38

Simulation of the Effects of Sensorineural Hearing Loss

by

Isaac John Graf

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1997

Author ....
Department of Electrical Engineering and Computer Science
August 22, 1997

Certified by ....
Louis D. Braida
Henry Ellis Warren Professor of Electrical Engineering
Thesis Supervisor

Accepted by ....
Arthur C. Smith
Chairman, Department Committee on Graduate Theses

Simulation of the Effects of Sensorineural Hearing Loss

by

Isaac John Graf

Submitted to the Department of Electrical Engineering and Computer Science
on August 22, 1997, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

By processing sounds with combinations of amplitude expansion and spectral smearing, listeners with normal hearing can experience the elevated thresholds, abnormal growth of loudness, and reduced frequency selectivity that are characterisic of sensorineural hearing loss. The effects of these simulations in isolation and in combination were investigated in a variety of psychoacoustic and speech reception tasks. In tests of frequency selectivity, expansion raises narrowband noise masking patterns and broadens the low frequency side of psychoacoustic tuning curves. Frequency smearing broadens the narrowband noise masking patterns and broadens both sides of the tuning curves. In tests of loudness matching, expansion substantially reduces loudness summation and frequency smearing has little effect. The effects of the simulations on consonant reception depend on SNR. Expansion causes a greater reduction in intelligibility at high SNR's than low SNR's, and spectral smearing is more significant at low SNR's than high SNR's.

Thesis Supervisor: Louis D. Braida
Title: Henry Ellis Warren Professor of Electrical Engineering

Acknowledgements

I thank my advisor, Professor Louis Braida, for the tremendous amount of help he has given me throughout this project. For every question and problem that I came across, however mundane or complex, he was always willing to listen and advise, and to provide moral support whenever necessary. I thank Dave Lum for all the time he spent to teach me about computers and to fix whatever was not working right. I thank Jeanie, Jay, Paul, and Ann. They are a great group to work with. I also thank my subjects TP, BS, JR, LP, KM, JK, and AC for their time and patience.

I thank my Speech and Hearing classmates- Diane, Janet, Joyce, Mark, Mike- as well as the Speech and Hearing faculty for contributing to so many thought-provoking, stimulating classes. And with sadness, I thank Ben Hammond. He was an inspiration for us and I hope we will always remember his overflowing enthusiasm and generosity.

Most importantly I thank my parents. I would not know where to begin or end listing my thanks so I won't try. To my brothers- Nathan, Solomon, and Benjamin- well, I had to stick you in here somehow, right? And to Esther, this has been a wonderful year for me.

Finally, I acknowledge the NIH, which has provided the funding for my studies in graduate school.

In memory of my great uncle, Dr. Peter Denes.

Contents

Chapter 1

Background

Recent improvements in digital hardware provide a good deal of flexibility in the kind of processing that can be provided in hearing aids. The main issue now is to determine what kind of processing to implement for different kinds of hearing impairment. A prerequisite for the development of better processing schemes is a greater understanding of sensorineural hearing loss.

There are two primary methodologies that can be used to learn about hearing loss-auditory physiology and psychoacoustics. Physiological studies of the inner ear investigate how the ear (usually of animals) responds to different acoustic stimuli. Psychoacoustics looks at how human subjects perceive different acoustic stimuli. Hearing loss simulation, which is the focus of this research, is an example of a psychoacoustic method. However, since it can be helpful to think of psychoacoustic experiments in terms of their physiological correlates, a brief discussion of auditory physiology will be presented here.

1.1 Physiology of Hearing Impairment

It has long been known that the cochlea carries out a frequency to place transformation on sound that arrives at the ear. At the base of the cochlea, the basilar membrane is stiff and tends to respond well to high frequencies. As it proceeds toward the apex (helicotrema) the basilar membrane becomes looser and responds better to lower frequencies. The frequency which elicits the largest response at a given place on the basilar membrane is termed the characteristic frequency (CF) or best frequency (BF) of that point on the membrane.

Besides being caused by the material properties of the basilar membrane, the response of the basilar membrane to sound is affected by the actions of outer hair cells (OHC's), which run along the length of the membrane. It is largely agreed that outer hair cells provide a frequency specific nonlinear amplification on the basilar membrane. That is, for a given point on the membrane, the OHC's will generate a large increase in response if the frequency of the stimulus matches the CF of that point and if the level of the stimulus is low. For higher levels and for off-CF stimuli, the OHC's seem to provide less and less amplification. Since the OHC amplification decreases as level increases, it is compressive. Figure 1.1 shows that the basilar membrane response is stronger and more compressive for frequencies close to CF and the response is more linear for frequencies away from CF.

It thus stands to reason that if the OHC's do not function normally (as may often be the case in sensorineural hearing loss) the cochlea does not have the amplifier and sharp frequency tuning that the healthy cochlea has. As a result of the lack of the cochlear amplifier, weak sounds will elicit much less response than normal (since the amplification is greatest for such sounds), whereas intense sounds will elicit a response similar to normal. This is likely the physiological basis for loudness recruitment, or abnormal growth of loudness with sound level (see Section 1.2). Due to the loss of sharp tuning on the basilar membrane, frequency selectivity will also likely be reduced: a given location on the basilar membrane will no longer respond so strongly to one frequency relative to all others.

One way to characterize frequency selectivity on a physiological level is through tuning curves of auditory nerve fibers (ANFs). A tuning curve specifies the level of a tone required to elicit a specified response (e.g., a minimum firing rate) in an ANF, as a function of frequency. It has been found that the shape of ANF tuning curves correspond well with the shape of basilar membrane tuning curves at the place on the basilar membrane corresponding to that ANF. Many investigators have shown that the tuning of auditory nerve fibers gets significantly broader when OHC's have been damaged. (See Figure 1.2).

1.2 Psychoacoustics of Hearing Impairment

One psychoacoustic phenomenon that is commonly associated with hearing loss of cochlear origin is loudness recruitment, or an abnormal loudness growth with sound level. Loudness recruitment can perhaps best be described by loudness balances obtained from people with unilateral hearing loss (see Figure 1.3). For the experiment, a tone is presented at a given level in the normal ear and a tone of the same frequency is also presented in the impaired ear. The subject adjusts the tone intensities until the sounds are equally loud. A log-log plot of the level presented to the normal ear as a function of the level presented to the impaired ear for equal loudness has a slope greater than 1 for intensities near threshold. Thus for a given change in sound level, the loudness change tends to be greater for a hearing impaired person than for a normal hearing individual.

Other typical psychoacoustic effects of sensorineural hearing loss are reduced frequency selectivity and reduced temporal resolution. Frequency selectivity can be measured by a variety of different psychophysical tests, such as narrowband noise masking patterns and psychoacoustic tuning curves (see Chapter 2). Temporal resolution can be measured by psychoacoustic tests such as gap detection and forward masked recovery times. It has been suggested (Moore, 1995) that poor temporal resolution in the hearing impaired as measured by these psychoacoustic tests is a result of loudness recruitment.

On a psychoacoustic level, one way to learn more about hearing loss is to conduct experiments on hearing impaired listeners and compare their abilities with those of listeners with normal hearing. Although this may be the most straightforward way to gain insight into hearing impairment, this method has the disadvantage of not being able to separate the effects of the different aspects of hearing impairment (e.g. elevated thresholds, recruitment, poor frequency and temporal resolution, etc.). Most hearing impaired listeners exhibit all of these phenomena to varying degrees, and correlational studies that attempt to link specific phenomena with speech reception have been rather inconclusive (Dreschler and Plomp 1980,1985). Another problem with comparing hearing impaired results with those of normal listeners arises from the level dependence of nearly all auditory abilities. The performance of the hearing impaired relative to normals depends on whether the two groups are compared at equal sound pressure levels, equal sensation levels, or equal loudness levels. As an example, hearing impaired ears are typically more sensitive to intensity differences than normal ears, when measured at equal sensation levels. When measured at equal sound pressure levels, however, they are generally less sensitive.

1.3 Simulating Hearing Loss

Another method that has been used by investigators to learn about hearing impairment is to simulate hearing loss. These studies have attempted to simulate one or more of three different phenomena that are commonly associated with hearing impairment: 1) elevated thresholds and loudness recruitment, 2) poor frequency selectivity, and 3) poor temporal resolution. (All phenomena may be frequency dependent).

1.3.1 Reasons for Simulating Hearing Loss

The goals of conducting experiments with simulating hearing loss may be enumerated as follows. 1) To test psychoacoustic models of hearing impairment: to evaluate the simulation models and compare them to ears with sensorineural hearing loss. 2) To separate the effects of recruitment, poor frequency resolution, and poor temporal resolution. 3) To gain insight into signal processing schemes that might be effective for alleviating the effects of hearing impairment. By separating the effects of different aspects of hearing loss, it should be possible to determine which aspects of hearing loss are most important to compensate for. 4) To allow normal hearing listeners to perceive what it is like to be hearing impaired. This is important for educating those who must deal with impaired listeners.

1.3.2 Previous Work on Simulating Hearing Loss

Following the work of Steinberg and Gardener (1937), early studies of simulated sensorineural hearing loss used spectrally shaped masking noise to simulate hearing loss. Additive noise simulations have been shown to elevate hearing thresholds and introduce loudness recruitment. These simulations have also been shown to produce consonant recognition scores similar to those of the hearing impaired (Zurek and Delhorne, 1987). There are several drawbacks to the noise simulation, however. First, the perceived sensation is not accurately simulated, since, except for tinnitus, the hearing impaired generally do not hear constant noise. Second, it is impossible to simulate severe impairments, due to the high levels of noise that would be required. Third, the physiological substrates of threshold elevation and recruitment may be different in the hearing impaired and noise-masked normals. Phillips (1987) suggests that recruitment caused by masking noise is only manifest in the higher auditory system, whereas the recruitment in the hearing impaired is likely due to OHC damage and is manifest in the firing patterns of the auditory nerve.

Dubno and Schaefer (1991) looked at the effect of noise simulations on several measures of frequency selectivity. They concluded that as determined by masked thresholds measured in broadband noise (critical ratios) the frequency selectivity was comparable in the noise masked normals and the hearing impaired. However, for forward masked psychoacoustic tuning curves and simultaneous narrowband noise masking pattern measurements (see Chapter 2), frequency selectivity is poorer for hearing impaired listeners than for noise-masked normals, even when threshold signal levels are equated.

In recent years, several investigators have used signal processing to simulate hearing loss. Elaborating on Villchur's idea of amplitude expansion, Duchnowski (1989), Moore et al. (1993), and Moore et al. (1995), have investigated the effects of elevated thresholds and loudness recruitment on speech reception. They used DSP implemented amplitude expansion on speech recognition in noise. Duchnowski found that the expansion simulation simulated the consonant and speech reception scores of moderate to severely impaired subjects in many but not all conditions. In particular, he found that normal hearing listeners who experience the expansion simulation achieve higher intelligibility scores than hearing impaired listeners in a high frequency emphasis condition. Moore et al. (1995) found that expansion has much more detrimental effects in interference produced by a single competing talker than in speech-shaped Gaussian noise, and that linear amplification is more effective in the latter case than in the former. They also observed that the effect of expansion depends on the type of hearing loss being simulated (i.e. flat vs. sloping, moderate vs. severe).

All these studies investigated the effects of simulated hearing loss on speech reception. However, speech reception can be degraded by a large variety of distortions, and thus even if a simulation can make speech reception comparable to that of the hearing impaired, it is not clear whether the reduced intelligibility results from perceptual phenomena that are similar to what the hearing impaired experience. In order to obtain a more careful assessment of the accuracy of a hearing loss simulation it is necessary to conduct psychophysical tests. Lum (1995a, 1995b) implemented the expansion processing in real time and observed its effects on a number of psychophysical tasks. He found that expansion can simulate reduced frequency selectivity as measured by forward masked psychophysical tuning curves fairly accurately. It produced little noticeable change in critical ratios, which is also in accordance with the results from the hearing impaired. However, the expansion simulation did not simulate reduced frequency selectivity as measured by simultaneous narrowband noise masking patterns. In addition, the expansion simulation preserves the finely varying microstructure of the audiogram of normal listeners (which is not characteristic of the hearing impaired), and it results in an ability to discriminate intensities that is considerably better than what is observed in hearing impaired listeners.

Other investigators have used DSP to smear the spectrum of speech in an effort to simulate the effects of poor frequency selectivity. For example, ter Keurs et al. (1992, 1993a) employed processing that smeared the spectral envelope along the frequency axis, while preserving the spectral fine structure. They found that speech reception in noise decreases as the amount of spectral smearing increases beyond a critical bandwidth. They also found that smearing affects vowel recognition more than consonant recognition. In their 1992 study, these authors observed that the difference in effects produced by single-talker and filtered noise interference decreased as the amount of spectral envelope smearing increased (a finding similar to that of Moore for amplitude expansion).

In a further study, ter Keurs et al. (1993b) found that the degree of smearing at which point speech reception in noise begins to suffer does not correlate with the frequency selectivity of the listener. They suggest that if reduced frequency selectivity were an important contributor to poor speech reception in noise, then the degree of smearing at which speech reception begins to suffer should be proportional to the width of the listener's auditory filters. They thus infer from their data that poor frequency selectivity has little effect on speech reception.

Moore and colleagues have also studied a type of spectral smearing, in which the short-term power spectrum rather than the spectral envelope is smeared (Baer and Moore, 1993). This processing preserved the short-term phase of the signal. They measured speech reception when different amounts of smearing were applied to the speech signal. They found, in agreement with ter Keurs et al., that mild smearing of quiet speech reduces speech scores negligibly. Performance degrades as the speech to noise ratio decreases and as the amount of smearing increases.

In recent studies, Nejime and Moore (1997) and Moore et al. (1997) have investigated the effects of combining amplitude expansion and frequency smearing on speech recognition for the normal ear of subjects with unilateral hearing loss. In their processing, they performed the frequency smearing first, followed by the amplitude expansion. Their findings indicate that smearing combined with expansion produces poorer results than expansion alone, but the combined simulations still yield better speech recognition than does unprocessed speech in the impaired ear.

1.4 Problem Statement:

In order to draw conclusions about the role of frequency selectivity in speech recognition from spectral smearing experiments, it is important to know how spectral smearing affects performance on psychophysical measures of frequency selectivity. It is also important to understand how the effects of smearing and expansion interact to affect the performance on different tasks. To this end, experiments were conducted to measure frequency selectivity and speech reception with the frequency smearing processing with and without amplitude expansion. These experiments focus on conditions that have not been tested in prior studies and that are expected to contribute significantly to our understanding of auditory processing. The experiments conducted include simultaneous narrowband masking patterns, forward masked psychophysical tuning curves, loudness summation, and consonant reception. In order to facilitate comparisons with previous work, for the two frequency selectivity experiments, the expansion conditions were chosen to match those used by Lum and Dubno and Schaefer, and for some of the consonant reception experiments, expansion conditions were chosen to match those of Duchnowski .

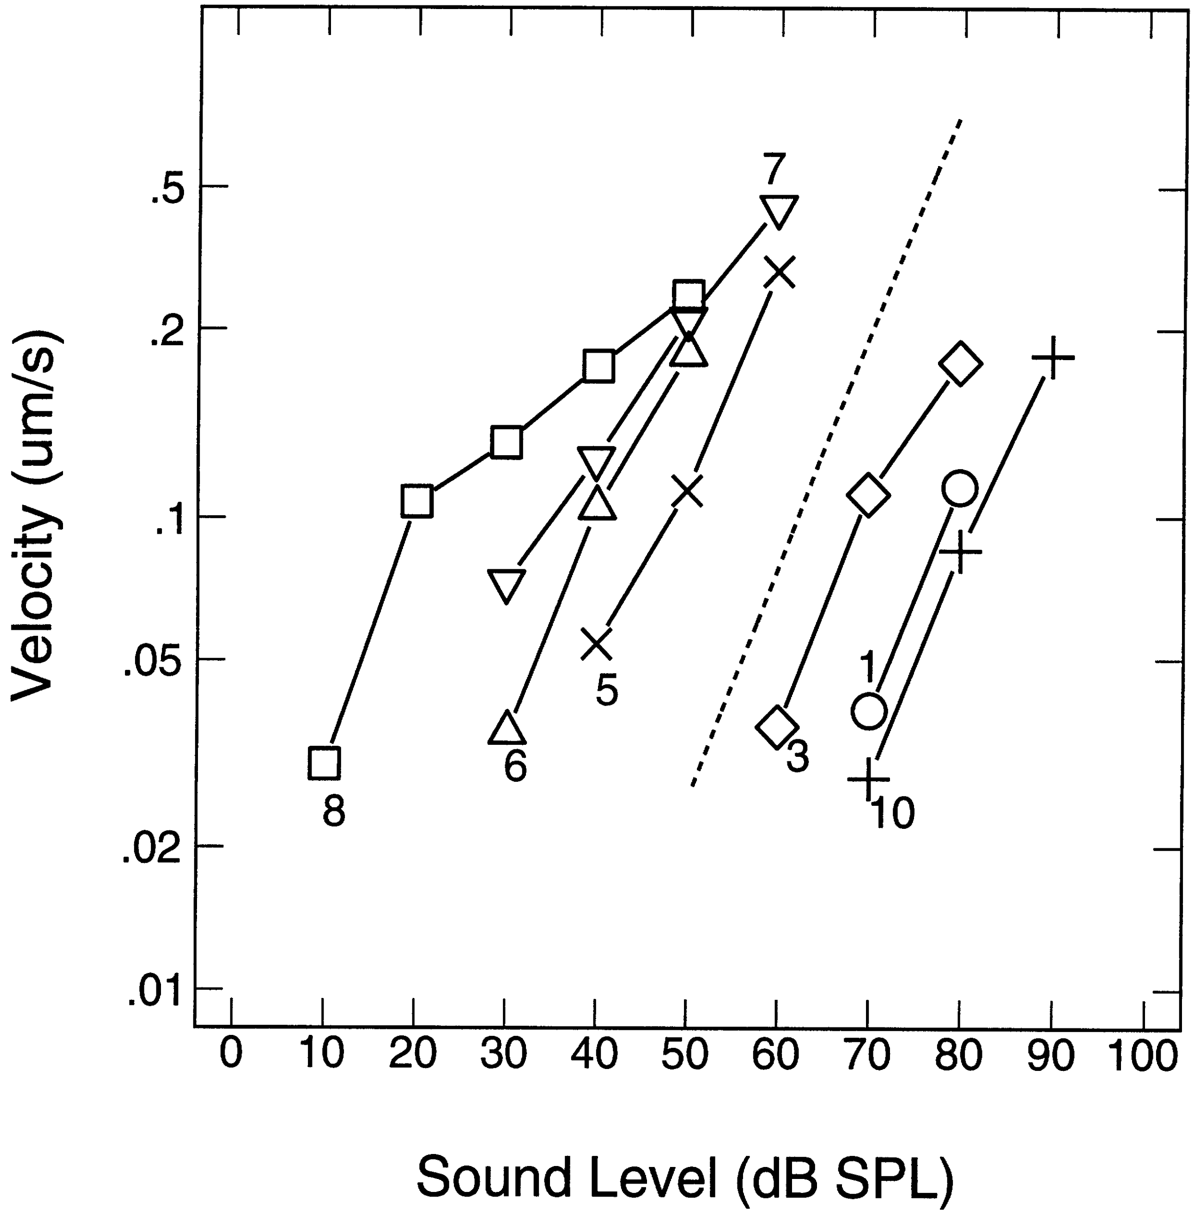

Figure 1-1: Basilar membrane velocity of a chinchilla at the point on the basilar membrane where CF = 8 kHz. The stimulating frequency is shown by the number close to each plot. The dashed line indicates what the slope would be if basilar membrane response were linear. Redrawn from Robles et al. (1986).

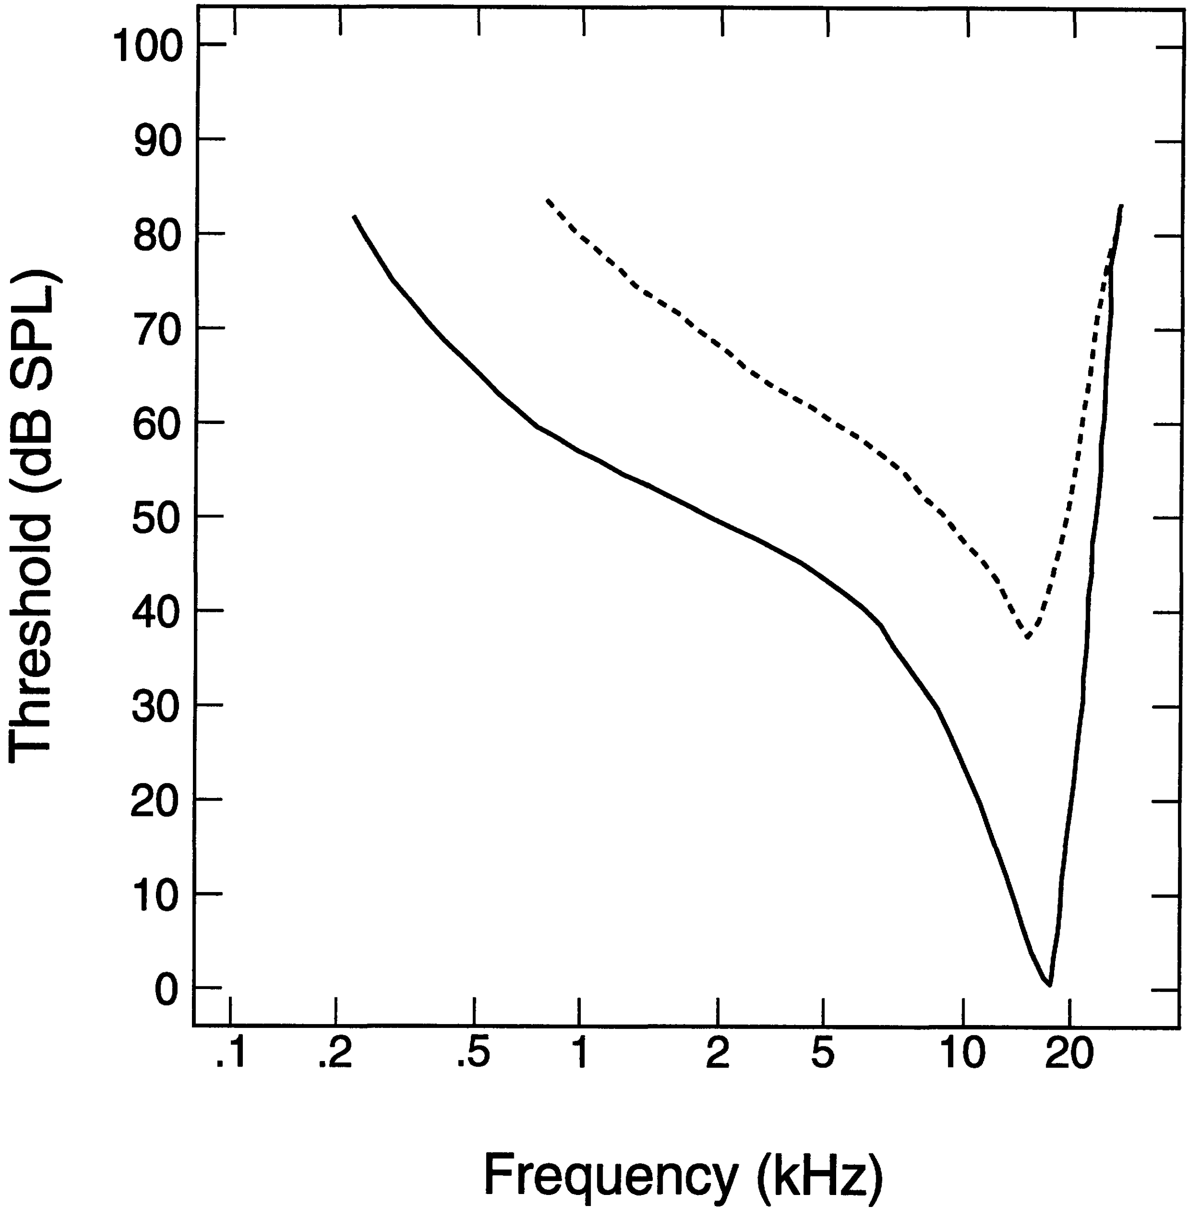

Figure 1-2: Typical examples of ANF tuning curves from a normal cat ear (solid line) and from an ear with damaged OHC's (dotted line). Liberman and Dodds, 1984.

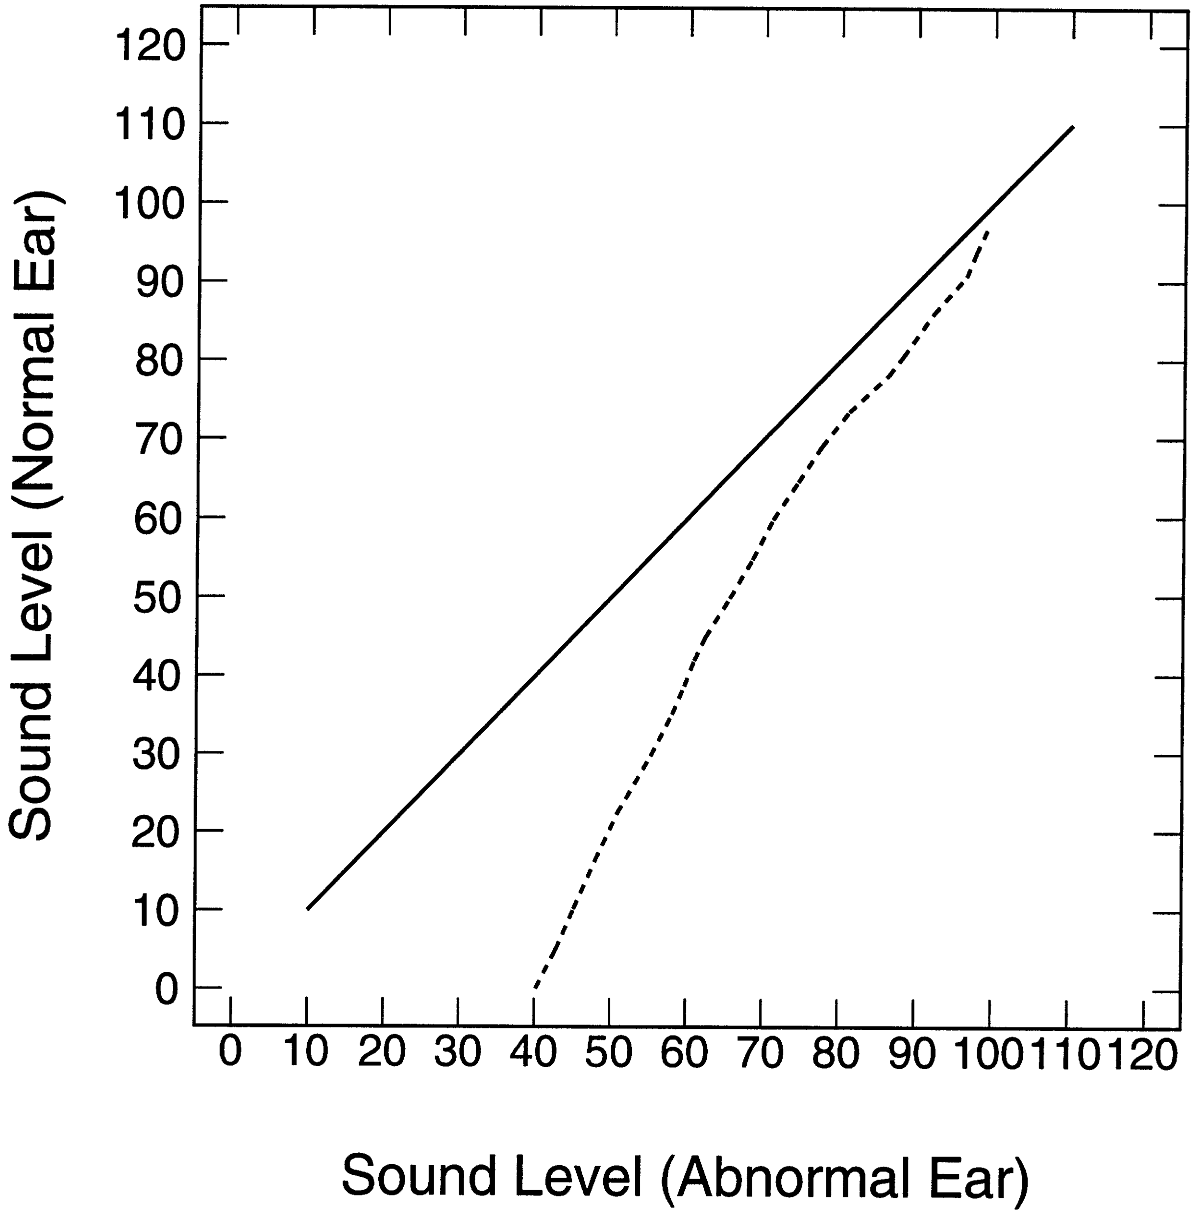

Figure 1-3: Loudness matching pattern for a subject with unilateral hearing loss (dotted line). Redrawn from Miskolczy- Fodor, 1960. A pattern of no recruitment is shown with the solid line.

Chapter 2

Experiments

The first two experiments that were conducted provide methods of characterizing frequency selectivity. The third experiment looks at loudness perception, and the fourth experiment studies consonant reception. In order to motivate testing these experiments on subjects with simulated hearing loss, this chapter presents results for these experiments from normal and hearing impaired listeners and discusses how these results might be interpreted.

2.1 Simultaneous Masking with Narrowband Noise

One measure of frequency selectivity is provided by narrowband noise masking patterns. These patterns are measured by presenting a narrowband of noise (200 Hz wide, for example) at a particular sound level and at a particular center frequency. The masked threshold of a tone is measured as a function of frequency for frequencies above, below, and in the bandpass of the noise (See Figure 2.1). If the bandwidth of the noise is equal to the critical bandwidth, then for frequencies in the band of the noise, the threshold of the tone is about the same level as the level of the noise in a critical band. As the frequency is moved further away from the noise passband, the threshold drops. Figure 2.1 shows that for a 200 Hz wide band of noise centered at 1200 Hz, at around 500 Hz on the low side and

4000 Hz on the high side the narrowband noise does not raise the threshold of the tone at all for normal hearing listeners.

For the hearing impaired, the narrowband noise masking pattern is broader than for normals. The on-frequency threshold is again about the same as the level of the noise and at about 500 Hz and 3000 Hz the masked thresholds match the unmasked detection thresholds, as is the case for the normals (though at 4000 Hz the masked threshold seems to rise again for the hearing impaired) (See Figure 2.2).

The width of the narrowband noise masking pattern can be viewed as a measure of frequency selectivity for the following reason. When a tone is presented at a frequency at a barely detectable level, it is likely that only a few auditory nerve fibers, with best frequencies near the frequency of the tone, are responding. On a more psychoacoustic level, it can be viewed as though one auditory filter is responding to the tone when the tone is presented at threshold level. When narrowband noise is present, these fibers (or auditory filter) may be responding to the noise. In this case, the level of the tone will have to be increased before their response to the tone plus noise is detectably greater than their response to the noise alone. The extent to which a fiber with a best frequency away from the passband of the noise responds to the noise is dependent on the frequency selectivity of the fiber. For example, if the passband of the noise is 1100-1300 Hz and we focus on a fiber with a best frequency of 2500 Hz, then if there was significant response to the noise, the fiber is not very frequency selective. On the other hand if a fiber with a best frequency of 1600 Hz does not respond at all to the noise, then this fiber is highly frequency selective.

It is not obvious how to compare the narrowband noise masking patterns of the hearing impaired with those of normals. From one point of view, the frequency selectivity may not be much different for the two sets of subjects. The masking, or difference between masked threshold and unmasked threshold, is roughly 10 dB at 2500 Hz for both normals and the hearing impaired, and at 1600 Hz the masking is larger for the normal listeners, implying greater spread of masking (poorer frequency selectivity) for normals. Looked at another way, frequency selectivity is a measure of the response of a fiber (or filter) to sound at its CF relative to its response to sound away from its CF. A high ratio corresponds to good frequency selectivity and a low ratio corresponds to poor frequency selectivity. In this sense, frequency selectivity is poorer than normal in the hearing impaired. For high frequencies, the masked thresholds for the hearing impaired are typically considerably above the unmasked thresholds even though the unmasked thresholds in the hearing impaired are higher than the masked thresholds for the normal listeners. If the fiber in the impaired ear with best frequency of 2500 Hz responded to the off-CF noise relative to the on-CF tone with the same ratio as for the normal cochlea, then the response of the fiber to the noise should be equivalent to the response to a 2500 Hz tone at around 15 dB SPL (the masked normal threshold). Since this response is well below threshold for the hearing impaired, the noise would be expected to cause no threshold elevation for a tone of 2500 Hz. (It should essentially be undetectable for a fiber with CF at 2500 Hz). However, if frequency selectivity is reduced for the hearing impaired, then the response of this fiber to the noise relative to its response to the on-CF tone will be increased compared with that of the normal cochlea. An additional argument that frequency selectivity is reduced in the hearing impaired is that when broadband noise is presented to a normal subject in order to raise thresholds to that of the hearing impaired, the narrowband noise masking pattern is much narrower than it is for the hearing impaired (see Fig. 2.2). Dubno argues that this finding suggests that the broadening of the masking pattern is not accounted for by elevated thresholds alone.

2.2 Forward Masked PTC's

A second measure of frequency selectivity in the auditory system is provided by psychophysical tuning curves (PTC's). A PTC shows the level of a masking tone that is required to mask a probe tone of a fixed level and frequency, as a function of frequency of the masker tone. Thus higher thresholds indicate greater frequency selectivity- the masker tone has less masking effect on the probe tone and lower thresholds indicate less frequency selectivity. Examples of forward masked PTC's for normal and hearing impaired listeners are shown in Figure 2.3. (Forward masked indicates that the masker tone is presented before the probe tone is presented).

The normal PTC's typically exhibit a minimum masker level at the frequency of the probe. This makes sense because the same fibers that are most sensitive to the probe are the ones being most affected by the masking tone. The slope of the PTC rises sharply on the high frequency side and less sharply on the low frequency side (on a logarithmic frequency axis). On the low frequency side, the slope of the curve tends to get smaller as the frequency is decreased further from that of the probe tone. The hearing impaired show a wide variety of shapes in their forward masked PTC's (see Figure 2.3). In all cases, though, the tuning curve rises more slowly on both the low and high frequency sides. In some cases, especially for subjects with mild hearing loss, the high frequency side of the tuning curve is still relatively sharp, while the low frequency side broadens more considerably. For more severe hearing losses, the PTC can be almost flat.

On a physiological level, the correlate of the psychophysical tuning curve may be viewed in a similar way to that of simultaneous narrowband noise masking. In order for the probe tone to be detected, it must elicit a sufficient change in response from the nerve fibers most sensitive to the probe. When a masking tone precedes the probe tone, if the cells used to detect the probe have responded vigorously enough to the masker tone, then it is likely that neural adaptation will have set in and they will not be able to sufficiently respond to the probe. It is assumed that the greater their response to a given frequency of masker, the lower the level of the masker that is required to mask the probe tone. (This situation is one of excitatory masking). Thus greater frequency selectivity is manifested by sharper PTC's. For a given masker frequency at a different frequency than the probe, if the fibers with CF at the probe are highly frequency selective then they will not respond much to the masker tone and a high level masker will be required to mask the probe as a result.

Forward masked PTC's in the hearing impaired (and simulated normals) provide an alternate method to characterize frequency selectivity that complements the narrowband noise masking experiment. In addition, forward masked PTC's are probably the easiest psychoacoustic correlate of auditory nerve tuning curves to measure. In agreement with physiological tuning curves, forward masked psychoacoustic tuning curves are sharp at the tip and exhibit a sharper slope on the high frequency side than on the low frequency side.

2.3 Loudness Summation

In normal hearing listeners, if the bandwidth of a noise is increased while the overall intensity is held constant, the perceived loudness remains constant while the bandwidth stays within a certain limit (called a critical band). As the bandwidth exceeds a critical band, the loudness increases with increasing bandwidth. Loudness summation is defined for a wideband and narrowband noise pair with the same center frequency as the difference in levels of the two noises that produces equal loudness. (The level of the narrowband noise (in dB) minus the level of the wideband noise). It has been found that loudness summation varies with the level of the narrowband noise, as seen in Figure 2.4. For hearing impaired subjects, loudness summation is markedly reduced (Figure 2.4).

Motivation for measuring loudness summation on people with simulated hearing loss derives from a model of loudness perception developed by Zwicker (Moore, 1995). In his model, the loudness of a stimulus is determined by the "excitation pattern" produced by the stimulus. (The excitation pattern may be thought of as an approximation to the amplitude of the response along the basilar membrane to a given acoustic stimulus). To calculate the excitation pattern, the input stimulus is passed through a bank of auditory filters, and the output of these auditory filters is the excitation pattern. Specifically, for each frequency, the value of the excitation pattern is determined by the output of the auditory filter whose center frequency is that frequency. The "specific loudness" is then calculated for each critical band, by passing the output of each auditory filter through a compressive nonlinearity. The area under the specific loudness pattern is then proportional to the predicted loudness.

Florentine et al. (1979) hypothesized that the reduction in loudness summation could be explained by incorporating higher thresholds with recruitment and wider auditory filters into the Zwicker model of loudness perception. They found that incorporating recruitment with elevated thresholds alone does not account for much of the decrease in loudness summation, but that when widened auditory filters were incorporated, the predicted values of loudness summation were reduced considerably and matched the hearing impaired data well (see Figure 2.5). In a later study, Florentine et al. (1997) found that incorporating poor frequency selectivity and recruitment into Zwicker's model was able to account well for the observed loudness growth for hearing impaired individuals with both gradual and abrupt high frequency hearing losses. Launer et al. (1997) also used Zwicker's spread of excitation model to predict loudness summation and loudness growth in the hearing impaired. Instead of incorporating broadened auditory filters into the model, however, they altered parameters in the transformation from excitation pattern to specific loudness. They found that altering one component (corresponding to the threshold in each auditory filter) in this transformation equation can account well for reduced loudness summation, and that the altering of a second component (the exponent in the transformation equation) accounts well for loudness growth in the hearing impaired. By evaluating the effects of simulated recruitment and reduced frequency selectivity separately and in combination, the question of whether reduced frequency selectivity is necessary to account for the reduced loudness summation in the hearing impaired can be tested.

2.4 Consonant Recognition

In addition to conducting psychoacoustic experiments it seems important to look at the effects of the simulations on speech reception, since speech reception is probably the most important issue for the hearing impaired. It has been observed that while the hearing impaired can understand speech in quiet very well as long as it is audible, they experience greater difficulty than normals in the presence of interfering sounds. In the presence of broadband background noise, the hearing impaired have a speech reception threshold (SNR at which 50% of words are correctly identified) that is between 2-7 dB higher than that of normals. When the background noise is a competing talker, the speech reception threshold is over 12 dB higher than normal (Festen and Plomp, 1990). In order to investigate the effects of expansion and smearing on consonant recognition, the simulations were applied to the speech sounds in isolation and in combination.

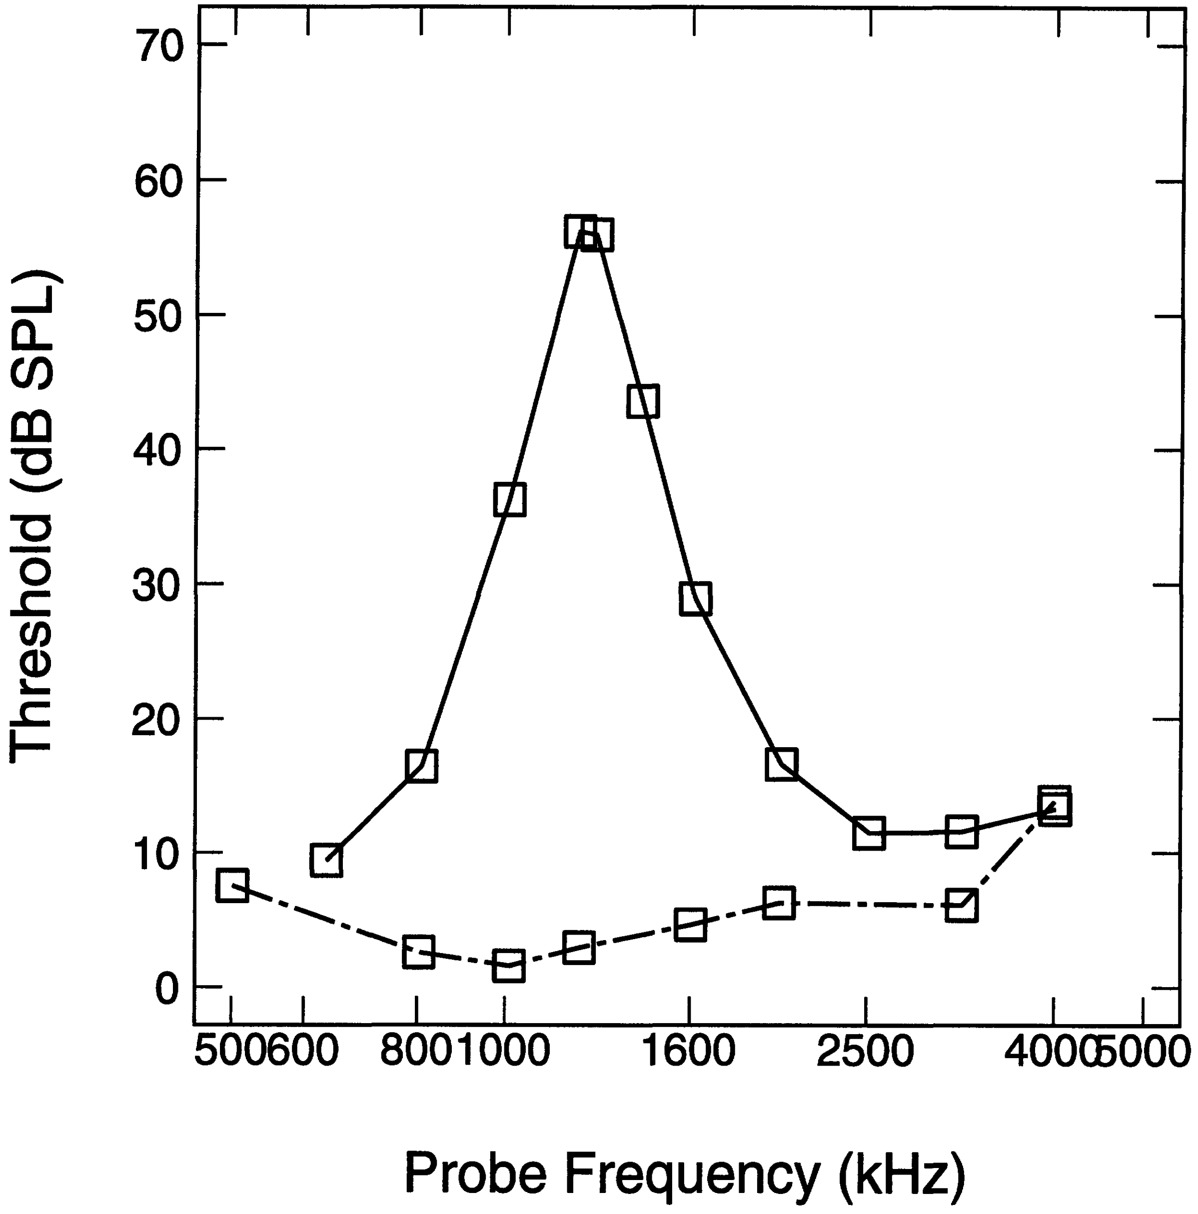

Figure 2-1: Narrowband noise masking pattern of normal hearing subject, shown by solid line (□). Noise band is 200 Hz wide, centered at 1200 Hz and at 63 dB SPL. Absolute thersholds are shown by broken line (□). Redrawn from Dubno and Schaefer, 1991.

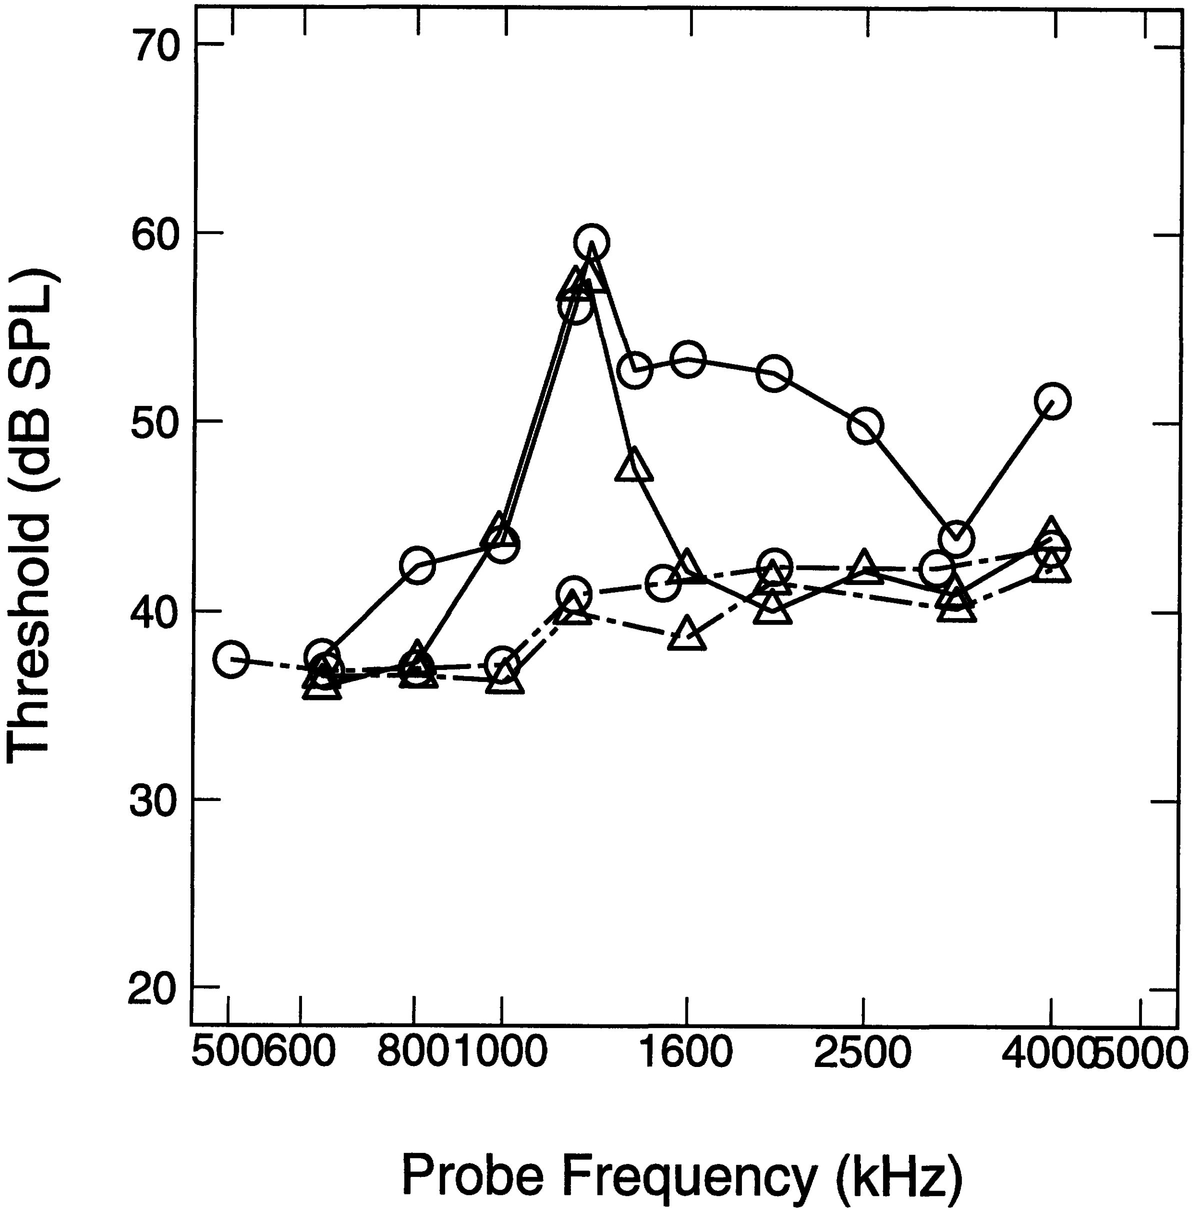

Figure 2-2: Narrowband noise masking patterns and thresholds for the hearing impaired (o) and normal listeners with hearing loss simulated by masking noise (△). Solid lines indicate masking patterns. Broken lines indicate absolute thresholds. Redrawn from Dubno and Schaefer, 1991.

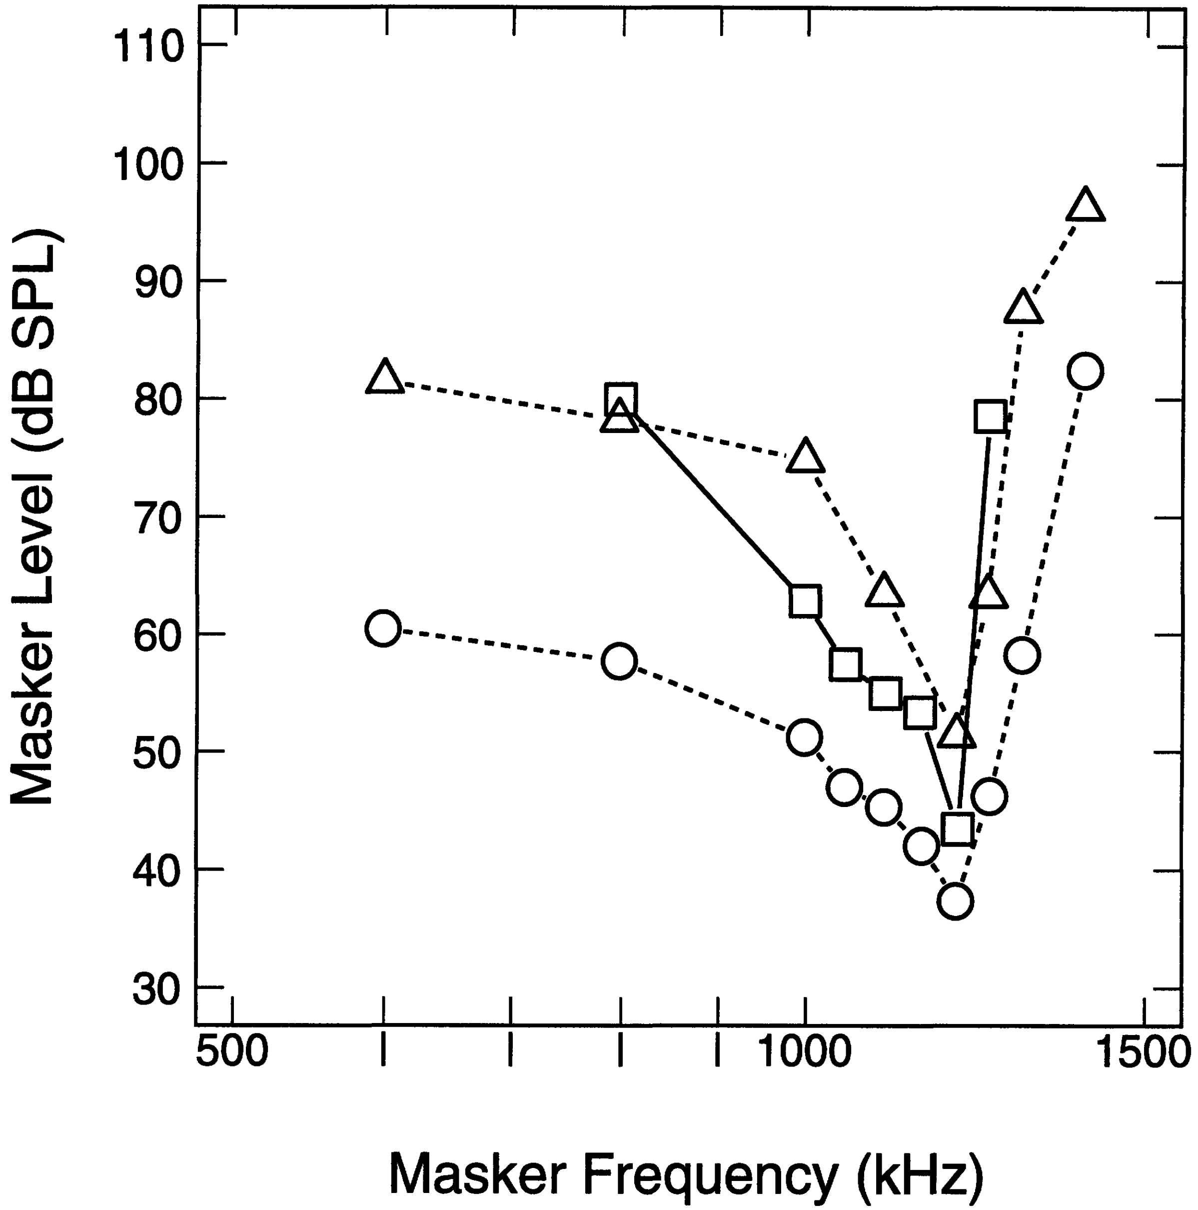

Figure 2-3: Forward masked psychophysical tuning curves for a normal subject (□, solid line) and two hearing impaired subjects (○, △, dotted lines). Probe tone is at 36 dB SPL. Redrawn from Dubno and Schaefer, 1991.

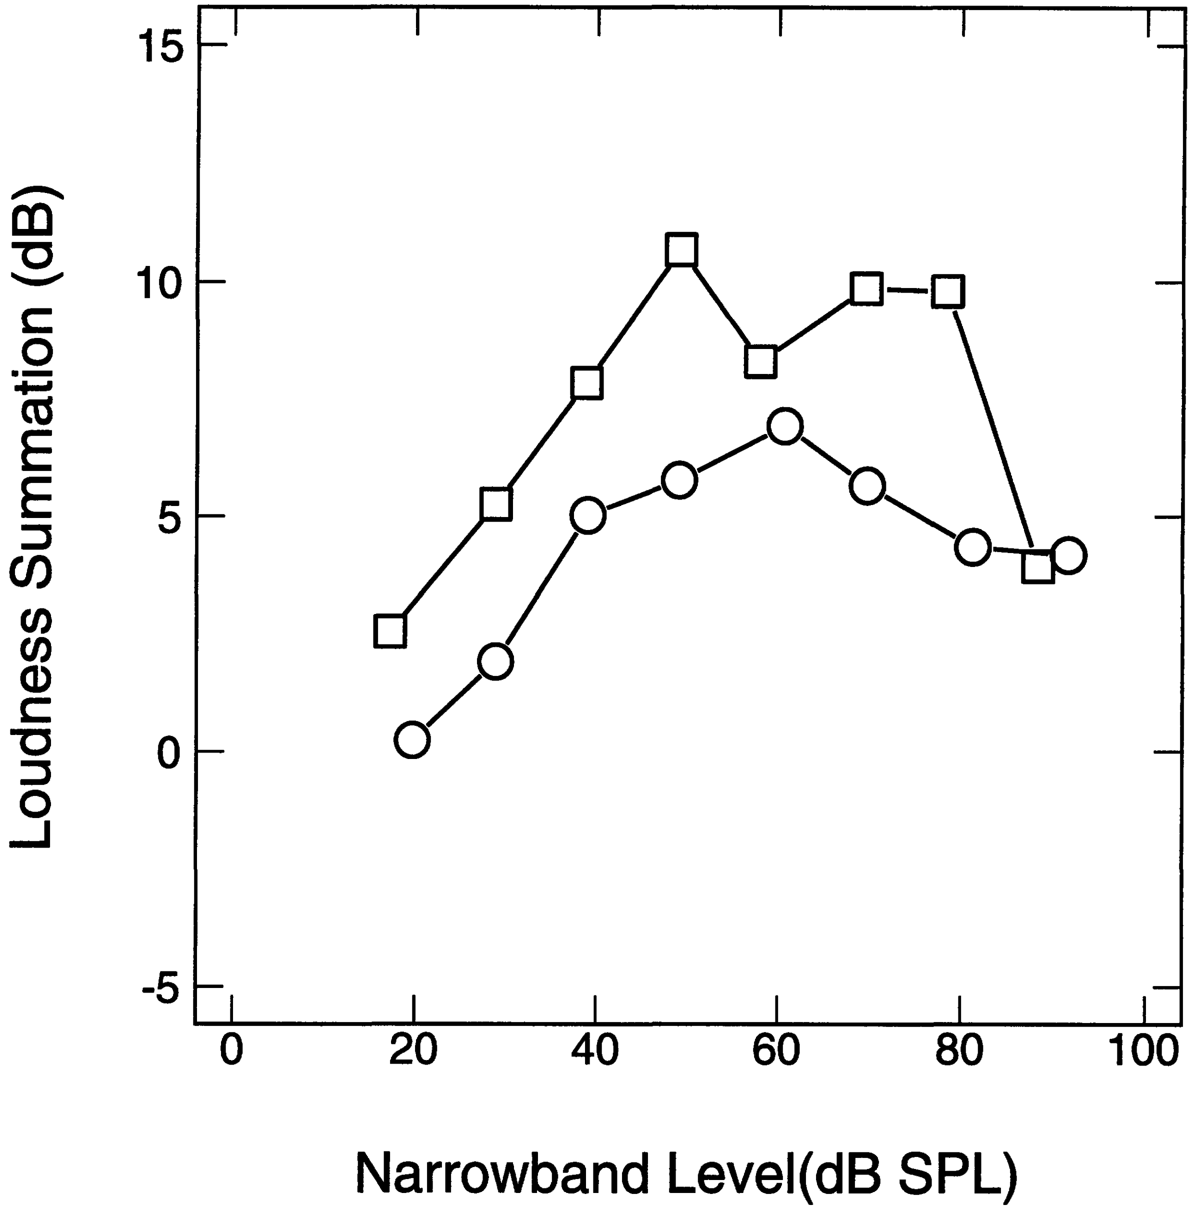

Figure 2-4: Loudness summation for normal hearing (□) and hearing impaired (○) subjects. Narrowband noise is in the band 3761-4380 Hz. Wideband noise is in the band 2000-7909 Hz. Redrawn from Florentine et al., 1980.

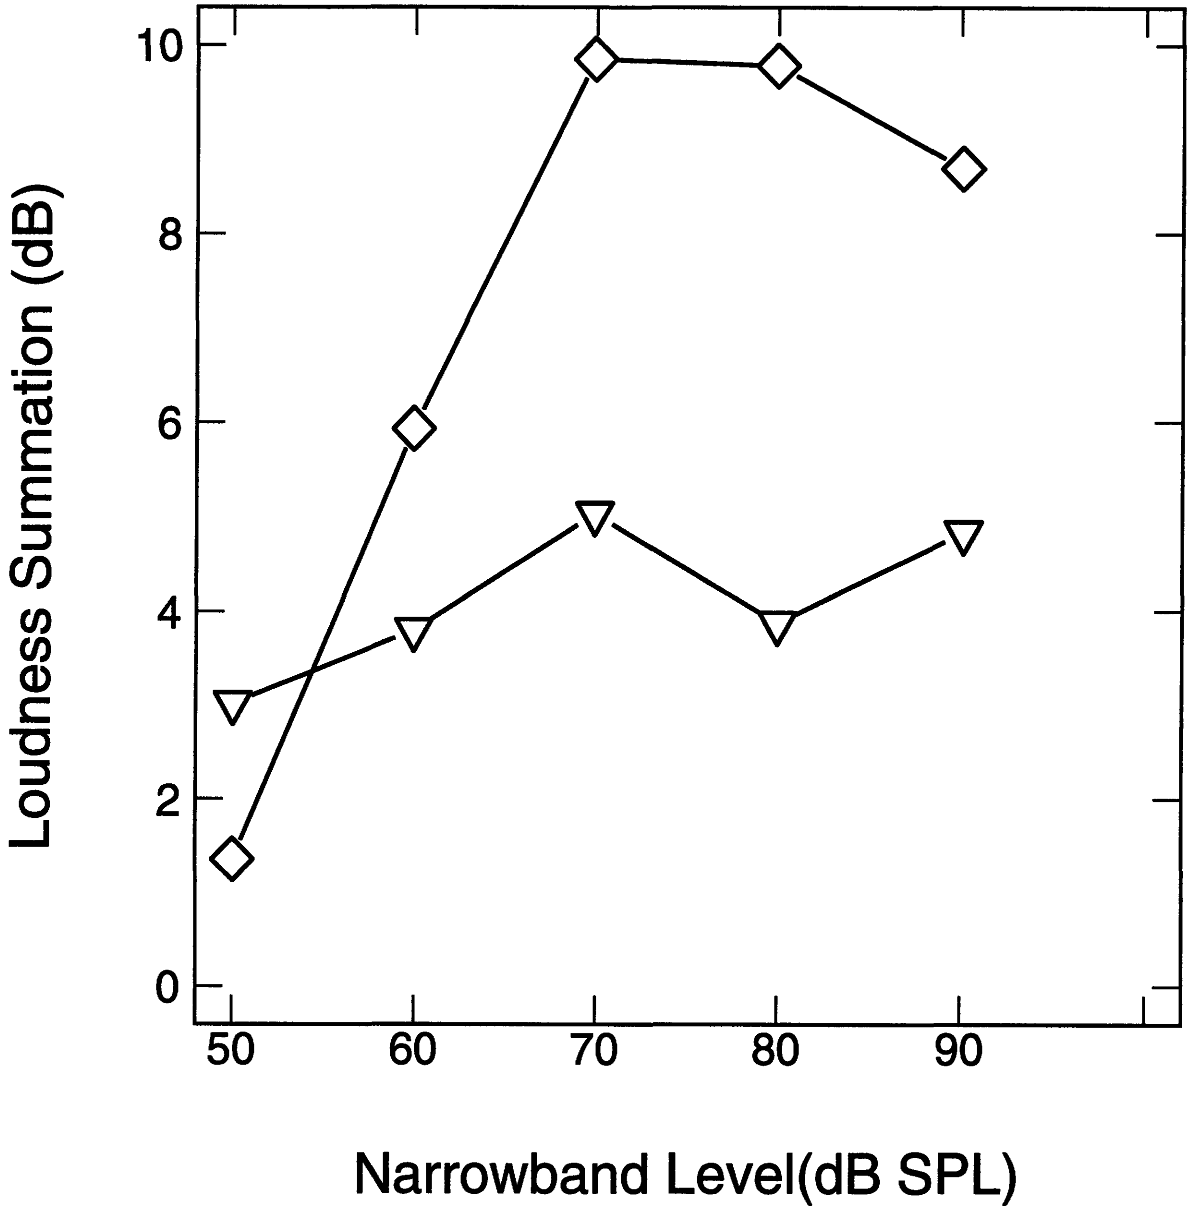

Figure 2-5: Predictions for loudness summation where recruitment is added to the model (◇) and where poor frequency selectivity in addition to recruitment is added to the model (▽). From Florentine et al., 1979.

Chapter 3

Signal Processing

3.1 Amplitude Expansion

The signal processing systems for the two types of hearing loss simulations were developed separately. The real-time amplitude expansion system that simulates loudness recruitment was developed by Lum (1995a). In Lum's implementation, the signal is first analyzed into fourteen frequency bands. The power in each band at time t is calculated by windowing each filtered signal with a non-causal 20 ms rectangular window centered at time t. The power calculations are updated every 0.45 ms. For each sample point, the output of each filter at time t is multiplied by a gain. This gain is determined from the power in the band at that point in time by the expansion input-output function (e.g., Figure 3.1). For weak input sounds, the gain is, in fact, a large attenuation. The value of the attenuation applied (in dB) decreases linearly with increasing input level inside the recruitment range. (For input levels between T_i and T_r). For levels below the recruiting range, a large attenuation is applied so that the signal is inaudible. For levels above the recruiting range, no attenuation is applied. T_i is the impaired threshold at the center frequency of the band, T_n is the normal threshold. The slope of the recruitment function is then determined by

$$T_r = \frac{tan(\alpha)(T_i - T_n)}{tan(\alpha) - 1} + T_n,$$

where

$$\alpha = 47 + 0.45(T_i - T_n)$$

in degrees. The Ariel DSP-96 digital signal processing board, with a Motorola DSP96002 DSP chip, was used for the signal processing. A detailed description of this system can be found in (Lum, 1995a; Lum, 1995b).

3.2 Frequecy Smearing

3.2.1 Real-time System

Poor frequency resolution was simulated by smearing the short time spectrum of the signal. Two systems were developed for this purpose. One operates in real-time, the other off-line. The real-time system uses hardware similar to that of the expansion system. The DSP algorithm was implemented in C, compiled via the Motorola DSP96 Optimizing C Compiler, and then downloaded to the DSP board. The code was supported by a library of subroutines written by Matthew Sexton and David Lum.

Algorithm

The smearing processing algorithm proceeded as follows (See Figure 3.2). The acoustic waveform is sampled at 11025 samples per second and windowed with an 8 ms Hamming window (88 samples). The resulting sequence is zero padded to 128 samples and the 128 point DFT is calculated. The DFT is separated into its power spectrum (magnitude squared) and phase. The power spectrum (from ω=0 to ω=π; corresponding to 64 frequency samples) is then warped by projection on to a mel-frequency axis that is linear up to 1000 Hz and logarithmic thereafter. The 64 sample warped spectrum consists of each sample of the original spectrum repeated twice up until 1000 Hz (corresponding to the first 24 samples of the warped spectrum). Above 1000 Hz, the number of samples that are repeated twice decreases with increasing frequency, while the number of samples repeated once increases with increasing frequency. Above 2500 Hz, some of the input samples are discarded. This warping is done to mimic the characteristic frequency distribution along the basilar membrane in the cochlea.

The mel-warped power spectrum is smeared by convolving it with a Gaussian-like function. This convolution is accomplished by multiplying the DFT's of the smearing function and the warped spectrum and taking the inverse DFT of the product. The smearing function is of the form

$$G(m) = \begin{cases} e^{-\pi(m/B_H)^2} & 0 \leq m < 64 \\ e^{-\pi(m/B_L)^2} & -64 < m \leq 0, \end{cases}$$

where B_L and B_H determine the width of the smearing function at low and high frequencies, respectively, and m refers to the discrete samples along the mel-frequency axis. B_L and B_H thus correspond to excessive upward and downward spread of excitation respectively.[1]

Since the smearing function is of length N=127 samples, and the power spectrum $(0 < \omega < \pi)$ is of length 64 samples, the convolved power spectrum is 64+N-1=190 samples in length. Only the 64 samples starting from the sample $(N+1)/2 = 64$ correspond to meaningful frequencies $(0 < \omega < \pi)$. These samples are then projected back to a linear frequency axis and scaled so that the total power is equal to that of the unprocessed

---

[1] It is hoped that this smearing will make auditory nerve fibers in the normal listener more responsive to off-CF sounds relative to on-CF sounds than normal (since off-CF energy is shifted to on-CF and on-CF energy is shifted off-CF); for sounds with spectral contrast, this should produce a broader response along the basilar membrane than normal, thereby simulating the supposed broader excitation patterns of the hearing impaired.

segment. The projection on to a linear axis was done by means of a lookup table created by hand that maps each of the 64 warped frequency samples to the corresponding unwarped frequency sample that is it's best match. For the low frequencies, many of the warped samples were discarded and for the high frequencies, many of the warped samples were repeated. Finally, the smeared, scaled power spectrum is recombined with the original (unprocessed) phase.[2] The inverse DFT of the resulting complex signal is then calculated to produce the time signal, and the first 88 samples of this time signal are used. (The extra 40 samples of the 128 point DFT are discarded). The 8 ms time window is shifted every 4 ms[3] and the frequency-smeared outputs are recombined via the overlap-add method. In an effort to reduce noise in the system arising from the A-D converter, when the smearing processing is first run, the system determines the DC value of the input it receives in the absence of external inputs. It then subtracts this DC value from all subsequent input waveforms (suggested by Lum) to reduce noise output of the smearing processor.

Characterization of the Processing

The frequency smearing processing was found to produce no noticeable distortion for speech or tonal inputs when $B_L = B_H = 0.001$ (essentially no smearing). As the smearing bandwidth increases, processed speech becomes increasingly muffled. The effects of applying different smearing bandwidths on the vowel /ah/ can be seen in Figure 3.3. It can be seen that spectral contrast is reduced as smearing bandwidth is increased.

Real time implementation of the frequency smearing processing obviates the need for pre-processing a large number of waveforms off-line. However, this processing produced undesired artifacts in the smeared signal when the input signal was a tone (a stimulus

---

[2] It is not clear how the phase should be altered to mimic the effect of the hearing impaired. Since the purpose of this study was to look at effects of a broader response pattern on the basilar membrane, no modification to the phase was implemented. This had effects on the smearing processing, however (see section 'Characterization of the Processing').

[3] The frame rate was 4 ms instead of 2 ms used by Baer and Moore , because of computational constraints imposed by the real-time processing.

used in many of the psychoacoustic experiments). The artifacts appear to be due to the linear phase spectrum that results from windowing a tone with a linear phase Hamming window. When the magnitude spectrum of the tone is smeared and then recombined with the linear phase spectrum, the energy in the resulting signal is very concentrated in time. When these time signals are overlap-added, the resulting time signal has a strong 'pitch' at the frame rate of the processing. In addition, if the period of the tone is a submultiple of the frame period, an additional artifactual periodicity (with the period of the frame rate) is introduced in the smeared signal. Thus the real time processing can be used for speech processing, where the windowed short-time spectrum does not have linear phase and is not periodic, so that the artifactual pitch problem and the frame-rate periodicity problem are absent, but it cannot be used to process tones. Figures 3.4 and 3.5 show the effects of applying the real-time smearing algorithm to a 1000 Hz tone. The artificial periodicity with the period of the frame rate can be seen in both the time and frequency domains.

The smearing processing has several potential drawbacks. One issue is that the actual smearing of the spectrum is less than is prescribed for by the bandwidth of the smearing Gaussian. This is due to the effects of overlapping the smeared short time segments that correspond to smeared power spectra combined with the unprocessed phase spectra. The resulting time signal has a short time spectrum that is less smeared than it would be if the unprocessed short time spectrum was smeared in isolation with the Gaussian function. The spectrum of the smeared waveform calculated with a 20 ms window (Figure 3.3) is similar to what is expected, but the spectrum calculated over an 8 ms window often exhibits more spectral detail than expected (Figures 3.6-3.10). The prescribed magnitude spectrum is shown by the dotted line in theses figures, and it changes little from frame to frame. The actual magnitude spectrum of the smeared waveform is much different than the prescribed magnitude spectrum, and it varies from frame to frame. The tendency of smeared spectra to retain detail was observed by Baer and Moore, who also used overlap adding of smeared short time power spectra recombined with the unprocessed short-time phases.

An additional issue of the smearing processing is that the gain in rms level of the output waveform relative to the input waveform was not specified. Although the total energy in each smeared frame equalled that of its input frame, the overlap adding removes this equality for the resulting output signal. It was observed that for a tone, the gain was unity when $B_L = B_H =.001$ and the gain dropped to -4 dB when $B_L = B_H = 4.0$.

The frequency response of the smearing system is also not very well defined. For example, a tonal input (with a impulse spectrum) will result in an output signal with a Gaussian spectrum, as desired. However, a noise input with a flat spectrum, will result in an output spectrum that peaks in the mid-frequencies (2750 Hz) and that tails off at the low and high frequencies. This spectral shaping is caused by the convolution of the smearing function with the flat input spectrum. (There are no frequencies below 0 Hz or above 5500 Hz with which to convolve the Gaussian). Thus for a noise input, the output spectrum will not be flat, although a flat output spectrum would presumably be desired. Figure 3.11 shows the output spectrum obtained by passing a flat spectrum through the frequency warping, Gaussian smearing, and frequency unwarping procedures. The drop in magnitude is more pronounced in the high frequencies than in the low frequencies because the smearing was done on a mel frequency axis, instead of a linear frequency axis.

3.2.2 Offline Frequency Smearing

Excitation Pattern Model

To avoid the distortions that result from processing sinusoids via the real-time system, an alternative type of processing was developed. The goal of this processing is to produce a signal that evokes the same excitation pattern in the cochlea of a normal hearing listener as the unprocessed signal evokes in an impaired ear (see Section 2.3). Excitation patterns were calculated based on a model developed by Moore, using auditory filters derived in experiments in which the masker was a noise with a spectral notch and the smeared signal was a tone. The auditory filter shapes are assumed to have the form:

$$W_c(f) = \begin{cases} (1+p_l f)e^{-p_l f} & 0 < f \leq c \\ (1+p_u f)e^{-p_u f} & c < f, \end{cases}$$

where $W_c(f)$ is the magnitude of the filter with center frequency c as a function of frequency. The scaling factors p_l and p_u determine both the bandwidth and the slope of the lower and upper skirts of the auditory filter, respectively. The values of p_l and p_u depend on the center frequency of the filter c, and the higher the values of p_l and p_u, the more sharply tuned the filter. Specifically, p_l and p_u are inversely proportional to the ERB (equivalent rectangular bandwidth) of the auditory filter (where $ERB = 24.7(0.00437c + 1)$), so the filters get less sharp as the center frequency increases. This seems to compensate for the fact that the filters are not defined along a logarithmic frequency axis, but rather a linear one. Figure 3.12 shows an example of a normal auditory filter and an impaired auditory filter with p_l and p_u decreased by a factor of 3.

In the alternate smearing processing system (implemented in Matlab), signals are analyzed by 500 auditory filters with the $W_c(f)$ shape, with $10 < c < 5000$, spaced every 10 Hz. The excitation pattern is the array of output powers of the 500 filters, where the output power for the filter of center frequency c is the integral of the power spectrum of the signal times $W_c(f)$. (See Figures 3.13-3.14 for the excitation patterns to a 2000 Hz tone for a normal ear and an impaired ear).

Algorithm

For a given stimulus, the smearing system attempts to derive a new stimulus that produces the same excitation pattern in the normal ear as the unprocessed stimulus produces in the hearing impaired ear. This excitation pattern E_i is equal to $H_i S$, where S is the input power spectrum (a 500 dimensional vector corresponding to frequencies from 10 to 5000

Hz), and H_i is the matrix of the 500 hearing impaired auditory filters (subscript i refers to the hearing impaired). Each row of H_i contains the samples of the magnitude of the impaired auditory filter for each of the 500 frequencies. The first row of the matrix has its center frequency at 10 Hz, the second row has its center frequency at 20 Hz and so on. The auditory filters of the normal ear are also calculated and put in a corresponding matrix H_n, where subscript n refers to normal listeners. The processed signal should thus have the power spectrum $\hat{S}$, whereby $H_n\hat{S} = E_i$, the excitation pattern produced in the ear with widened auditory filters. (The excitation pattern produced in the normal ear by processed stimulus $\hat{S}$ produces the same excitation pattern E_i that is produced by the unprocessed stimulus S in the impaired ear). Thus $\hat{S} = H_n^{-1}H_iS$. The resulting smeared power spectrum is then upsampled by a factor of 10 and convolved with a sinc function of length 20. This provides a 5000 sample smeared power spectrum instead of a 500 sample spectrum. [4] The resulting power spectrum is then square rooted to obtain the magnitude spectrum, and a phase is applied whereby the value of the phase for each sample of the DFT is generated randomly from a uniform distribution between -π and π. The Fourier transform is completed for the 5000-10000 Hz frequencies by concatenating 5000 samples of the original signal with flipped magnitude and negative phase on to the obtained 5000 sample signal. An IDFT is then taken, so that the resulting time signal is 10000 samples (1 second) long. Several stimuli were generated in this way with identical magnitude spectra. The phases of the different stimuli were generated separately, however.

In the case of a sinusoidal or narrowband noise input, the smeared signal is generally a narrowband noise and it has the property that its short time spectrum is consistent with the desired smeared spectrum at all points in time. The lower two plots in Figure 3.15 and 3.16 show the power spectra of a tone that has been smeared with filters whose values of p_l and p_u are $\sqrt{3}$, and 3 times smaller than normal. The top plots show the resulting power

---

[4] Though using a 5000 sample power spectrum would have avoided the convolution with the sinc operation, performing calculations with matrices of dimensions 5000 × 5000 was too computationally intensive.

spectrum when p_l is 3 times smaller than normal and p_u is normal. Figure 3.17 shows the power spectrum of smeared narrowband noise, where the unsmeared signal was a 200 Hz band centered at 1200 Hz and the values of p_l and p_u were 3 times smaller than normal.

The smearing processing for all tones and noises was carried out separately, even for the simultaneous masking experiment, where the noise and tone were presented together. This was done so that tonal stimuli could be attenuated separately from the narrowband noise in the adaptive psychoacoustic experiments (see section 4.1).

To combine the frequency smearing processing with amplitude expansion (simulated recruitment), the pre-processed smeared waveforms were passed through attenuators that were programmed to provide an appropriate attenuation for a given nominal sound level. The formula for the attenuation to be applied was similar to that used by Lum in his experiments with amplitude expansion. Specifically, the attenuation applied was $(80 - \mathit{input})/2$. This prescribes a 40 dB attenuation for a 0 dB SPL sound, where the attenuation linearly decreases with level (in dB) such that no attenuation is applied to input sounds of 80 dB or more. This simulates mild/moderate sensorineural hearing loss, where the threshold is raised about 40 dB and loudness grows twice as fast as normal.

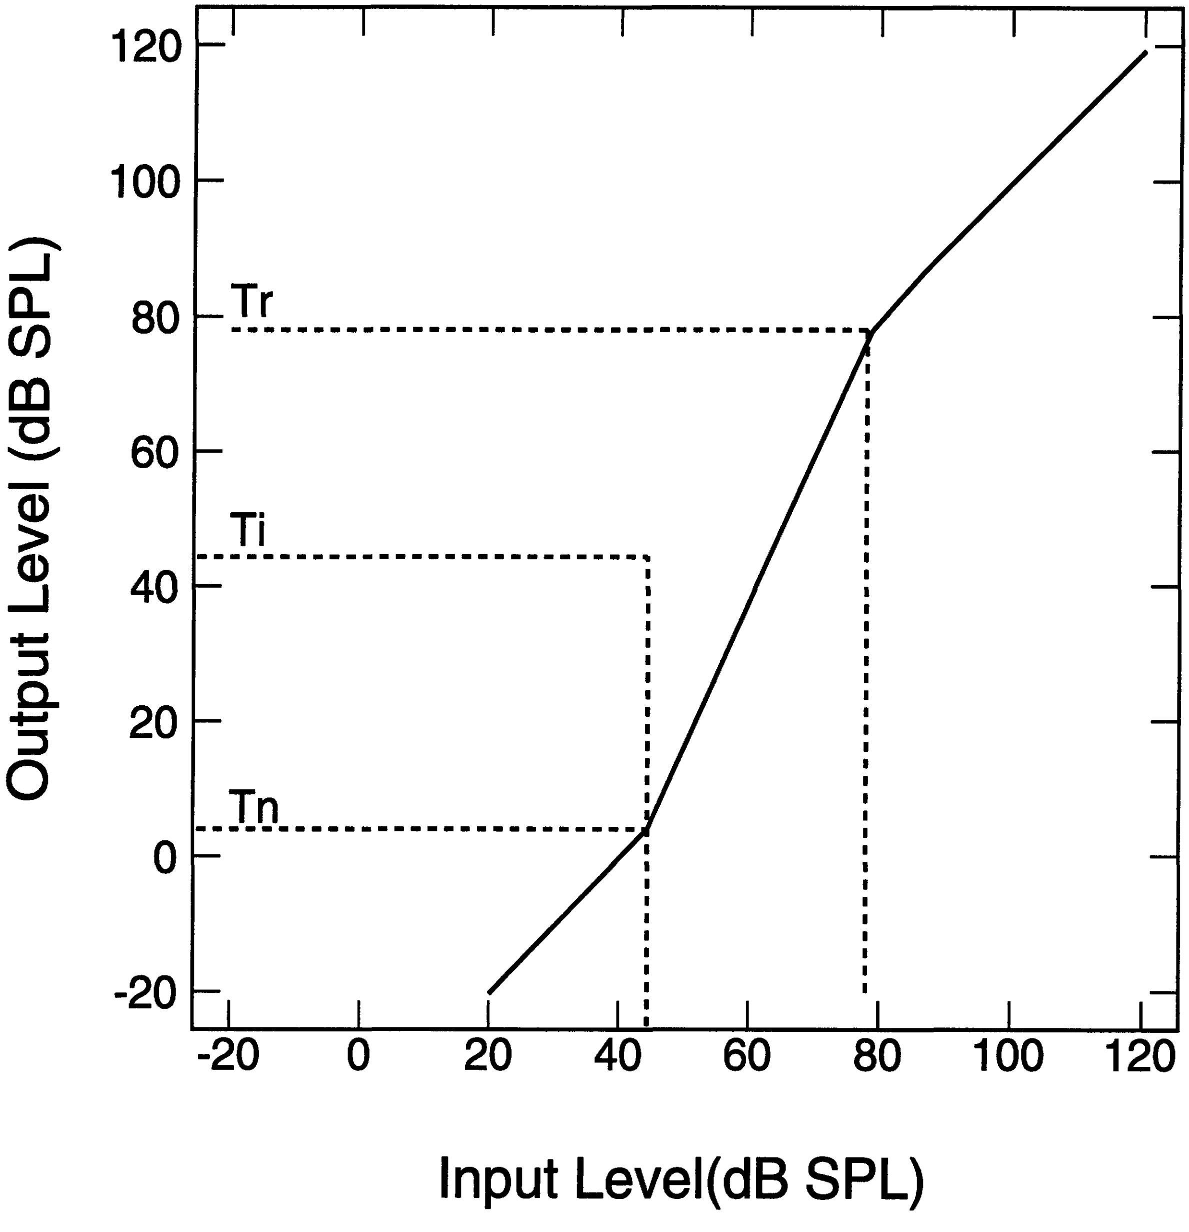

Figure 3-1: Example of an input-output function of the expansion simulator. Normal threshold is marked by T_n, impaired threshold by T_i. The threshold of complete recruitment is T_r.

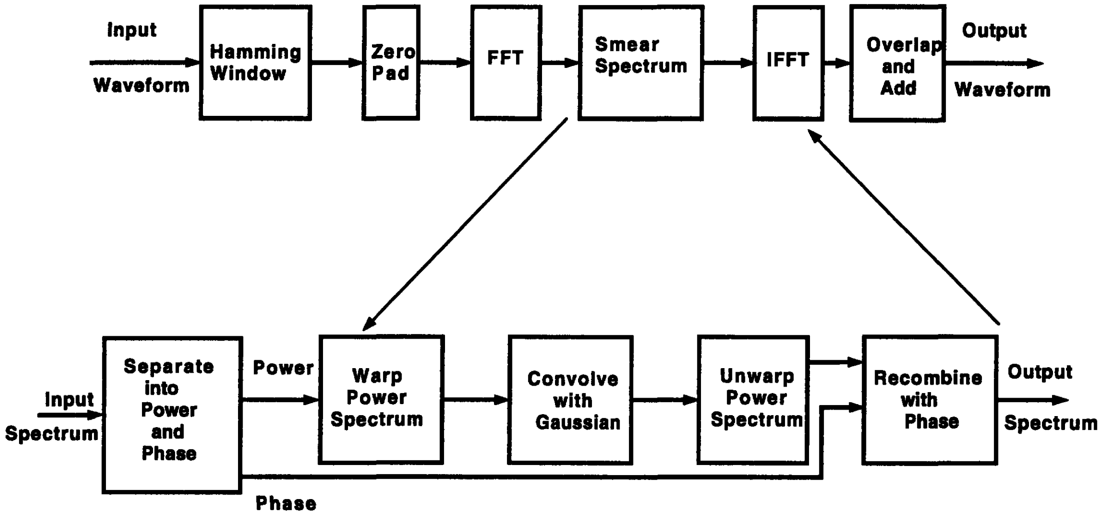

Figure 3-2: Block diagram of real time frequency smearing processing. The 'smear spectrum' block from the upper diagram is written out in detail in the lower diagram.

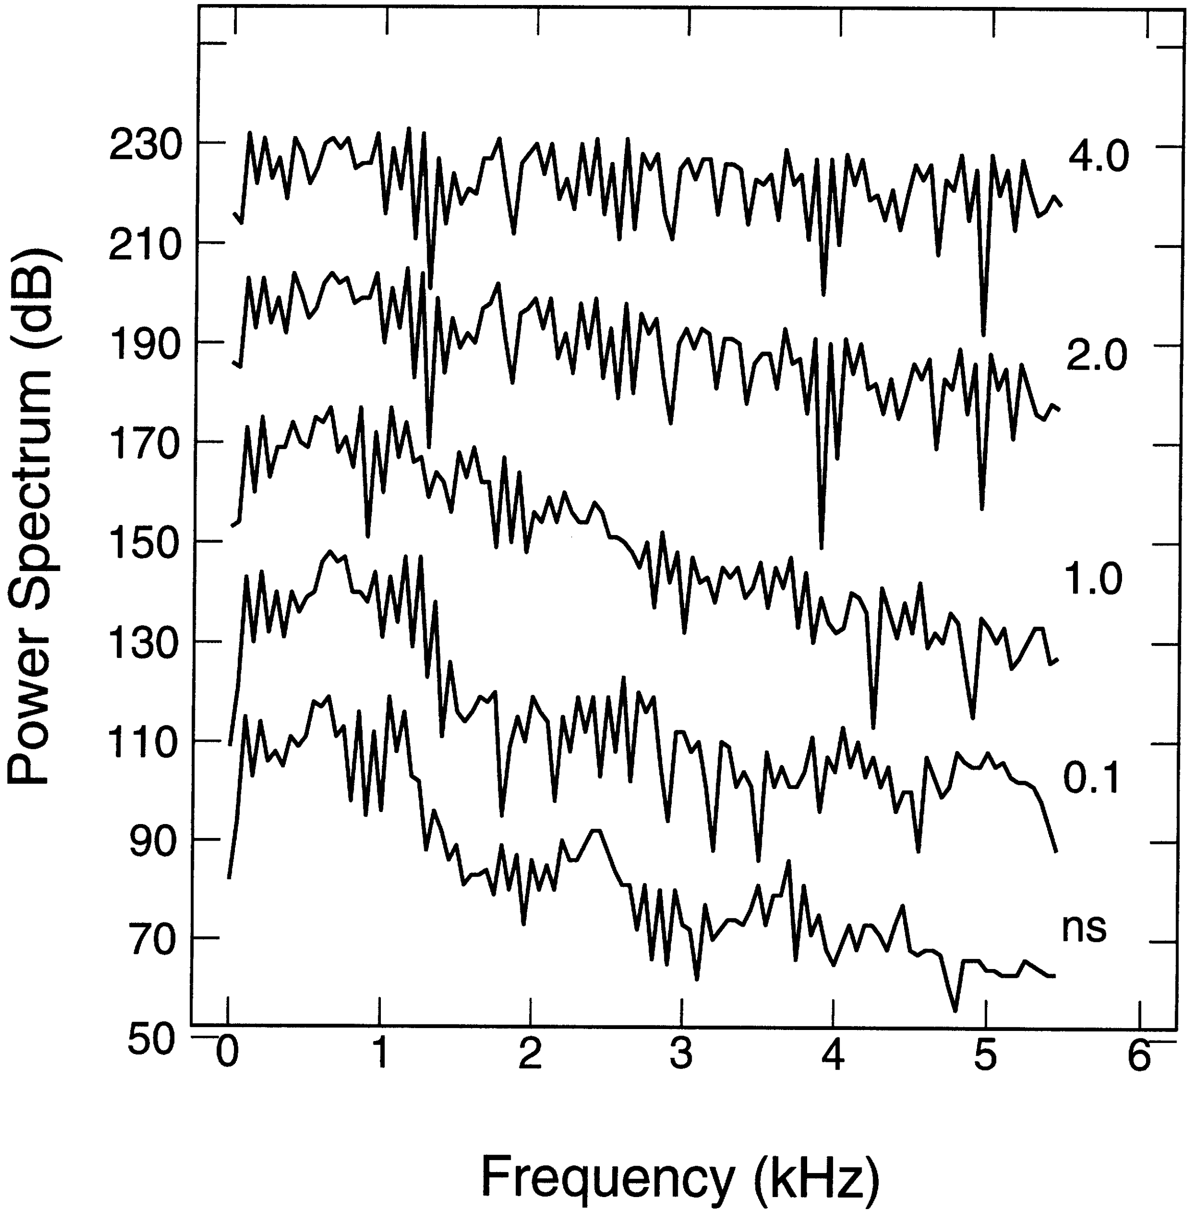

Figure 3-3: Power spectra of a 20 ms segment of the vowel /ah/, windowed with a rectangular window, and smeared with different smearing bandwidths. The bandwidth of the smearing function is indicated as a parameter ($B_L = B_H$), (ns = no smearing). Successive plots were offset by increments of 20 dB.

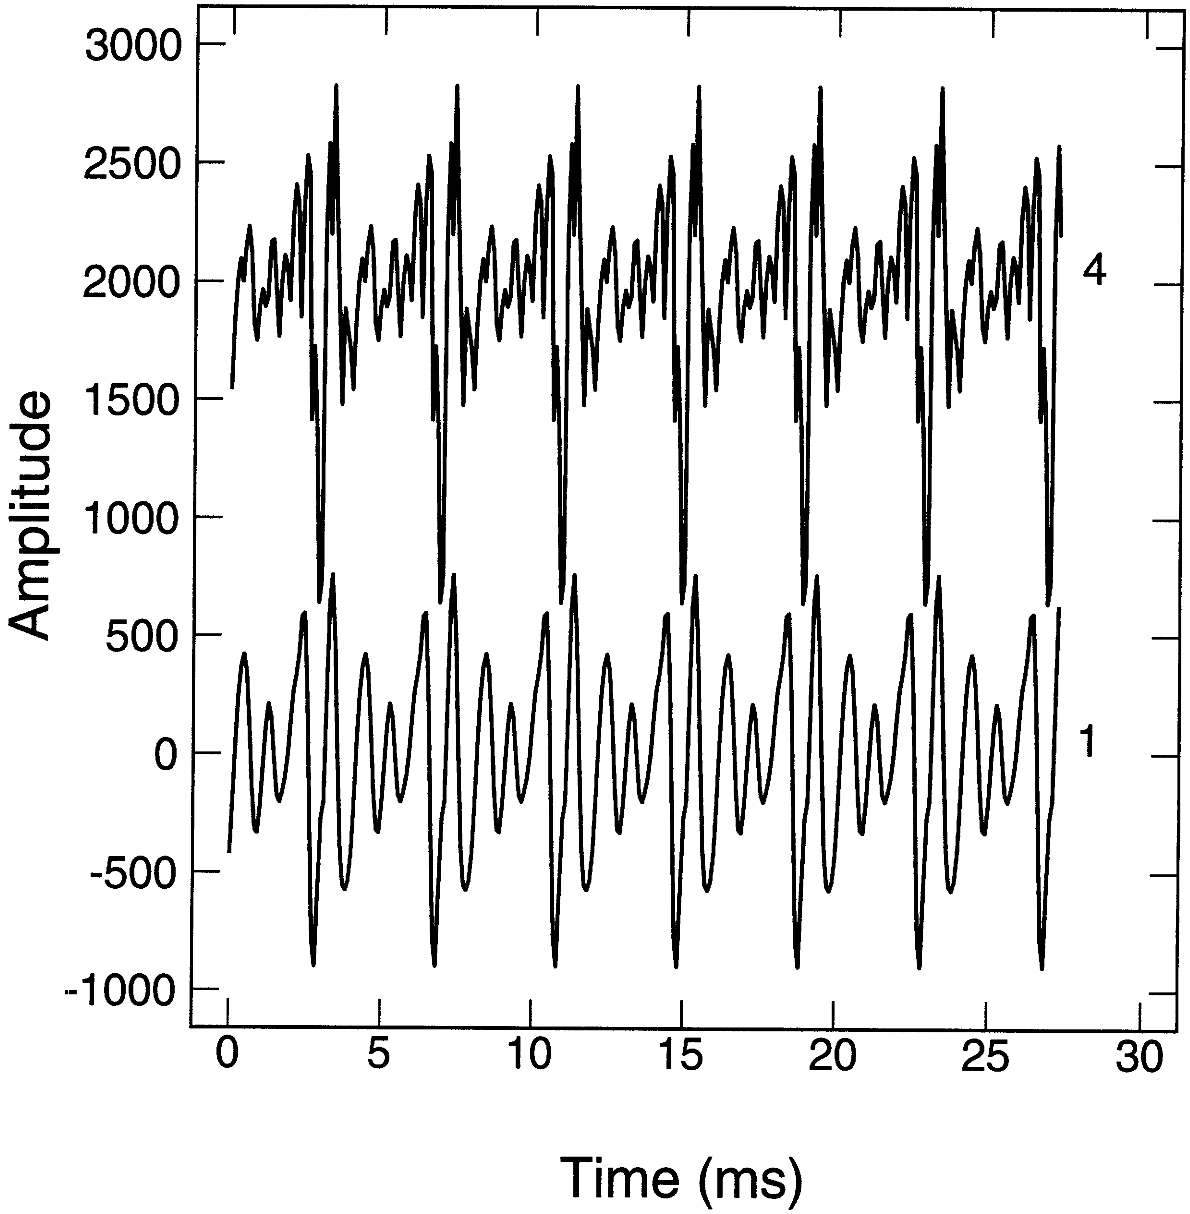

Figure 3-4: The time waveform of a 1000 Hz tone after being processed by the real time smearing system. Top waveform: B_L = B_H =4. Bottom waveform: B_L = B_H = 1. Amplitude is in arbitrary units. Top wavefrom is offset by 2000 units.

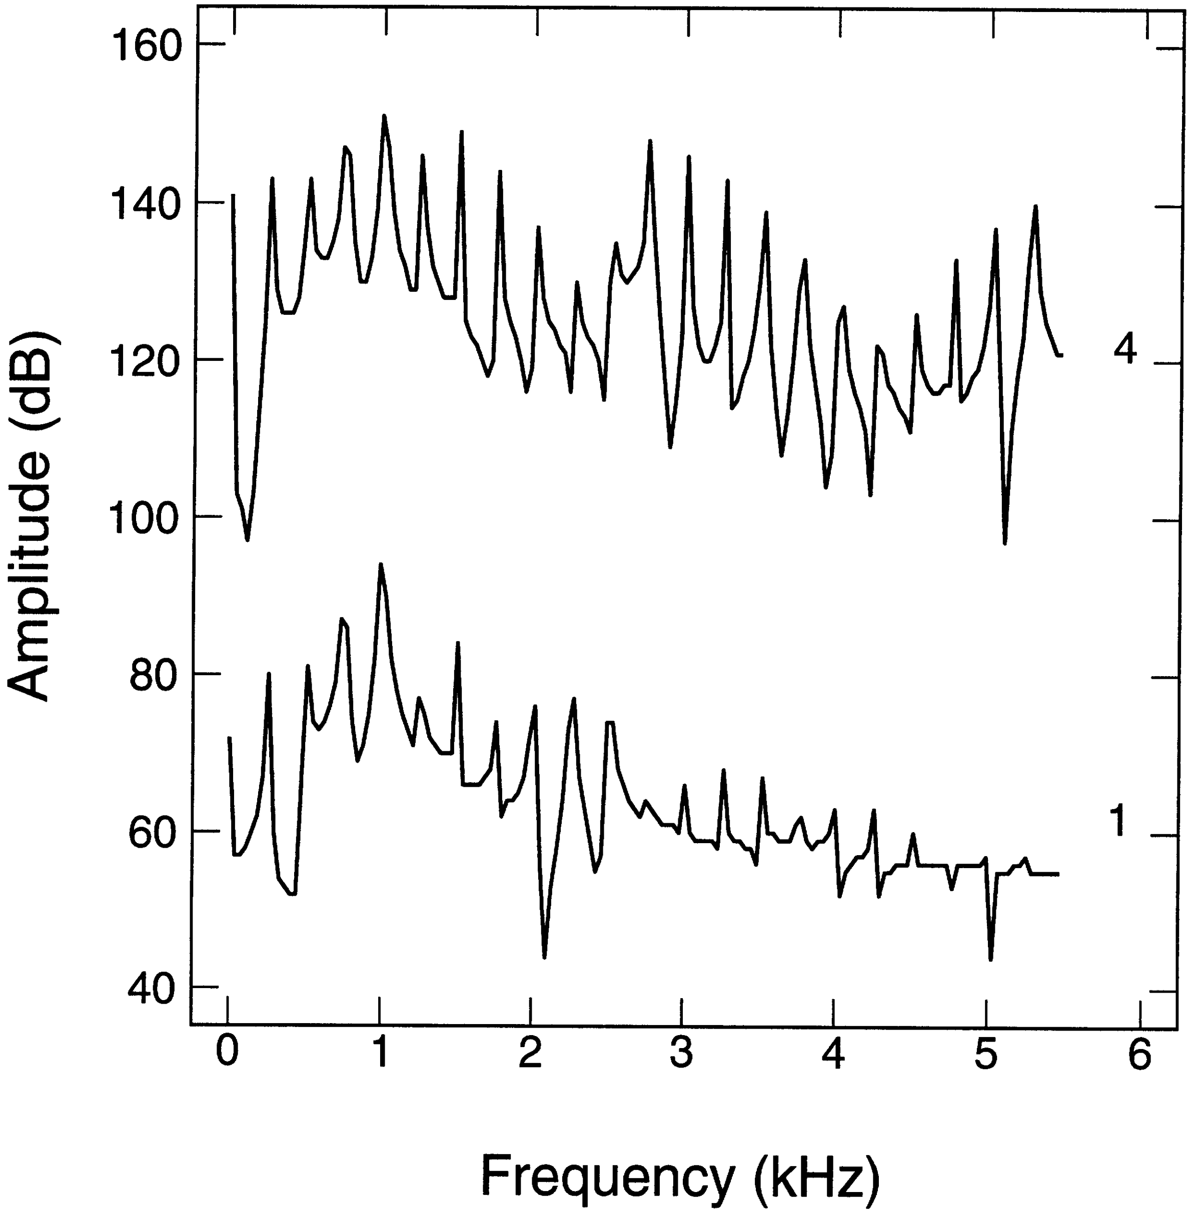

Figure 3-5: The power spectrum of a 1000 Hz tone after being processed by the real time smearing system. Top spectrum: $B_L = B_H = 4$. Bottom spectrum: $B_L = B_H = 1$. Top spectrum is offset by 60 dB.

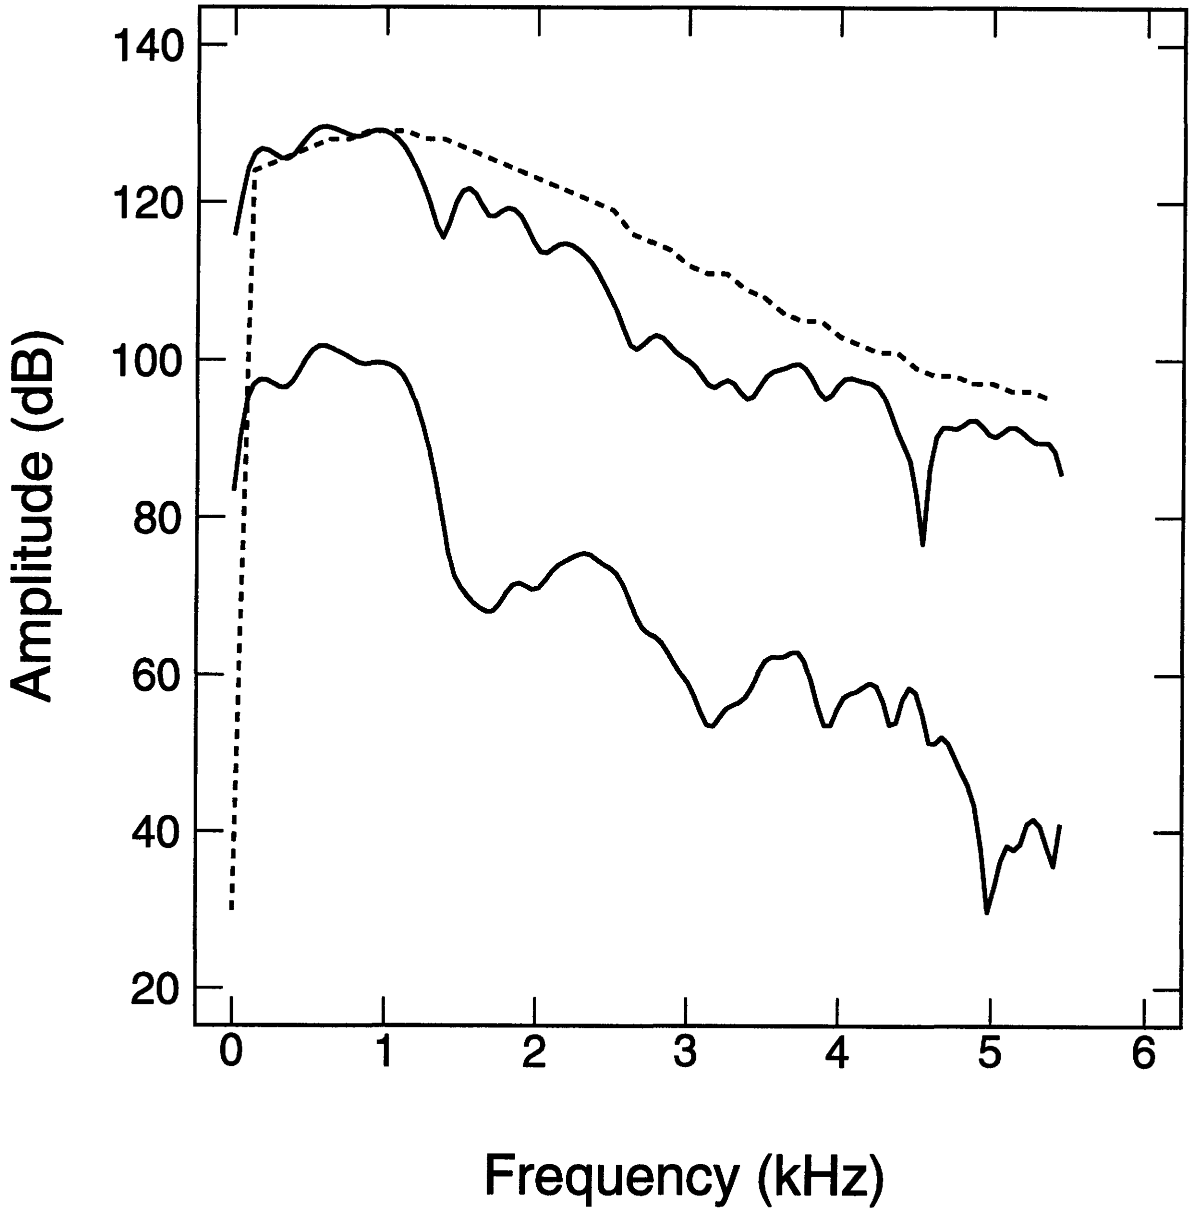

Figure 3-6: Example of the effect of the overlap add method on the synthetic vowel /ah/. Bottom Figure: Short time spectrum of an 8 ms segment of the unprocessed vowel /ah/, winowed with a Hamming window. Top Figure: Short time spectrum of the same 8 ms segment as in top figure, after the signal passed through the smearing and overlap adding processing. ($B_L = B_H = 1$). The dotted line shows the smeared spectrum that is prescribed by the smearing algorithm, without the effects of the overlap adding. Top spectrum is offset by 30 dB.

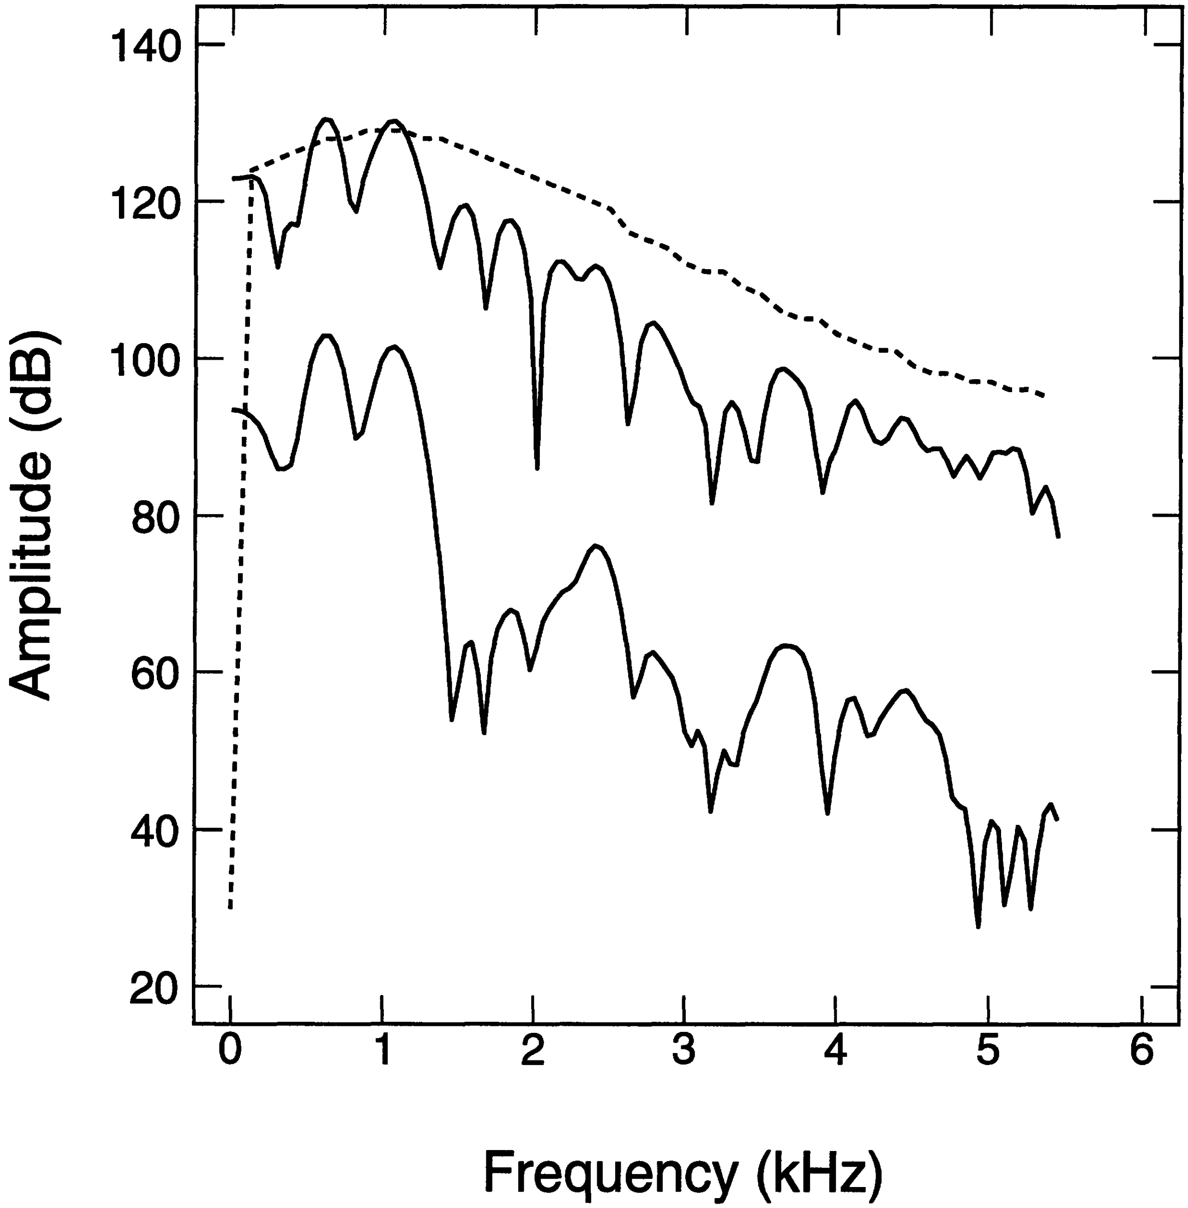

Figure 3-7: Same as Figure 3.6. One frame (4 ms) later.

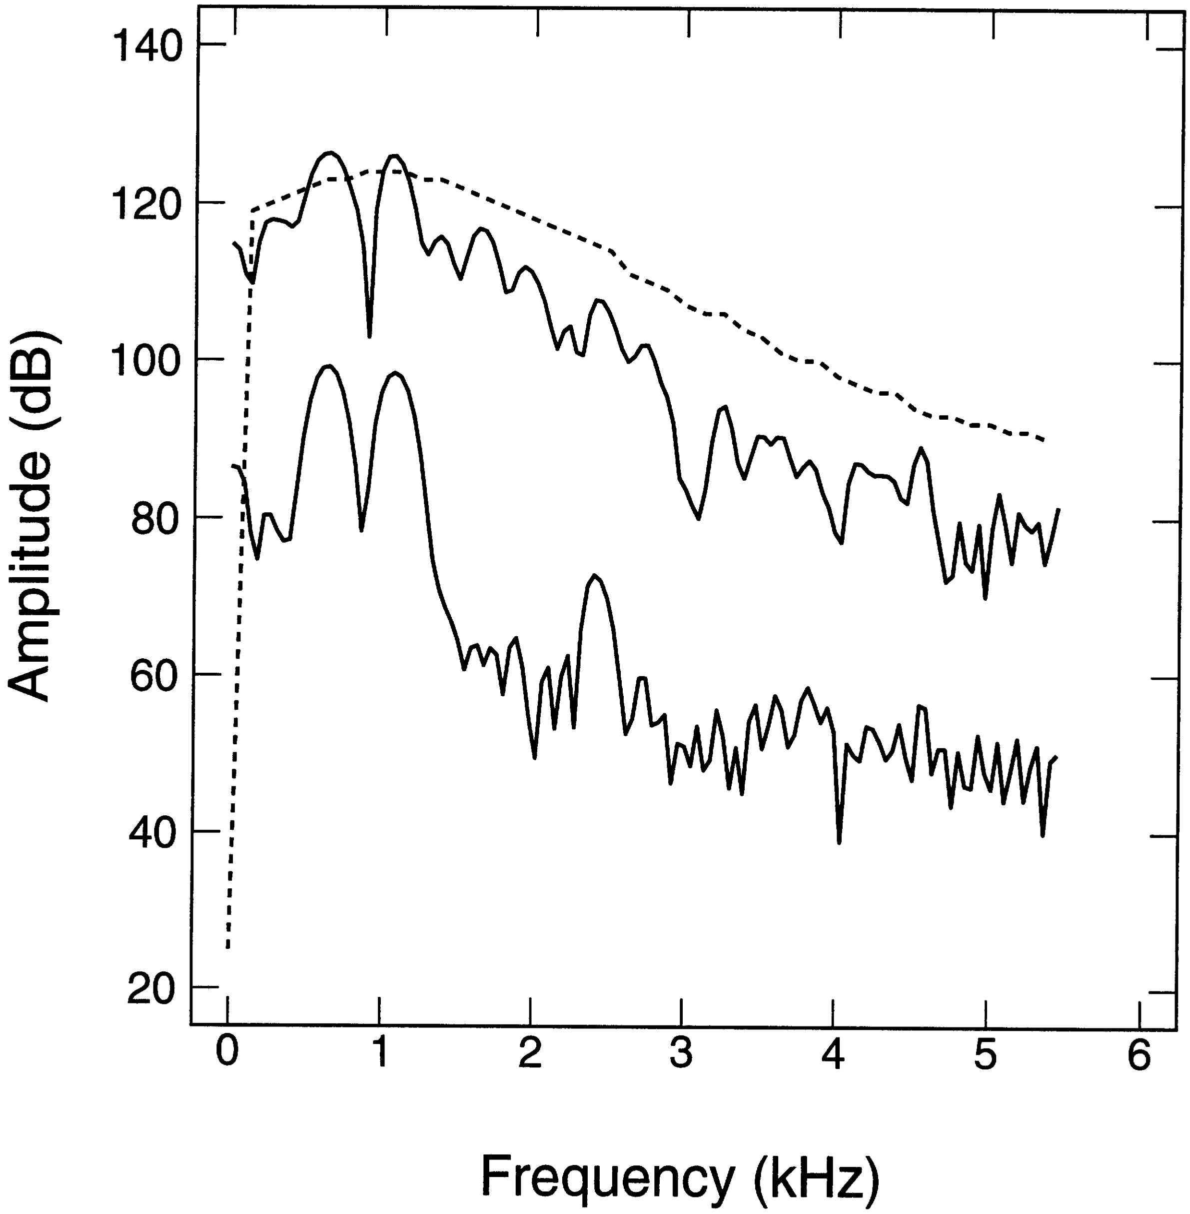

Figure 3-8: Same as Figure 3.7. One frame (4 ms) later.

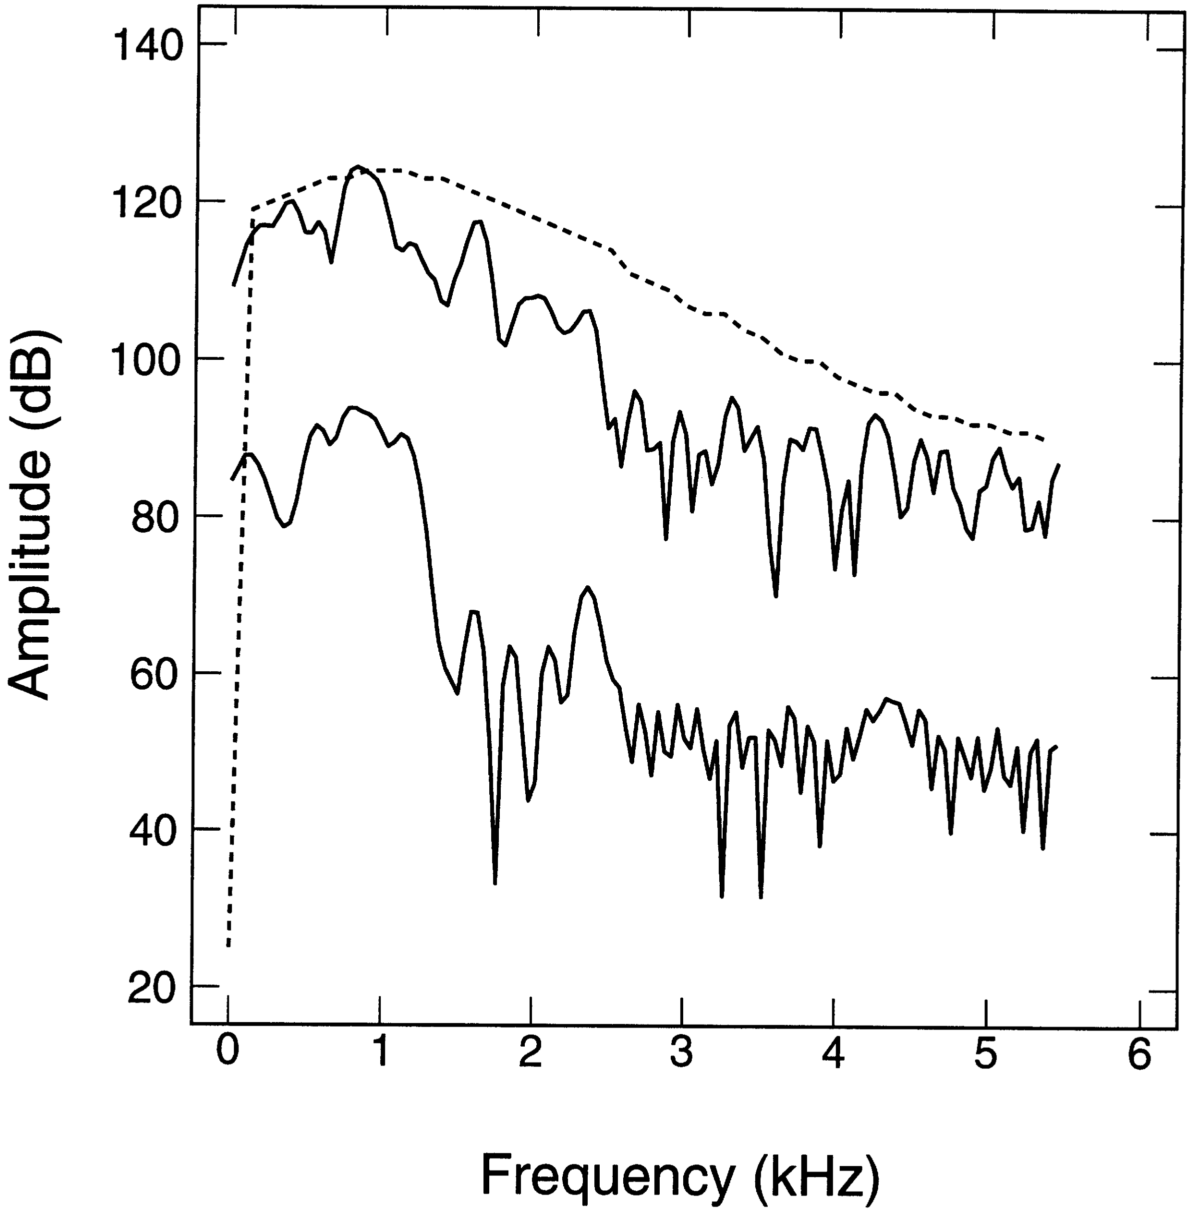

Figure 3-9: Same as Figure 3.8. One frame (4 ms) later.

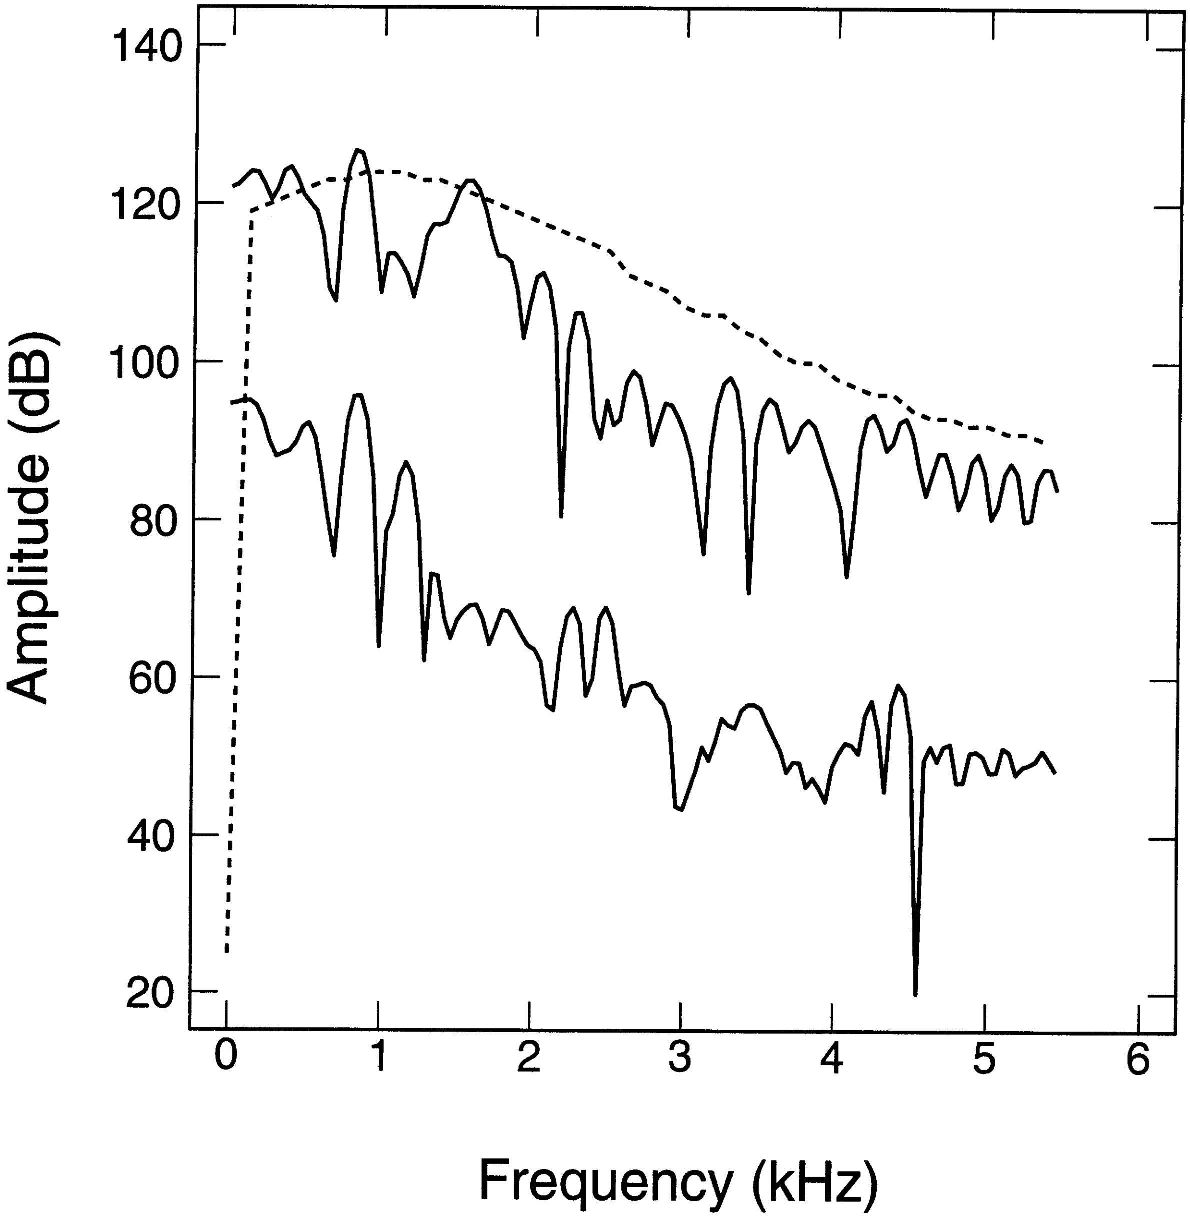

Figure 3-10: Same as Figure 3.9. One frame (4 ms) later.

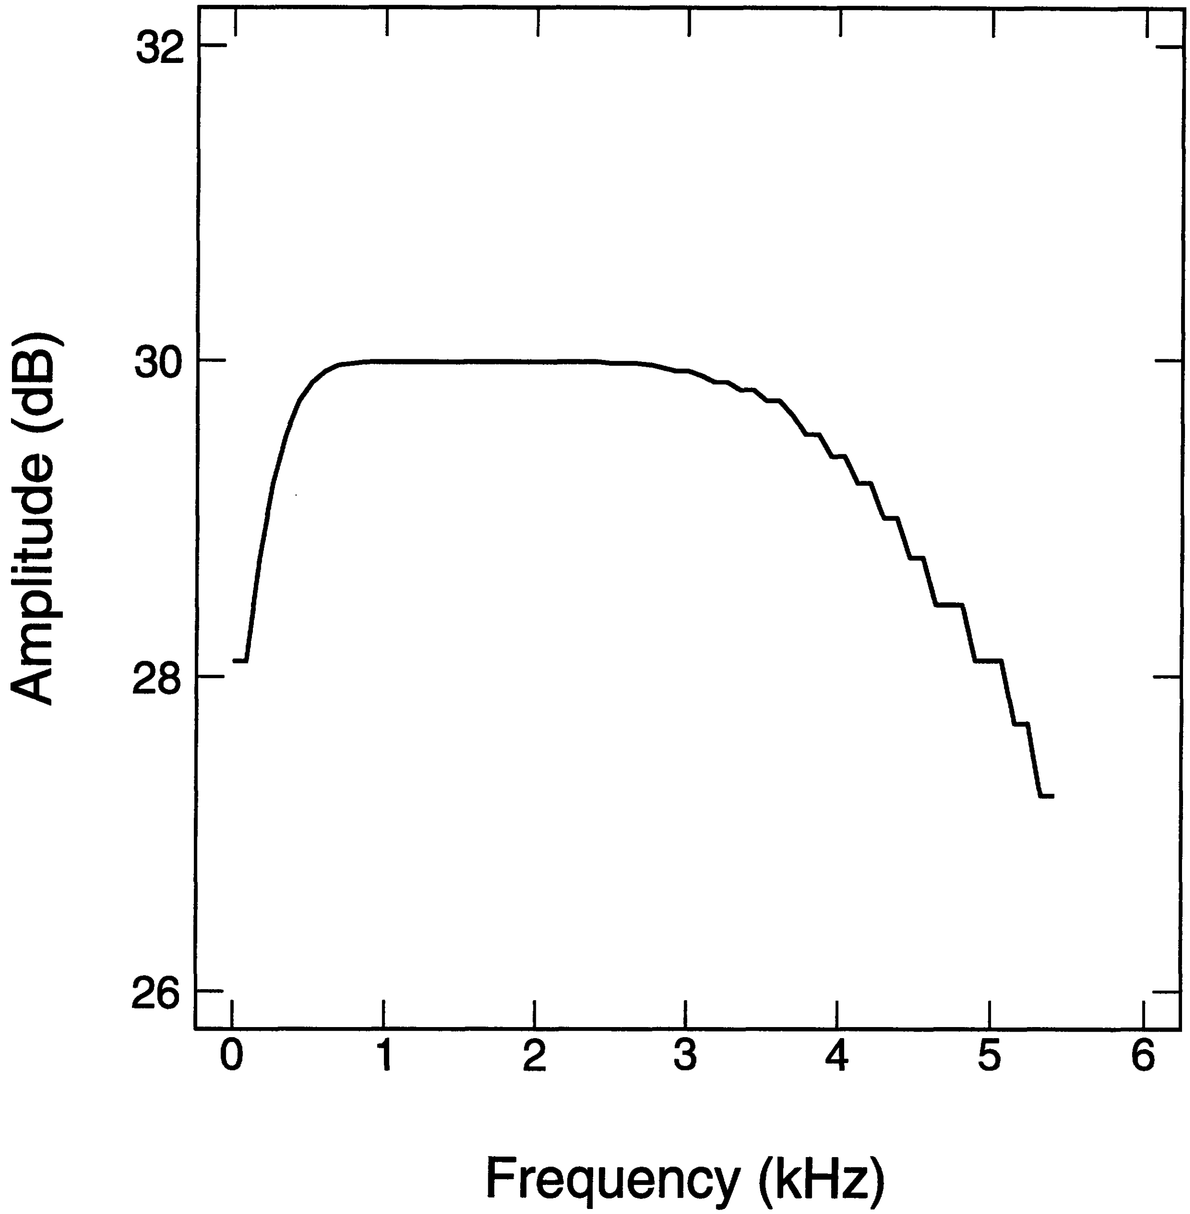

Figure 3-11: Result of passing a flat power spectrum through the frequency warping, smearing, and unwarping processes. The resulting power spectrum tails off at the edges due to the convolution of the Gaussian with nonexistent frequencies. The high frequency side is ragged due to the deleting and repeating of spectral samples in the warping procedures.

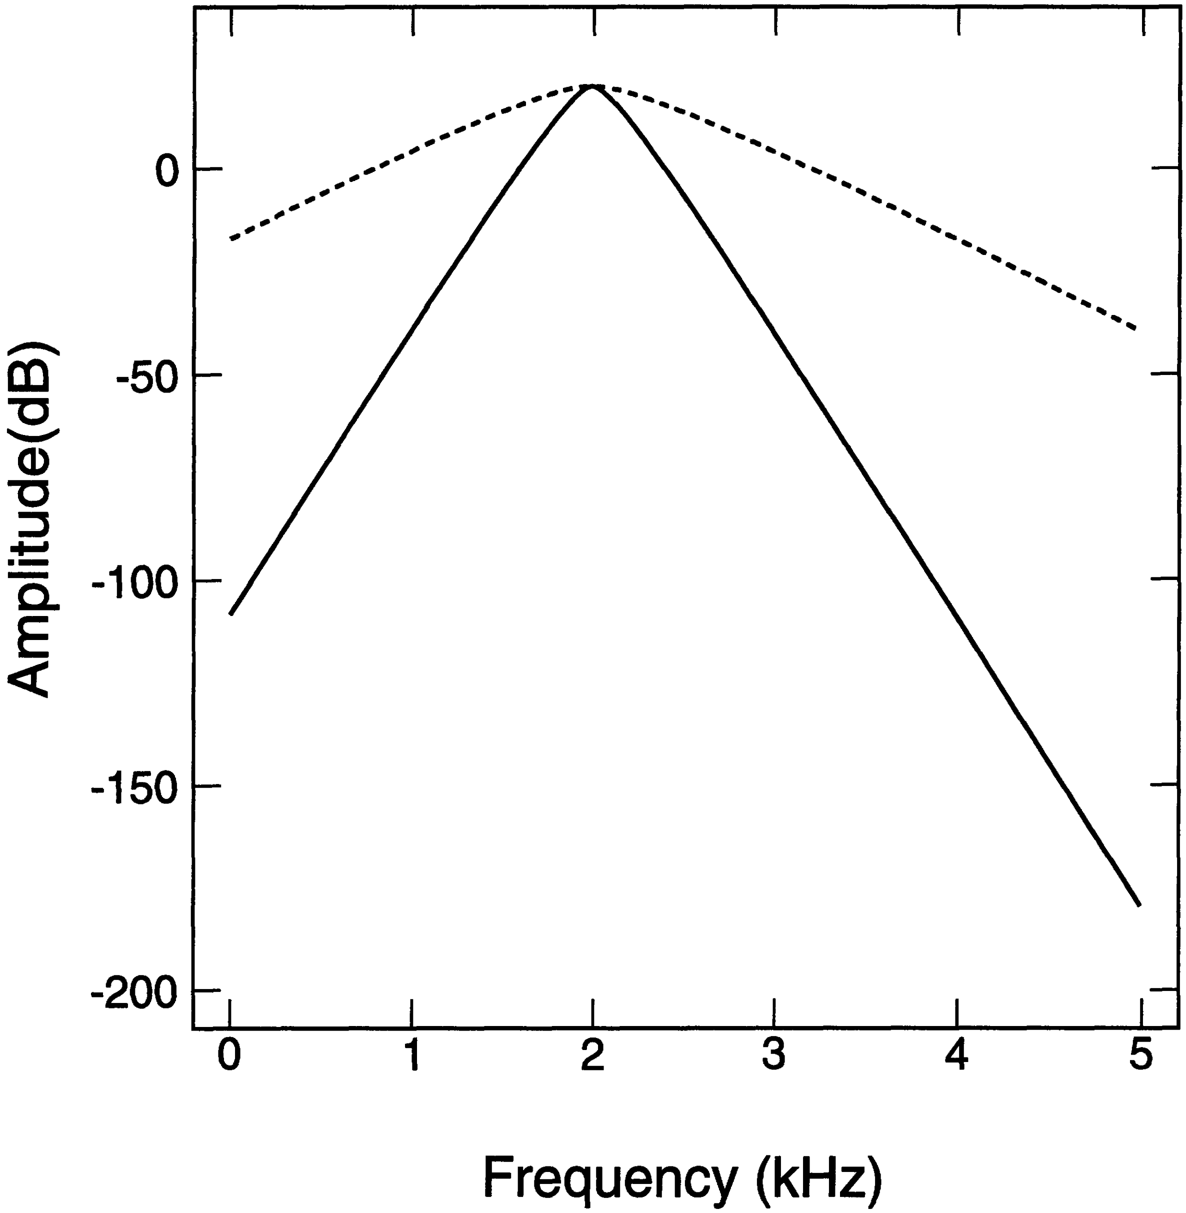

Figure 3-12: Comparison of normal auditory filter (solid line) with an auditory filter with symmetric broadening by a factor of 3 (dotted line). Center frequency of filter is 2000 Hz. Filters are of the roex form.

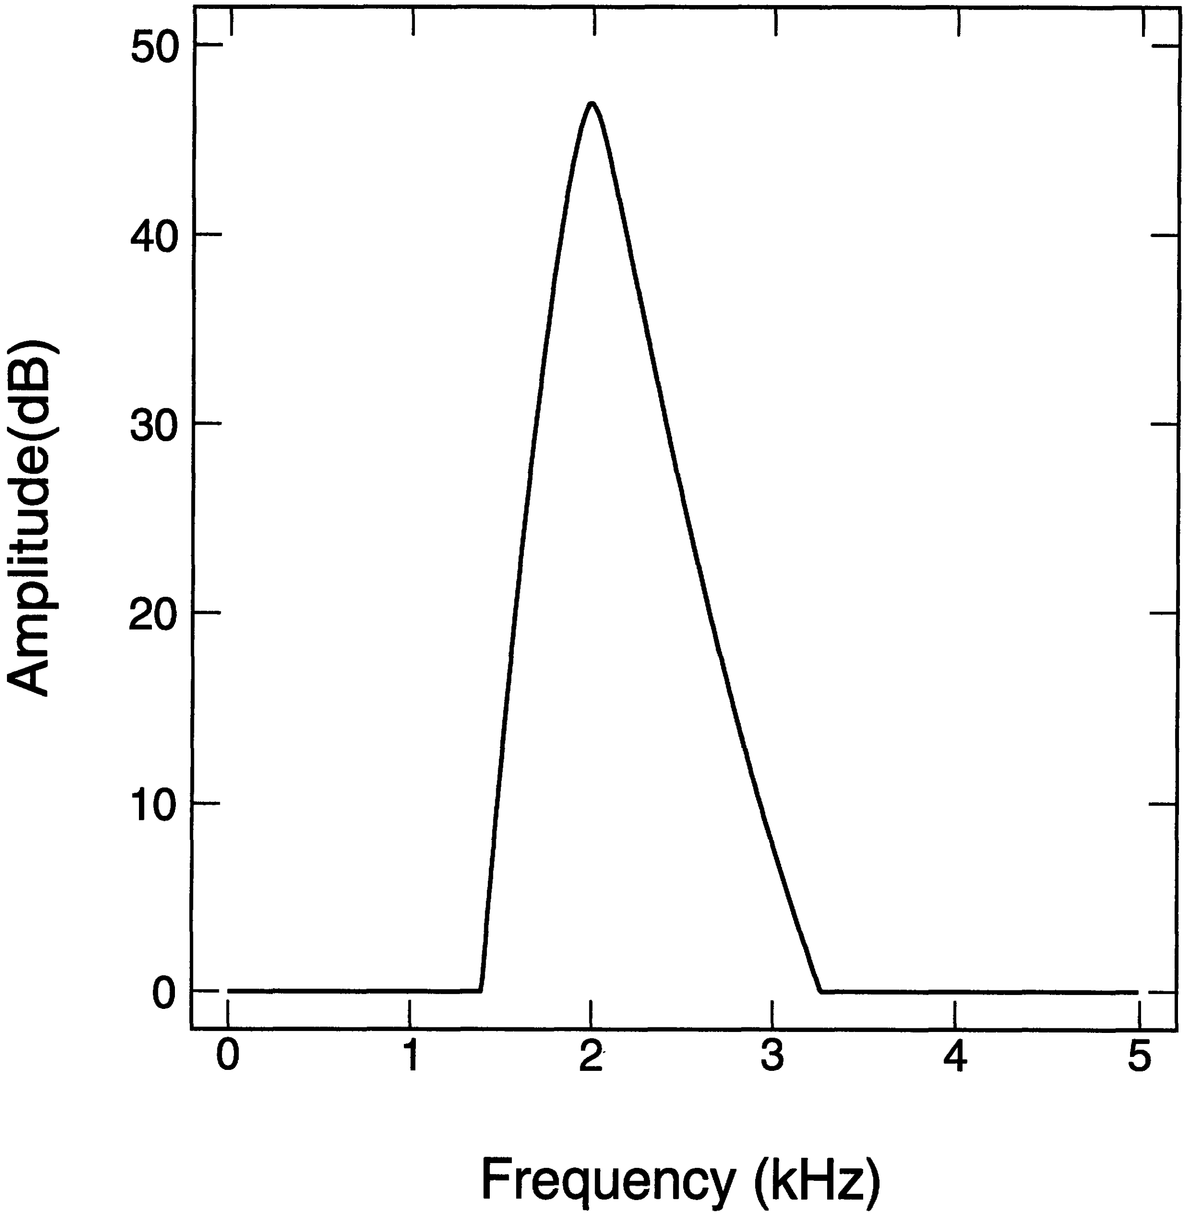

Figure 3-13: Excitation pattern for a 2000 Hz tone for a normal ear (Arbitrary scaling). Values below 0 dB were set to 0.

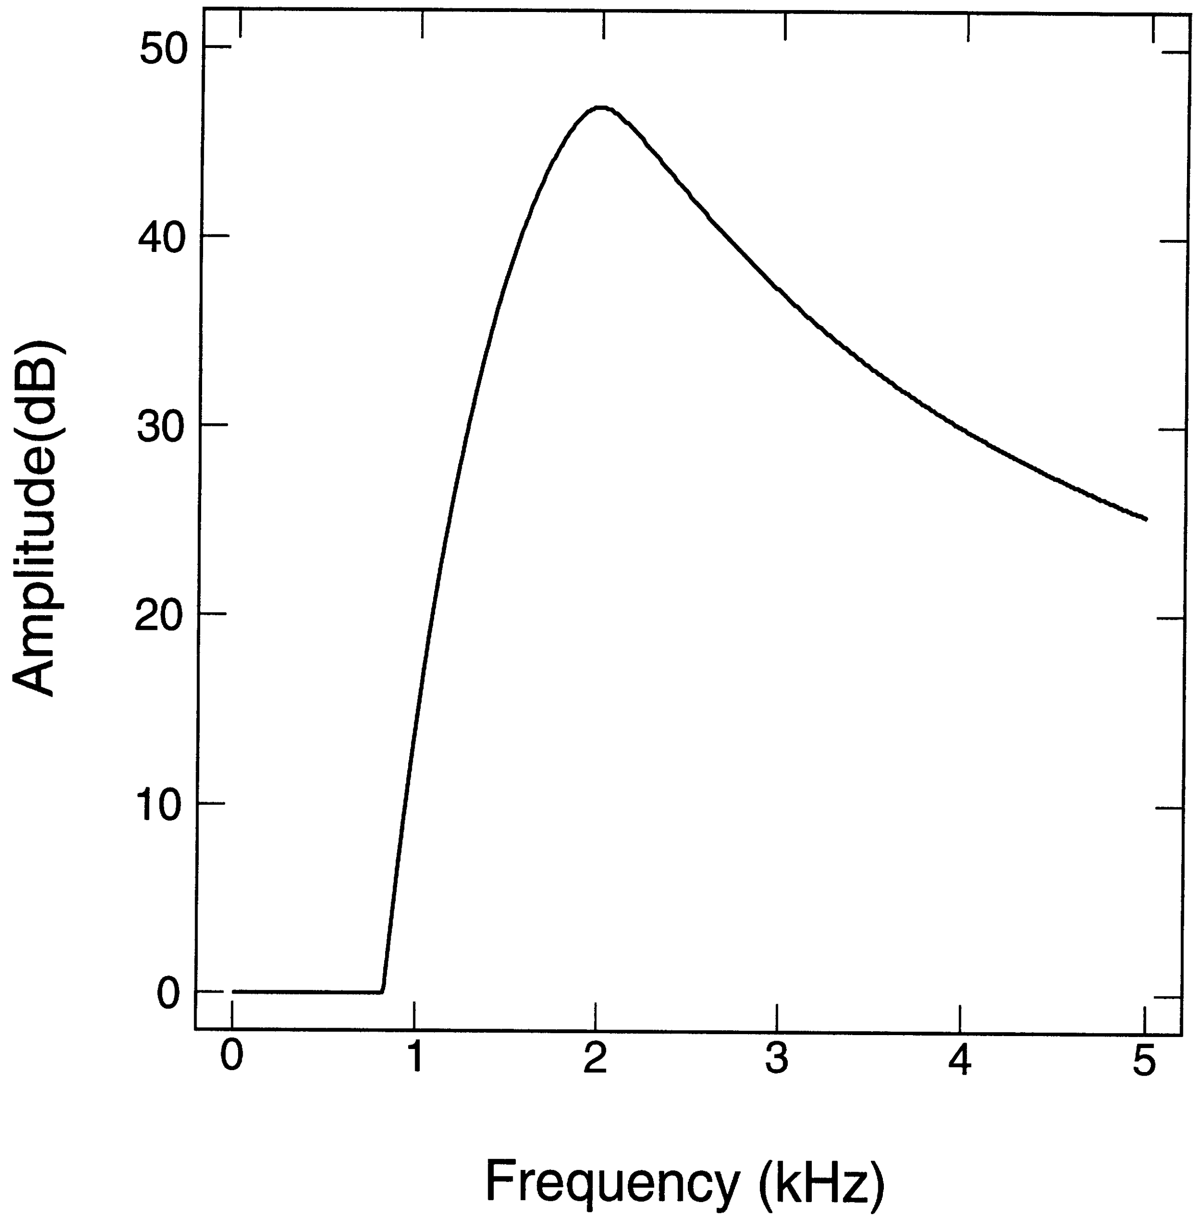

Figure 3-14: Excitation pattern for a 2000 Hz tone for an ear with auditory filters symmetrically broadened by a factor of 3. (Arbitrary scaling). Values below 0 dB were set to 0.

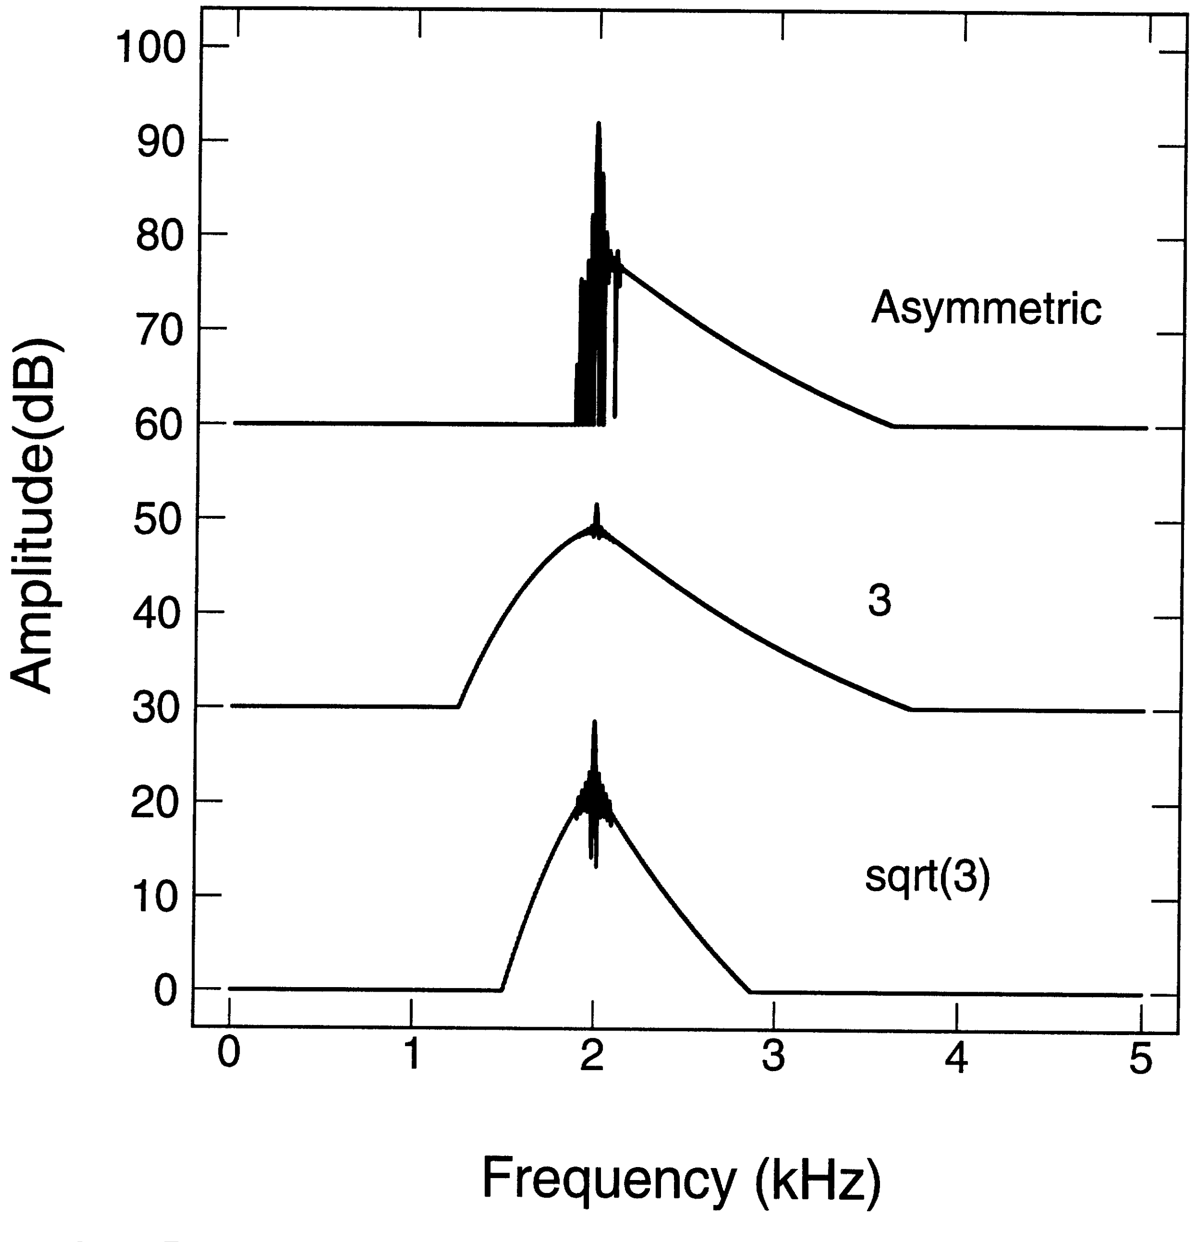

Figure 3-15: Power spectrum of 2000 Hz tone after being smeared by a factor of $\sqrt{3}$ (bottom), by a factor of 3 (middle), and by asymmetic smearing of a factor of 3 on the high frequency side (top). Values below 0 dB were set to 0. Successive plots were offset by increments of 30 dB.

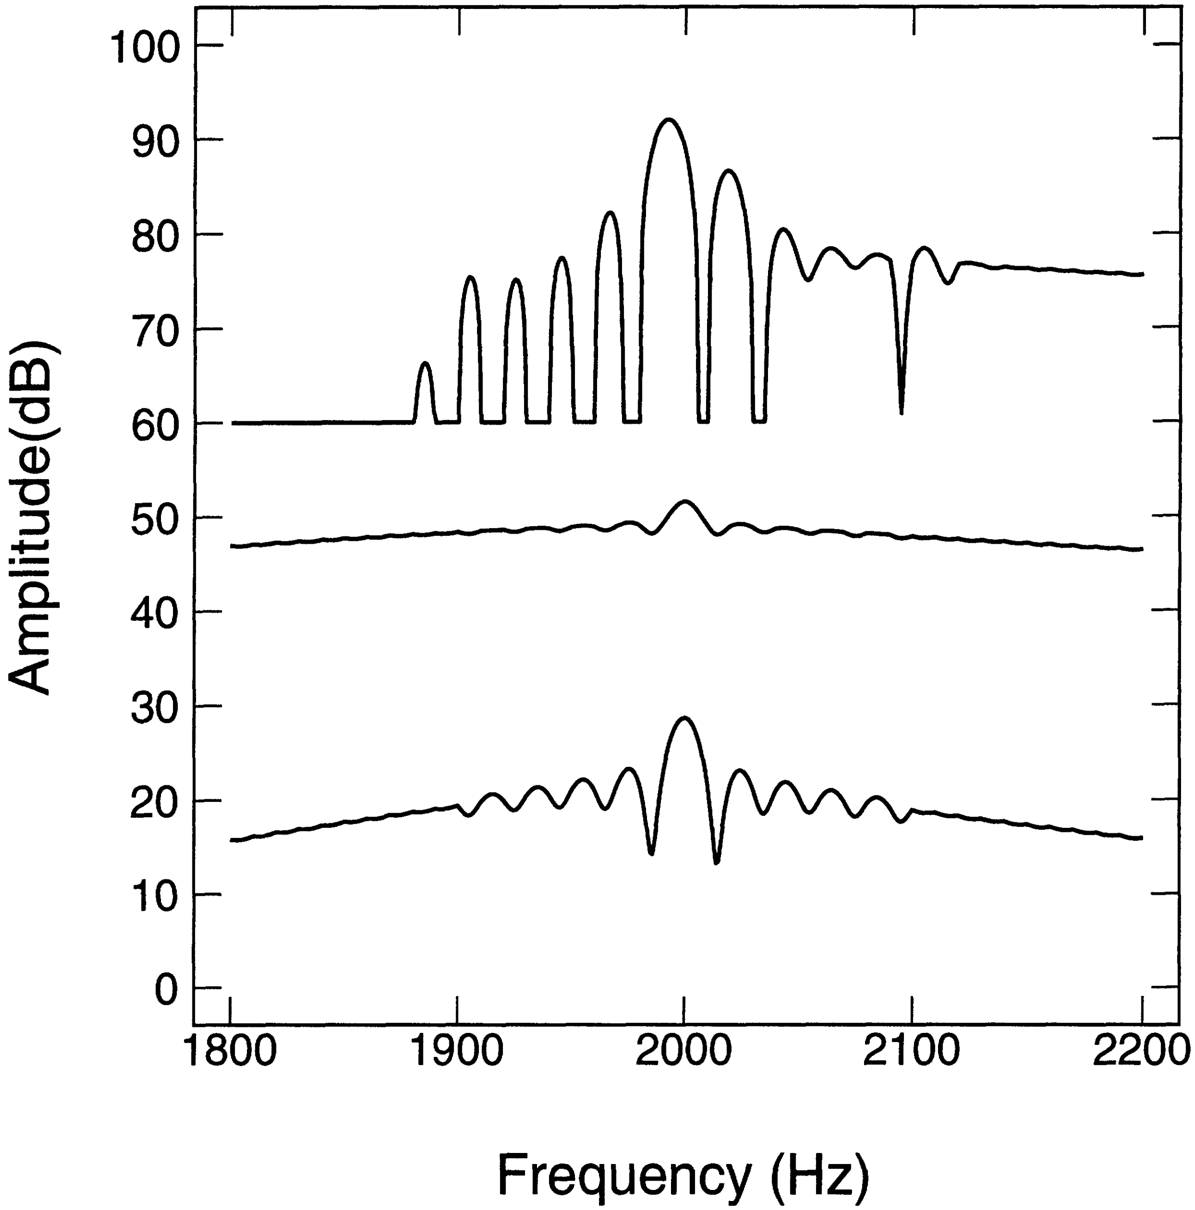

Figure 3-16: Blow up of Figure 1.15 in the frequency region 1800-2200 Hz.

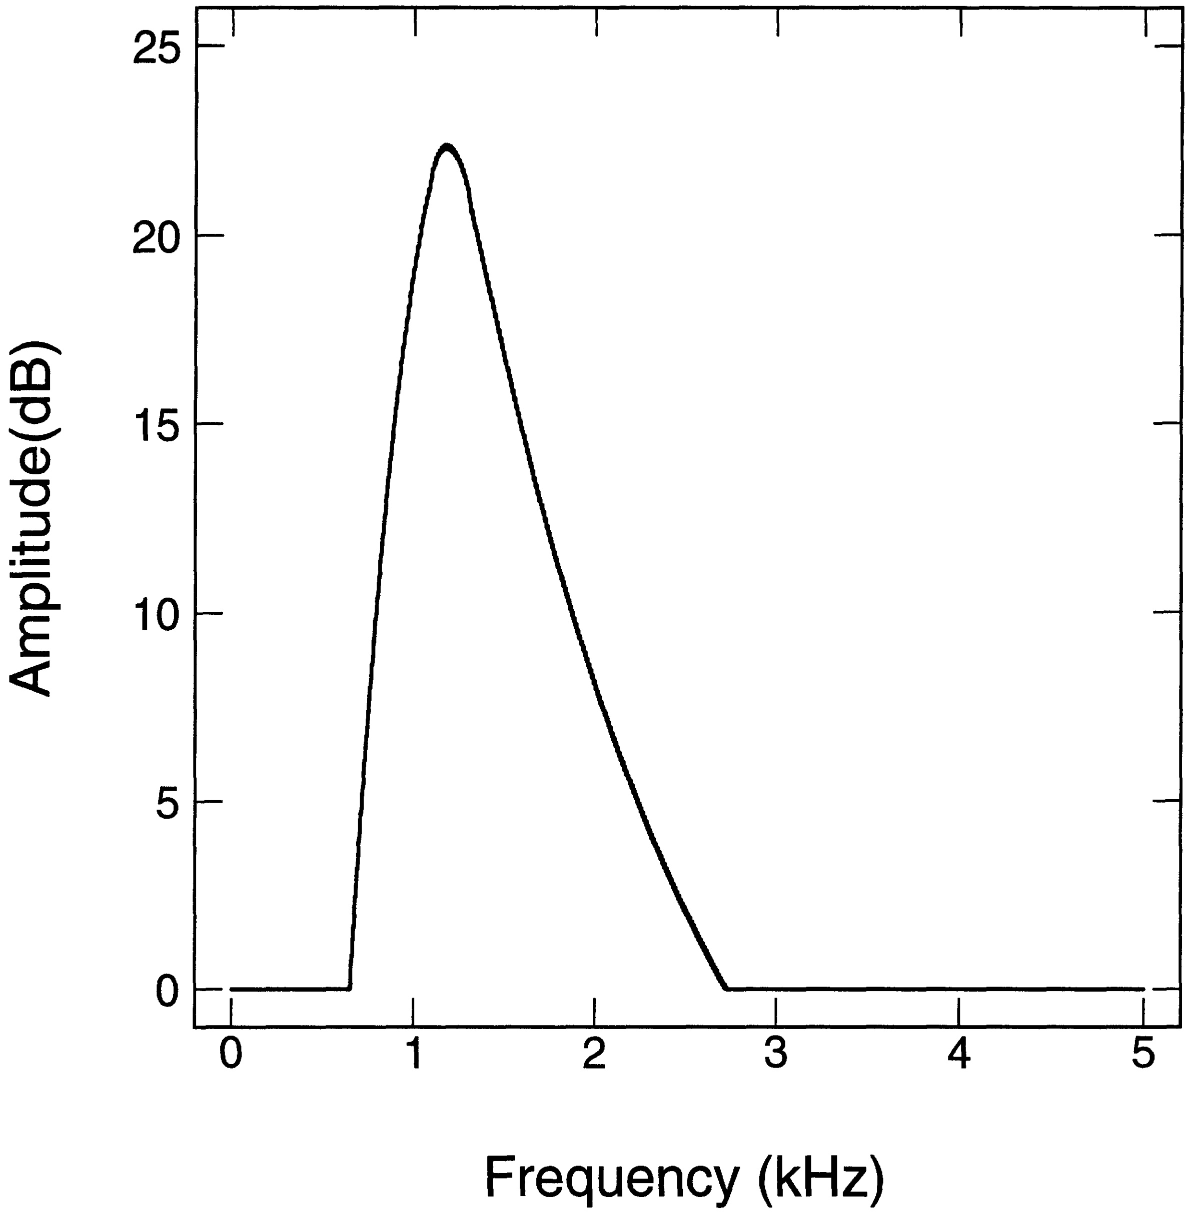

Figure 3-17: Power spectrum of a 200 Hz- wide narrow band of noise, centered at 1200 Hz after being smeared by a factor of 3. Values below 0 dB were set to 0.

Chapter 4

Experimental Setup and Parameters

Eight subjects were used for the experiments. All subjects had clinically normal hearing and were undergraduate or graduate students at MIT. They were paid for their time.

4.1 Simultaneous Masking

4.1.1 Stimuli and Procedure

For the simultaneous masking experiment, the tones and masking bandpass noise were preprocessed separately with the smearing processing in Matlab. The digital waveforms were scaled so that all waveforms (tones and noises) had the same rms value and that good advantage was taken of the 16 bit integer representation of the signal, without clipping. For each smearing condition, there were five versions of the smeared tone for each tonal frequency and five versions of the smeared bandpass noise. For a given condition, each version had the identical magnitude spectrum, but the phase for each waveform was different and was generated randomly. For each presentation of the stimuli, one of these five waveform pairs was chosen at random. Both the masker and signal had a 25 ms rise time and fall time. The signal rise time began 25 ms after the masker rise time (when the masker reached its steady state level) and the signal fall time began 25 ms before the beginning of the masker fall time. The steady state on-time for the masker was 350 ms, so the steady on-time for the signal was 300 ms. The duration between intervals was 300 ms.

The waveforms were played through a sound card on a PC. The output of each channel went to a separate programmable attenuator. The outputs of the attenuators were summed and the resulting signal was sent monaurally to the right earphone of a subject seated at a terminal in a sound-proof booth (see Figure 4.1). Following the presentation of the stimulus, the terminal prompts the subject to indicate in which interval he heard the signal. The programs used to control the sound card and adjust the attenuators were written in ESPUD, a LISP-like interpreter.

In order to carry out the adaptive psychoacoustic experiments, the transformed up-down procedure of Levitt was used. For each trial, the masker was played in both intervals, and the signal was presented randomly in one of the two intervals. The signal level was increased after each incorrect response and it was reduced if the subject was correct twice in a row at a given signal level. The stepsize of the level change was 20 dB for the first change and it was reduced by half for each subsequent change in level until the stepsize became 1 dB (it was not reduced below 1 dB). The run continued until a given number of direction changes occurred (usually 8). The threshold was then determined to be the average of the levels at the last 6 changes of direction.

4.1.2 Simulation of Recruitment

For a given experiment (run), the Levitt procedure determined the nominal level at which the signal should be presented each trial. (The nominal level of the narrowband masking noise remained constant throughout the run). If no recruitment is being simulated then the attenuation applied to the signal is the difference between the level of the signal at the earphones without attenuation and the nominal level (call this difference D). If recruitment is being simulated, then the formula

$$Attenuation = \begin{cases} (80 - input)/2 & input \leq 80 \\ 0 & input > 80, \end{cases}$$

is used to calculate the desired additional attenuation to be applied, where input refers to the sound level within a critical band (see Section 3.2.2). This attenuation (due to the simulation of recruitment) is added to D to obtain the total attenuation applied. If the tone is outside the passband of the masker, then the desired attenuations are calculated separately for the signal and masker. If the tone is in the frequency band of the simultaneous masking noise, then the nominal levels of the tone and masking noise are summed, and the resulting level is inputted to the above formula. The attenuation prescribed by this equation is then added to D for each of the signal and masker.

Since the unattenuated level of the signal at the earphones was measured to be 90 dB SPL, signals louder than 90 dB SPL could not be produced.

4.1.3 Experimental Parameters

The masking noise for this experiment had a 200 Hz bandwidth centered at 1200 Hz. To match the conditions of Lum and Dubno, the noise was presented at 40 dB SPL/Hz. For a 200 Hz wide band of noise, this meant that the nominal noise level was 40 + 10log10 200 = 63 dB SPL. The frequencies tested were 600, 800, 1000, 1200, 1250, 1600, 2500, 3000, and 4000 Hz. 3 subjects were used for this experiment. For all subjects, three smearing conditions were tested: no smearing, smearing to simulate symmetric auditory filters that are three times wider than normal, and smearing to simulate asymmetric auditory filters, in which the low frequency side is three times wider than normal but the high frequency side is normal width. The asymmetric condition was chosen because some hearing impaired subjects exhibit excessive upward spread of masking, corresponding to auditory filters that are broadened primarily on the low frequency side of the filters. The effect of the asymmetric smearing processing is to shift the frequency content of the signal toward higher frequencies. The first two smearing conditions were tested with and without amplitude expansion. The asymmetric condition was tested only with expansion. For the third subject, an additional smearing condition was tested in which the simulated auditory filters were symmetric with a width $\sqrt{3}$ wider than normal. The asymmetric condition was tested with and without expansion. For each condition (frequency of probe, degree of smearing, presence of expansion), three threshold estimates were obtained for each listener in most cases, each typically requiring about 50 trials. The subjects were trained 4-6 hours until their results stabilized, and they were tested in 6-8 sessions, where each session lasted 2-3 hours.

4.2 Forward Masked PTC's

4.2.1 Stimuli and Procedure

The setup for the forward masked PTC's was similar to that of the simultaneous masking experiment. The masking and probe tones were smeared offline and played out through the two channels of the sound card. A two interval, two alternative, forced choice procedure was used, whereby the masker tone was presented in both intervals, and for each trial the probe tone was presented following the masker in one of the two intervals. The rise and fall times for both the masker and probe were each 7 ms. The steady state on-time of the masker was 200 ms. and the on-time of the probe was 16 ms. The duration between the end of the masker and the start of the probe was 2 ms, and the duration between intervals was 300 ms. The outputs from the sound card were attenuated and applied to earphones. The subject was asked to identify the interval that contained the probe. The adaptive procedure used for the experiments was similar to that of the simultaneous masking. The masker level was decreased following an incorrect response, and it was increased following two correct responses at a given level. The initial step size was chosen to be be 8-10 dB and was halved after each level change until it reached 1 dB. The run was completed after 8 changes in direction, and the threshold was the average of the last 6 levels at which the direction changed.

Similar to the simultaneous masking case, the nominal levels for the signals were determined by the Levitt transformed updown procedure. The external attenuations applied to the signals were calculated in the same way as for the simultaneous masking paradigm. However, if recruitment were simulated, since the masker and probe were not presented simultaneously, their power did not need to be added to determine the attenuation when the masker and probe fell in the same critical band (since the simulated recruitment is fast acting).

4.2.2 Experimental Parameters

The frequency of the probe was 1200 Hz and the actual level of the probe at the earphones was 25 dB SPL, whether or not recruitment was simulated. In the expansion condition, the nominal level of the probe was 52.5 dB SPL. The frequencies of the masker tones were chosen to obtain a good idea of the shape of both sides of the PTC's, without a lot of redundant data points. The frequencies tested were 600, 1000, 1200, 1250, 1300, 1400, 1600, and 2500 Hz. In most cases, three estimates were carried out for each frequency for each simulation condition. Three subjects were employed for this experiment. Each subject was tested with no smearing, symmetrical smearing of three ERB's, and asymmetric smearing of three ERB's on the low frequency side, with no smearing on the high frequency side. Two subjects were tested with symmetric smearing of $\sqrt{3}$ aswell. Subjects trained 2-5 hours for this experiment. They came for 3-6 sessions for 2-3 hours per session.

Most measurements were made in the no-expansion condition. Given the model of recruitment used, it was redundant to repeat the runs with the expansion condition, since the probe tone was presented at exactly the same level with or without simulated recruitment. Since the masker thresholds at different frequencies were obtained without simulated recruitment, these levels can be determined using the recruiting equation

$$output = (input + 80)/2,$$

(where output = threshold with expansion and input = threshold without expansion in dB SPL), to calculate what levels of the corresponding thresholds would be in the presence of expansion. Some runs were conducted with recruitment being simulated to confirm that this method was reasonable.

4.3 Loudness Summation

4.3.1 Stimuli and Procedure

To measure loudness summation, narrowband and wideband noises were smeared offline and played out through the two different channels of the sound card. A two alternative, two interval, forced choice procedure was used in which the narrowband noise was presented in one interval and the wideband noise was presented in the other interval. As in the simultaneous masking experiment, the rise and fall times of the noise bursts were 25 ms, and the on-time of the bursts was 300 ms. The time interval between the presentation of the narrowband and wideband noises was 300 ms. The experimental set up was the same as for the other psychophysical experiments and the subject had to indicate the interval in which the noise sounded louder.

For each data point (corresponding to a particular level of the narrowband noise), the level of the wideband noise was varied from trial to trial. For each run, two transformed up-down procedures were run at the same time (Jesteadt, 1980). In one of the procedures the level of the wideband noise was decreased whenever the subject indicated that the wideband noise was louder, and the level of the wideband noise was increased if the subject responded that the narrowband noise was louder twice in a row. This procedure tended to result in a lower matching level for the wideband noise. In the second procedure, the level of the wideband noise was increased each time the subject indicated that the narrowband noise was louder, and the level of the wideband noise was decreased if the subject responded that the wideband noise was louder twice in a row. This procedure tended to result in a higher matching level for the wideband noise. Any given presentation of stimuli to the subject was taken from either up-down procedure at random. When one procedure was finished, all presentations came from the other procedure. For each procedure, the step size of the level changes started at 20 dB and was halved at each level change until it reached 1 dB. The procedure was finished after 10 direction changes had occurred, and the threshold of the procedure was found by averaging the levels at the last 8 directional changes. When both procedures had been completed, the results from the two procedures were averaged to arrive at a final matching level of the wideband noise for the run. In order to reduce the effects of biasing in the results due to the order of presentation, for every level of narrowband noise, separate runs (each run consists of combining the two above procedures) were carried out, where in one run the narrowband noise was always presented first, and in the other run the wideband noise was always presented first. The matched level of the wideband noise was then determined to be the average of the matched levels of the wideband noise indicated by these two runs. The expansion algorithm for this experiment was identical to that of the forward masked PTC's because the two noises were never presented simultaneously.

4.3.2 Experimental Parameters

The center frequency of both of the bands of noise used in this experiment was 1200 Hz. The narrowband noise had a bandwidth of 200 Hz and the wideband noise had a bandwidth of 1800 Hz. The ratio of bandwidths corresponded to the bandwidths used by Florentine and Zwicker (Florentine, Buus, Scharf and Zwicker, 1980) (0.17 times the center frequency for the narrowband and 1.52 times the center frequency for the wideband). The levels of the narrowband noise that were used for the no-expansion condition were 25, 45, 55, 65, 75, and 85 dB SPL. When recruitment was simulated, only 55, 65, and 75 dB SPL were used. (25 and 45 dB SPL would be inaudible and 85 dB SPL would not result in any recruitment for levels near the level of the narrowband noise). Two runs were carried out for each of the narrowband and wideband noises being presented first (each run consists of two up-down procedures). Thus a total of eight thresholds were averaged for each condition. Runs were conducted with all four combinations of expansion/ no expansion, and 3 * ERB smearing/ no smearing. Three subjects were used in this experiment. All had participated in one of the other psychophysical experiments and were provided with 2-3 hours of practice for this experiment. For each subject, the experiment lasted about four sessions, with 2-3 hours per session.

4.4 Thresholds

The transformed updown adaptive procedure was also used to calculate absolute thresholds. For the threshold paradigm, the tone (possibly smeared) was played in either the first or second interval at random. In the other interval, there was silence. The terminal informed the subject immediately before each interval began, so the subject would know whether the silence corresponded to the interval before or after the tone. There were 300 ms between intervals. The stimuli used for the threshold test depended on the subject.

For subjects who participated in the simultaneous masking experiment, the stimuli used were sinusoids of 300 ms duration. For subjects who participated in the forward masking experiment, the stimuli used to measure thresholds were the short 30 ms probe tones. The stepsizes and number of turning points were the same as for the simultaneous masking and forward masking experiments.

4.5 Consonant Recognition

4.5.1 Stimuli and Procedure

The consonant recognition experiment was modelled after Duchnowski (1989) and Zurek and Delhorne (1987) and it consisted of recognizing a set of 24 consonants. Each consonant was paired with three vowels (a, i, u) to make 72 consonant-vowel pairs. The 72 stimuli were presented to the subject in a random order. For each CV presentation the subject typed what consonant he heard at the terminal. In training runs, the subject received feedback as to what the correct consonant was; in the test runs, no feedback was given. Two subjects were used for this experiment. The waveforms used for this experiment were played through a DAL card on a PC. Speech shaped noise was played through the other channel of the DAL card. The speech and noise were passed through separate attenuators to adjust the SNR, and the outputs of the attenuators were summed. This output was routed to a DSP96 card in a second PC, which the smeared the signals in real time. For the no-smearing condition, this PC was bypassed. The resulting output was then sent to another DSP96 card inside a third PC, which carried out the real-time amplitude expansion. This PC was bypassed for the no-expansion condition. The resulting signal was then sent to the right earphone of the subject in a sound proof booth (see Figure 4.2).

4.5.2 Experimental Parameters

Two sets of expansion parameters were used in this experiment. The first set was similar to that used in the psychophysical experiments. Normal thresholds of 0 dB SPL and impaired thresholds of 40 dB SPL were specified in all bands, for all frequencies. This resulted in a T_r = 75 dB SPL for the recruiting function for all bands. The attenuator for the speech signal was set such that the speech would be played at a level of 70 dB SPL rms, before passing through the simulation processors. The bandwidths (B_L and B_H) of the smearing Gaussian in the smearing algorithm were set to 1.0 (symmetric). Tests were conducted under twelve different conditions. 3 SNR's were used (quiet, 6 dB, and 0 dB), and each SNR was tested with all four combinations of smearing/no smearing, expansion/no expansion. Training lasted about 5 hours, until the scores appeared to plateau and 3-5 runs were carried out for each condition.

The second set of expansion parameters matched the conditions that Duchnowski used to simulate a severely impaired subject. For this experiment, the normal thresholds were the SPL's corresponding to 0 dB HL in the center of each band (generally in the range of 5-10 dB SPL). The impaired thresholds were those of subject AL in Duchnowski's study, whose thresholds ranged from 95 dB SPL at the low frequencies to 60 dB SPL at the high frequencies. (See Figure 4.3). In this experiment the unprocessed speech with or without noise was passed through a filter before it passed through the simulations. This filter was either flat with no gain (a no hearing aid condition) or a high frequency emphasis filter, with a frequency response similar to what Duchnowski used (simulating a high frequency emphasis hearing aid). (See Figure 4.4). In either case, the attenuator for the speech signal was set so that the speech was presented at a level of 94 dB SPL, before passing through the simulations. The smearing bandwidth used was the same as for the previous experiment ($B_L = B_H = 1.0$). Tests were conducted under 8 different conditions: two frequency responses, and all four combinations of quiet/0 dB SNR and smearing/ no smearing. The expansion simulation was used in all runs. Training lasted about one hour, since subjects were previously trained on the consonant recognition experiment. 3-5 runs were carried out for each condition. The combined consonant recognition experiments lasted 9-10 sessions with 1-2 hours per session.

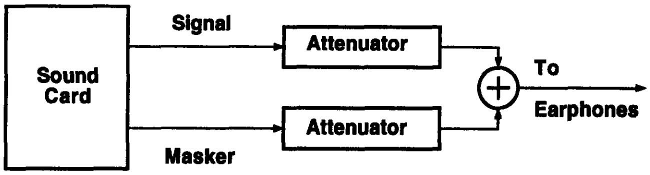

Figure 4-1: Setup diagram for psychoacoustic experiments that used pre-smeared stimuli.

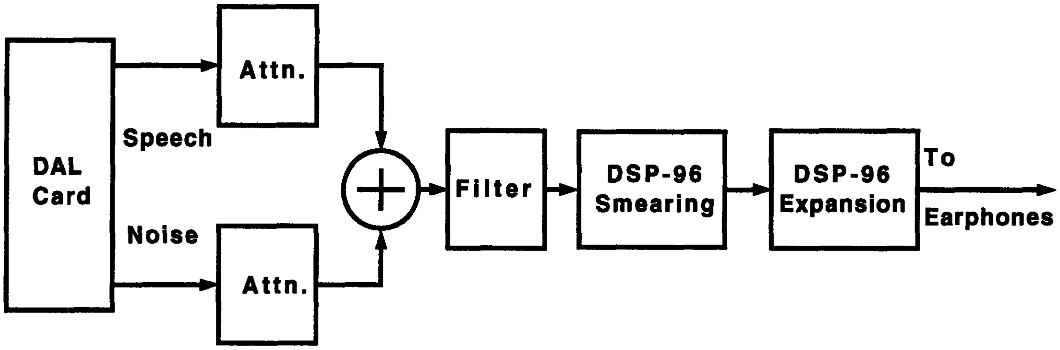

Figure 4-2: Set up diagram for consonant reception experiments. Consonant waveforms and speech shaped noise are played out through the DAL card.

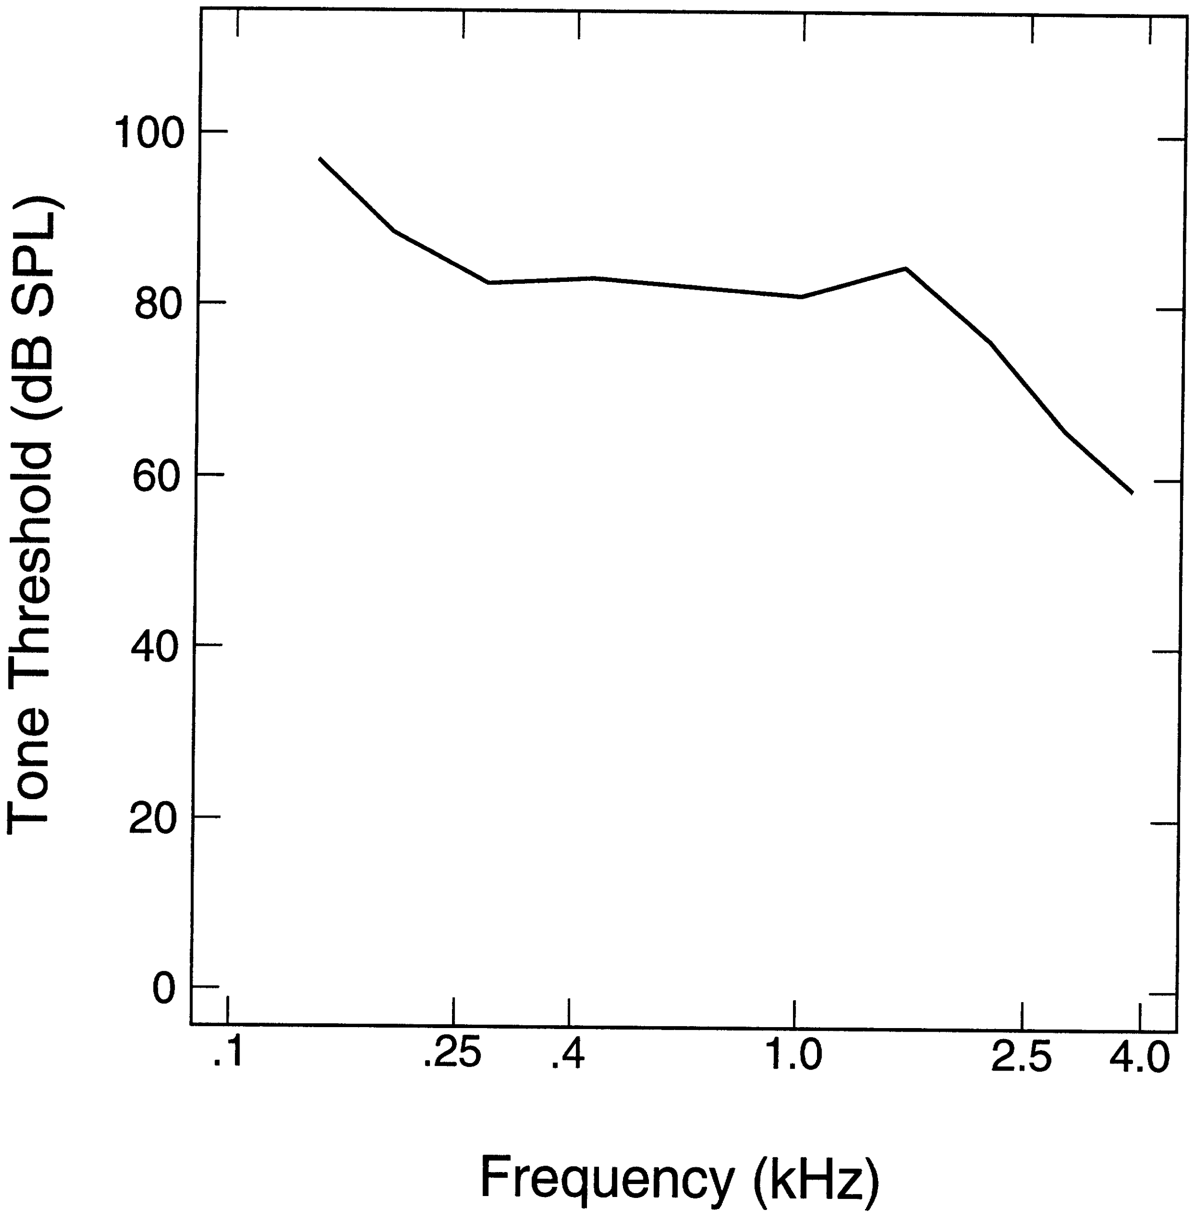

Figure 4-3: Absolute thresold of hearing impaired subject AL. Redrawn from Duchnowski, 1989.

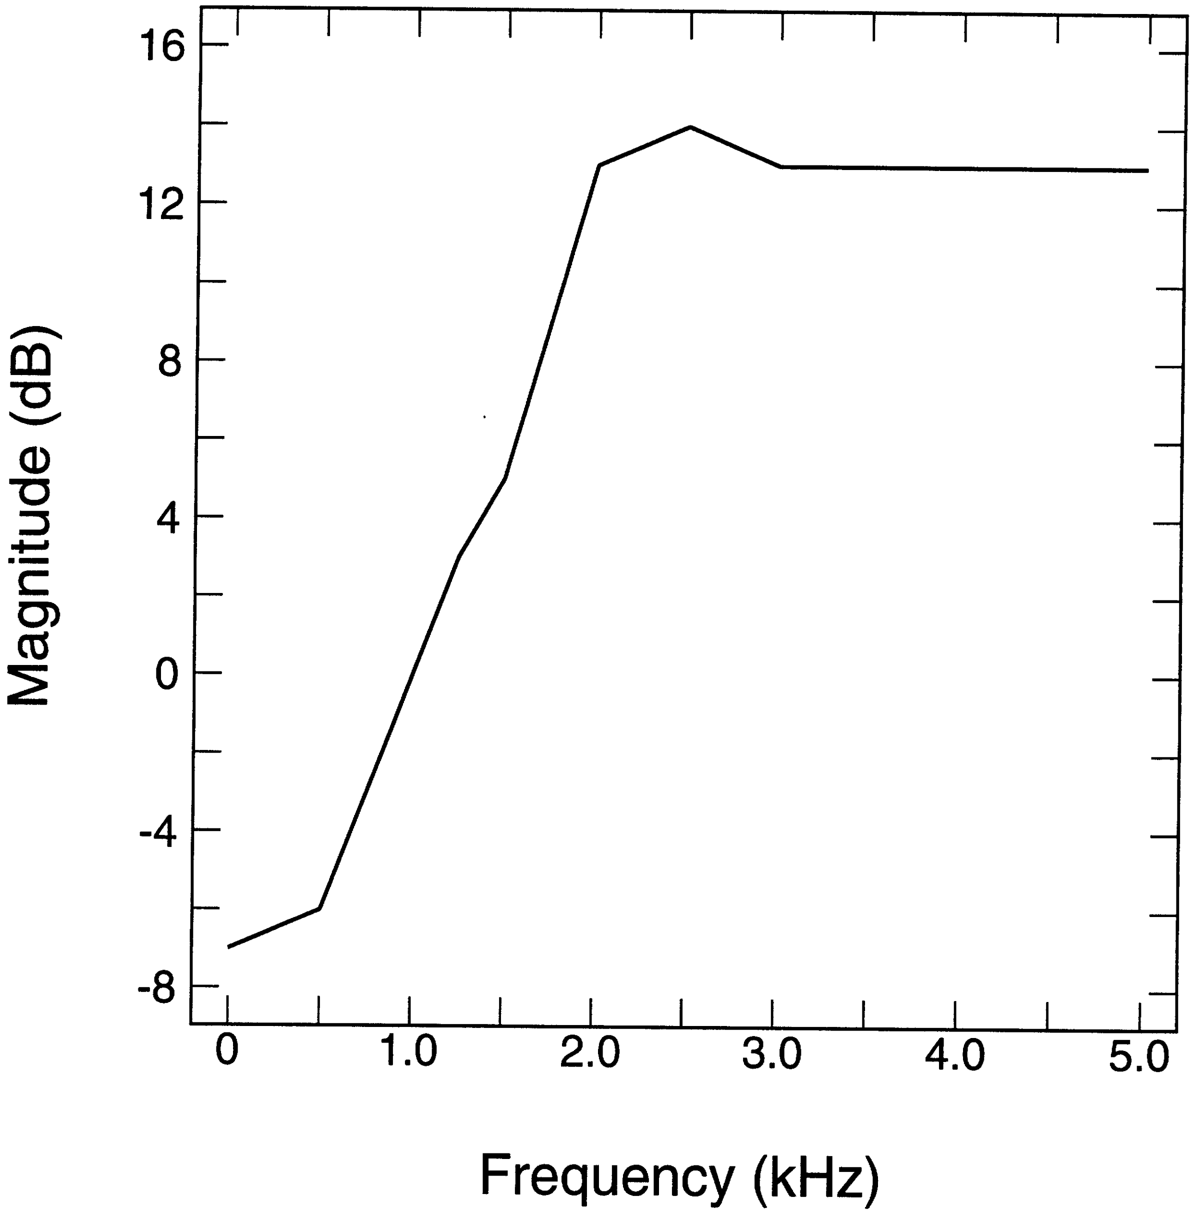

Figure 4-4: Frequency response of high frequency emphasis filter.

Chapter 5

Results

5.1 Narrowband Masking Noise

The masking patterns for the simultaneous narrowband masking noise experiment are shown in Figures 5.1-5.7. For each condition, an analysis of variance was performed with subject, frequency, and repetition as factors, that indicated that subject by frequency did not interact significantly. (The shape of the masking pattern is essentially the same for all subjects). [1] As a result, the data from all three estimates from all three subjects were averaged. For the condition where the smearing was $\sqrt{3}$, only one subject was tested, so the data is from that subject alone.

To obtain a feel for the variability of the data within each of the subjects the standard deviations of many sets of data were calculated. For each subject, the standard deviation (in dB) of the repeated data points were calculated for each condition of hearing loss simulation at each frequency. The SD's were then averaged across frequency for each hearing loss condition. These averaged SD's are shown in Tale 5.1. The numbering system

---

[1] A factor or interaction of factors was considered not to be significant if the probability that the apparent effect of the factor was as large as it was, under the hypothesis that the factor has no effect on the data was greater than .05. (In most cases, however, this probability was greater than .9).

of the conditions is as follows: 1) no simulation, 2) expansion, 3) smearing = $\sqrt{3}$, 4) smearing = $\sqrt{3}$, plus expansion, 5) smearing = 3, 6) smearing = 3, plus expansion, 7) asymmetric smearing plus expansion, 8) absolute thresholds, no smearing, 9) absolute thresholds, smearing =3.

In addition, the standard deviation of the average thresholds for the three subjects was calculated for each hearing loss simulation condition and each frequency. This measure provides information about the variability of the data across the three subjects. The SD's were then averaged across frequencies, and the resulting numbers are shown in the bottom row of Table 5.1.

Condition	1	2	3	4	5	6	7	8	9
Subject JK	1.4	1.5	3.1	2.0	2.6	2.4	2.9	2.0	2.2
Subject TP	2.1	1.6	-	-	2.3	2.0	3.1	3.3	2.4
Subject LP	3.3	2.3	-	-	3.3	2.9	1.5	2.6	2.4
AVG	4.2	2.5	-	-	2.6	2.2	2.1	6.5	6.1

Table 5.1: Standard deviations (in dB) from the narrowband noise masking data, averaged across frequencies.

5.2 Psychoacoustic Tuning Curves

The PTC's for the different conditions tested are shown in Figures 5.8-5.12. For each condition, an analysis of variance was performed with subject, frequency, and repetition as factors. The ANOVA indicated that subject and frequency did not interact significantly (using the same significance criteria as for the simultaneous masking experiment) for any conditions except for the no-smearing condition (with and without recruitment). [2]

---

[2] In the case of no simulation, the likelihood of the data given the hypothesis of no interaction was 2%.

Nevertheless, for simplification purposes, data from all estimates from the three subjects were averaged in all cases. For the condition of auditory filters broadened by $\sqrt{3}$, the data from two subjects were averaged.

The data for conditions with expansion were calculated from the no-expansion data for each smearing condition, using the recruitment equation (as described in 4.2.2). Some runs were conducted with expansion applied to confirm that these calculations were reasonable.

As with the narrowband noise masking data, the standard deviations of the repeated data points were calculated for each hearing loss condition and each frequency for each subject. The data are shown in Table 5.2. The numbering system for the conditions is the same as for Table 5.1.

The standard deviation (in dB) was then taken of the average thresholds of the three subjects for each simulation condition and each frequency. These standard deviations were then averaged across frequency (for each condition), and the resulting standard deviations are shown in the bottom row of Table 5.2.

Condition	1	2	3	4	5	6	7	8	9
Subject JR	2.9	-	5.0	-	2.5	-	2.1	1.4	2.1
Subject KM	3.5	-	-	-	4.3	-	2.9	0.5	1.7
Subject BS	2.4	-	4.5	-	2.6	-	2.0	1.7	1.9
AVG	6.1	3.2	2.9	1.7	4.7	2.3	2.9	2.5	3.8

Table 5.2: Standard deviations (in dB) from the PTC data, averaged across frequencies.

5.3 Loudness Summation

The loudness summation values calculated for each condition are shown in Figures 5.13-5.15 for each of the three subjects tested.

For the loudness summation data, the standard deviations for a given subject and a given simulation were calculated for the average of two runs calculated simultaneously-where one run tends toward higher levels and the other tends toward lower levels (see section 4.3.2). These standard deviations were then averaged over narrowband noise levels for each of the four simulation conditions. Table 5.3 shoes these averaged standard deviations for each subject and each condition. The numbering for the conditions is 1) no simulation, 2) expansion, 3) smearing = 3, 4) smearing = 3, plus expansion.

Condition	1	2	3	4
Subject BS	4.0	2.9	2.0	1.9
Subject JK	4.7	1.6	2.9	1.4
Subject LP	5.1	2.6	5.0	1.7

Table 5.3: Standard deviations (in dB) from the loudness summation data, averaged across levels of the narrowband noise.

5.4 Consonant Recognition

Figure 5.16 shows the percent correct for each of the two subjects used for each condition of the experiment in which mild hearing loss was simulated. Figure 5.17 shows the percent correct for each of the two subjects (or one subject) for all conditions used in the simulation of moderate/severe hearing impairment.

For each subject, the standard deviation in the percent correct data was calculated for each simulation condition. Table 5.4 shows these standard deviations for each condition for each subject. Conditions numbered 1-11 are from the mild/moderate hearing loss simulation. Conditions 12-19 are from the severe hearing loss simulation experiments. The numbering system for the conditions is 1) smearing, quiet, 2) expansion, quiet, 3) both, quiet, 4) no simulation, 6 dB SNR, 5) smearing, 6 dB SNR, 6) expansion, 6 dB SNR, 7) both, 6 dB SNR, 8) no simulation, 0 dB SNR, 9) smearing, 0 dB SNR, 10) expansion, 0 dB SNR, 11) both, 0 dB SNR, 12) HFE, no smearing, quiet, 13) HFE, smearing, quiet, 14) flat, no smearing, quiet, 15) flat, smearing, quiet, 16) HFE, no smearing, 0 dB SNR, 17) HFE, smearing, 0 dB SNR, 18) flat, no smearing, 0 dB SNR, 19) flat, smearing, 0 dB SNR.

In addition, the average scores of each subject for each simulation condition were used to compute the SD across subjects for each condition. The bottom row of Table 5.4 shows the standard deviations for each condition.

Condition	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
Subject AC	5.9	2.9	5.1	2.0	3.1	1.8	6.6	4.5	1.9	2.5	3.0	4.2	7.4	3.5	4.6	-	-	-	-
Subject IG	2.1	7.4	11.2	4.5	5.5	2.9	7.1	2.5	3.5	2.3	3.5	4.7	3.6	5.2	4.4	3.4	3.8	1.8	1.9
AVG	0.1	0.0	2.5	0.5	2.0	2.1	1.0	4.0	1.8	1.3	3.0	0.8	3.0	1.3	4.5	-	-	-	-

Table 5.4: Standard deviations (in % correct) from the consonant recognition data.

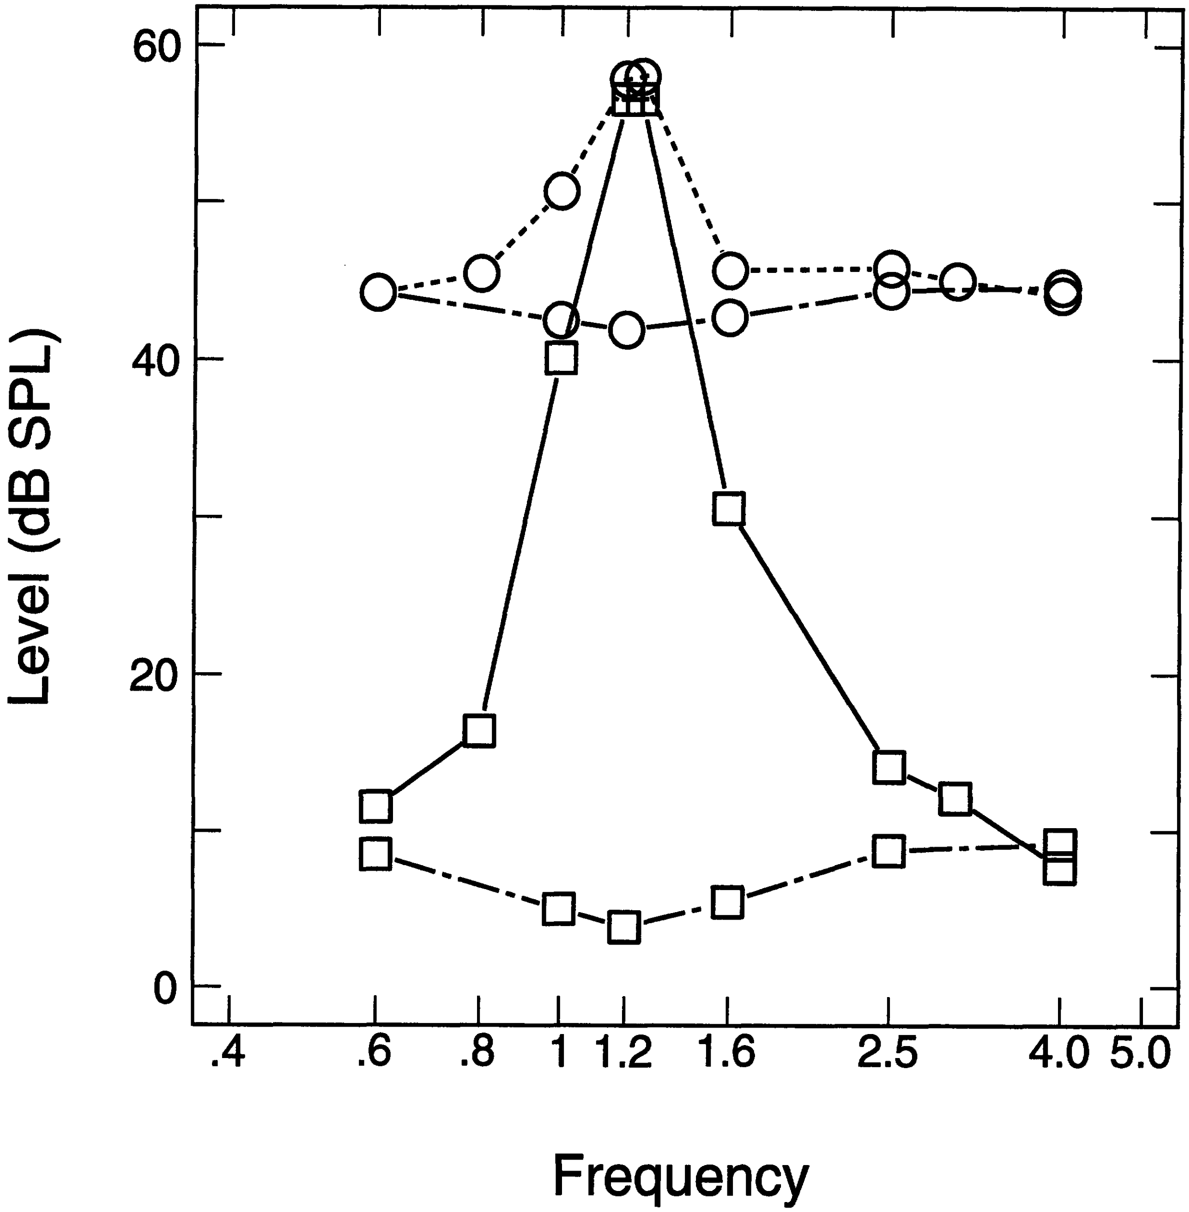

Figure 5-1: Effect of expansion on narrowband noise masking patterns. Solid line (□) shows no-simulation pattern. Dotted line (o) shows pattern with expansion. Broken lines show absolute thresholds in the absence of narrowband masker for no simulation (□) and the calculated absolute thresholds with expansion (o).

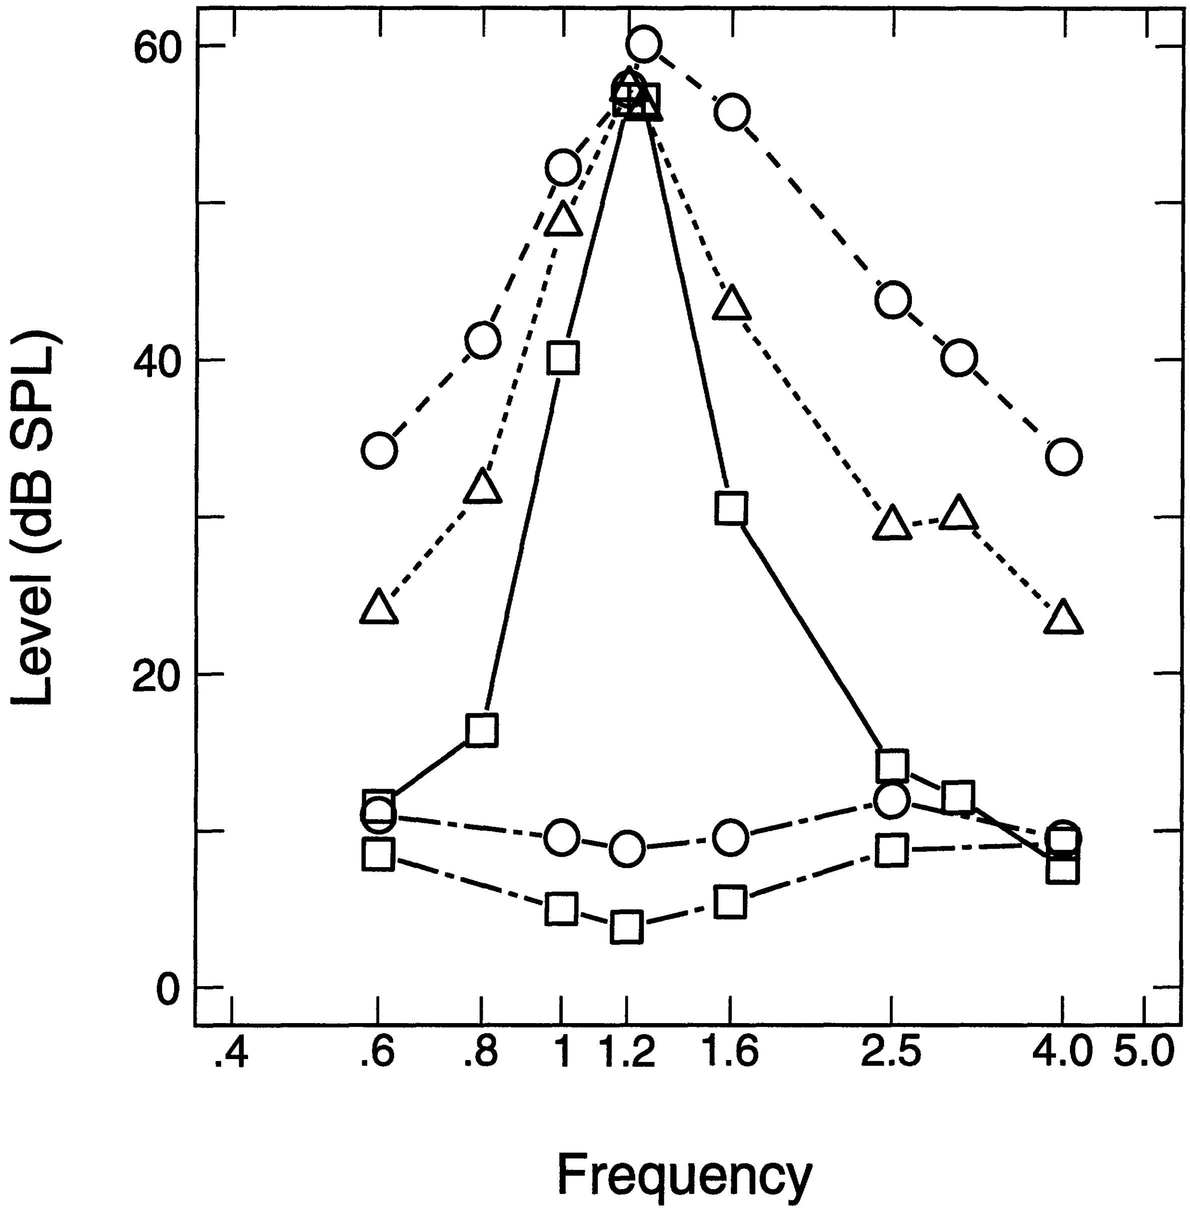

Figure 5-2: Effect of smearing on narrowband noise masking patterns. Solid line (□) shows the no-simulation pattern. Dotted lines show the data for smearing = $\sqrt{3}$ (△) and smearing = 3 (○). Broken lines show thresholds in the absence of masking noise for no smearing (□) and smearing = 3 (○).

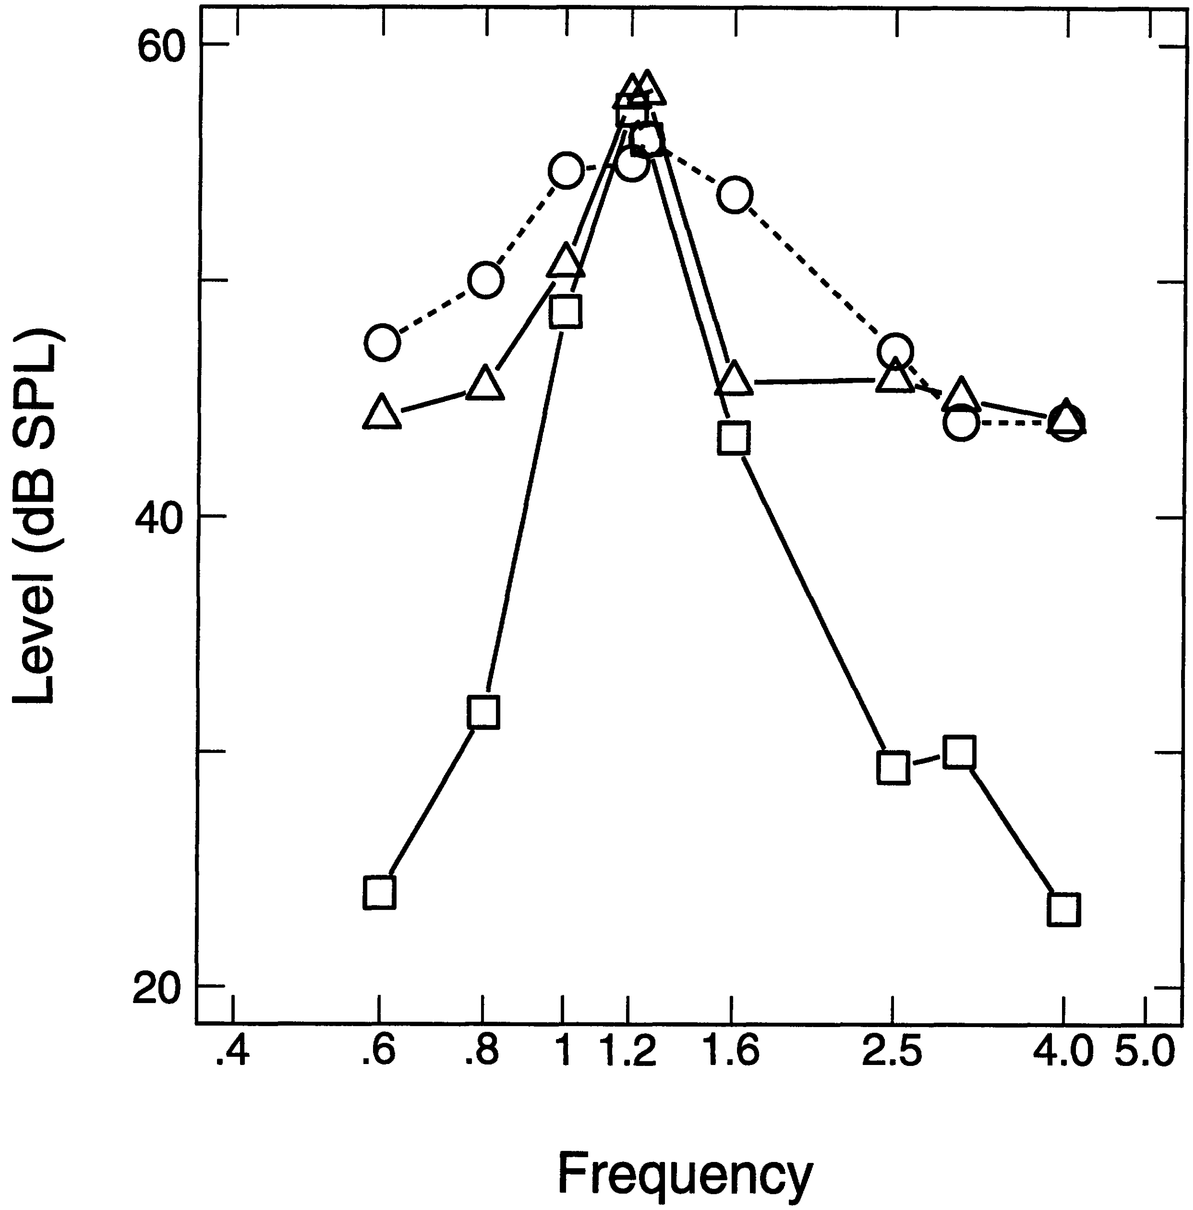

Figure 5-3: Comparison of simulations alone and in combination. Solid lines show smearing = $\sqrt{3}$ alone (□) and expansion alone (△). Dotted line (○) shows smearing = $\sqrt{3}$ with expansion.

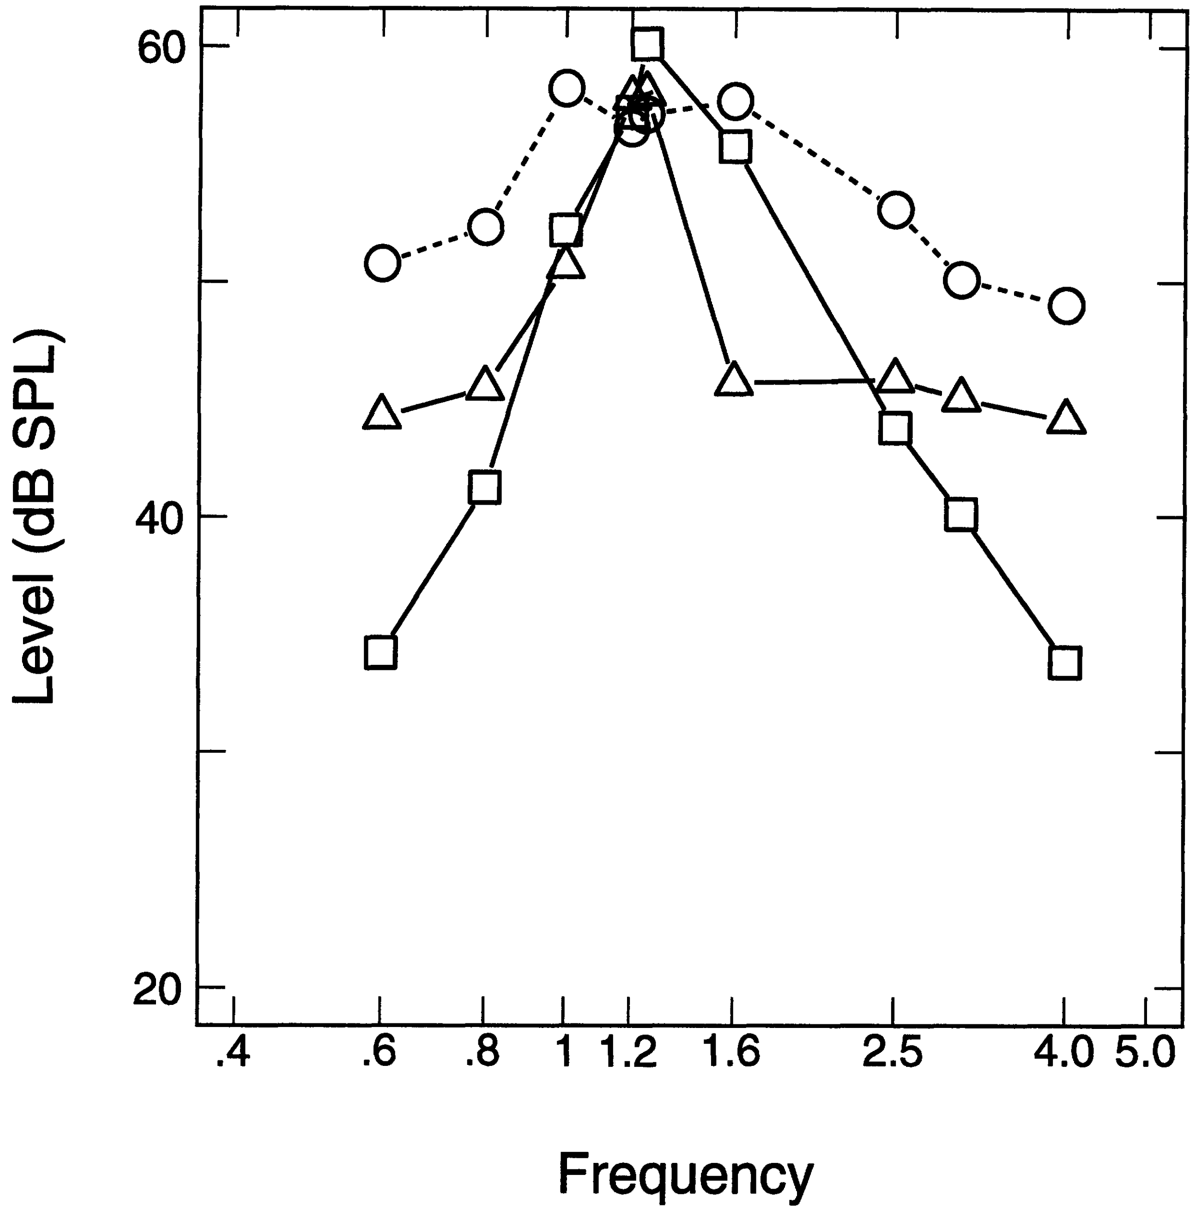

Figure 5-4: Comparison of simulations alone and in combination (continued). Solid lines show smearing = 3 alone (□) and expansion alone (△). Dotted line (○) shows smearing = 3 with expansion.

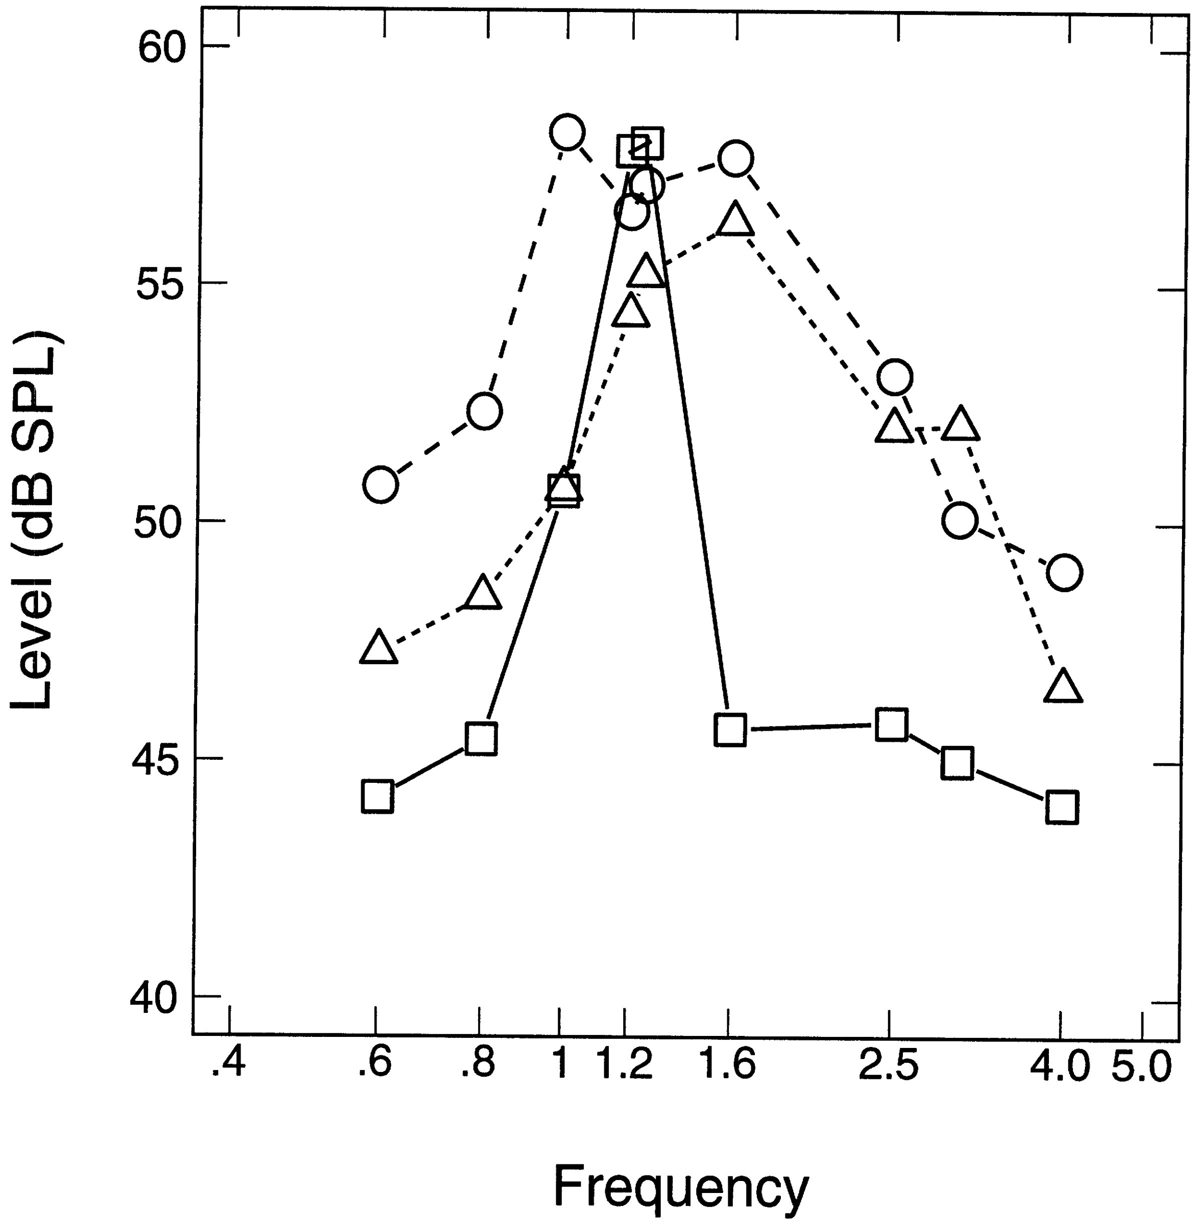

Figure 5-5: Effect of smearing on expansion. Symbols show the data from expansion alone (□), symmetric smearing = 3 plus expansion (○), and asymmetric smearing plus expansion (△).

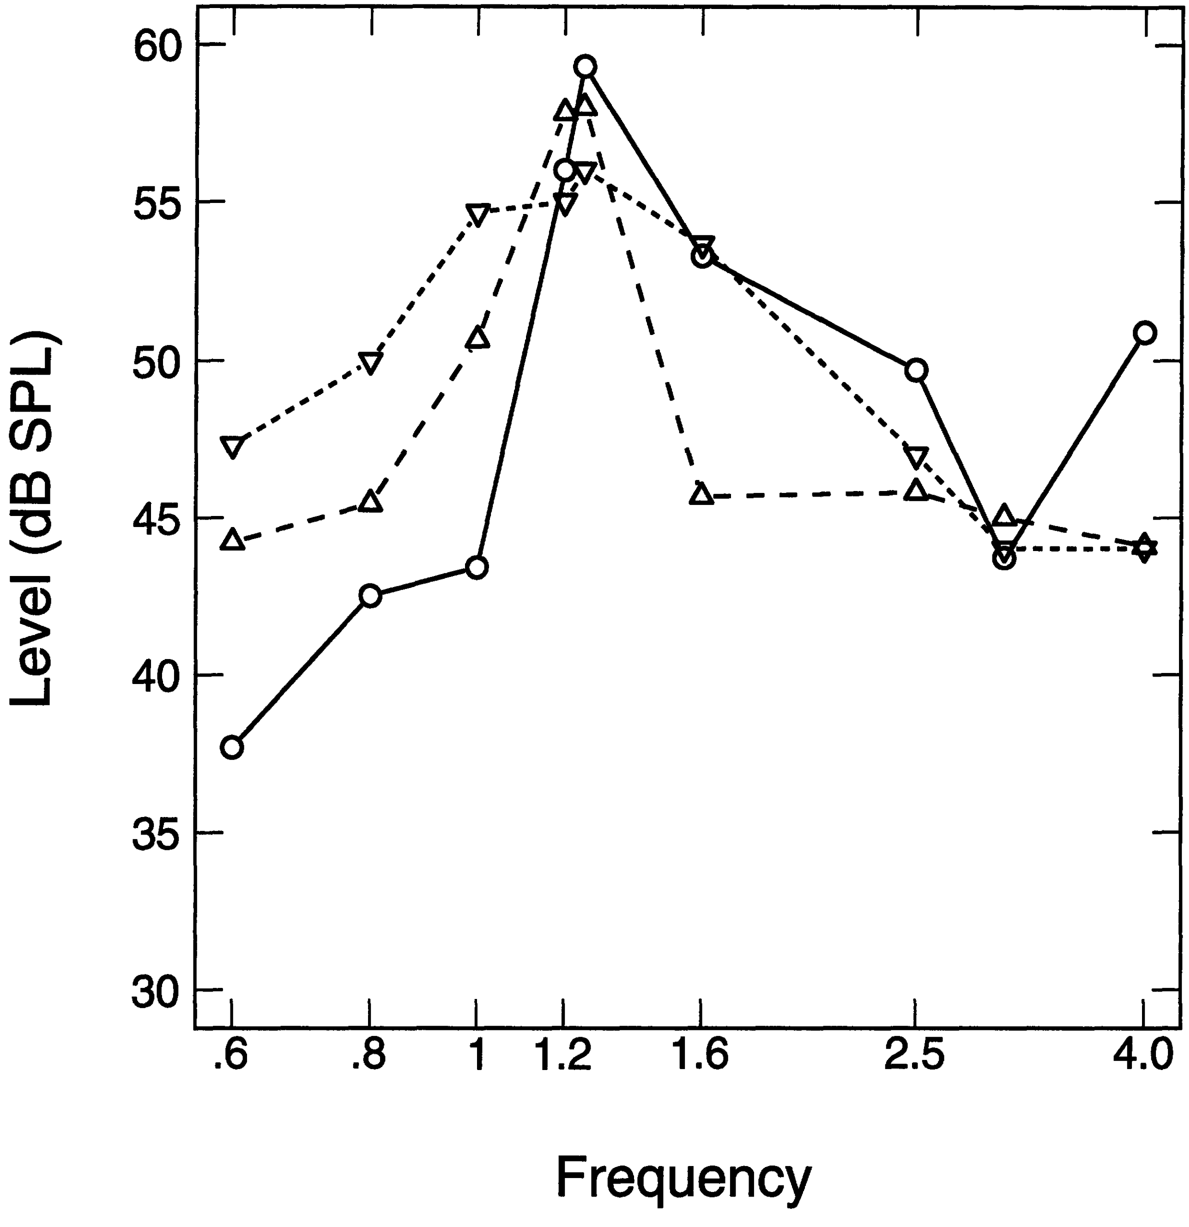

Figure 5-6: Comparison of simulations with hearing impaired data. Symbols show the data from hearing impaired data (Dubno and Schaefer, 1991) (o), expansion simulation alone (△), and smearing = $\sqrt{3}$ plus expansion (▽)

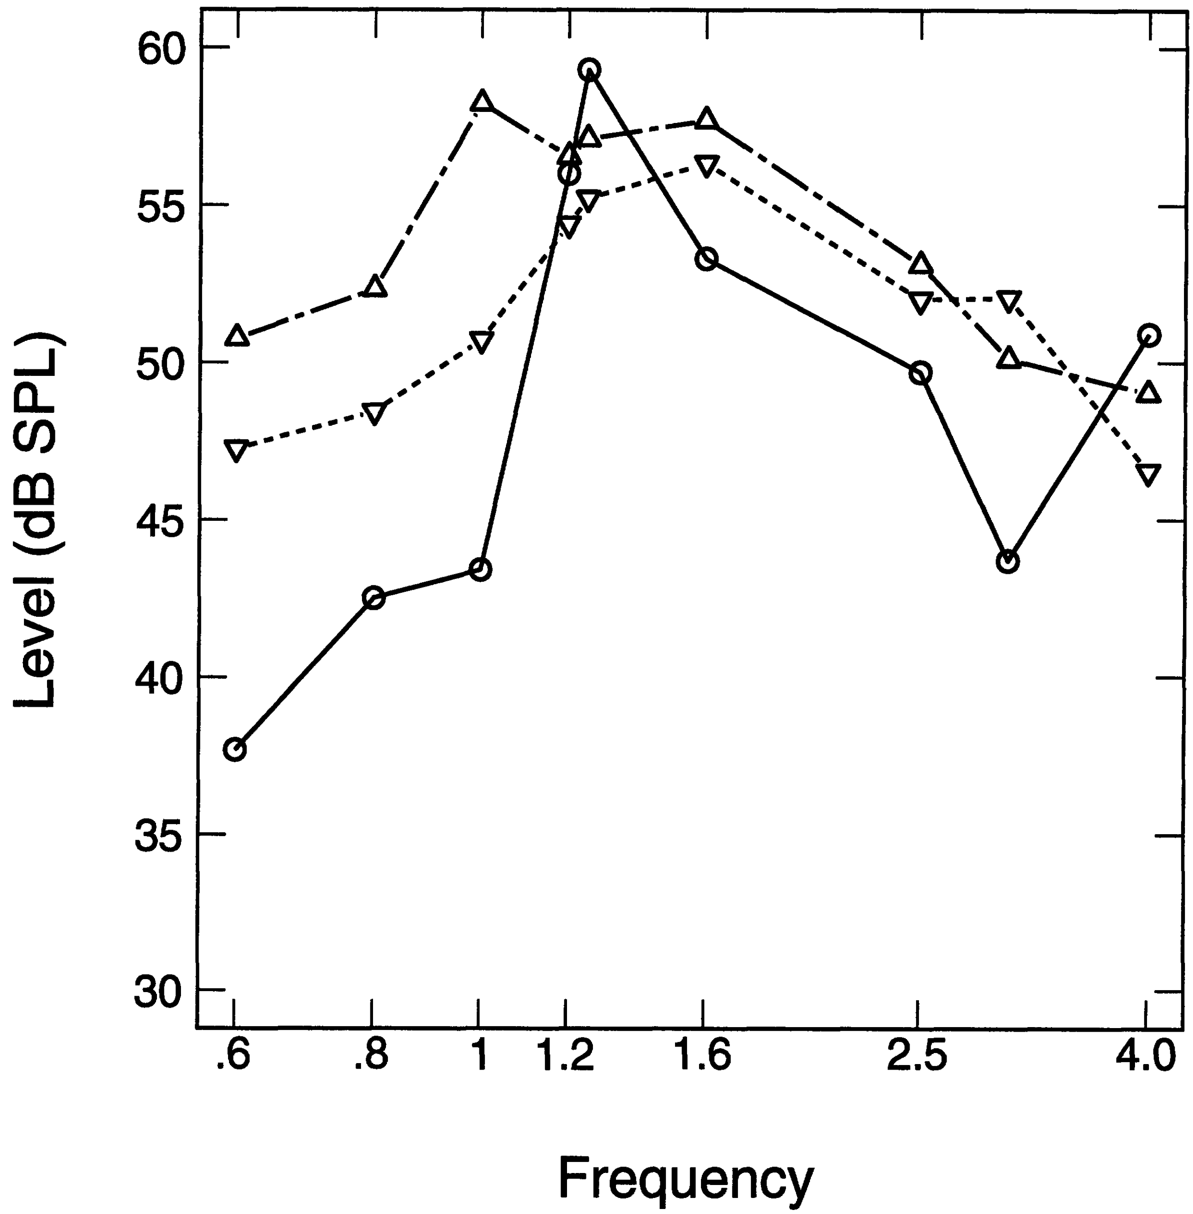

Figure 5-7: Comparison of simulations with hearing impaired data (continued). Symbols show the data from hearing impaired data (Dubno and Schaefer, 1991)(○), asymmetric smearing plus expansion (▽), and symmetric smearing = 3 plus expansion (△).

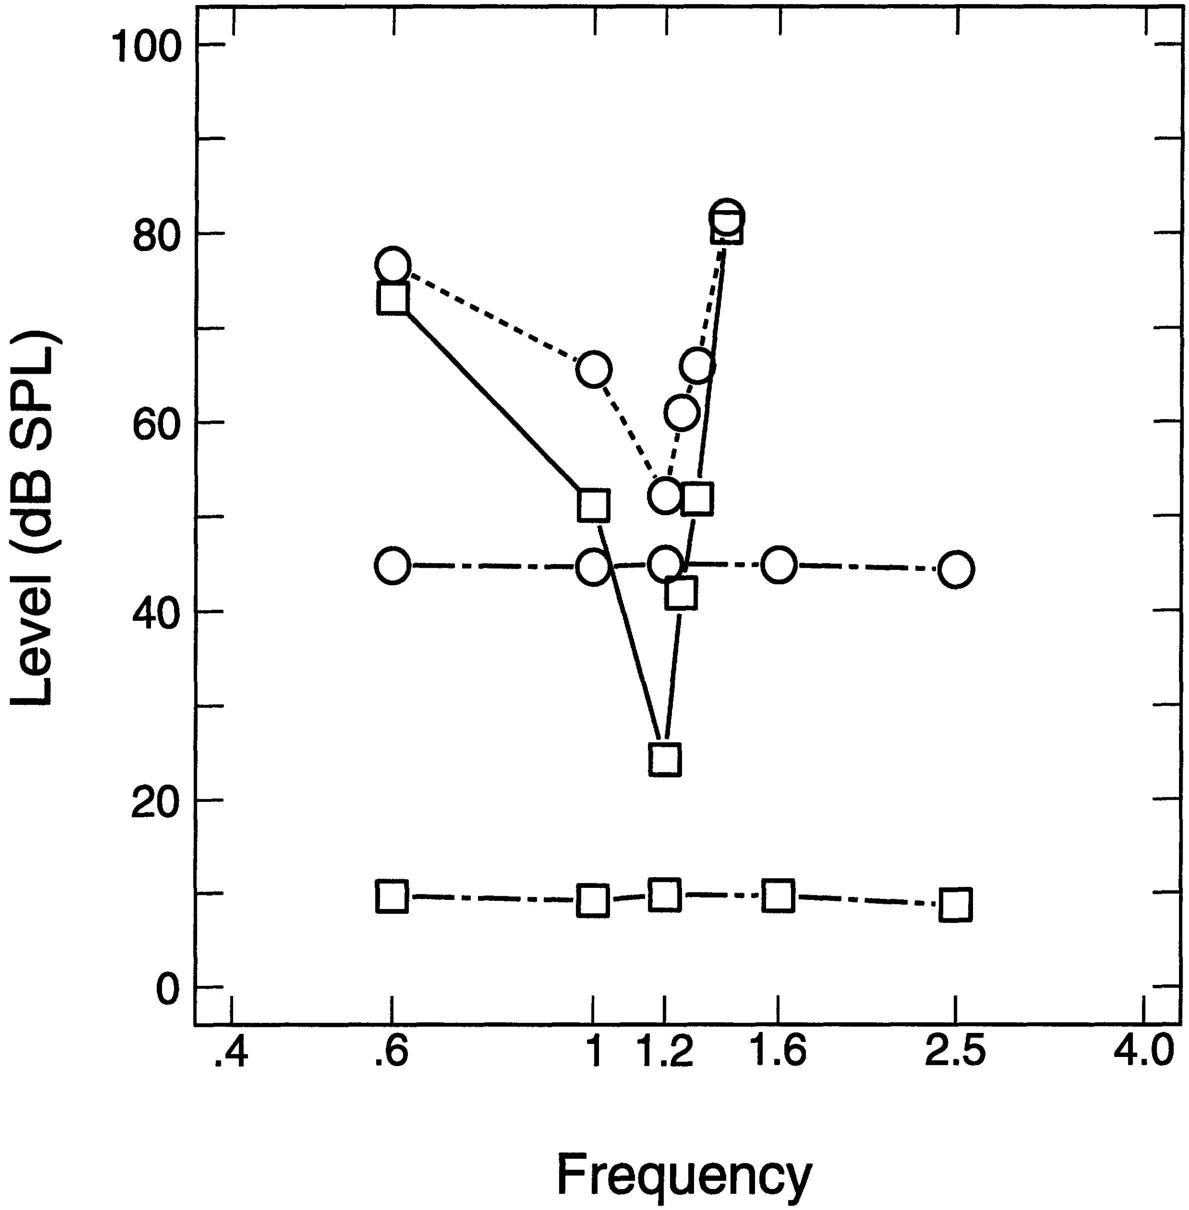

Figure 5-8: Effects of expansion on PTC's. Symbols with solid lines show the data for no simulation (□). Symbols with dotted lines correspond to expansion (○). Symbols with broken lines show the data for absolute thresholds for no simulation (□) and expansion (○). All expansion data was computed from the no-simulation data.

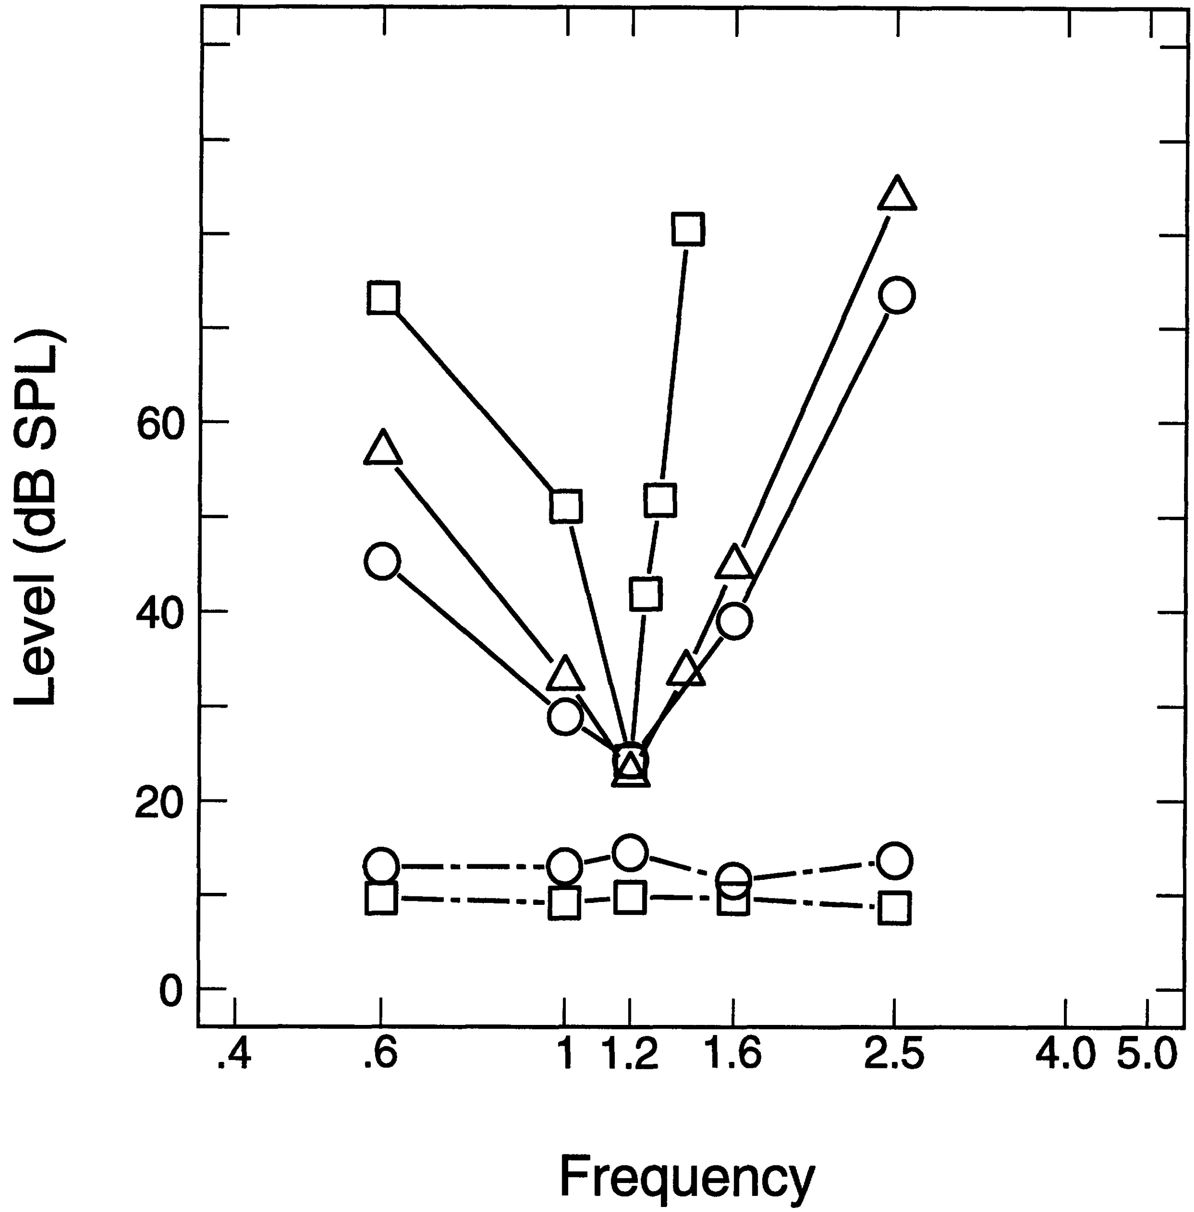

Figure 5-9: Effects of smearing on PTC's. Symbols with solid lines show the PTC's for no simulation (□), smearing = $\sqrt{3}$ (△), and smearing = 3 (○). Symbols with broken lines show the data for absolute thresholds with no simulation (□) and with smearing = 3 (○).

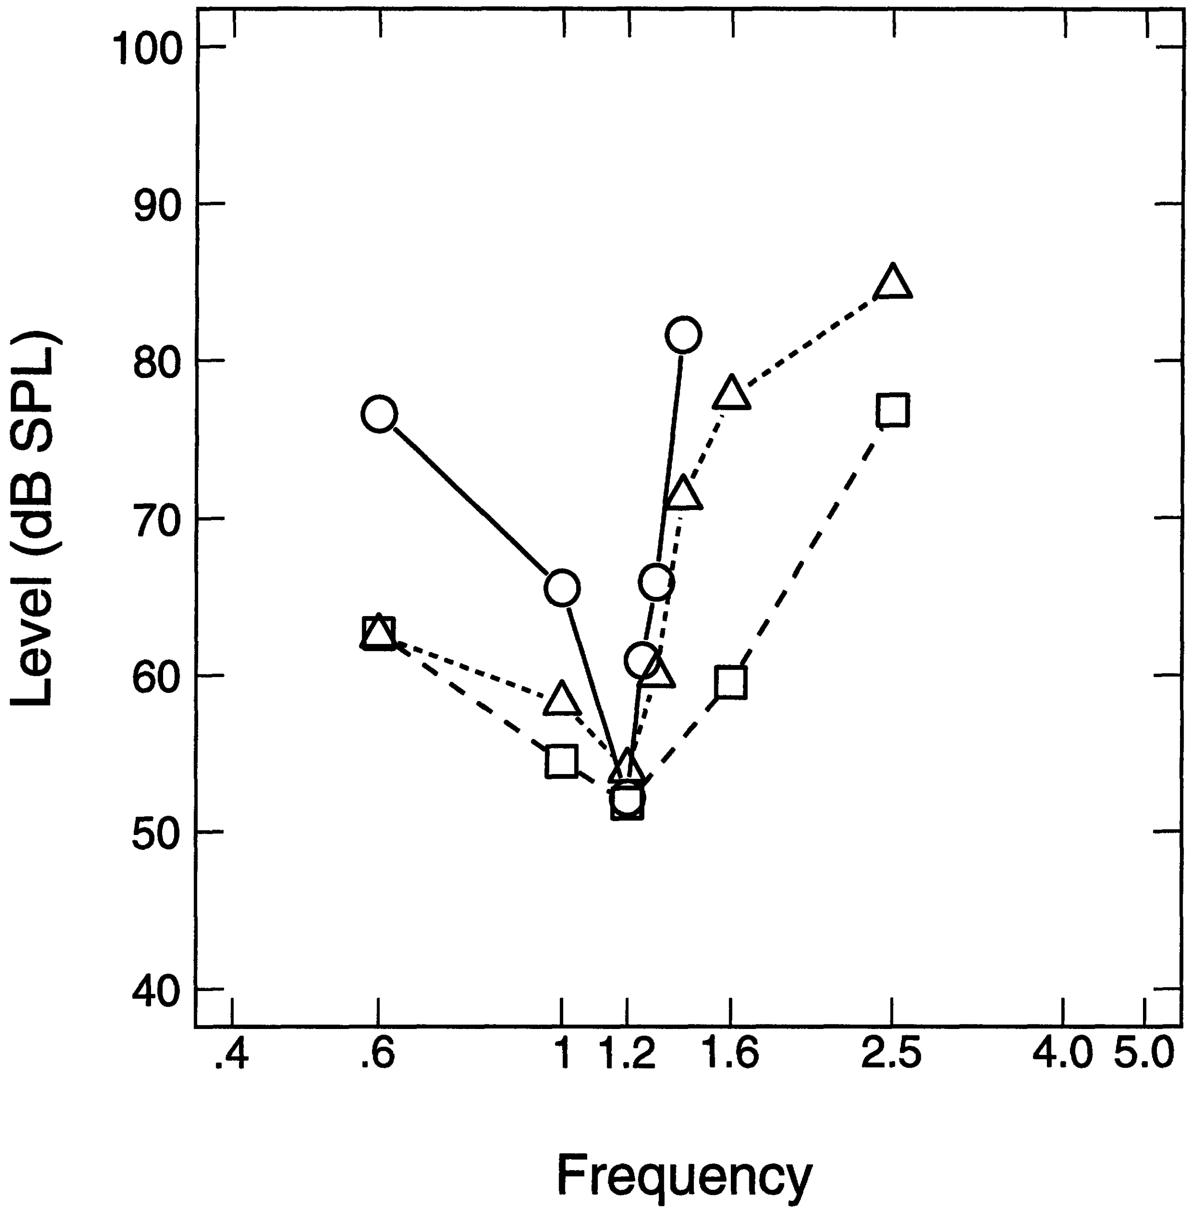

Figure 5-10: Effects of smearing on expansion. Symbols show the data for expansion alone (○), smearing = 3 plus expansion (□), and asymmetric smearing plus expansion (△).

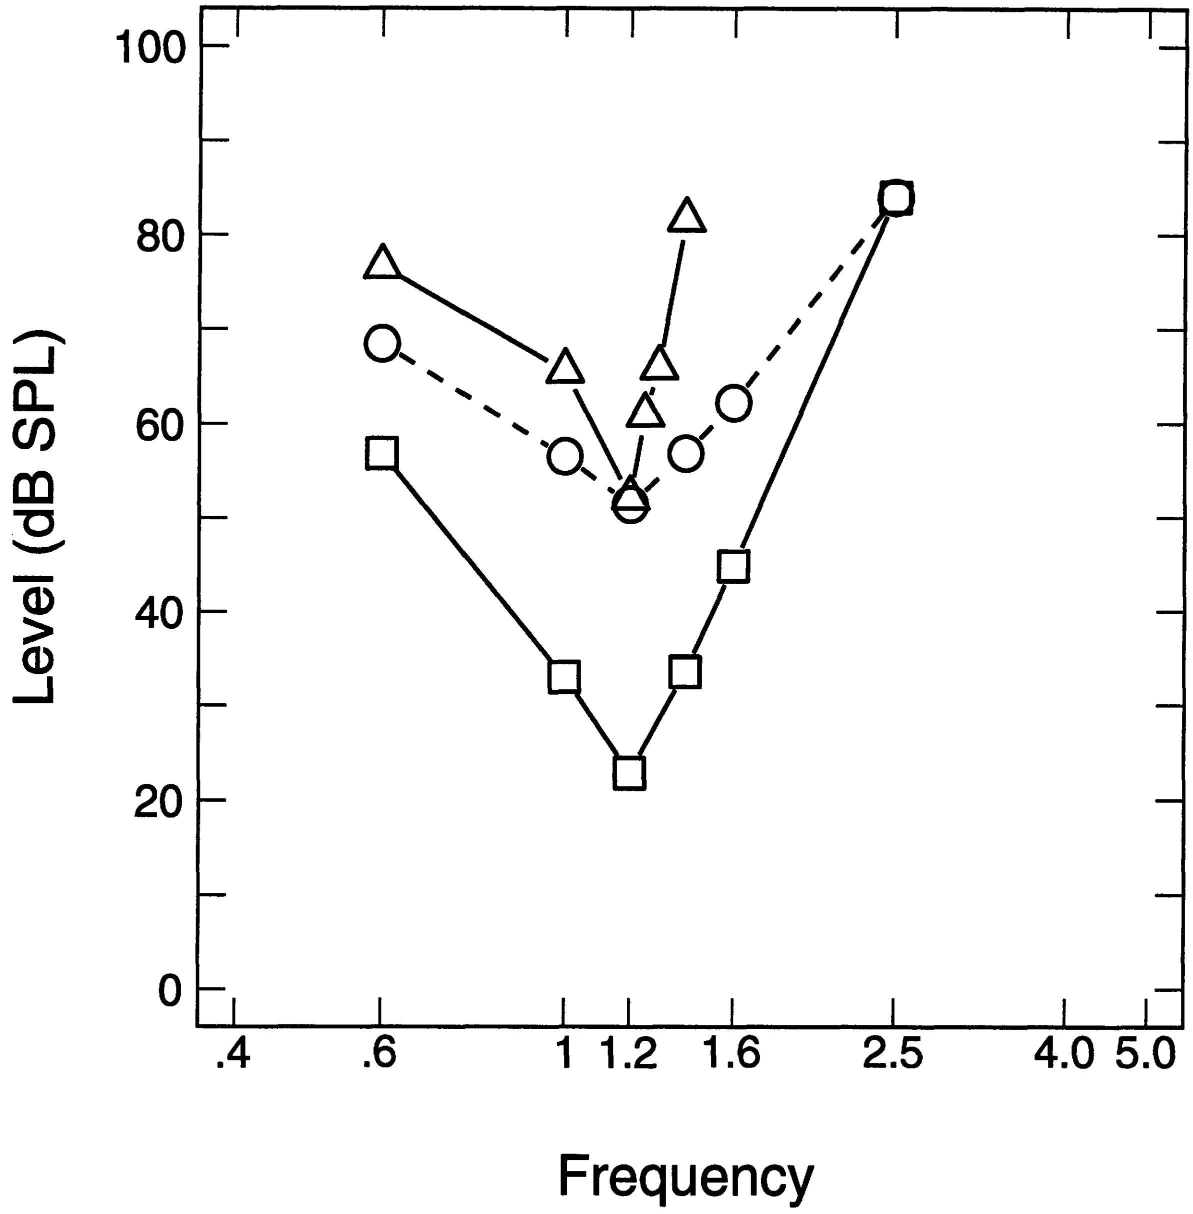

Figure 5-11: Comparison of simulations alone and in combination. Symbols show the data for smearing = $\sqrt{3}$ alone (□), expansion alone (△), and smearing = $\sqrt{3}$ with expansion (○).

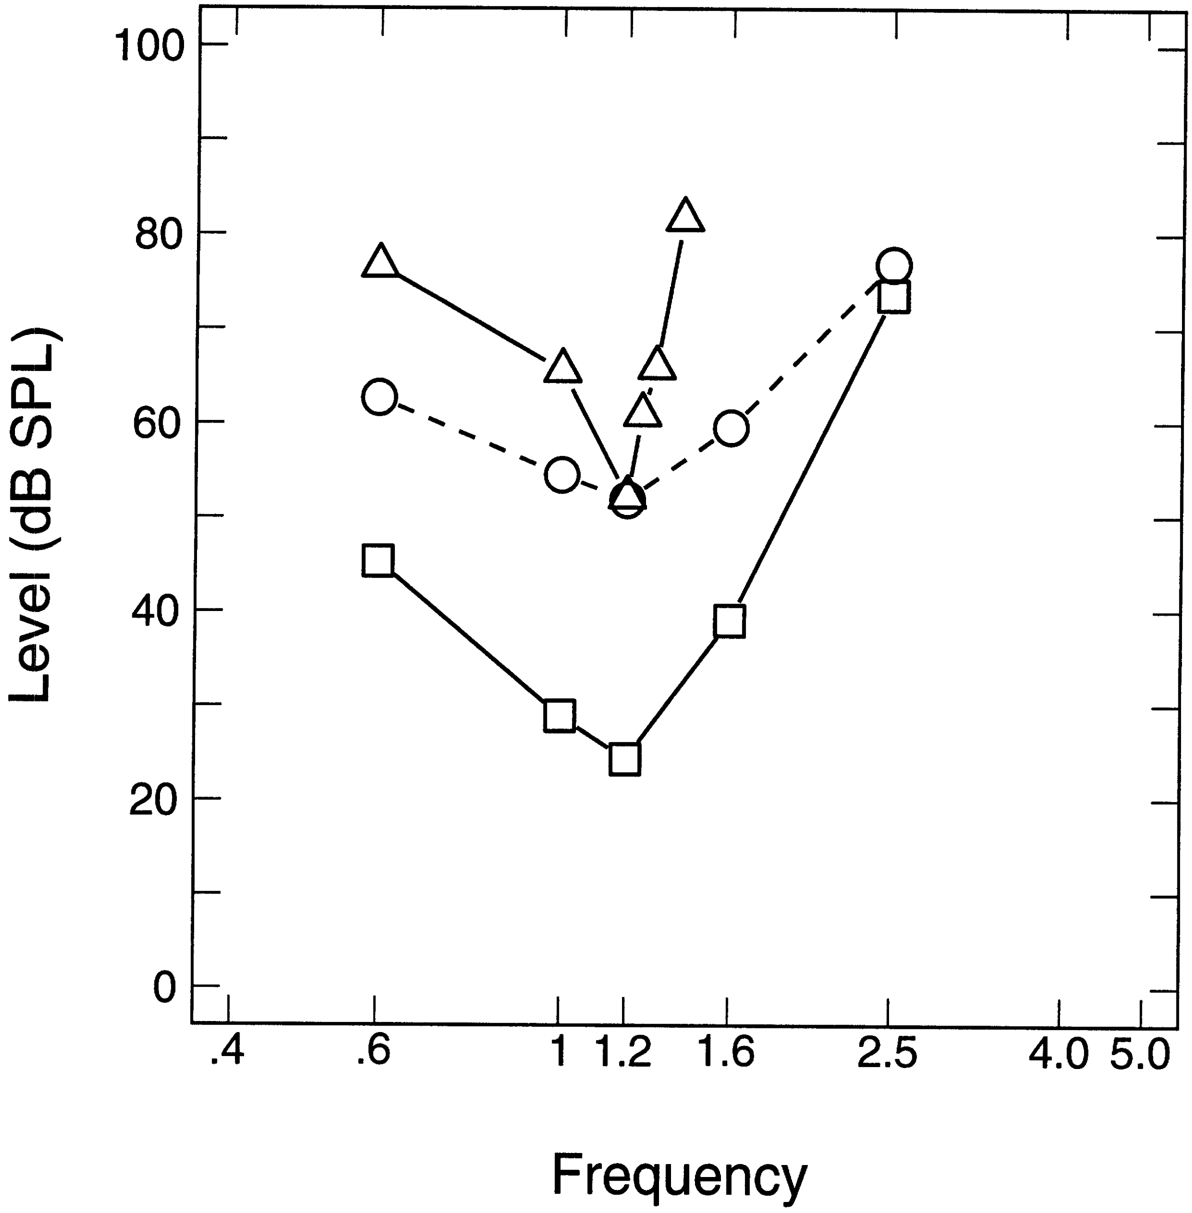

Figure 5-12: Comparison of simulations alone and in combination (continued). Symbols show the data for smearing = 3 alone (□), expansion alone (△), and smearing = 3 with expansion (○).

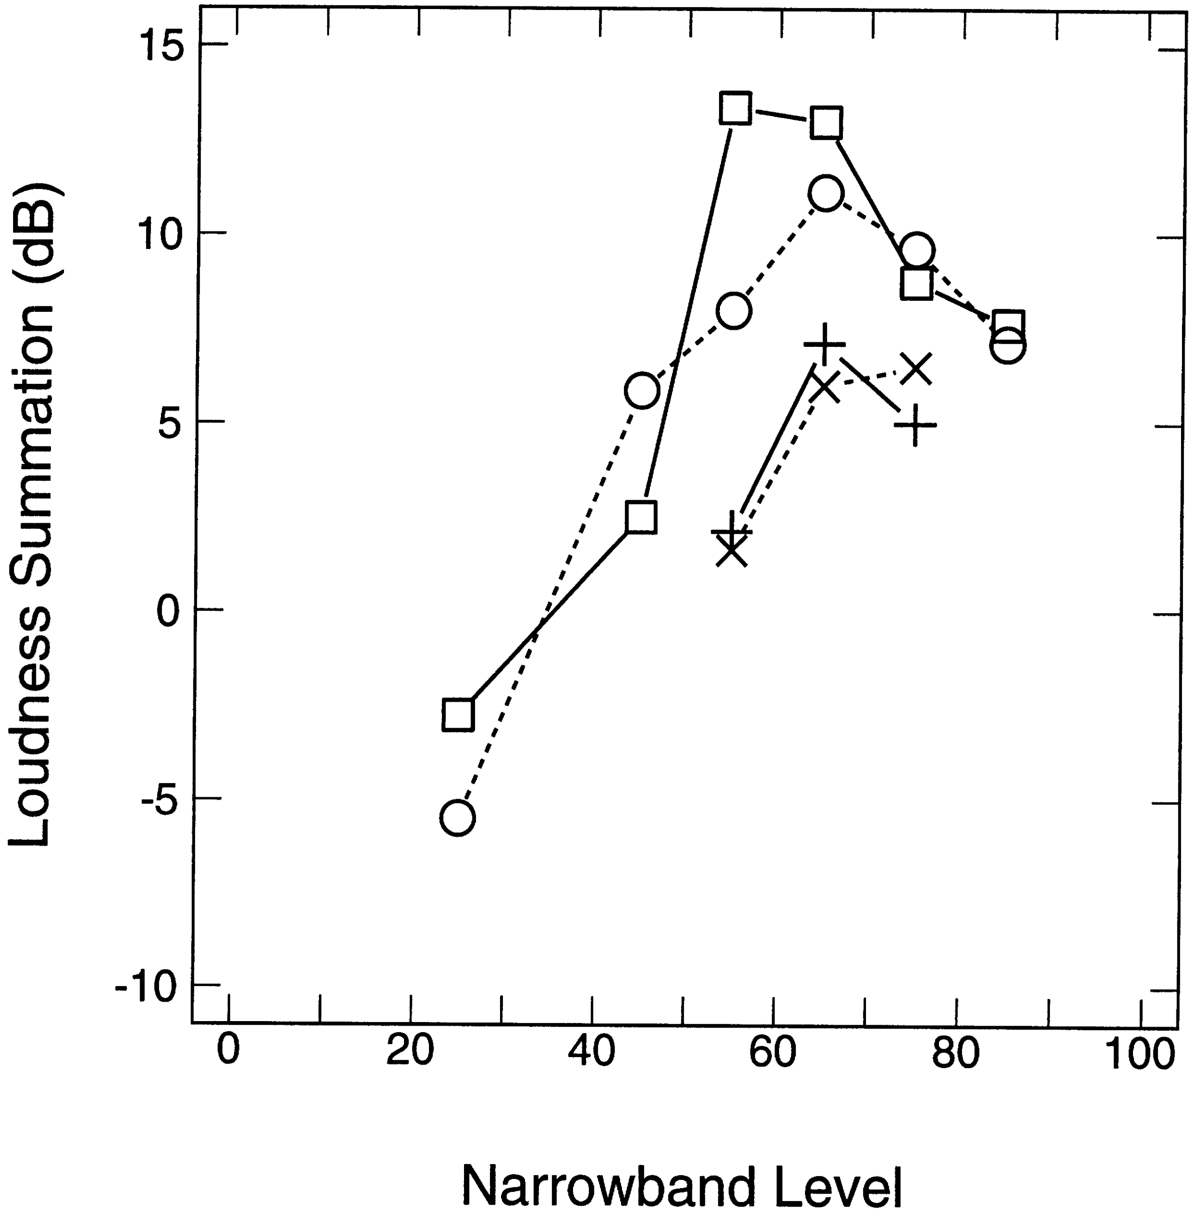

Figure 5-13: Loudness summation for subject LP. Symbols show the average matches for no simulation (□), expansion (+), smearing (○), and the simulations in combination (×).

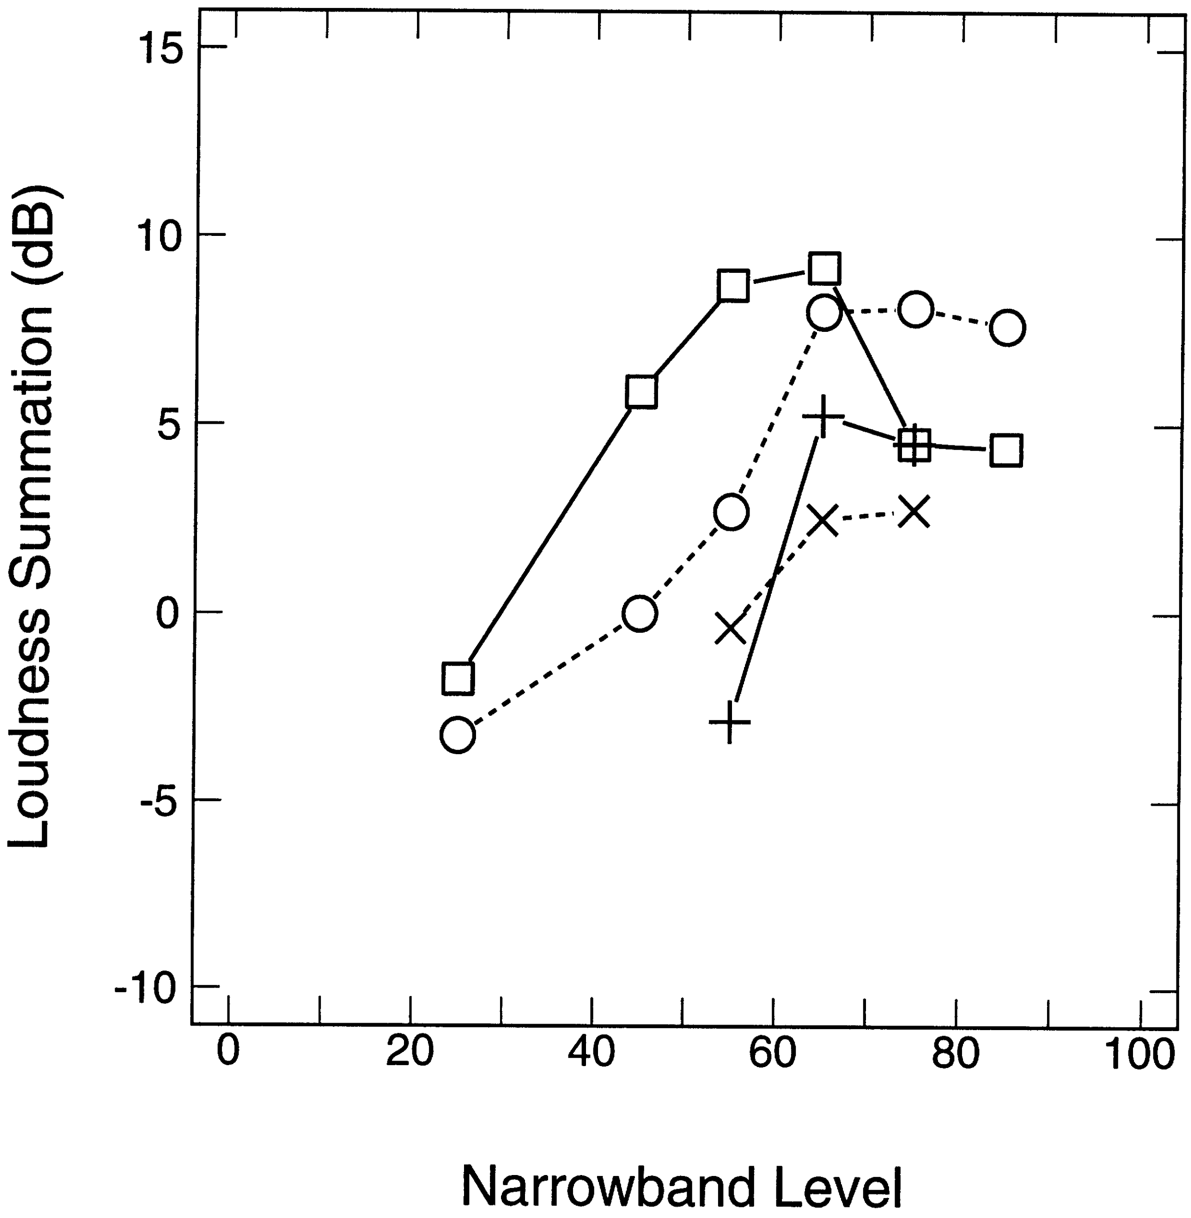

Figure 5-14: Same as Figure 5.13, for subject JK

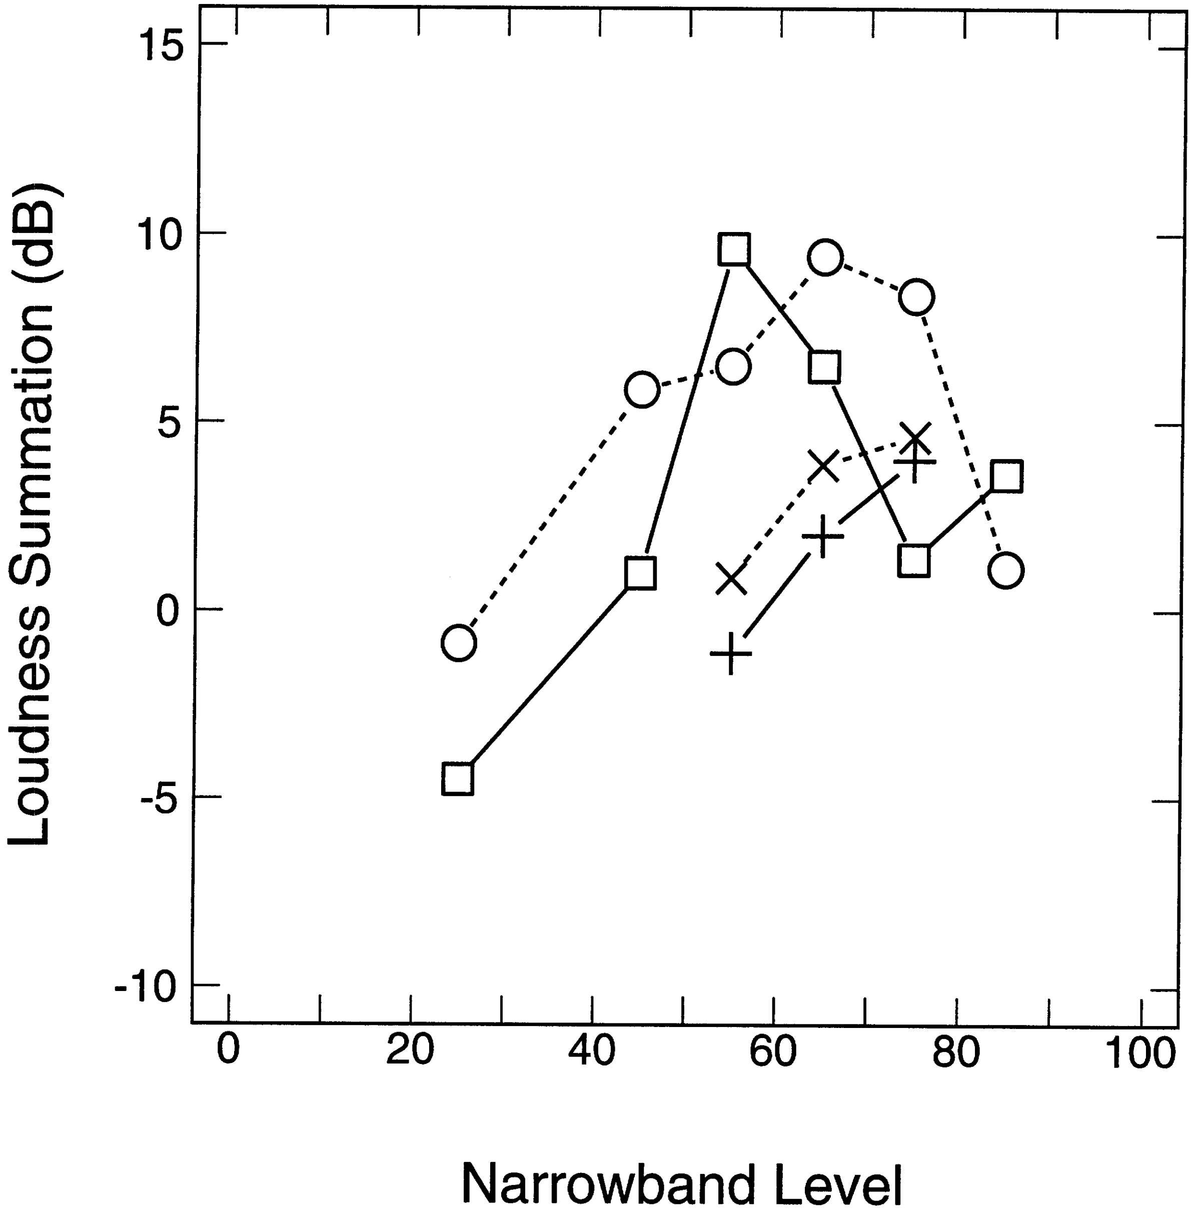

Figure 5-15: Same as Figure 5.13 for subject BS.

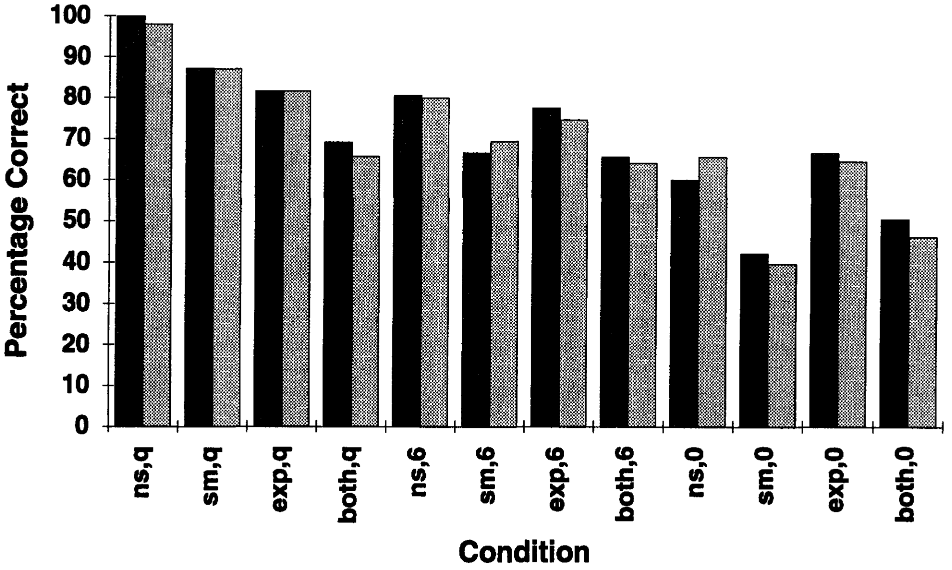

Figure 5-16: Consonant reception for both simulated normal subjects for all experimental conditions; simulation of mild/moderate hearing loss. Darker shade = subject IG. Lighter shade = subject AC. ns = no simulation, sm = smearing only, exp = expansion only, both = smearing plus expansion, q = quiet, 6 = 6 dB SNR, 0 = 0 dB SNR.

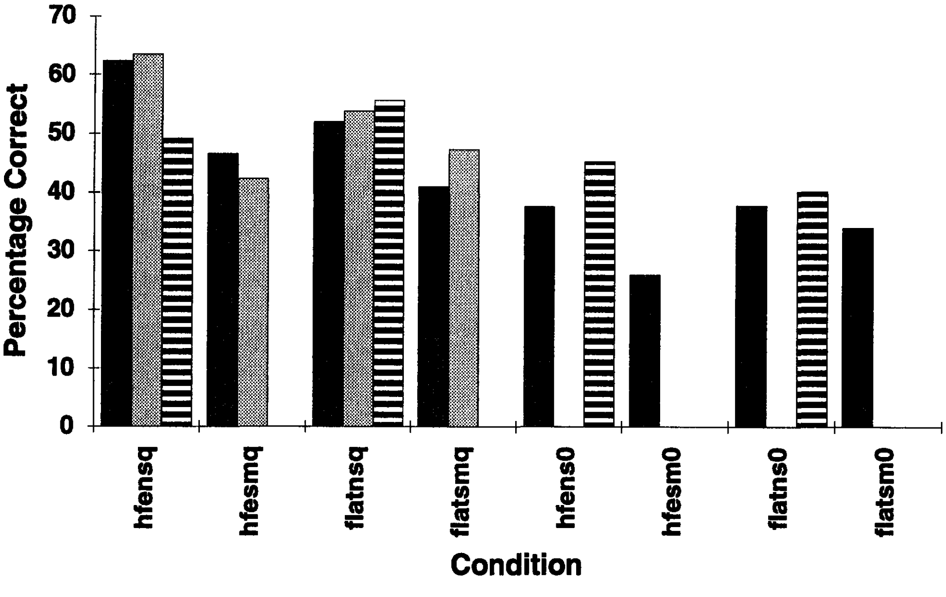

Figure 5-17: Consonant reception for simulated normal subjects and hearing impaired subject AL (from Duchnowski) for all experimental conditions; simulation of severe hearing loss. Dark shade = subject IG. Light shade = subject AC. Striped = subject AL. hfe = high frequency emphasis frequency response, flat = flat frequency response. ns = no smearing, sm = smearing, q = quiet, 0 = 0 dB SNR.

Chapter 6

Discussion

6.1 Simultaneous Masking

With expansion alone, narrowband noise masked thresholds are fairly constant below 800 Hz on the low frequency side and above 1600 Hz on the high frequency side (see Figure 5.1). This is basically in agreement with the data from Lum (1995a). It indicates that expansion has only small effects on frequency selectivity as measured by narrowband noise masking patterns. At 1000 Hz, however, the masking pattern is somewhat broadened. This result corresponds with the finding of Egan and Hake (1950) that the masking pattern on the low frequency side of the narrow band of noise broadens as the level of the noise is decreased.

From Figure 5.2 it is clear that frequency smearing significantly broadens the masking pattern for frequencies both above and below the masker. In the smearing alone condition, as the smearing bandwidth is increased from $\sqrt{3}$ to 3, the off-frequency masked thresholds (frequencies away from the noise band) are raised for both the high and low frequencies, but the increase is more pronounced in the high frequencies than the low frequencies. In both smearing conditions the thresholds continue to fall on both sides, as the frequency gets further away from the masking noise frequency. The unmasked hearing thresholds were in the normal range for all three subjects. Subject TP, however, had a minor head cold when quiet thresholds were measured that may have introduced a small conductive hearing loss. When the stimuli were smeared the quiet thresholds rose between 2 and 7 dB for almost all frequencies.

When expansion and smearing are combined, the effect is to raise the masked thresholds relative to either simulation alone. From one point of view, the effect of adding expansion to the smearing is to raise and flatten the masking pattern. (see Figures 5.3 and 5.4). When expansion is added, thresholds of frequencies far from the noise masker are raised more than thresholds of frequencies close to the noise masker on both the high and low frequency sides and for both degrees of smearing tested. From the other point of view, the effect of smearing on expansion is to raise considerably the thresholds close to the center frequency of the masker, and to raise the thresholds of frequencies distant from the masker to a lesser extent (see Figures 5.3- 5.5).

For the condition of asymmetric smearing with expansion, the masking pattern tended to resemble that of expansion alone on the low frequency side and expansion with smearing on the high frequency side. The low-frequency asymmetric masking pattern was a similar shape to that of the expansion simulation, but it was elevated a few dB relative to the expansion pattern. The high frequency masking pattern for asymmetric smearing plus expansion was within 3 dB of the thresholds from the symmetric smearing plus expansion case at all frequencies, but the asymmetric high frequency thresholds tend to be 1-3 dB below those of the symmetric smearing condition. The on-frequency thresholds (in the frequency region of the masking noise) do not seem to be affected much by any of the simulations, though they are slightly lower in the asymmetric smearing condition than in the other conditions.

Figures 5.6-5.7 show the masking patterns for several conditions plotted with data from Dubno and Schaefer (1991). Comparison of the masking patterns of the different simulations with that of the hearing impaired can be made both qualitatively and quantitatively. On a qualitative level, the shape of the low frequency masking pattern of the hearing impaired is not very well accounted for by any of the simulations. The hearing impaired pattern is flat from 1000 to 800 Hz and then drops from 800 Hz to 600 Hz. None of the simulation conditions produce this shape in the masking pattern. For probe frequencies in the band of the masker, masked thresholds of normal and impaired listeners are very similar, and all simulations appear to adequately account for this. On the high frequency side, the conditions that involve smearing plus expansion appear to give the best qualitative match to the hearing impaired data. Although none of the simulations provide the dip at 3000 Hz and subsequent rise at 4000 Hz, the broadening of the pattern at 1600 and 2500 Hz seems to be quite well accounted for in the conditions of symmetric and asymmetric smearing (with filter broadening by a factor of 3) with expansion. The pattern is too sloping for the filter broadening of $\sqrt{3}$ plus expansion (subject JK) and it falls much too quickly in the expansion alone condition.

A quantitative comparison of the simulation data with the hearing impaired data is complicated by several factors. For one, the expansion simulation carried out in these experiments was chosen to be similar to that used by Lum (1995a), rather than to match the characteristics of a specific hearing impaired subject. The resulting thresholds were elevated above normal by about 40 dB, to about 45 dB SPL. The hearing impaired in Dubno's study had thresholds lower than 40 dB SPL for frequencies below 1200 Hz. Thus the simulated masking patterns are expected to be too high, at least on the low frequency side. In addition, at 3000 and 4000 Hz the hearing impaired masking pattern dips and then rises sharply. It seems surprising that this result would be due to a uniform elevation and broadening of the tuning curves along the cochlea. (The simulations assumed a uniform filter broadening with frequency.) To obtain some measure of how well the different simulations matched the hearing impaired data, the squares of the differences between the simulated and hearing impaired thresholds were calculated and summed across two ranges of frequencies, (600-1000 Hz and 1600-2500 Hz). These numbers were divided by the number of terms in the sum (2 for the high frequency distances and 3 for the low frequency distances) and then square rooted. These distances, calculated for each subject separately for each condition, are rms measures (see Table 6.1).

Condition			Low freq. distance			High freq. distance		
E	S_L	S_H	JK	LP	TP	JK	LP	TP
+	1	1	5.5	4.0	8.8	8.3	6.1	4.1
-	$\sqrt{3}$	$\sqrt{3}$	10.6	-	-	16.1	-	-
+	$\sqrt{3}$	$\sqrt{3}$	9.5	-	-	2.0	-	-
-	3	3	6.1	3.5	8.3	6.5	3.4	4.1
+	3	3	11.4	11.8	15.4	2.8	3.4	5.8
+	3	1	8.0	5.9	9.9	1.6	2.4	3.1

Table 6.1: Distances relative to hearing impaired data; subjects JK, LP, and TP. E refers to expansion, S_L and S_H refer to broadening relative to normal on the low frequency and high frequency sides of the auditory filters, respectively. Low frequency distances correspond to the 600-1000 Hz region and high frequency distances correspond to the 1600-2500 Hz region.

For subject JK, the best low frequency condition (the smallest distance) was the expansion simulation alone. For the high frequency condition, the condition of asymmetric smearing plus expansion resulted in the smallest distance, though the filter broadening by $\sqrt{3}$ plus expansion was close behind. For subjects LP and TP, the best conditions were smearing without expansion for the low frequency side and asymmetric smearing with expansion for the high frequency side. This suggests that the asymmetric smearing plus expansion condition captures the masking pattern the best in the 1500 or so Hz band above the masker pass band. On the low frequency side, the thresholds of the conditions with smearing and expansion combined were too high relative to the hearing impaired data.

In all cases the best high frequency distance was considerably smaller than the best low frequency distance, suggesting that none of the simulations are doing a very accurate job simulating the low frequency simultaneous masking pattern.

An additional observation that can be made from this data concerns the additivity of the effects of the two simulations in combination relative to being in isolation. If the masking of a simulated condition is defined as the difference in the thresholds from the no-simulation condition and the simulated hearing loss condition, then it is clear that the sum of masking (in dB) produced by the two simulations individually is considerably greater than the masking due to the simulations in combination. This is true for all off-noiseband frequencies and for all subjects. If masking is measured in linear rather than logarithmic units, then the masking of the simulations in combination tends to be larger than the sum of the masking of the simulations in isolation. Neither model was able to describe how the simulations combine at the different frequencies.

6.2 Forward Masked PTC's

The PTC's in the unsimulated condition resemble normal PTC's found in the literature (Dubno and Schaefer, 1991). The high frequency side is very steep, and the low frequency side is very steep near the probe tone and flatter away from the probe frequency. The effect of expansion (Fig. 5.8) is to broaden both sides of the PTC, with more broadening on the low frequency side than the high frequency side. The basic shape of the PTC (concativity) is unchanged, however. Smearing also has a substantial broadening effect on both sides of the PTC (Figures 5.9, 5.10). It is clear that as the symmetric smearing bandwidth increased from $\sqrt{3}$ to 3, the broadening increased for both sides of the PTC. When expansion and smearing are combined (Figures 5.11, 5.12) the effect is to raise the tip and broaden both sides of the tuning curve. From the point of view of adding smearing to the expansion processing, the effect of smearing is to broaden the tuning curve considerably, while the tip stays in the same place (Figure 5.10). For the asymmetric smearing case, the PTC is greatly broadened on the low frequency side and it remains fairly steep on the high frequency side. The PTC tends to rise steeply close to the frequency of the probe, but flattens out at higher frequencies. In this sense, the asymmetric curve does not match the PTC of expansion alone on the high frequency side. (The PTC of expansion alone is steep on the high frequency side).

The absolute thresholds for the forward masked probe stimuli tend to be a few dB higher than for the simultaneous masking probe tones, most likely because the duration of the PTC tones is shorter. The thresholds for the smeared tones are generally a few dB higher than for the unsmeared tones. This was also found for absolute thresholds measured with the longer simultaneous masking probe tones.

6.3 Loudness Summation

As seen in Figures 5.13-5.15, all three subjects (BS, LP, JK) exhibit typical loudness summation curves in which the greatest loudness summation occurs at 55-65 dB SPL and that fall off on both sides of this level range. For all three subjects expansion reduces loudness summation at narrowband levels of 55 and 65 dB SPL. At 75 dB SPL, expansion seems to have little, if any, effect. The effect of smearing on loudness summation is not very clear from the data. For subject JK, smearing reduced loudness summation for levels below 65 dB SPL and increased loudness summation at levels above 70 dB SPL. With expansion, smearing had the opposite effect, increasing loudness summation at low levels and decreasing loudness summation at higher levels. For subject BS, smearing increased loudness summation at all levels except 55 and 85 dB SPL, and when expansion was simulated, smearing increased loudness summation slightly for all levels. For subject LP, smearing had very little effect at any levels, with and without expansion. However for the non-expansion case, smearing tended to reduce loudness summation slightly.

The finding that smearing has little effect on loudness summation contrasts with the notion (e.g. Florentine et al., 1979)) that widened filter bandwidths are necessary to explain reduced loudness summation.[1] The finding that expansion alone produces a significant reduction in loudness summation is in agreement with the model of Launer et. al (see section 2.3).

6.4 Consonant Reception

6.4.1 Mild/Moderate Recruitment Simulation

With the simulations of a 40 dB flat loss and/or a symmetrical spectral smearing bandwidth of 1, the scores for the two subjects tested are fairly similar (see Figure 5.16). In quiet, expansion had a more deleterious effect on reception than did smearing. Relative to no processing, scores declined by about 13% with smearing and by about 18% with expansion. The effects of the simulations were roughly additive: scores declined by about 32% for the simulations in combination.

When speech shaped noise was introduced at a speech to noise ratio of 6 dB, the effects of smearing became larger than the effects of expansion. Smearing reduced the scores by 14 percentage points, while expansion reduced the scores by only 4 percentage points. The effects of the simulations were not additive, since in combination, scores were almost identical to the scores for smearing alone. Thus although the reduction in intelligibility caused by expansion in the no-smearing condition was small, the reduction was even smaller in the smearing condition.

When speech shaped noise was added at an SNR of 0 dB, smearing had a larger effect. The effect of expansion was negligible in subject AC, while in subject IG, expansion

---

[1] It is possible, but unlikely, that the discrepancy with the model results from the protocol used for the loudness matching experiment. The loudness summation data of Florentine et al. were obtained by alternating the two noise stimuli continuously, whereas a 2 interval protocol was used in these experiments.

improved scores by 6 percentage points. The scores for the simulations in combination were higher (by 7-8 percentage points) than those for smearing alone. Thus for subject IG, the error scores were additive, whereas for subject AC the simulations in combination produced a higher score than would be produced if the simulations were additive. Figure 6.1 shows the effects of the different simulations, for each SNR and each simulation condition, averaged over the two subjects.

In order to analyze how the different simulations affected transmission of different acoustic features, a SINFA analysis (Wang and Bilger, 1973) was performed on the confusion matrices of the experiments. For each of the twelve experimental conditions (SNR (quiet,6,0), smearing (y/n), expansion (y/n)) a confusion matrix was generated for each subject. (Each consonant in the confusion matrix had 15-21 presentation for subject AC and 9 presentations for subject IG). SINFA was used to determine transmission of three features: voicing (2 categories), manner (5 categories), and place (6 categories). See Table 6.2.

	Consonant											
Feature	P	T	K	B	D	G	F	TH	S	SH	V	XH
Voicing	-	-	-	+	+	+	-	-	-	-	+	+
Manner	1	1	1	1	1	1	2	2	2	2	2	2
Place	1	4	6	1	4	6	2	3	4	5	2	3
	Z	ZH	CH	J	M	N	R	W	L	Y	H	WH
Voicing	+	+	-	+	+	+	+	+	+	+	-	-
Manner	2	2	3	3	4	4	5	5	5	5	2	2
Place	4	5	5	5	1	4	5	1	4	5	6	6

Table 6.2: Feature table for the set of twenty four consonants

Relative transmission scores, after an arcsine transformation, were subject to ANOVA to determine whether each simulation parameter had a statistically significant effect on transmission of each feature. First, ANOVA was run with subject as a factor in addition to smearing, expansion, and SNR. In order to determine whether subject is a significant factor, the null hypothesis of subject having no interaction with other factors was tested with a .05 confidence rating. The result of ANOVA indicated that the null hypothesis could not be rejected, indicating that the presentation conditions affected feature transmission similarly for subjects IG and AC. For the remaining analysis, the percent transmission scores for the two subjects were averaged for each feature for each experimental condition. An ANOVA of these averaged scores was carried out again using smearing, recruitment, and SNR as factors to determine which simulation conditions have significant effects on transmission of the three phonemic features. Using a significance criterion of 5% (as for the ANOVA above), the following observations were made: For voicing, SNR, expansion, and SNR × expansion are statistically significant factors; for manner, SNR, expansion, smearing, and SNR × expansion are statistically significant factors; and for place, SNR, smearing, and SNR × expansion are statistically significant factors. (See Table 6.3). Based on the 5% significance criterion, it can thus be concluded that smearing has an effect on transmission of place information, expansion has an effect on transmission of voicing information, and both simulations have an effect on the transmission of manner information.

Feature	SNR	Expansion	Exp. × SNR	Smearing
Voicing	0.007	0.042	0.036	0.122
Manner	0.001	0.030	0.012	0.015
Place	0.001	0.062	0.034	0.002

Table 6.3: Effects of the simulations on features of speech. Numbers indicate the probability, given the hypothesis that the factor has no effect on the transmission of the phonemic feature, that the scores were as likely or less likely than what they were. (Low numbers suggest more confidence in rejecting this hypothesis).

6.4.2 Severe Recruitment Simulation

Figure 5.17 shows the scores achieved by each subject and each condition that was used in the simulation of Duchnowski's hearing impaired subject AL. It also shows the scores from subject AL. For the quiet conditions, the scores from the two simulated-normal subjects agree with each other fairly well. For the 0 dB SNR condition, data was obtained from only one subject. It can be seen that in quiet, smearing reduces performance by 16-21 percentage points for the high frequency emphasis (HFE) condition and by 5-11 percentage points for the flat condition. Looked at a different way, in both simulated normal subjects, high frequency emphasis helps to improve the score by about 10 percentage points when no smearing is applied. However, when smearing is applied, the benefit of high frequency emphasis is greatly reduced: in subject IG, the improvement is only 5 points and in subject AC HFE lowers the score by 5 points.

When speech shaped noise was presented at 0 dB SNR, smearing reduced scores by 11 points when HFE is applied and by 4 points in the flat frequency response condition; thus it again appears that smearing is more significant in the high frequency emphasis condition than in the flat condition. For the 0 dB SNR condition, HFE does not improve scores when there is no smearing. When there is smearing, HFE reduces scores. This is in contrast to the case in quiet, where HFE improves scores with no smearing, but has little effect when smearing is applied (HFE is thus more beneficial in quiet than in noise). An ANOVA with the factors subject, frequency response, and smearing found no significant changes for the features voicing, manner, and place, or for percent correct for either subject AC or IG.

Comparing the data from the simulated normals with that of subject AL (Figure 5.17), it is seen that in the quiet, flat condition, subject AL performed comparably to the simulated normals when the simulation was expansion alone (without smearing). In the HFE case, scores for AL were lower than those of the simulated normals without smearing, but higher than those of the simulated normals when smearing was applied.

Scores for AL decreased by 5 percentage points when HFE was introduced. This exactly matched the results of subject AC in the smearing condition. In the 0 dB SNR condition, subject AL achieved higher scores with HFE than flat. This result was not observed in the data from subject IG in either the smearing or the no smearing condition. Thus smearing results in data that is closer to that of the hearing impaired in quiet than in 0 dB SNR noise.

In comparing the data from this experiment using expansion alone with the simulated normal data of Duchnowski's subjects, it was found that scores are similar in the two studies in quiet for both frequency responses, and for the 0 dB SNR condition with the flat frequency response. However, for the condition of 0 dB SNR and HFE, subject IG scored 38% in this study, whereas the simulated normals in Duchnowski's study averaged 53% (Table 6.4).

	Quiet		0 dB SNR	
	Flat	HFE	Flat	HFE
Subject AL	55.7	49.2	40.1	45.3
SN's (Duchnowski)	50.5	60.0	42.7	53.0
Subject AC	53.8	63.5	-	-
Subject IG	52.0	62.4	37.8	37.6

Table 6.4: Percentage correct for the simulated normals in this study (AC, IG) without smearing, those for the simulated normals (SN's) in Duchnowski's study (1989), and those for hearing impaired subject AL.

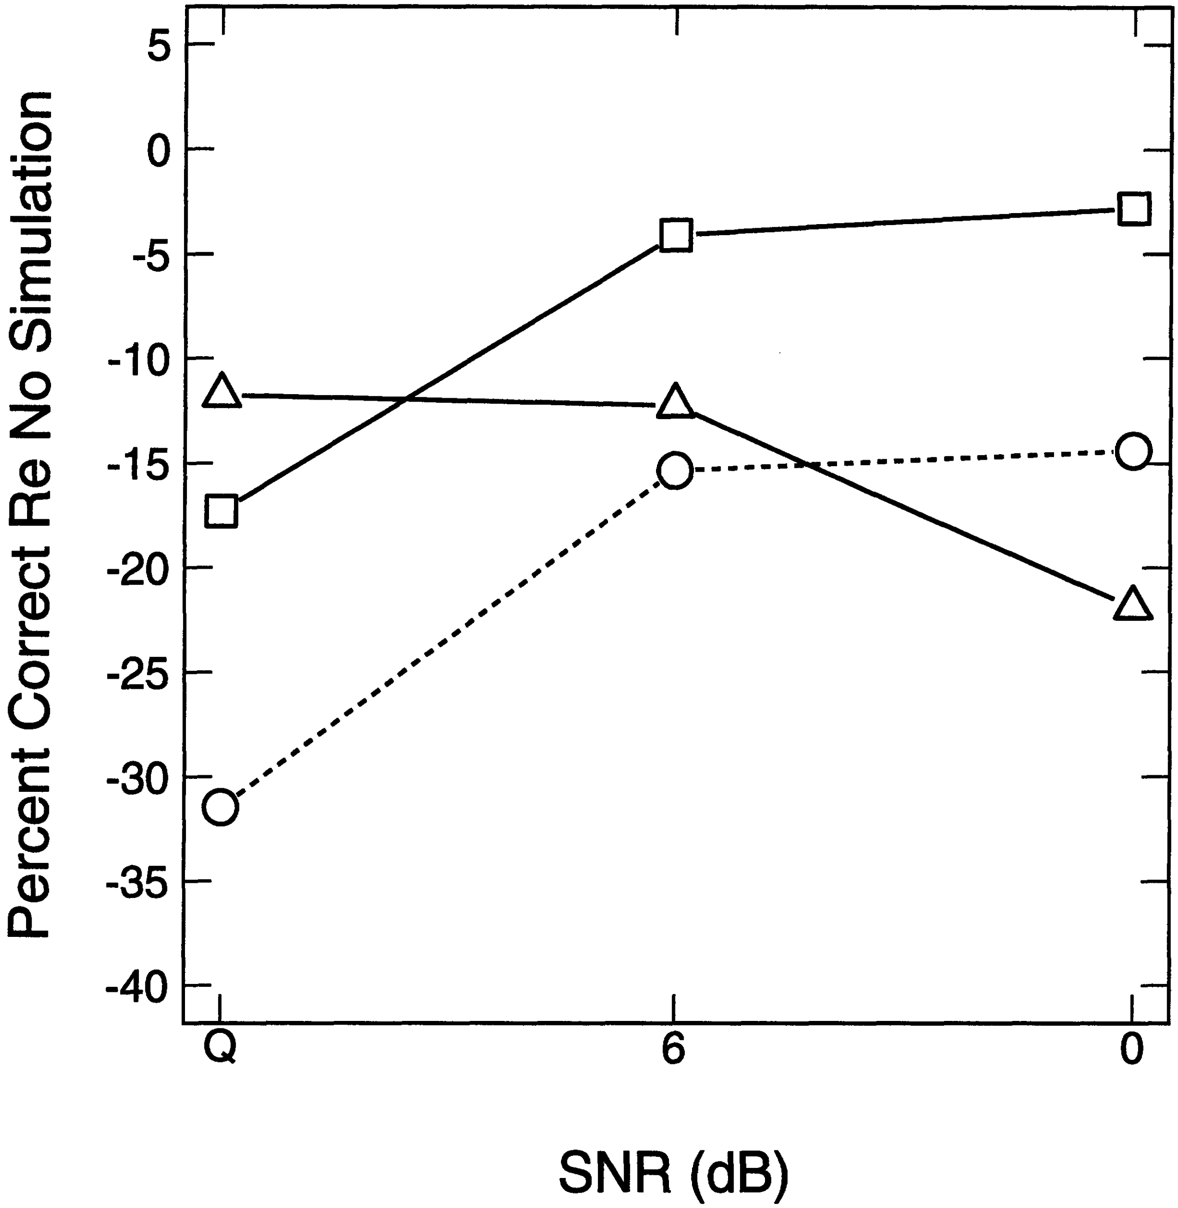

Figure 6-1: Effect of the simulations on performance. The plots show the difference in percent correct between the three different simulation conditons and the no simulation conditon. Low values indicate a large detrimental effect of the simulation on intelligibility. Symbols show the data for expansion (□), smearing (△) and expansion plus smearing (○).

Chapter 7

Conclusions

The focus of this section is to examine the degree to which the simulations studied in this thesis achieved the goals discussed in Chapter 1. These goals were 1) to compare the performance of simulated normals with those of the hearing impaired, 2) to observe the effects of different aspects of hearing loss in isolation, and 3) to suggest signal processing schemes for hearing aids.

7.1 Comparing simulated normal data with hearing impaired data

The results of the tests of frequency selectivity provided a test of Moore's notched noise model of auditory filters. If the assumptions of the model are correct and if the auditory filters of hearing impaired individuals are broadened by a factor of 3 (or $\sqrt{3}$), then the simultaneous masking data from the simulated normals should match very closely the data for the hearing impaired. The data indicates that the high frequency side of the masking pattern is fairly well matched up to 2500 Hz for the smearing plus expansion condition, but the low frequency side is not very well matched for any simulation condition.

The discrepancies in the results from the simulated normals and the hearing impaired may not be due entirely to inaccuracies in Moore's model. No attempt was made to simulate precisely the auditory filters of the impaired subject in the Dubno study. Rather, impaired filters 3 or $\sqrt{3}$ times the width of the normal filters were simulated. The expansion function used may not have corresponded exactly to that of the impaired subject. In addition, Moore's auditory filter model was simplified in these experiments in that level dependence of the filter shape was ignored. Nevertheless, the possibility remains that auditory filter data based on notched noise does not account well for narrowband noise masking data.

The consonant recognition experiments were intended to determine whether combining smearing and expansion would result in scores that more closely resembled those of the hearing impaired than expansion processing alone. The poorer performance in the HFE condition than the flat condition in quiet obtained for the hearing impaired subject AL was reproduced by the combination of simulations for normal hearing subject AC, but not for subject IG. The result that HFE provided little benefit to intelligibility was not obtained with simulated recruitment alone in either this study or Duchnowski's study.

In the 0 dB SNR condition, the simulations do not reproduce the improvement in performance that impaired subject AL achieved with HFE. In the HFE, quiet condition, the simulations in combination result in scores closer to those of impaired subject AL than did the expansion simulation alone. However, in the other conditions (flat-quiet, HFE-0 SNR, flat-0 SNR) expansion alone produced scores closer to those of subject AL than did the simulations in combination. Neither simulation captures the effect of AL's hearing impairment on different SNR and frequency response conditions very accurately.

7.2 Isolating the effects of Expansion and Smearing

Several observations can be made from the simulation's ability to separate the effects of poor frequency selectivity and recruitment. The loudness summation experiment separates these two effects to test Zwicker's model of loudness perception. The results of the experiment indicated that expansion does reduce loudness summation but that spectral smearing has little effect on loudness summation. This latter finding contrasts with the model of Florentine et al. (1979) which predicts that both loudness recruitment and reduced frequency selectivity contribute to reducing loudness summation in the hearing impaired. Rather, the data supports the claim of Launer et al. (1997) that elevated thresholds with recruitment account for the reduced loudness summation in the hearing impaired.

The consonant recognition experiment with a 40 dB loss simulated by expansion indicates that in quiet, both poor frequency selectivity and recruitment may adversely affect speech reception and that their effects are additive. As the SNR decreases the effects of frequency smearing become more pronounced, while expansion becomes less detrimental to intelligibility. It is possible, however, that the smearing parameters used for these experiments simulate a hearing impairment that is more severe than that simulated by the expansion simulation parameters used in the recruitment simulation.

The effects of expansion and smearing on frequency selectivity have been discussed previously (Chapter 6). Both expansion and smearing broaden the PTC's, though symmetric smearing has a larger effect on the high frequency side. For the narrowband noise masking experiment, the effect of expansion is to raise the masking patterns outside the noise band, whereas spectral smearing tends to broaden the masking patterns. The narrowband noise masking patterns also show that when expansion is present, smearing has a greater effect for signal frequencies closer to the masker frequencies than for frequencies far away from the band of noise (Figure 5.5). If expansion is not applied, then smearing has a significant effect at all signal frequencies (Figure 5.2). Conversely, when smearing is present, expansion has a greater effect for signal frequencies far from the masker frequencies than for frequencies close to the masker frequency (Figures 5.3, 5.4). If smearing is not present then expansion has a significant effect at all signal frequencies (Figure 5.1).

7.3 Relevance to Hearing Aid Processing

The above finding- that the effect of one aspect of hearing loss becomes significant at all frequencies instead of in a limited range of frequencies when another aspect of hearing loss is removed- may have implications for hearing aids. If a hearing aid is designed to compensate for one phenomenon (e.g. recruitment via amplitude compression), then poor frequency selectivity may now be more significant than it was with recruitment, for signal frequencies far away from the frequencies of any masker, or noise. Likewise, if a spectral enhancement scheme improved frequency selectivity then threshold elevation/recruitment would be more significant than it was before, for signal frequencies close to the masker frequencies.

An additional observation from the data that has potential implications for hearing aids is that expansion and frequency smearing affected different features of speech to different extents. Expansion has more effect on the perception of voicing, while smearing has more impact on the perception of place. In listening environments where the impaired listener is able to take advantage of lipreading, much of the information about place can be obtained visually, but voicing cannot. The data from this study suggests that in such a situation, loudness recruitment would have more detrimental consequences to the listener than would reduced frequency selectivity.

7.4 Future Work

Future work could proceed along the directions of 1) conducting more experiments with the existing hearing loss simulations, 2) improving the hearing loss simulations, and 3) observing the effects of different hearing aid processing schemes with the different simulations.

The first area of experiments could include studies of temporal resolution, intensity discrimination, and speech reception under conditions of different types of noise under the two simulations. The effects of expansion and smearing on temporal resolution could be investigated with such experiments as gap detection or forward masked recovery times. It might prove interesting to test the effects of the simulations on other auditory processes such as binaural phenomena and sound localization.

Work on improving the hearing loss simulations could involve explicitly simulating poor temporal resolution (Drullman, Festen and Plomp, 1994), or simulating poor frequency selectivity via an expansion algorithm that allows the gain in one channel to be influenced by the levels in the other channels (Duchnowski, 1989). In addition, most hearing loss simulation studies have essentially simulated outer hair cell loss, without simulating the loss of inner hair cells or auditory nerve fibers (in which case no acoustic information is sent to the brain in a given region of the basilar membrane). It is not clear how this type of pathology could be simulated. The third area of future work might involve comparing the effects of linear and compressive hearing aids on the different simulations, and perhaps also look at the impact of spectral enhancement algorithms on the simulations.

Bibliography

Baer, T. and Moore, B. C. J., (**1993**). "Effects of spectral smearing on the intelligibility of sentences in noise", *The Journal of the Acoustical Society of America*, **94**(3),1229–1241.

Dreschler, W. A. and Plomp, R., (**1980**). "Relations between psychophysical data and speech perception for hearing-impaired subjects i.", *The Journal of the Acoustical Society of America*, **68**(6),1608–1615.

Dreschler, W. A. and Plomp, R., (**1985**). "Relations between psychophysical data and speech perception for hearing-impaired subjects ii.", *The Journal of the Acoustical Society of America*, **78**(4),1261–1270.

Drullman, R., Festen, J. M., and Plomp, R., (**1994**). "Effect of temporal envelope smearing on speech reception.", *The Journal of the Acoustical Society of America*, **95**(2),1053–1064.

Dubno, J. R. and Schaefer, A. B., (**1991**). "Frequency selectivity for hearing impaired and broadband-noise-masked normal listeners.", *The Quarterly Journal of Experimental Psychology*, **43A**(3),543–564.

Duchnowski, P., (**1989**). *Simulation of Sensorineural Hearing Impairment*, Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Egan, J. P. and Hake, H. W., (**1950**). "On the masking pattern of a simple auditory stimulus", *The Journal of the Acoustical Society of America*, **22**(5),622–630.

Festen, J. M. and Plomp, R., (**1990**). "Effects of fluctuating noise and interfering speech on the speech-reception thershold for impaired and normal hearing", *The Journal of the Acoustical Society of America*, **88**,1725–1736.

Florentine, M., Buus, S., and Hellman, R. P., (**1997**). "A model of loudness summation applied to high frequency hearing loss", *Modeling Sensorineural Hearing Loss*, ,187–198.

Florentine, M., Buus, S., Scharf, B., and Zwicker, E., (**1980**). "Frequnecy selectivity in normally hearing and hearing impaired observers", *Journal of Speech and Hearing Research*, **23**,646–669.

Florentine, M. and Zwicker, E., (**1979**). "A model of loudness summation applied to noise-induced hearing loss", *Hearing Research*, **1**,121–132.

Glasberg, B. R. and Moore, B. C. J., (**1990**). "Derivation of auditory filter shapes from notched-noise data", *Hearing Research*, **47**,103–138.

Jesteadt, W., (**1980**). "An adaptive procedure for subjective judgements", *Perception and Psychophysics*, **28**(1),85–88.

Launer, S., Hohmann, V., and Kollmeier, B., (**1997**). "Modeling loudness growth and loudness summation in hearing- impaired listeners", *Modeling Sensorineural Hearing Loss*, ,175–185.

Levitt, H., (**1970**). "Transformed up-down methods in psychoacoustics", *The Journal of the Acoustical Society of America*, **49**(2),467–477.

Liberman, M. C. and Dodds, L. W., (**1984**). "Single-neuron labeling and chronic cochlear pathology iii. Stereocilia damage and alterations of threshold tuning curves", *Hearing Research*, **16**,55–74.

Lum, D. S., **(1995a)**. "Evaluation of a hearing loss simulator", Advanced Undergraduate Project, Massachusetts Institute of Technology.

Lum, D. S., **(1995b)**. *Further evaluation of a hearing loss simulator*, Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Lum, D. S., **(1996)**. "Lab guide to using the dsp-96 board under linux", Local lab documentation.

Miskolczy-Fodor, F., **(1960)**. "Relation between loudness and duration of tonal pulses iii. Response in cases of abnormal loudness function.", *The Journal of the Acoustical Society of America*, **32**(4),486–492.

Moore, B. C. J., **(1995)**. *Perceptual Consequences of Cochlear Damage*, *Oxford Psychology Series*, Oxford University Press.

Moore, B. C. J., Baer, T., Glasberg, B. R., and Vickers, D. A., **(1996)**. "Simulations of the effect of hearing impairment on speech perception", *Workshop on the auditory basis of speech perception, Keele University*, ,273–278.

Moore, B. C. J. and Glasberg, B. R., **(1987)**. "Formualae describing frequency selectivity as a function of frequency and level, and their use in calculating excitation patterns", *Hearing Research*, **28**,209–225.

Moore, B. C. J. and Glasberg, B. R., **(1993)**. "Simulation of the effects of loudness recruitment and threshold elevation on the intelligibility of speech in quiet and in a background of speech", *The Journal of the Acoustical Society of America*, **94**(4),2050–2062.

Moore, B. C. J., Glasberg, B. R., and Vickers, D. A., **(1995)**. "Simulation of the effects of loudness recruitment on the intelligibility of speech in noise", *British Journal of Audiology*, **29**,131–143.

Nejime, Y. and Moore, B. C. J., (**1997**). "Simulaton of the effect of threshold elevation and loudness recruitment combined with reduced frequency selectivity on the intelligibility of speech in noise", *The Journal of the Acoustical Society of America*, **102**(1),603–615.

Phillips, D. P., (**1987**). "Stimulus intensity and loudness recruitment: Neural correlates", *The Journal of the Acoustical Society of America*, **82**(1),1–12.

Robles, L., Ruggero, M. A., and Rich, N. C., (**1986**). "Basilar membrane mechanics at the base of the chinchilla cochlea i. Input-output functions, tuning curves, and response phases", *The Journal of the Acoustical Society of America*, **80**(5),1364–1374.

Sexton, M. G., (**1995**). "C language development on the dsp-96 dsp96002 floating-point attached processor board", Local lab documentation.

Steinberg, J. C. and Gardener, M. B., (**1937**). "The dependence of hearing impairment on sound intensity", *The Journal of the Acoustical Society of America*, **9**(1),11–23.

ter Keurs, M., Festen, J. M., and Plomp, R., (**1992**). "Effect of spectral envelope smearing in speech reception i", *The Journal of the Acoustical Society of America*, **91**(5),2872–2880.

ter Keurs, M., Festen, J. M., and Plomp, R., (**1993a**). "Effect of spectral envelope smearing on speech reception ii", *The Journal of the Acoustical Society of America*, **93**(3),1547–1552.

ter Keurs, M., Festen, J. M., and Plomp, R., (**1993b**). "Limited resolution of spectral contrast and hearing loss for speech in noise", *The Journal of the Acoustical Society of America*, **94**(3),1307–1314.

Villchur, E., (**1977**). "Electronic models to simulate the effect of sensory distortions on speech perception by the deaf", *The Journal of the Acoustical Society of America*, **62**(3),665–674.

Wang, M. D. and Bilger, R. C., (**1973**). "Consonants confusions in noise: A study of perceptual features", *The Journal of the Acoustical Society of America*, **54**(5),1248–1266.

Zurek, P. M. and Delhorne, L. A., (**1987**). "Consonant reception in noise by listeners with mild and moderate sensorineural hearing impairment", *The Journal of the Acoustical Society of America*, **82**(5),1548–1559.