Introduction to Audio Restoration using Computer Applications

by Michael Ljunggren (mike@physto.se)

1.0 Abstract

So you want to carry out some audio restoration? But you don’t know how it is done...? Ok, everybody knows what good music sounds like (except our parents...?) and we all know good audio quality from bad, but altering music to make it sound better is not a trivial task. Most of us can say what *sounds* wrong, but often don’t have any idea what to do about it.

This document is first of all written for people who are real beginners in the field of audio restoration. But hopefully intermediate and professional users will also benefit.

This paper addresses the use of desktop computers for audio restoration. There are many methods that can make an old recording sound better. The paper briefly describes how to use them, and why.

Lets start with the basics...

2.0 General approaches when restoring audio

You should really consider the next sentence. - If *you* were not the person creating the original sound or music, then the actual creator should be the person dictating how your restored result should sound! The goal should be to make the recording as good as new. This implies that “applying a lot of echoes and filters” will counteract the definition of “restoration”

And given that a “rule” is always only a “suggestion”:

Rule number one in audio restoration: *The restored result should sound just like the original sound*, just the way the creator intended it. Common mistakes are to apply strong filters to make the result more “hi-fi”, “cool” and “new”, or to place the music in a great cathedral-like room simulator. The only filter you should apply is one that brings the result closer to the original. The applied reverb should be such that it matches the original ambience.

Rule number two: *Don’t overdo it!* Today’s tools makes it possible to apply effects that are absolutely free from artefacts, therefore fooling the user into using too much. You *can* apply lots of effects, but *should* you? Many professional sound technicians tend to use an approach where you can “do all the restoration you want - as long as the original sound is preserved”. It does not take much until you start loosing details and feeling in the restoration. If you are more reserved in the restoration process, it may mean that you can still hear cracks & pops, and some noise in the background. *But you can also hear all the beautiful music*. It is an art to “remove everything unwanted and leaving everything wanted” at the same time. Why not go for “leaving everything wanted”? Most listeners tend to enjoy pieces of pristine music mixed with a little bit of noise, as opposed to a noise-free song that lacks feeling.

3.0 Recording different sources of audio

Well, the sound has to come from somewhere. Audio *restoration* assumes that we are dealing with *pre-recorded* material. Please refer to a book on recording for further information about producing original audio.

The only step where we “record” something is when transferring the original recording into the workspace of the computer. If the source is analog, we are *sampling*, and that is usually done by a soundcard or a dedicated AD converter.

Minimizing the chain of machines. Whenever you feed a signal through analog devices you *always* add noise to it. You should therefor use as few devices as possible, and have the volume as loud as possible through all stages.

The signal should be strong at the beginning of the chain so you can turn the volume down at the end. (This will also lower the volume of the noise.) If the signal is weak at the beginning (the difference between noise & signal volume is less), the noise will be more prominent if you have to amplify it in the end of the chain.

Try to have such a strong signal that you are forced to turn the volume down using the soundcard controls (most soundcards have a control panel for volumes and filters).

3.1 Vinyl

This is the oldest way of storing audio. Actually, when Thomas A. Edison first stored audio in 1877, he used cylindrical rolls of wax to engrave the movement of a vibrating needle. When the wax had hardened, sound could be reproduced by letting the grooves move the needle. Fantastic! Up until very recently this has been an ongoing method of storing sound waves, but instead of using cylinders, there have been platters, and instead of wax a plastic called *vinyl* has been used. Edisons recording, “Mary had a little lamb”, can be downloaded from the Internet at http://archervalerie.com/audio.html

Using vinyl for storing sound has both advantages and disadvantages. One drawback is that the quality depends on the durability of the physical material. Every crack and irregularity will be heard in the sound; and all vinyl record *have* cracks, it’s that simple. Another bad feature is that on top of the beautiful sound, there will be a very low frequency rumbling noise. And of course; all analog ways of storing sound have noise in them.

3.1.1 Sampling a vinyl record

It is not trivial to sample a vinyl record. You cannot simply plug the output of your turntable to your soundcard. The signal is far to weak, and must be amplified. You must also apply a *RIAA* filter to the signal. (A RIAA filter amplifies the bass and lowers the treble in a strictly specified way.)

If you have a dedicated phono amplifier, you are fortunate. If you don’t have one, you can use the one that is built into your home stereo.

Connect your turntable into your stereo. The inputs that are labelled ‘*Phono*’, or similar, imply that the signal is amplified sufficiently, and that a RIAA filter is applied. You then feed this signal into your soundcard which is done by taking the signal through the ‘Tape rec’ or a ‘Line out’ output of the stereo equipment. This signal will be matched with the ‘Line in’ input of a soundcard.

To summarize; you cannot plug the turntable into the soundcard directly. You will not hear anything because the signal is to weak. Secondly, the signal path is limited to three machines: the turntable, the stereo equipment and the soundcard. The volumes through these are probably already optimized.

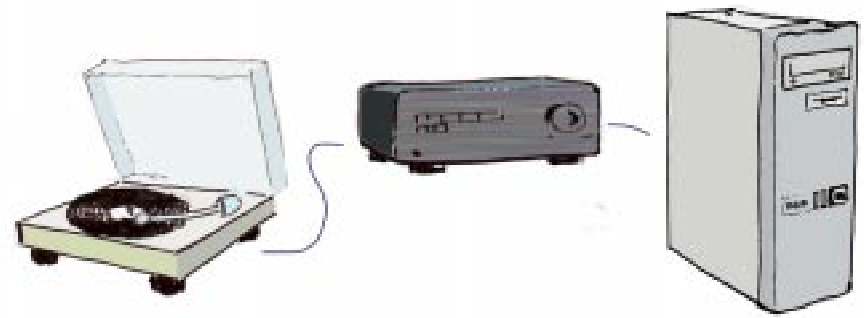

3.2 Tape

Analog tape consists of *many* extremely small magnets stored on a layer of plastic tape. When you record a sound onto the tape, you arrange the physical orientation of these magnets using the magnetism generated by the tape head. When you drag these magnets over the head during playback they will induce an electrical current, that is then heavily amplified by the electronics of the tape recorder.

An analog tape has substantial noise built into it. Unlike vinyl recordings, tapes must be used with some sort of noise reduction system to hide the noise. One other disadvantage is that the magnetized surface is not always consistent. Tiny dropouts will occur every now and then. The speed of the tape can also be hard to control and is dependant on the drive mechanism to a great extent.

To sample a tape, you can most often plug the output directly to the input of your soundcard. Find the ‘Line out’ or ‘Tape out’ on your machine. The signal path is minimal and involves only the tape recorder and your soundcard.

Be sure that the volume is high when playing the tape. Then reduce it using the soundcard controls.

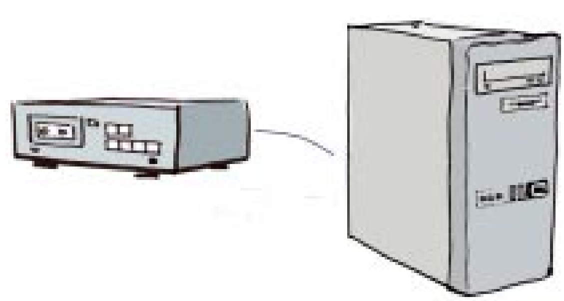

4.0 Means of restoring sound

This section assumes that you have sampled audio into some digital audio format in your computer. One to use is the *standard CD audio format*: 44100 samples per second, 16 bits per sample, stereo. If you intend to store the result on DAT (see section) you might also consider using a sampling frequency of 48000 Hz. Not all programs can handle more than 16 bits per sample, but if they can, you could use 24 or 32 bits.

It is recommended that you sample the audio in the same format that it will be saved in; converting between formats will not be perfect, although converting between 48kHz and 44.1kHz is pretty transparent.

4.1 Removing noise

But then, what *is* noise? In theory, everything *unwanted* can be considered to be noise, but in this section we can stick to “uniformly spread, random impulses”. A random signal sounds like noise. The behavior of noise comes from the nature of the recording but is always made up of tiny impulses.

Most often, noise can be heard all over the spectrum. If the noise is uniformly spread over the spectrum it is called “white noise”. If it is not uniform, it is called “colored noise”.

4.1.1 The principle of removing noise with *FFT’s*

This is a fantastic way of, not simply hiding but, *removing* noise. It has not been possible to use this technique in a practical way until today’s powerful computers became available, as one is totally dependant on *raw computational speed*. Put simply, this is not an easy task for a computer to do. Hundreds of millions of calculations are required to clean up a single second of sound! But still, with today’s desktop computers you can do this in real time.

Using this technique, the sound is divided into short sections (usually about 1/10 th of a second each). By viewing this piece of sound, one can see the shape of the waveform during 1/10th of a second. Sound represented like this is in the “time space” or “time domain”, meaning that you can see how the signal changes *in time*. But you can only guess what *frequencies* are hidden inside this little piece of sound. When looking at the frequency information of a sound you are using the “frequency space” or “frequency domain”. But by looking at the sound this way you have no idea how the sound changes in time, and this is partly why you have to divide the whole song into tiny 1/10th sec. pieces. An algorithm that is often used when transforming something from the “time space” into the “frequency space” is called “Fast Fourier Transform” or FFT. Never mind how it works, but it gives us the information about the frequencies inside...

The “time space”: shows the behavior of the sound over *time*.

The “frequency space”: shows the behavior of the sound *in terms of its frequencies.*

The ability to look at the *frequencies* of a sound gives us powerful means of restoring it. The ear, of course, perceives *frequencies* and not *wavy curves*.

Now we can affect one frequency without disturbing the others. We can also make assumptions about noise. Compared to the "music" in the recording, the noise is usually much weaker. Noise is also often uniformly distributed among all frequencies. (See the definition of white- and pink noise.)

One easy way to remove noise with the FFT method is to remove all frequencies that are weaker than a predefined value, and leave the stronger ones intact. (This is called thresholding.)

A problem with thresholding is that, even though noise exists across the whole spectrum, it is not equally loud everywhere. It is also usually not white, i.e. not uniform. Normally, the noise is concentrated in specific frequency bands, for example being stronger in the range of 1kHz to 10kHz. As we remove frequencies that are weaker than a predefined value, we will also remove parts of the music that have that volume as well. Usually, the treble of the sound is very weak and is therefore the first thing to go. Removing treble will remove the "distinctness" of the sound.

One way to deal with nonuniform noise is to analyze it first. To do that you must be sure that it is *only* noise that you analyze. When you know that noise is stronger (for example) in the mid-treble, you remove more there. Simple.

Many computer applications today can remove noise with FFT's (with better or worse results because all alterations will inevitably produce artefacts) Such programs work in two stages; first asking you to analyze a noise-only section, then removing that "image" of the noise from the rest of the material. The second step is very time consuming. The word FFT means Fast Fourier Transform; simply imagine the time required to do an *ordinary* Fourier Transform (DFT, Discrete Fourier Transform).

4.1.2 Hiding noise with filters

With equalizers and other filters you can often mask a *little* bit of the noise. But if you are dealing with audio material that was not prepared for this, you cannot automatically assume great success.

The ear is very sensitive to frequencies in the region of 1-5 kHz. When you hear noise, it is mainly in that region. If you apply a filter to remove frequencies in the 1-5kHz range, much of the noise will disappear. But so will the music, so take care. "Preparing" sound can be done by making frequencies stronger *before* storing, so you can mute them again during playback. This is how the most common noise reduction systems work on tape recorders. They assume two steps, one before recording and one after. But restoring sound assumes we haven't recorded the material ourselves.

4.1.3 Hiding noise by dynamically changing the volume

By altering the volume of a sound, you also alter the volume of the noise, obviously. One trick can be to compress the sound during recording and expand it during playback. But this assumes that we have recorded the sound ourselves. So what we *can* do is to lower the volume of the music in supposedly silent sections (between songs for example).

4.2 Removing clicks and pops ("impulses")

This section will not discuss the inner workings of the algorithms in use, but will guide you through the pitfalls of the methodology.

Pops and clicks are features of engraved music, like Edison's cylinders and vinyl records. As the material cracks and degrades, you can clearly hear the artefacts when the pickup of the turntable runs over the surface.

The nature of impulses is: "Extremely loud noise during an extremely short moments." The sum of this is a total artefact that is just as irritating as ordinary noise.

Unlike noise reduction, there is no "common" method of removing impulses in recordings. One way of doing this is to use "Kalman filtering", - an efficient way of finding unusual things in predictable environments (like impulses in music). But there are other methods that do not use Kalman filtering that work just as well.

One has to be aware that removing impulses from music implies two things: Finding the impulses, and then removing them. The problem with most algorithms is that sometimes even good music behaves as though it has pops and clicks in which case we would identify music as noise, and that is not good! This would occur in the case of a snare drum, for instance, which has a very strong and quick attack. The sound is loud and short which is the definition of an impulse.

When using most algorithms, we must be very sensitive when searching for pops & clicks, otherwise we will remove the attacks and transients of most sound.

When an impulse has been found, the next step is to remove it. There are numerous ways of *removing* something like this. One way is to replace the short moment with another piece of sound with the same characteristics. Other ways are to apply FFT noise reduction-techniques as described in the previous section. Interpolation is also a widely used technique and is a way of calculating and inserting missing information.

4.3 Frequency space editing

As described in the section about FFT noise reduction, the term "Frequency space editing" relates to the frequencies composing the sound. A "Space" relates to *how* we are seeing something. In the frequency space we see the frequencies of the sound whereas in the time space, we see the signal changing in time.

Now imagine that instead of *one* 1/10th of a second segment of the sound, we have hundreds, or thousands. We transform each of these into the frequency space, but because we have hundreds of them, we can also see how the frequencies change in time. This would be something like a 'Time-frequency space'. If we have a way of acting upon all of these short segments at the same time, we have something we can call a "Frequency space editor" or a "sonographic editor".

A frequency space editor gives us great control over the sound, *both* in time and frequency. All the information we have available can be arranged in a picture in which the horizontal view shows the time and the vertical view shows the frequencies. (These kind of pictures are very beautiful.) The strength of such an editor is that one can implement *secondary* functions, in the same way that a picture editor does.

Because the sound can be presented like a picture, one can use functions that are common to picture editing programs, for example pens and brushes.

What these pens & brushes actually *do* is up to the program to decide, but examples could be noise reduction and filtering, etc.

There are very few computer programs that implement frequency space editing.

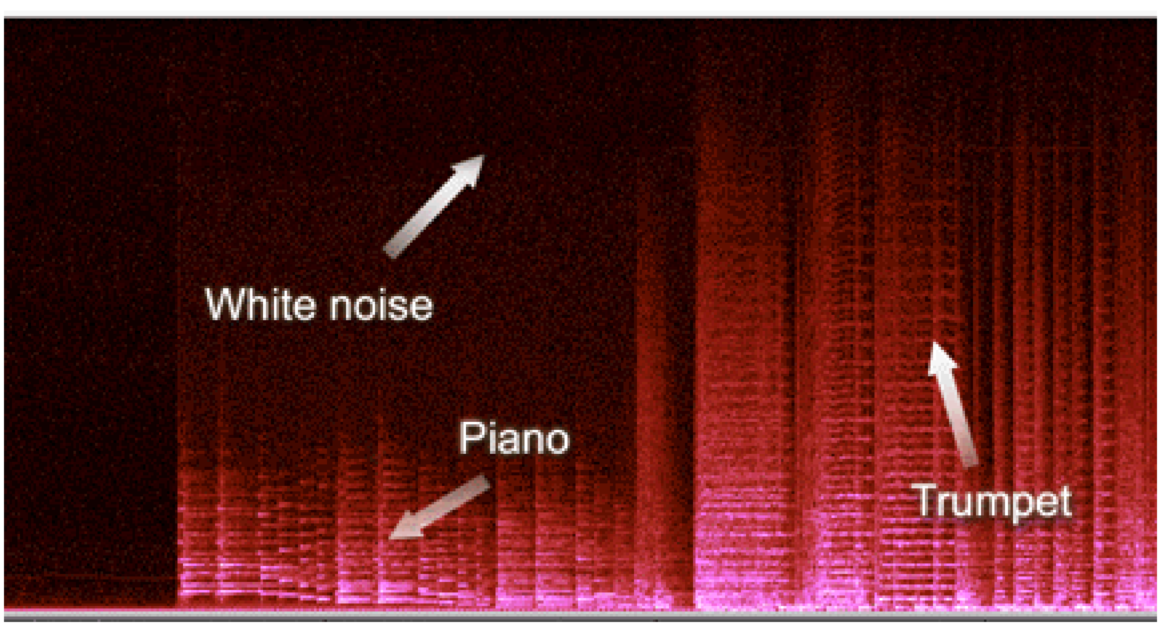

4.4 Sample by sample editing

Most audio tools can edit the waveform samples one by one. This is useful when you have to deal with very short periods of time. Sample by sample editing can be used with success when removing pops and clicks (short time noise), or when the music has short intervals of dropout and you can guess the form of the missing waveform. This way of editing the waveform gives you absolute control over the time, but no control at all over the frequencies. (See Figure 4.3, “Frequency space editing,” on page 7)

4.5 Reverbs

Reverbs mean echoes. Lots of echoes. A reverb tries, in some sense, to simulate the properties of a room, which in practice reflects thousands of tiny echoes from the walls. When restoring sound, one should be very careful in doing this, and remember that the goal of the operation is to make the music sound like the intention of the creator. Using lots of reverb “just because it will sound nice” will counteract this intention. Please remember that if *you* were not the one recording the music in the first place, the listener won’t necessarily want to be subject to your private ideas about what the music should sound like. If you are the only one who will listen to the restoration in the end, please remember that the original creator probably had something in mind that will take you a long time to realize, so why not stick with the original idea.

Use only such reverb that closely matches the reverb of the recording. This should include the size of the room, amount of reverb, timbre, and so on. You will realize that even a little tiny bit of reverb can do wonders for a bad recording, but too much will destroy it.

How do you find the ”right” reverb? You should try to find out what room size the creator intended you to experience. Step 2 is then to make the timbre of the reverb sound like the recording. Find out how much treble you can hear, then how the mid treble sounds. You also may want to add some extra bass to the reverb as this adds richness without affecting the overall impression.

When you are *searching* for the right reverb, remember to preview the result with a lot of reverb. When you think you have found the right combination, reduce the amount to what you feel is just about right - then reduce it some more! Remember that you are adding reverb on top of an identical room ambience. You shouldn’t overdo it - and you don’t have to.

4.5.1 Stereo broadening.

Why is everything we hear in stereo? Because we have two ears. No kidding. But, if you hear a singer or an instrument play alone in front of you, there is only one point of origin (the instrument). Even so, one can hear the “volume of the room” in the sound, in both ears. This is because of the reflections of the room you are in. The sound does not only go straight from the instrument to your ears, it also bounces off the walls and enters in your ear by indirect routes. Indeed, sound may bounce off the walls 100 times or more before you hear it. The total amount of reflection will build up into the "room ambience".

Recordings that give little impression of being in a room, *may* also lack a general feeling of distinctness in the mix. If you close your eyes and listen to a complex recording that is in mono, you will understand that this couldn't possible be real life, -the natural feeling of *size* is missing.

You should carefully consider whether the goal of your restoration is to create a *"new and improved"* version of the old recording, that "now is in blazing stereo". Perhaps you are best leaving it where it is? Or perhaps only adding a *little bit* of ambience; of course, not so much that you get dizzy. A good tip is that you add no more room ambience that you actually *can* hear when you stand a couple of metres away from the speakers.

How do you do stereo broadening? Probably the best way of doing this is to simply apply a reverb. A reverb often *statistically* mimics the properties of the walls. There are also special "room simulators" that simulate the actual path of the sound bouncing from wall to wall. This may not be quick but it is always very accurate. A reverb does a more "simplified simulation" that is quick but only "appears to be accurate".

When finding the right reverb, first find the size of the room you wish to imitate. Then try to understand the feeling of the room. Different walls absorb different kinds of sound. High treble sounds will be absorbed by any wall. If you have a lot of things (people and furniture) in the room, much treble will disappear, but bass and low frequencies are likely to be retained. The next step is to find what kind of filter you can apply to the reverb. (Apply the filter only to the reverb, - not the dry mix.) This filter would mimic the absorbance of the walls.

And again, don't use too much! Frequently compare your alterations with the original.

5.0 Combining two or more effects

There are lots of things you can do with any desktop computer program today. The possibilities may not be *endless*, but there are very many, and even if you know how to use all effects in question, things become complicated when starting to combine them. The most common functions that can be useful are:

Noise reduction
Pop & click reduction
Filters
Reverbs
Volume normalizing
Compression

But in what *order* do you apply these!? This is a very important question. If for example you choose to apply *all* the functions above, it <u>could</u> be in the following order:

Pop & click reduction ->
noise reduction ->
compression ->
filter ->
reverb ->
volume normalization

Pop & click reduction should be done before noise reduction, because the noise reduction algorithms work with segments of 1/10th of a second, and pop reduction works with sections of 1/1000 of a second. When undertaking noise reduction, you automatically affect sections that are far longer then the time of a 'pop'. Deal with the small details first and the longer details later.

Pops also interfere when you analyze a "noise-only" section described in Figure 4.1.1, "The principle of removing noise with FFT's," on page 5.

Noise reduction should be done before compression. When the noise-only section is first analyzed (Figure 4.1.1, "The principle of removing noise with FFT's," on page 5) it is assumed that the volume of the *noise of the recording is held on a roughly constant level*. When you compress something, a loud sound will lower the volume of that part, and therefor not keep the noise level constant. You will remove some of that loud sound instead. After compression, the noise will be weaker in parts of the song that contain loud sounds.

Compression should be applied before filtering. A compressor analyzes the volume of the different parts of the song, and if you apply a filter to this song you, will affect different instruments differently. Imagine a kick drum and a hihat with the same volume. If you apply a filter that makes the bass greater, you will raise the volume of the kick drum but not that of the hihat. This would compress the two instruments differently.

You should apply the reverb after compression, because if you compress a song that already has reverb in it, you will also compress the reverb and that may sound very unnatural. It would sound like the room ambience was constantly changing in time. (But that may be an interesting idea if you were to mix your *own* song?)

Volume normalization should be done as the last step. This is because whatever effect you use, the volume *will* be affected, therefore destroying any normalization you already have done. It is important to have a volume that is as great as possible throughout the process however. This reduces truncation errors, as a waveform has a limited number of bits, implying that it would <u>not</u> be wrong to normalize *every now and then*.

But remember that there are no dogmatic rules! Some situations need special methods, but the sequence described above can be used most of the time.

6.0 How to store the result

One way or the other, you need to store the restored result for later use or listening. You may not wanna put it on a cassette again, but you probably already guessed *that*... In fact, all analog media should be avoided, but luckily there are many ways of storing it digitally. You should keep some simple things in mind for storing sound.

The ear can hear frequencies up to 20.000 Hz. Discussions exist whether or not the human ear can perceive even *higher* frequencies than that, but in listening tests, such frequencies cannot be detected by most people. Please note that the new DVD-Audio standard permits frequencies up to 96kHz (Figure 6.2, "DVD-ROM," on page 14). Ear performance also degrades as time goes by. Children have better perception than their parents, and elderly people may have trouble perceiving frequencies over 10kHz.

When sampling a signal digitally, there must be *at least two samples to represent a cyclic event (like a vibration)*. Therefor, there must be twice as many samples per second as the frequency of the vibration. This implies that the sampling frequency for a sound containing a 20kHz signal must have at least 40.000 samples per second.

The ear has a wide dynamic range. This mean that you can hear *very wea k sounds and very loud* ones. However, the ear has problems hearing a weak sound *at the same time* as a loud one. This is because its sensitivity adapts to the environment which is changing constantly. Volumes are preferably measured in decibels, dB. A decent dynamic range should be well over 80dB. This implies that you can play loud sounds (0dB) and very weak sounds (-80dB). The dB scale is relative. You always need to compare a volume with something else. (If the loudest sound in some scale is 0dB, any sound that is weaker than that is negative. A -80dB sound is 80 dB weaker that the loudest, the 0dB sound.)

Cassette tapes without noise reduction have a dynamic range of about 60dB.

6.1 CD-R / CD-RW

Lately, the price of recording on CD media has dropped dramatically. It is now possible for most people to buy a CD burner themselves without selling the car or getting a loan from the bank. The quality of CDs is also good enough for every professional listener, and is always better than tapes and vinyl. CD's created on desktop computers can be played in common home stereo systems. You need a special CD burning program for storing data on CD's, but such programs are usually included with the machine.

An audio CD has only one format: 44.100 Hz, 16 bits and 2 channels. It is not possible to put another kind of waveform on a CD, unless it is stored as a .wav .aif .mp3 etc file, but then you lose the ability to play the waveforms on home stereo equipment.

CDs have a dynamic range of 96dB.

6.2 DVD-ROM

Perhaps the successor of the “ordinary” CD. The DVD is a compact disk that has a far greater density that the CDs of today. The enhanced ability to store information makes it an ideal medium for storing music with *better quality*. “Better” in this case means higher bit widths (up to 24) and higher sampling rates (up to 192kHz). 1998, the new standard for storing musical audio on the DVD was defined.

The ordinary Audio-CD standard allows only one format: **44.1kHz, 16 bits, 2 channels**.

The new standard for DVD permits: **44.1 / 88.2 / 176.4 / 48 / 96 / 192kHz sample rate, 16 / 20 / 24 bits, and up to 6 channels.**

The large amount of bits gives DVD-Audio a dynamic range of up to 144dB.

6.3 DAT

DAT is a little bit better than CD as it supports a sample frequency of 48kHz as opposed to 44.1kHz in CDs. DAT is a tape format. You usually don’t play DAT tapes in home stereo systems because the players are not cheap, but in professional studios the availability is greater.

DAT has a dynamic range of 96dB.

6.4 MiniDisc

The MiniDisc (MD) is a format that is becoming an alternative to the CD. MD uses *lossy compression* to store the digital data. “Lossy” means that you loose information in the sound. It is better than cassette tapes but not as good as CD or DAT. The compression scheme is somewhat like the common .mp3 file format which can reduce the filesize significantly.

The MiniDisc uses a compression format known as ATRAC which compresses the audio in the frequency space at a ratio of 1:5. A MD can store 140Mb of data (Assuming a 2-1/2 inch 74 minute disc.)

MD has a dynamic range that is comparable to 96dB.

7.0 Glossary

dB (decibel). The most common scale for volumes. The scale is relative (it is logarithmic). A scale that is relative must have a reference. When talking about the volume of a sound, you must always compare this volume with another volume, otherwise your explanation would be useless.

Usually the reference volume is 0dB. Anything louder than this would be on the positive scale, and anything weaker than this would be negative. A sound that is “half as loud” as this reference volume is -3dB. A volume that is twice as loud is +3dB.

Fourier transform. A mathematical algorithm for analyzing which frequencies are contained in a signal.

Impulse. A very short burst of noise. The word “impulse” actually means “the shortest disturbance possible”. In sampled audio, an impulse would be only one sample long. A pop or a click in vinyl audio is considered to be an “impulse” by the ear. Pops and clicks are short time disturbances on the surface of the material.

Frequency space. Information is arranged in a way where we only can see the frequencies of a signal.

Time space. Information is arranged so that we can see how the signal changes over time.

White noise. Randomness of the signal uniformly distributed over the whole spectral range.

Colored noise. As opposed to white noise, it is not uniformly distributed. In practice noise is mostly colored.

Sample frequency. When a continuos signal is transferred digitally, it is sampled this number of times every second. The more samples per second you have, the more detail there can be.

Sample bits. When a continuous signal is transferred digitally, every sample has this resolution. The more bits you have, the greater the dynamic range can be. The fewer bits you have, the more “noisy” the signal will appear to be.

8.0 References

[1] Friedrich Kittler, “Gramophone, Film, Typewriter” (Stanford: Stanford University Press, 1997)

[2] http://www.videodiscovery.com/vdyweb/dvd/dvdfaq.html

[3] Eckel, Gerhard: “A Signal Editor for the Ircam Musical Workstation”, ICMC: International Computer Music Conference, Septembre 1990.

[4] Déchelle, François, Maurizio DeCecco, Enzo Maggi et Norbert Schnell: “New DSP applications on FTS”, ICMC: International Computer Music Conference, 1996.

[5] Koch, S., H. Kaufman, J. Biemond: “Restoration of spatially varying images using multiple model extended Kalman filters”. Proceedings of the 1993 IEEE conference on decision and control, December 15-17, 1993. Page 1416.

[6] http://archervalerie.com/audio.html

9.0 Credits

- Thanks Eddie Alestedt for technical advises.
- Thanks Geoff Pickup and Sam Sundquist for linguistic conformance.
- Thanks “You” for reporting any new topics and / or errors in this text. Please send a mail to mike@physto.se