

High performance common-mode chokes

There can't be many components used by radio amateurs that are more misunderstood than the humble common-mode choke. Many amateurs still haven't even heard of it and even those who have can't agree on its name: is it a common-mode choke, an RF choke, a 1:1 current balun or a 1:1 Guanella balun? Since these are all alternative names for essentially the same device – the name changes mostly with the application – the best answer is probably: 'Any or all of the above'.

I shall generally refer to it as a common-mode choke (CM choke for short), and although very similar chokes can be used for control cables, phone lines etc, the scope of this article will be limited to chokes used in typical HF antenna systems.

This article is an attempt to shed some light on this misunderstood component, covering:

- What a CM choke does
- Why and where we might need one
- What properties it should have
- How to build a good one; and
- How to measure its performance.

But first we need to understand what 'common-mode current' is, and why we usually want to prevent it.

COMMON-MODE CURRENT.

Figure 1 shows a generator connected to a load impedance Z_L via a 2-wire transmission line. One side of the generator is connected to ground, but the load is completely 'floating'.

The basic concept of an electrical circuit tells us that the 'outgoing' current I_a and the 'return' current I_b are in fact the same current, flowing around the loop. There are no



PHOTO 1: A 7-turn broadband common-mode choke.




FIGURE 1: Balanced transmission line from a generator to a 'floating' load.




FIGURE 2: Common-mode current I_c creates an unwanted alternative return path.

alternative paths for that current to take. Because I_a and I_b are flowing along the two legs of the transmission line in opposite directions, there is a phase difference of 180° between them.

Equal and opposite currents in the two conductors of a transmission line are called

differential-mode currents. Since the two conductors are very close together, the external fields from the two opposing currents cancel out and the line does not radiate. These are the defining features of differential-mode currents and of a balanced transmission line.

Suppose we now change the conditions at the load by introducing a path to ground, Z_c , from the upper leg of the transmission line as shown in Figure 2. There are now two different paths that the return current can take. As well as the normal return current along the lower leg of the line, a new current I_c can now flow back to the generator via Z_c and earth. Notice that I_c flows into Z_c and then down to earth at the load end; and since currents flow in closed loops, that same current I_c will also emerge from the earth terminal at the generator as shown in Figure 2.

I_c is what we refer to as the common-mode current. Any current, I_c , that takes this alternative return path must result in the transmission line currents I_a and I_b being unequal – in other words, the transmission line is no longer balanced. To some extent, this is always the situation in real life.

To distinguish more clearly between differential-mode and common-mode currents that are coexisting on the same transmission line, it is often convenient to think of the unequal currents I_a and I_b as the superposition of a 'pure' differential-mode signal and a separate




FIGURE 3: Skin Effect makes coax behave as a THREE-conductor transmission line.

common-mode signal. The common-mode current I_c is simply equal to $(I_a - I_b)$. A little more algebra then shows that the differential-mode (balanced) currents in Figure 2 must each be equal to $(I_a + I_b)/2$.

If all that sounds strange and unfamiliar, it shouldn't, because Figures 1 and 2 also show some basic principles of UK domestic mains supplies. The upper leg of the transmission line is the 'Live' conductor and the lower line leg is Neutral. Figure 1 shows the normal situation where all the current returns in the Neutral leg – in other words, the two-conductor mains cable is intended to function as differential-mode transmission line. Figure 2 then shows a fault condition that allows a 'leakage current' I_c to flow directly from Live to ground rather than returning via Neutral.

For safety, we typically employ a Residual Current Device (RCD) to detect that leakage current by continuously monitoring for unequal Live and Neutral currents. Now that we are transmission line experts, we see that the RCD is really a common-mode current detector!

A defining feature of common-mode current in both mains and RF systems is that common-mode currents do not cancel out. Common-mode currents carried by a multiple-conductor line will always reinforce each other so that, regardless of the actual number of conductors, the cable behaves as if it were a single wire... and in RF systems, that means common-mode current will cause radiation. In order to return via the ground terminal at the generator, the common-mode current flowing through Z_c in Figure 2 will flow along any conductors it can find that lead towards ground. This makes unwanted

common-mode currents a prolific source of both radiated and conducted interference.

FUNNY STUFF, COAX! When we use coax for our transmission line, the two wires of Figure 1 and Figure 2 become the centre conductor and the shield (Figure 3). But because the centre conductor is totally enclosed by the shield, the electric and magnetic fields from the two conductors are extremely tightly coupled together, and physics then demands that the currents on the centre conductor

and the inside of the shield are equal and opposite. Unlike the parallel-wire line in Figures 1 and 2, the interior of a coaxial transmission line will only support differential-mode currents.

Something called the Skin Effect means that RF currents flow only along the surfaces of conductors and never through the thickness. More accurately, the current density flowing along a conductor is much higher at the surface than within its bulk; for example at 10MHz in a solid copper conductor the current density at a depth of just $20\mu\text{m}$ has fallen to 37% of its value at the surface. This means that in a coaxial line the transmission-line current is flowing almost exclusively on the outside surface of the inner conductor (I_a , shaded green in Figure 3) and on the inner surface of the braid (I_b , shaded blue).

Because RF current does not flow through the thickness of the shield, the interior of a coaxial cable is completely 'private' from the outside world – and that is what gives coax its unique tolerance to the way it is installed. Bend it, bury it, soak it in water, pass it through walls and metal bulkheads, or even wind the cable around a magnetic core – none of those actions makes any difference to the pure differential-mode currents in that private interior world. At first sight, therefore, we might think that coaxial line cannot possibly carry common-mode current... but unfortunately we'd be wrong!

The Skin Effect is a very important concept that unlocks many mysteries, but not all of its consequences are welcome. Because the shield has two surfaces, an outside as well as an inside, a completely separate current (I_c , shaded red in Figure 3) can flow along the outside surface.

Let me repeat that: the current, I_c , on the outside surface of the shield is completely independent of I_b on the inside surface. The two currents are completely separate and do not interact in any way. So, in fact, a coaxial cable is a 'three-conductor line' with a third independent conductor on its outside surface! This is hard to grasp for anyone who is new to RF engineering, but the Skin Effect is a dependable fact and so are all its consequences.

Comparing coax (Figure 3) with parallel-wire line (Figures 1 and 2), we now see that each of them can support both differential-mode and common-mode currents at the same time, but that they achieve it in different ways. The parallel line requires the two modes to share the same two conductors; but coax does it differently, by keeping the two modes completely separate – the pure differential mode on the inside and the common mode on the outside. We can always picture coax as a completely balanced internal pipeline for delivering RF power, overlaid by an entirely separate outside conductor that behaves as a single large-diameter wire. Because that's exactly what coax is. Once we grasp this key characteristic of coaxial cable, the solutions to many antenna common-mode problems become much more obvious.

Now we are ready to consider some practical engineering.

WHAT ARE CM CHOKES AND WHY DO WE NEED THEM? Quite simply, a common-mode choke suppresses or 'blocks' the common-mode current. And in doing so, it forces the currents at its output terminals to be balanced – equal in magnitude and opposite in phase – no matter how unbalanced the load or source to which it is connected may be. Because of this useful property, it frequently finds application interfacing balanced systems to unbalanced systems; no surprise then that it often gets called a balun!

Baluns in general come in many different types. CM chokes are the type that force current balance at their output terminals, so they are also sometimes called 'current baluns'. But there are also completely different types of balun that force voltage balance instead. Some baluns also perform an impedance matching function (for example 4:1 or 9:1) but this is not an essential part of what baluns do. Because our CM choke enforces current balance but does not perform an impedance matching function, within the wider family of baluns it can legitimately be called a '1:1 current balun'. And because Guanella first identified these properties it's also legitimate to call it a '1:1 Guanella balun'.

We will now focus on two typical amateur applications where we might need a CM choke: firstly, connecting a balanced dipole or beam antenna to unbalanced coaxial cable; and secondly, connecting the balanced feedline from a multiband doublet to an unbalanced tuner.

DIPOLE TO COAX. Figure 4 shows a detailed view of the connection between a half-wave dipole and its coaxial feedline, but with no balun.

The centre conductor of the coax is connected to the right-hand leg of the dipole and the shield to the left-hand leg. We have seen that the pure differential mode inside the shield means that the current flowing on the centre conductor of the coax (A) must be of equal magnitude and opposite sign to the current flowing along the inside surface of the braid. However, on emerging from the open end of the shield, that current from the inside of the shield is forced to divide itself between two possible paths: into the left-hand dipole leg (B) or back along the outside surface of the braid (C) as the separate common-mode current.

Kirchhoff's Current Law tells us that $(B+C)$ must equal the inside braid current, which is a differential-mode current equal and opposite to A. So if there is any common-mode current (C) at all, it must mean that the currents (A) and (B) in the dipole are unbalanced. That is why we need a balun – to block that common-mode current and force the dipole into a better state of balance.

The amount of common-mode current that flows back along the outside surface of the braid, compared to that flowing into the left-hand dipole leg, will depend on the relative impedances of those two different current paths. The impedance of the 'wanted' path into the antenna will be one-half of the dipole's feedpoint impedance, ie something in the region of 20-40Ω.

However, the impedance of the common-mode path impedance at the feedpoint could easily be quite similar, or even lower, depending on the combination of the coax length, its routing and its grounding. If the impedances are roughly equal, then about half of the current will be 'robbed' from the dipole and will flow down the outside of the coax instead. If the common-mode impedance is lower still, it's possible for current (C) to be much greater than (B) – resulting in gross imbalance of the dipole and severe radiation from the feedline.




FIGURE 4: Coax-fed dipole with no balun.




FIGURE 5: NEC model of a severely unbalanced coaxial feed.

Figure 5 shows a computer model of just such a situation, and the severely unbalanced current distribution that will result. Green wires 1 and 2 represent the two legs of a 20m dipole erected 36ft above ground; green wire 3 represents the outside of the coax shield, which is connected to the left-hand dipole leg. This coax shield is well grounded at the bottom end in the 'shack' (not shown).

The source of RF in this computer model is exactly where it is in real life: at the top of the feedline where the transmission-line currents emerge from the internal privacy of the coax. Then the mauve lines in Figure 5 indicate the magnitudes of the currents along the three wires. Because the length of the coax is close to a half-wavelength, the impedance of the common-mode path as seen at the feedpoint is very low, so most of the current that should have been flowing into the left-hand leg of the dipole is instead flowing back along the outside of the coax!

Does it matter? Well, yes it does – the coax shield is acting as a third radiating wire, inadvertently turning our dipole into a 'tripole'! And because the coax shield at the shack end is carrying a high surface current, it will radiate just like any antenna element; so we can expect RF feedback in the transmitter and also RF current flowing back along the domestic mains wiring, creating a strong possibility of interference to other nearby equipment. Of equal concern, during reception that unwanted 'third antenna wire' can pick up noise from the mains, and from other equipment in the shack, the house and the general surroundings, all of which then flows up the outside of the coax to the feedpoint and is injected into the signal path. This can raise the receiver noise level by large amounts (often above S9) and is a frequent cause of complaints about "my noisy QTH". Thirdly, the unequal currents flowing into the two dipole legs will distort the radiation pattern; whilst that




FIGURE 6: Three-terminal representation of a partly unbalanced doublet/ladderline system, as seen from the shack.

might not worry us unduly when using a dipole, it can be a serious problem when using a beam.

Ideally we want the common-mode current (C) in Figure 4 to be zero so that the currents (A) and (B) are equal and the dipole's balance is retained. We can achieve that by deliberately inserting a large impedance in the path of the common-mode current, forcing the dipole back towards its ideal balanced condition... and that is the role of the CM choke.

When we force a change towards a more ideal situation, it will also force a change in the voltages and currents at the feedpoint itself. Therefore you can expect the VSWR to change as well. This is completely normal – that's what almost always happens when you change from one antenna to another. And that is what you just did: the CM choke has literally changed the whole antenna, from the malfunctioning 'tripole' of Figure 5 into a properly functioning dipole that you have now.




FIGURE 7: Air-cored chokes have a narrow impedance peak, and are easily de-tuned by small amounts of stray capacitance.

DOUBLET – LADDERLINE – UNBALANCED TUNER. Next, let's look at the challenges of connecting an unbalanced tuner to a multiband doublet antenna fed with ladderline, for example a G5RV. Figure 6 represents the system as viewed at the shack end.

The components within the dotted rectangle are the equivalent circuit of the antenna system. If the antenna were in free space, the equivalent circuit would have been a simple two-terminal network representing only the two connections to the legs of the ladderline. However, no terrestrial antenna system can be completely independent of local ground, so a real-life representation requires the three-terminal network shown in Figure 6, the three terminals now being Leg a, Leg b and Ground.

If the doublet and ladderline were installed perfectly symmetrically with respect to ground, then of course we could expect that $Z_a = Z_b$. However, in real-world installations this is never the case. My own doublet/ladderline installation looks quite symmetrical but measurements showed that to be far from true. The measured values showed significant differences in the magnitudes of the three impedances Z_a , Z_b and Z_c , which were respectively 80Ω, 109Ω and 114Ω. But the differences in the vector ($R + jX$) impedance values were even more dramatic:

$$Z_a = (15.1 + j79)\Omega,$$

$$Z_b = (1.6 - j109)\Omega$$

$$\text{and } Z_c = (30.7 + j110)\Omega.$$

So as well as the doublet's feedpoint impedances Z_a and Z_b being highly asymmetrical, there was also quite a low common-mode impedance Z_c to ground. You probably shouldn't assume that your installation would be any better!

We have already seen that the ladderline currents I_a and I_b should be as well-balanced as possible, to reduce the common-mode component I_c that is mostly responsible for radiation from the ladderline and noise pick-up during reception. If we connect an unbalanced tuner directly to the ladderline as shown in Figure 6, this will destroy any system balance that might otherwise have existed because one of the two terminals (Leg b) is shorted directly to ground at the transmitter. It's clear that we must do something to prevent this.

As in the dipole/coax case, the common-mode current can be found flowing on the outside of the coax, so the best way to block it is by placing a CM choke at the transition between the coax and the ladderline. The high impedance of the CM choke then forces the external shield current towards zero, so that I_a and I_b are forced towards balance.

FOLLOW THE CURRENT! Notice that in both the dipole/coax and the doublet/ladderline examples we are concerned primarily with currents rather than voltages. In both cases it is unwanted common-mode current that we are aiming to suppress, and also differential-mode currents that we are aiming to establish on the ladderline. This need to focus on currents is true in nearly all antenna balun applications that we come across, indicating that a current balun or CM current choke is nearly always the appropriate choice.

A voltage balun is almost always the wrong choice because it fails to address the real problem. If I had used a voltage balun at terminals a and b of my own doublet/ladderline system, it would have forced equal voltages across the two unequal impedances Z_a and Z_b , and thereby guaranteed that the currents I_a and I_b would be unequal – so a voltage balun would have done exactly the opposite of what was needed! (Interestingly, this doesn't prevent many tuner manufacturers from incorporating a 4:1 voltage balun, which makes one wonder how well they are keeping up with the state of the art.)

WHAT MAKES A GOOD CM CHOKE?

A good CM choke will present a very high impedance to unwanted common-mode currents across its entire range of operating frequencies, without having any significant effect on its ability to deliver the desired differential-mode power to the antenna system. Which leads to the question: how high a common-mode (CM) impedance do we need? And the answer, of course, is: "It depends!"

In the dipole/coax case (assuming the feedpoint impedance is around 50Ω) we could easily achieve a reasonable current balance with a CM impedance as low as 250Ω. From that very limited point of view, 500Ω has often been presented as a value to aim for. But that isn't the whole story, for two important reasons.

First, in the multiband doublet/ladderline case the feedpoint impedance is likely to be very high on some bands. To achieve the same degree of balance as the easy 50Ω case, the CM impedance will need to be higher still – possibly several thousand ohms.

The second consideration is power dissipation in the choke. If the CM impedance is not sufficiently high, the very high voltages present on the ladderline on some bands can still manage to drive enough common-mode current through the choke that losses in any core material can lead to overheating. Assuming UK power limits, and a duty cycle of 50% or less, a CM impedance

of 4000Ω or more should ensure safe choke operation and is a realistic design aim.

There will certainly be some cases where a lower CM impedance will be good enough, so we don't have to be so particular about the type of CM choke. But usually we only discover that in hindsight. For reliable performance – to get it right first time – we still need to aim for a high CM impedance like 4000Ω.

AIR-CORED CHOKES – RARELY A GOOD CHOICE.

A commonly recommended method for making a CM choke is to take some extra length of coax and simply form it into a multi-turn coil. As I hope you'll understand, making the coax into a coil will obviously not affect differential-mode currents flowing inside the coax. In the outside world, making a coil will have some of the desired effect of impeding the common-mode current flowing along the outside surface of the shield. But the performance will rarely be as good as we might expect.

Figure 7 shows the CM impedance of a typical 'coiled-coax' choke comprising 5 turns of RG-58 on a 110mm (4.25-inch) diameter. Some very impressive CM impedances can be achieved in this way, as shown by the blue trace on the chart which reaches a peak value of over 25,000Ω. This high peak impedance is the result of the coil's inductance combining with the stray capacitance between the turns to form a high-Q parallel tuned circuit. Looking across a wider range of frequencies, at lower frequencies the choke behaves like an inductor, so increasing the frequency increases the impedance. Then as we come close to the self-resonant frequency (SRF) the impedance peaks sharply. As we go beyond the SRF the impedance falls sharply away again, and continues to fall because the choke is no longer behaving like an inductor; the dominant component is now the stray capacitance.

Impressive though the peak impedance is, this type of choke has several weaknesses. The high-Q nature of the choke means it only delivers high impedance over a relatively narrow range of frequencies; and secondly, any small change in stray capacitance




FIGURE 8: $\mu'S$ affects inductance while $\mu''S$ affects the loss resistance.

can shift the self-resonant frequency dramatically, as shown in the red trace in Figure 7 where a 2.2pF capacitor has been placed in parallel with the choke to represent a small change in the layout of the turns. Even such a small change in capacitance has detuned the SRF by 17%. Because the performance can be seriously affected by such small details, different constructors following the same design may well end up with significantly different results.

There is a further problem with air-cored chokes which is even more important but not often appreciated. Because of their low-loss characteristics, the common-mode impedance that these air-cored chokes can deliver is almost entirely reactive (inductive below the SRF, or capacitive above it)... and any reactance can be partly or even completely cancelled out by another reactance of opposite sign. So if a reactive choke is inserted into a common-mode path which, in the absence of the choke, is reactive and of opposite sign, it is perfectly possible for the choke to reduce the net impedance, allowing the CM current to become even worse!

By way of example, a computer model of a 20m dipole erected 30ft over average ground and fed with RG213 showed a common-mode path impedance back along the coax shield of (28-j200)Ω; the CM impedance is capacitive because the feedline is just under a half-wave long. If we now insert a choke at the feedpoint that has a reactance of +j200Ω and very little resistance, the reactances will cancel and will reduce the net CM path impedance to 28Ω!

That said, air-cored CM chokes can sometimes be useful in non-demanding cases, especially if you restrict them to monoband applications and try to make sure that their self-resonant frequency is close to the operating frequency. (But then, please don't spoil it all by taping the coil of coax directly onto a metal boom!)

FERRITE CHOKES FOR BROADBAND PERFORMANCE. Clearly, for multiband applications it would be preferable if the choke maintained a high CM impedance across a much wider bandwidth, and also if its impedance was predominantly resistive because, unlike reactance, resistance can't be cancelled out. Fortunately, the properties of some ferrite materials allow us to approach this ideal.

If we wind a CM choke using a high permeability core material such as ferrite, we can achieve a high CM impedance with relatively few turns, which makes the choke easier to wind and also keeps the inter-winding stray capacitance to acceptable levels. However, it's the complex nature of a ferrite's permeability that proves to be its most useful attribute.

Figure 8 shows a plot of permeability against frequency published by manufacturer Fair-Rite for a core made from their Type 31 ferrite material. The red line represents permeability as most folk envisage it – the magnetic property of the ferrite that magnifies the inductance of a coil wound around it. This aspect of the permeability is given the symbol $\mu'S$ (pronounced "mu-dash-S"). Notice that $\mu'S$ is frequency dependent; its value is constant at low frequencies but for this particular material the value falls rapidly above 1.3MHz.

But because this ferrite material is 'lossy' it also generates an equivalent resistance in any coil wound round it, and this property is reflected in the green line. Fair-Rite give this resistive property the symbol $\mu''S$ (mu-double-dash-S). Notice that $\mu''S$ rises with frequency before falling again, but that $\mu''S$ is always more important than $\mu'S$ at frequencies above 3.5MHz. We might therefore expect that a choke wound around this Type 31 material will be more resistive than reactive over a wide frequency range from 3.5MHz upwards, and we'd be right!

Note that the use of a lossy ferrite material is NOT a bad thing. The ferrite does not introduce any loss into the differential-mode signal path. It only affects the common-mode path, where we can put that loss to very good use! The resistive loss in the core (due to $\mu''S$) is what we have to thank for delivering a dependably high impedance across a broad bandwidth.

But now another warning: DO NOT use iron dust toroids for common-mode chokes! Their low permeability produces very poor chokes that have very low impedance in the HF range. (See the chart at the end of this article... and here again, we have to wonder about the level of knowledge among some suppliers to the amateur market.)

HOW DO WE BUILD A GOOD CM CHOKER? So at last we come to the practical details – but I make no apology for keeping you waiting. Unless you understand the technical background of what a CM choke is for, and what qualities make a good one, you are much less likely to get good results.

The key point is that the highest performing CM chokes use ferrite cores. One simple construction technique is to slip ferrite beads over a transmission line. That works, and the choke's impedance

is pretty much proportional to the number of beads used; but if instead we wind the transmission line around a ferrite toroid we can benefit from the so-called 'N-squared' effect whereby the impedance increases as the square of the number of turns, reducing the number of turns needed for a particular level of impedance. A multi-turn choke is also cheaper, lighter and smaller than a long string of beads; it also allows more design freedom to control the self resonant frequency.

Commonly-available ferrite toroids come in diameters up to 61mm (2.4 inches) with an inside diameter of 25mm (1.0 inches) so practical considerations mean the transmission line used for the winding is often small diameter coax (RG-58 size or smaller). As an alternative to coax, you can sometimes use a bifilar pair of parallel wires; the choice between the two is determined by the application. For example, if the choke is to be used at the feedpoint of a coax-fed, nominally 50Ω antenna, we would prefer to wind the choke with 50Ω coax so that it doesn't introduce an unwanted impedance transformation. On the other hand, a choke used at the output of a tuner has no such requirement because any unwanted impedance transformation can be tuned out, so we can wind it with a transmission line of any characteristic impedance




FIGURE 9: Optimised CM choke for 6-30MHz.

within reason. We would likely choose a bifilar pair of well insulated wires for this application because they can withstand higher differential-mode voltages than small diameter coax.

Photo 1 shows how a typical toroidal choke is wound. This one comprises 7 turns of coax on two stacked 61mm diameter toroidal cores. Plastic tie-wraps have been used to secure the ends of the coax. The crossover at the 4th turn is not essential and it has no major impact on the choke's performance; it's mostly a way of making the two ends of the coax emerge on opposite sides of the toroid, which can be handy in some installations. Every pass of the coax through the centre of the toroid – including the crossover – counts as one turn.

How the coax ends are terminated will depend on the application. The choke is often placed in a box with a coax connector at one end and a terminal pair at the other. It's important in this case that the box is plastic. Using a metal box can significantly alter the stray capacitance and change the choke's performance, and if you ground both ends of the coax to the metal box, the choke will stop working altogether! Plastic boxes are cheaper and better for this application, so why bother with metal?

Figure 9 shows the sort of CM choking impedance that can be achieved using these construction techniques. This choke is an optimized design using 9 turns of coax on three stacked FT240-52 toroids. It achieves the design goal of at least 4,000Ω impedance from 6MHz to 30MHz, and the impedance is above 10,000Ω from 13MHz to 27MHz.

The green trace indicates that the choke impedance is predominantly resistive over a large part of this bandwidth, and from 10MHz to 28MHz it achieves 4000Ω on resistance alone.

By careful choice of the ferrite mix, number of turns, and the number of toroids in the core stack, it is possible to produce a CM choke for most HF applications. **Figure 10** illustrates the performance of a number of useful designs, and the chart format allows easy comparison of the various options.

The CM impedance in **Figure 10** is represented by the coloured bars. Dark green is the highest impedance (>8kΩ) and red the lowest (>500Ω). Within these multi-coloured bars, the black line shows the frequency range over which the choke is predominantly resistive.

By way of comparison, a ferrite bead choke, an iron dust choke and an air-cored coiled-coax choke have been

included at the bottom of the chart.

Notice the very modest performance of the bead choke, the abysmal performance of the iron dust choke, and the very narrowband performance of the coiled-coax choke!

All the results in this chart were obtained using RG-58 for the windings, and using the cross-over winding arrangement shown in the lead photograph. Using a smaller diameter coax, a bifilar winding, or not using the cross-over will change the results somewhat, but with these fairly wideband chokes the effects are not major.

TYPES AND SOURCES OF FERRITE PARTS.

Throughout this article I have used the popular amateur designations for the ferrite parts; for example FT240-43 indicates a Ferrite Toroid with an outside diameter of 2.40 inches, made from the Fair-Rite type 43 grade of ferrite. Two points to note here:

- Those particular grades of ferrite are manufactured by only one company, Fair-Rite in the USA. The grades or 'mixes' that are most useful are 31, 43, 52 and 61. Other manufacturers may produce similar materials but their properties will be different (meaning: you're on your own!).




Figure 10: G3TXQ's recommendations for good ferrite chokes

- Fair-Rite do not use the 'FT240-43' style of part numbering. These numbers are added by resellers, primarily to the amateur market. Fair-Rite and their main dealers use the original part numbers, which are as follows:
 - 'FT240-31' (Fair-Rite part number 2631803802)
 - 'FT240-43' (Fair-Rite part number 2643803802)
 - 'FT240-52' (Fair-Rite part number 5952003801).

Note that the 3rd and 4th digits are the grade of material – 31, 43 or 52.

You can order Fair-Rite toroids from a number of suppliers serving the UK. Recommended suppliers include Farnell (www.farnell.com), Mouser Electronics (based in the USA but with a dedicated UK website <http://uk.mouser.com> and excellent shipping arrangements) and a number of amateur dealers (search on eBay for the FT- number).

By all means shop around for the best prices, but always check the total cost including shipping. Ferrites are heavy and fragile, which makes them genuinely expensive to ship, so take advantage of 'free shipping' offers wherever you can.

G3TXQ CHOKE CHART. The chart in **Figure 10** shows my recommendations for a range of high-performance ferrite chokes, between them covering all of the HF amateur bands. Be warned that the last three are examples that you should not follow!

Finally, my website includes a performance chart similar to Figure 10, but which includes a different range of choke designs (www.karinya.net/g3txq/chokes). It's worth taking a look from time to time because I add new designs as I make further measurements. The website also includes more detail about why reactive chokes can be undesirable.

Appendix: How to measure CM choke impedance

My early attempts at choke impedance measurement used a single-port vector impedance analyser (AIM4170) with the choke directly connected from the measurement port to ground. However, it's not ideal: very high choke impedances are outside the range where the analyser can be expected to be accurate. Despite careful calibration to the measurement plane, the analyser always added the equivalent of a few pF of parallel capacitance; this significantly shifts the self-resonant-frequency of higher-Q chokes.

More accurate results can be obtained by measuring the attenuation that the choke introduces when placed between a signal source and a load. So, for instance we could place the choke in series between a signal generator and an RF voltmeter; then, knowing the generator output and the RF voltmeter reading, we could deduce something about the choke's impedance. However this simple scalar measurement will not tell us anything about the choke's complex impedance ($R \pm jX$). As we have seen throughout this article, it is vital to know both the resistance and the reactance for a complete understanding of how well it will perform.

Fortunately, a 2-port vector network analyser (VNA) can measure both the magnitude and the phase of the attenuation introduced by the choke, and that allows us to fully determine the choke's complex impedance. I use an Array Solutions VNA2180 in the arrangement shown in **Figure 11**. This jig typically adds the equivalent of 0.2pF or less of parallel capacitance.

The ports of the VNA are connected to the test jig via two short lengths of coax – long enough that the choke is physically well-isolated from anything that could cause stray coupling and affect the measurements. The test jig comprises two BNC sockets mounted on a small piece of PCB material; crocodile clips soldered to the BNC centre pins allow the choke to be connected in line. A 50Ω resistor on the jig ensures that the impedance 'seen' by Port A of the VNA never becomes extreme.

The VNA is first calibrated with the two coax leads removed from the jig and connected back to back through a BNC double-female adapter. Then they are re-connected to the jig, the choke clipped between terminals A and B, and a VNA measurement scan made between the required frequencies. The resulting S21 amplitude and phase data is then transferred to a spreadsheet to calculate the choke's complex impedance, and produce the plots shown throughout this article.

Figure 12 shows the significantly different results that are obtained on a high-Q choke when using the one-port reflective measurement of an analyser like the AIM4170 compared to the two-port transmissive measurement of the VNA2180. Incidentally, this choke design is taken from the *ARRL Antenna Book* where the self-resonant frequency is stated as 14.3MHz, indicating an even greater error caused by uncorrected stray capacitance.




FIGURE 11: How to measure complex impedance using a two-port VNA.




FIGURE 12: Effect of measurement method on the apparent resonant frequency. The higher measurement, made as described here, is the more correct value.