

Headphone Essentials 7:

A parametric EQ primer

© Dale Cotton, 2021. All rights reserved. Version 1.1 (August 2021).

You may re-circulate this document. You may not claim authorship or copyright to it.
(Photographs are advertising copy in the public domain: copyright belongs to the original photographers.)



Note: this document is part of a instructional series. If you would like to start with more foundational information on the nature of sound and related material, see *Basics of Musical Sound*, *Basics of Headphone Sound* and *The Skinny on Headphone Frequency Response Graphs* at [Headphone Essentials](#).

There are three main types of equalization, or EQ, commonly available: 10-band graphic, 31-band graphic and multi-band parametric. This tutorial focuses on software-implemented parametric EQ. Also available on the internet are canned EQ settings for hundreds of headphone models. These are essentially EQ recipes that you plug into an EQ app. These are mostly aimed at changing the frequency response of a given headphone to match the Harman 2018 target curve for over-ear headphones and for a few different targets for in-ears. The Harman over-ear target is neutral/accurate except for a fairly strong (6 decibel) built-in bass boost that importantly tapers off before affecting the mid-range. If you're one of the iconoclastic minority who prefers a different bass quantity or even other departures, we'll cover customization in [Step 6: Tweaking the M50x pre-set](#), below.

You're certainly welcome to do your EQ'ing by ear, and we'll cover that in [Tweaking a frequency response from scratch](#), below. If your headphone sounds fine except for a screechy area in the treble, for instance, you can identify at what frequency that screech you hear peaks using a sine sweep, then apply a negative EQ to that frequency. Simply skip to [Step 3](#), follow along far enough to learn how to use the controls of a parametric equalizer, then take a look at the [tweaking](#) section. That said, many people find one tweak leads to another tweak leads to a bottomless rabbit hole of increasingly unsatisfactory twiddling. Many find the simplicity of using a pre-fab EQ to a known standard to be appealing.

Step 1: locating a pre-set EQ for your headphone model

For my money (actually, it's free) the place to start is Oratory1990's [reddit list](#). Oratory is the web handle for a transducer design engineer (a transducer is the sound-producing device in each headphone ear cup. He uses the ultra-expensive measurement equipment in his workplace during after-hours time to make the measurements needed to produce accurate EQs. He is also a headphone hobbyist and is active on reddit. Oratory uses a software program to generate a prototype EQ for each headphone, then tweaks the result by ear as needed.

A second major resource is Jaakko Pasanen's [AutoEQ project](#). Jaakko uses his superb programming skills to scour the internet for usable measurements, turns those into numeric values, then applies further programming to generate an EQ preset for each headphone model in the vast horde of information. There's no hand tweaking, but that's something you can do for yourself if it proves desirable. Also, if you're up to a bit of learning curve, it's possible to import a given AutoEQ pre-set into EqualizerAPO via the command line. We'll be using the manual input method here.

As an example to work with, we'll grab the PDF file from the Oratory list for the Audio-Technica ATH-M50x, given its long-standing popularity. So go to the Oratory reddit page and download the PDF for that headphone to follow along. The top of the PDF looks like this:

Basics of Headphone Sound

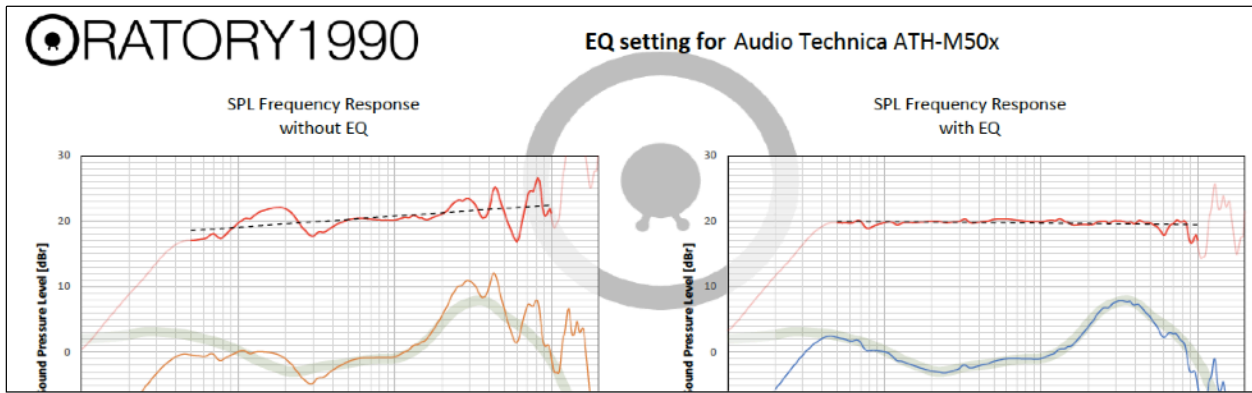


Fig. 1: top portion of Oratory PDF file for the Audio-Technica ATH-M50x

The information we need, however, is near the bottom:

Filter Settings					
	Filter Type	Frequency	Gain	Q-Factor	BW
Band 1	PEAK	40 Hz	-3,1 dB	0,6	2,19
Band 2	LOW_SHELF	105 Hz	5,5 dB	0,71	
Band 3	PEAK	165 Hz	-4,6 dB	0,9	1,53
Band 4	PEAK	315 Hz	5,7 dB	1,4	1,01
Band 5	HIGH_SHELF	2000 Hz	3,0 dB	0,71	
Band 6	PEAK	2800 Hz	-5,5 dB	1,7	0,84
Band 7	PEAK	4340 Hz	-5,6 dB	4,5	0,32
Band 8	PEAK	7050 Hz	-3,6 dB	6,0	0,24
Band 9	PEAK	8000 Hz	-5,0 dB	5,0	0,29
Band 10	HIGH_SHELF	11000 Hz	-11,0 dB	0,71	


Fig. 2: Oratory1990 EQ settings for Audio-Technica M50x headphone

Step 2: downloading the parametric EQ app for your computer or device

Parametric EQ is available on Mac and Android computers, but the majority of headphone owners seem to be running Windows computers and so have access to the king of all software equalizers, [Equalizer APO](#). So we'll use it to illustrate the process. If you're using a different app, the interface may look different but it has to have the same set of controls we'll be using if it's a parametric EQ tool.

Click on the link above, download the zip file, unzip it, then run the exe file to install — but you don't need to open it. Equally important, now download and install the front-end interface called [Peace GUI](#). Peace is the software you will actually interact with for doing EQ, so keep its icon handy. Open the Peace app.

Step 3: learning your way around a parametric EQ app interface

Fig. 3 shows the default Peace interface on the left with the Equalizer Default configuration selected from the list at the bottom left of the main Peace window. I've also clicked on the graph window icon  located centre right to bring up a second window, then resized it to fit in the available space. Finally, at the bottom I have the Oratory1990's filter settings table for the M50x visible in the background to work from:

Basics of Headphone Sound

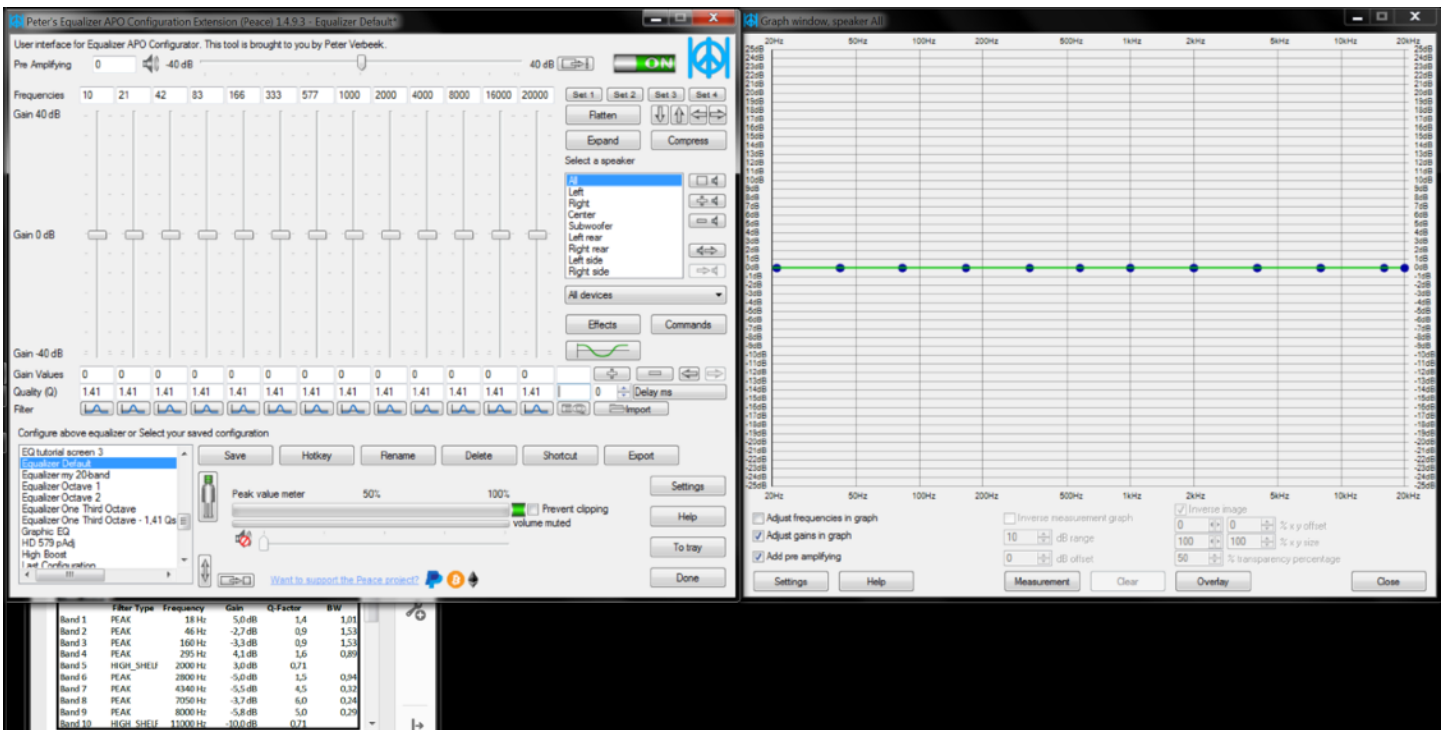


Fig. 3: suggested layout for using Peace GUI for this tutorial

This is the arrangement I recommend working with. If you have two monitors and can put the graph window in its own display, all the better.

The heart of the action takes place here:




Fig. 4: filter properties input area of Peace GUI main screen

In Fig. 4 I've highlighted the first ten filters in blue and indicated the various types of input boxes you'll be working with in black. The gain slider and gain input box are alternate controls for accomplishing the same thing. We'll use the input box pretty much exclusively. The top row

Basics of Headphone Sound

labelled **Frequencies** is pre-supplied with the numbers 10, 21, 42, etc. These may be useful in some contexts, but they are meaningless for the task at hand and will be freely overwritten. The goal is to get the first 10 filters in Peace to match the values shown in the Oratory table in Fig. 2. (The final three filters can be ignored or deleted. As long as a filter has a gain value of 0 it has zero effect on your audio output.)

Note: I vaguely remember initially being only able to input gain in whole and half number values. This is fine, rounding Oratory's more exact decimal values to .0 or .5 will be well within the margin of error that exists due to manufacturer unit variation, not to mention the limits of human hearing. But if you want to work with more exact values, click on the Settings button then snoop around until you find the appropriate adjustment. While you're there, notice there is a huge number of other interface tweaks available.

Step 4: getting a handle on the underlying concepts

But what does all this mean? Here is a portion of the first graph from the M50x PDF file:

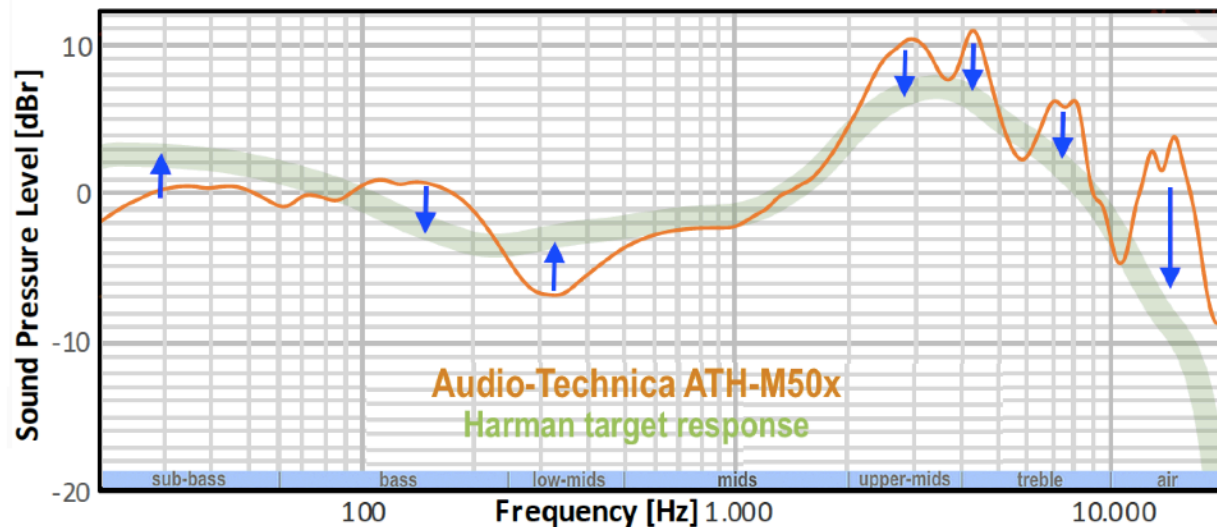


Fig. 5: M50x frequency response compared to target response

Note: the frequency scale in Fig. 5 uses the European convention of a period instead of a comma as the thousands separator.

The graph shows pitch (called *frequency*) from bass to treble left to right, and loudness up and down. The broad green curve is the desired frequency response, or “tuning”, that would remove the inaccuracies of the M50x’s audio playback (again, except that a carefully boosted bass is incorporated). The orange line is how a typical M50x natively performs without help from EQ. Where the orange line rises above the green band the M50x is too loud in those frequency areas. Where the orange line falls below the green band, the M50x is too quiet. The blue arrows show areas where corrections need to be made to get the sound output from the M50x to match the desired loudness output. I also added a legend in pale blue at the bottom showing how all this relates to the familiar bass-mids-treble spectrum of sound that we can easily identify. The green curve is wide in order to show that there is a range of acceptable variance — we don’t need to achieve an exact overlap between the corrected M50x and the centre-line of the target. (Technically, the width of the green curve is a somewhat arbitrary 1 decibel above and 1 decibel below the exact target curve.)

Specifically, we’ll be entering ten corrections, each of which affects an area of the frequency response to its left and right:

Basics of Headphone Sound

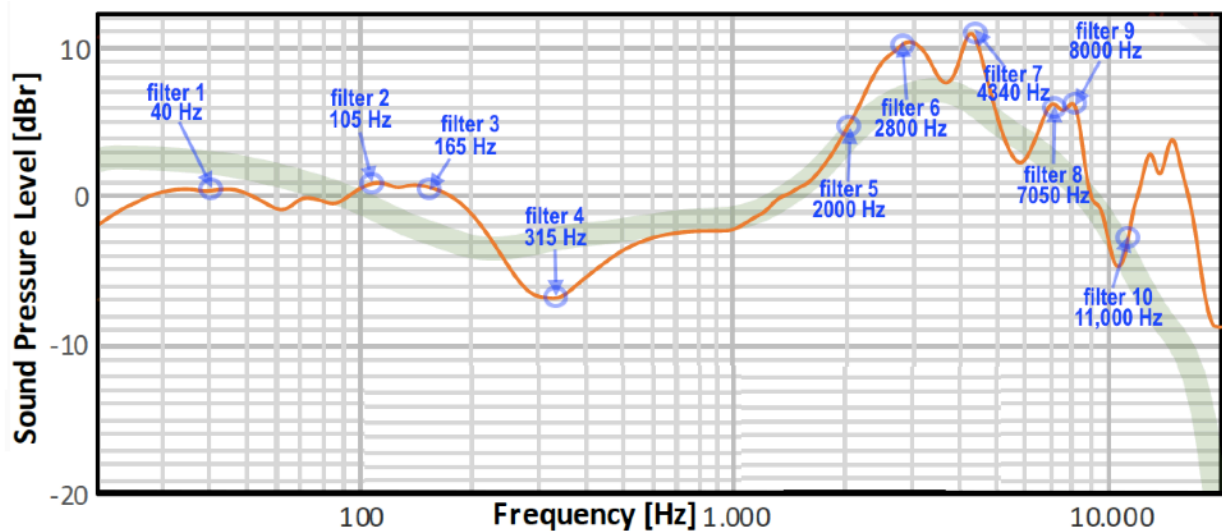


Fig. 6: ten corrections to the M50x frequency response

Vocabulary 1: The use of the word *filter* in this context undoubtedly has a wonderful origin in audio history but is problematic and may be confusing if taken literally. It might be better to think of these as mathematical functions. In real life a filter invariably removes some things and lets others pass. But in EQ land a filter can either add to or subtract from the input sound signal, simply depending on a plus or minus sign of one variable called gain.

And Fig. 7, taken from the Oratory PDF, shows the range and shape of the influence of each filter:

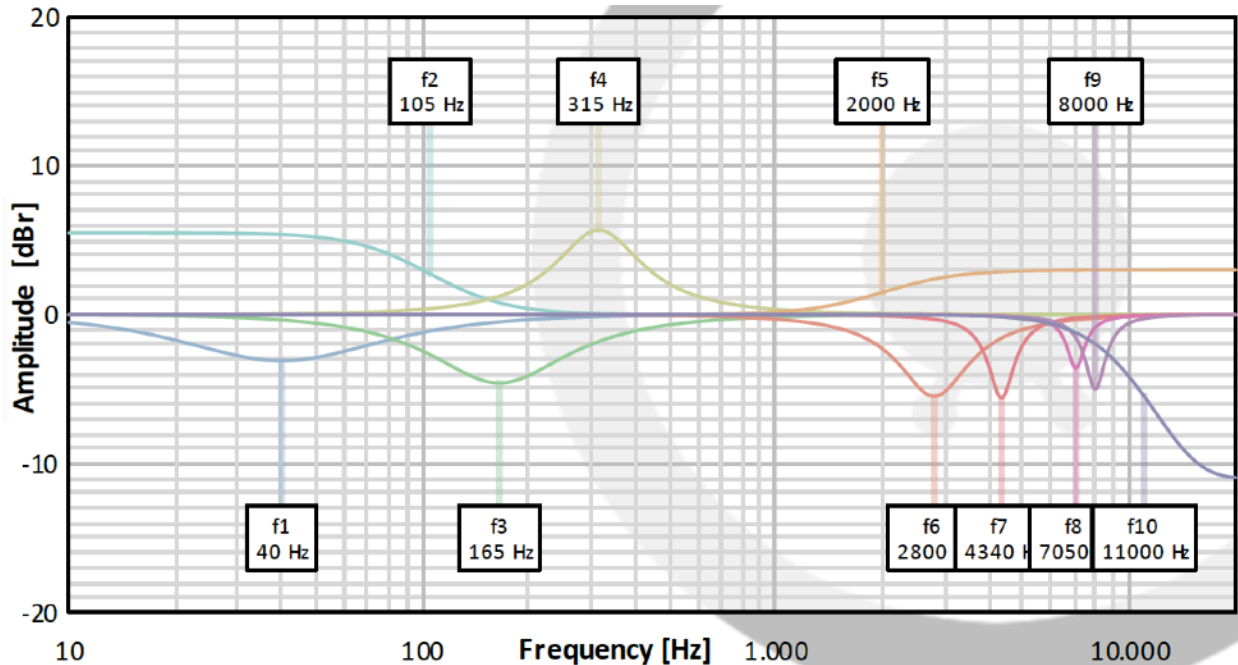


Fig. 7: range and degree of loudness change imposed by each filter

Notice that f2 and f5 both extend like plateaus. F10 does the same thing but most of it is cut off from view to the right. All the other filters are hill/valley shaped, depending on whether they point upward, representing a loudness increase, or point downward, representing a loudness decrease. If we look back at the table in Fig. 2, we see that f2, f5 and f10 all have the designation *high* or *low shelf*. Equally, the other seven filters have the designation *peak*.

Basics of Headphone Sound

Another correspondence we can see between Fig. 2 and Fig. 7 is that the numbers in the **Frequency** column of the Fig. 2 table increase going down row by row, while the same numbers are presented in the Fig. 7 graph going from left to right. A third correspondence between Fig. 2 and Fig. 7 is that the numbers in the **Gain** column in the Fig. 2 table correspond to the maximum heights or minimum depths of the ten filter curves shown in Fig. 7. A final correspondence is less obvious. The larger the **Q-factor** number in the Fig. 2 table, the narrower the corresponding peak or valley is in Fig. 7.

Vocabulary 2. From the above considerations we deduce that *gain* is just audio-speak for signal strength which in turn changes the subjective experience we call loudness. Another vocabulary puzzle is *quality* (Q), which in normal usage just means a distinctive attribute of something. That's helpful ... not. Hiding behind this particular vocabulary snafu is that a change in Q is a change in the range of frequencies to the left and right of the centre frequency that the filter has influence upon. But if we're de-mystify vocabulary, why *frequency*? How often your offspring uses the phrase "I mean" per day is a frequency value but probably not what's intended here. To clear up the audio-speak use of frequency, I refer you to the first explainer in the Headphone Essentials series: [Basics of Musical Sound](#).

Step 5: Creating an example EQ

We're finally ready to get down to work. Click in the upper left Frequencies input box shown in Fig. 8:




Fig. 8: Entering values for first filter
(with the first table row from the Oratory M50x PDF as an inset outlined in green)

Type in 40 to overwrite the default number. Now press the tab key. This takes you down to the **Gain Values** input box directly below, as shown by the red oval. Type in -3.1 and press tab again. Finally, type 0.6 into the **Quality (Q)** input box. You've now specified the three numeric parameters from the first row of Oratory's table shown in Fig. 2. But the first row of the table also includes the **Filter Type** value of *PEAK* at the bottom. Peak is the default and most common value. Plus, it's also the correct value for filter 1, so you don't need to do anything. But the button just below the red oval is there to select the filter type value when needed.

Basics of Headphone Sound

Switch over to the graph window. The left side of the graph should show the same curve as that labeled **f1** in Fig. 7, although it may be vertically stretched or compressed depending on the scale and how you've resized that window.

Explore: Before proceeding further to complete the input of the table in Fig. 2, instead play with the input values you've just entered for filter 1. Move the filter one gain slider located directly above the red oval in Fig. 8. Notice how the curve changes as you do so. The larger the positive number for gain the taller (and therefore louder) the resulting graph peak. Changing the Q value to smaller amounts makes the slope on either side of the peak gentler. Increasing the Q number makes the slope of the peak steeper. Changing the frequency number moves it to a new location on the graph. When all that is clear, restore the filter one values to 40, -3.1 and 0.6.

Audio: At this time, if you're an Equalizer APO newbie, it's a good idea to check that your Equalizer APO plus Peace set-up is actually working. Start a piece of music or other audio playing on the same computer. While you're hearing that audio, grab the gain slider for the eighth band with a frequency of 1000. Move the slider up and down carefully. This is typically an easy frequency to hear so you should be able to hear the changes you're inducing as they happen. If nothing seems to happen, try the 577 slider. If you're still not hearing a change to the playback, there may be a configuration or permissions issue to debug. If you need help, check the [Peace GUI forum](#) to see if a solution is available.

Using the tab-and-type input method, fill in the values from the second row of the table in Fig. 2 into the second column in Peace:

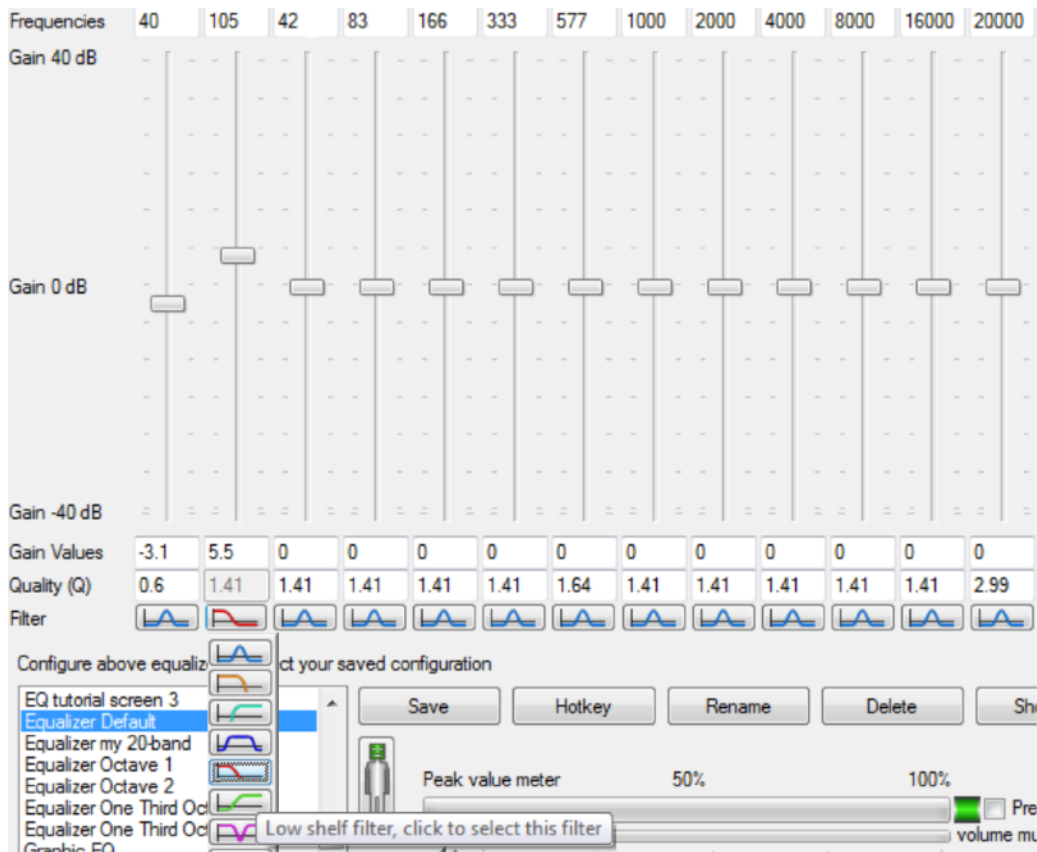


Fig. 9: selecting a filter type for the second filter

But this time, after you enter the gain value, click on the **Filter** selection button, then choose the fifth option, as shown. Notice that in Peace the ability to enter the **Quality** value is now

Basics of Headphone Sound

greyed out. This is because the shape of the low shelf filter is pre-determined by its type. The 0.71 value in the Fig. 2 table is there for other parametric equalizer apps that operate differently. (Also, the last column labeled **BW**, for *bandwidth*, is there for equalizer apps that use bandwidth instead of the inversely related number, Q. So for EqualizerAPO, you just ignore the **BW** numbers.)

At this point you have all the input concepts in place. Go ahead to enter the other eight filters as shown in Fig. 2:

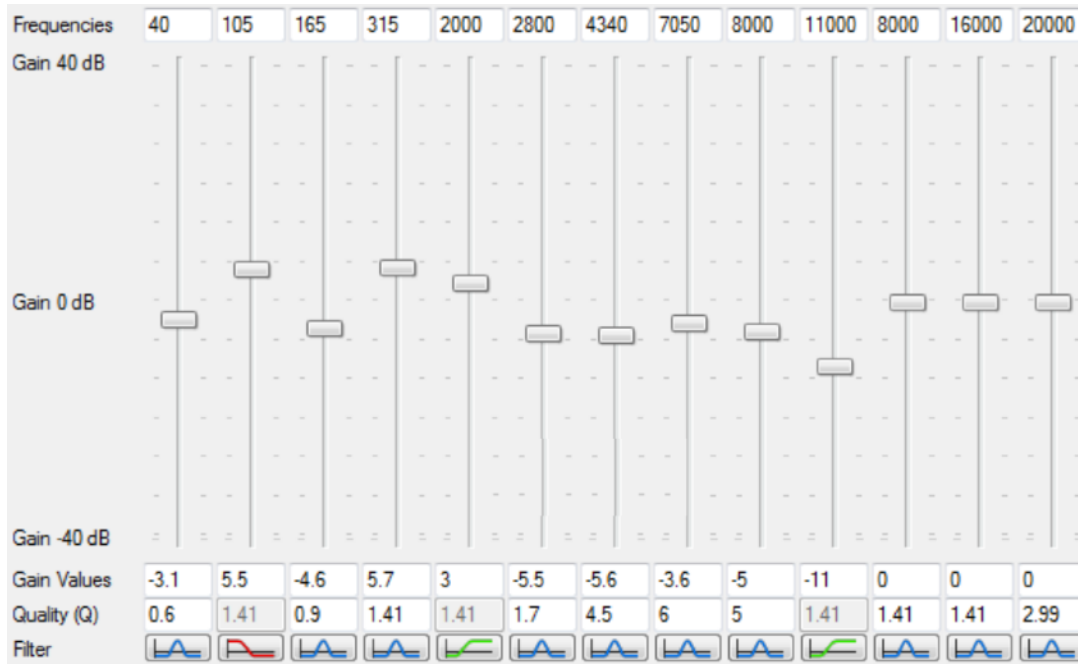


Fig. 10: completed data entry for the M50x EQ

Notice that Peace provides three more filter columns than the ten we actually need. In fact, you can add or delete any number of filter columns using the + and - buttons just to the right of the gain input boxes.

Now check the graph window and compare it against the EQ graph provided in the PDF:

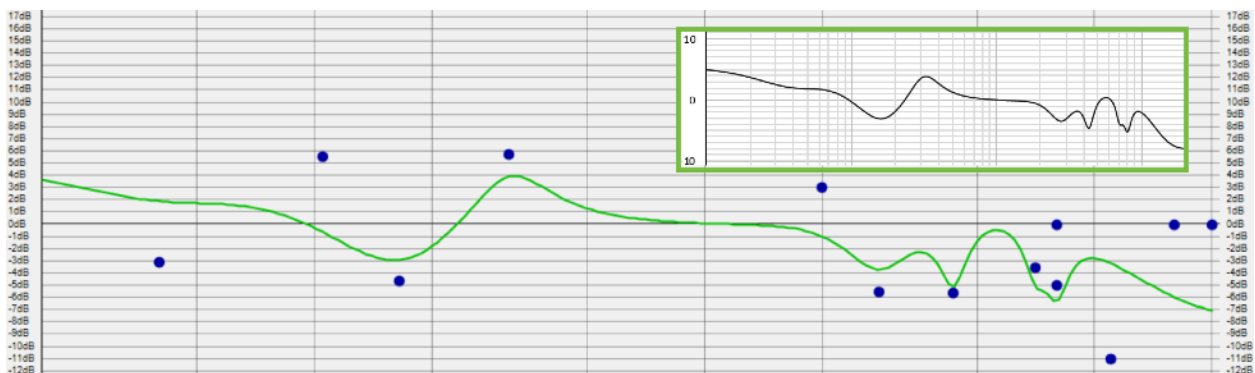


Fig. 11: resulting graph display of entered EQ (with PDF graph as inset upper right for comparison)

But why this particular shape? To answer this we need to take a detour. Oratory1990 maintains another [web presence](#) focused on the measurements themselves rather than EQ. For Fig. 12 I chose a presentation of the same measurements we see in Figs. 5 & 6, but this time showing the green curve stretched out as a flat line. Now the red measurement curve more simply shows its departure from the neutrality target:

Basics of Headphone Sound

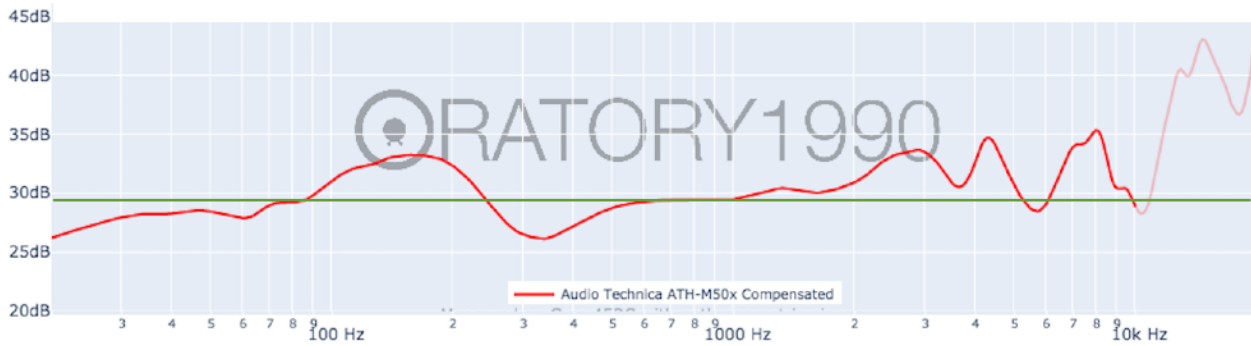


Fig. 12: M50x measurements compensated to the Harman target as a straight line (green)

This is called a compensated frequency response graph, but could equally be called an error curve.

If we simply flip Fig. 12 upside down we get:

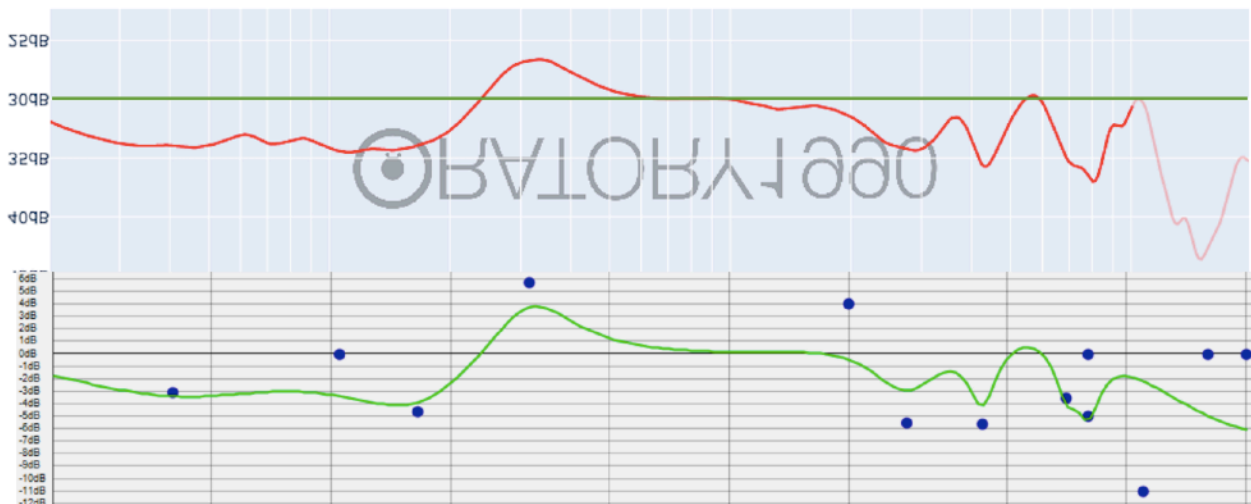


Fig. 13: Fig. 12 flipped upside-down and compared to Fig. 11

As we see in Fig. 13, the curve in Fig. 11 simply supplies the inverse of the M50x's departures from the Harman target. The EQ curve applies this inverse loudness profile to the source sound signal before sending the sound signal to the M50x. Doing this cancels out the M50x's departures from the Harman target.

Note: the match between Fig. 11 and Fig. 13 is remarkable when considered overall. But there are detail differences when examined more closely. These differences are on the same order as the sonic differences from one unit of a given headphone model to another unit. In fact, they are also on the same order as the typical sonic differences between the left side and right side of any one headphone.

The final input step is the most likely one to be a nuisance. At the bottom centre of the PDF is a little box:

Enter that number (notice the minus sign!) in the input box in the upper left corner of the Peace window:

Preamp gain:
-5,0 dB

Basics of Headphone Sound

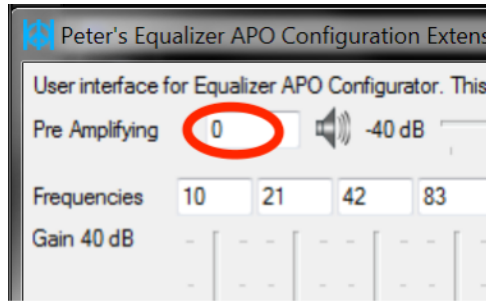


Fig. 14: pre-amp value input

This is a safety feature that prevents the loudness increases in the EQ you've just created from pushing the output of your sound device into dangerous territory. In fact, it just forces you to turn up the volume on your sound output device to compensate. But now the combination of any EQ loudness increase plus your own preferred listening level is less likely to do damage. If you listen at anything like sane loudness levels, the pre-amp correction is very likely unnecessary. Where it comes into play is the occasional accident. For example, if you are using a headphone that requires a higher volume setting, then switch to a different headphone that requires much less power to achieve the same loudness - but you forget to adjust the volume knob before doing so. The result may blast your ears to oblivion, but it is unlikely to blow out your headphone.

The actual final step, is to click the **Save** button then assign this EQ a name, such as *M50x Oratory1990 EQ to Harman target*.

And that may well be all you need to know. Most likely, the headphone you actually want to EQ is not an M50x. So all that remains is to follow the above process from step 1, but this time substitute the filter table for your own headphone in place of the M50x filter table we used above. The following material, steps 6 and 7, may or may not ever be useful to you. You might want to stop reading here until such a time occurs in which you want to modify the sound of your headphone beyond an available pre-set.

Step 6: Tweaking the M50x pre-set

At the bottom right of the Oratory PDF is the following:

Adjust gain of band 2 to preference (bass)
Adjust gain of band 5 to preference (treble)
Adjust gain of band 10 to preference (airiness)

Fig. 15: recommended adjustment options

I've written an entire 20 page [explainer](#) on the Harman and other frequency response targets. The research behind the Harman target showed enormous differences in people's preferences for bass level. The bass level built into the Harman target by default is simply the majority preference (64% of tested listeners). That doesn't make it right for you. So feel free to add or lower the bass level by changing the gain value on the second filter, which is given as 5.5 decibels. Beyond bass, there is remarkably little divergence between people's preferences. However, there are individuals who are especially treble sensitive. And there are also people who regularly listen to treble-emphasized music, such as the high-pitched wail of rock and heavy metal guitar solos. So feel free to play with filters 5 and 10 if you find you're experiencing any high frequency irritation.

Basics of Headphone Sound

It's even possible that you may actually desire a fully accurate sound from your headphone. This pretty much simply means removing the built-in bass boost of the Harman target. In other words, setting filter 2's gain to 0:




Fig. 16: A) Oratory Optimum HiFi compensated M50x measurements inverted, B) Peace graph with filter 2 gain set to 0

However, the M50x is a closed-back portable headphone. It is frequently used in noisy environments, such as a commuter bus or train. In that context, the default bass boost is probably serving to counteract the bass rumble of the vehicle. But for quiet home or office listening, if you want fully accurate playback, Oratory provides an alternate pre-set titled Optimum HiFi for the M50x, which is fully neutral/accurate. This is a fairly recent development (in 2021 as I write this) — hopefully, Oratory will provide Optimum HiFi pre-sets for more headphone models as time passes. That said, there is very little difference between Oratory Optimum HiFi and Harman with 0 bass boost, as we see from Fig. 16.

Alternative approach: Tweaking a frequency response from scratch

Above, I mentioned that another approach to EQ is simply to remove any annoying aspects of the headphone's sound and leave the rest alone. If you own a headphone that does not have a pre-set available, this may well be your only option. For years I lived with a Sennheiser HD 579 for which no useful frequency response measurements were available. I had very limited success tackling its combination of boomy upper bass + lower mid-range plus its shrill upper mid-range on my own. Mostly, it sat un-used. Then, a few years ago, [Crinnacle](#) posted a measurement made with his earlier equipment, which helped a great deal. More recently still, he acquired an industry standard rig and his current HD 579 graph pretty much nails it. But this is an especially problematic case. If, for example, no measurements for the beyerdynamic DT 1990 existed, it would still be easy to tame its lower treble peak (which bothers relatively few people, but those it bothers it bothers a lot.)

To do this, first use the [Online Tone Generator](#) to slowly move through the treble frequency range from 4,000 hertz on up to the limits of your hearing. Obviously, set the loudness fairly low. Cheating, I can show you what you'd hear as you did this, since DT 1990 measurements do exist:

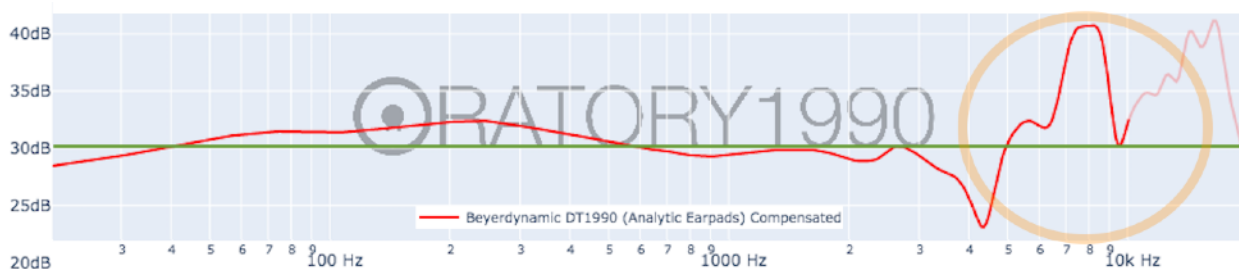


Fig. 17: Oratory Optimum HiFi compensated DT 1990 measurements with treble circled in orange.

Basics of Headphone Sound

From about 6500 to about 9500 hertz the loudness shoots up, briefly plateaus, then plummets. This is a fairly obvious fault to spot using the Online Tone Generator, as is the dip that precedes it. More problematic is the lazy elevation on the left side of the graph from 40 to 500 hertz. This is easy to miss or mis-interpret when listening. In any case, going back to the fiction that we're working entirely by ear, we'd have no trouble identifying the rise from roughly 6500 to 9500 hertz. That's a 3000 hertz span, so the centre point would be at $6500 + 1500$ or 8000 hertz. So, moving to Peace we would start with a new Equalizer Default work area and choose any filter column. But presumably, since the 11th column already has 8000 entered as its frequency, we'd simply use that (small red oval in Fig. 18):

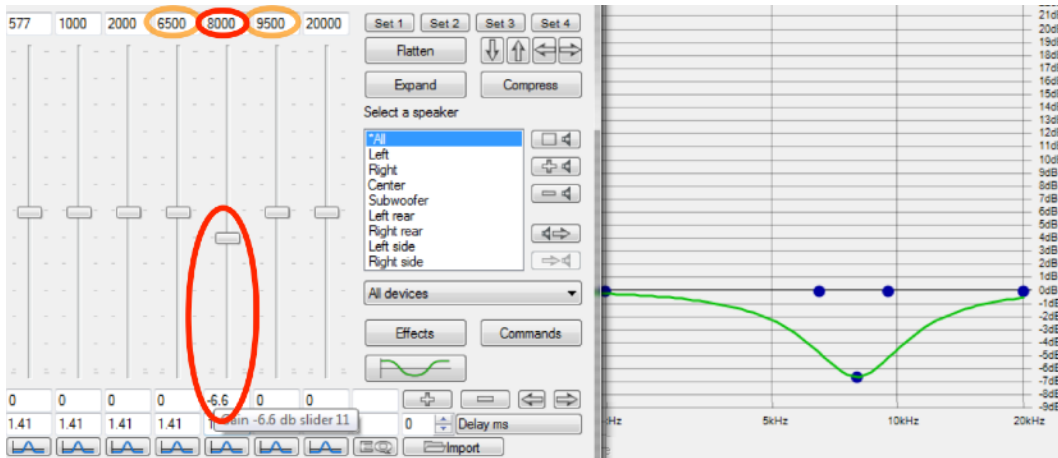


Fig. 18: 8000 hertz correction attempt 1

As we see in Fig. 18, I simply moved the gain slider down by a moderate amount (large red oval) then looked at the result on the graph. The spread of the valley that change created was clearly too broad, since there was still some distance between the green line and the black centre line at 5,000 hertz and 10,000 hertz. Thinking for a moment, I changed the frequency numbers *but nothing else* on the filters immediately to the left and right to 6500 and 9500 respectively (orange ovals in Fig. 15). This caused the dark blue dots in the graph to become visual markers for where I wanted to confine the change to.

So, going back to the 8000 hertz filter, I set the Q value to a variety of higher numbers to find a value that would create a steep enough slope on either side:

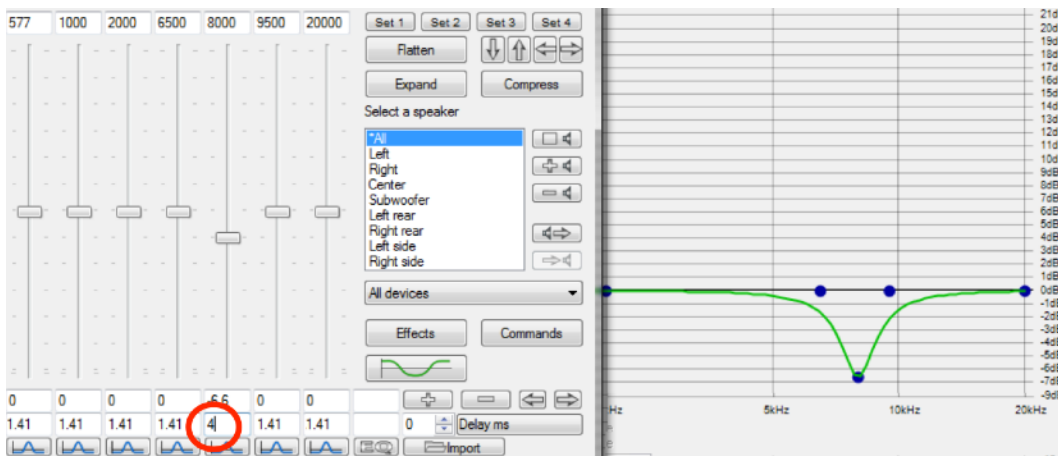


Fig. 19: 8000 hertz correction attempt 2

A Q of 4 seemed to do the job. However, this results in a pointy crevice that doesn't nicely keep the 6500 and 9500 hertz range at the same loudness of loudness. So back to the proverbial drawing board. I switched from using a single central peak filter to using two peak filters on either side of 8000 hertz:

Basics of Headphone Sound

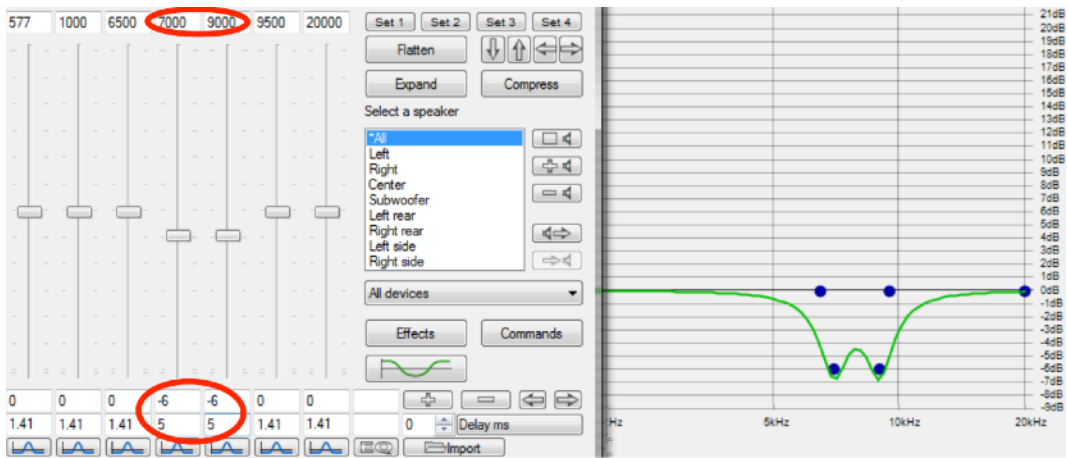


Fig. 20: 8000 hertz correction attempt 3

This is arguably an improvement, but the W-shaped bottom is less than elegant. Nothing for it but to introduce yet another control point:

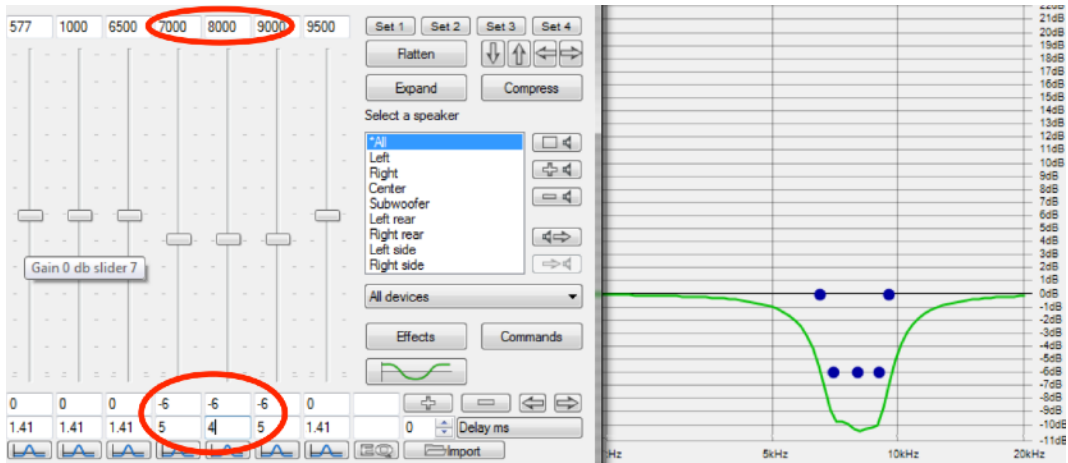


Fig. 21: 8000 hertz correction attempt 4

Better, but now the entire dip has dropped from 6,000 hertz down to 10,000. If I actually want a -10 decibel reduction, great — time to find out. So I switch back to the Online Tone Generator, slowly sweep through the treble frequency range again and decide that much loudness reduction is too extreme. Solution:

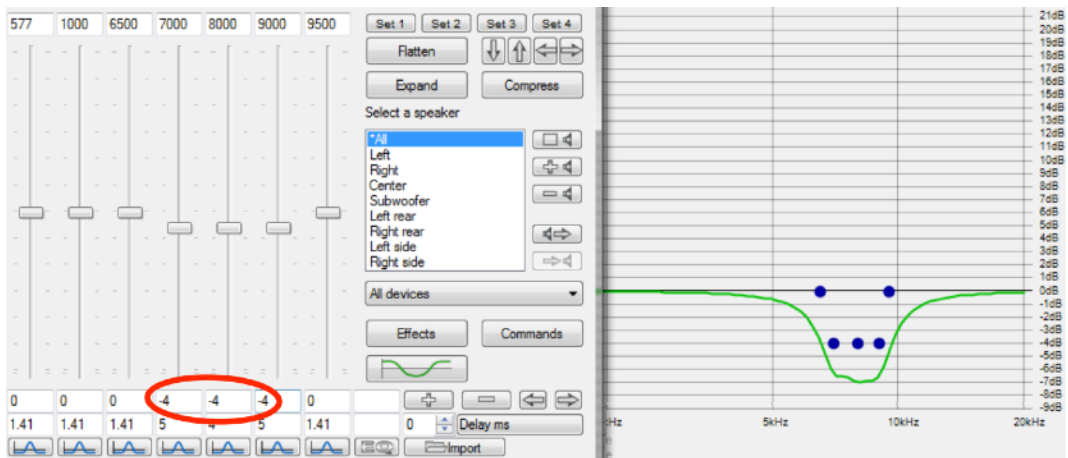


Fig. 22: 8000 hertz correction attempt 5

Basics of Headphone Sound

Changing all three of the gain numbers from -6 to -4 does the job. I could now tweak the gain individual values to create a prettier result, but the decimal place adjustments would be too small for my hearing to perceive. The reason I needed to make the -6 to -4 correction, however, is crucial. Overlapping filters accumulate where they overlap when of the same gain sign — such as all plus or all minus. Overlapping filters partially or completely cancel when of opposite gain signs, depending on the relative strength of their gain and Q numbers. This is both a strength and weakness of parametric EQ. In some contexts this is helpful, in others, a nuisance. When doing EQ by ear it generally means a lot of trial and error work to hone in on the correction shape you're trying to achieve.

The solution in Fig. 22 gets the job done, but it's based on using the default peak filter type. So next I experimented with combinations of high and low pass plus high and low shelf filters. Nothing I tried seemed to produce an equal or better result in fewer steps, so I choose to stick with original the all-peaks approach and call it a day

A few caveats:

- Even the best frequency response graphs, like Oratory's, are likely to be inaccurate for your own hearing in the higher frequencies. This is not the fault of the measurements but simply due to the fact that human ear anatomy is highly divergent, which in turn impacts our high frequency hearing.
- Even the best frequency response graphs, like Oratory's, are likely to be slightly inaccurate when applied to your own headphone. Manufacturing tolerances of the sound generating components of headphones are less than perfect. So your own headphone unit may have some audible departure from the available graphs. Whether this will be audible or not is another question.
- Current thinking is that a dip in the frequency response in the vicinity of the 9,000 to 10,000 hertz range is there for a reason — to provide an important spatial localization cue.
- You may be a proud member of the Order of the Bass Cannon, but headphone models differ greatly in how much added bass they can actually handle without seriously degrading the resulting sound. Many an open-back, dynamic driver headphone model comes from the factory with close to as much bass elevation as it can handle.



And that about wraps up this transfer of parametric knowledge from my cramped little cranium to your hopefully spacious and comfortably appointed one. I hope you've enjoyed the ride! The next [Headphone Essentials](#) episode looks at what to do when you can't use parametric EQ.