British Medical Bulletin (1987) Vol. 43, No 4, pp. 887-908

Psychophysics of normal and impaired hearing

Brian C J Moore
Department of Experimental Psychology, University of Cambridge, Cambridge

A central theme in the perception of sound by normally hearing and hearing-impaired listeners is the frequency-analysing capacity of the auditory system. This capacity plays a role in our ability to detect signals in masking noise, to identify the timbre of speech and musical sounds, and to perceive the pitch of complex tones. The basic properties of these abilities can be understood by conceiving of the peripheral auditory system as containing a bank of bandpass filters (the auditory filters), whose centre frequencies cover the whole audible range. The filters at different centre frequencies are characterised by the value of the auditory filter bandwidth or critical bandwidth (CB). The basic frequency-analysing mechanisms seem to be well established at the level of the cochlear nerve. Information about stimulus frequency, intensity, and spectrum may be carried both in the distribution of activity across nerve fibres and in the temporal patterns of neural firing. Temporal patterns may be particularly important in the perception of pitch. Damage to the inner ear results in an impairment in the frequency-analysing mechanisms. Thus the ability to detect and discriminate signals in noise, to identify the timbre of sounds, and to perceive the pitch of complex sounds, may all be impaired. Temporal resolution and sound localisation ability may also be reduced in cases of cochlear hearing loss, and loudness perception may be abnormal. These disabilities are not corrected with a conventional hearing aid.

The commonest way of assessing hearing in the clinic is to measure the lowest intensity at which pure tones (sinusoids) can be detected, as a function of the frequency of the tones, giving what is called the pure tone audiogram. The measurement is obtained separately for each ear using headphones. However, in everyday life people are rarely required to listen to pure tones at very low intensities through one ear only. Rather, the function of the auditory system is to analyse and discriminate sounds which typically contain many sinusoidal components, whose intensity may vary over a very wide range, and which may come from many different directions in space. This paper is concerned primarily with the perception and discrimination of sounds at intensities typical of those encountered in everyday life, both for those with normal hearing, and for the hearing impaired.

THE ACTION OF THE EAR AS A FREQUENCY ANALYSER

It is a central characteristic of the auditory system that it acts as a limited-resolution frequency analyser; complex sounds are broken down into their sinusoidal frequency components. The initial basis of this frequency analysis almost certainly depends upon the tuning which is observed in the cochlea (see Wilson & Russell, this issue). Indeed, it is possible that the tuning observed in the cochlea is sufficient to account for the frequency-analysing capacity of the entire auditory system.[1,2] Largely as a consequence of this analysis, people can hear one sound in the presence of another sound with a different frequency. This ability is known as *frequency selectivity* or *frequency resolution*.

Measurement of the ear's frequency selectivity

Important sounds are sometimes rendered inaudible by other sounds, a process known as 'masking'. Masking may be considered as a failure of frequency selectivity, and it can be used as a tool to measure the frequency selectivity of the ear. One conception of masking, which has had both theoretical and practical success, assumes that the auditory system contains a bank of overlapping band-pass filters, with adjacent, ordered centre frequencies.[3] In the simple case of a sinusoidal signal presented in a background noise, it is assumed that the observer will make use of the filter whose output has the highest signal-to-masker ratio. The signal will be detected if that ratio exceeds a certain value. In most practical situations the filter determining listening performance will have a centre frequency close to that of the signal.

The characteristics of the auditory filter have been investigated by many people. Fletcher[3] was one of the first to apply the auditory-filter concept to the masking of sinusoidal tones by broad-band noise. He assumed that to predict threshold it would be reasonable to approximate the auditory filter's response with a simple rectangle having a flat top and vertical edges. Thus all frequency components falling within the flat top or passband would be passed equally, whereas components outside the passband would be rejected. He called the width of this passband the *critical bandwidth* (CB).

Although this approximation of the auditory filter as a simple rectangle works quite well for signals in broad-band noise, it does not work well for other types of maskers. One demonstration of this is the *psychophysical tuning curve* (PTC)—an analogue of the method used by neurophysiologists in determining a neural tuning curve (see Palmer, this issue). The signal used is a sinusoid which is presented at a very low level, say 10 dB above the absolute threshold. It is assumed that this will excite only a small number of nerve fibres with *characteristic frequencies* (CFs) close to that of the signal. Thus, to a first approximation, only one auditory filter will be involved in detecting the signal. The masker is either a sinusoid or a narrow band of noise.

To determine a PTC the signal is fixed in frequency and level, and the level of the masker required to mask the signal is determined for various centre frequencies of the masker. If it is assumed that the signal will be masked when the masker produces a fixed amount of activity in the neurones which would otherwise respond to the signal, then the curve mapped out in this way is analogous to the neural tuning curve.[4] Some examples are given in Figure 1. Returning to the concept of the auditory filter, we can think of the points on a PTC as representing the family of masker levels required to produce a fixed output from the filter centred at the signal frequency. Normally we determine a filter characteristic by plotting the output as a function of frequency for an input fixed in level. However, if the filter is linear the two methods are equivalent. Thus the filter characteristics can be obtained simply by turning the tuning curve upside-down.

An alternative method of determining the auditory filter shape has been described by Patterson.[6] The method is illustrated in

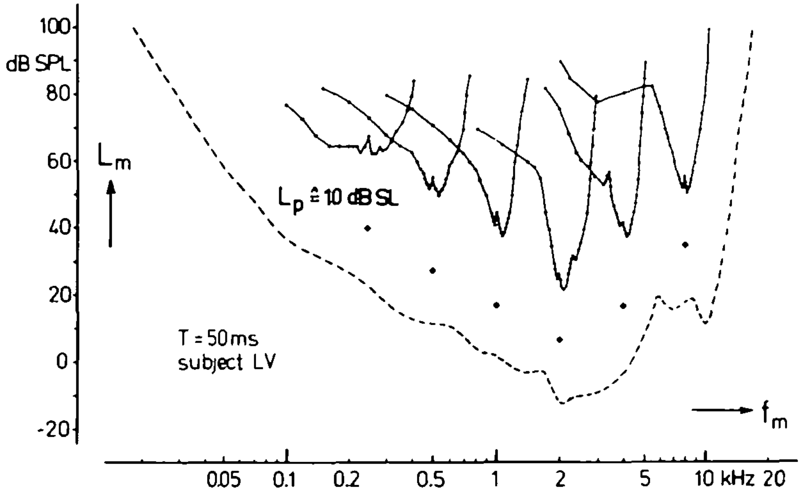

Fig. 1 Psychophysical tuning curves (PTCs) determined in simultaneous masking using sinusoidal signals 10 dB above absolute threshold (called 10 dB SL). For each curve the solid diamond below it indicates the frequency and level of the signal. The masker was a sinusoid which had a fixed starting phase relationship to the 50-ms signal. The masker level required for threshold is plotted as a function of masker frequency. The dashed line shows the absolute threshold for the signal. From Vogten,[5] by courtesy of the author.

Figure 2. The signal is fixed in frequency, and the masker is a noise with a bandstop or notch centred at the signal frequency. The deviation of each edge of the notch from the signal frequency is denoted by Δf. The threshold of the signal is determined as a function of notch width. Usually the notch is symmetrically placed around the signal frequency and the analysis assumes that the auditory filter is symmetric on a linear frequency scale. This assumption is not unreasonable, at least for the top part of the filter and at moderate sound levels.[7]

As the width of the notch is increased, less and less noise will pass through the auditory filter; thus, the threshold of the signal will drop, i.e. improve. The amount of noise passing through the auditory filter will be proportional to the area under the filter in the frequency range covered by the noise. This is shown as the shaded areas in Figure 2. Given our assumption that threshold corresponds to a constant signal-to-masker ratio at the output of the auditory filter, then the change in threshold with notch width tells us how the area under the filter varies with Δf. The area under a function between certain limits is obtained by integrating the

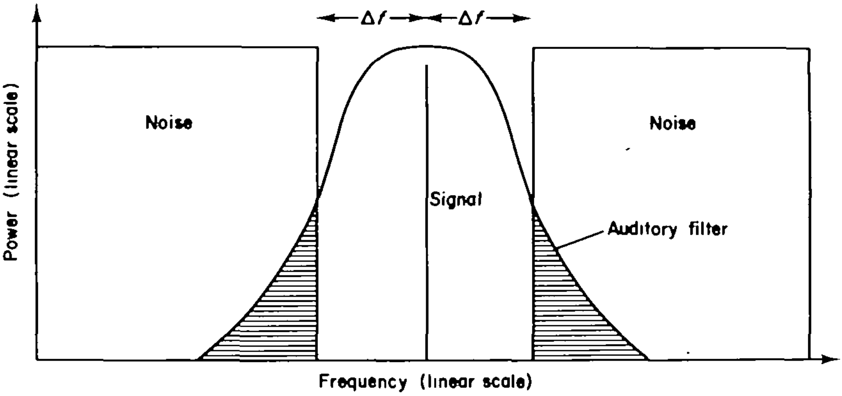

Fig. 2 Schematic illustration of the method used by Patterson[6] to determine auditory filter shape. The threshold of the sinusoidal signal is measured as a function of the width of a spectral notch in the noise masker. The amount of noise passing through the auditory filter centred at the signal frequency is proportional to the shaded areas. From Moore.[24]

value of the function over those limits. Hence by differentiating the function relating threshold to Δf, the shape of the filter is obtained. In other words, the attenuation of the filter at given deviation, Δf, from the centre frequency is proportional to the slope of the function relating signal threshold to notch width at the value of Δf. The method can also be extended to the case where the filter is not assumed to be symmetric, provided certain assumptions are made about the general form of the filter.[7]

The auditory filter shape obtained using Patterson's method varies somewhat depending on the level of the masker. Figure 3 shows the shape of the filter for various noise levels ranging from 20 to 90 dB SPL. The low-frequency side of the filter becomes less steep with increasing level, while the high-frequency side becomes somewhat steeper. The bandwidth of the auditory filter increases with increasing centre frequency. However, when expressed as a proportion of centre frequency, the bandwidth tends to be narrowest at middle to high frequencies. Over the range 100 to 6500 Hz, and at moderate sound levels, the *bandwidth* (BW) is well approximated by the expression

$$BW = 6.23F^2 + 93.39F + 28.52,$$

where F is frequency in kHz.[8]

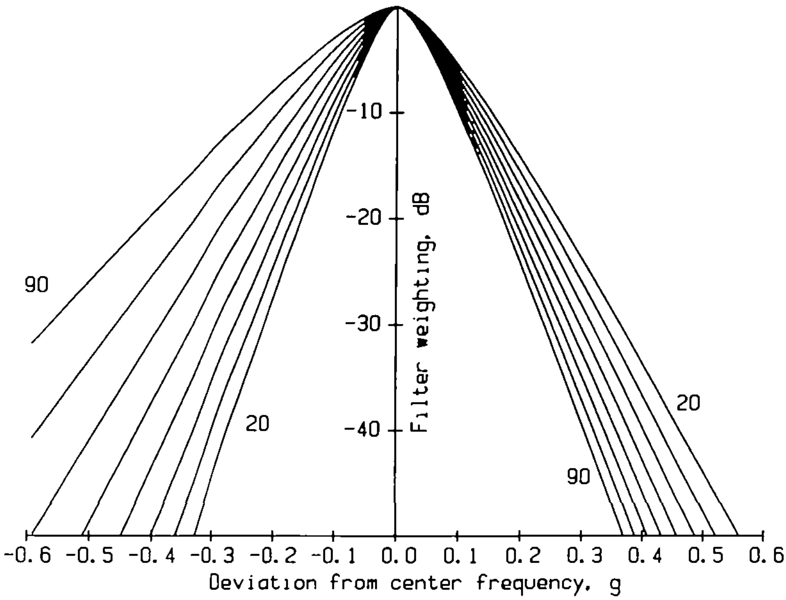

Fig. 3 Auditory filter shapes as a function of masker level (20–90 dB SPL) for a centre frequency of 1 kHz. The filter shapes are plotted as a function of the deviation from the centre frequency divided by the centre frequency. With increasing sound level the upper skirt of the filter becomes slightly steeper and the lower skirt becomes considerably less steep. From Moore.[1]

Frequency selectivity in the hearing impaired

There is now considerable evidence that in listeners with hearing impairments of cochlear origin there is a loss of frequency selectivity. In general, greater threshold elevations tend to be associated with broader auditory filters. However, the following cautions should be observed:

1. There can be considerable variability among patients, even when the elevation in absolute threshold is similar. Although the bandwidth of the auditory filter is correlated with the threshold elevation,[9,10] some patients have broad filters and almost normal thresholds, while some have elevated thresholds but almost normal filters.

2. The auditory filter becomes broader at high sound levels even in normal listeners (see Fig. 3). Since measurements with patients usually have to be made at high sound levels, part of the broadening may be attributed to the same level effect as occurs in normal hearing.

One way around these problems is to use subjects with unilateral impairments. Each subject then acts as his/her own age-matched control. Figure 4 shows a comparison of auditory filter shapes obtained separately from each ear of six patients with unilateral cochlear hearing losses (data from Glasberg & Moore[10]). The upper panels show filter shapes for the normal ears and the lower panels show filter shapes for the impaired ears, which had threshold elevations at the test frequency (1 kHz) ranging from about 40 to 60 dB. Losses were relatively flat as a function of frequency. A notched-noise masker was used, as described earlier, and the same noise spectrum level (50 dB) was used for testing all ears, so the results are not subject to the difficulty discussed in (2) above.

It is clear that the auditory filters are considerably broader in the impaired ears. The most obvious feature is that the lower skirts of the filters are consistently and considerably less sharp in the impaired ears. This implies a high susceptibility to the upward spread of masking from low frequencies. This appears to be a common feature in cases of cochlear impairment, even in patients with relatively flat losses as a function of frequency (see Stephens, this issue). It may partially account for the fact that hearing aids are often most effective when their gain is greater at high frequencies than at low. A rising frequency-gain characteristic will help to alleviate the effects of the upward spread of masking.

Effect of impaired frequency selectivity on the masking of speech by noise

A consequence of reduced frequency selectivity is a greater susceptibility to masking by interfering sounds. The frequency content of a sound being listened to will usually differ from that of other sounds in the environment. By making use of auditory filters tuned close to the signal frequency or frequencies, the signal will be passed but much of the background noise will be attenuated. When the auditory filters are broader than normal, the rejection of background noise will be much less effective. Thus background noise will severely disrupt the detection and discrimination of sounds, including speech. Indeed, difficulty in understanding speech in noise is one of the commonest complaints among people with hearing impairments of cochlear origin (see Summerfield, this issue). The ability to understand speech in noise decreases as frequency selectivity decreases.[11-13]

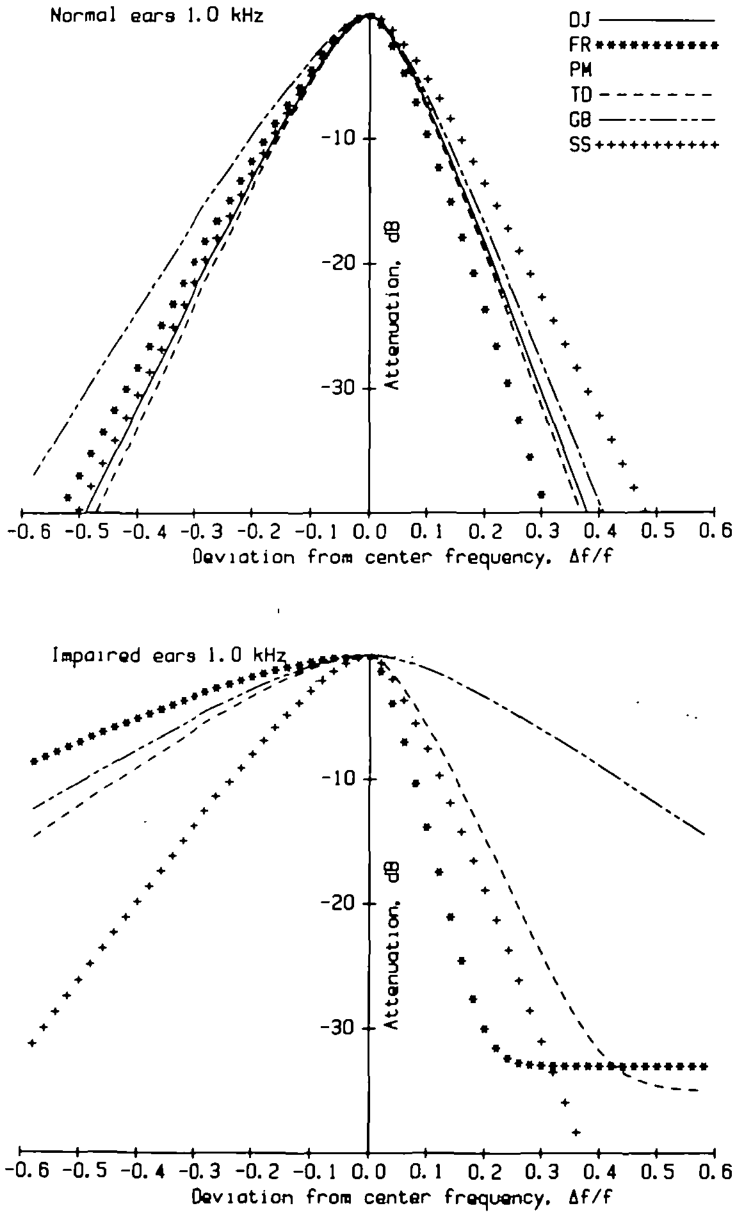

Fig. 4 Auditory filter shapes for the normal ears (top) and the impaired ears (bottom) of six subjects with unilateral cochlear impairments. The filter shapes are plotted as a function of deviation from the centre frequency divided by the centre frequency. The impaired ear of subject DJ had too little frequency selectivity for a filter shape to be determined. From Glasberg & Moore.[10]

THE PERCEPTION OF LOUDNESS

Equal-loudness contours

In describing the perception of sound it is useful to have some kind of scale which allows one to compare the loudness of different sounds. A first step towards this is to construct equal-loudness contours for sinusoids of different frequencies. Say, for example, we take a standard tone of 1 kHz, at a level of 40 dB SPL, and ask the listener to adjust the level of a second tone (say, 2 kHz) so that it sounds equally loud. If we repeat this for many different frequencies of the second tone, then the sound level required (plotted as a function of frequency) maps out an equal-loudness contour. The level of the 1-kHz standard sound defines the loudness level, in phons. If we repeat this procedure, for different levels of the 1-kHz standard tone, then we will map out a family of equal-loudness contours. Such a family is shown in Figure 5.

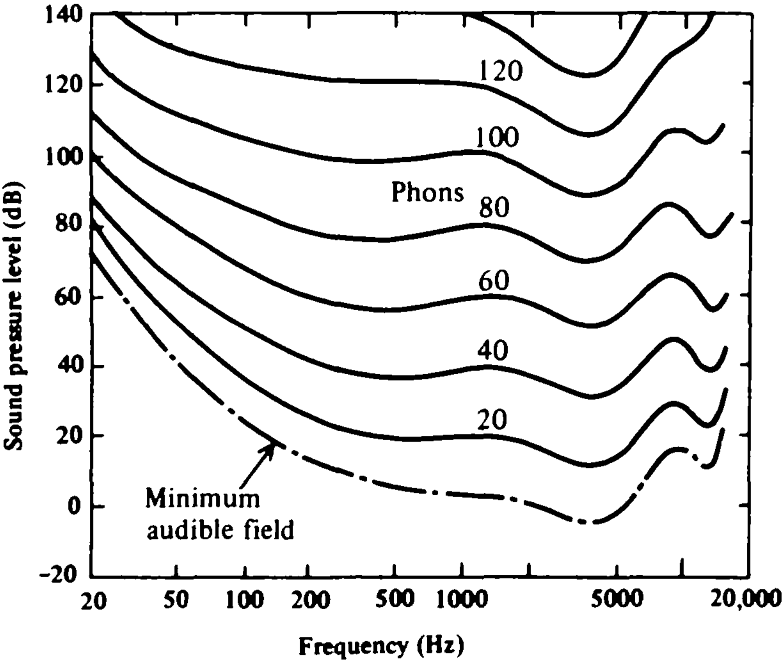

Fig. 5 Equal-loudness contours for various loudness levels, as indicated on each curve. The dashed-dotted curve shows the absolute threshold (minimum audible field)—from Moore.[24] *Original data from:* Robinson DW, Dadson RS. A redetermination of the equal-loudness relations for pure tones. Br J Appl Physics 1956; 7: 166–181.

Note that the contours resemble the absolute threshold curve (lowest curve in the figure) at low levels, but tend to become flatter at high levels. This finding has been utilised in the design of sound-level meters, which weight the power at different frequencies according to the shapes of equal-loudness contours. At low sound levels, low-frequency components contribute little to the total loudness of complex sound, and so as 'A' weighting (approximating the equal loudness contours between 30 and 60 phons) is used, which reduces the contribution of low frequencies to the overall meter reading. At high levels (above 90 phons), where the equal-loudness contours are flatter, a more nearly flat weighting characteristic, the 'C' weighting, is used. A 'B' weighting is used for intermediate levels (60–90 phons).

The role of frequency selectivity in determining loudness

It has been known for many years that if the total intensity of a complex sound is fixed, its loudness will depend on the frequency range over which the sound extends. The basic mechanism underlying this seems to be the same critical band or auditory filter as is revealed in masking experiments. Consider as an example a noise whose total intensity is held constant while the bandwidth is varied. The loudness of the noise can be estimated indirectly by asking the listener to adjust the intensity of a second sound, with a fixed bandwidth, so that it sounds equally loud. The two sounds are presented successively. When the bandwidth of the noise is less than a certain value the loudness is roughly independent of bandwidth. However, as the bandwidth is increased beyond a certain point, the loudness starts to increase. This is illustrated in Figure 6, for several different overall levels of the noise. The bandwidth at which loudness starts to increase is known as the *critical bandwidth for loudness summation.* Its value is approximately the same as the bandwidth of the auditory filter. A model which explains this effect is described by Moore & Glasberg.[14]

Loudness perception and recruitment in impaired ears

Cochlear hearing loss is usually associated with an abnormality of loudness perception known as loudness recruitment. Although the absolute threshold may be elevated, the rate of growth of loudness with intensity is more rapid than normal, so that at high intensities a sound will appear as loud in an impaired ear as it would in a

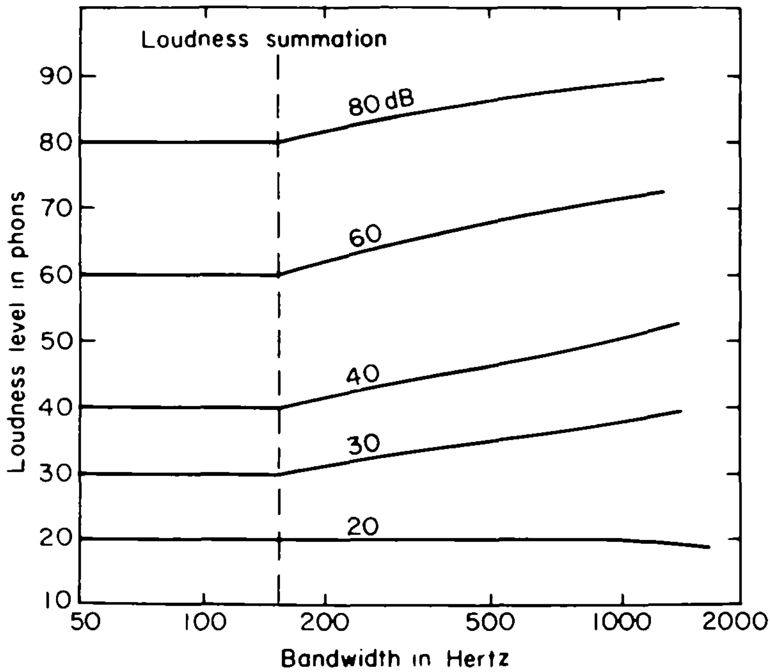

Fig. 6 The loudness level of a band of noise centred at 1 kHz, measured as a function of the width of the band. For each of the curves the overall level was held constant and is indicated in the figure. The dashed line shows that the bandwidth at which loudness begins to increase is the same at all levels tested (except that no increase occurs at the lowest level). *Adapted from:* Feldtkeller R, Zwicker E. Das Ohr als Nachrichtenempfänger. Stuttgart: S Hirzel (1956) by courtesy of the authors and publisher.

normal ear. The effect is most easily demonstrated when only one ear is affected, since then loudness matches can be made between the two ears, but it can be measured in other ways. The presence of recruitment can limit the usefulness of conventional hearing aids, since if the gain of the aid is set so as to make sounds of low intensity clearly audible, sounds of high intensity will be uncomfortably loud. Hearing aids incorporating 'compression', especially multi-band compression, can be useful in alleviating this effect.[15,16]

Evans[17] suggested that recruitment may be a consequence of the impaired frequency selectivity which is commonly associated with it. If tuning curves (or auditory filters) are broader than normal, then as the intensity of a tone is increased above threshold the activity will spread across the nerve fibre array (or to adjacent auditory filters) more rapidly than it would in the normal ear. This rapid spread of neural activity could account for the rapid growth of loudness with intensity. Moore et al.[18] tested this idea by measuring how the rate of growth of loudness of a tone in recruiting ears was affected by presenting that tone in a band-stop noise. Such a noise should mask neural activity evoked by the tone at CFs remote from the tone frequency. Thus, if Evans' suggestion were correct, the noise should substantially reduce the rate of growth of loudness of the tone. They found that the noise had little effect on the loudness of the tone, and concluded that an abnormally rapid spread of activity across the nerve fibre array is not the main cause of recruitment. The underlying cause of recruitment remains uncertain, but the input/output characteristics of individual nerve fibres may be involved (see Palmer, this issue).

THE PERCEPTION OF PITCH

The pitch of pure tones

Pitch is defined as that attribute of auditory sensation in terms of which sounds may be ordered on a musical scale, i.e. that attribute in which variations constitute melody. For sinusoids (pure tones) the pitch is largely determined by the frequency; the higher the frequency the higher the pitch. One of the classic debates in hearing theory is concerned with the mechanisms underlying the perception of pitch. One theory, called the place theory, suggests that pitch is related to the distribution of activity across nerve fibres. A tone with a given frequency will produce maximum activity in nerve fibres with CFs close to that frequency, and the 'position' of this maximum is assumed to determine pitch. Shifts in frequency will be detected as changes in the amount of activity at the place where the activity changes most.

The alternative theory, which we will call the temporal theory, suggests that pitch is determined by the time-pattern of neural spikes. For frequencies up to about 5 kHz, neural spikes are phase-locked to the stimulus, so that the time-intervals between successive spikes carry information about the stimulus frequency (see Palmer, this issue).

One major fact which these theories have to account for is humans' remarkably fine acuity in detecting frequency changes. This ability is called frequency discrimination, and is not to be confused with frequency selectivity. For two tones of 500 ms duration presented successively, a difference of about 3 Hz (or less in trained subjects) can be detected at a centre frequency of 1 kHz. It has been suggested that tuning curves (or auditory filters) are not sufficiently sharp to account for this fine acuity in terms of the place theory.[14] A further difficulty for the place theory is that frequency discrimination worsens abruptly above 4–5 kHz.[19] Neither neural measures of frequency selectivity (such as tuning curves) nor psychophysical measures of frequency selectivity (such as PTCs or auditory filter shapes) show any abrupt change there. These facts can be accommodated by the temporal theory. Changes in frequency discrimination with centre frequency (and with tone duration) can be predicted from the information available in inter-spike intervals.[20] The worse performance above 4–5 kHz corresponds well with the frequency at which the temporal information ceases to be available.

Studies of the perception of musical intervals also indicate a change in mechanism around 4–5 kHz. Below this, a sequence of pure tones with appropriate frequencies conveys a clear sense of melody. Above this, the sense of musical interval and of melody is lost, although the changes in frequency may still be heard. Therefore the evidence supports the idea that, for pure tones, pitch perception and discrimination are determined primarily by temporal information for frequencies below 4–5 kHz, and by place information for frequencies above this. The important frequencies for the perception of music and speech lie in the frequency range where temporal information is available.

The pitch perception of complex tones

In general, any complex sound which is periodic will have a pitch, provided that the waveform repetition rate lies in the range 20–16 000 Hz. The pitch is related to the repetition rate, in the same way that it is related to frequency for pure tones. The pitch assigned to a complex tone is generally taken to be the frequency of a sinusoid which has the same pitch, but unlike the position with sinusoids, the physical properties determining the pitch of complex sounds are not straightforward. Theories of the pitch perception of complex tones have evolved considerably in recent years.

Periodic complex tones can be analysed into a series of sinusoidal components; a fundamental component whose frequency equals the repetition rate of the sound, and a series of harmonics whose frequencies are integral multiples of that of the fundamental. Although the pitch heard is usually the same as that of the fundamental component, the fundamental does not have to be present to generate this pitch; this is called the 'phenomenon of the missing fundamental'.

Most modern theories of pitch perception assume a two-stage process. In the first stage the frequencies of the lower harmonics are determined. This analysis depends on the ear's filtering mechanism, but the time structure of the output from each filter, as represented in the temporal patterns of neural discharge at each CF, is probably also important.[21] In the second stage some form of pattern recogniser determines a fundamental frequency whose harmonics match those of the stimulus as closely as possible.[22,23] The perceived pitch corresponds to the frequency of this internally determined fundamental. Notice that for these theories the resolution of individual harmonics is critical; hence frequency selectivity should play a key role in pitch perception. For a more comprehensive review of recent data and theories on pitch perception the reader is referred to Moore & Glasberg.[14]

Pitch perception in impaired hearing

The data on pitch perception in the hearing impaired are sparse, and the results have been quite variable. Most people with cochlear hearing losses do have impaired frequency discrimination, but a few have almost normal discrimination. This applies to both pure and complex tones. When frequency selectivity is impaired, quite good discrimination might still be possible on the basis of temporal information alone. Only when temporal processing is also disrupted will performance become very poor. Thus some subjects can show very broad PTCs, and almost normal pitch discrimination for complex tones.[25] For a review of this topic see Rosen & Fourcin.[26]

THE PERCEPTION OF TIMBRE

Timbre may be defined as the characteristic quality of sound that distinguishes one voice or musical instrument from another. Timbre depends on several different physical properties of sound, including:

1. Whether the sound is periodic (giving a tonal quality for repetition rates from about 20 to 16 000 Hz), or irregular (giving a noise-like quality).
2. Whether the sound is continuous or interrupted. For short

sounds the exact way in which the sound is turned on and off can play an important role. For example, in the case of sounds produced by stringed instruments, a rapid onset (a fast rise time) is usually perceived as a struck or plucked string, whereas a gradual onset is heard as a bowed string.

3. The distribution of energy over frequency (i.e. the spectrum), and changes in the spectrum with time. This is the correlate of timbre which has been studied most widely.

For steady-state periodic sounds it is possible to use the more restricted definition of timbre given by the American Standards Association: 'that attribute of auditory sensation in terms of which a listener can judge that two steady-state complex tones having the same loudness and pitch, are dissimilar'. Timbre defined in this way depends primarily on the distribution of energy over frequency.[27] For example, sounds containing predominantly high frequencies have a 'sharp' timbre, whereas those containing mainly low frequencies sound 'dull' or 'mellow'. This is another example of the action of the ear as a frequency analyser. The components in a complex sound will be partially separated by the auditory filters, and the distribution of activity at the output of the filters, as a function of filter centre frequency, will determine timbre.

Timbre perception in impaired hearing

As a consequence of the reduced frequency selectivity which is associated with cochlear hearing loss, the ability to hear changes in timbre is also impaired. Thus it will be more difficult for the impaired listener to tell the difference between different vowel sounds or to distinguish musical instruments. For a review see Rosen & Fourcin.[26]

THE TEMPORAL RESOLUTION OF THE EAR

The auditory system is particularly well adapted to detecting changes in sounds as a function of time. The limits of this ability reflect the temporal resolution of the ear. One measure of this requires the subject to detect a brief gap in a relatively long duration sound. Many gap-detection experiments have used wideband noise as a stimulus, since introducing a temporal gap in such a noise does not change its spectrum. The results generally agree quite well, the threshold value being 2 to 3 ms.[28,29] More recently, gap thresholds have been measured for band-limited noises, to determine whether gap threshold varies with centre frequency. Unfortunately, when a noise band is abruptly switched off and on, to produce the gap, a change in spectrum occurs. Energy is spread or 'splattered' to frequencies outside the nominal bandwidth of the noise. In order to prevent the detection of this 'spectral splatter', which could give rise to artificial low thresholds, the noise bands have been presented with complementary band-reject noise.[30–32] Some results from Shailer & Moore[32] are plotted in Figure 7.

The value of the gap threshold increases monotonically with decreasing centre frequency. At high frequencies the gap threshold is similar to that found for wideband noise, suggesting that subjects make use primarily of high frequencies when detecting gaps in broadband noise. The increase in gap threshold at low frequencies may be connected with the temporal response of the auditory filter. When the input to a narrowband filter ceases abruptly, the filter continues to 'ring' for some time; everyday examples are the sound produced by a wine glass or a tuning fork when they are tapped. Ringing in the auditory filters could partially fill in a brief gap in a signal, thus limiting gap-detection performance. In general, the narrower the bandwidth of a filter,

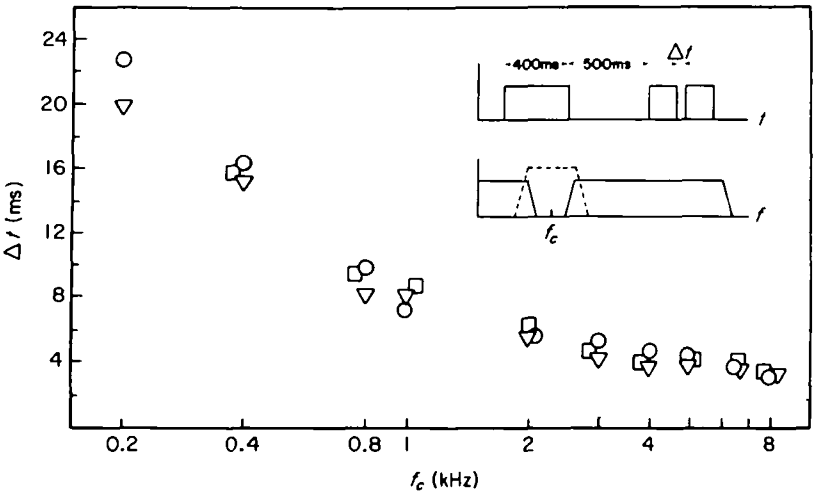

Fig. 7 Thresholds for the detection of a temporal gap (Δt) in a bandpass noise stimulus, as a function of the centre frequency, f_c, of the noise. Each symbol shows results for a different subject. The insets schematically illustrate the time course and frequency spectra of the stimuli. From Shailer & Moore.[32]

the longer the time for which it rings. Thus the increase in gap threshold at low frequencies may be explained by the decrease in the bandwidth of the auditory filter at low frequencies.

Temporal resolution in impaired hearing

Measurements of temporal resolution in the hearing-impaired show that considerable individual variability exists. Some listeners with cochlear impairments show almost normal temporal resolution, while others show marked impairments. Again, some caution is necessary in interpreting the results. We have seen that when wideband stimuli are used, normal-hearing subjects make use of high frequencies to perform the task. A subject with a loss primarily at higher frequencies will not be able to use high-frequency information effectively, and this alone could give rise to impaired performance; it is as if the impaired subject were listening to a low-pass filtered version of the stimulus.[33]

Several workers have measured gap-detection for band-limited noise to determine whether listeners with cochlear impairments have reduced temporal resolution when compared with normal subjects at similar centre frequencies.[30,34,35] On average, the impaired subjects did show reduced temporal resolution, and this impairment was observed whether the comparison with normal subjects was made at equal sound pressure level or at equal amounts above absolute threshold. Tyler et al.[34] and Moore & Glasberg,[36] have shown that increased gap thresholds are associated with a reduced ability to understand speech in noise.

THE LOCALISATION OF SOUNDS

Binaural cues

It has long been recognised that slight differences in the sounds reaching the two ears can be used as cues in sound localisation. The two major cues are differences in the time of arrival at the two ears and differences in intensity at the two ears. For example, a sound coming from the left will arrive first at the left ear and be more intense in the left ear. For steady sinusoidal stimulation, a difference in time of arrival is equivalent to a phase difference between the sounds at the two ears. However, phase differences are not usable over the whole audible frequency range. Experiments using sounds delivered by headphones have shown that a phase difference at the two ears can be detected and used to judge location only for frequencies below about 1500 Hz. This is explicable when we realise that, at high frequencies, the wavelength of sound is small compared to the dimensions of the head, so the listener cannot determine which cycle in the left ear corresponds to a given cycle in the right; thus phase differences become ambiguous and unusable at high frequencies. On the other hand, at low frequencies listeners' accuracy at detecting changes in relative time at the two ears is remarkably good; changes of 10–20 μs can be detected, which is equivalent to a movement of the sound source of 1–2° laterally.

Intensity differences between the two ears are primarily useful at high frequencies. This is because low frequencies bend or diffract around the head, so that there is little difference in intensity at the two ears whatever the location of the sound source. At high frequencies the head casts more of a 'shadow', and above 2–3 kHz the intensity differences are sufficient to provide useful cues. For complex sounds, containing a range of frequencies, the difference in spectral patterning at the two ears may also be important. For a discussion of the physiology of the localisation cues. (See D Moore, this issue.)

The idea that sound localisation is based on interaural time differences at low-frequencies and interaural intensity differences at high frequencies has been called the 'duplex theory' of sound localisation, and it dates back to Lord Rayleigh.[37] However, it has been realised in recent years that it is not quite correct.[38] *Complex* sounds containing only high frequencies (above 1500 Hz) can be localised on the basis of interaural time delays, provided that they have an appropriate temporal structure. For example, a single click can be localised in this way no matter what its frequency content. Periodic sounds containing only high-frequency harmonics can also be localised on the basis of interaural time differences, provided that the *envelope* repetition rate is below about 600 Hz.[39] Since many of the complex sounds we encounter in everyday life have envelope repetition rates below 600 Hz, interaural time differences will be used for localisation in most circumstances.

The role of the pinna

Binaural cues are not sufficient to account for all of our localisation abilities. For example, a simple difference in time or intensity will not define whether a sound is coming from in front or behind, or above or below, but people can clearly make such judgements. In recent years it has been shown that the pinnae play an important role in sound localisation.[40] They do so because the spectra of sounds entering the ear are modified by the pinnae in a way which depends upon the direction of the sound source. This direction-dependent filtering provides cues for sound source location. The pinnae are important not only for localisation, but also for judging whether a sound comes from within the head or from the outside world. A sound is only judged as coming from outside if the spectral transformations characteristic of the pinnae are imposed on the sound. Thus sounds heard through headphones are normally judged as being inside the head; the pinnae do not have their normal effect on the sound when headphones are worn. However, sounds delivered by headphones can be made to appear to come from outside the head if the signals delivered to the headphones are prerecorded on a dummy head or synthetically processed (filtered) so as to mimic the normal action of the pinnae. Such processing can also create the impression of a sound coming from any desired direction in space. The pinnae alter the sound spectrum primarily at high frequencies. Only when the wavelength of the sound is comparable with the dimensions of the pinnae is the spectrum significantly affected. This occurs mostly above about 6 kHz.

The precedence effect

In everyday conditions the sound from a given source reaches our ears by many different paths. Some of it will arrive via a direct path, but a great deal may only reach the ears after reflections from one or more surfaces. However, people are not normally aware of these reflections or echoes, and they do not appear to impair the ability to localise sound sources. The reason for this seems to lie in a phenomenon known as the precedence effect.[41] When several similar sounds reach the ears in close succession (i.e. the direct sound and its echoes) the sounds are perceptually fused into a single sound, and the location of the total sound is primarily determined by the location of the first (direct) sound. Thus the echoes have little influence on the perception of direction, although they may influence the timbre and loudness of the sound.

The precedence effect only occurs for sounds of a discontinuous or transient character, such as speech or music, and it can break down if the echoes are sufficiently long delayed or intense compared to the direct sound. However, in normal conditions the precedence effect plays an important role in the location and identification of sounds in reverberant conditions. It seems to be primarily a binaural phenomenon. When one ear is blocked room echoes become more noticeable; sounds then generally appear 'boomy' or 'muddy'.

Sound localisation in the hearing impaired

Most hearing losses result in some degradation in sound localisation.[42] However, there may be considerable individual differences even in patients with similar audiograms. In general, acoustic neuromas lead to greater localisation problems than cochlear losses. Most patients show a reduced ability to use interaural time and intensity differences. In addition, people with high-frequency hearing losses are generally unable to make use of the directional information provided by the pinnae. Hearing aid users also suffer in this respect, since, even if the microphone is appropriately placed within the pinna, the response of most aids is limited to frequencies below 6 kHz.

REFERENCES

1 Moore BCJ. Parallels between frequency selectivity measured psychophysically and in the auditory nerve. In: Cianfrone G, Grandori F, eds. Cochlear mechanics and oto-acoustic emissions. Scand Audiol 1986; Suppl 25: 139–152
2 Pickles JO. The neurophysiological basis of frequency selectivity. In: Moore BCJ, ed. Frequency selectivity in hearing. London: Academic Press 1986; pp. 51–121
3 Fletcher H. Auditory patterns. Rev Mod Phys 1940; 12: 47–65
4 Zwicker E. On a psycho-acoustical equivalent of tuning curves. In: Zwicker E, Terhardt E, eds. Facts and models in hearing. Berlin: Springer-Verlag 1974; pp. 132–141
5 Vogten LLM. Pure-tone masking: A new result from a new method. In: Zwicker E, Terhardt E, eds. Facts and models in hearing. Berlin: Springer-Verlag 1974; pp. 142–155
6 Patterson RD. Auditory filter shapes derived with noise stimuli. J Acoust Soc Am 1976; 59: 640–654
7 Patterson RD, Moore BCJ. Auditory filters and excitation patterns as representations of frequency resolution. In: Moore BCJ, ed. Frequency selectivity in hearing. London: Academic Press 1986; pp. 123–177
8 Moore BCJ, Glasberg BR. Suggested formulae for calculating auditory-filter bandwidths and excitation patterns. J Acoust Soc Am 1983; 74: 750–753
9 Pick GF, Evans EF, Wilson JP. Frequency resolution in patients with hearing loss of cochlear origin. In: Evans EF, Wilson JP, eds. Psychophysics and physiology of hearing. London: Academic Press 1977; pp. 273–281
10 Glasberg BR, Moore BCJ. Auditory filter shapes in subjects with unilateral and bilateral cochlear impairments. J Acoust Soc Am 1986; 79: 1020–1033
11 Patterson RD, Nimmo-Smith I, Weber DL, Milroy R. The deterioration of hearing with age: Frequency selectivity, the critical ratio, the audiogram and speech threshold. J Acoust Soc Am 1982; 72: 1788–1803

12 Dreschler WA, Plomp R. Relation between psychophysical data and speech perception for hearing-impaired subjects. J Acoust Soc Am 1980; 68: 1608–1615
13 Festen JM, Plomp R. Relations between auditory functions in impaired hearing. J Acoust Soc Am 1983; 76: 652–662
14 Moore BCJ, Glasberg BR. The role of frequency selectivity in the perception of loudness, pitch and time. In: Moore BCJ, ed. Frequency selectivity in hearing. London: Academic Press 1986; pp. 251–308
15 Laurence RF, Moore BCJ, Glasberg BR. A comparison of behind-the-ear high-fidelity linear hearing aids and two-channel compression aids, in the laboratory and in everyday life. Br J Audiol 1983; 17: 31–48
16 Moore BCJ, Laurence RF, Wright D. Improvements in speech intelligibility in quiet and in noise produced by two-channel compression hearing aids. Br J Audiol 1985; 19: 175–187
17 Evans EF. The sharpening of frequency selectivity in the normal and abnormal cochlea. Audiology 1975; 14: 419–442
18 Moore BCJ, Glasberg BR, Hess RF, Birchall JP. Effects of flanking noise bands on the rate of growth of loudness of tones in normal and recruiting ears. J Acoust Soc Am 1985; 77: 1505–1513
19 Moore BCJ. Frequency difference limens for short-duration tones. J Acoust Soc Am 1973; 54: 610–619
20 Goldstein JL, Srulovicz P. Auditory-nerve spike intervals as an adequate basis for aural frequency measurement. In: Evans EF, Wilson JP, eds. Psychophysics and physiology of hearing. London: Academic Press 1977; pp. 337–346
21 Moore BCJ, Glasberg BR, Shailer MJ. Frequency and intensity difference limens for harmonics within complex tones. J Acoust Soc Am 1984; 75: 550–561
22 Goldstein JL. An optimum processor theory for the central formation of the pitch of complex tones. J Acoust Soc Am 1973; 54: 1496–1516
23 Terhardt E. Pitch, consonance and harmony. J Acoust Soc Am 1974; 55: 1061–1069
24 Moore BCJ. An introduction to the psychology of hearing, 2nd edn. London: Academic Press 1983
25 Hoekstra A, Ritsma RJ. Perceptive hearing loss and frequency selectivity. In: Evans EF, Wilson JP, eds. Psychophysics and physiology of hearing. London: Academic Press 1977; pp. 263–271
26 Rosen S, Fourcin A. Frequency selectivity and the perception of speech. In: Moore BCJ, ed. Frequency selectivity in hearing. London: Academic Press 1986; pp. 373–487
27 Plomp R. Aspects of tone sensation. London: Academic Press 1976
28 Plomp R. Rate of decay of auditory sensation. J Acoust Soc Am 1964; 36: 277–282
29 Penner MJ. Detection of temporal gaps in noise as a measure of the decay of auditory sensation. J Acoust Soc Am 1977; 61: 552–557
30 Fitzgibbons PJ, Wightman FL. Gap detection in normal and hearing-impaired listeners. J Acoust Soc Am 1982; 72: 761–765
31 Fitzgibbons PJ. Temporal gap detection in noise as a function of frequency, bandwidth and level. J Acoust Soc Am 1983; 74: 67–72
32 Shailer MJ, Moore BCJ. Gap detection as a function of frequency, bandwidth and level. J Acoust Soc Am 1983; 74: 467–473
33 Bacon SP, Viemeister NF. Temporal modulation transfer functions in normal-hearing and hearing-impaired listeners. Audiology 1985; 24: 117–134
34 Tyler RS, Summerfield Q, Wood EJ, Fernandes MA. Psychoacoustic and phonetic temporal processing in normal and hearing-impaired listeners. J Acoust Soc Am 1982; 72: 740–752
35 Glasberg BR, Moore BCJ, Bacon SP. Gap detection and masking in hearing-impaired and normal-hearing subjects. J Acoust Soc Am 1987; 81: 1546–1556

36 Moore BCJ, Glasberg BR. Relationship between psychophysical abilities and speech perception for subjects with unilateral and bilateral cochlear impairments. In: Schouten MEH, ed. Psychophysics and speech perception. The Hague: Nijhoff 1987; In press
37 Rayleigh Lord. On our perception of sound direction. Phil Mag 1907; 13: 214–232
38 Hafter ER. Spatial hearing and the duplex theory: How viable? In: Edelman GM, Gall WE, Cowan WM, eds. Dynamic aspects of neocortical function. New York: Wiley 1984; pp. 425–448
39 Neutzel JM, Hafter ER. Lateralization of complex waveforms: spectral effects. J Acoust Soc Am 1981; 69: 1112–1118
40 Batteau DW. The role of the pinna in human localization. Proc Roy Soc B 1967; 168: 158–180
41 Wallach H, Newman EB, Rosenzweig MR. The precedence effect in sound localization. J Exp Psychol 1949; 27: 339–368
42 Durlach NI, Thompson CL, Colburn HS. Binaural interaction in impaired listeners. Audiology 1981; 20: 181–211