

The Future of Music

Part One: Tearing Down the Wall of Noise

By Suhas Sreedhar

Posted 1 Aug 2007 | 18:48 GMT

You're listening to your favorite Pink Floyd CD on your home stereo when you accidentally hit the "change CD" button on the control panel. All goes quiet for a bit as your CD player urgently shifts to play whatever is in the next tray. With dread, you desperately reach for the volume knob, but it's too late--your speakers blast the latest Green Day album. Reacting like you were just pricked by a pin, your hand jolts to the volume knob and turns it down. You breathe a sigh of relief. But that's not the end of it. Ten minutes later you feel that something isn't right. Even though you love this album, you can't listen to it anymore. You shut it off, tired, puzzled, and confused. This always seems to happen when you switch from a classic album to a modern one. What you've just experienced is something called overcompression of the dynamic range. Welcome to the loudness war.

The loudness war, what many audiophiles refer to as an assault on music (and ears), has been an open secret of the recording industry for nearly the past two decades and has garnered more attention in recent years as CDs have pushed the limits of loudness thanks to advances in digital technology. The "war" refers to the competition among record companies to make louder and louder albums. But the loudness war could be doing more than simply pumping up the volume and angering aficionados--it could be responsible for halting technological advances in sound quality for years to come.

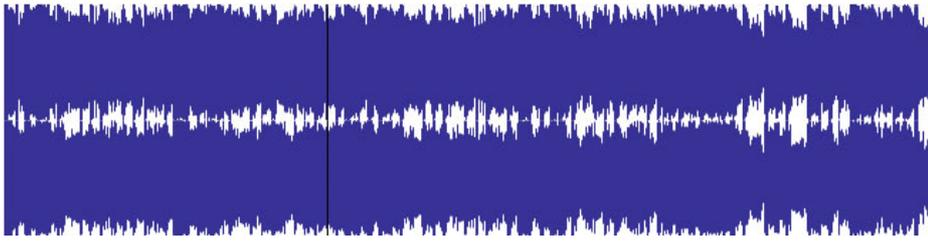
Overcompression

The smoking gun of the loudness war is the difference between the waveforms of songs 20 years ago and now. Here is an example:



[\(/images/aug07/images/wave1.jpg\)](/images/aug07/images/wave1.jpg)

A waveform from the late 80s / early 90s. *Click on image for larger view.*



[\(/images/aug07/images/wave2.jpg\)](/images/aug07/images/wave2.jpg)

A waveform from now. *Click on image for larger view.*

The second waveform not only has a higher amplitude than the first but is also highly compressed--there is very little difference between its highest points and the average level. In other words, the new song has a drastically reduced dynamic range--the difference between the loudest parts (the peaks) and quietest parts of the sound.

Music, like speech, is dynamic. There are quiet and loud moments that serve to accentuate each other and convey meaning by their relative levels of loudness. For instance, if someone is talking and suddenly shouts, the loudness of the shout, in addition to the content, conveys a message--be it a sense of urgency, surprise, or anger.

When the dynamic range of a song is heavily reduced for the sake of achieving loudness, the sound becomes analogous to someone constantly shouting everything he or she says. Not only is all impact lost, but the constant level of the sound is fatiguing to the ear. So why is achieving greater and greater loudness so important that the natural ebb and flow of music has been so readily sacrificed?

The answer goes back to the beginnings of recorded music.

The Vinyl Era

Loudness has always been a desirable quality for mainstream popular music. The louder a song is overall, the more it stands out from ambient noise and the more it grabs your attention. Studies in the field of psychoacoustics, which investigates how humans perceive sound, show that people judge how loud a sound is based on its average loudness, not its peak loudness. So even though there might be two songs whose loudest parts reach the same loudness level in decibels (dB), the one with the higher average level is generally perceived as louder.

As far back as the early 1960s, record companies began engaging in a loudness battle when they observed that louder songs in jukeboxes tended to garner more attention than quieter ones. To maintain their competitive edge, record companies wanted to keep raising the loudness of their songs. But the physical properties of vinyl records limited engineers' ability to perpetually increase loudness.

A vinyl record consists of a lacquer into which small V-shaped grooves--vibrational transcriptions of analog sound--are cut. Creating a record in the studio involves a process called mastering, where songs are sonically adjusted and placed into an appropriate order to fit the given medium's requirements. Mastering for vinyl was always a balance between loudness and playing time. The louder you wanted a song to be, the wider the groove needed to be in order to accommodate the larger amplitude of the transcription. Since there's only a limited amount of usable surface area per vinyl disc, gaining loudness meant sacrificing playing time, especially on a long playing (LP) record where upwards of six songs were often fit on each side of the disc.

In order to save the cost of manufacturing an excessive number of vinyl discs per album, playing time usually won out over loudness. Live music typically has a dynamic range of 120 dB, peaking at about the same loudness of a jet engine (though some concerts have gone even louder). Vinyl records tend to have about 70 dB of dynamic range. This meant that in order to fit a song onto a record, it either needed to have its overall amplitude reduced or it needed to be compressed--have its peaks brought down to a lower level--to fit within the given range. How much of each was done varied from record to record and defined the art of mastering. However, the tools of analog signal processing limited the amount of compression that was possible.

Analog compressors of this era were basically voltage control amplifiers that varied the level of an output signal based on a control voltage, similar to the devices that regulate signals in AM radio. Such compressors were typically used on individual instrument tracks (vocals, guitar, etc.) to add clarity to a sound or to change the sound of an instrument for effect. However, in some cases, such as with the hit singles put out by Motown Records (for example, "Want Ads" by Honey Cone), compressors were used to boost the loudness of songs to higher-than-average levels. Mastering engineers accomplished this by reducing the dynamic range of a song so that the entire song could be amplified to a greater extent before it pushed the physical limits of the medium. This became known as "hot" mastering and was typically done on singles where each side of the record only contained one song. Generally, however, the average level of songs and albums stayed relatively the same throughout the period.

"CD" Behavior

"The invention of digital audio and the compact disc became a new fuel for a previously existing loudness race," says Bob Katz, a renowned mastering engineer and one of the first outspoken critics of dynamic-range overcompression. "The reason is that the analog media did not permit what we would call 'normalization' to the peak level."

When the compact disc (CD) was introduced in the early 1980s, there was much for audiophiles to be happy about. Digital audio removed many of the physical restrictions vinyl had imposed, such as concerns about surface noise (caused by dust, scratches, the lacquer itself, and so on) and limited dynamic range. The CD was capable of supporting a dynamic range of about 96 dB. For most of the 1980s, when CDs were still high-end products and mastering engineers largely did not have access to digital signal-processing technologies, albums released on CD tended to make use of this better dynamic range.

Unlike vinyl, which had varying loudness limits due to its physical characteristics, the CD had a definitive peak loudness limit due to its specified digitizing standard, a form of pulse code modulation (PCM). PCM had previously been used in telephony as a method of digitizing an analog signal. When an analog signal is sampled for digitization, each level of the signal is quantized (stored as a number in binary). How frequently samples of the signal are taken is specified by the sampling rate, and the total number of unique quantization levels capable of being stored is determined by the number of bits. When Sony and Philips specified the standard for CD audio, they determined that the sampling rate would be 44.1 kHz with 16 bits per sample. Using the rule of thumb, the approximation of 6.02 dB of dynamic range per bit gave CD audio roughly 96 dB of dynamic range. The highest loudness level (16 bits of all 1's) was designated as 0 decibels full scale (dBFS). Lower levels were assigned negative numbers.

In the 1980s, CDs were mastered so that songs generally peaked at about -6 dBFS with their root mean square (RMS)--or average levels--hovering around -20 dBFS to -18 dBFS. As multidisc CD changers began to gain prominence in households toward the end of the decade, the same jukebox-type loudness competition started all over again as record companies wanted their CDs to stick out more than their competitors'. By the end of the 1980s, songs on CDs were amplified to the point where their peaks started pushing the loudness limit of 0 dBFS. At this point, the only way to raise the average levels of songs without having their loudest parts clipped--the digital equivalent of distortion, where information is lost because it exceeds the bit capacity--was to compress the peaks.

While analog compressors had been limited in the extent to which they could reduce peak levels, digital compressors were much more powerful. As mastering engineers began to get hold of digital signal-processing tools, they were able to "hot" master songs even more. The process was similar to what had been done on some vinyl singles--peak levels were brought down by a certain amount, and then the entire waveform was amplified until the (now reduced) peaks once again reached 0 dBFS. The result? The average level of the entire song increased.

The 1990s saw average amplitude levels go from around -15 dBFS to as much as -6 dBFS in extreme cases. Most songs in this decade, however, remained at around -12 dBFS. The 2000s saw the loudness war reach its height, with most current songs having an average level of -9 dBFS or higher. From the mid 1980s to now, the average loudness of CDs increased by a factor of 10, and the peaks of songs are now one-tenth of what they used to be. The loudness war is also not just confined to the big four record companies (Warner Music Group, EMI, Sony BMG, and Universal Music Group). Overcompression is now widespread and performed by independent labels and international record companies.

The CD Is Dead; Loudness Continues

The biggest change from 15 years ago to today is how people consume music. With more than 100 million iPods sold worldwide as of early this year, more and more people are listening to music on the go rather than at their home stereos. Physical media like CDs are on their way out. And yet overcompression continues to plague the music world.

Even though the CD might be in its death throes, most digital music available online was mastered for CDs. Popular formats like MP3, AAC, and Free Lossless Audio Codec (FLAC) merely use data-compression techniques (not to be confused with dynamic-range compression) to reduce the amount of data a song encoded in PCM takes up. As long as the specter of CDs continues to haunt the online world, downloaded songs will still be subject to overcompression.

But the problem doesn't just lie on the production end. If people are listening to songs in a noisy environment--such as in their cars, on trains, in airport waiting rooms, at work, or in a dormitory--the music needs to be louder to compensate. Dynamic-range compression does just that and more. Not only does it raise the average loudness of the song, but by doing so it eliminates all the quiet moments of a song as well. So listeners are now able to hear the entire song above the noise without getting frustrated by any inaudible low parts.

This might be one of the biggest reasons why most people are completely unaware of the loss of dynamics in modern music. They are listening to songs in less-than-ideal environments on a constant basis. But many listeners have subconsciously felt the effects of overcompressed songs in the form of auditory fatigue, where it actually becomes tiring to continue listening to the music.

"You want music that breathes. If the music has stopped breathing, and it's a continuous wall of sound, that will be fatiguing," says Katz. "If you listen to it loudly as well, it will potentially damage your ears before the older music did because the older music had room to breathe."

Some audiophiles find relief by going back to the past. A few musicians still continue to release their albums on vinyl records (in addition to CDs and online formats). Because vinyl cannot support the loudness that CDs can, these modern vinyl releases are much quieter than their CD counterparts. But they are often less compressed as well, and, in some instances, remastered in a way that is as dynamic as albums released in the 1960s and 1970s.

One of the most prominent examples of this is the recent Red Hot Chili Peppers album *Stadium Arcadium*, which was remastered for vinyl by mastering engineer Steve Hoffman with the intent of providing full dynamic sound. Hoffman is one of the few mastering engineers who have actually refused to take certain jobs because he's been

asked to overcompress music. "[It happens] all the time," says Hoffman. "At least once a week."

But turning to vinyl for uncompressed music might not always provide salvation. In order to save the cost of remastering, record companies might simply take the compressed master of a song, reduce the overall loudness, and place it on vinyl. Katz warns, "You could take the Red Hot Chili Peppers recording and put it onto vinyl just as it came from CD, and it would sound just as fatiguing. [The only difference is] you'd just have to turn the volume control up because you couldn't get the peak level the same."

Tearing Down the Wall

Audiophiles looking to the future for relief from overcompression see a cloudy picture. DVD-Audio and Super Audio Compact Disc (SACD) are two high-fidelity formats that were thought to be solutions to the loudness war. Both formats offer not only a greater dynamic range than CD but also higher sampling rates. This allows for frequencies higher than what most humans are capable of hearing to be encoded onto the medium, addressing a common complaint by people who prefer analog over digital because they claim they can hear these frequencies.

DVD-Audio uses PCM encoding that can support 24-bit, 192 kHz stereo sound (contrasted with the CD's 16 bit, 44.1 kHz) yielding 144 dB of dynamic range, 14 dB over the human threshold of pain. SACD, like the CD, was developed by Sony and Philips and uses a form of pulse-density modulation (PDM) encoding branded as Direct Stream Digital. Basically, instead of having 16-bit samples at a frequency of 44.1 kHz, it takes 1-bit samples at 64 times that rate (2.82 MHz). It has a dynamic range of about 120 dB. Additionally, both SACD and DVD-Audio are capable of high-fidelity five-channel surround sound.

Since their introduction in 2000, however, neither format has taken hold. An overwhelming majority of releases have been of the classical music genre, which has generally not been subject to overcompression to begin with. So even if audiophiles wanted to spend upward of \$300 for a DVD-Audio or SACD player, chances are they won't be able to buy their favorite popular albums in either medium.

Since music has gone online, the possibility of having high-fidelity digital files remains, and formats such as FLAC are capable of supporting 24-bit audio. Slim Devices, a company acquired last year by Logitech, has created two products--the Squeezebox and the Transporter--that wirelessly stream digital files from a computer or the Internet to high-end stereo receivers. Both are capable of handling 24-bit audio, but the problem, says Sean Adams, former CEO of Slim Devices, is lack of content.

"If we're going to go to higher levels of sound quality, the real problem is actually getting the content out. Right now, unfortunately, the industry has kind of gone backward from CD quality. When MP3 came out, [it was called] CD quality when it really wasn't," says Adams. "We've made some improvements since then with better [compression techniques], but it's really a function of people demanding better sound quality. That has to happen first before the [recording] industry's going to start producing it."

Overcompression, however, seems to be one of the biggest obstacles to overcome. With music being compressed to have smaller and smaller dynamic ranges, the need for the next high-fidelity audio format vanishes. If record companies aren't making use of the full dynamic capability of CDs, then why bother moving to another format with even more potentially unused capability? And with the average consumer being either completely unaware of, or only subconsciously irritated by, the current state of overcompressed music, there is little incentive for sound quality to progress. Consequently, all the potential benefits of higher-quality audio--lifelike dynamic range, greater frequency response, and multichannel surround sound--remain unseen, even though the technology exists today. Audiophiles are forced to return to vinyl and analog recordings that should have been obsolete 20 years ago.

But there might still be hope for getting out of the loudness war. RMS (average) normalization algorithms, such as Replay Gain, have been implemented in many digital audio players and work to bring all songs in a digital library to the same average level. With Replay Gain enabled, songs originating from many CDs are processed and played back at a consistent average level of loudness. This helps listeners because they no longer have to adjust their volume each time they go from one album to another. And while such normalization cannot undo the compression of music (it amplifies or reduces the song in its entirety), it counteracts any efforts that were put in to make one song louder than another, essentially nullifying the loudness war altogether.

Many hope that widespread implementation of technologies like Replay Gain will make record companies see that further and further compression in the name of competitive loudness is a feckless task, and slowly but surely popular music will begin to return to a dynamic, less-compressed state. In fact, many digital audio players have caught on; Winamp uses Replay Gain, and iTunes has its own normalization option called Sound Check, which also works on iPods.

Whether the loudness war can end and give rise to the next generation of high-fidelity audio depends heavily on the attitudes of consumers. Unlike the CD and DVD video, there is no overwhelming industrial push toward the next level of sound quality. How songs and albums will sound will depend entirely on whether or not the listener actually cares about the intricacies of the music.

For a multimedia version of this article click [here \(http://spectrum.ieee.org/aug07/5480\)](http://spectrum.ieee.org/aug07/5480)

To go to Part 2 of this article click [here \(http://spectrum.ieee.org/aug07/5499\)](http://spectrum.ieee.org/aug07/5499)