. ilters based on frequency-dependent negative resistors offer the performance of *LC* filters but without the bulk and expense. And component intolerance is lower than for other lf filters types. Ian Hickman explains.

FILTERS

using negative resistance

When it comes to filters, it's definitely a case of horses for courses. At rf the choices are limited; for .unable filters covering a substantial percentage bandwidth, it has to be an *LC* filter If the tunable elements are inductors, you have a permeability tuner; alternatively, tuning may use variable capacitance, or varactors. Fixed frequency filters may use *LC*s, quartz crystals, ceramic resonators or surface-acoustic-wave (SAW) devices, while at microwaves, the 'plumbers' have all sorts of ingenious arrangements.

At audio frequencies, *LC* filters are a possibility However, the large values of inductance necessary are an embarrassment, having a poor Q and temperature coefficient, apart from their bulk and expense

One approach is to use '*LC*' circuits where the inductors are active circuits simulating inductance There is a number of these, and **Fig. 1** is an example. For high-pass filters, synthetic inductors with one end grounded, Fig. **1a**), suffice, but for low-pass applications, rather more complicated circuits simulating floating inductors are required, Fig. **1b**).

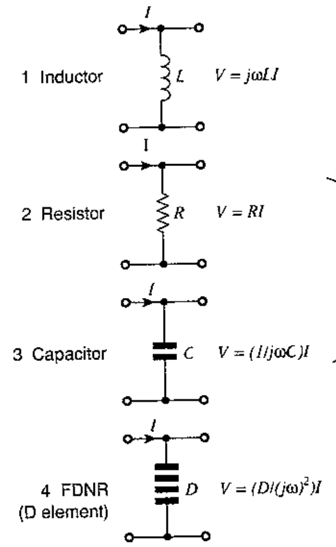

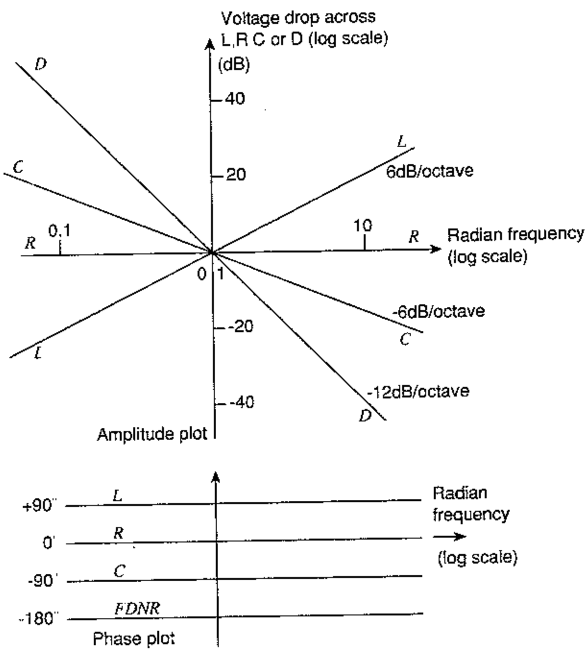

Fig. 2. Resistance (reactance?) of an fdnr – also known as a supercapacitor or D element – varies with frequency.

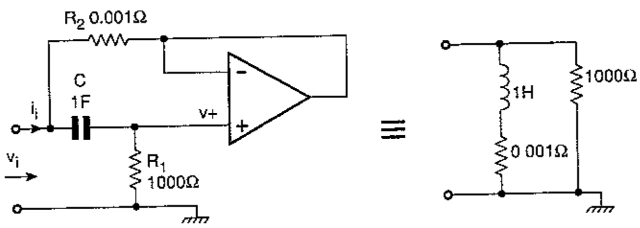

Fig. 1. Synthetic inductors.
a) A 1H 'inductor' with one end grounded. Q is 10 at 0.00159Hz, and proportional to frequency above this. Below, it tends to a 0.001Ω resistor, just like the corresponding real inductor.

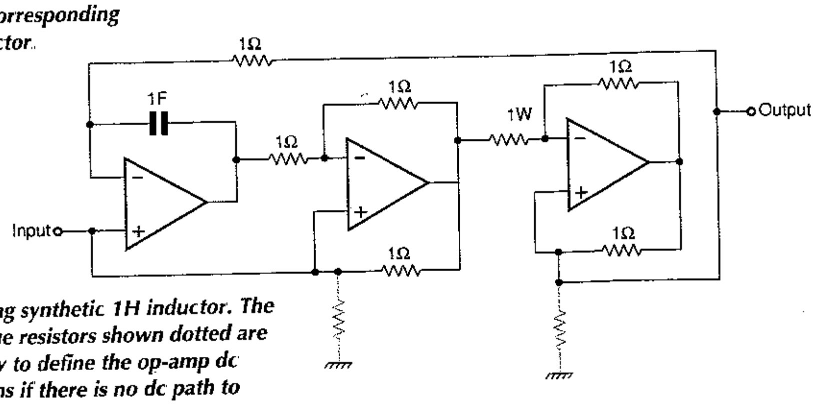

b) Floating synthetic 1H inductor. The high-value resistors shown dotted are necessary to define the op-amp dc conditions if there is no dc path to ground via input and output.

More recently, switched-capacitor filters have become available, offering a variety of filter types, such as Butterworth, Bessel and Elliptic. These vary in complexity up to eight or more poles.

For narrow band-pass applications, a strong contender must be the *N*-path filter. This scheme uses switched capacitors but is not to be confused with switched capacitor filters; it works in an entirely different way. However, both switched capacitor and *N*-path filters are time-discrete circuits, with their cut-off frequency determined by a clock frequency. Hence both types need to be preceded by an anti-alias filter – and usually followed by a low-pass filter to suppress clock frequency hash That is the downside: the upside is that tuning is easy – simply change the clock frequency

The cut-off or centre frequency of a switched capacitor filter scales with clock frequency, but the bandwidth of an N-path filter does not.

Where a time-continuous filter is mandatory, various topologies are available, including Salen and Key and Rausch. An interesting and useful alternative to these and to LC filters, with either real or simulated inductors, is the fdnr filter, which makes use of frequency-dependent negative resistances.

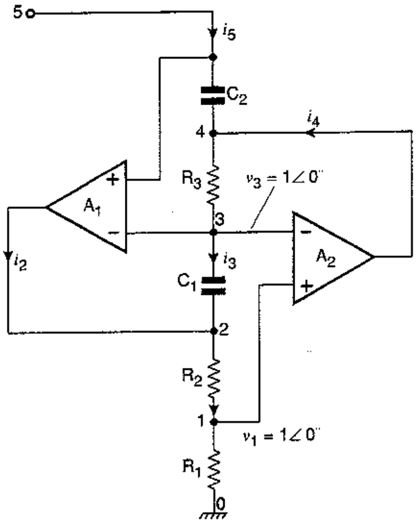

Fig. 3a). Circuit diagram of an fdnr. If V_1 is the voltage at node 1, etc., then V_1=V_3=V_5. Also, i_1=i_2+i_3 and i_3=i_4+i_5.

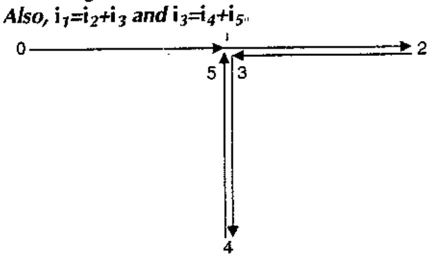

Fig. 3b). Voltage vector diagram for (a) when R_1=R_2=R_3=R, C_1=C_2=C and f=1/2πCR.

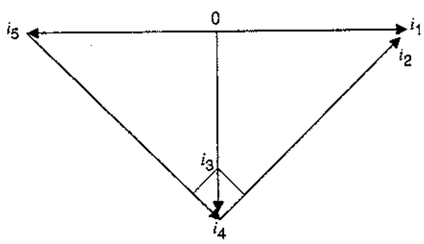

Fig. 3c). Current vector diagram for (a), for the same conditions as (b).

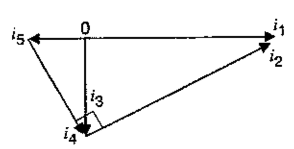

Fig. 3d). As (c) but for f=1/4πCR. Note that i_2 and i_4 are always in quadrature.

What is an fdnr?

A negative resistance is one where, when you take one terminal positive to the other, instead of sinking current, it sources it – pushes current back out at you.

As the current flows in the opposite direction to usual, Ohm's law is satisfied if you write $I=E/-R$, indicating a negative current in response to a positive potential difference, or pd. This would describe a fixed, or frequency independent, negative resistance. But fdnrs have a further peculiarity – their resistance, reactance or impedance, call it what you will, varies with frequency. Just how is illustrated in **Fig. 2**.

With inductors, the voltage leads the current by 90°; with capacitors, it lags by 90°. Combining these with resistive terminations, where the voltage leads/lags the current by 0°, you can make filters. Such filters may be high-pass, band-pass, low-pass, or whatever you want.

It was pointed out in a famous paper[1] that, by substituting for L, R (termination) and C in a filter, components with 90° more phase shift and 6dB/octave faster roll than the L, R and C, exactly the same transfer function could be achieved.

Referring to Fig. 2, L, R and C are replaced on a one-for-one basis by R, C and fdnr respectively. An fdnr can be realised with resistors, capacitors and op-amps, **Fig. 3**.

So how does an fdnr work?

Analysing Fig. 3 provides the answer. Looking in at node 5, you see a negative resistance; but what is its value?

First of all, note that the circuit is dc stable. At 0Hz, where you can forget the capacitors, A_2 has 100% negative feedback via R_3, and its non-inverting input is referenced to ground.

Likewise, A_1 has its non-inverting input referenced to ground, assuming there is a ground return path via node 5. It also has 100% negative feedback; A_2 is included within this loop.

The clearest and easiest way to work out the ac conditions is with a vector diagram; just assume a voltage at node 1 and work back to the beginning. Thus in Fig. 3, assume that $V_{1,0}$, i.e. the voltage at node 1 with respect to node 0 or ground, is 1V ac, at a frequency of 1 radian per second (1/(2π) or 0.159Hz), and that R_1=R_2=R_3=1Ω, C_1=C_2=1F. Thus the voltage at node 1 is represented in Fig. **3b)** by the line from 0 to 1, of unit length, the corresponding current of 1A being shown as i_1 in Fig. **3c)**.

Straight away, you can mark in, in b), the voltage $V_{2,1}$, because R_1=R_2, and node 1 is connected only to an (ideal) op-amp which draws no input current. So $V_{2,1}$ equals $V_{1,0}$ as shown.

But assuming A_2 is not saturated, with its output voltage stuck hard at one or other supply rail, its two input terminals must be at virtually the same voltage. So now $V_{3,2}$ can be marked in, taking one back to the same point as node 1. Given $V_{3,2}$, the voltage across C_1, whose reactance at 0.159Hz is 1Ω, the current through it can be marked in as i_3 in Fig. 3c).

Of course, the current through a capacitor leads the voltage across it, and i_3 is accordingly shown leading the voltage $V_{3,2}$ by 90°. Since i_1=i_2+i_3, i_2 can now be marked in as shown.

As i_3 flows through R_3, $V_{4,3}$ can now be marked in, and as the voltages at nodes 5 and 3 must be equal, $V_{5,4}$ can also be marked in. The current i_5 through C_2 (reactance of 1Ω) will be 1A, leading $V_{5,4}$ as shown. Finally, as i_3=i_4+i_5, i_4 can be marked in, and the voltage and current vector diagrams (for a frequency of $1/2\pi CR$) are complete.

The diagrams show that $V_{5,0}$ is 1V, the same as $V_{1,0}$, but i_5 flows in the opposite direction to i_1 – the wrong way for a positive resistance. Fig. 3d) shows what happens at f=$1/4\pi CR$, half the previous frequency. Because the reactance of C_1 is now 2Ω, i_3 is only 0.5A, and therefore $V_{4,3}$ is only 0.5V. Now, there is only ½V ($V_{5,4}$) across C_2, but its reactance has also doubled. Therefore i_5 is now only 0.25A; not only

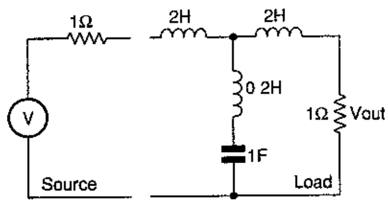

Figs. 4a), above, and b), below. A low component count elliptic low-pass filter with a minimum attenuation of 36dB from twice the cut-off frequency upward, the price being as much as 1dB pass-band ripple. The minimum capacitor design of 4a) is more convenient than 4b) for conversion to an fdnr filter.

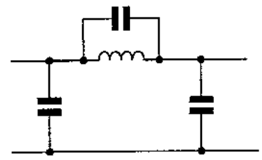

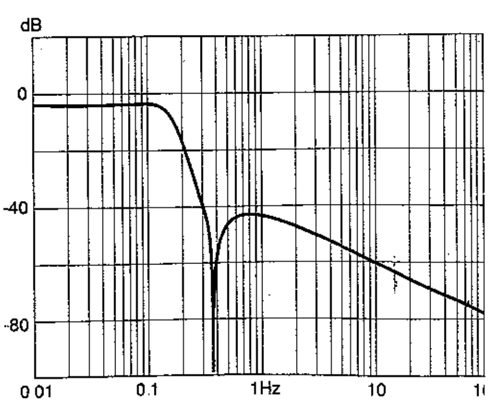

Fig. 4c). The frequency response of the filter.

is the current negative (a 180º phase shift), it is inversely proportional to the square of the frequency, as shown for the fdnr in Fig. 2

Pinning down the numbers

Looking in at node 5, then, appears like a -1Ω resistor at 0.159Hz, but you need to know how this ties up with the component values. The values of the vectors can be marked in, on Figs. 3b) and 3c), starting with $V_{1,0}$=1V. Then $V_{2,1}=R_2/R_1$, and $V_{3,2}=-R_2/R_1$. It follows that $i_3=(-R_2/R_1)/(1/j\omega C_1)=-j\omega C_1 R_2/R_1$. Voltage $V_{4,3}=R_3 i_3=-j\omega C_1 R_2 R_3/R_1$, and $V_{5,4}=-V_{4,3}$.

So $i_5=-V_{4,3}/(1/j\omega C_2)=j\omega C_1 j\omega C_2 R_2 R_3/R_1$. Looking in at node 5 the resistance is $V_{5,0}/i_5=V_{1,0}/i_5$, where $V_{1,0}$=1V. So finally the fdnr input looks like,

$$fdnr=R_1/(j\omega C_1 j\omega C_2 R_2 R_3)$$
$$=-R_1/(\omega^2 C_1 C_2 R_2 R_3) \qquad (1)$$

With 1Ω resistors and 1F capacitors, this comes to just −1Ω at omega=1rad/s, or 0.159Hz. To get a different value of negative resistance at that frequency, clearly any of the *R*s or *C*s could be changed to do the job, but it is best to keep all the *R*s equal – at least roughly – and the same goes for the *C*s.

As a cross check on equation (1), note that it is dimensionally correct. The units of a time-constant *CR* are seconds, while the units of frequency are 1/seconds, be it cycles or radians per second. Thus the units in the denominator cancel out and, with a dimensionless denominator, the expression has the units of the numerator R_1, which is ohms.

A practical example

Designing an fdnr filter starts with choosing an *LC* prototype. Consider a simple example – a low-pass filter with the minimum number of components, which must reach an attenuation of 36dB at little more than twice the cut-off frequency. This is a fairly tall order, but a three-pole elliptic filter will do the job, if you allow as much as 1dB pass-band ripple.

A little experimentation with a CAD program came up with the design in **Fig. 4a)**. This has nice, round component values, although the cut-off frequency is just a fraction below the design aim of 1 radian per second, but never mind, it will do for starters.

If you were designing an *LC* filter as such, you would certainly choose the π section of Fig. **4b)**, rather than the tee section, as the π section is the minimum inductor version. But for an fdnr filter, the minimum capacitor version is preferable, as the *C*s become fdnrs (fairly complicated), whereas the *L*s become *R*s and are therefore cheap and easy.

But before passing on to consider the fdnr, note that the computed frequency response of the normalised 1Ω impedance *LC* filter is as shown in Fig. **4c)**.

The low frequency attenuation shows as 6dB rather than 0dB. This is because the 1Ω impedance of the matched source (a 2V emf ideal generator behind 1Ω) is considered here as part of the filter – not as part of the source. To the 2V generator emf, which is what the CAD program models as the input, the source and load impedance appear as a 6dB potential divider.

The fdnr version of the filter is shown in **Fig. 5**. Not only do the *L*s become *R*s and the *C*s fdnrs, but the source and terminating resistors become capacitors. In an *LC* filter, the source and terminating resistors would usually be actually part of the source and load respectively. But an fdnr filter at audio frequencies will be driven from the 'zero' output impedance of an op-amp and feed into the nearly infinite impedance of another. As a result, you must provide the terminations separately if you want the response to be the same as the prototype *LC* filter.

In Fig. 5, the inductors have been replaced with resistors on an ohm-per-henry basis, and the *R*s and *C*s converted to *C*s and fdnrs similarly. As it happens, the required fdnr value is −1Ω, so values of 1Ω and 1F in the circuit of Fig. 3 will do the job.

If one had used the tabulated values for a 1dB ripple, 35dB A_s three-pole filter, e.g. from Ref. 2, **Fig. 6**, the required value of C_2 in the tee-section version, would have been 0.865F.

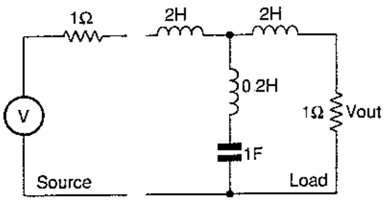

Fig. 5. An fdnr version of the low-pass filter.

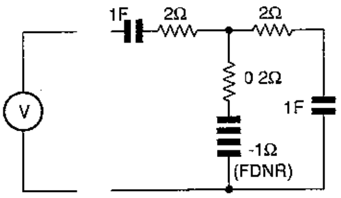

Accordingly, from equation (1), R_1 in Fig. 3 would become 0.865Ω, or you could change R_2 and/or R_3 to achieve the same effect. Alternatively, you could scale C_1 and/or C_2, but it is best to leave them at 1F – the reason for this will become clear later.

Having arrived at a 'normalised' fdnr filter design, i.e. one with a 0.159Hz cut-off frequency, the next step is to denormalise it to the desired cut-off frequency – let's say 10kHz in this case.

There is no need to change the *R*s at this stage, but to make the fdnr look like −1Ω, or −0.865, or whatever, at 10kHz, the capacitor values must be divided by 2π times 10,000. And since the termination capacitors must also look like 1Ω at this frequency, they must be scaled by the same ratio.

You now have a filter with the desired response and cut-off frequency, but the com-

Fig. 6. Tabulated normalised component values for three-pole 1dB pass-band ripple elliptic filters with various values of A_s at Ω_s. Ω here means the same as ω elsewhere in the article.

Ω_s	A_s (dB)	C_1	C_2	L_2	Ω_2	C_3
1.295	20	1.570	0.805	0.613	1.424	1.570
1.484	25	1.688	0.497	0.729	1.660	1.688
1.732	30	1.783	0.322	0.812	1.954	1.783
2.048	35	1.852	0.214	0.865	2.324	1.852
2.418	40	1.910	0.145	0.905	2.762	1.910
2.856	45	1.965	0.101	0.929	3.279	1.965
Ω_s	A_s (dB)	L_1	L_2	C_2	Ω_2	L_3

© 1958 IRE (now IEEE)

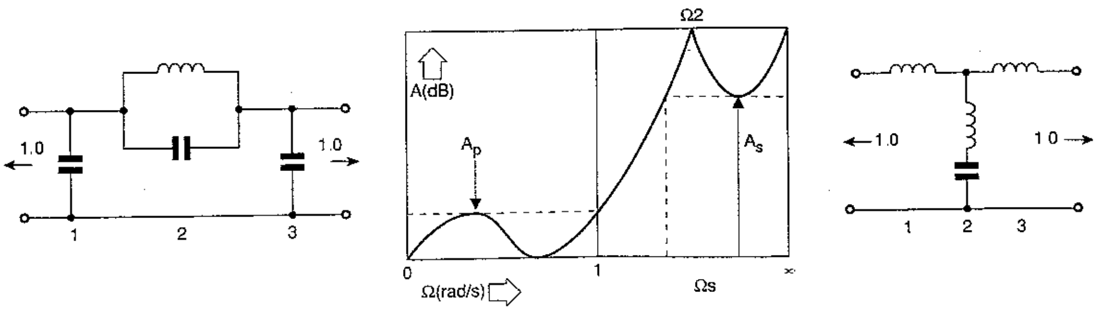

ponent values shown in round brackets in **Fig. 7**, are a little impractical. This is easily fixed, by a further stage of scaling.

Since resistors are more easily obtainable in E96 values and 1% selection tolerance, it pays to scale the 15.9µF capacitors to a round value, such as 10nF. So all impedances must be increased by this same ratio: N=1590 – the resistors multiplied by N and the capacitors divided by N.

Conveniently, the *C*s in the fdnr are the same value as the terminating capacitors, if, as recommended, any change in the required normalised fdnr negative resistance was effected by changing the *R* values only. The resultant practical component values are shown in square brackets in Fig 7.

One peculiarity of an fdnr filter is due to its use of capacitive terminations. The impedance of these varies with frequency and, notably, becomes infinite at dc (0Hz). Thus any practical fdnr filter would have infinite insertion loss at this frequency.

This is remedied by connecting resistors in parallel with the terminating capacitors, to determine the 0Hz response. They are shown in Fig. **7a)** and have been chosen, taking into account the two 3,180Ω resistors, to provide 6dB attenuation at 0Hz. This is done to match the filter's pass-band 6dB loss. With the addition of these, Fig. 7a) is now a practical, fully working low-pass filter, the computed frequency response of which is shown in Fig **7b)**

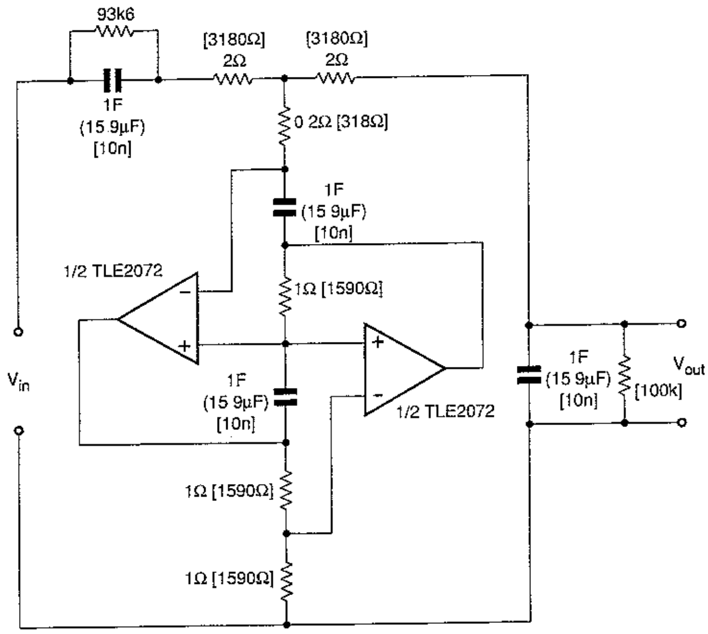

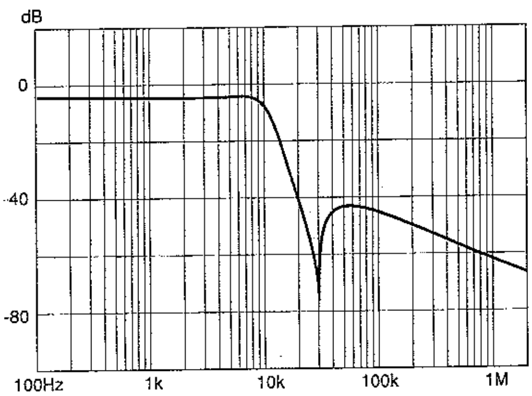

Fig. 7a). Complete fdnr filter with 10kHz cut-off frequency. The values in round brackets are a little impractical, but are easily scaled to more sensible values in square brackets.
Fig. 7b). Computed frequency response of the above filter. The cut-off frequency (at –1dB) is a shade below the intended value, as was that in Fig. 4c).

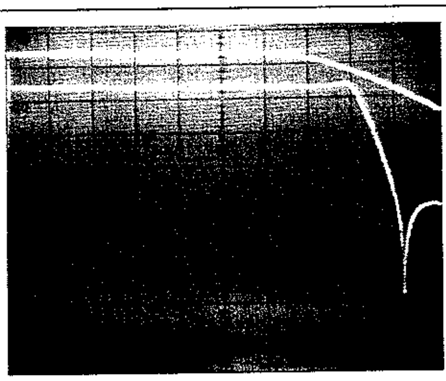

Fig. 8. Actual frequency response of the circuit of Fig. 7. The horizontal scale is logarithmic frequency, the left-hand vertical being 20Hz, the third, sixth and ninth vertical graticule lines representing 200Hz, 2kHz and 20kHz respectively. Horizontal graticule lines are at 10dB intervals.
Upper trace, generator reference level top of screen, representing the source emf. This trace was recorded with the shunt leg of the filter open circuited with the 318Ω resistor removed. Lower trace, response of complete filter withthe 318Ω resistor replaced. Reference level has been moved dowan one graticule division for clarity.

Log sweeps and IF bandwidths

The response shown in **Fig. 8** was taken using the log frequency base mode of the *HP3580A* 0-50kHz spectrum analyser. In this mode, the spot writes the trace across the screen at a steady rate, taking about 6 seconds to sweep from 20Hz to 44.3kHz. Thus the sweep rate in hertz per second increases greatly as the spot progresses across the screen.

This means that if a resolution bandwidth narrow enough to resolve frequency components encountered near the start of the sweep (e.g. 1Hz or 3Hz bandwidth) is used, then near the end of the sweep the analyser will be passing through any signals far too fast to record their levels even approximately.

On the other hand, if a bandwidth such as 300Hz – wide enough to accurately record signal amplitudes in the 20kHz region – is used, the zero frequency carrier breakthrough, response will extend half way across the screen. So, when using log sweep mode to record the amplitudes of stationary signals, compromises must be made.

But this is not the case in Fig. 8, for here the only signal of interest is the output of the tracking generator, to which the analyser is, by definition, always tuned. So the analyser is at no time sweeping through a signal and in principle it might seem that the 1Hz bandwidth could be used. There is a restraint on the bandwidth, however, set by the rate at which the signal amplitude changes. This can get quite fast in the vicinity of a notch, and accordingly the trace in Fig. 8 was recorded with a 30Hz resolution bandwidth. At 10Hz bandwidth, the notch appeared shunted slightly to the right and its full depth was not recorded.

On the other hand, at a 100Hz bandwidth, the notch response was identical to that shown, but the left hand end of the trace, representing 20Hz, was elevated slightly, due to the zero frequency carrier breakthrough response. If, due to a fortuitous conjunction of component tolerances, the actual notch depth had been much deeper than it actually was, the 100Hz bandwidth would have been necessary to capture it. In that case, it would be better to switch back to linear frequency base mode, and make the notch measurement at a span of 100Hz, or even 10Hz, per horizontal division

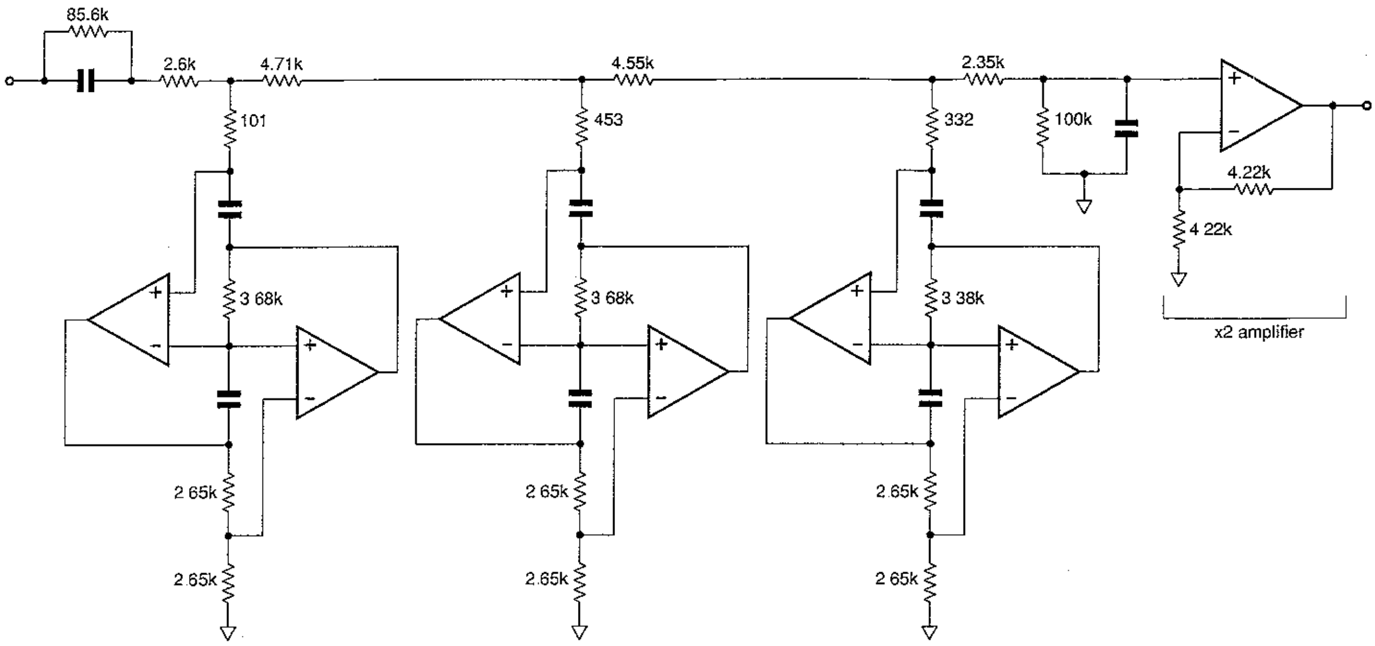

This is all fine in theory, but does it work in practice?

Proof of the pudding

Ever of a pragmatic – not to say sceptical – turn of mind, I determined to try it out for real. So I made the circuit of Fig 7a) up almost exactly as shown, and tested it using an *HP3580A* audio frequency spectrum analyser.

The circuit was driven from the *3580*'s internal tracking generator. There were minor differences. Whereas the plot of Fig 7b) was modelled with *LM318* op-amps, these were not to hand, so a *TLE2072CP* low-noise, high-eed j-fet input dual op-amp was used. This is a handy Texas Instruments device with a 35V/μsec slew rate and accepting supplies in the range ±2 25V to ±19V.

The required resistor values were made up using combinations of preferred values, e.g. 82kΩ+12kΩ for 93.6kΩ, 270Ω+47Ω for 318Ω, etc, all nominal values thus obtained being within better than 1% of the exact values. 100kΩ+12kΩ was used for the terminating resistor, to allow for the 1MΩ input resistance of the spectrum analyser in parallel with it. The resistors were a mixture of 1% and 2% metal film types, except the 47Ω, which was 5%. The four 10nF capacitors were all 2.5% tolerance polystyrene types.

Although the circuit worked, its response was not exactly as hoped, due to being driven from the *3580A*'s 600Ω source impedance. So a *TLE2027* single op-amp – not to be confused with the *TLE2072* dual device used for the main circuit – was used as a unity gain buffer to drive the filter from a near-zero urce impedance. Its output level was set at the top of the screen, **Fig. 8**

First, the filter action was disabled by removing the 318Ω resistor, leaving a straight-through signal path. The upper trace shows the 6dB loss due to the terminations mentioned earlier. It also shows a first-order roll-off due to the effect of the terminating capacitor at the load end, with the two 3180Ω resistors.

Response of the complete filter, with the 318Ω resistor replaced, is shown in the lower trace. The reference level has been moved down one graticule division for clarity. The –1dB point is at two divisions in from the left, which, given the horizontal scaling of three divisions per decade, corresponds to 9 3kHz – pretty close agreement with the predicted performance of Fig 7b)

In logarithmic frequency mode, the analyser's bandwidth extends only up to 44.3kHz. But this is far enough to see that the notch frequency and the level of the return above it, 36dB below the l.f. response, also agree with the computed results

Others like it, too

Various applications have been found for fdnr filters, especially in measuring instruments. The advantage here is that the response is predictable and close to the theoretical.

Some other active filter sections, when combined to synthesise filters of a higher order, show a higher sensitivity to component tolerances. This is a disadvantage where the filters are used in the two input channels of an instrument, which requires close matching of the channel phase and amplitude responses. For this reason, fdnr filters were used in the input sections of the *HP5420A*, **Fig. 9**[3] ■

Fig. 9a). Normalised seven-pole elliptic LC prototype filter, and b), below, derived fdnr input antialiasing filters used in the HP5420A.

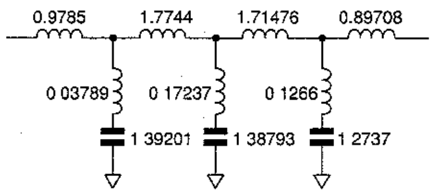

References

1. Bruton, L T. 'Network Transfer Functions Using the Concept of Frequency Dependent Negative Resistance', *IEEE Transactions on Circuit Theory* Vol CT-16, pp405-408, August 1969.

2. Hickman, I, 'Newnes Practical RF Handbook', 1993, ISBN 0 7506 0871 4, p245.

3 'Front-end Design for Digital Signal Analysis Patkay'. Chu and Wiggers, *Hewlett Packard Journal*, Vol 29, No 22, October 1977. p9