

Logitech Squeezebox Boom Audio Design

Caleb Crome
Sr. Hardware Engineer
Logitech, Inc.



Introduction

Ever since I bought my first Squeezebox from Slim Devices several years ago, I wanted a version with built-in speakers. I was not alone. Members of the extremely active Squeezebox community have built quite a few homebrew Squeezebox "boom boxes" over the years, mashing together Squeezebox hardware with amplifiers, speakers and power supplies from a variety of sources. After Logitech acquired Slim Devices and I was hired on, I was excited to be part of the team that would build the Logitech® Squeezebox™ Boom all-in-one network music player.

Over the past months, our team has worked extremely hard to build a compact, self-contained, high-performance network audio system. With advanced digital signal processing, a high-quality bi-amplified speaker design, an easy-to-use user interface, line input and subwoofer output, Squeezebox Boom is a system that can go in any room of the house and sound great.

After a quick tour of the high-level architecture and disassembly photos of the Squeezebox Boom, this paper will describe the audio architecture starting from a digital PCM signal (after any MP3, OGG, FLAC or other decoding), and will follow the signal through the digital signal processing (DSP) chain, digital-to-analog converters (DACs), power amplifiers, the speaker drivers and acoustical enclosure.

The DSP signal flow in Squeezebox Boom contains several processing stages that optimize the sonic experience. The primary DSP stages that will be discussed are: volume processing, bass management, StereoXL™ stereo enhancement, woofer-tweeter crossover, subwoofer processing, and driver protection.

Although we believe Squeezebox Boom sounds great today, because of the extremely flexible and upgradeable nature of the Squeezebox architecture, we can roll in new features on an ongoing basis to meet our customers' needs. Some of the future enhancements could include: automatic loudness filter for low-volume listening; dynamic range compression for low-volume listening; multi-band equalization; dynamic range compression for high-noise environments; and whatever else we can think up to help improve the overall experience.

System Overview

The diagram below (Figure 1) shows a simplified block diagram of the Squeezebox Boom audio system. For the purposes of this paper, we'll break it down between CPU section and audio section. The CPU section contains all the traditional Squeezebox components, such as the CPU, Ethernet, display, IO, and user interface. This paper will not cover the Squeezebox CPU and IO in any depth, but will be focused on the audio architecture and processing.

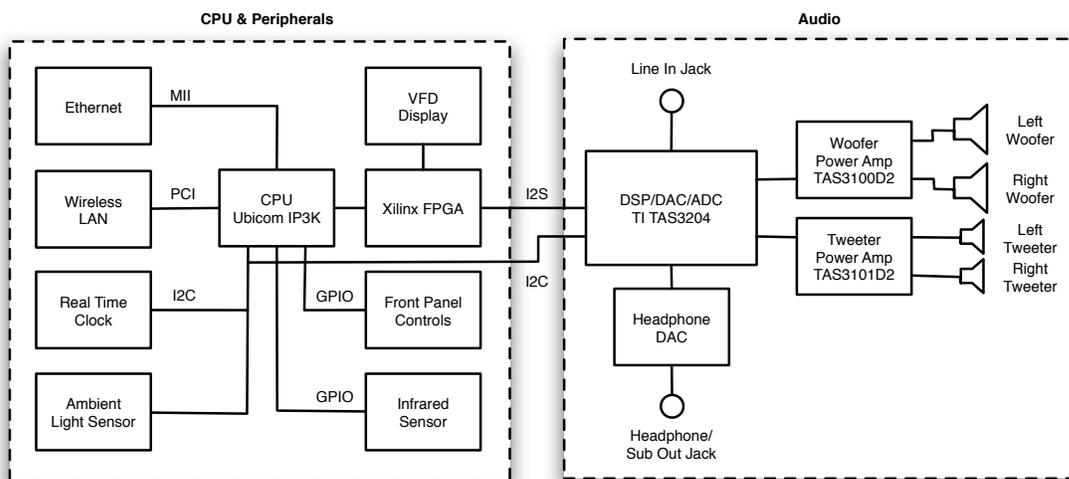


Figure 1: Squeezebox Boom Block Diagram. This paper will discuss the audio design, starting at the I2S & I2C interface between the CPU and Audio sections and follow the signal path through to the speakers, line-in and subwoofer-out.

As seen in Figure 1, the audio section includes the primary digital audio processor chip (TI TAS3204), the subwoofer/headphone DAC (Wolfson WM8501), the power amplifiers (TI TPA3100D2 and TPA3101D2), and the loudspeakers themselves.

Digital Audio Processor

The TI TAS3204 digital audio processor is a high-performance DSP optimized for audio applications, combined with high-performance (over 100 dB signal-to-noise ratio) DACs and analog-to-digital converters (ADCs) built in. It also can send and receive multiple channels of I2S to talk to secondary DACs or ADCs. The DSP processor itself is a 135 MHz, 48-bit DSP, with 28-bit coefficients and a 76-bit accumulator.

Secondary DAC

The secondary DAC, the Wolfson WM8501, is used to drive the subwoofer/headphone port. It gets its digital signal from the TAS3204 via an I2S interface. The DSP software is configured to process the signal differently depending on whether there is a subwoofer or a headphone plugged in to the output jack. The user can select from the user interface the type of plug-in they

are going to use.

Power Amplifiers

High-performance class-D amplifiers power the woofers and tweeters and can easily deliver full power to the speaker drivers with minimal distortion. We chose high-quality class-D amplifiers for both woofers and tweeters. It may be a bit atypical to use class-D amplifiers for tweeters, but we found there to be no significant sonic difference between the class-D and class-AB amplifiers we tested.

Hardware/Mechanical Design

We designed Squeezebox Boom, working closely with our industrial and mechanical designers, to simplify the assembly process and minimize risk, yet maintain top-notch audio performance in an attractive package. The basic configuration is a sealed enclosure consisting of a rear cup-shaped case and a front panel assembly, where the speaker wires pass from inside to outside through a single rubber grommet.

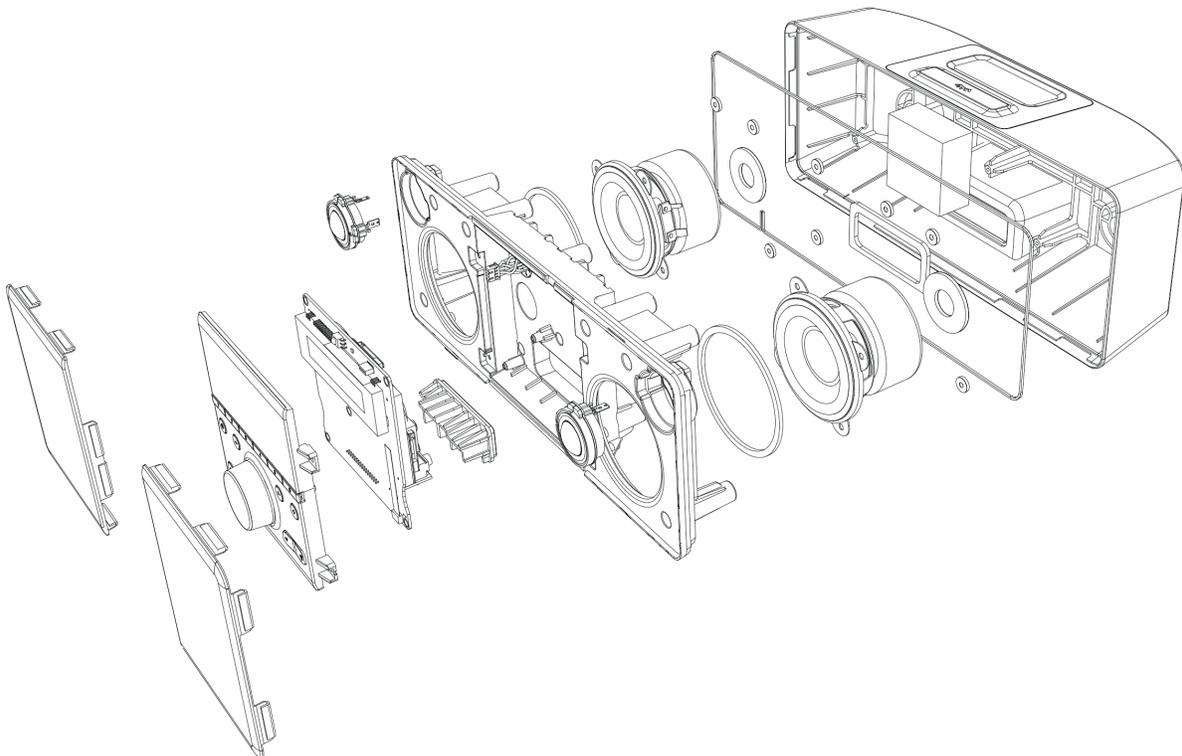


Figure 2: Exploded schematic view of Squeezebox Boom assembly

Loudspeaker Drivers

The drivers we chose were custom developed by Logitech's audio engineers to produce the best sounding products while maintaining reasonable costs. The woofers are 3" long-throw drivers with woven cloth cone and a rubber surround. They have a flat frequency response of between 100 Hz and 4 kHz. The tweeters are 3/4" soft-dome drivers that have a flat response of between 1200 Hz and 20 kHz. There is almost no signal loss all the way to 20 kHz. The woofer-tweeter crossover is set at 2 kHz.



Figure 3: Grills and front panel removed.



Figure 4: Main board removed, exposing the rear of the main board and acoustical chamber.



Figure 5: Acoustic housing disassembled. The woofers share a rear acoustic chamber. Production units have foam around the wires to prevent buzzing.



Figure 6: Loudspeaker Drivers. Each channel uses a woofer and tweeter

Audio Design

Squeezebox Boom is a bi-amplified design, using digital crossovers and independent DACs for each speaker, with a second independent crossover for the subwoofer output. The crossovers and equalization are implemented in software on a digital signal processor (DSP). This is the same technology that's found in high-end studio monitor speakers. Obviously, the Squeezebox Boom doesn't compete in bass performance with high-end studio monitors, but because of its advanced signal processing capabilities combined with very high-quality drivers, we believe we have created one of the best sounding products in its class.

Typical desktop speaker systems will be a 2.0, or occasionally a 2.1 system. Very few desktop speaker systems use true tweeters, and thus the high end will either be nonexistent or it will 'beam' with much more energy coming from the front of the system than off axis. This is a fundamental property of sound propagation. For the best quality sound, it's critical that loudspeakers be as omnidirectional as possible.¹ The result is more unified and balanced sound than can be achieved with other architectures.

¹ A Multiple Regression Model for Predicting Loudspeaker Preference Using Objective Measurements, Sean E. Olive, AES Convention 116, paper numbers 6113 and 6190. <http://aes.org/e-lib/browse.cfm?elib=12847>

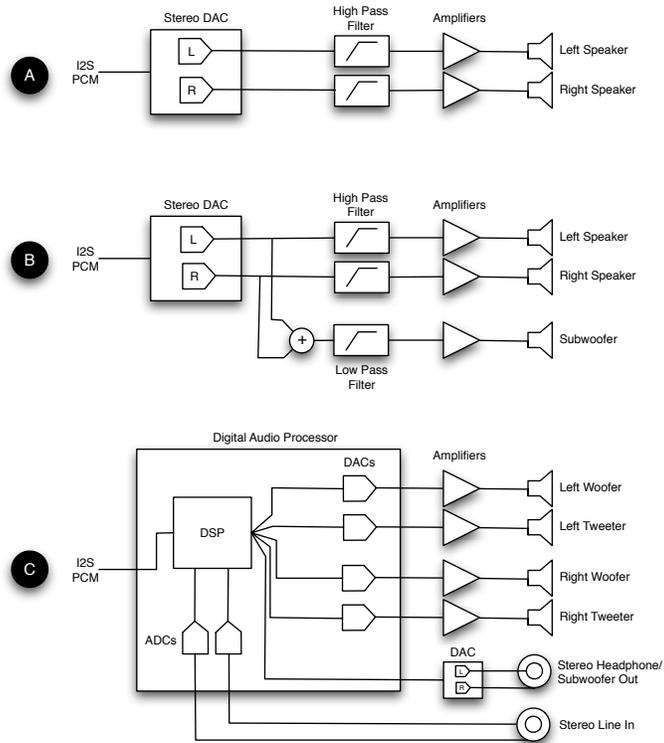


Figure 7: Comparison of different desktop speaker architectures.

In the figure above (Figure 7), Section A shows a typical 2.0 desktop speaker system with single drivers for each channel, and two power amplifiers. This design will generally compromise low- and high-frequency response.

Section B shows a desktop 2.1 system with stereo satellite speakers plus a subwoofer. This type of system will make up the low-end and sound full, but it will lack a flat frequency response to 20 kHz

Section C shows the design of Squeezebox Boom. Four separate speakers, four separate amplifiers, and six DACs all digitally controlled with a high powered DSP provide the ultimate in acoustic signal integrity and can produce great sound through the entire audio spectrum. Without a subwoofer, the Squeezebox Boom goes from a -3 dB response at 50 Hz (at low volume settings) to about 85 Hz (at high volume settings), all the way to 20 kHz. With the addition of a subwoofer, the entire audio band, from 20 Hz to 20 kHz, is covered.

The Squeezebox Boom DSP

The Squeezebox Boom uses a 48-bit-by-24-bit, 135 MIPS (million instructions per second) DSP core. With such significant DSP horsepower there are many ways to improve the sound quality; good crossovers are just the beginning.

One may wonder why we need a 48-bit data path, providing 289 dB of dynamic range, when 16-bits and 24-bits (96 dB to 144 dB) is good enough for most sound systems. The answer is that when performing mathematical operations (e.g., DSP) on a signal, the signal is often multiplied, divided, and added to many times. With a 16-bit data signal, any 16-bit signal processing will inevitably result in loss of signal fidelity, either by causing overflow, saturation, or extra quantization noise. Quantization noise, saturation and overflow are undesirable.

In order to perform DSP effectively, the DSP processing must have significantly more precision than the signal being operated on. Of course, a 16-bit processor can perform 32-bit, 48-bit, or higher precision math, but with corresponding CPU performance penalties. The 48-by-24-bit data path of the Boom DSP processor allows us to simply design audio processing algorithms, and minimize excess digital noise into the system – we can easily maintain greater than 100 dB of dynamic range in the digital domain and thus guarantee we don't degrade the original signal beyond the noise floor of the DACs. The large bit-depth of the DSP we chose makes DSP programming relatively straightforward.

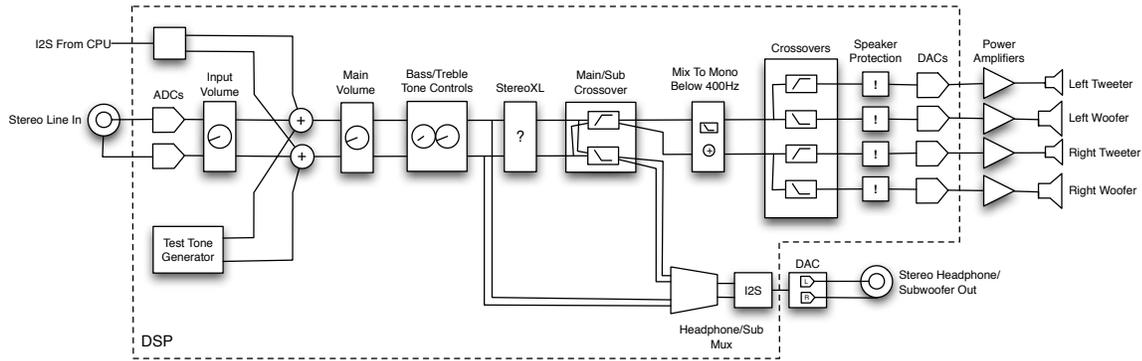


Figure 8: Block diagram of the Squeezebox Boom DSP flow (with other components for reference)

Volume Control

Regular DACs implement digital volume control by simply changing the gain of a signal from one sample to the next. Better DACs will detect a zero-crossing of the signal, and apply volume changes there. In both cases the volume change is very abrupt and audibly detectable.

The volume control on the Squeezebox Boom is implemented as a true digital ramping volume control so volume changes are completely click and pop free. Figure 9 shows the envelope (about 40 ms) that is applied to a signal as the volume changes from one setting to another. Regular DACs implement digital volume control by simply changing the gain of a signal from one sample to the next. Better DACs will detect a zero-crossing of the signal, and apply volume changes there. In both cases the volume change is very abrupt and audibly detectable. Using zero crossing detection is much better than without, but it's not perfect.

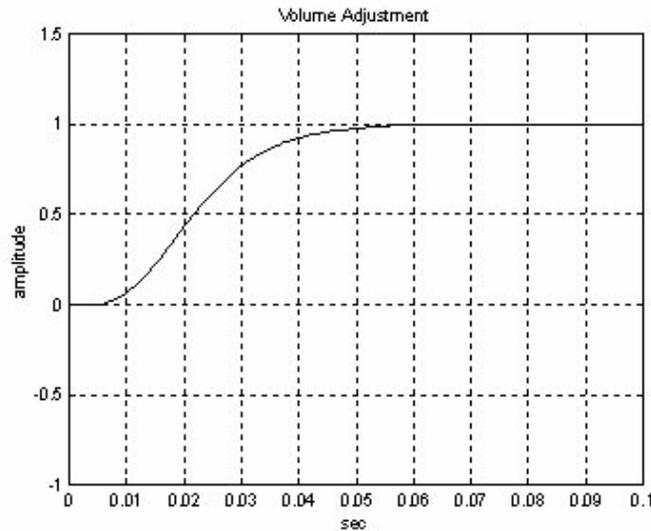


Figure 9: Click-less volume changes in Squeezebox Boom. This volume ramping profile shows how the volume gently changes from one volume to another. This shows a switch from 0 to 100% volume, but any other volume change follows a similar.

StereoXL Technology

StereoXL™ is a proprietary stereo expansion technology developed by Logitech audio engineers. It provides a significantly enhanced stereo sound field, expanding the sound stage beyond the physical extents of the speakers themselves.

I have always been wary of this type of processing for fear that it would wreck the audio quality in order to achieve the effect of stereo widening. However, once the engineers from the Logitech Audio business unit helped implement StereoXL in the Squeezebox Boom, the results were very surprising. The audio quality wasn't diminished at all, but the sound seems to come from everywhere.

That said it can be overdone, and the quality depends on the track used and the encoding used. In order to allow for varying user preferences, and track encodings there are 3 settings for StereoXL. The best one is typically in the middle (medium).

Crossovers

To produce outstanding audio quality in such a small form factor, we needed to optimize the two frequency extremes: low frequencies and high frequencies. This necessitated using two speakers (the woofer and tweeter) for each channel. Any time there is more than one speaker per channel, a crossover of some type is required to steer the high-frequency energy to the speakers with better high-frequency response, and low-frequency energy to the low-frequency speakers.

There are many different types of crossover “alignments,” with various advantages and disadvantages. For the Squeezebox Boom, we chose a very successful crossover type known as a 4th-order Linkwitz-Riley crossover.² This was chosen because of its high performance and

² Active Crossover Networks for Noncoincident Drivers; Siegfried Linkwitz, JAES Volume 24 Issue 1 pp. 2-8, February 1976. Available online here <http://aes.org/e-lib/browse.cfm?elib=2649>.

proven track record. The extreme flexibility of the networked music player plus a completely software controlled audio architecture allows us to continue to improve the product over time and make these updates easily available to customers.

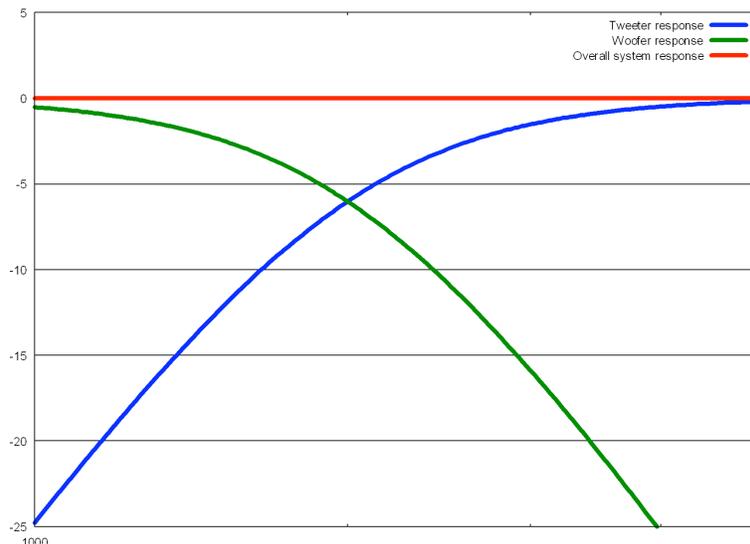


Figure 10: 4th order woofer/tweeter crossover implemented in Boom.

The Squeezebox Boom architecture is unique among the desktop systems we have evaluated in its ability to be optimized up to, and beyond its manufacture date. We were willing to spend more in component cost to bring the absolute best system to market possible and meet our time-to-market demands. We could have saved cost by using analog (passive or active) crossovers and eliminated the DSP processing all together; instead, we chose to build the best system we could while meeting our industrial design and budget requirements.

Bass Extension

To make the Squeezebox Boom sound fantastic at moderate to low listening levels, we need to extend the bass response. A typical bass control is a very different concept from true bass extension. The red line in Figure 10 shows the measured bass response of the Squeezebox Boom, without any bass extension.³

³ In the absence of a large anechoic chamber, it is difficult to make free-field acoustic measurements down to 20 Hz. For low-frequency measurements we use a method developed by Richard Small (of Thiele-Small parameter fame) that uses a microphone inside the sealed enclosure. For more information on this technique, see Simplified Loudspeaker Measurements at Low Frequencies, JAES Volume 20 Issue 1 pp. 23-33; February 1972.

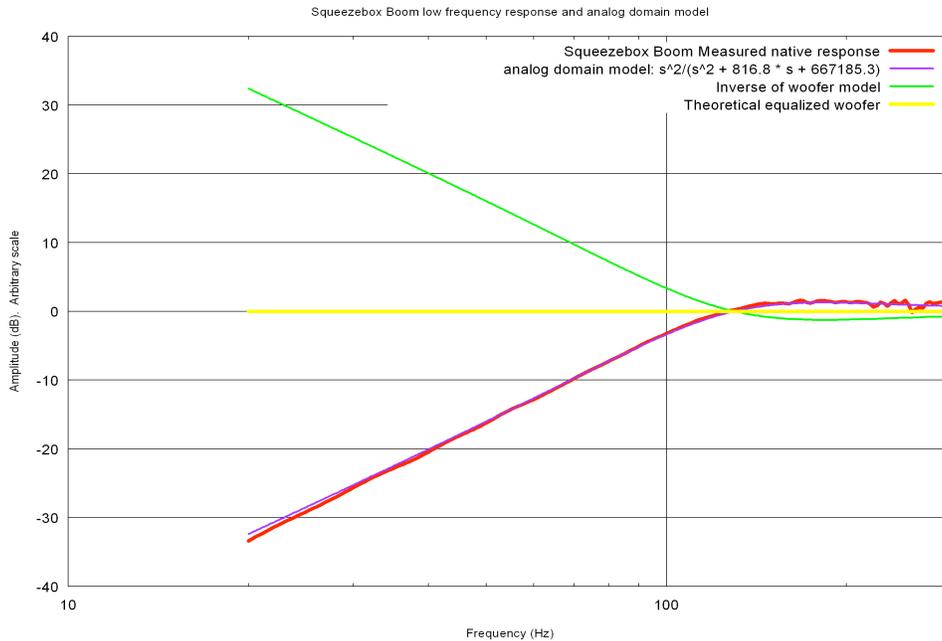


Figure 11: Frequency response of Squeezebox Boom, with theoretical low frequency compensation filter. Red: Measured response; Blue: theoretical model; Green: compensating filter; yellow: Theoretical frequency response of the equalized system.

As can be seen from the plot, -3 dB point is about 100 Hz. Squeezebox Boom comes in a very small form factor, so one wouldn't really expect a much lower cutoff frequency but you sure do *want* it to go lower. With signal processing, we can extend the bass response much lower at moderate listening levels by doing some clever signal processing.

The low frequency roll-off of a closed box (i.e. non-ported) loudspeaker is a 2nd-order roll-off of 12 dB/octave. (Take a look at the measured low frequency response, and one can see that in fact, the drop between 100 Hz and 50 Hz is in fact about -12 dB).

By modeling the low-frequency response mathematically, we can create an inverse filter that compensates for the speaker roll-off, and then we can put a new low frequency roll-off at whatever frequency is desired. Of course, the tradeoff is that as the bass response extends lower, it's necessary to drive much more power into the speakers at low frequency, risking distortion. At low volumes, this can be done very effectively without excessive distortion. Of course, as the volume goes up, the excess boost at low frequencies causes either the power amplifier or the loudspeaker itself to go beyond its linear region, causing significant distortion, system shut down, or even physical damage to the speakers.

The blue line of Figure 11 shows the model of the woofer response. By converting this to a digital domain representation and effectively inverting it, we get the compensation filter shown in green in Figure 11, and the overall response shown in yellow (i.e. perfectly flat down to 20 Hz).

In general, it would not be wise to implement the green digital filter of Figure 11 without any other filtering, since it would quickly over-drive the speakers and power amplifiers. This 32 dB boost at 20 Hz corresponds to more than 1000 times more energy at 20 Hz than at 200 Hz. This is not reasonable for any but the lowest listening level. To make matters worse at low listening levels, an additional signal boost is needed at low frequencies to compensate for the loss of sensitivity to low frequencies at low SPL levels.

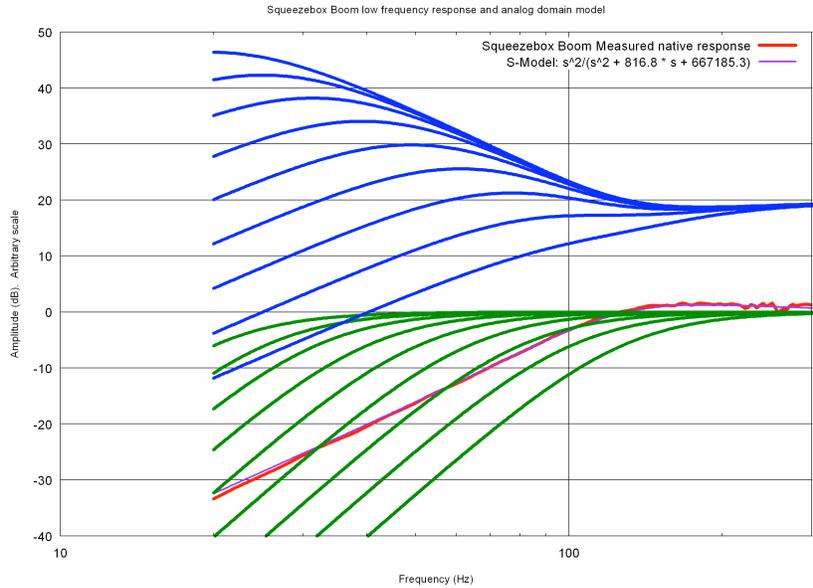


Figure 12: Family of low-frequency equalization curves (blue) and simulated low-frequency system response (green). The curves show cutoff frequencies ranging from 20Hz to 137Hz in 1/3-octave steps. For clarity, the blue curves are offset by +20 dB. The green curves correspond to the final overall frequency response of the system (simulated in this figure). They are 4th-order Linkwitz-Riley high pass filters.

The method we chose for bass extension with Squeezebox Boom is to make the overall low frequency roll-off a 4th-order Linkwitz-Riley filter. The reason for that choice will become clear in the subwoofer section of this paper. By applying the compensation of Figure 12 in series with a variable frequency 4th-order Linkwitz-Riley filter, we can choose any low frequency cutoff frequency we desire as seen in Figure 12, which shows an entire family of filters with cutoff frequencies ranging from 20Hz to 137 Hz in 1/3 octave increments. The choice of cutoff frequency is based primarily on the final signal strength going to the speakers, and whether or not there is a subwoofer plugged in. The true beauty of the Squeezebox Boom architecture is that we can dynamically and silently switch among any of the filters as volume settings and/or signal strength changes. Currently, we implement a sliding cutoff frequency based on volume setting.

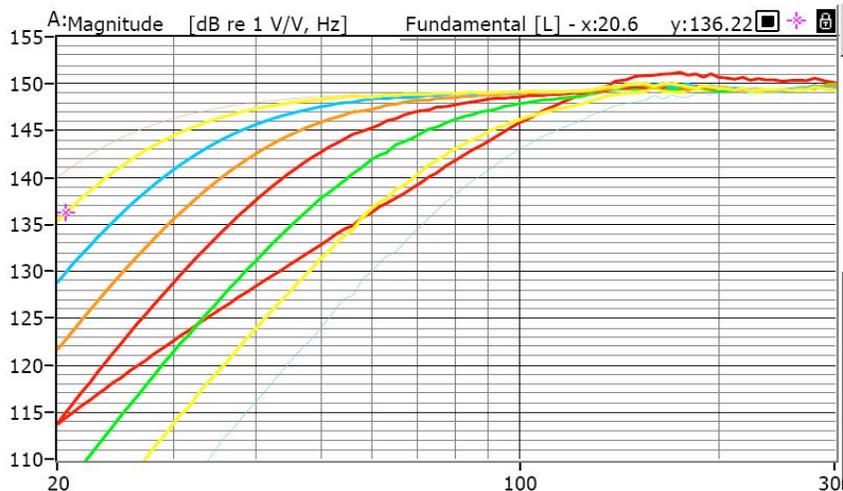


Figure 11: Actual in-box frequency response measurements of the family bass extension filters used in Squeezebox Boom.

The difference between with and without bass-extension was surprising to me. At low listening levels, with significant extension, the listener can easily be fooled into thinking a subwoofer is attached.

Subwoofer/Headphone Output

The Squeezebox Boom has a subwoofer/headphone output. A Wolfson WM8501 DAC, the same chip that is use in Squeezebox Receiver, drives this output. It is in a slightly different configuration optimized for headphone and subwoofer output, as opposed to line out. It is capable of driving about 1V RMS into 16-ohm load, or about 1.7V into a high impedance load. The WM8501 DAC is driven from the same DSP chip that drives the speaker DACs, and it can be controlled completely independently. This means that in headphone mode, the output will get full stereo output with full frequency response.

When the device is configured for a subwoofer, we automatically slide the low frequency cutoff filter for the woofers up to 100 Hz, and create a complementary low pass filter at 100 Hz for the subwoofer output. Both filters are also 4th-order Linkwitz-Riley filters, which create a very nice crossover between the Squeezebox Boom and the stand-alone subwoofer. The beauty here is that the user only needs to turn the 'frequency' knob of the subwoofer up all the way – no need to fine-tune the frequency knob since the Squeezebox already knows the optimal cutoff frequency and will automatically implement the best possible filter.

Line-In

The Squeezebox Boom has a line-level input that can be put to many uses. Since it goes into the built-in ADCs in the DSP, we can perform virtually any processing on it we want and sent it back out the speakers, or headphone/sub-output. Some potential uses are:

- Play music from a friend's portable MP3 player.
- Connect to a computer's audio output and use the Squeezebox Boom as a PC speaker system. This allows you to play streaming music and simultaneously mix in the PC sounds.
- Connect an amplified microphone to it, install the community-developed Karaoke plugin and use the Squeezebox Boom as a karaoke machine.

Conclusion

Does it all really work? The short answer is yes! By taking into account and managing the many factors that push and pull on a product design (cost, size, time-to-market, features, and performance), we believe we have built the best streaming audio system available in its class.

There are several future enhancements that we would like to make to the Squeezebox Boom in future software releases. Some of the future enhancements could include: automatic loudness filter for low-volume listening; dynamic range compression for low-volume listening; multi-band equalization; dynamic range compression for high-noise environments; and whatever else we can think up to help improve the overall experience.

Most importantly, feedback from our customers will guide us toward what enhancements, fixes, or corrections need to be implemented, and these can be added in regular software updates. For some changes, there won't be any DSP or firmware related changes necessary, since DSP configuration commands can be sent directly from the streaming server.