BUILDING A SYNTHESISER

Nigel Gresley breaks out the soldering iron and other bits and pieces in order to put together a synthesiser, bought from Wood and Douglas in kit form. Was it as easy as he first thought? Apart from a minor problem or two, everything went as planned. Here, our tame engineer talks us through the build from start to finish.

Amateur RADIO HOME PROJECT

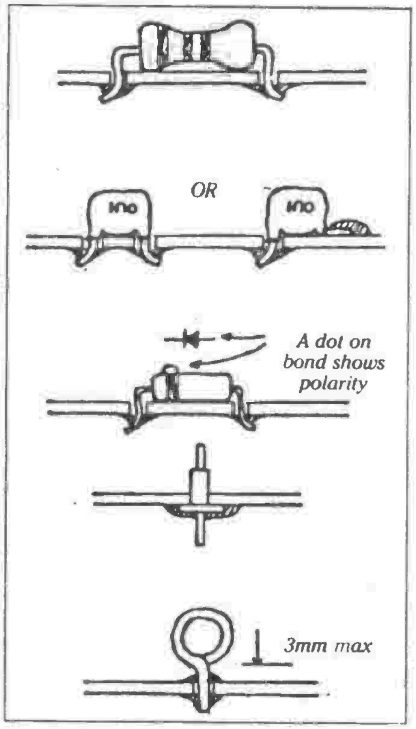

Above: Methods of attaching capacitors to the pcb. Miniature coax cable should be used for all RF signal connections. Below: 70FM05R series – conversion to external oscillator drive.

Everyone's heard of synthesisers. They feature in practically every rig from the Land of the Rising Sun these days. But almost no-one understands them. Learned journals tend to talk glibly about things like jitter, phase noise and sidebands as though they were as familiar as morning coffee, but most people don't really have a clue as to what it's all about. Well, what we're going to do is to describe in simple terms how the beast is supposed to work and then we'll show you how to build one from a kit – in this case it's from Wood and Douglas, whose UHF FM transmitter and receiver we had a look at a while back.

So, what's a synthesiser anyway? It's basically a way of generating frequencies for amplification and modulation and what-have-you later on in the wireless, and it's a sort of 'different' VFO. A frequency synthesiser is designed around the principle of something called a phase-locked loop. Now don't panic, chaps, we'll get to that in a bit – and the rough idea is to lock a free-running oscillator to one which is crystal-controlled and therefore a lot more stable.

When a phase-locked loop is "in lock", the two oscillators – the "reference" or crystal-controlled one and the "variable" – are running at the same frequency. If you try to alter the frequency, say by waving your mit in the general direction of the coil of the variable oscillator, you'll find that the loop will compensate for this and, if you like, "correct" the variable oscillator to keep it on the right frequency.

What happens is that the reference frequency is compared with the variable frequency in something called a "phase comparator" circuit. This consists of some kind of mixer which uses transistors or diodes as the non-linear element which every mixer has to have otherwise it won't mix. Now the output of the phase comparator is a voltage – don't worry about the finer detail for now, just take our word for it otherwise we'll be here all day and you'll start nodding off – and this

Metal mounting screws

voltage varies with the difference between the reference and the variable oscillators. This is very handy because if you amplify it, it can provide an error signal which can be used to control the frequency of the variable oscillator and tell it to get back on frequency pretty smartish or it'll smash the variable oscillator's transistors over the head. So one thing we need to do is to make a variable oscillator whose frequency is controllable by a voltage – this is dead easy and indeed from here on in we'll refer to the variable oscillator as a **voltage controlled oscillator,** or VCO as it's usually known.

This is an important thing to twig because all the clever magazines airily refer to VCOs as though they were given away free with the morning paper and lots of people(we hear on the wireless)

70FM05R Series – Conversion to External Oscillator Drive

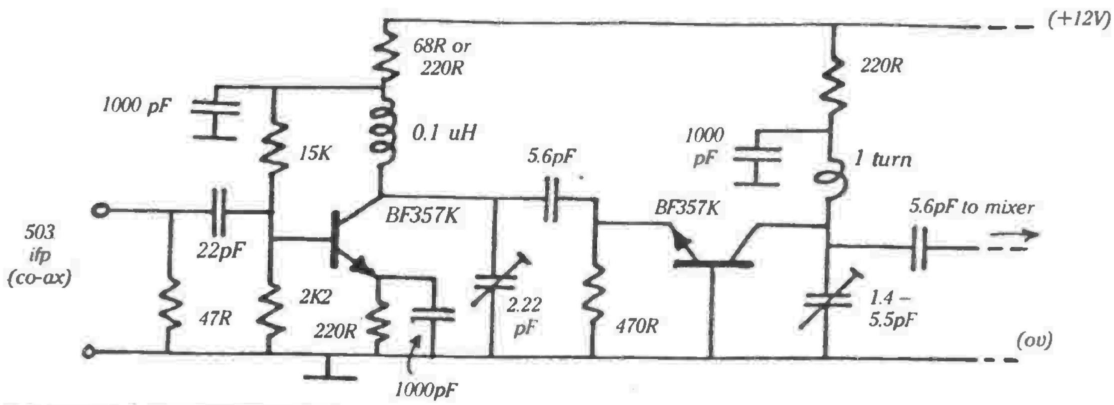

don't know what one is and are too shy to admit it. So don't forget – VCO stands for voltage controlled oscillator, and it just means an oscillator whose frequency can be controlled by sticking some volts in the right place.

The usual way to do this, by the way, is with varicap diodes, which have the rather handy property of changing their capacitance according to the voltage across them. Gresley's Fourteenth Law says "the more the volts the less the puffs". Actually – please don't take that the wrong way . . .

Crystals are costly

Anyhow, the story so far is that we've got a VCO whose output is mixed with the output of a reference oscillator so as to provide a DC error voltage which is fed back to control the frequency of the VCO. Pause for the roll of drums and hey presto, one phase-locked loop, sir; would you like chips or saute with it? We have to keep it simple at Bicester otherwise we get all messed up and have heavy interviews with the MD . . .

In fact, of course, there's a lot more to phase-locked loops than that, as you'll find out later on in your career, but don't panic – it'll all fall into place by and by.

Well, this is all very well, we hear you say, but why not just use a crystal for each frequency you want and get rid of all this VCO and phase comparator stuff? If you did you'd need one crystal for every output frequency you wanted in your wireless and this can get more than somewhat costly – as many of us found out in the early days! The next really clever bit is to introduce an offset into our phase-locked loop (which we'll refer to as a PLL from now on to save on ink – we have to make a profit, you know) so that the thing can produce outputs on different frequencies.

This is kind of tricky because it'd need to be extremely accurate and stable and also variable; hmm, quite a tall order. However, if you could do such a thing you'd end up being able to make the VCO lock up to provide lots of different output frequencies, which, if you think about it, would be very handy. Men, it can be done – but this is where we leave off the fancy RF stuff such as oscillators and suchlike and start coming over all digital. Don't panic – it'll all come . . .

What you do first of all is to take the output of the VCO and smash it into some suitable circuitry so that you get more or less a square wave out of the end. You then stick this into a piece of logic which is called a **divider,** basically because it divides, Brian. Now, this has to be very stable and divide by some nice whole number like 100 or 101 or whatever – it mustn't ever, ever, divide by anything in between even if you want it to. The idea is that if you poke 100 pulses into it, or 101, or however many it happens to be, you get just one out – this stability, in fact, is an inherent property of a logic-type divider and it's just what the doctor ordered for synthesisers as we'll see.

Now then. The next sneaky thing is that it's possible to get dividers where the division ratio is **programmable,** as the Clever Dicks say, or adjustable, as we say in Bicester. By division ratio, we simply mean, for instance, 100:1 or 101:1 or whatever, as in the example. Note that the division is always in whole numbers and never in nasty fractions – nice easy stuff. Bearing in mind the principle of the PLL which we discussed a while ago, the VCO is forced to lock up to the reference oscillator, because the error voltage tells it to, only now we've stuck this divider in the way.

Gresley's 14th Law says: the more the volts the less the puffs . . .

What happens then? Let's take an example and look at it – let's imagine that the reference frequency is set by a crystal to 25kHz and that the VCO can operate anywhere between 2 and 3MHz: it'll go wherever the error voltage tells it to, remember. If we then arrange for our divider to divide by 100, which it will do exactly if we so desire, then the frequency which emerges from it after we've stuffed the VCO output into it will be somewhere between 20 and 30kHz depending on wherever the VCO happens to be frequency-wise at the time. Now the phase comparator will detect the error between the 25kHz reference and the divider output and it will come over all bossy and change the VCO frequency in order to minimise the error. The loop will lock up when the divider output is 25kHz, or in other words when it's the same as the reference frequency, and so the actual frequency of the VCO wil be 25kHz times 100 which is 2.5MHz. Amazingly smart, isn't it?

Technical bod examines the Wood and Douglas diagram, then smiles sweetly (?) when he finally understands what's in store!

The state of play is that we have a VCO which is locked up at quite a high frequency, which is possibly more useful than that of the basic PLL. But wait! Things get even cleverer by and by; if we then tell the divider to change to a different division ratio such as 101, there will then be a new error signal present at the output of the phase comparator. The VCO gets told to do something about it, because of the voltage changing its frequency, and the system will lock up on to a new frequency which is 25kHz times 101 ie 2.525MHz. This, of course, is 25kHz higher than the old frequency, and it just so happens that the reference oscillator's frequency is 25kHz . . . can this be leading somewhere? Yes, folks, it is. If you try any division ratio which the divider can handle without bursting into flames or blowing all the house fuses, you'll find that the difference between consecutive divisions will always be 25kHz – which is, of course, the reference frequency.

The Clever Dicks will have sussed already that if you were to mix this VCO output up to VHF you'd end up with a very useful transmitter with 25kHz channel spacing; in practice, you'd choose

BUILDING A SYNTHESISER

very much higher VCO frequencies so that getting up to the final output frequency was a bit easier and to keep any spurious outputs down to a lower level. All synthesisers, by the way, have some of those – it's the usual RF bugs and you can't get completely shot of them, although some synthesisers do much better than others. You can't go too high, though, because programmable dividers only go up to a few tens of MHz – well, most of them anyway. If you want to go any higher you have to use something called a prescaler, which is usually a fixed divider dividing by 10 or 100, and the output of this then drives your fancy lower-frequency programmable divider.

Another method which you'll come across is to mix the frequency of the VCO down to a more reasonable amount, and this is the method which Wood and Douglas use in their 70cm machine. Here, the VCO does its stuff at about 140MHz and it gets mixed with the output of a crystal oscillator at 69MHz which is doubled up to 138MHz. 140 minus 138 is 2MHz, which is easy for a low-cost programmable divider to handle (why fork out lots of cash on expensive ICs when a much cheaper and just as effective approach is available, even though you need to be a bit more careful on the RF side?). In fact, there are advantages and disadvantages in all these techniques although we'll leave them for now so that we've got some brain cells left for later on!

Just to finish this bit of the article, can any Smart Alec see the easy way to modulate the VCO with FM? I thought as much. If you make the response time of the PLL fairly slow, or in other words make it take a while to lock up again when you change the division ratio, and then stick some audio on to the error signal, the VCO will follow this and hence it'll be frequency modulated. Mind you, if the response of the PLL is fast it'll treat the audio as an error and correct it – hey presto, no FM!

Another interesting point is that if you have a PLL running at the IF of your receiver – say 10.7MHz, for instance – and you stick that IF of the receiver into the phase comparator where the reference usually goes, you'll find that if you're receiving an FM transmission and the loop is fast enough, you'll find perfectly demodulated audio sitting on top of the error signal! What happens is that the VCO is made to follow the incoming signal by the fast loop, and since the reference itself is wobbling about (because it's an FM carrier) the VCO wobbles about in sympathy – as we know, in order to make the VCO wobble the error signal itself must likewise wobble and if you take a high impedance amplifier and apply the error signal to it, the wobbling gets retrieved as audio. By jingo, Carruthers, clever stuff . . .

What you get for your money

Enough of this wobbling about; what about the Wood and Douglas machine? It's their Model 70SY 25B, and along with it comes the 70VC05 which is the transmit and receive oscillator board, ie the bit that the synthesiser controls. The first impression of the kit is that that they've sent you two of them by mistake – there are a hell of a lot of bits and pieces. Wood and Douglas say that this is a very advanced project and in a way they're not kidding; reading through the info provided suggests fairly straightforward constructional techniques but it's when you get to the setting up instructions that mild hysteria sets in! Muttering "it's bound to work first time" under my breath I decided to wade in and see what happened.

Two boards are provided; the synth itself measures about 4in by 3in and the VCO 3½in by 2in. The boards themselves are fairly well made and had been roller-tinned, which was a Good Thing. Judging by the PCB tracks it looked as though the VCO board was going to be the more demanding; the component density

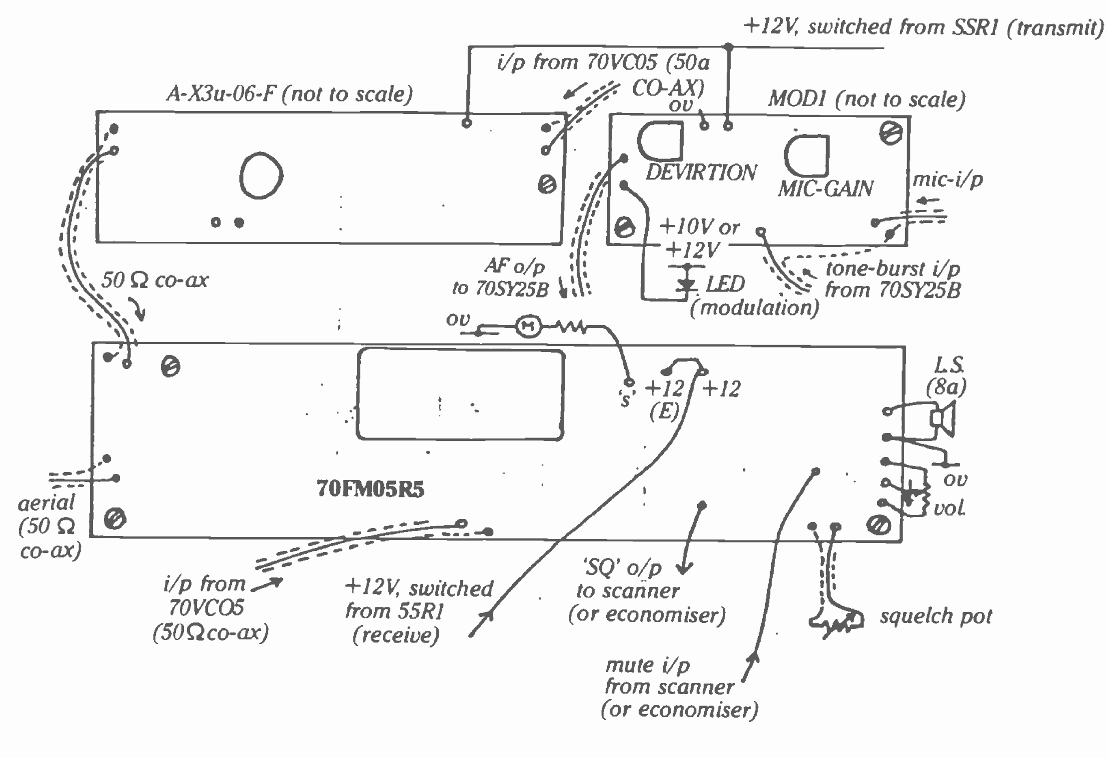

Interconnection diagram for receiver, transmitter and modulator in 70cm synthesised transceiver.

looked very high and indeed most of the resistors on the layout plan were standing upright so as to squeeze them all in (sounds like the commuters on my train into work). As with the Tx and Rx which we looked at recently, there aren't any component markings on the board but a bit of thought and care, a lot of referring to the layout plan and gallons of coffee should make for success (I told myself). The boards are double-sided and connections to both sides are quite frequent; there are also wire links connecting one side of the board to the other and I found out the hard way that it was easy to miss these on the layout plan.

However, eventually I screwed up what courage I had left at the end of a day and plugged in the iron. I used a miniature 25watt iron and 22swg solder and there's no way you can use the trusty old kettle-mender on this board; the distance between many of the tracks is less than 1/32in so if you haven't got a small iron you'll need one for this kit.

Two minutes into the project and one very tedious point started to stare me in the face. The layout plan has all the components numbered, as does the circuit diagram, and this means that to find out which R or C or whatever goes in which 'oles you have to look it up on the chart provided. Well, that doesn't cause earthquakes and tidal waves but it's fatally easy to misread the chart and put the wrong bit in, especially when you've been concentrating for a couple of hours

"The VCO worked perfectly as soon as the appropriate volts were applied - which had to be good"

and the old brain cells are getting a bit tired without you realising it. First you discover that you haven't got any 100 Ohm resistors left and you wonder why; you definitely need one for R6. You then discover that you used it for R7, which should have been 10K – fine, mystery solved, but now try getting the 100 Ohm out without (a) frying your finger and (b) wrecking the tracks on the PCB. Oh brother, that ain't easy, the thin tracks are very likely to lift off the board as you attack them with your iron and melt the adhesive, and you need either desoldering wick or a solder sucker if you're not going to end up linking across the broken bit with pieces of wire.

I would have sold my granny into slavery to have seen component values as well as numbers on the circuit diagram – you can get a much better idea of what's supposed to be happening. Maybe a couple of copies of the layout plans, one with numbers and the other with values would have been even better.

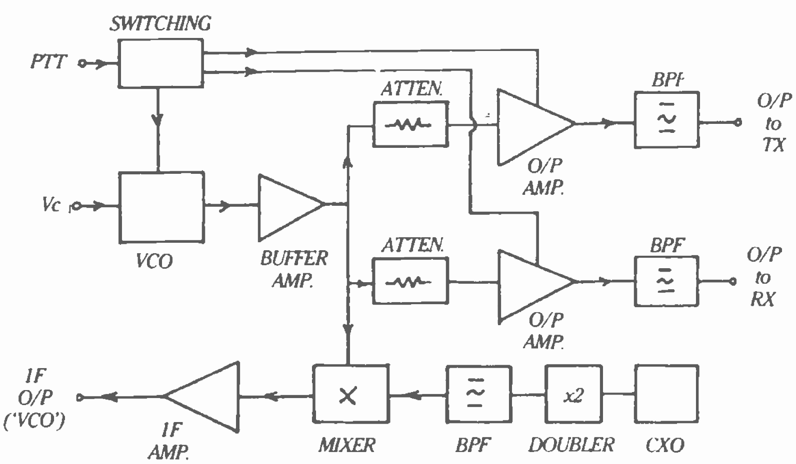

BLOCK DIAGRAM OF 70VC05 BOARD

Block diagram of the 70VC05 board. All the diagrams illustrated on these pages come with the building instructions from Wood and Douglas.

Stop wingeing, boy, and get on with it. I did resist the temptation to "do it my way", and so slavishly followed the instructions, although I damn near marked up the diagrams beforehand. Anyway, I built and tested the VCO first because I figured that I could use that to test the synth itself without a signal generator; the instructions themselves were clear and simple and there was no real problem. But then – another tedious point. The photocopied sheets are all done double-sided, presumably to save a shekel or two, but this means you're continually flipping loose pages backwards and forwards instead of being able to lay everything in front of you with all the relevant info being visible. As it was, the step-by-step instructions were on page 8 and the component value chart on page 7 – oh well, it'll give me some exercise . . . !

However, glory be, the VCO worked perfectly as soon as the appropriate volts were applied, which had to be good. It took a bit of time to set it up since the trimmers are fairly coarse in adjustment, but it all came together. I'd reckon you need a counter at the very minimum, good up to 150MHz, and a scope would be helpful although it isn't essential. The two VCO buffers, one for Tx and the other for Rx, did their stuff OK, so that was another bit sorted out. The final section of the board mixes the VCO output with a local oscillator on 138MHz, and the product of this mixer is an IF of about 6MHz on transmit and 2MHz on receive. This is supposed to be amplified by a broadband amp to provide about five volts peak-to-peak for the synth, but as it was I wasn't getting more than about two out of it; it was definitely coffee-and-fiddle time. After a lot of messing about, guess what – I'd put the MOSFET mixer in round the wrong way . . . With bated breath, I swopped it round and, thankfully, out popped about seven volts P-P – whoosh. I'd missed the indent mark on the source lead of the FET and thought from previous experience that on the drawing of the device it was being shown uppermost. Ah well, never again (joke).

Next it was time for the synth board and amazingly enough it seemed to work first time. The instructions call for quite a lot of logic checking, but when this was done and the synth board connected up to the VCO the system ought to be in a position to lock up and it should be possible to check the output frequencies against 432MHz channel numbers – they give you a chart for this. The only thing here is that you need to provide a binary code to the divider in order to tell it what division ratio you're after and the only really easy and straight-forward way to do this is to use a couple of BCD thumbwheels at about £3 each – in fact, unless you're prepared to build up some sort of logic-based controller with a keypad or thereabouts, or hard-wire some rotary switches (which is distinctly naff), thumbwheels are about the only practical solution. W & D don't provide any with the kit,

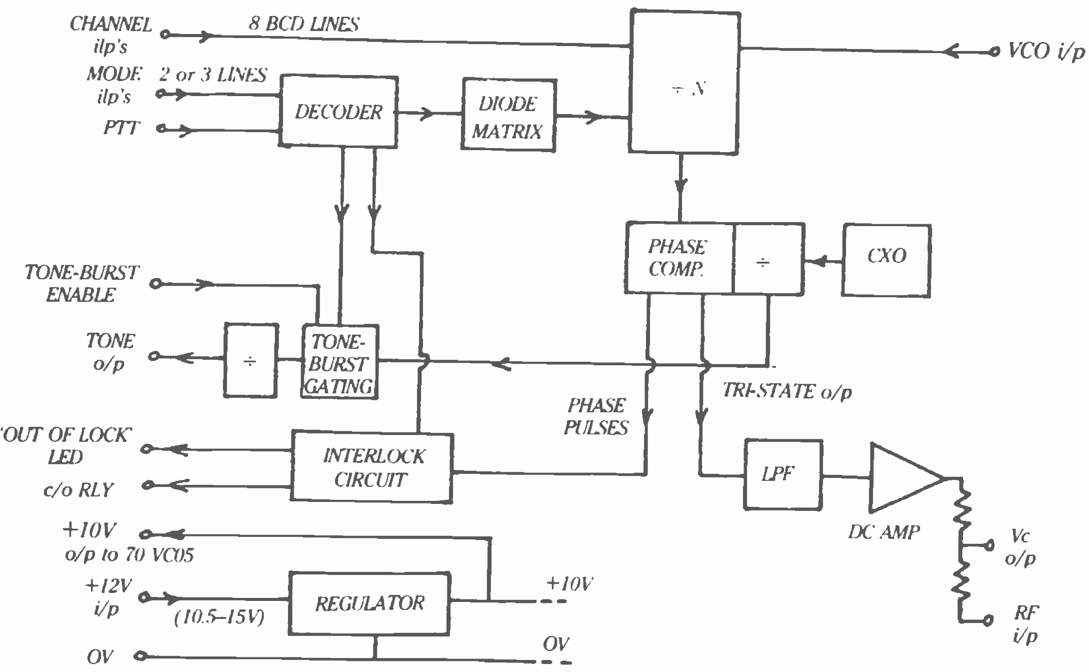

BLOCK DIAGRAM OF 70SY25B BOARD

Block diagram of the 70SY25B board.

although they will flog you some separately. There's a drawing provided with the kit which shows you how to connect them up, although there's also a lot of stuff on the drawing that's to do with their "Proscan" channel scanner and this is a touch misleading because it shows diodes in series with some of the outputs from the thumbwheels. These are completely unnecessary if you aren't using the scanner, unless I've had a total cerebral collapse, and I reckon a dedicated drawing would have been a better idea.

So having got it all going, I thought I'd see how it worked when connected up to a receiver; in this case the W & D 70FM05RS that we used last time. You need to modify the local oscillator a bit to provide the necessary interface, and this requires an R and C and also a choke; this latter was a pig because it needed to be a miniature one to fit on the board and it wasn't supplied. However, W & D do sell them separately. Having got it all together, I switched on and it worked! It all seemed to be very good, and GB3HR came storming on through six inches of wire, although I could hear a bit of quieting on some frequencies which must have been due to stray harmonics from various bits of the VCO. A bit of screening, in the shape of a die-cast box or thereabouts, ought to sort that one out although I hadn't got one handy at the time.

The last giant step was to try it out with the TX – this was the same as the one we used last time without the oscillator and with just a bit of gain at 144MHz, a tripler and a PA on 432MHz. The synth output is designed to drive this directly, and after mistakenly tuning up to 288 instead of 432MHz it worked like the proverbial charm, no sweat. I was ever so pleased with myself.

Impeccable performance

Synthesised 70cm FM transceiver kits don't exactly grow on trees, so Wood and Douglas deserve ten marks out of ten for this lot. I was very pleased indeed with the finished article, both physically and technically; the design was obviously repeatable and the technical performance was impeccable. It's all well laid out, if you're fairly familiar with small-scale soldering and not afraid of miniature circuitry, and there seems not the slightest reason why anyone with no previous experience with CMOS logic couldn't get it all to work beautifully.

A bit of RF experience would certainly help, since that's undoubtedly the trickiest part of it, and the only real criticism I have is the way in which they've laid out the instructions. Let's hope they find it in their hearts to breathe on them a bit, because this really is a superb little kit and it deserves to sell like the proverbial hot cakes.