10 Quartz crystal filters

10.1 INTRODUCTION

Electronic filters are signal-processing devices whose function is to discriminate in some way between the different frequency components of an input signal. *Lowpass* and *highpass* filters are, respectively, filters that transmit signals at frequencies below or above a defined cut-off frequency and attenuate those at frequencies above or below the cut-off. *Bandpass* filters transmit all frequencies between defined upper and lower limits, and attenuate frequencies outside those limits. *Bandstop* filters attenuate frequencies between upper and lower limits and transmit all other frequencies. *Allpass* filters transmit all frequencies, but introduce time delays depending on the frequency of the individual signal components.

Ideally, filters would have zero attenuation in their *passband(s)*, that is those frequency intervals in which signals are to be transmitted, and infinite attenuation in their *stopbands*. Additionally, the phase shift produced by the filter for signals in the passband would be proportional to the frequency so that the *group delay*, which is the derivative of the phase with respect to the frequency, would be constant in the passband. This would then allow the undistorted transmission of a signal made up of components with frequencies lying entirely within the passband. In practice, these characteristics of an ideal filter are not realizable with a finite number of components, so that real filters have to be defined in terms of the following parameters (or some equivalent set):

(1) maximum attenuation allowed in the passband;
(2) minimum attenuation required in the stopband;
(3) width of the *transition region* from those points in the passband where the maximum attenuation allowed is last achieved, to those adjacent points in the stopband where the minimum attenuation required is first achieved;
(4) allowed deviation from a linear phase or constant group delay characteristic in the passband.

Figures 10.1 and 10.2 illustrate these parameters for the simple case of a lowpass filter.

Generally speaking, the complexity, that is the number of elements, needed in a filter increases with the level of attenuation required in the stopband for a

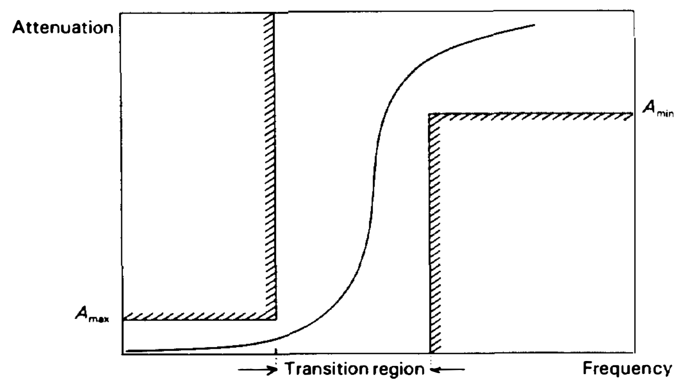

Fig. 10.1 Amplitude specification, LP filter.

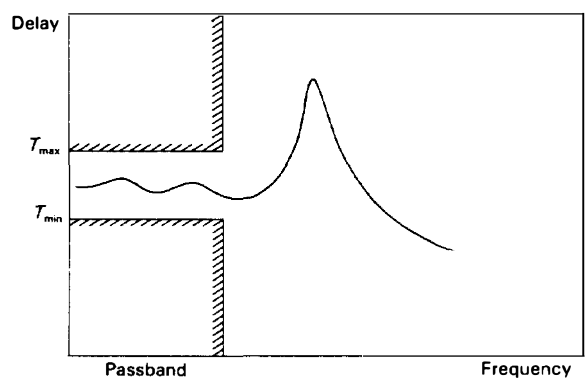

Fig. 10.2 Group delay specification, LP filter.

given transition region, or alternatively, for a given level of stopband attenuation, the complexity increases as the width of the transition region is reduced. Additionally, for the common class of filters known as *minimum-phase* filters, the phase or group delay response is not independent of the amplitude response. Specifically, for such filters, a narrow transition region between pass and stopbands is accompanied by large variations in the group delay near the passband edges, whereas, on the other hand, a flat group delay response implies a rounded passband and a wide transition region. To overcome this requires either additional complexity in the filter, such as the use of allpass sections to *equalize* the group delay response, or else the use of non-minimum phase filters, which implies greater complexity in the filter design process.

Electronic filters can be realized in many different forms, ranging from the classical *LC filter*, utilizing inductors and capacitors, through *mechanical*, *ceramic* and *crystal* filters, various forms of *microwave* filters, *surface acoustic wave* (SAW) devices, up to active filters, again realized in various different forms, such as *active RC* filters or *switched-capacitor* devices. *Crystal filters* are passive electro-mechanical devices in which crystal resonators are used in place of *LC* resonant circuits. Crystal filters can themselves be divided into two categories, *discrete* and *monolithic* (or *polylithic*). The former use individual crystal resonators as circuit components, whereas the latter use resonators that are acoustically coupled by virtue of their fabrication in close proximity on the same crystal blank. Although resonator materials other than quartz have been used in the past, and are still actively being investigated, quartz remains the most commonly used, thanks to its unique combination of very low losses and high stability with respect to both time and temperature.

The specific applications in which quartz crystal filters come into their own are those requiring narrowband highly selective devices. There have historically been two major application areas, the first being the use of crystal filters in telephone *frequency-division multiplex* (FDM) systems. In such systems, many telephone channels are transmitted simultaneously over a wideband transmission line, such as a coaxial cable, and separated at the receiving end by banks of *channel filters*, each bank consisting of several filters with a nominal bandwidth of approximately 3 kHz and centre frequencies equally spaced across the frequency range occupied by the group of channels. In order to prevent interference between the channels, the filters must be extremely selective, and the high Q of quartz crystal resonators that allowed the construction of such filters was a key factor in the successful implementation of FDM telephone systems. The application of crystal filters in this field was well established by the 1940s and accounted at that time for the vast bulk of crystal filters produced. However, both the introduction of *time-division multiplex* (TDM) systems and the use of *mechanical* rather than crystal channel filters has meant that the importance of the latter has decreased.

The second major application area for crystal filters, and currently most important, is in radio communication and electronic navigation systems. In such systems the prime use is as *intermediate frequency* (IF) filters, although there are some secondary applications such as 'front-end' or *antenna* filters, and occasional applications for bandstop filters. Sometimes also crystals are 'imbedded' in wider band networks to provide sharp attenuation peaks, but neverthless the overwhelming majority of crystal filters are used in narrowband, bandpass roles. In VHF and UHF radio systems, they are used to provide the necessary discrimination between closely spaced communications channels. In HF systems, besides this function, crystal filters are also used to separate the sidebands in SSB transmissions and to act as 'roofing' filters in double conversion systems.

10.2 ELEMENTARY CIRCUITS: DISCRETE CRYSTAL FILTERS

The general problem of the design of a crystal filter falls into two distinct parts. On the one hand, there is the problem of designing a filter network capable of accommodating the particular equivalent circuit of a crystal resonator, and on the other, the problem of designing resonators with the appropriate values for the elements of the equivalent circuit as determined by the network design. Usually, the network design problem is handled with the tools developed for the synthesis of *LC* filters, which in itself is a highly developed subject. The original approach to the design of *LC* filter circuits, the *image parameter* method, has in modern practice been superseded by 'exact' methods which in their full generality rely heavily upon the use of digital computers to carry out the numerical calculations involved.

[Note: there is a very large literature on the subject of network synthesis, although relatively little has been written on the particular subject of crystal filters. For references to the original work in the field and to modern developments, the books by Zverev (1967), Humpherys (1970), Weinberg (1962) and Temes and Mitra (1973) should be consulted. Mention should also be made of the collection of reprints edited by Sheahan and Johnson (1977), which contains several important papers on crystal filters, and of the extensive bibliography on crystal, mechanical and SAW filters contained in Gerber and Ballato (1985).]

The 'exact' methods of filter synthesis can be shown to yield 'optimum' designs, with 'optimum' being defined in terms of specific performance characteristics. However, they do not in general give the designer the same degree of control over the final structure of the filter as would have been implicit in the old image parameter method. Thus it is often the case that in designing a crystal filter, the network provided by an exact synthesis technique only provides a starting point, with more or less extensive network transformations being required to produce a form in which the equivalent circuit of the crystal can readily be accommodated. However, once the network is obtained, its analysis can be straightforward, and the operation of the standard crystal filter configurations can be understood on basic principles, even if the specific choice of element values has to be based on complex calculations.

Considering first discrete crystal filters, that is, filters using individual crystal resonators to replace *LC* resonant circuits, a broad distinction can be made between narrow, intermediate and wideband filters. (In the overall context of electronic filters, all these categories would be treated as narrowband: the present distinction is made strictly within the context of crystal filters.) *Narrowband* units are those that can be realized in principle by crystals and capacitors only; *intermediate* band filters use inductors either wholly or partially to compensate for the crystal static capacitances and other stray capacitances; and *wideband* units use inductors as integral elements of the filter.

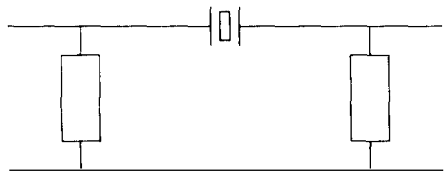

Fig. 10.3 Crystal in series arm of ladder circuit.

Narrowband crystal filters can be realized as either ladder or lattice structures. When a crystal is used in a series arm of a ladder structure as in Fig. 10.3, it is clear that the series resonance frequency ofthe crystal must lie in or close to the passband of the filter, whereas the anti-resonance frequency must lie in the stopband. For a ladder with a crystal in a shunt arm (Fig. 10.4), the reverse is true, and in either case it follows that the bandwidth of the filter, that is the width of the passband, must be substantially less than the separation of the series and anti-resonance frequencies of the crystal. From Chapter 6, if the motional capacitance and static capacitance of the crystal are C_1 and C_0, respectively, this spacing is given by

$$\Delta f/f_s = C_1/2C_0 \tag{10.1}$$

For a fundamental mode AT-cut crystal, the capacitance ratio C_0/C_1 is approximately 200, so that from Eqn (10.1) the fractional bandwidth of a narrowband filter using AT-cut fundamentals must be substantially less than 1/400 or in percentage terms substantially less than 0.25%. For overtone crystals, the capacitance ratio reduces with the square of the overtone order, so that with third overtone crystals for example, the maximum narrowband ladder bandwidth must be less than 0.025%.

This bandwidth restriction is a severe limitation on the ladder structure realization. In addition, simple ladders with crystals either in the series arms only, or in the shunt arms only, necessarily have attenuation peaks either above or below the passband which although possibly of use in SSB applications are generally an embarrassment. Both disadvantages are avoided in the lattice arrangement of Fig. 10.5. Although, following convention, only two lattice arms are shown in the figure, the full lattice arrangement is in fact fully symmetrical and balanced, being equivalent to a bridge network with the filter input being applied to one pair of nodes and the output taken

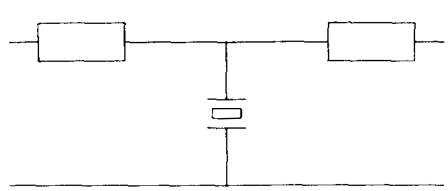

Fig. 10.4 Crystal in shunt arm of ladder circuit.

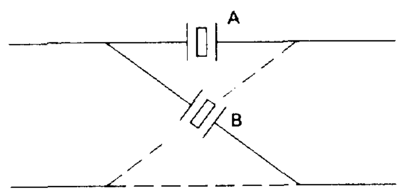

Fig. 10.5 Crystal lattice.

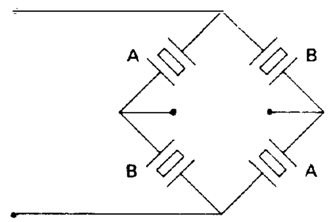

Fig. 10.6 Lattice redrawn as a bridge.

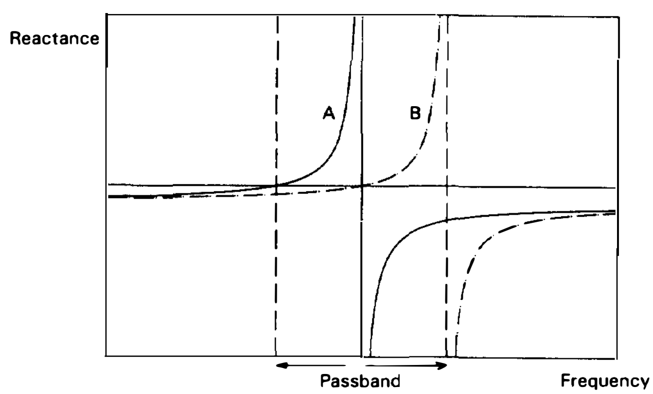

Fig. 10.7 Reactance plot for lattice filter.

from the other pair, as in Fig. 10.6. Clearly, when the lattice arms A and B have equal impedances, the bridge is balanced and there will be no output, that is there will be a pole of attenuation.

An image parameter analysis shows that if the lattice arms A and B have impedances $Z_A = jX_A$ and $Z_B = jX_B$, and if the lattice is terminated in the image impedance $(Z_AZ_B)^{1/2}$, then the attenuation of the lattice is zero when X_A and X_B have opposite signs. If each lattice arm contains a single crystal and the crystal resistances are neglected, then the arm reactances X_A and X_B will be as drawn in Fig. 10.7, where the series frequency of crystal B has been chosen to coincide with the anti-resonance frequency of crystal A. Clearly X_A and X_B have opposite signs from the series frequency of A to the pole frequency of B, so that the bandwidth of the lattice extends to twice the pole-zero spacing of the crystals, and therefore more than twice the bandwidth of a corresponding ladder filter.

In the stopband region, the attenuation of the section depends on the ratio of X_A and X_B, going to infinity when the ratio is unity. In the far stopband, the crystal impedances essentially reduce to the impedances of their static capacitances, so that if these are made the same, then the attenuation in the stopband will tend monotonically to infinity, or in practice, to some large finite value determined by circuit strays and component tolerances. If on the other hand the static capacitances are chosen to have some ratio other than unity, attenuation peaks appear on either side of the passband, their precise location depending on the ratio of the static capacitances. Figure 10.8 shows typical examples of the image attenuation due to a single lattice section.

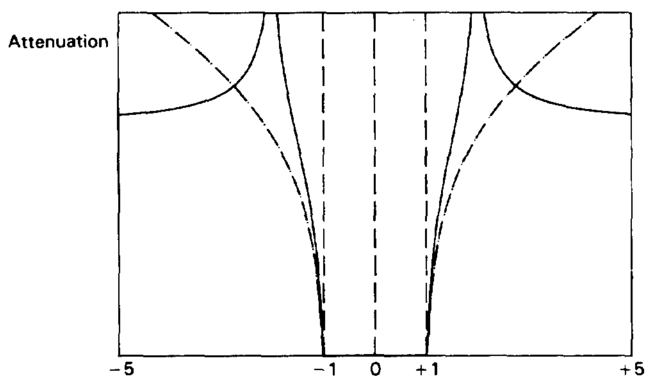

Fig. 10.8

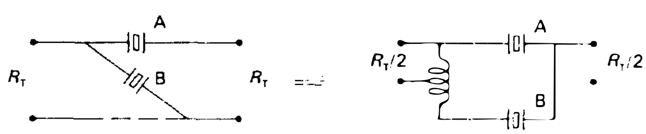

Fig. 10.9 'Half-lattice' transformation.

The full lattice of Fig. 10.5 contains four crystals in two identical pairs. Apart from the manufacturing difficulties in producing matched pairs of crystals, it would clearly be a great advantage if the need for identical crystals could be eliminated. This can in fact be done in several different ways, but the most commonly used is the 'half-lattice' transformation of Fig. 10.9,

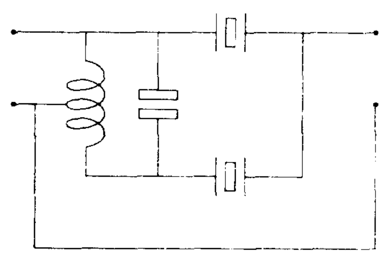

Fig. 10.10 'Half-lattice' with tuned transformer.

whereby the number of crystals is halved at the expense of the introduction of an ideal centre-tapped transformer. In practice, the transformer is realized as a tightly coupled tuned transformer, tuned to the centre frequency of the filter, as in Fig. 10.10. Because of the very narrow bandwidths involved, the effect of the tuned circuit on the filter response can be neglected, except insofar as the finite Q of the inductor introduces losses that add to the losses due to the motional resistances of the crystals.

The flexibility of the lattice circuit as compared to the ladder circuit is such that it is generally preferred for all but the narrowest bandwidth applications. The main disadvantage of the single lattice section is that it is difficult to achieve high levels of stopband attenuation because of the difficulty in maintaining the balance between the lattice arms over wide frequency ranges and also over wide operating temperature ranges. This is usually overcome by cascading a number of lattice sections.

The maximum bandwidth available from the two crystal lattice discussed above is twice the pole-zero spacing of the component crystals. For AT fundamentals it is therefore in percentage terms about 0.5%. To achieve wider bandwidths, it is common in *intermediate band* designs to introduce inductors as shown in Fig. 10.11. By the elementary network equivalences for the extraction of series and parallel elements from the arms of a lattice structure, these inductors are effectively in parallel with the crystal static capacitances. Consequently they can be used to reduce the effective values of the static capacitances and therefore increase the pole-zero spacings of the crystals and the maximum attainable bandwidth. The penalties to be paid for this increase in bandwidth are, firstly, the increased losses due to the finite

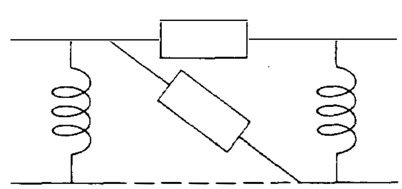

Fig. 10.11 Intermediate band lattice.

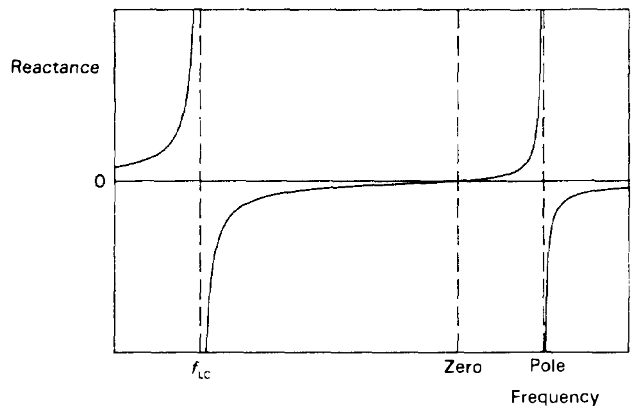

Fig. 10.12 Reactance plot for crystal with shunt inductor.

coil Q, and, secondly, the introduction of an additional pole in the frequency response of the lattice arms. This is shown in Fig. 10.12 as f_{LC} and provided the added inductance L is not too small occurs at frequencies considerably below the passband of the filter. The effect of the additional poles is to produce an unwanted passband, which can in some cases significantly affect the overall performance of the filter.

This technique by itself is only useful for achieving marginally larger bandwidths than can be achieved witout the additional inductors. It is characteristic of the two-crystal lattice as so far described that its image impedance increases from very low values at the lower edge of the passband to very high values at the upper edge. In a filter built as a cascade of crystal lattice sections with additional inductors to broaden the bandwidth, the loss resistance of the inductors appears across the signal path and has substantially more effect on the high impedance side of the passband than on the low impedance side. It is therefore typical of such filters to show a pronounced distortion of the passband, the attenuation increasing markedly from the lower to the upper passband edge.

To avoid this distortion, it is possible to choose the inductance L to resonate with the crystal static capacitances at the centre frequency of the filter. In such a case, each lattice arm impedance has a zero at the series resonance frequency of the crystal, together with two poles which are symmetrically placed either side of the zero, as shown in Fig. 10.13. The pole-zero spacing is given by

$$\Delta f/f = (C_1/C_0)^{1/2}/2 \tag{10.2}$$

Provided the required filter bandwidth is substantially less than this spacing, the poles can, to a first approximation, be ignored, so that in the neighbourhood of the centre frequency, the lattice arm impedances can be regarded as

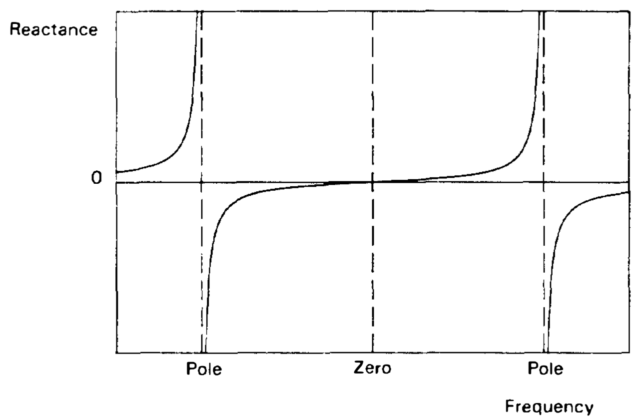

Fig. 10.13 Reactance plot for crystal with C_0 tuned out.

simple series LC circuits. Reference to Fig. 10.14 then shows that the two lattice arm reactances will have opposite signs in the frequency interval between their series resonances, so that the upper and lower edges of the filter passband will be defined by the crystal frequencies. Moreover it is clear that the image impedance $(Z_A Z_B)^{1/2}$ vanishes at both ends of the passband, giving rise to a *symmetric impedance* characteristic. Now when such lattices are cascaded, the effects of the inductor losses are spread more uniformly over the passband, leading to less relative distortion.

Although the poles of the lattice arm impedances will have little effect on the passband response if the bandwidth is sufficiently small compared to the pole-zero spacing, the poles will nevertheless affect the stopband response by

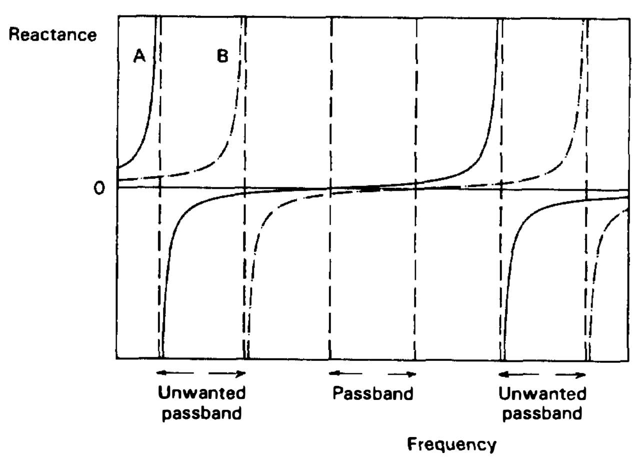

Fig. 10.14 Reactance plot for symmetric impedance lattice.

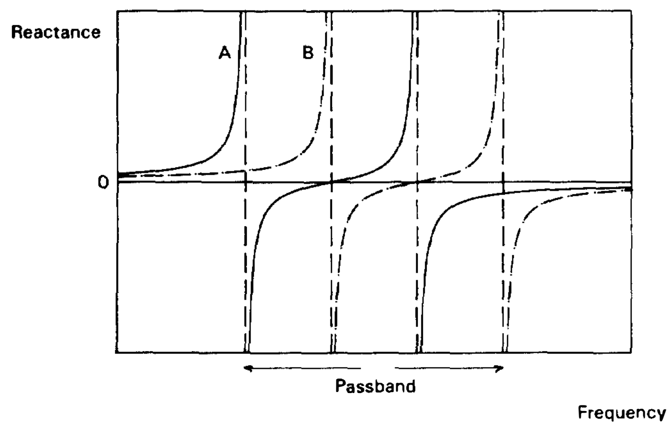

Fig. 10.15 Reactance plot for wideband lattice.

introducing unwanted passbands at the frequencies defined by Eqn (10.2). In *wideband* designs, this problem is avoided by making use of the poles to increase the bandwidth still further. This is illustrated in Fig. 10.15, from which it is clear that with the choice of pole and zero frequencies shown, the bandwidth extends from the lowest pole of one lattice impedance to the highest pole of the other lattice arm. The fractional bandwidth is then given by

$$\Delta f/f = 3(C_1/C_0)^{1/2}/2 \tag{10.3}$$

For AT fundamentals, this is approximately 10%, but it should be noted that despite the apparent attractiveness of the wideband approach, its application is restricted by the Q of available inductors to centre frequencies not much above 1 MHz.

In the lattice realizations so far discussed, the lattice arms have each contained just one crystal unit. In all cases, but particularly so in the intermediate and wideband structures, more selectivity can be obtained by utilizing the extra degrees of freedom made available by using two or more crystals in each lattice arm. This is achieved by choosing the crystal parameters in such a way that the lattice arm impedances have equal values at certain predetermined frequencies in the stopband; since the lattice is then balanced at these frequencies, attenuation peaks are thereby created.

10.3 ELEMENTARY CIRCUITS: MONOLITHIC CRYSTAL FILTERS

Since their introduction in the middle 1960s, monolithic crystal filters have become the most widely used form of crystal filter. Strictly speaking, the term

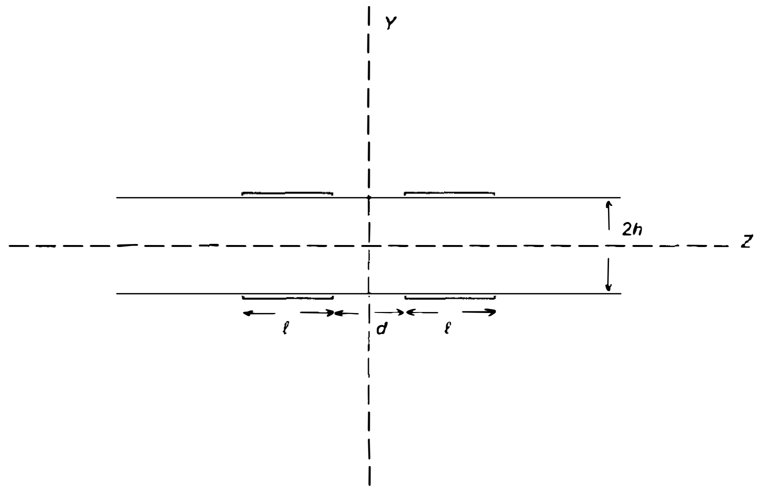

Fig. 10.16 Cross-section of monolithic dual.

monolithic refers to a filter fabricated on a single quartz plate, and incorporating two or more acoustically coupled resonators. Each resonator is formed by a pair of opposing electrodes on the major surfaces of the quartz blank, and the size, thickness, and separation of the pairs of electrodes determine the electrical response of the filter. In practice, because of manufacturing difficulties in producing sufficiently large plates, and more importantly because of the degradation of the stopband response caused by unwanted vibrational modes, it is usual to restrict the number of resonators on a plate to two only. The resulting device is termed a monolithic dual, and filters of higher degree than two are produced as a cascade of duals. The proper term for such a cascade is a 'polylithic' filter.

The operation of the monolithic dual can be explained in terms of the energy trapping theory discussed in Chapter 3. Figure 10.16 shows a cross section of a dual fabricated on a plate of thickness $2h$ in the y direction, with two pairs of parallel strip electrodes of length l along the z axis and infinite width in the x direction. The two pairs of electrodes are separated by a distance d. As in Chapter 3, thickness twist (TT) waves, propagating along the z direction, can be trapped under the electrodes as a result of the lowering of the cut-off frequency in the electroded region caused by the mass loading effect. Provided the spacing d is sufficiently large, each pair of electrodes can be regarded as constituting an independent trapped energy resonator, with there being no interaction or coupling between the two resonators. If, as shown in Fig. 10.16, the electrode lengths l are equal and if also the mass loadings are equal, then the TT mode frequencies of the two resonators will be equal. The spacing d required for there to be no interaction is determined by the rate of decay of the evanescent waves in the unelectroded regions,

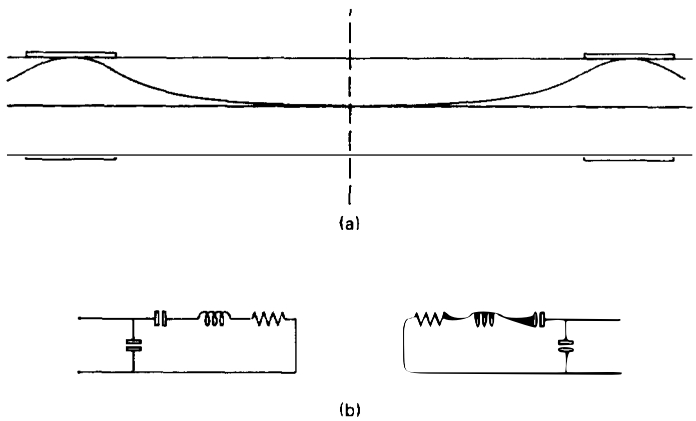

Fig. 10.17 Uncoupled resonators: (a) amplitude of vibration; (b) equivalent circuit.

specifically in the region between the electrodes. Figure 10.17(a) shows schematically the amplitude of vibration as a function of z when there is no coupling between the resonators. In this case each resonator can be represented by the conventional crystal equivalent circuit, as in Fig. 10.17(b).

As the spacing of the electrode pairs is reduced, it is clear that at some stage the evanescent wave associated with one resonator will extend into the region occupied by the other, and vice versa (Fig. 10.18a). Then by a mechanism entirely analogous to quantum mechanical 'tunnelling' through a potential barrier, energy can be transmitted from one resonator to the other, resulting in *acoustic coupling* between the resonators. In equivalent circuit terms, this coupling can be represented in various ways, one possibility being shown in Fig. 10.18(b), with a lattice equivalent in Fig. 10.18(c).

The equivalent circuit of the monolithic dual in this or alternative forms is already in one of the standard forms used in the design of narrowband coupled resonator filters. It has, in fact, precisely the form of a simple two-pole filter, with a bandwidth directly proportional to the coupling between the resonators. As already indicated, the coupling, and hence the bandwidth, is a function of the spacing d between the electrode pairs, decreasing exponentially with d. As the mass loading on the electrodes is increased, the difference in the cut-off frequencies between the electroded and the unelectroded regions increases. This results in a faster rate of decay of the evanescent waves, and consequently a decrease in the coupling and bandwidth. Thus in the manufacture of duals, both the electrode spacing and the mass loading can be adjusted to obtain the desired couplings.

A single dual has limited use as a filter, but it is straightforward to cascade duals to obtain more selective characteristics. The simplest case is that where the only additional components used are capacitors to ground at the

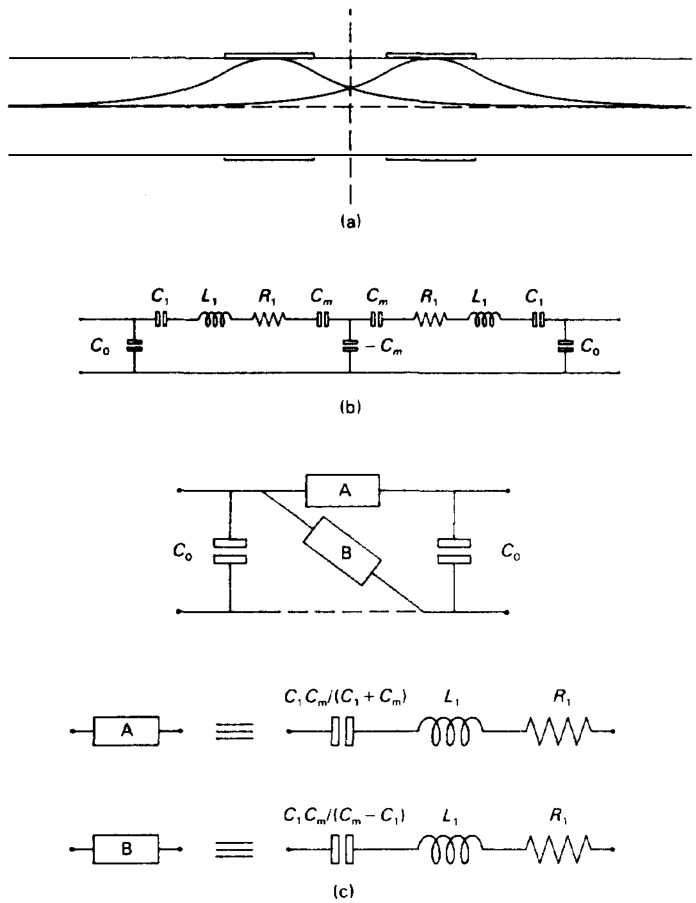

Fig. 10.18 Coupled resonators: (a) amplitude of vibration; (b) equivalent circuit; (c) lattice equivalent for monolithic dual.

junctions between adjacent duals, as in Fig. 10.19. This is the *narrowband* case, and just as for discrete crystal filters, there is an upper limit on the bandwidth achievable. For a given response type, the total capacitance required at each junction, including the static capacitances of the resonators, is inversely proportional to the bandwidth. Since the total capacitance cannot be reduced below the crystal shunt capacitances plus circuit strays without using inductors, this determines the upper limit. Just as in discrete crystal, *intermediate* band designs, the use of inductors allows wider bandwidths to be achieved, but the penalties of increased losses and some passband distortion follow. Since in practice one of the most attractive features of monolithic construction is precisely the ability to produce filters without bulky wound

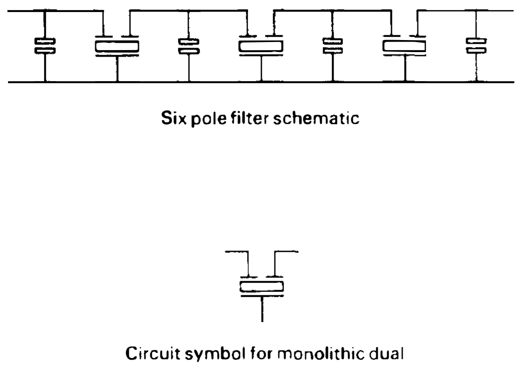

Fig. 10.19 Six-pole narrowband polylithic filter.

components, the vast majority of monolithic or polylithic filters produced are of the narrowband type. (But note that inductors are frequently used at input and output to provide impedance matching; in some cases these are included in the filter package, but in miniature types the user has to provide the necessary inductive terminating reactance.)

The typical polylithic filter, as in Fig. 10.19, is only capable of realizing 'all-pole' filter responses in which all attenuation poles occur at infinity. Such responses include the Butterworth and Chebyshev types and various linear-phase characteristics, but not for example the 'elliptic function' type of response. It can be shown that for given (equiripple) levels of passband and stopband attenuations, and for a given number of resonators, an elliptic function filter has the sharpest cut-off of all possible filters, that is the narrowest transition region between pass and stopbands. This is achieved by the optimum placing of attenuation poles at finite frequencies in the stopband. To realize filter characteristics that have such poles using monolithic

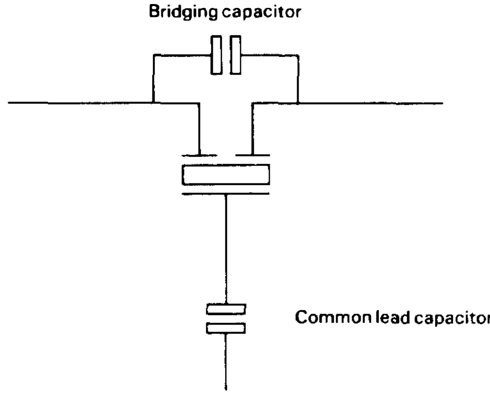

Fig. 10.20 Dual with bridging and common lead capacitors.

duals requires additional components, for example as shown in Fig. 10.20. This arrangement provides a pair of poles symmetrically disposed about the centre frequency, and can be cascaded with similar sections or standard 'all-pole' sections to give various responses. Further specialized characteristics can be achieved by adding discrete crystals to an otherwise polylithic structure, for instance to provide attenuation poles at single frequencies.

10.4 LIMITATIONS ON CRYSTAL FILTER PERFORMANCE

10.4.1 Bandwidth limitations

Crystal filters, whether discrete or monolithic, are often the only choice for narrowband applications in the range from one to several hundred megahertz. However, in specifying filters there are some fundamental limitations. It has already been pointed out that there is an upper limit on the bandwidth available, from up to 10% of the centre frequency for wideband, low frequency designs, down to fractions of a percent for designs using overtone mode crystals or duals. There are also lower limits on the bandwidth achievable.

Firstly, as the percentage bandwidth decreases, the resonator Q necessary to achieve acceptable insertion loss and passband distortion increases rapidly. For AT-cut crystals, the intrinsic material Q is approximately 10^6 at 10 MHz and is inversely proportional to frequency so that at 100 MHz it is only 10^5. Consequently the lower bandwidth limit imposed by the need for high Q is roughly proportional to frequency. Secondly, as the percentage bandwidth decreases, the problems of temperature stability and manufacturing tolerances on the resonator frequencies become more critical. If, for example, the operating temperature range of the filter is such that a frequency–temperature tolerance of ± 10 ppm on the resonator frequencies is the best that can be economically achieved, and if similarly a tolerance of ± 10 ppm is the best that can be achieved in manufacturing, then these tolerances have to be allowed for in the filter design.

For purposes of illustration, suppose the specification requirement, to be met over the operating temperature range, is for a 3 dB bandwidth of ± 2.5 kHz and a 60 dB bandwidth of 12.5 kHz at a centre frequency of 50 MHz. Since the resonator frequencies can be expected to vary a total of ± 20 ppm or ± 1 kHz due to the combined causes of temperature changes and unit to unit variations, it is necessary to allow margins of 1 kHz on either side of the passband, thus making the design bandwidth ± 3.5 kHz rather than ± 2.5 kHz, a 40% increase over the nominal. Clearly, it is not possible to decrease the *design* bandwidth indefinitely and still provide for the necessary tolerances.

The need to allow margins also has a major impact on the apparent complexity of narrowband filters. The nominal shape factor in the example above is 5:1 from 60 to 3 dB, which, as reference to any of the standard filter texts will show (for example, Zverev, 1967), is easily obtainable from a four pole Chebyshev filter with 0.5 dB ripple. However on taking into account the necessary margins on the passband, and also similar margins on the stopband, the shape factor reduces to 11.5/3.5 or approximately 3.3:1. This is *not* achievable with a four pole 0.5 dB Chebyshev filter, making either a five pole Chebyshev design or else a design with finite attenuation poles necessary. In either case, the resonator Q requirement will be more severe than in the simple four pole Chebyshev case.

10.4.2 Unwanted responses

A further major limitation on crystal filter performance is the presence of *unwanted responses* in the resonators, whether discrete or monolithic. The most troublesome responses are generally the *inharmonic* modes discussed in Chapter 3. These form a group of resonances located at frequencies just above the main or 'wanted' response. In energy-trapped designs, proper dimensioning of the electrodes and choice of mass loading can produce resonators relatively free from these unwanted or 'spurious' modes, but especially at higher frequencies and overtones the energy-trapping criteria are difficult to satisfy while still maintaining acceptable values of resonator motional inductance and resistance. Typically the mass loading cannot be reduced below a limit set by the need to maintain an electrode film thick enough to retain good electrical conductivity, and then the electrode diameters dictated by energy trapping theory turn out to be very small. As a consequence, the motional inductance, which is inversely proportional to the electrode area, is high, in conflict with the filter designer's usual need for as low an inductance as practically possible.

The presence of unwanted modes can affect the filter response in two ways. In relatively narrowband designs, the unwanted modes will typically all fall into the stopband of the filter, appearing as sharp 'spikes' that degrade the ultimate attenuation of the filter in narrow frequency intervals on the high frequency side. This is probably the least troublesome manifestation of 'spurious' modes. As the bandwidth increases, the filter impedance will generally increase, and the unwanted modes will both increase in level, and appear relatively closer to the centre frequency, until they actually occur in the transition region. Further increase in bandwidth finally leads to a situation where the spurious responses begin to appear in the passband of the filter, showing up as sharp notches superimposed on an otherwise smooth response. The occurrence of spurious modes in the passband can very often limit the maximum practical bandwidth achievable to values considerably less than those implied by the overall design limitations discussed in Section 10.2.

10.4.3 Group delay characteristics

With the rapid increase in data transmission over radio communications networks, the group delay performance of crystal filters is becoming increasingly important. As mentioned previously in Section 10.1, for the class of filters known as *minimum-phase* filters, the delay and amplitude characteristics are not independent. For such filters, high selectivity implies rapidly changing group delay at the passband edges, whereas, flat delay characteristics imply poor selectivity. Nearly all the crystal filter realizations currently used fall into this class, and so cannot satisfy independent specifications on delay and amplitude. This situation is in contrast to that existing for SAW filters, where generally the phase and amplitude characteristics can be independently controlled. However, crystal filters have been developed using unconventional realization techniques which can satisfy simultaneous phase and amplitude specifications and these show promise for the future.

10.4.4 Non-linear effects

Crystal filters containing such known nonlinear components as ferrite cores can be expected to show the same effects as other devices using the same components. Hence ferrites and other nonlinear materials should be avoided in the manufacture of filters where intermodulation or other amplitude dependent phenomena are to be avoided. Assuming this to be the case, the residual non-linearities in crystal filters can generally be assigned to the individual resonators. As such, a fuller discussion is given in Chapter 7, but it can be pointed out here that in the case of intermodulation, there appear to be two separate mechanisms acting according to whether the test tones applied are in the passband or in the stopband.

To be specific, consider two test tones at frequencies $f_0 + df$ and $f_0 + 2df$, with f_0 being the centre frequency of the filter. Third-order intermodulation then gives rise to an intermodulation product (IMP) at the centre frequency f_0. If the frequency offset df is such that both test signals are in the filter stopband, then the resonator currents can be expected to be small, since the crystals are being driven at frequencies far from resonance. On the other hand, if df is such that both signals fall within the filter passband, then the resonator currents can be expected to be much larger as the test frequencies will then be closer to the crystal resonance frequencies.

In the latter case, it appears that the intermodulation products are the result of intrinsic nonlinearities in the resonator material, specifically in the elastic stress-strain relationship. This case has been treated theoretically by Tiersten (1974, 1975), with the result that the intermodulation ratio, that is the ratio of the input power in the test tones to the power delivered to the load at the intermodulation frequency, turns out to be proportional to the square of the overtone order and to the square of the electrode area (Smythe, 1974). That is for the in-band case, filters using overtone resonators are generally to be preferred to those using fundamental mode resonators.

In the case where the test tones are in the stopband, the intermodulation mechanism has been found to be largely process dependent. The key factors appear to be the same as those discussed in Chapter 7 in relation to the increase in crystal resistance and the change in resonator frequency sometimes observed at low drive levels, that is surface finish and cleanliness and the quality of the electrode films. Consequently, this type of intermodulation is difficult to quantify, and may well not depend on the test tone power levels in the expected way, often being relatively worse than expected at low power levels and better than expected at higher power levels.

Other non-linear effects that may be observed in crystal filters at low input power levels can also be associated with the process factors mentioned in the previous paragraph. At very high power levels other effects can be observed that result from shifts in resonator frequency and resistance due to the generation of thermal gradients and stresses in the crystal blanks, but such power levels should be avoided.