: The frequency spectrum of the mains voltage.

commonly-used method to defeat this problem is shown **Figure 8**, and there are many designs which can be found using the Internet search words 'DC blocker'. The website at **[4]** is very detailed. It should be noted that the working voltage of the capacitors need not be large, as the voltages across them are small, but large currents will flow through them.

There are commercial products which claim to cure this problem, some at exorbitant prices.

Mains voltage variations must be considered as the UK mains voltage is set to be 230V minus 6% plus 10%, ie 216V to 253V. **Figure 9** shows my measurements of the mains voltage over a 6h period, which varied between 218V and 234V (I have seen larger variations). I called the electricity network operator, the 'DNO' **[5]**, some years ago when the mains voltage dropped below the statutory minimum. After investigation, the supplier discovered that a temporary redirection of the local 33kV line had not been implemented correctly. Your situation may well differ; my son, who lives in London, has a remarkably-stable supply voltage for example. But we are all aware that, in the event of power-supply problems at times of high load (typically in the winter), 'brown outs' may occur. These are periods when the supply voltage is allowed to drift lower than the statutory minimum. Such events might usefully be considered within the design brief.

The mains waveform itself is usually less than ideal, not just in its voltage variability, but also in its spectral content. A PC-based oscilloscope was used along with a x100 probe to produce the waveform shown in **Figure 10**, and the frequency spectrum in **Figure 11**, up to 12kHz. As can be seen, the waveform was quite flat-topped in this measurement, which is much to be expected, and this produces spectral products above 2kHz. I am advised that there can also be significant RF content, as well as mains spikes of several times the nominal voltage.

The mains frequency here in the UK is very stable, as can be seen from **[6]** and **[7]**. There appears to be only 50 or 60Hz in use throughout the world, although some countries do have trouble providing a reliable electrical offering. I am aware of only one country, namely Japan, that uses both 50Hz and 60Hz supplies, for purely historical reasons.

RSGB handbooks from Edition 3 onwards have diagrams showing the relationships between LC and RC networks and the suppression of ripple, but note that PSUDII, pspice, and other programs are able to produce accurate results with ease.

Both common-mode and differential-mode chokes are readily available for installation in the feed to the transformer primary. The idea here is that a signal on the mains, which might be RF that is generated in the shack and present on both L and N, would then be reduced. I can't personally vouch for their efficacy, but companies such as Wurth Electronik **[8]** produce a large range. It is probably less likely that differential chokes are needed.

CURRENT THROUGH THE DIODES

The advantage of modelling is that otherwise-obscure or difficult-to-digest topics become visible. In the example shown in Figure 1, the current through the transformer is nearly 8.8A peak, even though the load draws only 3A. If the capacitor after the diodes is increased in value, this current will increase. Measuring such a current could be achieved by use of a current clamp, but I believe it's easier to model it.

WHERE TO INSTALL FUSES

Mains plugs are fused, as they should be, with as low a value of fuse as the system permits, rather than leaving the default value at 13A. Be aware, however, that fuses become 'tired' if used near their rating, and after some time they may therefore become open circuit. A fuse should ideally be placed on, or near, the output of the power supply, in order to protect the power supply from failure of equipment connected to it. However, there can be reasons to install additional fuses. Consider the risk of diodes becoming short circuit. This happened to me and, as luck would have it, I was in the room at the time and smelt that classical transformer burning smell, and I switched off in time to save a very expensive transformer. Consideration of failure modes (even if unlikely) can save a lot of money and frustration. In my example, two semiconductor diodes went short-circuit at the same time, not likely but it happened. When using a bridge rectifier circuit, I put a fuse in one wire from the transformer to the diodes; that may be overkill, but a fuse is cheaper than a new transformer! Fuses themselves come in fast and slow varieties. It is often not possible, purely by inspection, to spot the difference, and indeed some types of fuse have a ceramic interior which prevents inspection. Slow-blow types typically have a 'T' on one of the end caps to indicate a 'timed' delay (ie not fast). The actual over-current required can be found in the data sheets; buy from a reliable source so that you can rely on the information. Data sheets on the 'Rapid Electronics' web site provide extensive characteristics (provided by the manufacturers). Some even give a time versus over-current graph to indicate how long an over-current situation need arise before the fuse blows.

I have not covered 'crowbar'-type circuits, simple as they are, as I consider them as providing protection for connected devices. I always add a crowbar circuit in line to my expensive amateur-radio equipment connected to the typical 13.8V supply, as I feel sure that one day the supply will fail with a high output voltage.

HOW TO TAME INRUSH CURRENT

This is not usually a problem with valve rectifiers as they turn on slowly, but is a potential problem with semiconductor or mercury rectifiers, because the first capacitor, when using a capacitor input design, presents a short circuit to the transformer at the instant of turn-on. When using large transformers coupled with large capacitors, it is quite common to hear a 'bump' on switch-on. In the relatively recent past, it was considered prudent to employ zero-switching turn-on methods (often using a thyristor) until it was observed that this method increased the risk of causing diode failure. Simple soft-start circuits are available from outlets such as Ebay, and can be used to good effect; they work by inserting a resistor in series with the mains for a second or so before switching the resistor out with a relay. A resistor of, say, 100Ω limits the inrush current to 2.5A, although consideration needs to be given to its power and current rating. A 10W rating might be in order to pass that current without self-destruction. Some experimentation as to the correct value of the resistor may to be required. The value of the resistor needs to be selected such that the circuit is softly started, and yet the system voltage is as near as possible to its final value. This may sound difficult, but there is no need for extensive experimentation.

TRANSFORMERS

Clearly there are suppliers of sound and robust transformers, but beware as some available from Ebay are, to be honest, little more than junk. I purchased a pair of R-core transformers for a 13.8V supply and, when I had tested them, they were consigned to the bin!

EI and toroidal transformers have quite different characteristics. In general, toroidal transformers have a high bandwidth (one of the reasons they are used at RF, although constructed using powdered iron or ferrite). I have measured some EI mains transformers with a -3dB bandwidth of about 30kHz and significant output at 100kHz, whereas a similarly-rated toroidal transformer's bandwidth was over 400kHz, with significant output at over 1MHz. There are also C-core and R-core transformers; the advantage of a C-core transformer is that the inter-bobbin capacitive coupling should be almost zero.

Toroidal transformers are often employed as they hum a lot less that other types (although they may need to be rotated in order to reduce hum pickup by the surrounding circuit). Even so, they can hum for a variety of reasons (including DC on the mains as mentioned above): other reasons include being made down to a low price point so that quality has been sacrificed. (The legal definition for quality is 'fitness for purpose including implied conditions') The transformer may induce hum in a metal chassis. I use thin MDF under EI transformers as it has some flexibility and then secure the transformer with Nyloc nuts and bolts (not over-tightened). Toroidal transformers tend to be shipped with springy top- and bottom- washers which have a similar effect. A commercial 30A linear supply designed for the amateur radio market hummed very badly, even at minimal current usage. I reduced the noise significantly using this approach.

The better EI transformers usually have an electrostatic screen (ES) between the primary and secondary windings. Connect this to earth as its purpose is to screen the secondary from some of the 'crud' on the primary. Low-cost transformers do not usually have this screen. Toroidal transformers may have a grain-orientated silicon steel (GOSS) band round the outside but not, in my experience, an ES winding. The purpose of the GOSS band is to reduce magnetic field from the transformer, and as such it does not need to be earthed.

Bespoke transformers are readily available, albeit at an increased cost and extended lead times. I have used both Majestic Transformers and JMS Transformers, and both companies are reliable, although all the transformers I have received from Majestic have been fully vacuum-impregnated with varnish, something that does not happen on cheaper products. The idea behind impregnation is that it provides a strong mechanical bonding, and offers protection from environmental conditions. Using a vacuum during this process eliminates moisture and provides deeper penetration of insulation into transformer cavities. If you are going to order a bespoke transformer, only order one designed for 60Hz (it may use less iron) if you know for certain that it will not be used on a 50Hz supply.

TRANSFORMER HEAT

Transformers get hot in normal usage because they are not 100% efficient. If a transformer is 95% efficient (possible for a toroid) then approximately 5% of the transformer power delivered will be wasted in heat. Note this is not the power drawn by the load, but the power that the transformer delivers to the circuit. A transformer rated at 20V using capacitor input rectification will produce approximately 26V after rectification, and if the load current is 2A, then the power delivered by the transformer will be 52W. Very small transformers are significantly less efficient (possibly only 75%) which implies that they will get hotter, all other things being equal. If in the above example you purchased a 40VA transformer (2A that you will draw and a 20V winding), then beware you are overrating the transformer. Design criteria for modern transformers are more relaxed than in days of old, as enamelled copper wire is useable to a very high temperature, and the cost of all metals has increased significantly in recent years, which encourages cost-cutting of transformer-core sizes and wire diameters. Magnetizing currents through the laminations generate heat even under no-load conditions. Once again, small transformer get hotter, all other things being equal.

ASSESSING THE INTERNAL TEMPERATURE OF A TRANSFORMER

Sometimes you may wish to know the interior temperature of a transformer, but it is not possible to measure it directly. Logically, the winding nearest the core will usually be hotter than the outer windings. The recommended method (suggested to me many years ago by Dr Sowter of Sowter Transformers, now deceased) was to measure the resistance of the windings when cold, then measure it again once the transformer has been running such that its operating temperature is stable. Knowing that the temperature coefficient of resistance of copper is 3930 parts per million per degree C, (let's call it 4 parts per thousand for simplicity), here is an example calculation:

If the cold resistance is 100W then, for every 1 degree rise in temperature, the resistance will increase by 100 x 4/1000, or 0.4W. If the resistance is 112W when hot, then the temperature is 12/0.4 or 30°C above ambient temperature. Simple when you know how.

RINGING IN THE TRANSFORMER SECONDARY WINDING

Transformers are coils. A transformer secondary winding has, in the rectifiers, a device which turns on and off. Once again, the 1968 4th Edition RSGB handbook (omitted in later editions) indicates most clearly that the effect of the turn-on and turn-off of a transformer secondary winding is the generation of a transient voltage, as shown **Figure 12**. Any transient voltage with a fast rise time (which this most certainly has) will include frequency components extending to very high frequencies. Such components will tend to bypass LC smoothing networks (and to a lesser extent RC networks) as well as regulators, as we have seen their rejection of high frequencies declines quite rapidly. How do we reduce this effect? There were, some time ago, some useful references on the Internet about methods and tools to eliminate this effect. Old references to 'Quasimodo', a pc-based tool to determine the resistance and capacitance required to damp a transformer, still exist but have not been updated for many years. Jensen Transformers (a US firm which contributed a few chapters in the Handbook for Sound Engineering) has also made some most useful contributions. In essence, damping can only be achieved by the placement of a resistor in series with a capacitor across the secondary of a transformer. The 4th Edition of the RSGB Handbook dated 1968 at Figure 17.5 shows how to calculate values (see **Figure 13**). However, I have found these formulas to be of limited use. A good-quality oscilloscope can see the ringing, but it is usually only for a very limited time just after diode conduction, and will require the use of a delayed time base in order to 'walk' along the waveform. Appropriate adjustment of the resistor value (start with, say, 100W and increment by a factor of 5 each time) can be relatively readily achieved. Useful, but very technical, articles can be found at **[9]**, **[10]**, and **[11]**.

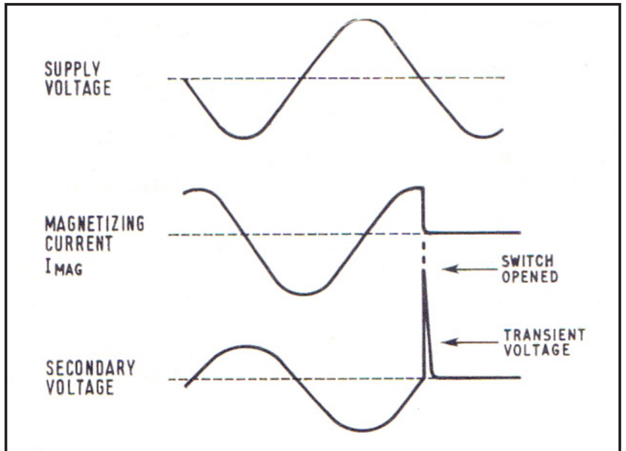

Fig. 17.4. Switch-off transient generated from transformer or choke.

FIGURE 12: Illustrating the generation of the switch-off transient voltage spike.

For use across the secondary:

$$C = \frac{225\ \text{Imag}\ T^2}{V}\ \mu F$$

$$R = \frac{200}{C_{\mu F}}\ \text{ohms}$$

where

Imag = Magnetising r.m.s. current.
For this purpose a value of 10 per cent of the primary or secondary as appropriate may be used.

V = Transformer primary r.m.s. voltage.

T = Transformer $\frac{\text{primary}}{\text{secondary}}$ r.m.s. voltage.

FIGURE 13: Equations for calculating a 'snubber' resistor and capacitor.

RECTIFICATION

This is a surprisingly-complex area, so let's take a simplified view. Clearly, full-wave rectification of a pure sine wave produces a waveform which is high in harmonic content. A pc-based oscilloscope was used to measure the waveform and its frequency spectrum when using a small transformer, bridge rectifier, and low-value capacitor (see

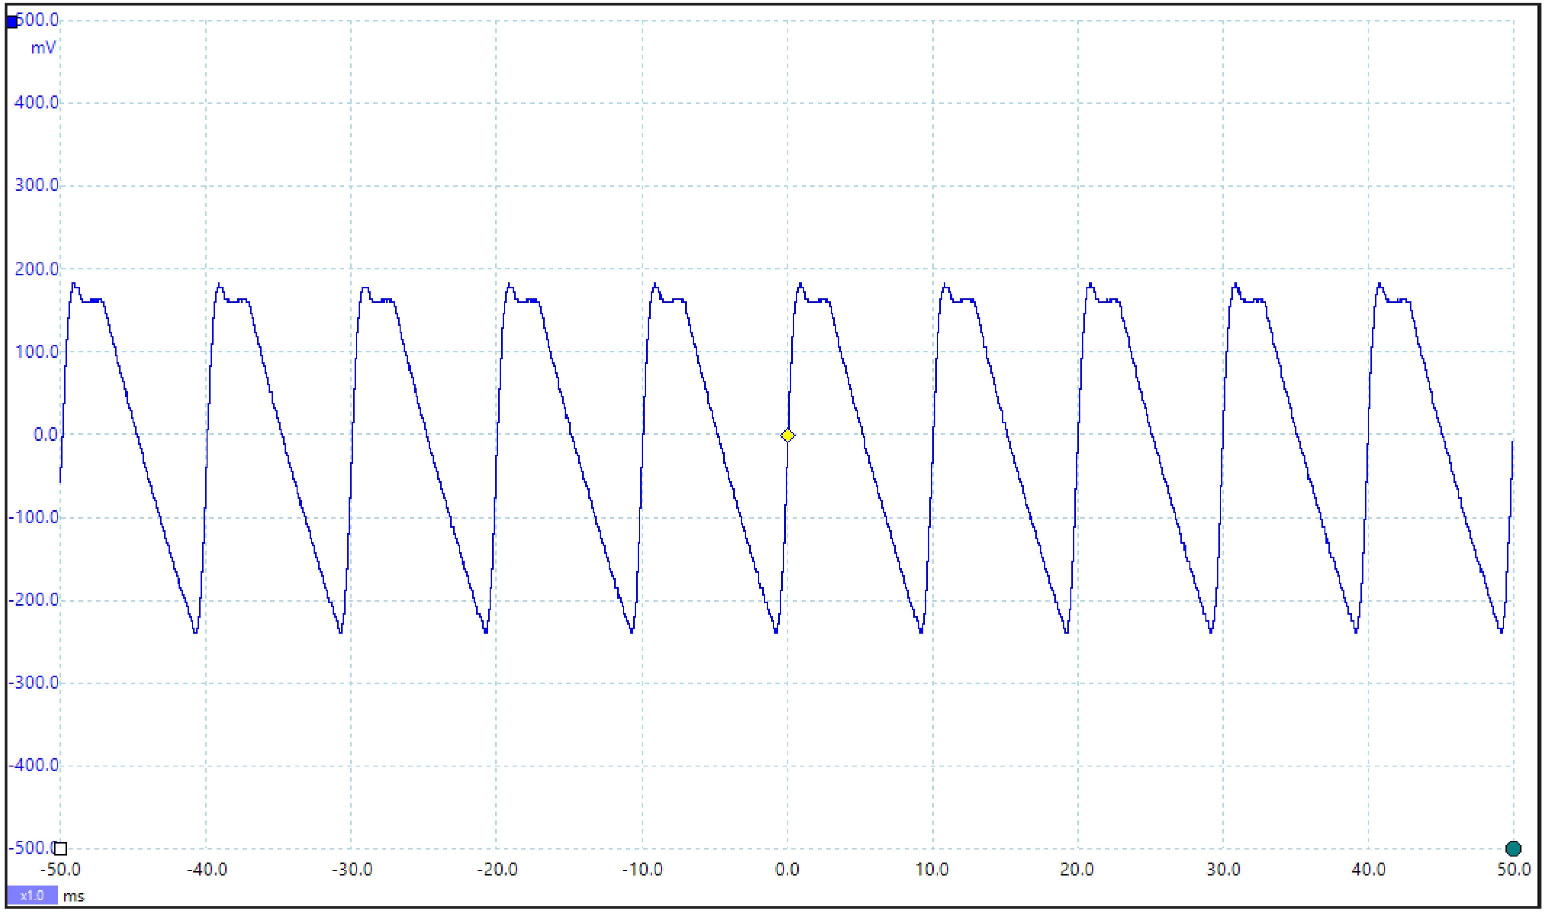

FIGURE 14: A post-rectification waveform.

Figure 14 and **Figure 15**). As can be seen, the waveform is very flat topped (but with a nasty little spike), similar to the mains input waveform earlier. However, the process of rectification produces sharp and pointy-downward peaks, which imply high-order harmonics. There is much power, not just at the 100Hz expected for full-wave rectification, but also up to 6kHz or so. This must all be filtered by subsequent LC and/or RC smoothing circuits. Laminations of transformers also pick up very significant mains hum at 50Hz and its multiples. For this reason, as well as for safety, transformers and chokes need to be earthed wherever possible. I was surprised, when monitoring the body of a transformer, to see the amount picked up as shown in **Figure 16**.

All metal parts that you may possibly come into contact with must also be earthed for safety reasons.

Unsurprisingly, the addition of even a single stage of RC filtering significantly reduces the upper harmonics and general ‘crud’. It is not just the fundamental that is reduced because each RC filter acts as a single-pole low-pass filter, and a choke-capacitor filter acts as a two-pole filter. If practicable, use a low-value series resistor (modelled in PSUDII of course) and a following capacitor as an additional stage of filtering, rather than relying on the input capacitor alone. First decide how much voltage drop you can afford and work from there. As can be seen in the PSUDII plot in **Figure 17**, the waveform of the ripple is not just reduced in amplitude but also its shape is changing, becoming less triangular and hence lower harmonic content. The red line shows the voltage across the reservoir capacitor and the green line that across the output capacitor.

THE EFFECTIVENESS OF LC AND RC SMOOTHING NETWORKS

It appears to be taken as read that a choke reduces AC in a fairly straightforward manner as the frequency rises. Unfortunately, this omits a key characteristic of all chokes, namely their self-resonant frequency. Any choke below its self-resonant frequency will be inductive, and above it will be capacitive, and hence will be less effective as the input frequency rises. So, what might the self-resonant frequency of a choke be and how do we measure it? Let us take two examples:

The Maplin nominal 10H choke, but produced by Danbury, has a measured inductance of 6H (there will be some variability between different units because of winding variability), a resistance of 160Ω, a self-resonant frequency of 6kHz, and hence a self-capacitance of about 120pF.

The Sowter 8970c nominal 10H input choke, has a measured inductance of 8.5H (there will be some variability between different units because of winding variability), a measured resistance 137Ω, a self-resonant frequency of about 5kHz, and hence a self-capacitance of about 120pF.

I made these measurements using my Marconi TF2700 bridge; they are subject to some errors, and may not necessarily replicate

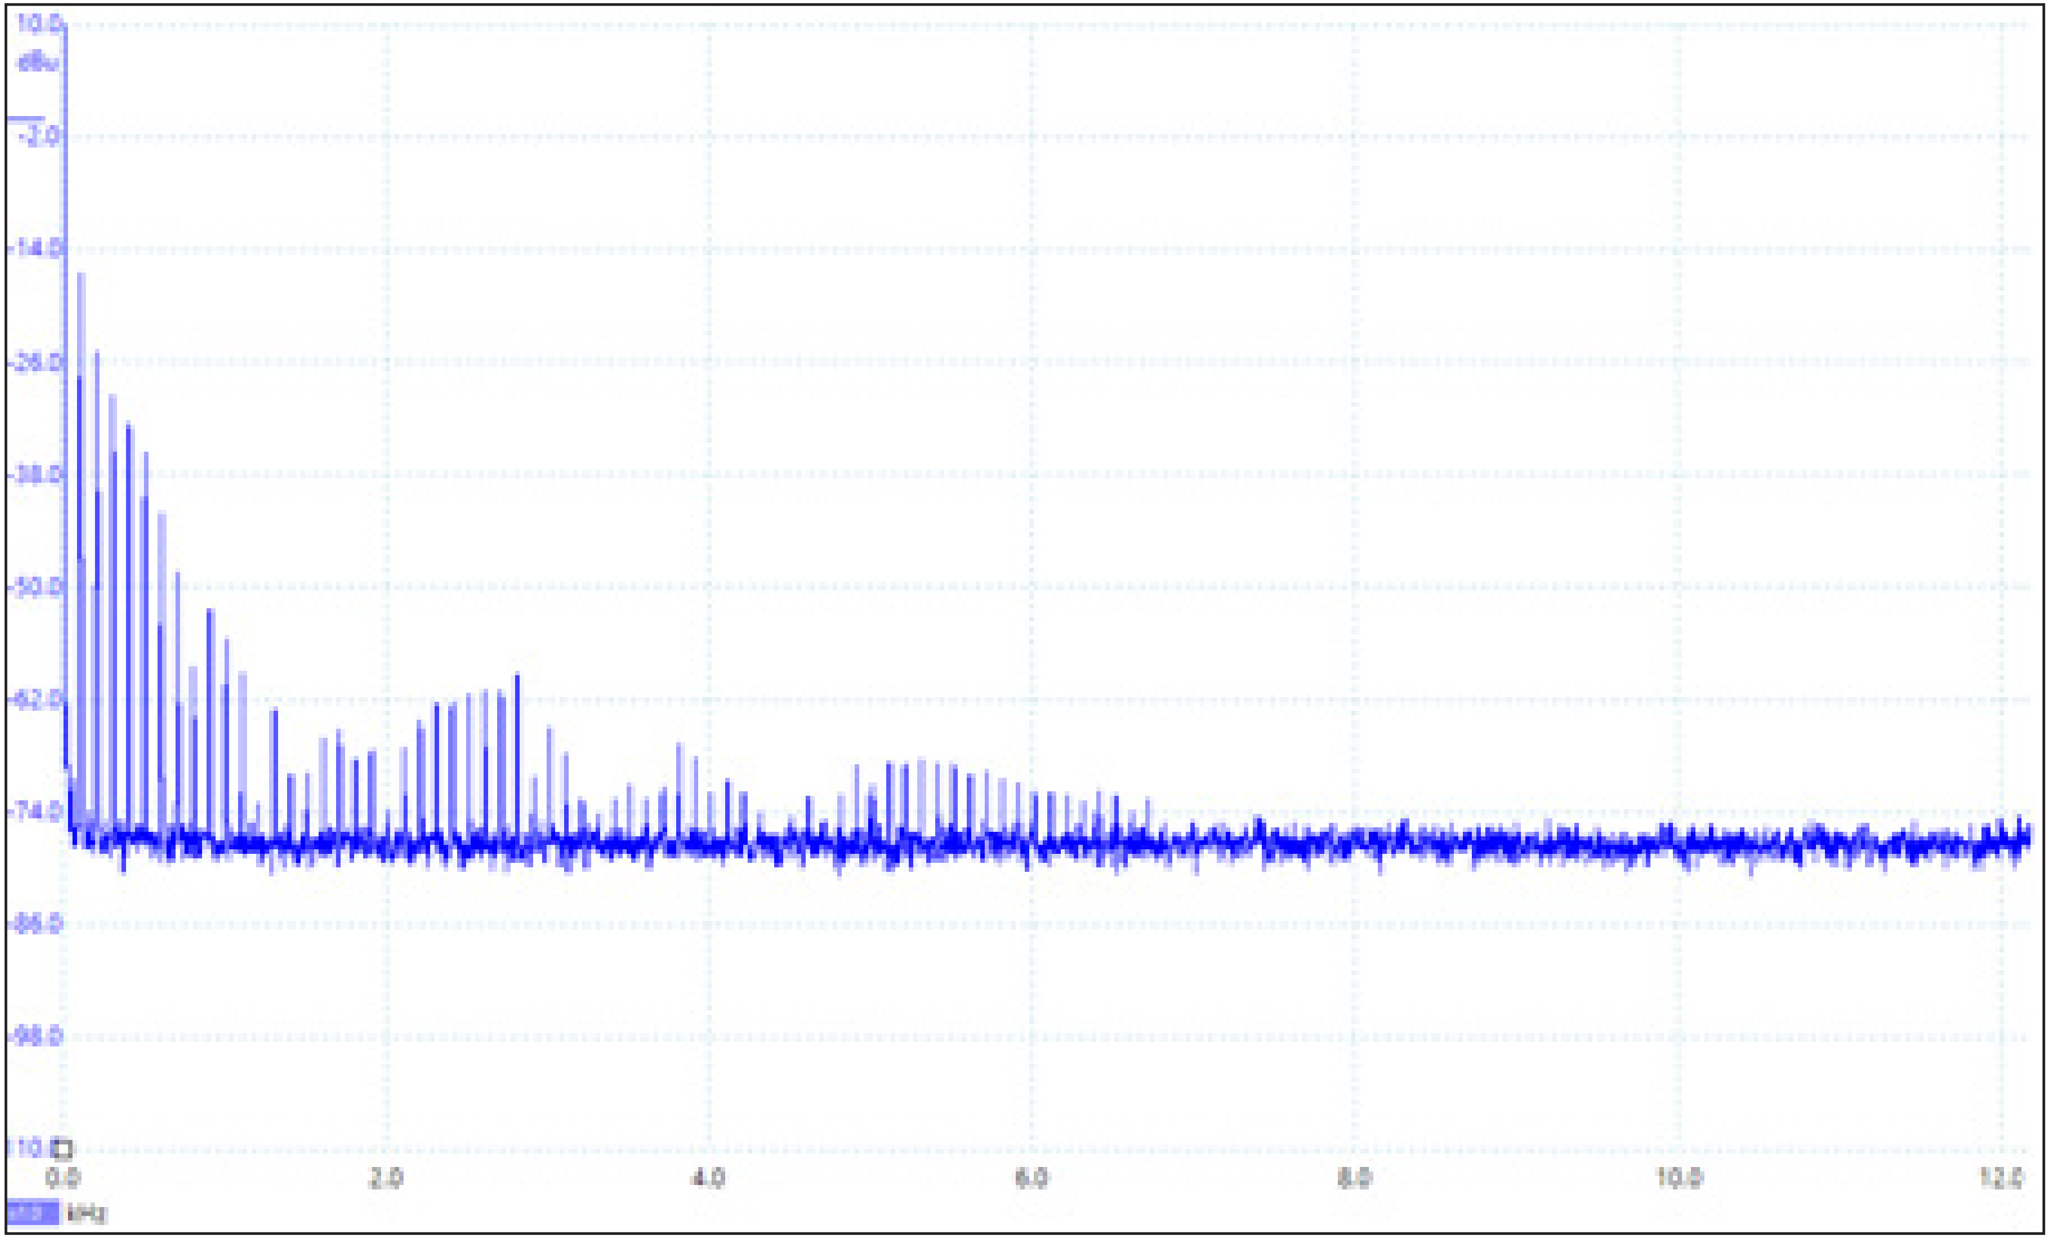

FIGURE 15: The frequency spectrum of the waveform shown in Figure 14.

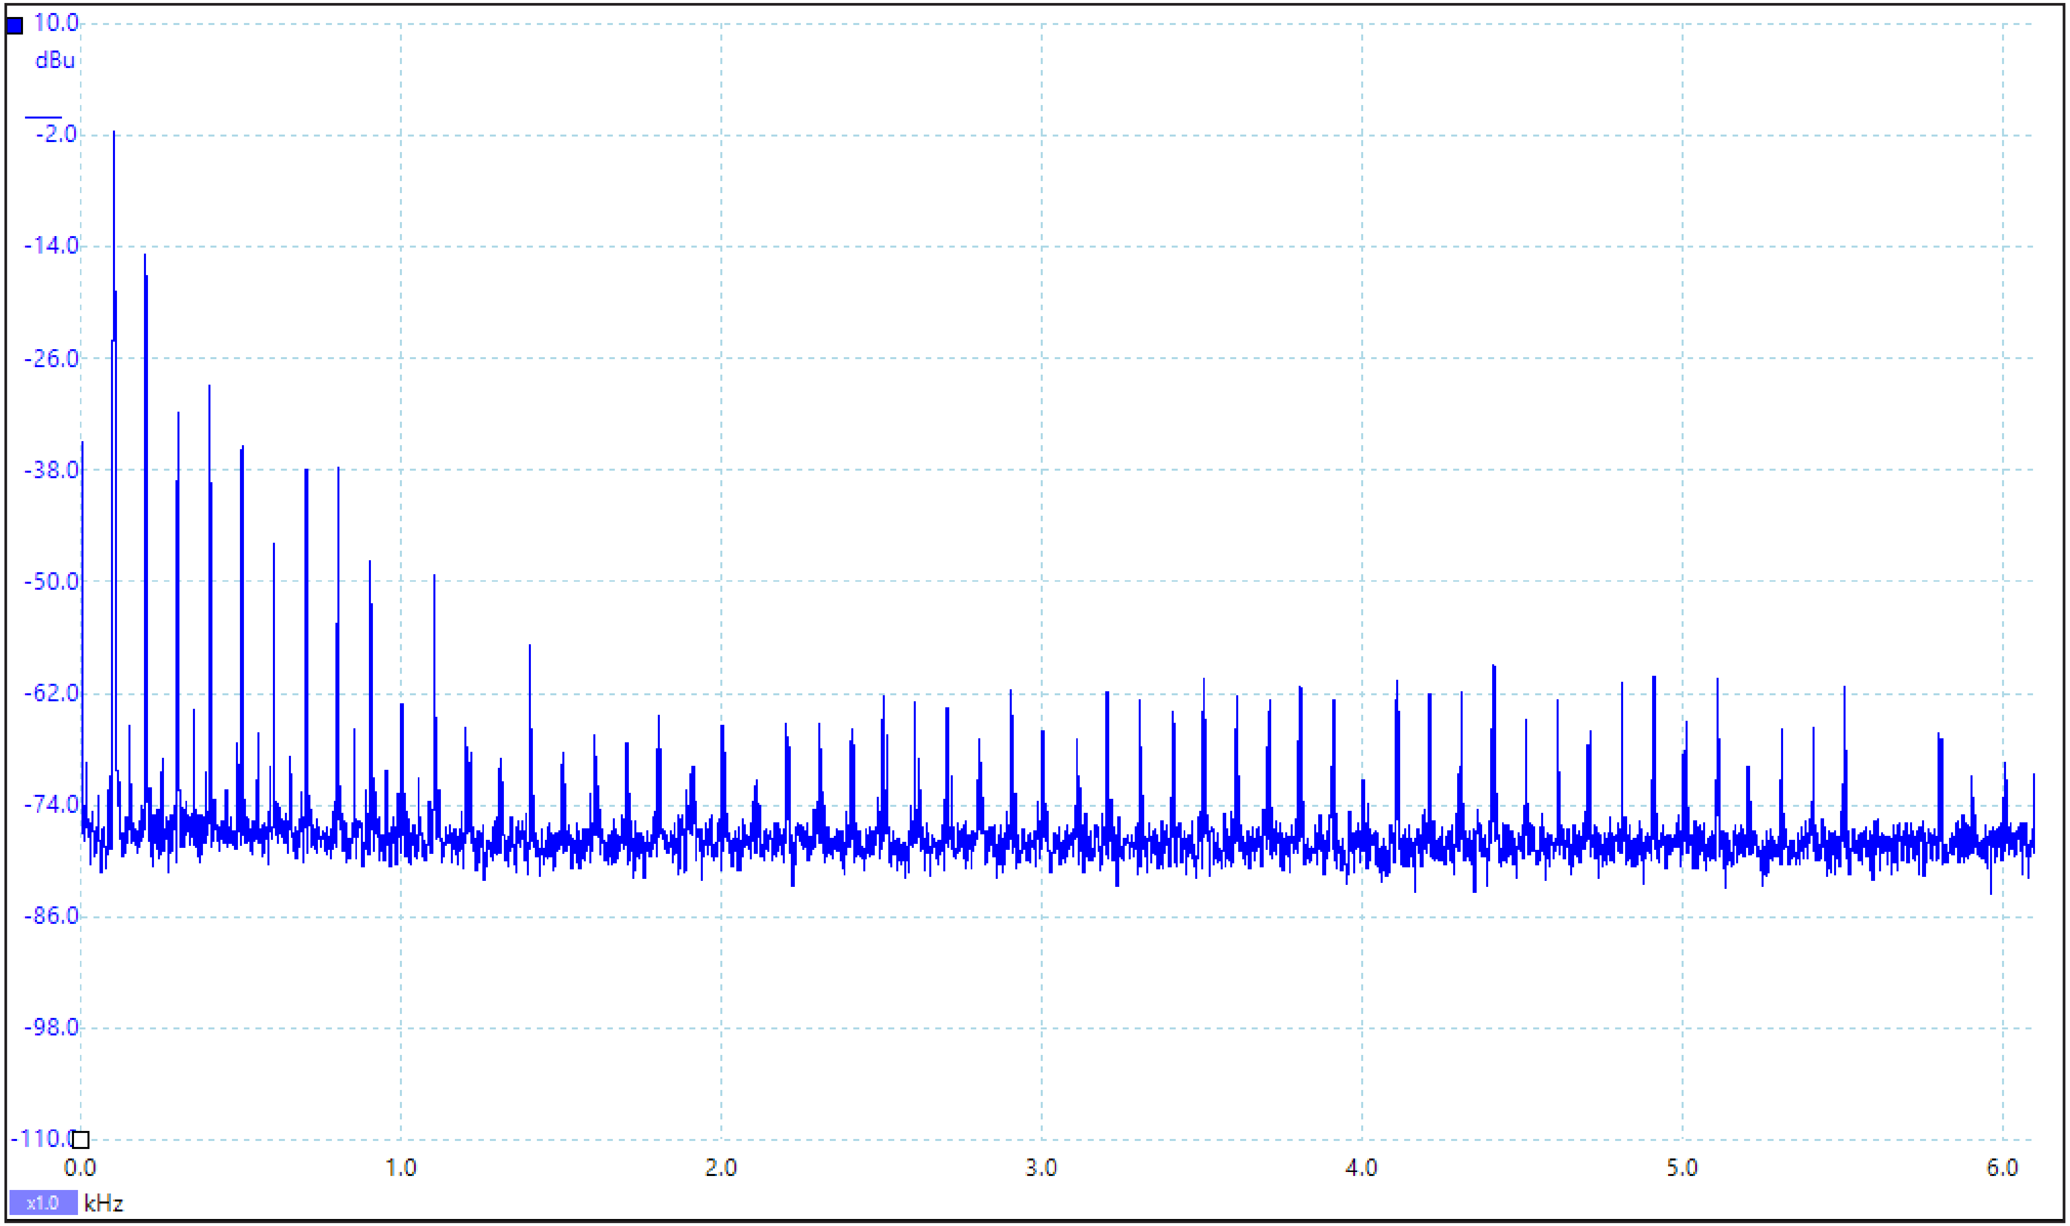

FIGURE 16: The spectrum of a signal picked up from the body of a transformer.

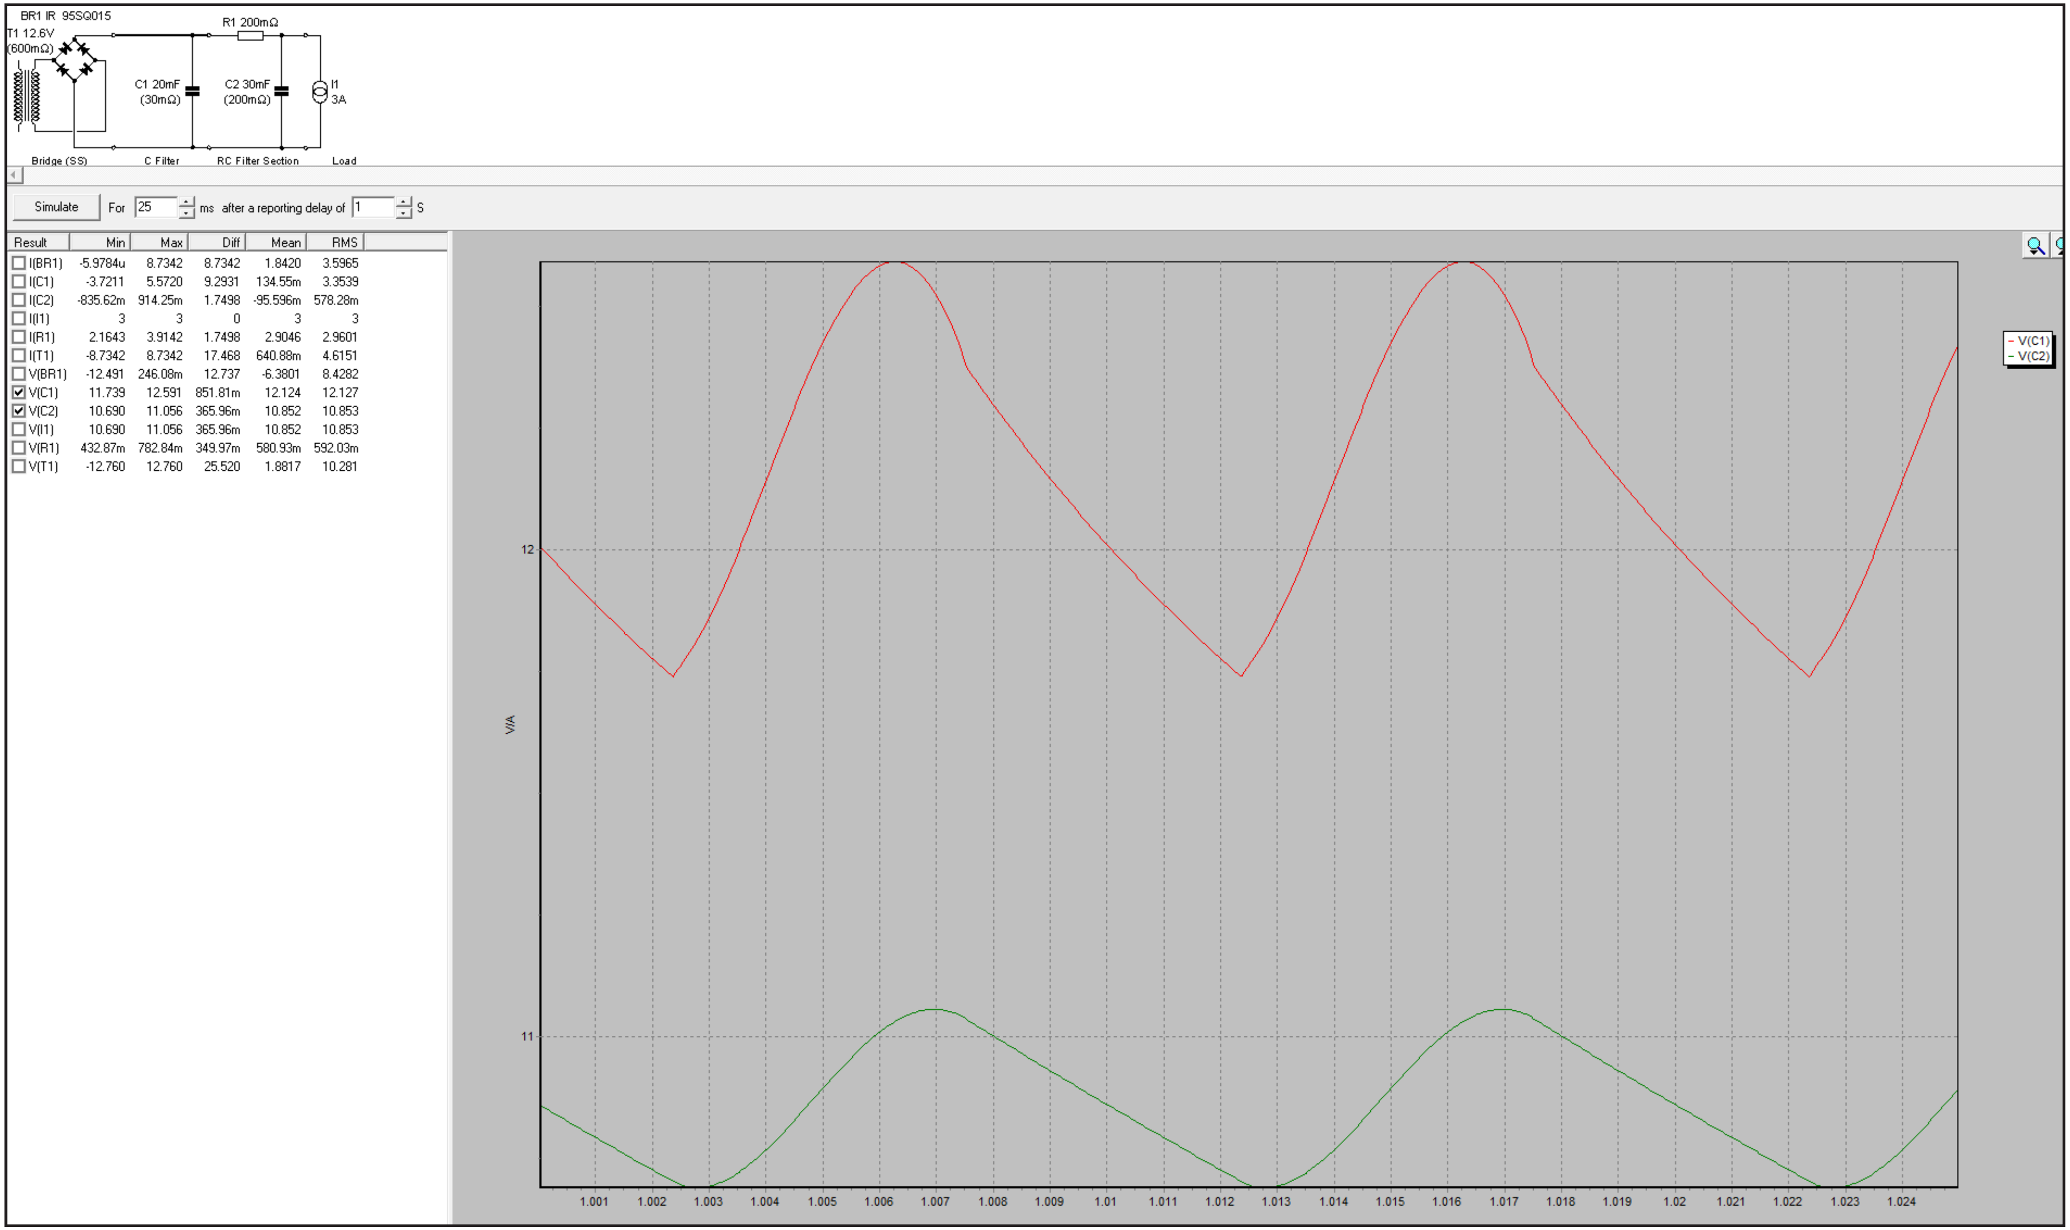

FIGURE 17: The effect of one additional RC network post rectification.

the values when current is passed through them. I have had occasion to measure a few chokes, and found that they all were resonant in the range 2kHz to 10kHz, excepting one that was specially wound. The series distributed capacitance implies that at some relatively low frequency the ripple rejection will have been significantly reduced. In situations where RF or other unwanted frequencies are present on the supply this may need to be considered.

It is very simple to measure the self-resonant frequency approximately; no great precision is required as the resonance curve is very broad because of the low Q as a result of the resistance of the winding. **Figure 18** shows the method. Sweep the audio signal generator to find the response, and note the frequency.

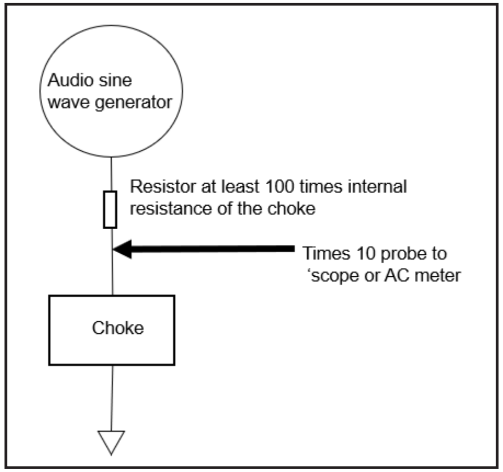

FIGURE 18: Measuring the resonant frequency of a choke.

A capacitor input filter produces an output voltage of approximately the peak AC output voltage (rms value times 1.4) minus any voltage loss across the rectifiers and corresponding transformer resistance. However, when high currents and/or low capacitor values are employed, the ripple voltage can be substantial, and hence the useable voltage (ie the minimum voltage at the last capacitor) will be somewhat less than this. That is why modelling is important to ensure that the lowest output voltage is still sufficient for the correct operation of the following circuit.

Another type of rectification and smoothing circuit, namely the choke input filter, where the rectifiers are followed by a choke THEN a capacitor, has very different operating characteristics. The RSGB Handbooks version 3 onwards sets out in more detail the operating characteristics of this type of circuit, which once again can be modelled successfully in PSUDII. So, when would one use such a circuit? Historically they were used where a 'stiffer' supply rail was required eg the HT supply in valve modulators and valve SSB transmitters. They have better regulation of the output voltage versus output current, ie they are 'stiffer', but that comes with significant cost namely, including the financial cost of the choke and its corresponding weight/size, and the reduced output voltage which is typically only 0.9 times the applied AC input voltage excluding the voltage drop caused by the resistance of the choke. That is quite a penalty! Mathematical derivation of the 0.9 figure is beyond the scope of this article. Modelling will show that the current through the rectifiers is almost constant, rather than the strong peak currents typical of capacitor input filters. That might be important if the AC supply has limited peak current capability. One caveat is that this circuit requires a minimum current to function as a choke input design, and bleeder resistors were commonly used to ensure this condition. Failure to provide the critical current results in the circuit behaving like a slightly-modified capacitor input filter with the resultant higher output voltage. Another factor to be aware of is that the choke may see much higher AC voltages across it than usual, and the choke should be selected as designed for choke input filters. This is particularly the case for higher voltages such as found in valve equipment. Further information on the operation of this type of circuit may be found at https://www.ee-diary.com/2022/11/how-choke-input-filter-works.html amongst other sources.

Type SLPX 85 °C Snap-In Aluminum Electrolytic

Best Value 85 °C Snap-In Type

Cap (µF)	3000 h @ 85 °C Catalog Part Number	Max 25 °C ESR (Ω) 120 Hz	Max 25 °C ESR (Ω) 20kHz	Max 85 °C Ripple (A_{rms}) 120 Hz	Max 85 °C Ripple (A_{rms}) 20kHz	Nominal Size (DxL) (mm)
63 Vdc (79 Vdc Surge)						
3300	SLPX332M063A7P3	0.121	0.091	2.72	3.40	22 x 40
3300	SLPX332M063C3P3	0.121	0.091	2.74	3.43	25 x 30
3300	SLPX332M063E1P3	0.121	0.091	2.78	3.48	30 x 25
3900	SLPX392M063A4P3	0.102	0.077	3.09	3.86	22 x 45
3900	SLPX392M063C5P3	0.102	0.077	3.13	3.91	25 x 35
3900	SLPX392M063E3P3	0.102	0.077	3.09	3.86	30 x 30

Cap (µF)	3000 h @ 85 °C Catalog Part Number	Max 25 °C ESR (Ω) 120 Hz	Max 25 °C ESR (Ω) 20kHz	Max 85 °C Ripple (A_{rms}) 120 Hz	Max 85 °C Ripple (A_{rms}) 20kHz	Nominal Size (DxL) (mm)
80 Vdc (100 Vdc Surge)						
8200	SLPX822M080H4P3	0.040	0.030	5.89	7.36	35 x 45
10000	SLPX103M080H9P3	0.033	0.025	6.63	8.29	35 x 50
100 Vdc (125 Vdc Surge)						
820	SLPX821M100A1P3	0.324	0.243	1.86	2.33	22 x 25
1000	SLPX102M100A3P3	0.265	0.199	2.02	2.53	22 x 30
1200	SLPX122M100A3P3	0.221	0.166	2.12	2.65	22 x 30

FIGURE 19: Capacitor properties table.

SAFETY

As mains voltages are present in the primary circuits, the usual safety precautions apply about handling high voltages. Be aware, too, that high-current secondary windings and supplies can deliver very nasty burns as they can get very hot. I recommend the removal of watches and rings before carrying out any live measurements. Multiple amps through a finger ring could give you a painful injury.

HOW TO READ CAPACITOR INFORMATION

Figure 19 shows an example of a section of a capacitor data sheet. The important columns are the capacitance, working voltage and, of particular concern when modelling using pspice or PSUDII, the ESR, as we have seen that this value has a significant impact on ripple.

TEMPERATURE AND ITS EFFECT ON LONGEVITY

The ARRL handbook 2019, Chapter 4.11, is devoted to heat management, and I recommend that you read it and other sources of advice. As the temperature of a semiconductor device rises, its lifespan decreases. Manufacturers use accelerated life test at high temperatures, and then predict the approximate life expectancy in normal use. The graph shown in **Figure 20**, taken for a GaN device, shows how important the management of device temperature is to its longevity, or the mean time to failure (MTTF). As a rule of thumb, the lifespan halves for every ten degree C increase in temperature. Therefore, good thermal management is critical.

SIZING A HEAT SINK, AND WHY HEATING EFFECTS ARE IMPORTANT

As I can testify, having demonstrated the effects of heat sinks and enclosures; this is a topic not well understood by many radio amateurs. Let us make a few simple and crude experiments to illustrate the effects of an enclosure and air movement on temperature. In the experiments all temperatures were measured using a contact platinum type thermometer.

In practice we need to design for the hottest surrounding temperature that we are likely to meet in practice. Given that temperatures as high at 40°C and more can occur during the summer, designing for a maximum ambient

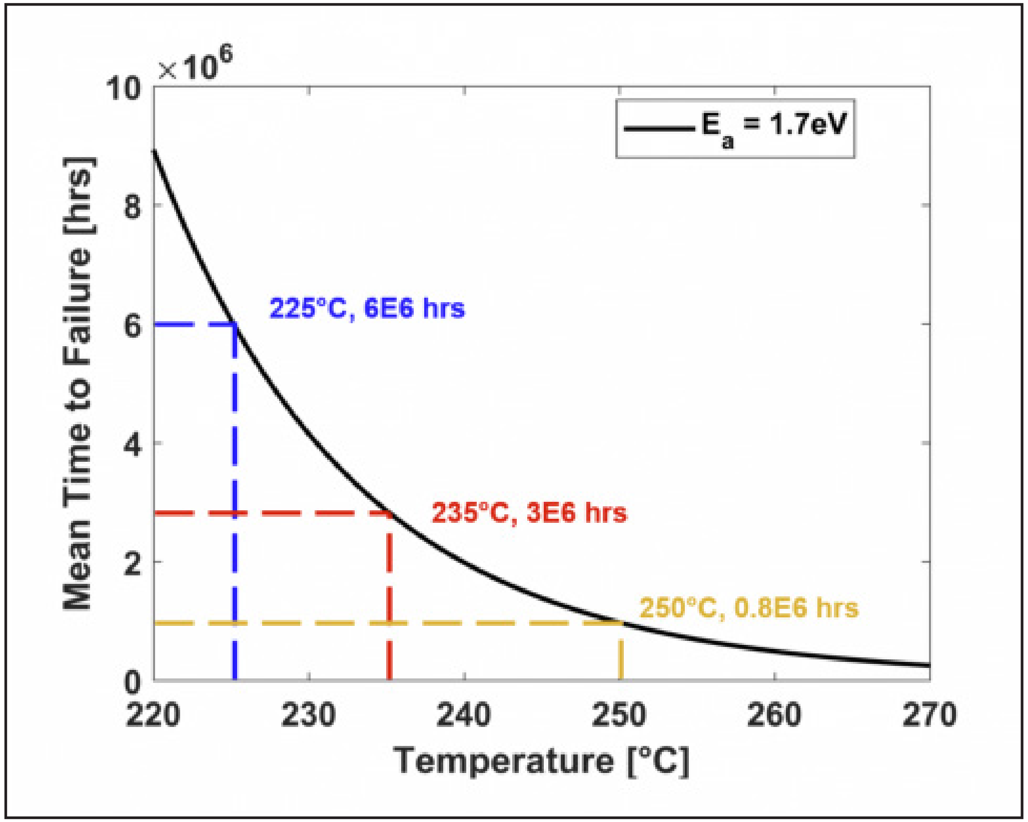

FIGURE 20: The longevity of a GaN device as a function of operating temperature.

Typical Thermal Resistance

Package	θ_{ja}
TO-92	132°C/W
TO-220	29°C/W
TO-243AA (SOT-89)	133°C/W

FIGURE 21: DN2540 temperature specifications.

temperature of 50°C would seem appropriate.

It is critical to remember that Silicon has a low melting point such that devices can melt at 180°C to 200°C . The table in **Figure 21** sets out the thermal conduction characteristics for the three package types of the constant-current source based on the FET type DN2540, which I used for the tests. The subscript JA is the junction-to-ambient thermal resistance as measured under the standard conditions as set out by standards body JEDEC **[12]**. Whilst this represents a standard and defined method for testing, it is not meant to be representative of actual usage. Indeed, a graph in [12] shows that, for two devices in a given situation, JA can vary between 20°C/W and 100°C/W, dependent upon the size of the pcb! Other measures such as JC (junction-to-case) are available, but for a given device either JA or JC should be specified.

Let's begin by understanding the numbers, then move onto heat sinks. Clearly this may not represent our use case where the device will probably be mounted on a pcb and in an enclosure. For the DN2540 TO220 package, the figure states that, for every watt dissipated, the junction temperature will increase by 29°C in this ideal JEDEC defined situation. Smaller packaged versions, such as the TO-92 and SOT-89, are considerably worse such that dissipation of significantly less than 1W is likely to result in device failure in just about any practical situation.

An AVIS SW38-4 heat sink, mounted on a small pcb, was used for additional tests. The data sheet shown in **Figure 22** describes how the rating changes as air-flow increases. Unfortunately, we are unlikely to know the air speed inside an enclosure; indeed, it would probably be difficult to measure in any practical situation. A rating of 10.2°C/W is given again under a somewhat idealized situation. Heat sinks are available with ratings in the low single figure °C/W but they will be large, heavy, and expensive. In practice, forced-air cooling is usually a cheaper and more-effective way of achieving the desired goal of significant heat dissipation. My Alinco DM340-MW 30A linear power supply uses just this approach. A constant current circuit was used (it was in the junk box) such that, at 15V, a current of 78mA was passed, giving a dissipation of approximately 1.2W. The DN2540 and heat sink on the bench gave a heat-sink temperature (measured close to the TO220 tab) some 20°C above ambient. When the circuit was place in a small cardboard box, the temperature of the heat sink was 48°C above ambient. The, when a 40mm Noctua fan was directed to produce air flow across the heat sink, its temperature was only 7°C above ambient. This indicates the importance of air flow in reducing temperatures and Noctua fans are ideal as they are very quiet.

A constant current circuit was used (it was in the junk box) such that, at 15V, a current of 78mA was passed, giving a dissipation of approximately 1.2W. When not attached to a heat sink, the tab temperature was a 68°C above ambient, much hotter than I would wish to run this type of device. Remember also that the junction temperature will exceed the tab temperature.

As can be seen, the impact of a fan is very significant, so much so that a fan is almost certain to be the best approach, rather than spending more money, real estate, and weight on a larger heat sink. Remember also that the diodes, chokes, and resistors in the smoothing network and/or stabilizer pass elements, will also dissipate heat which adds to the overall heat and hence the temperature burden in an enclosure. Enclosures take some time, possibly an hour or more, to reach a stable operating temperature, so it is important to allow enough time to reach equilibrium before taking temperature measurements. The temperature is also dependent on the material of the box, metal being a much better conductor of heat than plastic. Therefore, the results taken for a heat sink inside a particular enclosure cannot be extrapolated to other enclosures.

A Noctua fan was used, once again in the junk box, as these are the 'goto' fans I use for all projects where a fan is required. They are incredibly quiet, far more so than the typical fan. Indeed, I usually replace fans in my networking/ PC equipment with Noctua products in order to have peace and quiet in the shack.

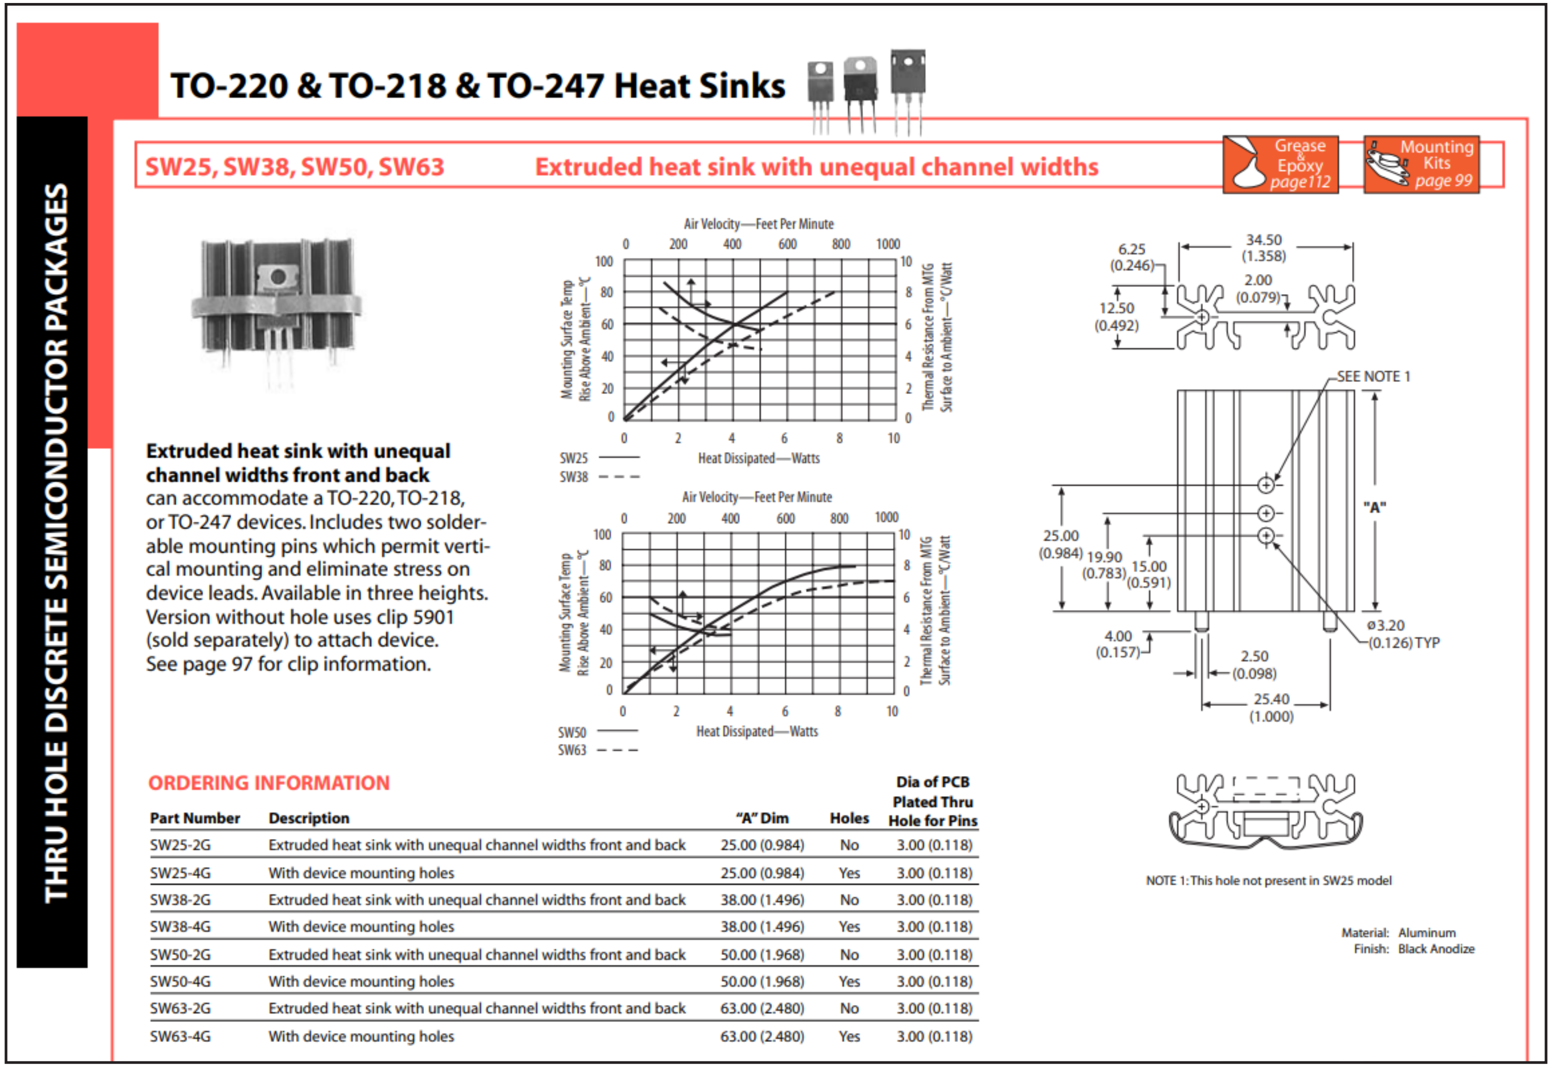

Extruded heat sink with unequal channel widths front and back can accommodate a TO-220, TO-218, or TO-247 devices. Includes two solderable mounting pins which permit vertical mounting and eliminate stress on device leads. Available in three heights. Version without hole uses clip 5901 (sold separately) to attach device. See page 97 for clip information.

ORDERING INFORMATION

Part Number	Description	"A" Dim	Holes	Dia of PCB Plated Thru Hole for Pins
SW25-2G	Extruded heat sink with unequal channel widths front and back	25.00 (0.984)	No	3.00 (0.118)
SW25-4G	With device mounting holes	25.00 (0.984)	Yes	3.00 (0.118)
SW38-2G	Extruded heat sink with unequal channel widths front and back	38.00 (1.496)	No	3.00 (0.118)
SW38-4G	With device mounting holes	38.00 (1.496)	Yes	3.00 (0.118)
SW50-2G	Extruded heat sink with unequal channel widths front and back	50.00 (1.968)	No	3.00 (0.118)
SW50-4G	With device mounting holes	50.00 (1.968)	Yes	3.00 (0.118)
SW63-2G	Extruded heat sink with unequal channel widths front and back	63.00 (2.480)	No	3.00 (0.118)
SW63-4G	With device mounting holes	63.00 (2.480)	Yes	3.00 (0.118)

FIGURE 22: Heatsink air-flow data.

One last point to consider is the electrical isolation of the heat sink. Thermal washers may need be applied in order that the heat sink does not short out the supply to the chassis (which must be earthed if it is possible to touch it in normal usage) or, worse still, come into contact with a high voltage.

CONCLUSION.

I hope that this article gives readers the confidence to design their own power supplies with greater certainty that the result will live up to their expectations.

REFERENCES

1. https://groups.io/g/duncanampspsud
2. https://vivadifferences.com/7-difference-between-center-tapped-and-bridge-rectifier/
3. https://www.allaboutcircuits.com/technical-articles/fast-ultrafast-soft-standard-schottky-selecting-the-right-rectifier/
4. https://pinkfishmedia.net/forum/threads/whats-the-best-diy-dc-blocker-circuit.261765/
5. https://www.energynetworks.org/customers/find-my-network-operator
6. https://www.mainsfrequency.com/
7. https://gridwatch.co.uk/frequency
8. https://www.we-online.com/en/components/products/WE-CMB
9. http://www.hagtech.com/pdf/snubber.pdf
10. https://www.diyaudio.com/community/threads/simple-no-math-transformer-snubber-using-quasimodo-test-jig.243100/
11. file:///D:/Downloads/Quasimodo_jig_revA.pdf
12. https://www.allaboutcircuits.com/technical-articles/theta-ja-junction-to-ambient-thermal-resistance-thermal-performance-calculation/