In Practice

Cost-effective ferrite chokes and baluns

PHOTO 1: Clockwise from left: low-bands ferrite choke, mid-bands choke, high-bands choke, the ferrite cores.

NEW DESIGNS. This month I will be describing some new designs of ferrite loaded chokes for suppressing unwanted common-mode currents at HF. The same basic designs can be used in several different ways:

- As choke baluns ('current baluns') for coaxial feedlines
- In the shack, applied to various coaxial, mains and data cables
- Applied to consumer electronics to suppress interference on antenna feedlines, audio/video and mains cables.

These are all different kinds of EMC problems, and the same designs of RF chokes can often be used in a number of situations by simply changing the type of cable involved. An expanded version of this article on the RSGB Members Only website explains much more of the technical background about common mode currents **[1]**. This shorter version tells the other part of the story, the search for RF chokes that have high performance but don't use large and expensive ferrite cores.

How good does an RF choke need to be? To make a good transmitting balun? To suppress RF in the shack? To suppress RF interference to (or from) consumer electronics? The problems of EMC engineering are that every situation is different and there is never enough technical information to be sure of success. We can meet the first challenge by always using high-performance chokes and filters that are capable of handling *almost all* EMC problems. We meet the second challenge by making maximum use of whatever facts we do know. In EMC, and in much else besides, our best and most reliable friend is Ohm's Law.

When we use an inline RF choke to suppress unwanted RF current, we are inserting some additional impedance between two impedances Z1 and Z2 that are already present in the system. **Figure 1a** is highly simplified but it captures the essential features of almost every EMC situation. Looking upstream of where you're going to insert the choke, the unwanted common mode current has some kind of source which we can represent as V1 with an impedance of Z1. Looking downstream, that current is almost certainly 'trying to find earth', along a pathway that has a series impedance Z2. The only things that change between one case and another are the values of V1, Z1, Z2 and of course the unwanted common mode current itself, I_{CM}.

The aim of the RF choke is to reduce I_{CM} to some much lower level that the affected equipment can tolerate. To achieve this (**Figure 1b**) the impedance of the RF choke will obviously need to be much higher than Z1 and Z2 combined – but how much higher does Z_{CHOKE} need to be, to be certain that it will dominate the situation? As I said above, EMC has no universal answers so we have to apply a combination of engineering and experience.

MUST DO BETTER. To provide dependable solutions for a much wider range of practical EMC problems, experience shows that RF chokes need to have an impedance of at least a few *thousand* ohms, maintained across a wide bandwidth [1]. Many existing types of cable chokes fail to meet these criteria, so there are some EMC problems where they fail to work. Air-wound chokes and ferrite loaded chokes have different weaknesses, so I will discuss each kind in turn.

AIR-WOUND CHOKES. These are the simple coils of cable that are often suggested as choke baluns. We tend to think of these coils as inductors, but their high-frequency performance is actually dominated by the distributed capacitance between the turns. For example, take about 2.2m of thin coax like RG8X or RG58 and wind it into a five-turn bundle of about 125mm average diameter (**Figure 2**, inset). This has an inductance of about 6μH, but the capacitance between the turns is equivalent to about 9pF in parallel with the 6μH. So instead of an inductor, what we actually have is a high-Q parallel resonant circuit with the measured impedance characteristics of Figure 2.

This parallel resonant circuit does not make a dependable RF choke. The impedance is only high around the resonant frequency, and much lower elsewhere. The resonant frequency is also quite sensitive to small changes affecting the capacitance between the turns, even how tightly the turns are taped together. But the fatal flaw of these chokes is that their performance is very dependent on the situation in which they're being used. This is because the impedance of the choke consists almost entirely of either inductive or capacitive reactance, at all frequencies except the very narrow region close to resonance as shown in Figure 2. Going back now to Figure 1, the reactive impedance of the choke is in series with the upstream and downstream impedances, Z1 and Z2… which also have inductive or capacitive reactances of their own. You never know from one situation to the next whether Z1 and Z2 are going to reinforce the impedance of the choke or cancel it. "Are you feeling lucky today?" is not my idea of good RF engineering!

In practice, reactive (air-wound) chokes often provide enough impedance to handle 'soft' EMC problems; but they don't have the broad bandwidth that is often claimed, and they can sometimes let you down badly. In a word, reactive chokes are not *dependable* EMC solutions.

FERRITE LOADED CHOKES. To overcome the problem that reactive impedance can sometimes shift or disappear, the impedance of a dependable RF choke needs to be both large and predominantly resistive. The advantages of resistive impedance are that it cannot be cancelled out and it also tends to broaden the useful bandwidth of the choke. Any practical choke will also have some reactance, which is nice in situations where it works for you, but resistive impedance is the only solid foundation for dependable performance.

The only way to create a high resistive impedance is to carefully engineer a certain amount of loss into the choke… and that is why we need the ferrite. Don't panic about 'loss': unlike many other situations, resistive loss in an RF choke *is a very good thing*. We just need to make sure that it appears as a very high value of R in the series impedance, $Z_{CHOKE} = (R \pm jX)$. The resistive (heat) loss in the choke equals I_{CM}^2R, where I_{CM} is the residual level of common-mode current that remains *after* the choke has been inserted. If the choke has successfully suppressed the common mode current (and thus solved the EMC problem) then the residual value of I_{CM} will be very low and you'll be unlikely to notice significant heating in the ferrite. This is why we're aiming for an R value of several thousand ohms, rather than a lower value like 500Ω which experience has proved to be inadequate (see the expanded article).

Ferrite chokes with a resistive impedance less than 1000Ω are at much greater risk of underperforming and overheating. Many of these chokes were designed to meet that inadequate target of 500Ω, and some commercial examples have also suffered further cost-cutting, eg by using smaller quantities of ferrite and failing to use the correct materials. If a ferrite loaded choke begins to overheat, the ferrite may reach the Curie temperature at which its magnetic permeability collapses, allowing I_{CM} to increase and causing further overheating – the choke will almost literally 'crash and burn'. As I said above, these poorly performing chokes may work for 'soft' EMC problems but they don't have enough impedance to handle anything challenging.

CAN DO BETTER. To make a really good ferrite choke, you need to do two things:

1. Choose the right grade of ferrite, one that actually has some loss at the operating frequency.
2. Construct your choke to create just the right amount of coupling between the ferrite material and the magnetic field around the cable.

Neither of those things will happen by blind luck. There are hundreds of different grades of ferrite with widely differing magnetic properties. They all look the same so you have to know exactly what grade you're using. That will mean buying 'named ferrite' from a reliable source. With the help of a Vector Network Analyser, it then becomes quite easy to develop some effective ferrite-loaded chokes **[2]**. But if you don't have access to that level of test equipment, the only route to dependable performance is to copy someone else's designs.

Most of the published designs [1] originate in the USA and use ferrite cores manufactured by the Fair-Rite Corporation. None of these cores are cheap, and here in Europe they will cost about twice as much due to shipping, VAT and all the other markups and 'handling charges'. We can reduce costs a little by shopping carefully, combining orders with other amateurs and taking advantage of special offers; but the cost of ferrite is always going to be a much bigger consideration on this side of the Atlantic.

Jim Brown, K9YC has been particularly active in developing designs for high performance ferrite loaded chokes, and his PDF papers and PowerPoint presentations are essential reading [1]. To make sure that his chokes can handle even the most stressful applications at power levels up to 1500W, K9YC aims for very high values of resistive impedance (preferably 5000Ω or even more). However, that superb performance is achieved by using large ferrite cores, sometimes four or five at a time, which are not affordable at European prices. Thus we are forced to look for alternative designs that cost a lot less but can still handle the large majority of balun and EMC problems. In other words, we're looking for cost-effectiveness.

Strings of ferrite beads are definitely not cost-effective. Ferrite beads can usually take only one 'turn' of cable (one pass through the centre hole = 1 turn) and each individual bead generates quite a low impedance, so a high impedance will need an awful lot of beads in series. Ten or 20 beads will only give enough impedance to handle the easy, soft problems; for dependable performance, think 40 or 50 large beads and then work out the cost! As K9YC and many others have pointed out, the cost-effective way to achieve a high impedance is to use multiple turns through the same core, because the impedance will then increase with the number of turns *squared*. But multiple turns of thick coax, mains or rotator cable will require a large core… and we're straight back into the problem of expensive ferrite.

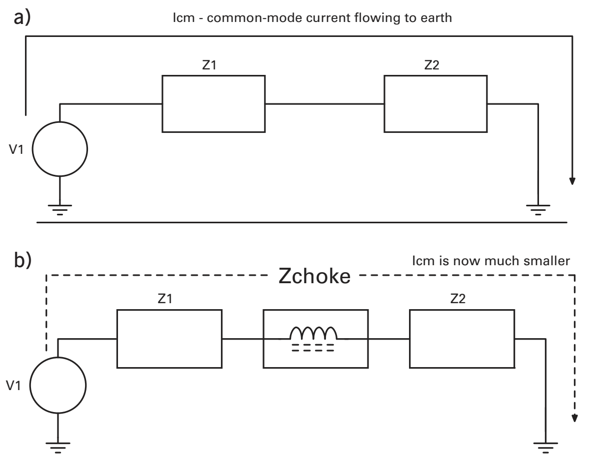

FIGURE 1: An effective common mode choke must dominate the upstream and downstream impedances, Z1 and Z2.

TABLE 1: Dimensions of the three HF ferrite chokes

	Turns	Mean diameter	Cores
Low bands	5	125mm	3
Mid bands	4	85mm	3
High bands	3	Close wound	2, glued side-by-side

All ferrite cores are Fair-Rite 2643167851 = Farnell 1463420. No substitutes allowed!

For some years I have been looking for a way out of this, and inspiration came with the new 2010 *ARRL Handbook*. The Transmission Lines chapter features some new choke designs that use a small number of relatively low-cost ferrite cores threaded onto a coil of cable (**Photo 1**). These cores have an oval central hole, 26 x 13mm, which will take several turns of thin transmitting coax like RG8X, or similar-sized cable of any other type. And although they are made by Fair-Rite in the USA, these particular cores don't have to be specially imported; they are readily available as stock items from Farnell UK at about £2.70 each [1].

That ARRL design concept has opened the way to a range of cost-effective ferrite chokes that can tackle the large majority of balun and other EMC problems across the HF spectrum. The three chokes in Photo 1 are only examples of what can be done; each choke delivers a high resistive impedance over at least a 2:1 frequency range using only two or three of the oval Fair-Rite cores. The performance isn't as good as Jim Brown's biggest and best, but they are a major advance over most of the balun and EMC chokes that we're using at present. Because all these chokes use the same ferrite cores, the most cost-effective strategy is to keep a stock of the bare cores, and then quickly run up a suitable choke for any cable that needs one.

The key dimensions for the three HF-band chokes are given in **Table 1**, and further construction details are on the 'In Practice' website [1].

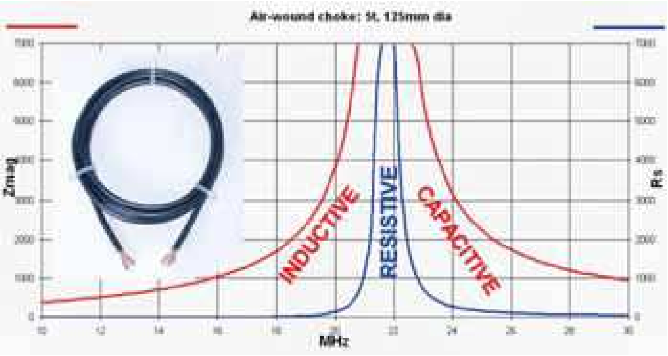

FIGURE 2: Performance of an air-wound choke: notice the very sharp resonance at 21MHz.

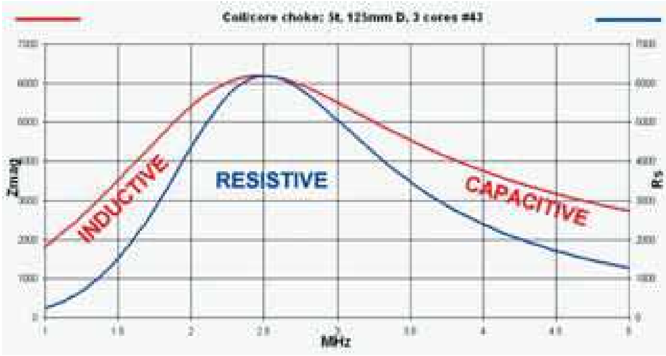

FIGURE 3: Performance of the low-bands choke with three ferrite cores.

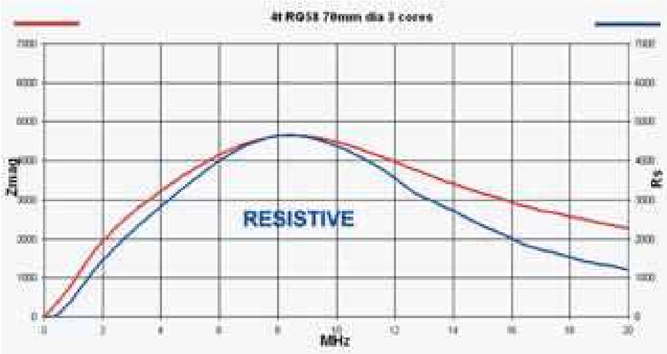

FIGURE 4: Performance of the smaller mid-bands choke with three ferrite cores.

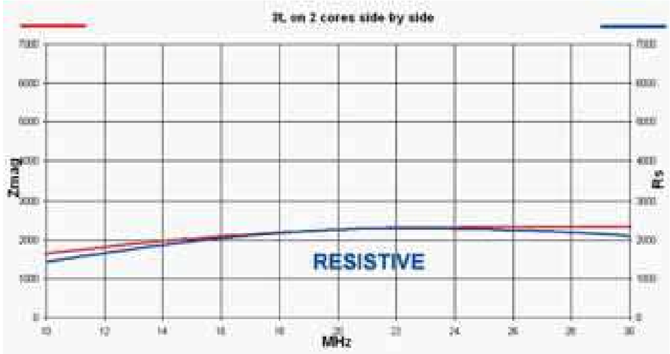

FIGURE 5: Performance of the high-bands choke, wound on two ferrite cores side by side.

LOW BANDS. When two or three of our ferrite cores are threaded onto the flat 5-turn coil that was described earlier (**Photo 1, left**) the narrowband 21MHz choke from Figure 2 is transformed into a broadband choke covering 1.8–3.8MHz. **Figure 3** shows the measured performance. The blue trace is the resistive part of the impedance, which is about 4000Ω on Top Band and 3000Ω on 80m. The total impedance (red trace) includes some additional inductive reactance at lower frequencies and capacitive reactance at higher frequencies, but like Jim Brown I only regard this a bonus – nice to have, but we're not actually depending on it for good performance. Despite the drive to reduce ferrite costs, I found that three cores gave a worthwhile increase in the resistive part of the impedance, compared with the two cores used in the ARRL design.

As you see from Figure 3, the two amateur bands are actually on the skirts of the resonance peak, so that peak needs to be positioned fairly accurately to produce similar performance on both bands. To obtain the correct amount of distributed capacitance between turns of the coil, you'll need to follow the detailed assembly instructions on the website with care.

MID BANDS. To cover 5, 7 and 10MHz, reduce the coil diameter and the number of turns but still use three cores (**Photo 1, top right**). **Figure 4** shows excellent performance across all three bands, and this same choke may also be usable for easier EMC problems down to 3.5MHz and up to 14MHz. For optimum wideband coverage it is essential that the turns of cable are stacked vertically inside the cores with no crossovers, exactly as shown in Photo 1.

HIGH BANDS. For 14 – 30MHz coverage, this design concept is somewhat running out of steam but we aren't beaten yet. If two of the same cores are superglued together side-by-side as shown in **Photo 1, lower right**, three turns will make quite a respectable choke for a 20 – 10m beam. The impedance (**Figure 5**) isn't quite as high as the lower-frequency chokes at their very best, but it is substantially resistive across the whole 14–30MHz range. In terms of 'value for ferrite' this two-core choke will at least equal a straight string of *40 to 50* ferrite beads!

By the way, if you want more impedance or a wider bandwidth, you can cascade any of these chokes in series along the cable. The interactions are quite mild and the impedances always seem to reinforce each other (rather than destroying each other, as always happens with reactive air-wound chokes). Measurement results are available in the expanded article [1].

FURTHER DEVELOPMENTS. We are definitely onto something: this design concept is already delivering high performance at an affordable cost, and more development work will surely drive both factors closer to the optimum.

A number of experimenters are already finding other ways to make feedline chokes using these same cores **[3]** and they also show great promise for other applications such as filtering the mains supply to the shack (see the May 2009 column).

The expanded version of this article contains some preliminary conclusions about the relationships between coil dimensions, the numbers of cores and the resulting impedance and bandwidth. This is all wide open for experimentation, using these and possibly some other types of ferrite cores – so if you have the necessary test gear [2], go to it.

On the other hand, if you need some high-performance RF chokes right now, go directly to Table 1. Please check the expanded article and the 'In Practice' website for more details about construction [1], and let me tell you one last time: you must use *only* the specified ferrite cores!

NOTES AND REFERENCES

[1] An expanded version of this article is available on the RSGB Members Only website: www.rsgb.org/membersonly/publications/radcomplus/index.php
There are further notes and web links on the 'In Practice' website: http://tinyurl.com/inpractice
[2] This kind of development work requires a Vector Network Analyser that can compensate for the stray inductance and capacitance of the test jig, and can measure high impedances with verifiable accuracy. Unfortunately this is beyond the capabilities of R-X antenna analysers like the MFJ-259B.
[3] Special thanks to G3TXQ, M0JEK, VK1OD, VK4OQ and K6MHE.