A human nonlinear cochlear filterbank

Enrique A. Lopez-Poveda[a)]
Centro Regional de Investigación Biomédica, Facultad de Medicina, Universidad de Castilla-La Mancha, Campus Universitario, 02071 Albacete, Spain

Ray Meddis
Centre for the Neural Basis of Hearing, Department of Psychology, University of Essex, Colchester CO4 3SQ, United Kingdom

(Received 18 April 2001; revised 10 July 2001; accepted 10 September 2001)

Some published cochlear filterbanks are nonlinear but are fitted to *animal* basilar membrane (BM) responses. Others, like the gammatone, are based on *human* psychophysical data, but are linear. In this article, a human nonlinear filterbank is constructed by adapting a computational model of animal BM physiology to simulate human BM nonlinearity as measured by psychophysical pulsation-threshold experiments. The approach is based on a dual-resonance nonlinear type of filter whose basic structure was modeled using animal observations. In modeling the pulsation threshold data, the main assumption is that pulsation threshold occurs when the signal and the masker produce comparable excitation, that is the same filter output, at the place of the BM best tuned to the signal frequency. The filter is fitted at a discrete number of best frequencies (BFs) for which psychophysical data are available for a single listener and for an average response of six listeners. The filterbank is then created by linear regression of the resulting parameters to intermediate BFs. The strengths and limitations of the resulting filterbank are discussed. Its suitability for simulating hearing-impaired cochlear responses is also discussed.
[DOI: 10.1121/1.1416197]

PACS numbers: 43.64.Bt [LHC]

I. INTRODUCTION

Computational models of human peripheral frequency selectivity have a long history and a wide range of potential uses. They are a fast way to estimate excitation patterns along the cochlear partition in response to an arbitrary acoustic stimulus and have potential uses in, for example, signal compression algorithms, automatic speech recognition devices and simulations of individual patterns of hearing loss. Historically, the computations have used banks of linear gammatone filters (Patterson *et al.*, 1992) but more recently researchers have attempted to use nonlinear filters (Giguère and Woodland, 1994; Lyon, 1982; Carney, 1993; Goldstein, 1990, 1995; Lopez-Poveda *et al.*, 1998; Meddis *et al.*, 2001). Much of what we know about the nonlinear aspects of peripheral frequency selectivity is derived from animal studies, particularly from direct observation of the response of the cochlear partition to acoustic stimulation (Rhode, 1971; Rhode and Recio, 2000; Robles *et al.*, 1986, 1991; Sellick *et al.*, 1982; Yates *et al.*, 1990; Ruggero *et al.*, 1992). Nevertheless, this information is consistent with recent psychophysical observations showing changes with level of the filter's best frequency (BF) (McFadden and Yama, 1983) and bandwidth (BW) (Glasberg and Moore, 1990; Rosen *et al.*, 1998; for a review see Moore and Glasberg, 1987, and Moore, 1998). In this work we present a computational algorithm that was originally designed to accommodate animal observations and demonstrate how it can be adapted to simulate the hearing of a single human listener and the average hearing of a group of listeners.

The new filterbank aims to simulate the basilar membrane (BM) response using an array of point models (Lopez-Poveda *et al.*, 1998; Meddis *et al.*, 2001). It is similar in many respects to gammatone filterbanks that are in current use. The main difference is that each individual gammatone filter is replaced by a dual resonance nonlinear (DRNL) filter unit (to be described in detail later in this work). It has already been shown that the system can be used to model animal observations. The purpose of the present study is to show that, given suitable parameter changes, it can also be used to simulate human psychophysical data. The parameters of the DRNL units vary with respect to position along the cochlear partition but are fixed with respect to the intensity of the stimulus. Nevertheless, it is an emergent property of the system that the effective BW of each unit changes with signal level. The primary aim of the present study is to explore a new methodology for using human psychophysical data to identify parameters of the DRNL units. A second aim of the study is to identify mathematical functions that allow us to specify the parameters of DRNL units at any arbitrary BF.

Plack and Oxenham (2000) have recently published estimates of BM nonlinearity in normal listeners at BFs between 250 Hz and 8 kHz. These data will be used to help define the parameters of the DRNL units. They used a ''pulsation threshold'' paradigm (Houtgast, 1972), which is based on an auditory illusion whereby an interrupted sound is perceived as being continuous if there is sufficient energy from another sound during the interruptions. Plack and Oxenham presented a stimulus consisting of a signal tone at BF rapidly alternated on three occasions by a low-frequency (LF) tone

a)Electronic mail: ealopez@med-ab.uclm.es

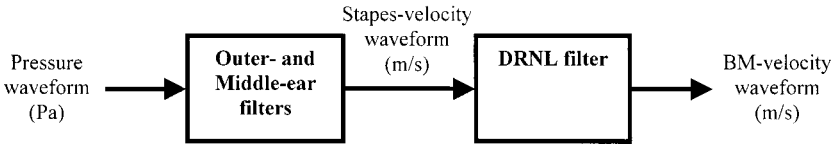

FIG. 1. Model diagram.

(where LF is defined as 0.6×BF). For every level of the BF tone, the level of the LF tone was adjusted to find the ''pulsation threshold,'' that is the level at which the perception of the signal changes from ''pulsed'' to ''continuous.'' They argue that this is the level at which the two tones are generating equal levels of BM activity in the auditory filter tuned to the BF tone (see also Houtgast, 1972). Their method assumes that the response to a signal as a function of intensity will be compressively nonlinear at the BF of a filter but linear in the low-frequency (LF) tail of the filter. This assumption has received support from both animal (e.g., Sellick *et al.*, 1982) and human (e.g., Oxenham and Plack, 1997; Nelson and Schroder, 1999) studies. As a result, a plot of the level of the LF tone at threshold versus the level of the BF tone provides an estimate of the BM response at BF. If the auditory system responds linearly to both tones, the pulsation threshold should increase linearly with increases in the level of the BF tone. In fact, the threshold rises only slowly as a function of the BF tone. This result is consistent with the idea that the system is nonlinear. Specifically, it shows that excitation rises more slowly as a function of the intensity of the BF tone when compared with the response to the LF tone. In addition, their results show that the amount of nonlinearity varies as a function of BF.

The aim of the modeling study is to find parameters that will allow a DRNL unit to simulate these results quantitatively as well as qualitatively. In what follows, we describe the model and the method used to choose appropriate parameters. Then, it is shown how well the model can be made to fit the data. By allowing the parameters to vary as a function of BF, it is shown that the data can be simulated quantitatively at six different BFs. It is shown that the parameters required to fit the data vary approximately linearly as a function of $\log_{10}$(BF). It is proposed that functions of this type could be used as a basis for defining complete DRNL filterbanks to represent the hearing of an individual listener as well as the average response of six listeners.

II. THE MODEL

The overall structure of the model is shown in Fig. 1. It consists of two main stages: (*i*) an outer/middle-ear filter function, which transforms a headphone-delivered sound pressure waveform into a stapes velocity waveform, and (*ii*) a DRNL filter that simulates the BM velocity of vibration in response to stapes velocity.

A. The outer/middle ear stage

A sound pressure waveform produces vibration of the tympanic membrane, which, in turn, induces the vibration of the stapes. Each of these two processes of the peripheral hearing is modeled by means of a linear-phase, 512-point, finite impulse response (FIR) filter. The coefficients of each FIR filter were obtained from empirical frequency responses by applying an inverse fast Fourier transform routine with MATLAB (The Mathworks, Inc., ver. 5.3).

The outer-ear (headphone pressure-to-eardrum pressure) frequency response is taken from Pralong and Carlile [1996, Fig. 1(e)] and is shown in Fig. 2(a). It corresponds to a typical human outer-ear pressure-gain function measured at a point close to the tympanic membrane when the stimulus is delivered through a pair of Sennheiser HD-250 Linear headphones. A response measured with a Sony MDR-V6 headset would have been preferred, as this is the model that Plack and Oxenham employed to collect their pulsation-threshold data. Unfortunately, such response is unavailable. However, both the MDR-V6 and HD-250 are circumaural headphones and have a fairly flat frequency response in the range of

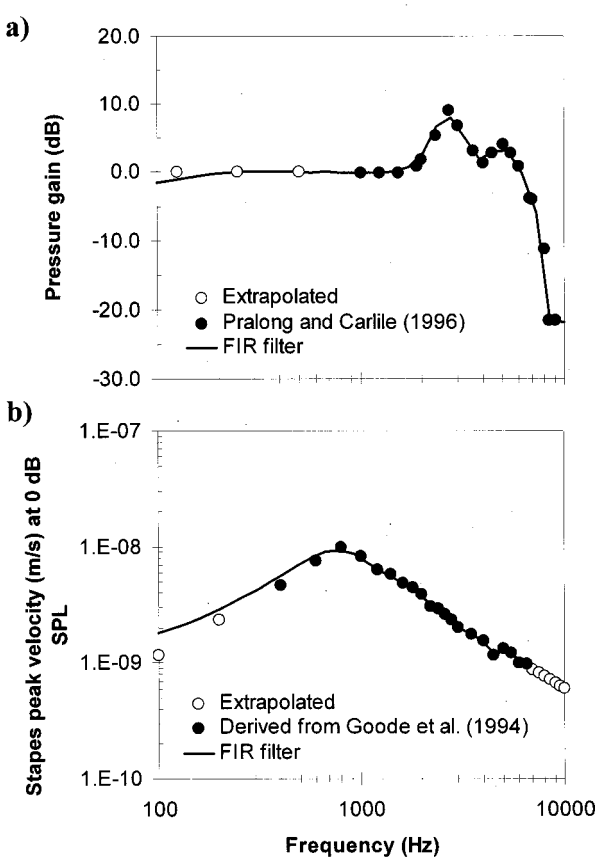

FIG. 2. Stage 1 of the model: human outer/middle ear filter functions. Filled symbols represent experimental data points. Open symbols represent extrapolation points used for evaluating the model over a wider frequency range. Lines represent the actual frequency response of the linear-phase, 512-point, FIR filters used in the model. (a) A typical human headphone-to-eardrum sound pressure gain (from Pralong and Carlile, 1996, Fig. 1E). (b) Stapes peak velocity (m/s) as a function of frequency for a sound pressure input of 0 dB SPL. The velocity values have been derived from the peak-to-peak stapes displacement data measured in cadavers by Goode *et al.* (1994, Fig. 1, set: 104 dB SPL). It is assumed that stapes velocity is linearly related to pressure.

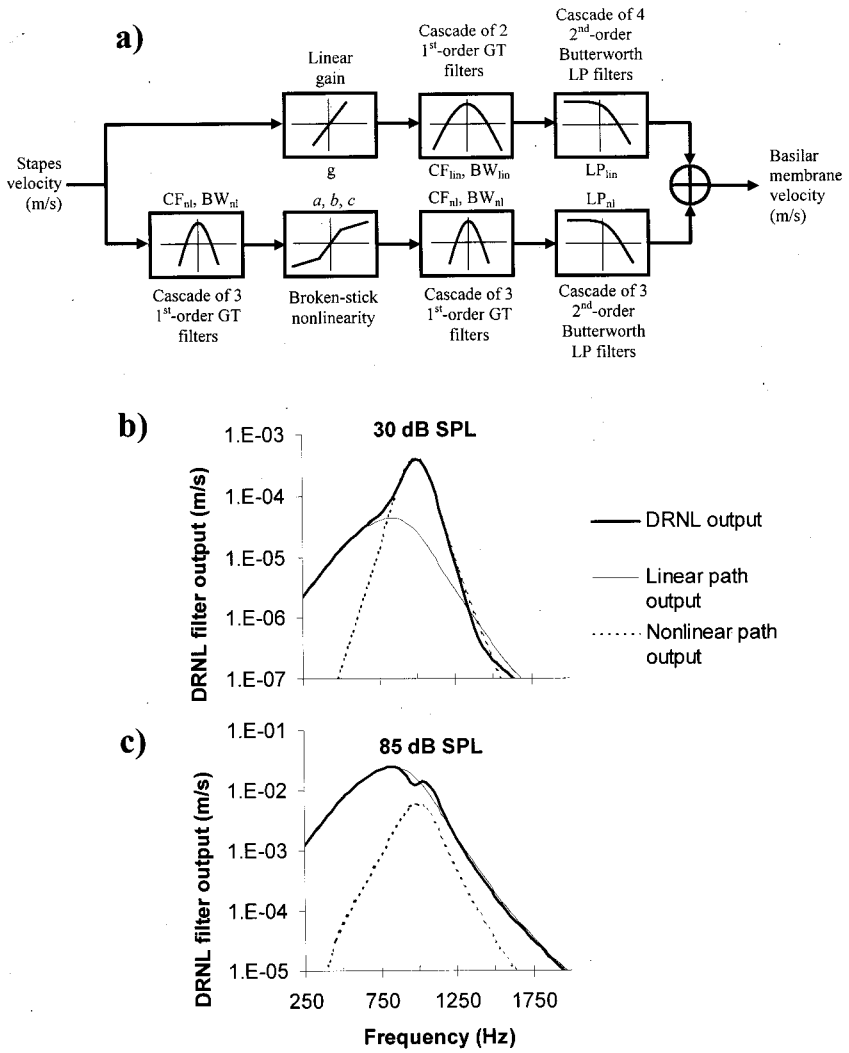

FIG. 3. (a) Stage 2 of the model: The DRNL filter (Meddis *et al.*, 2001). The parameters of each block are shown in the space between the linear (top) and the nonlinear (bottom) paths. The output signal from the DRNL filter is the sum of the signal coming out of each path. (b) Isointensity response of the linear (thin continuous line), and nonlinear (thin dotted line) filter-paths for an input level of 30 dB SPL. At this low intensity, the summed response of the DRNL filter (thick continuous line) is dominated by the response of the nonlinear path. (c) The same as (b) but for an input level of 85 dB SPL. In this case, the summed response is dominated by the response of the linear path. See text for details.

interest (100–8000 Hz). Hence, it is reasonable to assume that the outer-ear response of Fig. 2(a) is a good approximation to that measured with the MDR-V6 headset.

The middle-ear response (stapes velocity as a function of stimulus frequency) is shown in Fig. 2(b) for a stimulus level at the eardrum of 0 dB SPL. The data is derived from stapes displacement measurements in cadavers by Goode *et al.* (1994, Fig. 1) after sound pressure stimulation near the tympanic membrane. Consistent with the observations of Goode *et al.*, peak stapes velocity is assumed to increase linearly with stimulus pressure. The range of empirical data points has been extrapolated from 400–6500 Hz to 100–10 000 Hz [see Fig. 2(b)] in order to be able to evaluate the model over a wider frequency range. The extrapolation is consistent with the measurements of Kringlebotn and Gundersen (1985).

The same outer- and middle-ear filters have been used throughout the modeling work described next.

B. The DRNL filter

Stapes motion transmits energy to the intracochlear fluid, which induces, in turn, motion of the BM. This process is modeled by a DRNL filter (Meddis *et al.*, 2001) which simulates the velocity of vibration of a given site along the BM in response to a given stapes velocity waveform. Its structure and parameters are shown in Fig. 3(a). The input signal follows two independent paths, one linear and one nonlinear. In the linear path, a gain, g, is applied and then the signal is filtered through a cascade of (two or three, see later in this work) first-order gammatone (GT) filters (parameters: CF_{lin} and BW_{lin}) followed by a cascade of four second-order low-pass filters. In the nonlinear path, the input signal is filtered through a cascade of three first-order GT filters (parameters: CF_{nl} and BW_{nl}) followed by a nonlinear gain (see later in this work), followed by another cascade of three GT filters having the same parameters (CF_{nl} and BW_{nl}). During parameter estimation, the CF_{nl} is set to the frequency of the probe signal being studied and is not a free parameter. However, the CF of the linear path (CF_{lin}) is different and typically below CF_{nl} (see later in this work).

The nonlinear gain function is

$$y(t) = \mathrm{sign}[x(t)] \cdot \min[a|x(t)|, b|x(t)|^c], \tag{1}$$

where $x(t)$ and $y(t)$ are the input and the output signals of the nonlinearity, respectively, and a, b, and c are parameters of the model. The details of the time-domain digital implementation of the DRNL filter are given in the Appendix.

TABLE I. DRNL parameters at various signal frequencies used throughout to simulate the characteristics of subject YO's normal hearing. In principle, this set of parameters is ''subject specific.'' The bottom rows show the actual BF and BW_{3dB} of the DRNL filters.

Signal frequency (Hz)	250	500	1000	2000	4000	8000
			DRNL linear path			
No. cascaded GT filters	2	2	2	2	2	2
CF_{lin} (Hz)	235	460	945	1895	3900	7450
BW_{lin} (Hz)	115	150	240	390	620	1550
g	1400	800	520	400	270	250
LP_{lin} cutoff (Hz)	CF_{lin}	CF_{lin}	CF_{lin}	CF_{lin}	CF_{lin}	CF_{lin}
No. cascaded LP filters	4	4	4	4	4	4
			DRNL nonlinear path			
No. cascaded GT filters	3	3	3	3	3	3
CF_{nl} (Hz)	250	500	1000	2000	4000	8000
BW_{nl} (Hz)	84	103	175	300	560	1100
a	2124	4609	4598	9244	30274	76354
b $[(m/s)^{(1-c)}]$	0.45	0.280	0.130	0.078	0.060	0.035
c	0.25	0.25	0.25	0.25	0.25	0.25
LP_{nl} cutoff (Hz)	CF_{nl}	CF_{nl}	CF_{nl}	CF_{nl}	CF_{nl}	CF_{nl}
No. cascaded LP filters	3	3	3	3	3	3
DRNL filter BF (Hz)	260	508	1002	2006	3978	7720
DRNL filter BW_{3dB} (Hz)	47	70	118	210	415	755

The output from the DRNL filter is the sum of the output from the linear and the nonlinear paths. For consistency with its original description by Meddis *et al.*, it is assumed to represent BM velocity.

Although the DRNL unit is simple in construction, the effects of a change in signal level on its properties are not immediately apparent. To help visualize these, two extra figures are supplied. Figure 3(b) shows separately the filter output as a function of signal frequency for the linear and the nonlinear path for a 30 dB SPL input signal. These functions are generated using the parameters to be used later for the 1-kHz site (subject YO, Table I). The combined output of the DRNL unit is shown as the thick continuous line and is simply the sum of the two filter functions. Note that the nonlinear function dominates the output and is the main determinant of the shape of the summed output function. Figure 3(c) shows the same set of functions but this time for an 85 dB SPL signal. The linear filter function has grown considerably while the nonlinear filter function has grown very little. This is because the nonlinearity is compressive. As a consequence, the output from the linear function dominates the aggregate output. The thick line representing the summed output of the DRNL unit is now a much wider filter function than that shown in Fig. 3(b). In addition, we can see that the BF of the DRNL unit has shifted to a lower frequency.

At very low signal levels, the DRNL unit operates linearly. This is because the nonlinear path is linear for low signal inputs [see Eq. (1)]. At very high signal levels, the DRNL also operates essentially as a linear filter. This is because the linear filter comes to dominate the output. These properties of the DRNL filter are consistent with the data of Plack and Oxenham (2000) which often show a linear response at low signal levels, followed by a compressive nonlinearity and then followed by a return to linearity at 80 dB SPL. A similar return to linearity is also sometimes found in observations of animal BM response, although its significance is disputed (Ruggero *et al.*, 1996, Fig. 1; Johnstone *et al.*, 1986, Fig. 5; Rhode and Cooper, 1996, Fig. 7). Note that the effect of signal level on filter width, BF and nonlinearity are emergent properties of a level-independent set of model parameters.

III. MODELING HUMAN BM NONLINEARITY FOR NORMAL-HEARING SUBJECTS

A. Method

The model was tuned to simulate human BM nonlinearity as estimated by psychoacoustical experiments of pulsation threshold (Plack and Oxenham, 2000). Plack and Oxenham presented subjects with a stimulus consisting of a number of interleaved segments of a signal tone and masker tone. Each individual segment was ramped up and down (2-ms raised-cosine ramps) and the frequency of the masker was 0.6 times the frequency of the signal. They measured signal frequencies of 250, 500, 1000, 2000, 4000, and 8000 Hz. For any given signal level, the task was to measure the masker threshold level at which the subject reports the signal to sound ''pulsating'' as opposed to continuous. Background noise was used to prevent off-frequency listening. Figure 4 symbols show their experimental findings for one subject (YO, filled circles) together with the average response of the six subjects considered in their experiment (open circles).

Here, it is assumed that pulsation threshold occurs when the signal and the masker produce the same amount of excitation at the place of the BM of maximum excitation to the signal frequency. For this reason, the paradigm for evaluating the model was a simplified version of that used by Plack and Oxenham. Instead of using a pulsating stimulus, the signal tone and the masker tone were passed independently through the model. Signal and masker frequencies were the same as those used by Plack and Oxenham. Both tones had a duration

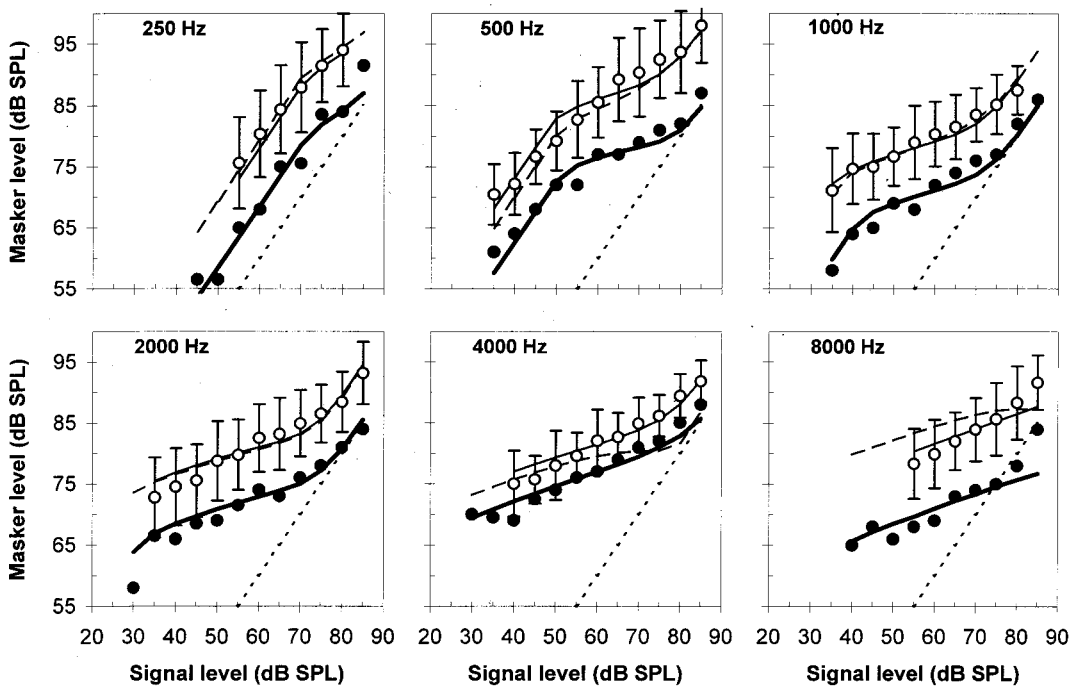

FIG. 4. Comparison between the experimental pulsation threshold data (Plack and Oxenham, 2000, Fig. 2) and the model. Filled symbols represent the data for subject YO. Open symbols represent the average data for the six listeners in the study. Error bars represent an interval of ±1 standard deviation. The thick continuous line represents the model behavior with the optimum set of parameters for subject YO (Table I). The thin continuous line shows the model behavior with the optimum set of parameters for the average data (Table II). The thin dashed line shows the response of the filterbank parameters calculated with the regression lines in Table III, average response. The dotted line illustrates a linear behavior. See text for details.

of 84 ms and were ramped up and down with 2-ms raised-cosine ramps. At each frequency, the signal level was varied within the ranges used by Plack and Oxenham. The sampling frequency (f_s) was 64 000 Hz.

At each BF, the model's peak response to the signal (O_S) and the masker (O_M) was measured during the last half of the stimulus duration in order to avoid any effects that may occur at the onset. For each signal level (L_S), the task was to find out a masker level (L_M) such that the ratio O_S/O_M was equal to (or just exceeded) a value of one.

B. Parameter optimization

Two sets of model parameters were chosen so as to optimize the match between the model results and the two sets of experimental data shown in Fig. 4. However, the model contains a large number of parameters. Not all of them were allowed to vary. Some parameters were fixed by introducing constraints to the search space. These constraints are based on our current understanding of the cochlear response as explained later in this work.

CF_{nl} was fixed at the probe signal frequency. Consistent with Meddis *et al.* (2001), the cut-off frequencies of the low-pass filters in the linear and the nonlinear paths (LP_{lin} and LP_{nl}) were fixed equal to the CF of the GT filters in their respective paths (CF_{lin} and CF_{nl}).

The compression exponent, *c*, was also fixed at 0.25 across BF. This may appear to be surprising in view of the fact that the slopes of the functions in Fig. 4 change with signal frequency. However, the slope of the psychophysical functions in Fig. 4 is only indirectly related to the compression function in Eq. (1). The psychophysical functions, in the model at least, are based on the summed activity of two components, only one of which is nonlinear. The linearity of one component dilutes the nonlinearity of the other in the output of the DRNL unit. We had previously found (Meddis *et al.*, 2001) that a fixed value of the compression exponent, *c*, was consistent with the variation observed in the slopes of the input/output functions in the animal data. By adopting the same strategy on this occasion, it was possible to produce a quantitative match to the data while benefiting from the need to estimate one parameter less.

Parameter *a* of the nonlinearity is responsible for the "sensitivity" at the tip of filter. It, therefore, determines the BM response at the absolute hearing threshold (AHT). There is evidence that AHT occurs at a BM velocity in the region of 5×10^{-5} m/s (Ruggero *et al.*, 1997). For this reason, *a* was fixed so that the DRNL filter output peak velocity is 5×10^{-5} m/s at the subject's hearing threshold.

Parameter CF_{lin} determines the BF of the DRNL filter at high levels, when the linear filter path dominates the DRNL output (see Fig. 3). There is physiological (e.g., Rhode and Recio, 2000) and psychophysical (McFadden and Yama, 1983) evidence of a shift in BF as a function of level. For this reason, an important initial constraint was set on CF_{lin} by requiring it to be lower than CF_{nl}. The shift found by McFadden and Yama from 65 to 95 dB SPL gives a ratio BF_{65}/BF_{95} of approximately 1.1 to 1.4. This does not imply directly that the ratio CF_{nl}/CF_{lin} must be set within that range in the model. Because the DRNL filter is the sum of two components, each of which consists of a number of cascaded gammatone filters, followed by low-pass filters, its BF is not equal to either CF_{lin} or CF_{nl}. However, by trial and error, we

TABLE II. DRNL parameters at various signal frequencies used throughout to simulate the average pulsation-threshold data. The bottom rows show the actual BF and BW_{3dB} of the DRNL filters. Note that, in this case, the number of cascaded GT filters in the linear path is 3, whereas it was 2 for subject YO (Table I; see text for details).

Signal frequency (Hz)	250	500	1000	2000	4000	8000
			DRNL linear path			
No. cascaded GT filters	3	3	3	3	3	3
CF_{lin} (Hz)	244	480	965	1925	3900	7750
BW_{lin} (Hz)	100	130	240	400	660	1450
g	1150	850	520	410	320	220
LP_{lin} cutoff (Hz)	CF_{lin}	CF_{lin}	CF_{lin}	CF_{lin}	CF_{lin}	CF_{lin}
No. cascaded LP filters	4	4	4	4	4	4
			DRNL nonlinear path			
No. cascaded GT filters	3	3	3	3	3	3
CF_{nl} (Hz)	250	500	1000	2000	4000	8000
BW_{nl} (Hz)	84	103	175	300	560	1100
a	2194	5184	7558	9627	22288	43584
b $[(m/s)^{(1-c)}]$	0.450	0.280	0.130	0.078	0.045	0.030
c	0.25	0.25	0.25	0.25	0.25	0.25
LP_{nl} cutoff (Hz)	CF_{nl}	CF_{nl}	CF_{nl}	CF_{nl}	CF_{nl}	CF_{nl}
No. cascaded LP filters	3	3	3	3	3	3
DRNL filter BF (Hz)	258	508	998	2006	3978	7720
DRNL filter BW_{3dB} (Hz)	50	68	118	210	415	755

have been able to conclude that the observed BF_{65}/BF_{95} ratio can be modeled by setting the ratio $CF_{nl}/CF_{lin} \geqslant 1.025$ approximately (see Tables I and II).

Another important constraint was set by requiring the DRNL filter functions near threshold to have 3-dB-down bandwidths (BW_{3dB}) that are consistent with the human psychophysical observations specified in the formula given by Glasberg and Moore (1990) (see later in this work). This criterion constrains BW_{nl} as the nonlinear filter path dominates the DRNL filter function at threshold. It does not mean, however, that BW_{nl} is set equal to the value given by Glasberg and Moore. Since the DRNL filter function at threshold is the result of a number of cascaded GT filters, the effective DRNL filter BW is somewhat lower than BW_{nl}, the BW of the individual GT filters. As a result, BW_{nl} must be greater than the human psychophysical bandwidths (see Tables I and II).

A weaker constraint was introduced on BW_{lin} so that the effective DRNL filter BW at high levels (when the linear filter dominates the DRNL filter output) is larger than at threshold, in agreement with psychophysical human filter functions. It was decided not to set any stronger numerical constraints on BW_{lin} as it deeply influences DRNL filter output to the masker and hence the psychophysical data functions in Fig. 4. These functions can be characterized (from left to right) as an initial linear, a compressed, and a final linear section. When concentrating on these separate features, the modeler was able to use the following rules when varying the remaining parameters (b, g, and BW_{lin}) to optimize the fit. The gain b was varied to set the height of the compressed section in the functions of Fig. 4. Finally, once parameter a has been fixed (see earlier in this work), the gain g together with BW_{lin} were set so that the height of the initial linearity could be reproduced and so that O_M (the response to the masker) is approximately equal to O_S (the response to the signal) at 85 dB SPL, in agreement with the psychophysical data of Fig. 4.

With the above mentioned constraints, the variable parameters were varied manually to optimize the fit by minimizing the Euclidean distance between the psychophysical data and the model results. Parameters to fit the data of subject YO (Table I) were found in the first place. These were then adjusted to find optimum parameters to fit the average pulsation thresholds in the Plack and Oxenham study (Table II). It is important to notice that across BFs the experimental average pulsation thresholds are 8 to 9 dB higher than those of subject YO (in Fig. 4, compare open and closed symbols). Based on our main assumption that pulsation threshold occurs when the masker and the signal produce equal levels of BM excitation (see also Houtgast, 1972, and Plack and Oxenham, 2000), this result implies that the auditory filters corresponding to the average response are necessarily steeper in their low-frequency tail (at $0.6\times$BF) than those of subject YO. In the model, this effect can be achieved in two ways: either by reducing BW_{lin} or by increasing the number of cascaded gammatone filters in the linear path. The necessary adjustments to BW_{lin} would require its value to be lower than BW_{nl}, hence violating one of the constraints set previously. It was, therefore, decided to allow the number of cascaded gammatone filters in the linear path to vary from two (for subject YO) to three (for the average response) in order to account for the ''vertical'' variability in pulsation thresholds. The issue is further discussed in Sec. IV.

C. Creating a human filterbank

A computational model like the one presented here is expected to be used as part of larger models of the auditory periphery. This sort of application usually requires a filterbank rather than a discrete number of filters.

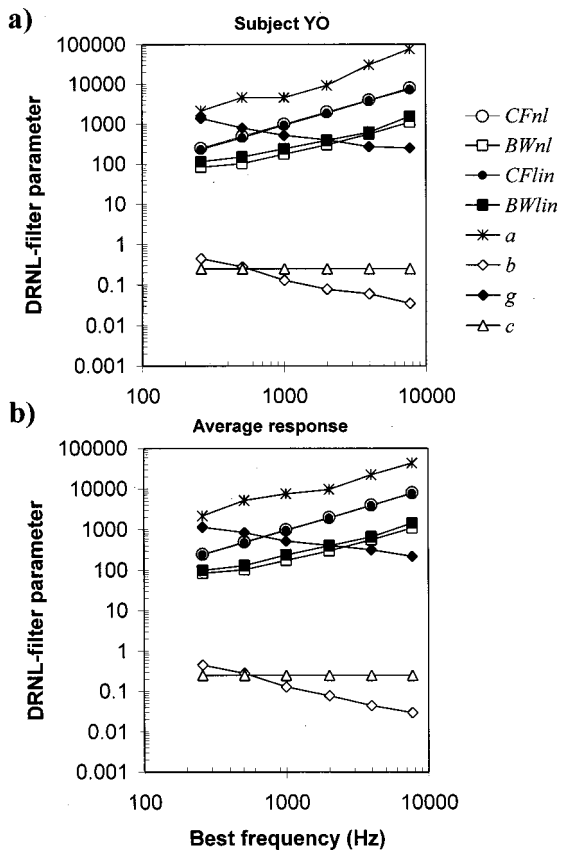

FIG. 5. Optimum model parameters as a function of BF. (a) Parameters for subject YO, Table I. (b) Parameters to fit the average data of six listeners, Table II. Note that the relationship between $\log_{10}$(parameter) and $\log_{10}$(BF) can be approximated as linear in both cases. Note also that the most variable parameter from a linear regression is *a*, which is adjusted to fit the subjects' threshold.

Figure 5 shows the variation of the two sets of model parameters as a function of BF. In both cases, it shows that the logarithm of each parameter as a function of the logarithm of BF can be approximated by a linear relationship as follows:

$$\log_{10}(\text{parameter}) = p_0 + m \log_{10}(\text{BF}), \tag{2}$$

where p_0 and m are the regression coefficients. This result suggests a possible way of creating a filterbank by linear regression of each parameter at intermediate BFs. The regression coefficients, p_0, and m, have been estimated for each parameter using least squares regression and are reported in Table III. With these regression equations, it is possible to specify the values of the free parameters for a DRNL representation of YO's or the average-listener's hearing at BFs between 250 Hz and 8 kHz. In what follows, only the results of the average-response filterbank are shown as those for subject YO's filterbank are similar.

D. Results

1. Input/output curves

Figure 4 compares the experimental data (Plack and Oxenham, 2000, Fig. 2) and the model results for both data sets (subject YO and average responses). For the most part, the fit is quantitatively very good in both cases. At BF=250 Hz, the model masking function is mainly linear, while at 1000 Hz, the model shows a three stage function (linear, compressed, then linear again). At 8000 Hz, the model replicates the almost wholly compressed nature of the function except for the return to linearity at high signal levels. The only significant failures occur at BF=2000 Hz, where the model fails to simulate the linearity at 30 dB SPL for subject YO, and at BF=8000 Hz, where the model fails to simulate the linearity at 80–85 dB SPL. The failure at 2000 Hz is linked to one of the parameter constraints that fixes the tip sensitivity parameter, *a*, to agree with the YO's or the average hearing threshold. It may also be a consequence of the fact that the outer-ear or the cadaver-based, middle-ear function do not correspond to subject YO's. The failure at 8000 Hz is linked to the deep notch in the outer-ear response [see Fig. 2(a)]. Pralong and Carlile (1996) pointed out that the greatest variability in outer-ear responses from trial to trial occurs precisely at around 8000 Hz. Their result is consistent with the findings of Kulkarni and Colburn (2000) and Møller *et al.* (1995). This suggests that at this particular frequency the outer-ear function shown in Fig. 2(a) is inappropriate to model the response of the subjects with the headphones used by Plack and Oxenham. Although not shown here, we did find an almost perfect match to the 8-kHz data assuming a flat outer-ear response at this frequency.

TABLE III. Regression-line coefficients p_0 and m for creating the filterbank assuming a relationship of the form: $\log_{10}(\text{parameter}) = p_0 + m \log_{10}(\text{BF})$, with BF expressed in Hz.

	Subject YO		Average response	
DRNL filter parameter	p_0	m	p_0	m
CF_{lin}	−0.102 05	1.023 85	−0.067 62	1.016 79
BW_{lin}	0.184 58	0.744 70	0.037 28	0.785 63
CF_{nl}	−0.059 12	1.018 32	−0.052 52	1.016 50
BW_{nl}	−0.037 39	0.775 79	−0.031 93	0.774 26
a	0.788 80	1.017 92	1.402 98	0.819 16
b	1.439 70	−0.751 38	1.619 12	−0.818 67
c	−0.602 06	0.000 00	−0.602 06	0.000 00
g	4.305 06	−0.509 42	4.204 05	−0.479 09
$LP_{lin\ cutoff}$	−0.102 05	1.023 85	−0.067 62	1.016 79
$LP_{nl\ cutoff}$	−0.059 12	1.018 32	−0.052 52	1.016 50

The dashed thin line in Fig. 4 shows the fit to the average data using the filterbank parameters calculated by linear regression. The overall behavior is maintained. The filterbank response is within one standard deviation of the average response for the most cases. The fit, however, gets worse, as would be expected from using a set of parameters that is not optimum for the specific data set.

2. Thresholds and filter bandwidths

Figure 6(a) shows a good match between the experimental and the model thresholds (assumed to occur at an output velocity of 5×10^{-5} m/s) both for subject YO and the average data. This justifies the values selected for parameter *a*.

The figure also shows the thresholds estimated with the average-filterbank parameters. The larger deviations between experimental and filterbank values are less than 5 dB. They

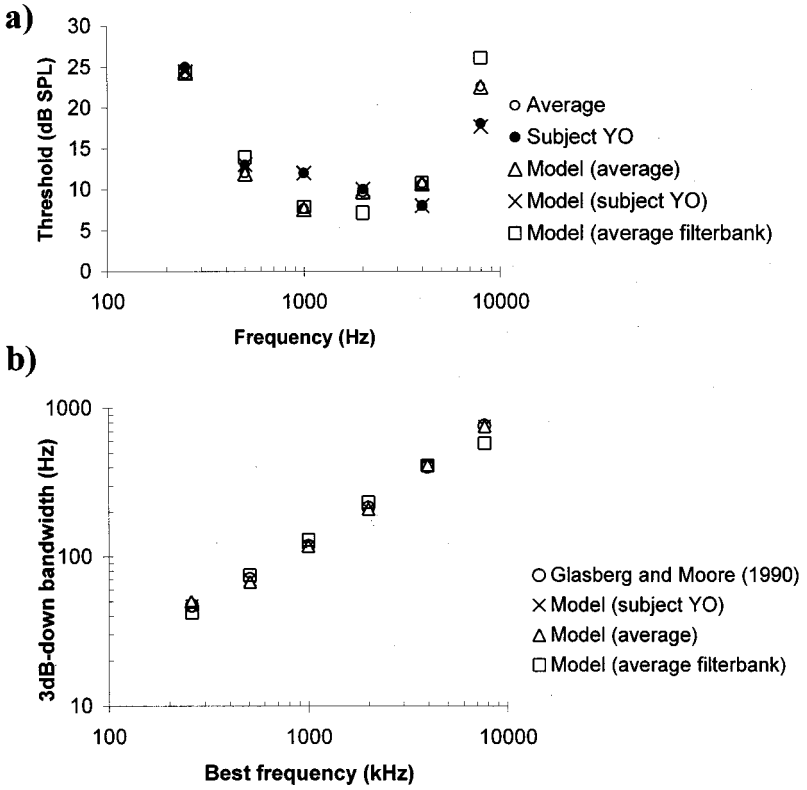

FIG. 6. (a) Experimental and model thresholds. Absolute threshold is assumed to occur at a BM velocity 5×10^{-5} m/s. Subject YO's thresholds (filled circles) match the corresponding model thresholds (crosses) estimated with the parameters in Table I. Likewise, the average threshold of the six listeners (open circles) match the model response (open triangles) calculated with the parameters in Table II. The thresholds estimated using the average-filterbank interpolated parameters (squares) are a reasonable match to the experimental average thresholds. The larger discrepancies at 0.5, 2, and 4 kHz are less than 5 dB and are the result of the deviation between the optimum (Table II) and the regression values for parameter *a*. (b) A 3-dB-down bandwidth of the modeled filters compared with the human data as predicted by the Glasberg and Moore (1990) formula. Note the good match between them. The largest discrepancy occurs for the filterbank at 8 kHz. See text for details.

occur at 0.5, 2.0, and 8.0 kHz and correspond to the largest deviations of the optimum values of parameter *a* over its regression line [Fig. 5(b)].

Figure 6(b) shows the approximate BW_{3dB} of the modeled filters compared with the human as calculated by the Glasberg and Moore (1990) formula:

$$BW_{3dB}\cong 0.89\times ERB\cong 0.89\times[24.7\times(4.37BF+1)], \quad (3)$$

where BF is expressed in kHz. The modeled BW_{3dB} were calculated directly from the isoresponse curves and not by applying the *critical band* paradigm used by Glasberg and Moore. For this purpose, the isoresponse curves were evaluated at a number of stimulus frequencies around the filter's BF enough to determine the 3-dB-down cutoff frequencies with reasonable accuracy (±2 Hz). It can be seen that the DRNL filter BWs are a good match to the psychophysical data. The average filterbank BW_{3dB} match the human bandwidths except for 8 kHz, where the filterbank BW_{3dB} is around 150 Hz lower.

3. Isointensity response

Figure 7(a) shows the isointensity response of the model with the set of parameters given in Table II, optimized to fit the average pulsation threshold data (the response with the subject YO parameters of Table I is similar). The curves correspond to stimulus levels from 30 to 70 dB SPL in steps of 10 dB and have been normalized to give a gain of 0 dB at their peak value. The psychophysical filtershapes derived by Baker *et al.* (1998) are shown in Fig. 7(b) for comparison. Finally, Fig. 7(c) shows the filter shapes of the average filterbank (Table III).

The model curves show distinct notches below BF that are not reported in the psychophysical literature. In the model, the notches are the result of phase cancellation between the outputs of the linear and nonlinear paths. Similar notches have been consistently observed in physiological measurements of the cochlear partition in chinchilla (e.g., Rhode and Recio, 2000) and may be the explanation for ''Nelson's notches'' (Kiang *et al.*, 1986, Fig. 5) occasionally reported in auditory nerve rate/intensity functions. If human listeners monitor a number of adjacent filters simultaneously and pay attention to the filter with the best signal/noise ratio, it is unlikely that these notches will be observed psychophysically as they offer a relatively poor signal to noise ratio.

The model filter function gets wider as the level of the stimulus increases. This is an important property of the DRNL filter consistent with the physiological (e.g., Rhode and Recio, 2000) and psychophysical behavior (e.g., Baker *et al.*, 1998; Rosen *et al.*, 1998; Glasberg and Moore, 2000).

Although not seen in Fig. 7, from 65 to 95 dB SPL there is a shift in the model BF towards lower BFs such that the ratio BF_{65}/BF_{95} ranges from about 1.10 (at BF=4 kHz) to 1.24 (at BF=8 kHz). The shifts are even larger (1.15 to 1.52) with YO's optimum parameters. Such shifts are consistent with the physiology (e.g., Rhode and Recio, 2000) and have also been observed psychophysically when adequate nonsimultaneous masking paradigms are employed (e.g., McFadden and Yama, 1983).

IV. DISCUSSION

The modeling exercise had two main aims. The first aim was to show that DRNL filters could be used to simulate psychophysical estimates of filter width and compression. The second aim was to show that the DRNL parameter changes are orderly with respect to BF. The close qualitative

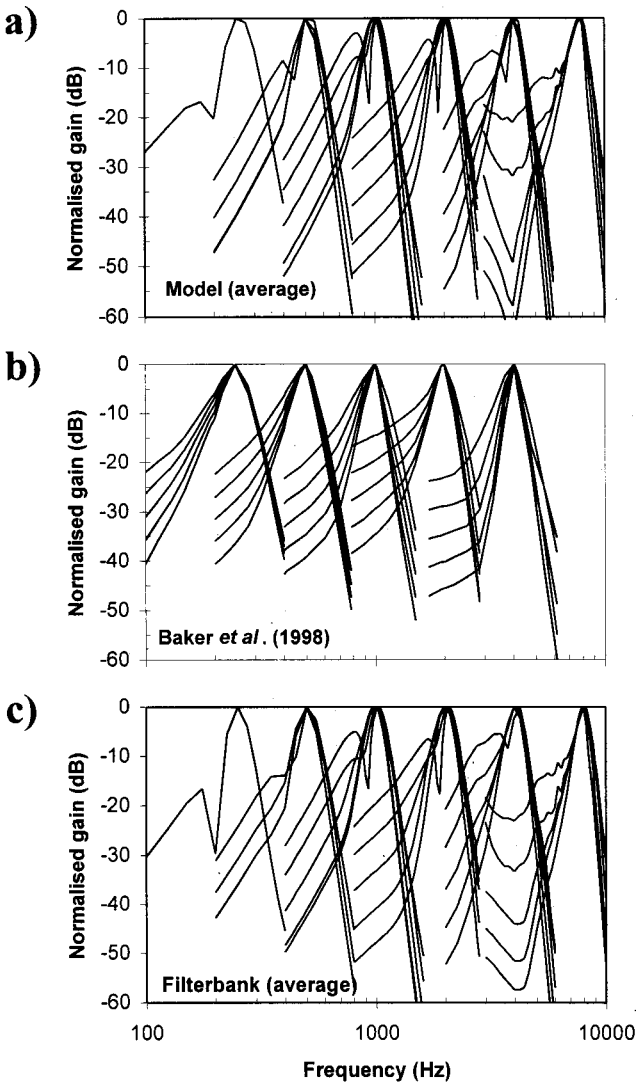

FIG. 7. (a) Isointensity curves modeled with the optimum parameters to match the average data (Table II). The response has been normalized to give a gain of 0 dB at its peak value. (b) Psychophysical filter shapes by Baker *et al.* (1998, Fig. 1). In both sets of curves, the stimulus level ranged from 30 to 70 dB SPL. (c) Isointensity curves modeled with the average-filterbank parameters (Table III).

and quantitative fits of the modeling results to the psychophysical data in Fig. 4 indicate that the first aim has been met with respect to both a single listener, YO, and the average response of six listeners. Similarly, the good fit of the regression functions in Table III suggests that the relationship between the parameters and BF is orderly enough to allow a filterbank to be generated for any point along the basilar membrane. This appears to be the case at least for the region representing a range of BFs between 250 Hz and 8 kHz.

The fitting process depends critically on three important assumptions. First, it assumes that the DRNL architecture gives an adequate account of the response of the cochlear partition to acoustic stimulation. Arguments in favor of this assumption are presented in Meddis *et al.* (2001). Second, it assumes that the pulsation threshold really does represent a condition in which the excitation of the masker, O_M, is equal to that generated by the probe signal, O_S. Arguments in favor of this assumption are presented by Houtgast (1972) and Plack and Oxenham (2000). However, the data of Plack and Oxenham also show that subjects with similar absolute hearing thresholds have pulsation thresholds that differ by as much as 20 dB (cf. subjects AO and YO in the original study). If the assumption was strictly correct, it would imply that their filters would be as much as 20 dB different at the masker frequency. Such intersubject variability seems rather large. Moore (1998, Fig. 3.19) gives an example of intersubject variability at BF=1 kHz, where the filter shapes of four subjects at 600 Hz (=0.6×BF) differ by no more than 10 dB. Therefore, it is possible that the pulsation threshold occurs at ratios O_S/O_M less than one for some subjects. The topic needs further investigation.

The third assumption is that listeners are able to attend selectively to the output of a single filter. This assumption is made in many psychophysical experiments and has proved to be of pragmatic value. To reduce the risk associated with this assumption, Plack and Oxenham (2000) took steps to minimize off-frequency listening.

The study does have some technical weaknesses that will need to be addressed in future studies. The outer ear function used in the model was based on measurements made on different subjects wearing different headphones from those in the Plack and Oxenham study. Furthermore, the middle-ear function used in the model was based on measurements made in cadavers. Ideally, we should have used measurements of the stapes response to headphone-delivered acoustic stimulation for the same subjects, with the same headphones as used in the psychophysical study. This is, of course, not practical and it is not immediately clear how this function is to be best estimated. Although the subjects' audiograms were available, they could not be used as a substitute for an outer/middle ear function because this would imply that no processes subsequent to the outer/middle ear contributes to the audiogram.

A second weakness of the model is that it does not address the issue how the filter BF changes with signal level. Animal studies show that BF does shift in this way. If this is also the case with human listeners, then it follows that listeners must redirect their attention to a different site when the probe signal level changes. We have made no attempt here to model this effect. In effect, the procedure previously described assumes that adjacent filters have very similar characteristics. This is despite that fact that the study as a whole shows that the characteristics of the filters are changing along the cochlear partition. Future studies will need to use a fine grain filterbank with facilities for taking the output from the filter whose *current* BF corresponds to the probe signal frequency.

The DRNL model is a point model in that there are no connections between the individual filter units. As a consequence, distortion products generated in one DRNL unit do not propagate to other units. This is unlikely to give rise to fatal difficulties when modeling the response to two tones at well-spaced frequencies, as is the case the Plack and Oxenham data. However, one of the attractions of a nonlinear filterbank is its potential for modeling such phenomena as two-tone suppression and combination tones. In this respect, while the DRNL units of the filterbank proposed earlier are

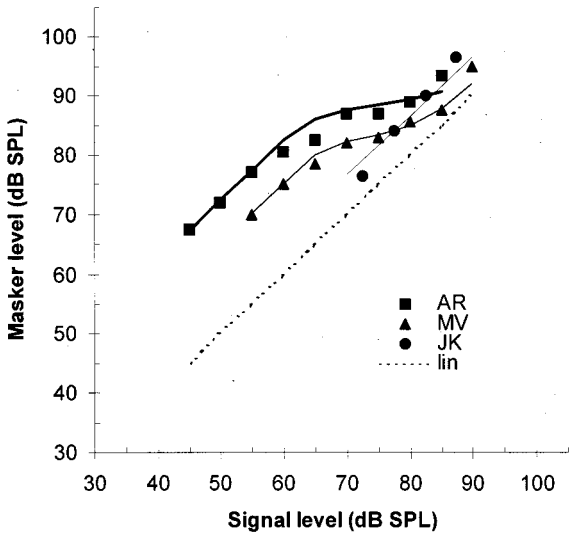

FIG. 8. Model fits (continuous lines) to hearing-impaired responses of the three different subjects (symbols). Experimental data from Oxenham and Plack (1997, Fig. 5). The dotted line illustrates a linear behavior. The frequency of the signal and the masker are 2 and 1 kHz, respectively. Note that subjects MV and AR show mild hearing loss (thresholds of 44 and 37 dB SPL, respectively). Subject JK, however, suffers from severe hearing loss (threshold=67 dB SPL) and hence shows no compression. The parameters in Table I (2 kHz) where used, except for the gains (*a*, *b*, and *g*), which were reduced to model impaired hearing, and BW_{lin} which, in the case of subject JK, was also reduced to model severe impaired hearing. The parameter modifications with respect to those in Table I are as follows: for subject AR, $a=480$, $b=0.06$, and $g=95$; for subject MV, $a=220$, $b=0.035$, and $g=95$; for subject JK, $a=0$, b=irrelevant given that a equals zero, $g=30$, and $BW_{lin}=320$.

able to generate ''local'' suppression and distortion products (Meddis *et al.*, 2001), a more complete model is needed to simulate effects that are not local to the probe tone. For instance, suppression has been shown to occur for suppressor tones with frequencies well below the probe (Abbas and Sachs, 1976; Duifhuis, 1980). Models alternative to the DRNL filter exist that claim to be able to account for this phenomenon (Goldstein, 1990, 1995). Future work is required to address this issue as far as the DRNL is concerned, as well as to make a detailed comparison concerning the relative merits of the various nonlinear filters that have been published recently.

The modeling study above was careful to use the data of a single subject (at least for the compression data). This is because DRNL units may prove useful for modeling the impaired hearing of individuals. If this proves to be feasible on a routine basis, it could be used to optimize the characteristics of hearing aids before supplying them to patients. Figure 8 shows an example of how the model can be tuned to fit hearing-impaired data for three different subjects (Oxenham and Plack, 1997, Fig. 5). In this case, the signal frequency was 2 kHz and the masker frequency 1 kHz. The model was evaluated using the same paradigm as described earlier. The parameters were identical to those in Table I (2 kHz) for normal-hearing subject YO, except for the gains (*a*, *b*, and *g*), which were reduced to model impaired hearing. For subject JK, with severe hearing loss, it was also necessary to reduce BW_{lin} (see Fig. 8 caption). Further studies will be required to evaluate the potential of this methodology.

ACKNOWLEDGMENTS

The authors acknowledge the advice and support of Chris Plack. The authors also thank two anonymous reviewers for constructive comments on an earlier version of this article. The development of the DRNL filter has progressed over many years and has involved contributions from a number of collaborators. The authors would like to thank, in particular, Michael Hewitt, Trevor Shackleton, and Mike Stone for substantial assistance at various times. Author EALP's work was supported by the Junta de Comunidades de Castilla La Mancha (Consejería de Sanidad, Ref. 2001-01044) and by the Universidad de Castilla–La Mancha.

APPENDIX: DIGITAL IMPLEMENTATION OF THE DRNL FILTER

The DRNL filter was implemented digitally in the time domain by implementing each of its filters and gains as a digital component. The implementation of each building component was done as follows.

1. The gammatone filters

The GT filter has an impulse response of the form (Patterson *et al.*, 1992)

$$h(t)=kt^{(n-1)}\exp(-2\pi Bt)\cos(2\pi f_c t+\varphi) \quad (t\geqslant 0), \tag{A1a}$$

$$h(t)=0 \quad (t<0), \tag{A1b}$$

where n is the order of the filter, B is its bandwidth, f_c is its center frequency, φ is its phase, and k is a gain.

The DRNL filter uses several cascades of first-order ($n=1$) GT filters only (see Fig. 3 and main text). They were implemented digitally as an infinite impulse response filter as follows (M. Stone, Exp. Psychology Lab., Cambridge, UK, personal communication, see also Slaney, 1993):

$$y[i]=a_0\cdot x[i]+a_1\cdot x[i-1]-b_1\cdot y[i-1] - b_2\cdot y[i-2], \tag{A2}$$

where $[i]$ refers to the ith sample of the digital signal, x and y are the input and output signals to/from the filter, respectively, and the coefficients are calculated as follows:

$$a_0=\left|\frac{1+b_1\cos\theta-jb_1\sin\theta+b_2\cos(2\theta)-jb_2\sin(2\theta)}{1+\alpha\cos\theta-j\alpha\sin\theta}\right|, \tag{A3a}$$

$$a_1=\alpha\cdot a_0, \tag{A3b}$$

$$b_1=2\alpha, \tag{A3c}$$

$$b_2=\exp(-2\phi), \tag{A3d}$$

where

$$\theta=2\pi f_c dt, \tag{A3e}$$

$$\phi=2\pi B dt, \tag{A3f}$$

$$\alpha=-\exp(-\phi)\cos\theta, \tag{A3g}$$

and $j=\sqrt{(-1)}$, and dt is the sampling period of the digital signal.

2. The low-pass filters

The DRNL filter implementation includes several cascades of second-order Butterworth lowpass filters (see Fig. 3 and main text). These were implemented digitally as follows:

$$y[i]=C\cdot x[i]+2\cdot C\cdot x[i-1]+C\cdot x[i-2]-D\cdot y[i-1]-E\cdot y[i-2], \tag{A4}$$

where the coefficients are

$$C=\frac{1}{1+\sqrt{2}\cot\theta+\cot^2\theta}, \tag{A5a}$$

$$D=2C(1-\cot^2\theta), \tag{A5b}$$

$$E=C(1-\sqrt{2}\cot\theta+\cot^2\theta), \tag{A5c}$$

and

$$\theta=\pi f_c dt, \tag{A5d}$$

where f_c is the 3-dB-down cut-off frequency of the low-pass filter, and dt is the sampling period of the digital signal.

3. The linear gain

The gain in the linear path of the DRNL filter was implemented digitally in the time domain as follows:

$$y[i]=g\cdot x[i], \tag{A6}$$

where $[i]$ refers to the ith sample of the digital signal, and x and y are the input and output signals to/from the linear gain stage, respectively.

4. The nonlinearity

The time domain digital implementation of the ''broken-stick'' nonlinearity was as follows:

$$y[i]=\text{sign}(x[i])\cdot\min(a|x[i]|,\ b|x[i]|^c), \tag{A7}$$

where $[i]$ refers to the ith sample of the digital signal, x and y are the input and output signals to/from the nonlinearity, and a, b, and c are parameters.

Abbas, P. J., and Sachs, M. B. (**1976**). ''Two-tone suppression in auditory-nerve fibers: Extension of stimulus response relationship,'' J. Acoust. Soc. Am. **59**, 112–122.

Baker, R. J., Rosen, S., and Darling, A. M. (**1998**). ''An efficient characterisation of human auditory filtering across level and frequency that is also physiologically reasonable,'' in *Psychophysical and Physiological Advances in Hearing*, edited by A. R. Palmer, A. Rees, A. Q. Summerfield, and R. Meddis (Whurr, London), pp. 81–88.

Carney, L. H. (**1993**). ''A model for the responses of low-frequency auditory-nerve fibers in cat,'' J. Acoust. Soc. Am. **93**, 401–417.

Duifhuis, H. (**1980**). ''Level effects is psychophysical two-tone suppression,'' J. Acoust. Soc. Am. **67**, 914–927.

Giguère, C., and Woodland, P. C. (**1994**). ''A computational model of the auditory periphery for speech and hearing research. I. Ascending path,'' J. Acoust. Soc. Am. **95**, 331–342.

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

Glasberg, B. R., and Moore, B. C. J. (**2000**). ''Frequency selectivity as a function of level and frequency measured with uniformly exciting notched noise,'' J. Acoust. Soc. Am. **108**, 2318–2328.

Goldstein, J. L. (**1990**). ''Modeling rapid wave form compression on the basilar membrane as multiple-band-pass-nonlinearity filtering,'' Hear. Res. **49**, 39–60.

Goldstein, J. L. (**1995**). ''Relations among compression, suppression, and combination tones in mechanical responses of the basilar membrane: data and MBPNL model,'' Hear. Res. **89**, 52–68.

Goode, R. L., Killion, M., Nakamura, K., and Nishihara, S. (**1994**). ''New knowledge about the function of the human middle ear: development of an improved analog model,'' Am. J. Otol. **15**, 145–154.

Houtgast, T. (**1972**). ''Psychophysical evidence for lateral inhibition in hearing,'' J. Acoust. Soc. Am. **51**, 1885–1894.

Johnstone, B. M., Patuzzi, R., and Yates, G. K. (**1986**). ''Basilar membrane measurements and the travelling wave,'' Hear. Res. **22**, 147–153.

Kiang, N. Y. S., Liberman, M., Charles, S. W. F., and Guinan, J. J. (**1986**). ''Single unit clues to cochlear mechanisms,'' Hear. Res. **22**, 171–182.

Kringlebotn, M., and Gundersen, T. (**1985**). ''Frequency characteristics of the middle ear,'' J. Acoust. Soc. Am. **77**, 159–164.

Kulkarni, A., and Colburn, H. S. (**2000**). ''Variability in the characterization of the headphone transfer function,'' J. Acoust. Soc. Am. **107**, 1071–1074.

Lopez-Poveda, E. A., O'Mard, L. P., and Meddis, R. (**1998**). ''A revised computational inner-hair cell model,'' in *Psychophysical and Physiological Advances in Hearing*, edited by A. R. Palmer, A. Rees, A. Q. Summerfield, and R. Meddis (Whurr, London), pp. 112–119.

Lyon, R. F. (**1982**). ''A computational model of filtering, detection, and compression in the cochlea,'' IEEE Proceedings, pp. 1282–1285.

McFadden, D., and Yama, M. F. (**1983**). ''Upward shifts in the masking pattern with increasing masker intensity,'' J. Acoust. Soc. Am. **74**, 1185–1189.

Meddis, R., O'Mard, L. P., and Lopez-Poveda, E. A. (**2001**). ''A computational algorithm for computing nonlinear auditory frequency selectivity,'' J. Acoust. Soc. Am. **109**, 2852–2861.

Møller, H., Hammershøl, D., Jensen, C. B., and Sørensen, M. F. (**1995**). ''Transfer characteristics of headphones measured on human ears,'' J. Audio Eng. Soc. **43**, 203–217.

Moore, B. C. J. (**1998**). *Cochlear Hearing Loss* (Whurr, London).

Moore, B. C. J., and Glasberg, B. R. (**1987**). ''Formulae describing frequency selectivity in the perception of loudness, pitch and time,'' in *Frequency Selectivity in Hearing*, edited by B. C. J. Moore (Academic, London).

Nelson, D. A., and Schroder, A. C. (**1999**). ''Forward masking recovery and peripheral compression in normal-hearing and cochlear-impaired ears,'' J. Acoust. Soc. Am. **106**, 2176–2177(A).

Oxenham, A. J., and Plack, J. C. (**1997**). ''A behavioral measure of basilar-membrane nonlinearity in listeners with normal and impaired hearing,'' J. Acoust. Soc. Am. **101**, 3666–3675.

Patterson, R, Robinson, K., Holdsworth, J., McKeown, Zhang, C., and Allerhand, M. (**1992**). ''Complex sounds and auditory imaging,'' in *Auditory Physiology and Perception*, edited by Y. Cazals, K. Horner, and L. Demany (Pergamon, Oxford, England), pp. 429–443.

Plack, C. J., and Oxenham, A. J. (**2000**). ''Basilar-membrane nonlinearity estimated by pulsation threshold,'' J. Acoust. Soc. Am. **107**, 501–507.

Pralong, D., and Carlile, S. (**1996**). ''The role of individualized headphone calibration for the generation of high fidelity virtual auditory space,'' J. Acoust. Soc. Am. **100**, 3785–3793.

Rhode, W. S. (**1971**). ''Observations of the vibration of the basilar membrane in squirrel monkey using the Mössbauer technique,'' J. Acoust. Soc. Am. **49**, 1218–1231.

Rhode, W. S., and Cooper, N. P. (**1996**). ''Nonlinear mechanics in the apical turn of the chinchilla cochlea *in vivo*,'' Aud. Neurosci. **3**, 102–121.

Rhode, W. S., and Recio, A. (**2000**). ''Study of the mechanical motions in the basal region of the chinchilla cochlea,'' J. Acoust. Soc. Am. **107**, 3317–3332.

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,'' J. Acoust. Soc. Am. **80**, 1364–1374.

Robles, L., Ruggero, M. A., and Rich, N. C. (**1991**). ''Two-tone distortion in the basilar membrane of the cochlea,'' Nature (London) **349**, 413–414.

Rosen, S., Baker, R. J., and Darling, A. (**1998**). ''Auditory filter nonlinearity at 2KHz in normal hearing listeners,'' J. Acoust. Soc. Am. **103**, 2539–2550.

Ruggero, M. A., Rich, N. C., and Recio, A. (**1996**). ''The effect of intense acoustic stimulation on basilar membrane vibrations,'' Aud. Neurosci. **2**, 329–345.

Ruggero, M. A., Robles, L., and Rich, N. C. (**1992**). ''Two-tone suppression in the basilar membrane of the cochlea: Mechanical basis of auditory-nerve rate suppression,'' J. Neurophysiol. **68**, 1087–1099.

Ruggero, M. A., Rich, N. C., Recio, A. Shyamla Narayan, S., and Robles, L. (**1997**). ''Basilar-membrane responses to tones at the base of the chinchilla cochlea,'' J. Acoust. Soc. Am. **101**, 2151–2163.

Sellick, P. M., Patuzzi, R., and Johnstone, B. M. (**1982**). ''Measurements of basilar membrane motion in the guinea pig using the Mössbauer technique,'' J. Acoust. Soc. Am. **72**, 131–141.

Slaney, M. (**1993**). ''An efficient implementation of the Patterson-Holdsworth auditory filter bank,'' Apple Computer Technical Report #35, Apple Computer, Inc.

Yates, G. K., Winter, I. M., and Robertson, D. (**1990**). ''Basilar membrane nonlinearity determines auditory nerve rate-intensity functions and cochlear dynamic range,'' Hear. Res. **45**, 203–220.