Digital audio page

Index

General information

Digital audio is the most commonly used method to represent sound inside a computer, many audio processing device and modern audio storage devices (like CD, MD, DVD).

Digital audio technology is a method of representing audio signal using binary numbers. An analog audio signal is converted to digital by the use of an analog-to-digital (A/D) converter chip by taking samples of the signal at a fixed time interval (sampling frequency). Binary number are assigned to these samples. This process is called sampling. This means that the audio data is stored as a sequence of samples taken from the audio signal using constant time intervals. A sample represents volume of the signal at the moment when it was measured.

This digital stream of data is then recorded onto storage media (magnetic tape, optical disk, hard disk or computer memory) or transmission path (telecommunication network, Internet, digital satellite, digital TV transmission). Upon playback, a digital-to-analog (D/A) converter chip reads the binary data and reconstructs the original analog signal. This process virtually eliminates generation loss as every digital-to-digital copy is theoretically an exact duplicate of the original.

In uncompressed digital audio each sample require one or more bytes of storage. Number of bytes required depends on number of channels (mono, stereo) and sample format (8 or 16 bits, mu-Law, etc.). The length of this interval determines the sampling rate. Normally used sampling rates are between 8 kHz (telephone quality) and 48 kHz (DAT tapes).

In digital audio, bit depth (like 16 bits, 24 bits etc.) describes the potential accuracy of a particular piece of hardware or software that processes audio data. In general, the more bits that are available, the more accurate the resulting output from the data being processed. Bit depth is frequently encountered in specifications for analog-to-digital converters (ADCs) and digital-to-analog converters (DACs), when reading about software plug-in, and when recording audio using a professional medium such as a digital audio workstation or a Digital Audio Tape machine. Bit depth is the number of bits you have in which to describe something. Each additional bit in a binary number doubles the number of possibilities. When you have a 16-bit system, there are 65,536 possible levels of audio signal. When you have a 24-bit process or piece of 24-bit hardware, there are 16,777,216 available levels of audio.

16 bits is a typical resolution (bit depth) of many digital audio systems nowadays. The audio on CD disk is stored at 16 bit resolution. Modern PC soundcards operate at 16 bits resolution. In some professional applications, there is need for more bits. In some digital professional audio systems you can encounter 20-bit or 24-but resolutions. Some devices can process soemtimes the data more accurately (for example at 32-bit resolution), but the information is generally transferred and conmmunicated at maximum of 24 bit resolution. 24-bits is practically enough for any real life audio application (pretty much on the performance limits of bers A/D converter) and it usually the maximum resolution supported by digital audio interfaces (AES/EBU in professional audio world, S/PDIF in consumer products).

There is universal definition of resolution: Resolution is the ratio between the largest representable signal and the basic ambiguity in representation, whether that ambiguity is due to random, unpredictable noise or quantization. This amobiguity exists simply because at low enough signal levels, you do not know whether the signal change is due to the signal or to the noise in the system (and, by "you" we mean "your ear" or "your test instrument" equivalently).

It should be further noted that properly implemented digital system do not sample the music in discrete "steps" but rather have the same unpredictable random ambiguity as conventional analog systems.

In some instances you might have heard of 1-bit audio systems. You might wonder how 1-bit can present audio signals, because 1 bit can present only onle signal level. The magic is in that those 1-bit samples are generated using a special sampling process which produces a series of 1-bit samples that represent the audio signal accurately after played back to playback device. The simplest 1-bit system you can think of is delta modulation (DM), where each bit presents should the output signal voltage increased or dicreased a bit. This was used in some telephone systems. Sigma-delta modulation (SDM) was developed in 1960s to overcome the limitations of delta modulation. Sigma-delta systems quantize the delta (difference) between the current signal and the sigma (sum) of the previous difference. SDM is also known as pulse density modulation (PDM). The maximum quantizer range of SDM is determined by the maximum signal amplitude and is not dependent on signal spectrum.

For the 1-bit samples to represent the audio signal accurately, this kind of 1-bit samples needs to be taken at very high rate (very much higher than samples with higher resolution are generally taken in audio applications). To convert a maximum amplitude 16-bit word, a 1-bit modulator (delta modulator) would have to perform 216 toggles per conversion period; with a sampling frequency of 44.1 kHz, this would demand a toggle rate to approximately 2.9 GHz, an impossibility with today's technology. For comparision with SDM with an audio band 22,1 kHz and 64 times oversampling, the internal sampling frequency rises to 2.8224 MHz, thus quantization noise is spread from dc to 1.4112 MHz. However, sigma-delta modulation adds noise-shaping benefits.

Nowadays SDM (sigma-delta-modulation) is used in the DA converters of many CD players. The 16-bit 44.1 kHz data is requantized and oversampled with SDM encoder and converted to analog line level signal with SDM decoder. 1-bit converter is in some cases easier and cheaper to implement and the most important benefit is its better tolerance for small variances in component values compared to the high demands of the conventional weighted-resistor and R-2R ladder high-bit DACs.

The physical device that converts analogue audio to digital audio is called ADC (Analog to Digital Converter) and the device which converts digital audio to analogue audio is DAC (Digital to Analog Converter).

Sampling parameters affect quality of sound which can be reproduced from the recorded signal. The most fundamental parameter is sampling rate which limits the highest frequency than can be stored. It is well known (Nyquist's Sampling Theorem) that the highest frequency that can be stored in sampled signal is at most 1/2 of the sampling frequency. Sample encoding limits dynamic range of recorded signal (difference between the faintest and the loudest signal that can be recorded). In theory the maximum dynamic range of signal is number_of_bits * 6 dB . This means that 8 bits sampling resolution gives dynamic range of 48 dB and 16 bit resolution gives 96 dB.

- 1-bit A/D and D/A Converters
- A Graphical Explanation involving Sampling
- Audio receivers tune in to binary broadcasts - Digital audio is more than the music played on PCs, PDAs, or solid-state Walkmans strapped to joggers' hips. Sonic-supplied Internet appliances can transport the tunes to every room in the house.
- Basics of Digital Recording - digtal audio introduction
- Consequences of Nyquist Theorem for Acoustic Signals Stored in Digital Format - published in proceedings from Acoustic Week in Canada 1991 - CAA Conference, Edmonton, Alberta, Canada, October 7 - 10, 1991, from Digital Recordings
- Digital audio gets an audition part one: lossless compression - Theory's fine, but reality's better. In this two-part series a PC, stopwatch, oscilloscope, and spectrum analyzer are used to get answer your audio-compression questions.
- Digital audio gets an audition. Part two: lossy compression - Lossless compression can shrink file sizes down a fair amount, but for serious weight loss, you need to permanently discard some of the data. Find out how the lossy codecs work and whether you can hear the differences between them.
- Digital audio lectures
- Digital Audio: Theory and Reality
- Destination distortion - High-resolution audio strides toward an unclear future. Advanced audio formats' success won't necessarily come from their high-resolution features. Advantages of the latest audio formats are apparent for silicon, software, and technology providers; system manufacturers; and media producers. But benefits for consumers are less obvious. Couple that with the formats' diversity and incompatibility, and they all may end up as niches.
- FAQs (Frequently asked questions with answers) on Digital Signal Processing
- Introduction to Digital Recording Techniques - paper was published in Canadian Acoustical Association Journal, 1989, from Digital Recordings
- Level Practices in Digital Audio
- Principles of digital audio
- Security scheme doesn't hold water(marking) - Audio watermarking inserts media identification data within the bit stream, using techniques that its advocates claim are inaudible, but they might not be that reliable and unaudible as claimed
- Signal Processing information page
- S/PDIF IEC958 (named IEC60958 at 1998)
- The Video Engineer's Guide to Digital Audio - full book on-line
- White Paper: Digital Audio Distribution Systems - paper examines the suitability of several computer networking technologies to the task of carrying real-time digital audio

Analogue to digital conversion technology

An analog-to-digital converter (also known as an ADC or an A/D converter) is an electronic circuit that measures a real-world signal (analogue audio) and converts it to a digital representation of the signal (digital audio). A/D-converter compares the analog input voltage to a known reference voltage and then produces a digital representation of this analog input. The output of an ADC is a digital binary code. By its nature, an ADC introduces a quantization error. This is simply the information that is lost, because for a continuous analog signal there are an infinite number of voltages but only a finite number of ADC digital codes. By increasing the resolution of the ADC, the number of discrete steps is increased, which reduces quantization errors.

A typical analog-to-digital process consists of two stages: discrete-time sampling followed by amplitude quantization. The sampling takes samples of the incoming signal voltage at teh predetermined rate (sample rate). Amplitude quantization convert those voltage samples to numeric value with limited resolution (determined by the number of bits used). Each of these two steps has requirements that must be met. Discrete-time sampling demands that the signal must be band limited to the appropriate Nyquist band before sampling to prevent aliasing. The signal must also be properly dithered before quantization to prevent the introduction of signal-correlated artifacts (usually described as "grain and grunge"). If you do not add dithering, there is some nasty stuff that happens to an undithered signal as it approaches the bottom of the bit bucket, and the rapidly mouonting severe distortion due to that. One of the big problems with quantisation distortion is that it isn't so many dBs below signal level, but is pretty much at a constant level regardless of the level of the programme material.

To avoid those problems, dither is used practically always in some form. The application of the correct amount of dither will completely eliminate all signal-correlated artifacts. Properly-dithered digital contains no nonlinear distortion of the signal being quantized. With any normal form of dither (technically, with any wordlength reduction to quantized PCM) you end up with an entirely decorrelated noise floor, which mimicks the behavior of ordinary noise. Dithered noise floor is not true Gaussian noise, indeed there are continuing discussions among pro-audio guys as to which form of dither is best for total inaudibility.

Dither can can applies in many ways and sometimes you might not need any spearate ditner noise source at all. For example when music is being played full-out there is generally enough non-harmonic, noise-like content to do an effective job of dithering on its own without need for a specific dither source. In a 24-bit recording system the intrinsic analogue noise floor will almost certainly be greater than 1 bit in amplitude, so a specific dither signal would be unnecessary. Of course it would be needed during 16-bit truncation.

In a dithered system the the definition of resolution is is not always very lear. With an un-dithered quantised system, resolution is simple - the size of the step. There can be no argument about that. When you introduce dither, things change. You no longer have that step, and resolution turns into a rather moot subject which is more to do with how small a signal you can discern in the presence of noise. This is no longer simple. By choosing appropriate bandwidths you can discern as small a signal as you like. So resolution no longer has any easily defined meaning. It is replaced by signal to noise ratio.

This was the conventional architecture. Nowadays different A/D conversion architectures tend to blur the implementational division between the two steps.

A/D-conversion in audio systems is a sampling process. All the sampling processes are limited by Nyquist limit. The Nyquist limit is defined as half of the sampling frequency. The Nyquist limit sets the highest frequency that the system can sample without frequency aliasing. In a sampled data system, when the input signal of interest is sampled at a rate slower than the Nyquist limit (fIN > 0.5fSAMPLE), the signal is effectively "folded back" into the Nyquist band, thus appearing to be at a lower frequency than it actually is. This unwanted signal is indistinguishable from other signals in the desired frequency band (fSAMPLE/2). Usually the signals are prefiltered before they enter the A/D-converter to avoid too high frequency signal components which can cause this kind of unwanted signals.

Typical good quality modern digital audio systems use at least 16 bit resolution (typically 16-24 bits) and sample rate of 44.1 kHz (CD sample rate), 48 kHz (DAT sample rate) or higher (96 kHz, 192 kHz etc.).

Sometimes you might wonder what is the dynamic range of a digital audio system. Hard to use accurate definition is that the total amount of uniquely representable information in the system is, indeed, proportional to the dynamic range in amplitude multiplied by the bandwidth.

But, if you accept the definition of dynamic range as " the ratio between the largest possible undistorted signal and the smallest unambiguous change in signal," then it's easy.

Actually, for run-of-the-mill PCM data, dynamic range only depends on the number of bits. To do the bare calculation, the dynamic range in dBs is 20 log (2^n) where n is the number of bits. But as I say, that is far from the whole story. The generally accepted rule of thumb is 6 dB per bit. So a 10-bit system has 60 dB, a 16-bit system 96 dB, etc. Please note that this value is for a single sample. Through common averaging techniques, it is possible, just like in non-digital systems, to encoded and detect signals that are well below the noice floor. The 96 dB figure is the worst-case, non-dithered instantaneous dynamic range of a 16-bit system.

When talking about A/D-conversion, things get more complicated. 0dB is easy - it is the biggest sine wave you can fit into the digital space. You are driving the ADC end-to-end. The other end is more problematic. You find it by measuring the noise level. Unfortunately this depends on the measuring bandwidth.

Every time you halve the measuring bandwidth, you reduce the noise by 3dB. Of course you can simply include the entire bandwidth, but in the case of audio that isn't really fair, because we don't hear that way. Particularly at low levels, we hear predominantly the middle range frequencies, and the extreme lows and highs disappear. At threshold levels even "A" weighting gives unfairly pessimistic levels of noise.

Digital to analogue conversion technology

A digital-to-analog converter (also known as a DAC or a D/A converter) is an electronic circuit that converts a digital representation of a quantity into a discrete analog value. The input to the DAC is typically a digital binary code, and this code, along with a known reference voltage, results in a voltage or current at the DAC output. The word "discrete" is very important to understand, because a DAC cannot provide a continuous time output signal; rather, it provides analog "steps." The steps can be lowpass-filtered to obtain a continuous signal.

In D/A conversion process the output of D/A converter is fed through a filter which will remove the image-frequency information (signal higher than 1/2 of sampling frequency) from the output signal. This image-frequency information can distort the output signal. Two methods exist for removing unwanted image signals from the DAC output to prevent alising in a following ADC. First approach is to use a high-performance lowpass filter (data -> DAC -> high-order lowpass filter). For low pass filtering usually a sixth-order lowpass filter is enough. The second methos is to use digital-interpolation filters and a simple analogue filter (data -> oversampling digital-interpolation filter -> DAC -> low-order lowpass filter).

- 1-bit converters
- A Suggested Explanation For (Some Of The) Audible Differences Between High Sample Rate and Conventional Sample Rate Audio Material - AES paper in pdf format
- Effects in High Sample Rate Audio Material - document in pdf format
- On Jitter, the S/PDIF Standard, and Audio DACs - a discussion of how the implementation of the S/PDIF standard used in consumer digital audio can effect the fidelity of the recovered analog signal
- Bit Wars Listening Tests - Can You Hear the Difference between 16- and 24-bit recording ?
- Everything you always wanted to know about jitter
- Resolution, Bits, SNR and Linearity - document in pdf format
- Timing Errors and Jitter - document in pdf format
- Understanding the origin and function of audio A/D converters
- Why Jitter matters in high resolution digital audio systems - white paper from Axon in pdf format

Digital audio formats

Format comparisions

- A Comparison Between Audio Formats - DAT, LD, AC-3, radio, TV, LP, phone

CD

Compact Disc is a stereo digital audio system based upon a 12cm single sided disc made from a polycarbonate material with an internal reflective layer of aluminium or occasionally gold. The information is read by sensing the presence or absense of reflected light from a tightly focused Laser beam pointing at the pits and bumps in the reflective surface within the disc. The digital audio signal is sampled to 16 bits per channel at a rate of 44.1KHz. The major developers were Philips electronics of Holland and Sony of Japan. In other words the audio format in audio CD is two-channel, signed 16 bit two's complement integer samples, 44.1 kSamples/second, linear PCM. This is defined in an industry-standard publication known as the "Red Book", the first in a series of "rainbow" books that define other formats (DVD, karoke, CD-R, etc. etc.).

According to Red Book specifications, a standard CD is 120 mm (4.75 inches) in diameter and 1.2 mm (0.05 inches) thick and is composed of a polycarbonate plastic substrate (underlayer - this is the main body of the disc), one or more thin reflective metal (usually aluminum) layers, and a lacquer coating. CDs are divided into a lead-in area, which contains the table of contents (TOC), a program area, which contains the audio data, and a lead-out area, which contains no data. An audio CD can hold up to 74 minutes of recorded sound, and up to 99 separate tracks. Data on a CD-DA (audio CD format) is organized into sectors (the smallest possible separately addressable block) of information. The audio information is stored in frames of 1/75 second length. 44,100 16-bit samples per second are stored, and there are two channels (left and right). This gives a sector size of 2,352 bytes per frame, which is the total size of a physical block on a CD. CD data is organized into frames (consisting of 24 bytes of user data, plus synchronization, error correction, and control and display bits) which are intricately interleaved so that damage to the disc will not destroy any single frame, but only small parts of many frames. In most cases error correction (reed-salomon) or error-concealment techniques can still solve those errors.

Some technical data on CD system:

```
Playing time:          74 minutes, 33 seconds maximum (can be pushed to 78-80 minutes)
Number of channels:    2 channels
Quantization:          16-bit linear
Sampling frequency:    44.1 kHz
Data bit rate:         2.0338 Mb/sec


The actual audio signal data from this data is
44,100 samples/second * 16 bits/sample * 2 channels = 1,411,200 bits per second.


Here is some data on CD disc itself:


Disc diameter:         120 mm
Modulation system:     Eight-to-fourteen Modulation (EFM)
Channel bit rate:      4.3218 Mb/sec (data rate stored to disk)
Error correction code: Cross Interleave Reed-Solomon Code (25% redundancy)
Laser wavelength:      780 nm (infrared)
Lser diode power:      Typically 3-5 mW
Reading signal power:  Typical CD laser optics put out about .3-1 mW at the objective lens
Laser focus point:     1 to 3 mm from lens front surface
Disc material:         Any acceptable medium with a refraction index of 1.55
Pit length:            0.833 µm - 3.56 µm
Pit depth:             ~0.11 µm
Pit width:             ~0.5 µm
Rotation:              Counter-clockwise when viewed from readout surface
Track speed:           1.2-1.4 m/sec
Disc rotation speed:   from about 500 rpm, down to about 200 rpm
```

The disc has a diameter of 120 mm with a 15 mm centre hole. The track extends between the 46 mm and 117 mm radii. The track pitch is constant at 1.6 um, making 22188 tracks across the radius of the disc, for a total length of 5.7 km.

The compact disc player contains two main subsystems: the audio data processing system and the servo/control system. The servo, control, and display system orchestrate the mechanical operation of the player and include such items as the spindle motor, auto-tracking, lens focus, and the user interface. The audio data processing section covers all other player processes.

Many audio CD players can play besides normal CDs also audio recordings done to CD-R disks. But not all. Sometimes people ask why. The answer is the following: All CD-R's have lower reflectivity than "real" CD's, and many older drives made before CD-R's were widely available don't have enough laser output (especially as they get older and the laser wears out) or enough range in their laser AGC system to increase the laser output enough to read CD-R's. A few older drives that had plenty of spare laser output available have no trouble reading CD-R's, while most modern drives deliberately have an AGC system for the laser output with enough range to cope with a variety of media types. If your player can't play bacd CD-R's there is not much you can do about it, except to try different brand and colour CD-R's - the various different formulations have different amounts of reflectivity - so some are better than others in marginal drives. Typically the the Blue/Silver disks to work best (blue dye, silver substrate, made by mitsubishi) and he Green/Gold ones (Green dye, gold substrate) to cause more trouble, but different drives sometimes prefer different colour CD-R's. In some cases it has been reported that the speed that you write the disk at can influence how easy it is for older drives to read the disk. The popular wisdom on this is that writing the disk slower - even down as low as single speed - makes a disk that is more easily read by older drives. If you have a marginal CD player you want still use with CD-Rs, get a bunch of different brands and keep burning with different settings until you find one that works. There are not many regular CD players that have problems with CD-Rs. Best CD-R media can help reduce those rare problems.

- A fundamental Introduction to the Compact Disk Player
- An Exploded CD player as Teaching Aid
- CD player features - a summary of experiences and opinions concerning CD specs and the utility of specific player features
- CD-player mods - CD player modification ideas for propable better sound quality
- CDplayer's Project - information to make and upgrade a CDPlayer at home
- CD Subcode Formats
- Compact Disc Formats - explanation of different types of CD disks
- Compact Disc: The Inside Story - articles on the technical details of the Compact Disc audio recording system originally published as a series in the magazine MAC Audio News
- Differences between CD-R/CD-RW discs and standard CD
- Digital Domain - site is designed to help audio engineers and musicians make better Compact Discs and CD ROMs
- EE498 Project Reports - exploded CD player, 1-bit D/A, CD testing, karioke device
- How CDs Work
- Introduction to CD player - how CD player works
- CD Player Laser Diodes - This document give you introduction to lasers used in CD player.
- Red Book - The Red Book is the 1980 document that provides the specifications for the standard compact disc (CD) developed by Sony and Philips. According to legend, the document was in a binder with red covers, originating the tradition for subsequent adaptations of CD specifications to be referred to as variously colored books.
- Subcode on the CD - check also Q-Subcode Decoding and Display project
- The Compact Disc - good introduction to CD technology by Philips
- The Compact Disk - how does it work
- Transferring LPs to CDR: Some Advice
- Understanding the CD - How CDs Work

CD copy protection

The legal and illegal copying of CDs has increased during the last few years due to the wide availability of low cost CD-recorders and media. This has resulted in a fall in the demand for legitimate pressed CD albums in some countries, which has made music companies to react to this situation. The new discs now making their way into record stores in the United States and Europe contain countermeasures that prevent playback on computers and, in some unintended cases, normal CD players as well. CD copy protection is designed to work it should play in your home CD player but not in your CD ROM drive. The reason being that CD ROM drives enable you to easily copy the digital audio data and burn copies of it.

The controversial new anti-copying technology introduces minute errors to the CDs or changes the location of data on the discs to prevent them from being played back on computers. In theory, most consumer CD players (but not all) can correct the errors and decipher the structure, unlike the more finicky computer CD drives (which uusally try to read more details than normal audio CD players). All of the big major record labels are experimenting with ways to block consumers from "ripping" or transforming their CD songs into MP3 files and distributing them widely online though on-line file sharing services.

This "CD copy protection" has been a well discussed topic because some record companies have started to use it on some music CDs. Copyright protected cd's do not allow you to replicate them in a cd burner nor do they allow you to rip the audio tracks "digitally" (although can still be done through analog). Copyright protections systems currently in use have downsides: not all CD players (especially computer CD-ROMs and some car plyers) will play those CDs.

The current copy protection techniques generally either corrupt the CD table of contents or actual audio data. The protection techniques that modify table of contents, modif it in in such a way that readers that know about data CDROMs get a tummy ache and see only some trivial data files, and not the audio tracks. This means that the protected discs can't be played on CD players that use CDROM drives for transports, of which there are quite a few. One way to make a CD act like is is to add a CD-Extra file track to the end of the CD, but making that CD-Extra track somehow broken (computer tried to read this track, but fails with error which causes the reading of whole CD to fail in many systems). Another method does corrupt the data. The plan is to put some deliberately-corrupted blocks, corrupted enough that the ECC will fail, into the audio data. It's transparent to analog output because the corrupted samples are in places where the interpolation done by the player when it can't do the digital ECC will end up with the same answer (or very close to it) as was in the original samples prior to corruption. Most PC CD-ROMs can detect the errors caused by data corrupptio, but most CD-ROM drive controllers lack the ability to pass this information of where error back up to the PC or do the error-concealment algorithm in firmware, prior to passing the data to the PC. A lot of the better CD-ROM drives do the error concealment on "ripped" data, so they aren't bothered at all by these "protected" discs. Cheaper CD-ROM drives may not do the error concealment.

Anyway this copy protection might make the copying of audio CD harder, but it does not stop you from doing that if you really want to do this. If you can play back the audio CD through your HIFI system, then there is always a way to copy it some way (through digital or analogue means). When you play back a CD on your normal CD player, you wll always get the correct audio signal from the CD player analogue output and CD player digital audio output (no matter if that CD is protected or not). No CD copy protection can stop that (or it will stop normal playback with normal CD players as well).

There are industry views that this kind copy protection techniques are actually a violation of the "red book" standards, to which everyone who presses music CDs has contractually agreed to comply. According some megazine article Philips has been putting these companies on notice that they may not advertise or label such products as "Compact Discs", and cannot use the CD logo or claim CD compatibility, as these discs deliberately violate the Red Book standard and are thus a technical violation of the license agreements for the use of the logos, etc. Sony is also a co-licenser of the CD system and they are very much in favor of copy-protection. So that leaves us in a sort of twilight state of not knowing what will happen next.

At the very least, consumers have the right to insist that any such products be accurately identified (warning stickers, etc.) so that we can make an informed decision about buying them (or not!). The only power we consumers have is to bring back the defective CD that does not play properlu and ask for a refund or non-copyprotected replacement.

It seems that music companies are now trying to latch the door on the barn after the horse escaped. What happens with many of the copy-protected discs is that many users will have problems in playing them properly (lots of angry customers and returned CDs to shops) and that someone breaks the copy protection and distributes the tracks through a file-sharing service ? Once a track enters the domain of a file-sharing service, there's no stopping its distribution. People copy those music files at higher priority, simply because they were copy-protected (also those people who would have bought a CD but it did not work for them). This means that applying copy protection which causes problems to use punishes the users which are willing to buy the original CD. If the original CDs do not work reliably anymore, those users might be tempted to go to the route of copying the music instead (in this way they get their own CD copy which works for sure on many different players including their home CD, DVD and PC).

Protected audio CD's can be backed up using a variety of methods but usually an easy 1:1 backup is not possible. This is exactly what they want, making it harder to make a backup so that people will buy the original. It is hard to say if audio CD Protections have really helped in any way (the history of game CD protections has not been a good one). In some cases CD protection even can work against you as it limits the use of the CD. You can't play protected CDs on CD-ROM players and even some normal CD players can not handle the protection. So every CD-protection scheme faces a crucial obstacle: making CDs unrippable onto CD-ROMs also makes them unplayable on some CD players -- a feature guaranteed to anger customers.

The CD copy protection reduces the usability value of CDs. The user's don't get anymore what they used to get when they bought it. They can't play it on certain devices and maybe can't make legal copies to other media devices they want to use (laptop hard disk, MD player etc.). This reduced usability combined with the same price can easily reduce the sales of the CDs (reduced usability with same price generally reduces sales in free market). For example many users who buy a new protected CDs are quite pissed off when they find out that their high quality car CD player refuses to play that CD. And they can't either listen it at work with their computer CD-ROM.

Combined this with the fact that those copy protections do not stop pirate CD makers on making pirate copies of copy protected CDs, those copy protections do not stop CD piratism. They can even lead some people to prefer pirated CDs, because non-copy protected pirate CDs work better for them than the real ones. This applies to both bought pirate CDs and ones burned to CD-R from MP3 files downloaded from net. So those user who have problems with CD protection can easily start to think "Why should I pay high price of CD to a large media company which I don't like and which just supplied expensive flawed products which I can't use ?" and look for altenative solutions.

Other concerns are that CD copy protections can break the chain necessary to preserve creative works. It will make them readable for a limited period of time and not be able to be moved ahead as media deteriorates or technologies change.

From CD end user point of view, those CD copy protection schemes are a very bad idea which should not be used.

- CD Audio Copy Protection - Copy protection systems are now available to reduce the threat of illegal copying. What they are and how they work ?
- CD creator burns copy-protection efforts - As major record labels unveil a new breed of CDs designed to prevent Napster-style piracy, Dutch consumer electronics maker Philips, the co-creator of the CD, is refusing to play along.
- CD Protections - The last few years game developers more and more are protecting their games with commercial CD Protections. Now the music industry has hopped the bandwagon and now is selling protected audio CD's. These CD's can be backed up using a variety of methods but usually an easy 1:1 backup is not possible. This page contains a list of all known CD Protections which are currently on the market, how they can be detected and how to by-pass many of them.
- Corrupt audio discs, aka "Copy-Protected CDs" - Summary of Sony's new set of non-"Red Book" audio CD releases
- Corrupt audio discs, aka "Copy-Protected CDs" - Here is a list of known the bad CDs available in UK.
- Is Universal Music's Copy Protection a Joke? - At end of 2001 Universal Music introduced copy protection for its line of audio CDs using Midbar Tech's Cactus Data Shield (CDS).
- Kopierschutz mit Filzstift knacken - Copy protection with felt-tip pen crack is discussed here. The text is in german, try automatic translation to English.
- Labels loosening up on CD copy locks - Fearful of consumer backlash, major record labels in the United States have slowed controversial plans for making CDs more difficult to copy.
- Lawmaker: Is CD copy-protection illegal? - Record companies' efforts to protect CDs against digital copying are beginning to draw scrutiny from lawmakers concerned that the plans might violate the law in USA.
- Marker pens, sticky tape crack music CD protection - Music disc copyright protection schemes such a Cactus Data Shield 100/200 and KeyAudio can be circumvented using tools as basic as marker pens and electrical tape.
- Philips moves to put 'poison' label on protected audio CDs - Netherlands giant Philips Electronics has lobbed a grenade into the audio copy protection arena by insisting that that CDs including anti-copying technology should bear what is effectively a plague warning. They should in Philips' view clearly inform users that they are copy-protected, and they shouldn't use the "Compact Disc" logo because they are not, in Philips' considered view, proper compact discs at all.
- Traitors In Our Midst - Music lovers have begun to fear that record companies purposely corrupt the data on audio CDs in an effort to restrict their use as a source for copies or MP3 files. Unfortunately, those fears are coming true with Cactus Data Shield digital-music-use restriction technology.
- Universal copy-protected CD shuns players - On the eve of Universal's kickoff of its long-awaited Pressplay online music subscription service, the record label is distributing its first copy-protected CDs in the United States, adding a few new twists to the controversial idea.

HDCD

HDCD® (High Definition Compatible Digital®) is a patented encode/decodeprocess for delivering on Compact Discs and DVD-audio better than normal CD audio quality. HDCD encoded CDs are encodedso that they fit 20 bits of information onto the CD while remaining completely compatible with the existing CD format.

- HDCD Home Page

DVD audio

DVD audio is a multichannel(PCM) audio format which stores many hours of high quality (better than CD) audio to a DVD disk which can be played on "DVD-Audio" capable DVD players. This audio format makes it possible to store more than 74 minutes of sound in two channels using the highest sampling frequency 192 kHz with 24-bit quantization, or six channels at a 92-kHz sampling frequency with 24-bit quantization. That's better than twice the performance of the CD audio specifications. Thus, dynamic range is 144 dB and frequency response is up to 96,000 Hz. The maximum transfer rate is 9.6 Mbits/second. The DVD-Audio format was originally intended to be finalized in spring 1996, but company infighting stretched that time line.

Even though though version 1.0 of DVD-Audio specifications (the PCM format descibed above) was published on the end of 1998, there is still some fuel on the format fight. Sony and Philips are also behind a competing system, Super Audio CD(SACD), using a one bit sampling 2- and multichannel format called Direct Stream Digital(DSD). A basics SACD disk will not play on normal CD player (you will need to have a SACD player to play it). It is possible to make so called hybrid CD/SACD, which will play back on most CD players (but not all). This hybrid disck has two data layers: one for normal CD data (that normal CD player reads) and another for SACD data (for SACD players). The SACD data in the disk is 1-bit bit-stream sampled data that is sampled at 2822.4 kHz sample rate. SACD support 1 - 5.1 audio channels. The maximum sound playback length of SACD is 109 minutes when stereo sound is stored to it. SACD has it's own copy protection mechanism that is built into the standard from the beginning and is supported by all SACD players.

A DVD-Audio discs can also be played on almost any DVD-Video player if it in the Video-section includes a track that follows the DVD-Video audio specs (some DVD-Audio disks have this as an option for compatibility). An optional CD layer can give DVD-Audio software compatibility with some players for CD and DVD-Video.

The record industry is only interested in SACD/DVD-A because of the built-in copy-protection features. Philips and Sony are interested in SACD because of the royalties they'll receive for many years.

- Digital Versatile Disc (DVD) - new laserdisk standard for audio and video
- DVD-Audio spec unveiled
- Guideline of Transmission and Control for DVD-Video/Audio through IEEE1394 Bus

DCC

Digital Compact Cassette (DCC) is dying digital tape format from Philips. This system was able to storage quite high quality (comparable to MD) MPEG compressed digital audio to a compact cassette format cassette (basically same mechanics, just better tape in it). DCC players were also capable of playing back normal analogue compact cassette (C cassette) tapes. Nowadays you can't buy new DCC devices. This technology lost the battle against MD and CD-R/CD-RW technolologies.

- DCC FAQ
- DCC-L Mailing List Homepage - information about DCC, mailing list subscription information, and archives from the list, includes DCC FAQ

Minidisc

MiniDisc (MD) is a small audio disc format. The disc used in the system is 7 cm (2.75") x 6.75 cm (2 21/32") x 0.5 cm (3/16"), the disc inside is 64 mm in diameter. One MiniDisc give around 160MB storage for 74 minutes in audio mode and 140MB in Data Mode. MiniDiscs are available as ready made records to buy and as re-recordable discs for making your own recording. The audio data in mini disc is compressed using ATRAC algorithm.

MiniDisc is suitable format for small portable audio listening devices and also for making quite high quality audio recordings.

- ATRAC: Adaptive Transform Acoustic Coding for MiniDisc - information on audio coding in use in MiniDisc system
- MiniDisc Community Pages - MiniDisc information pages which are mirrored all over the world, also mirrored in Finland
- MiniDisc DIGITAL Recording FAQ

- Minidisc FAQ: MDLP (``Long-Play'') Mode Topics - Minidisc Long Play is a new encoding method for audio on Minidisc that offers two modes: one gives 160 minutes stereo ("LP2"), the second gives 320 minutes stereo ("LP4").
- Minidisc Frequently Asked Questions
- MiniDisc Manager - software allows computer control of Sony MiniDisc recorders that have a compatible slink control port
- The MiniDisc Appreciation page - information about mini disk and how to use it in different applications
- US MiniDisc page

DAT

Since Sony introduced DAT (digital-audio tape) in 1987, it has been the preferred format option of many journalists, musicians, and live-concert and lecture-audience members, who want to capture a sonic snapshot of the events they attend.

Record-label-piracy concerns and consequent copy-protection and royalty complications, though, have curtailed DAT's acceptance in the consumer-electronics world.

Nowadays recording studios are moving to hard-drive-based setups, so the DAT's future is on the wane. DAT gear seems to be rapidly disappearing from manufacturer warehouses and store shelves.

- DAT-heads - home page for the DAT-heads mailing list
- DAT Tapes page - DAT links
- DAT World
- How to maintain your DAT
- Live taping to DAT
- Sony TCD D-7 and TCD D-8 resource page
- Sound savings: Portable audio recorder takes on tape, part 1 - DAT gear is rapidly disappearing from manufacturer warehouses and store shelves. What will replace it? Core Sound, along with a symphony of worldwide development partners, delivers a harmonic proposal.
- Tascam decs FAQ

Other digital audio formats

- PC-Based Multitrack Recording Frequently Asked Questions (FAQ)
- Unofficial Tascam DA-88 web page

Digital audio broadcasting

- DMX, Digital Music Express - digital music channels from cable
- UK Digital Radio Forum

Networked audio

- Networking for Audio
- Routing Audio Via Ethernet White Papers

Papers and FAQs

- AAC audio analysis - AAC has substantially fewer artifacts and artifact magnitudes than its MP3 codec predecessor.
- Audio Compression - very good introduction to audio compression part of Simon Fraser University CMPT 365 Course notes
- Digital Audio - Material for the lectures
- Digital Audio Transmission - Why Jitter is Important
- Everythign you always wanted to know about jitter but were afraid to ask - jitter is often misunderstood among recording engineers and audiophiles
- The Secrets of Dither - or How to Keep Your Digital Audio Sounding Pure from First Recording to Final Master

Speech and music coding

Information on coding

- AC-3 specification - in pdf format
- Audio Compression Introduction
- Comp.speech Frequently Asked Questions web page
- GSM 06.10 digital speech compression
- Human auditory masking and perception based coding - introduction to this topic
- Philips MPEG audio compression technology - documents, technology, software, components, links
- Speech Coding - introduction to speech codecs

Software

- Hyperreal Tool Shed - audio tool programs for computers
- Source code for audio codecs in C - ADPCM, G.72x, LPC, GSM, CELP

Digital audio interfaces

AES/EBU and S/PDIF

At the consumer level there are basically one digital S/PDIF format which is transported using two different interfaces. One commonly used interface option is electrical coaxial interface, using an RCA-type jack and plug (or sometimes a BNC bayonet connector) and a shielded, 75 ohm impedance cable. Other option or consumer S/PDIF interface is Toslink optical, which transmits the digital data in the form of light down a plastic fiber-optic cable (1 mm diameter fiber center). Nowadays the official name for S/PDIF is standard IEC958 "Digital audio interface". Professional audio equipments use AES/EBU format defined by standard IEC-958, which uses a 3 pin, balanced XLR-type connector and 3 conductor, 110 ohm impedance cable (allowed range 88-132 ohms) or multimode fiber optics (AT&T or ST-Glass).

Both AES/EBU and S/PDIF interface support a wide varity of sample rates. Here is table of sample rates and needed bandwidth.

```
Sampling Rate  Bandwidth
 32.0 kHz       4.096 MHz
 38.0 kHz       4.864 MHz
 44.1 kHz       5.645 MHz
 48.0 kHz       6.144 MHz
 88.2 kHz      11.289 MHz
```

```
 96.0 kHz      12.228 MHz
192.0 kHz      24.576 MHz
```

The consumer devices with S/PDIF interface have generally supported sample rates only up to 48 kHz. But the introduction of higher resolution audio formats have added the sample rates up to 192 kHz.

The maximum sample resolution supported is 24 bits. Many consumer S/PDIF implementations are though limited to only 16 bits of resoltion.

The cables uses in the digital audio connections vary between digital interface selected. The cable that connects to the coaxial interface is designed to be 75 ohm ciaxial cable. It is the same cable type as used in video industry to transfer video signal. For example the video industry standard RG-59 coaxial cable is very ideal to carry coaxial S/PDIF signal for long distances (the maximum supported distances). For shorter runs practically any 75 ohm coaxial cable will do nicely. For very short runs (up to 2 meters) the needs of coaxial cable are not so critical, and a normal "RCA audio cable" usually works well. For anything more than few meters run, use with 75 ohms coaxial cable and things work for sure.

Many new, high-quality Mini Disc, pro-audio, DAT (Digital Audio Tape), CD, DVD, and laser disc players, as well as digital amplifiers, DSS satellite receivers, and computer sound cards, are manufactured with digital optical output connectors. These connectors attach to optical cables, which are constructed with a PVC jacket and a plastic core. The cables transfer information accurately over short distances via digital light signals. Digital optical cable with plastic-core construction is less expensive than fiber optic cable with a glass core, but it still provides the benefits of optical transmission over short distances. The two types of connectors associated with digital optical transmission are TOSLINK®, a Toshiba® trademark, and the 3.5-mm Mini Plug connector.

Sometimes there is questions on hifi forums raised should I use coaxial or (Toslink) optical interconnection ? They're both digital, so as long as those conversion processes are reliable, nothing will be lost. When properly implemented their performance is equal. Sometimes people recommend coaxial connection over optical because of less processing on the way. But the processing between optical and electrical signal presentations does not change the data contents. This means that there can't be any difference in frequency response, amplitude distortion or noise. The properties of the signal are encoded in the digital data. That data is unchaged by the conversion process. If errors happen (baddly operating devices), the system will notice those very quicly (S/PDIF has parity bits to check data integrity) and mute the audio output.

One might also see arguments about optical connections introducing jitter. There could be a difference in jitter, but comprehensive measurements fail to show any significant differences. Jitter is way overplayed as an audio issue. As long as you are in the digital domain, jitter is utterly irrelevant. It is only at the moment of conversion to/from the analog domain that jitter is relevant. If you have a situation where the jitter is demonstrably worse with an optical connection, the fault lies not with the optical connection; it's in the D/A converter itself, and is the result of poor clock design in the D/A converter. In some cases, there are some bad implementations of optical. Combine that with a BAD implementation of a D/A converter and you have a system that does not work well. But, take solice in the knowledge that as bad as some people can put together a bad implementation of an optical output, there are others who can put together an equally incompetent implementations of coaxial output as well. Some such incometently equipment can be quite sensitive to the interconnecting cable. The quality of the sending and receiving units in your components will have a far greater impact than the cable type you choose.

The only benefit I can think of from an optical (Toslink) connection over a coax is the galvanic isolation between the two equipments, which can avoid ground loops. This could be relevant in some audio systems where ground loops are a problem. The downside is that Toslink optical connection has shorter distance limits than coax and optical cable usually costs more than coaxial cable. Optical cable is more fragile, more expensive, and the connectors are typically flimsier.

The general advice is use whatever (optical or electrical) you have easily available and be happy. In the digital domain, you can even convert between optical and electrical all you want without any loss (potentially, if your hardware does everything right).

Optical could be technically better for short runs (less than 10 meters) because it resists interference and ground loops. Optical cables are usually a lot more fragile than coax cables and usually cost more, so optical cable is not best for all. Coax for long runs because Optical is rated only out to 10 meters or so. There should be no difference between optical and coaxial transmission of digital data. Within their recommended lengths, they sound the same. If there are problems, the most likely symptoms are interrupted sound, clicks, pops, or no sound at all.

The main problem with Toslink is not sound quality. There is no "quality" difference between Toslink (optical) and coaxial (metalic). Ones are ones and zeroes are zeroes in the digital realm. The main issues with Toslink are durability, maximum length, and installability (Toslink cables aren't exactly supple). On the other hand Toslink optical connection is a great solution for ground loops and EMI pickup or creation.

In coaxial interconnection sometimes should I use a special "digital cable" or will a normal "RCA audio cable" work ? The right thing to use in any length is 75 ohm coax with appropriate connectors, such is used for RF and composite video. A "normal RCA cable" does not usually have the right impedance, because those cables can have their impeance anywhere in 45-75 ohms range and the signal attenuation usually much higher than proper coaxial cable. In short lengths like 1 meter, just about anything that is conductive and look like a cable will work about as well as anything else. A standard 1 meter audio interconnect can do the job just well.

On professiona AES/EBU interface the standard cabling solution is to use 110 ohm shielded twisted pair wire terminated to three pin XLR connectors as the standard defined. Most digital audio cables utilize foam polyethylene to minimize the cable's size. The advent of digital microphones requires AES/EBU cable designs with added flexibility, Although AES/EBU specifications require shielding on each channel of data, datagrade UTP .Category 5. can easily meetthe common mode balance requirements (-30 dB) without being shielded. You might ask can I use normal analogue audio cable (shielded twisted pair) for AES/EBU connection ? The answer is yes, but with limitation. The actual length is determined by the error correction and jitter tolerance of the receiving device. The impedance of most analog cables ranges from 40 ohms to 70 ohms. This large mismatch from the nominal 110 ohms results in signal reflection and jitter causing bit errors at the receiver. Also, the high capacitance of analog cables greatly increases the rise time of the digital square wave.

Another option to carry AES/EBU signal through 75 ohm coaxial cable wiring. Generally this is done using baluns. The transmission of digital audio over 75 ohm coax requires the use of baluns unless the device contains unbalanced coax AES inputs or outputs. In case your equipment has only balances inputs, you can use baluns to convert the unbalanced coax signal and a 110 ohm balanced transmission from one format to another. Much greater transmission distances are obtainable over coax as compared to twisted pair. The same coax used for digital video is ideal for digital audio. The coax used should have a pure copper center conductor (no copper covered steel or aluminum) and have good braid coverage (90% or more).

- AES-3 distribution: Coax or twisted pair?
- AES/EBU and S/PDIF digital audio interfaces - zipped ascii text file and GIF pictures
- AES and S/PDIF Recommended Transformers
- Digital Format Converters MSB, Midiman & Fostex Review - converters for S/PDIF conversion between Toslink optical and coaxial electrical
- Digital Studio Cable Guide
- Expert opinion: which SPDIF, opical or digital?
- FAQ for AES/EBU and S/PDIF interfaces - posts from the USENET newsgroup rec.audio.pro collected to one zip file
- Overview of Digital Audio Interface Data Structures - This document is a short summary of the most important parts of data format information from IEC958 and new AES3-199x and TV84 documents.
- S/PDIF Interface - a description of data format used in CD player digital output and some exaple circuits
- SPDIF overview
- The Jitters and Its Cure

MADI

MADI (Multichannel Audio Digital Interface) is a professional multichannel version of AES/EBUstandard for transmitting up to 56 channels of digital audio data over a single coaxial cable terminated with BNC connectors. MADI uses a second cable for word clock, with a fixed data rate of 100Mbps used on large, open-reel digital multitracks. Optical MADI implementations are available.

MADI (Multichannel Audio Digital Interface), also known as AES-10 standard, allows interconnection of two devices to transmit up to 56 channels of 24-bit digital audio with a single coaxial cable or via optical link. It is a standard interface to digital multitrack machines and mixing consoles like the Studer D941 On-Air Mixing console, Studer 950 Digital mixing system, Neve Capricorn or Sony PCM3348HR Digital Tape Recorder. Some manufacturer specific MADI extensions allow up to 64 channels in each direction and can operate at 96 kHz (offering up to 32 channels) or even up to 192 kHz (with 16 channels on one single connection each way).

Other

- Audio Distribution and Control using the IEEE 1394 Serial Bus - this paper describes a system for transmitting hundreds of channels of high quality digital audio, control & monitoring protocols, and digital video over an IEEE 1394 serial bus
- I2S enhanced - preliminary specification
- New Olympic Stadium Audio System Runs Via Ethernet
- SCSI Musical Data Interchange (SMDI)
- TDA1543 DUAL 16-BIT DAC Data Sheet - includes description of I2S interface
- The D7 7-pin connector
- The Kurzweil KDS 8+2 Channel Digital Audio Interface

Audio compression

The issue of using compression is the matter of reduction of the amount of data. A direct measure of the total amount of data or information in audio is its bitrate. By using the compression the audio signal can be transported with lower bit rate or stored to a smaller file size.

A typical audio file where we start from is CD quality audio. Audio Compact Disks store the audio data in files on the disk. The audio data is in a PCM (Pulse Code Modulation) format. Each minute of recording time consumes about 9 MB (Megabytes) of file storage space. A three minute song would occupy about 27 MB of file storage space, and a 5 minute song would occupy about 45 MB of file storage space. A 650 Megabyte Compact disc can contain up to 74 minutes of PCM audio.

The sheer size of PCM audio data files made them unpopular as a storage medium for audio on a computer or the fole format when the audio is transported through slow modem connection. Digital compression came along and changed all that. Digital compression offers possibilities for a huge reduction in the amount of file space required to store audio data files. This reduction in file size also reduced the time required to transmit them electronically dramatically.

Audio compression can be one of two categories, lossless or lossy:

- Lossless: An intermediate representation of the audio data is created where, in most cases, redundant information is removed. The data rate, in the case of CD audio, would be less than 1,411,200, but because redundant information was removed, and we can know how to put it back exactly, there is no loss of data when we decode it at the other end.
- Lossy: Data which is not deemed to be audibly significant is removed. Then, redundant data is removed. The resulting data rate is also less than 1,411,200 bits/sec, but, on decoding, while the redundant data is restored, the audibly insiginificant data is not. Thus, the decoded data, while it is the same data rate, is different than the original.

Lossless digital compression is commonly used to reduce the size of computer files for electronic transmission. In order for the files to be useable on a computer the files that are extracted from a compressed data file must be identical to the original file (before it was compressed). Lossless compression is great because it makes perfect copies but it doesn't yield very high compression ratios. That means it doesn't save huge amounts of disk storage space. ZIP, ARC, TAR, and SIT are some of the acronyms or formats of Lossless Compression commonly used on computers. Lossless compression sounds wonderful, however it does not yield an extremely high compression ratio. Typically, lossless routines give a ratio no higher than 2:1.

Lossy compression algorithms offers much higher compression ratios than lossless algorithms but in order to achieve this they need to discard some of the original data. Lossy Audio Compression can yield a variety of different rates of compression based on the ability of your hardware and software to encode and decode music audio data. Lossy compression is only suitable for use on audio or graphical data. The audio or graphics are reproduced but at a lower overall quality than they had before they were compressed. In some cases the difference is difficult to perceive. The compression ratio can usually be adjusted so the quality level can vary widely. Audio that is compressed at a 20:1 or 10:1 ratio will certainly sound inferior to audio that was compressed at a 2:1 ratio. There are various compression schemes to achieve different results; for example, many codecs remove portions of the audio signal that human hearing is less sensitive to. To some peoples ears the resulting audio has distinguishing artifacts or characteristics that change the listening experience. The final "quality" level is largely a matter of personal perception. Each different encoder/decoder (CODEC) has different strengths and weaknesses. Typically, encoding takes a long period of time and large amounts of processing power. decoding can usually be made real time or faster. MPEG, MP3, AAC, RA, WMF, JPEG, QT, and DivX are some of the acronyms or formats of Lossy compression commonly used for audio and video. The most commonly used and most talked about audio compression in typical computer and consumer audio enviroment is MP3.

Brief descriptions of some different audio compression formats:

- ADPCM: Adaptive Delta Pulse Code Modulation (ADPCM) is popular since is returns a high quality signal with very little processing power required for fast decoding. ADPCM comes in a few different versions. The Microsoft algorithm found in Windows ADPCM codec returns a compression scale of 4:1 (one fourth the original file size).
- MPEG: The Moving Pictures Experts Group (MPEG) put a considerable amount of research into ways of compressing video and audio to make movies fit onto compact discs. It hs released several versions of MPEG audio coding standards. Software routines for MPEG audio decoding are fast to decode with relatively low quality loss, but generally speaking it takes a LONG time to encode. MPEG audio compression loses quality quickly at compression levels higher than a 6:1 compression ratio. MPEG Audio compression continues to degrade until it becomes absolutely dismal at 24:1 compression. An extension to MPEG-1 and MPEG-2 audio called 'Layer 3' (commonly known by its file extension, MP3) has been developed allowing greater compression levels and more BIT DEPTH options. Since MPEG audio uses a perceptual model, the exact perceived amount of lost quality varies for each individual listener. MPEG-4 has been ratified as a standard, and new codecs will become available in the not-too-distant future.

The compression ratio for MPEG compression can be adjusted across a wide range of quality levels, however sound quality drops as the compression is increased. Trying to maintain the highest quality (without extremely noticeable signal loss) yields the following compression ratios for different MPEG coders:

```
Layer 1 = 1:4 (384 kbps for stereo)
Layer 2 = 1:6 or 1:8 (256 kbps or 192 kbps for a stereo signal)
Layer 3 = 1:10 or 1:12 (128 kbps or 112 kbps for a stereo signal)
```

With the increase in inexpensive mass storage and the increases in connection speeds we are rapidly approaching a point in time when compression of audio may not be as important of an issue as it was in the past.

General audio compression links

- Compression on a Windows PC

MP3

The MP3 format is a compression system for music. MP3 (MPEG-1, Layer 3): a compression standard that creates relatively small digital audio files with high-fidelity sound. MP3 is simply a file format that compresses a song into a smaller size so it is easier to move around on the Internet and store. The MP3 format for digital music has had, and will continue to have, a huge impact on how people collect, listen to and distribute music. The MP3 movement is one of the most amazing phenomena that the music industry has ever seen: Web's ability to advertise and distribute MP3 files is huge. Very many people are now downloading, listening to and saving MP3 files onto CDs!

The MP3 format helps reduce the number of bytes in a song without hurting the quality of the song's sound. The goal of the MP3 format is to compress a CD-quality song by a factor of 10 to 14 without noticably affecting the CD-quality sound. With MP3, a 32-megabyte (MB) three minute song on a CD compresses down to about 3 MB. This allows you to download music at useable speeds from the network and store hundreds of songs on your computer's hard disk without taking up that much space.

MP3 is lossy compression algorithm. When the music is compressed to MP3 format, some part of the original sound is lost. MP3 encoders create smaller files by getting rid of unnecessary audio information. When you convert .wav files to MP3, an encoder filters out data representing sounds outside the average person's hearing range. Because some of the original audio data is lost, this technique is called lossy compression. MP3 format uses characteristics of the human ear to design the compression algorithm: there are certain sounds that the human ear hears much better than others and certain sounds that the human ear cannot hear (especially when played with other sounds). Using facts like these, certain parts of a song can be eliminated without significantly hurting the quality of the song for the listener. When you are done creating an MP3 file, what you have is a "near CD quality" song. The MP3 version of the song does not sound exactly the same as the original CD song because some of it has been removed, but it's very close.

Lossy compression has its downside: The more data the encoder throws out to shrink the file size (a setting you can control), the worse the finished product sounds when you play it back. The result is that not all MP3 files sound alike. Also the encoding software can have effect on the sound quality, because the human hearing model used in then can be slightly different, so different encoders can throw out sligtly different parts of the sound, thus generating slightly different results. The differences in quality are usually largest at lowest bit rates (where most compression is used).

MP3 industry claims that music encoded at 128 kbps will produce "CD quality" MP3 files -- and at that quality, a one-minute audio file fills about 1MB of disk space. But regular MP3 users disagree, saying that they hear significant distortion unless the file has been coded at 160 kbps, which produces files that take up slightly more space but sound much better. There are some people that say that they need much higher conding rate. In some special application lower than 128 kbps rates can be used. MP3s coded at 96 kbps end up as far smaller files but have considerable noise.

Once you have an encoded MP3 file, you need a decoder or player -- which can be either software or hardware -- to convert encoded MP3 data into something you can listen to. You can download most software MP3 players for free from the Web. Maybe the best known MP3 playback program is Winamp.

Many people think that the MP3 standard is free and open, and that the ISO reference source code is also. But it is not entirely free. MPEG development is almost open. When an MPEG format is approved by ISO, it contains several parts. It defines the bitstream syntax that must be used to be compliant, operations done by the decoder, and it provides an informative only source code for encoding and decoding, which is freely available. But this source code is purely indicative, as MPEG never defines any encoding rule. The Fraunhofer Institute has been the main developer of MPEG audio Layer-3, and the MP3 standard that has been approved is mainly based on their work, which Fraunhofer has protected by several patents (a portfolio of 18 patents related to MP3). This patent portfolio owned by Fraunhofer and Thomson Multimedia is very extensive, and cover various aspects of MP3 encoding. Practically you can not use any MP3 implementation without using some parts of those patents. other companies could have some patents also covering parts of MP3. MP3 standard is established since 1992, and no one except Fraunhofer/Thomson Multimedia asked for fees.

- How MP3 Files Work
- How MP3 works
- How MP3 Works: Inside the Codec - So what's the trick? How does the MP3 format accomplish its radical feats of compression and decompression, while still managing to maintain an acceptable level of fidelity to the original source material? The process may seem like magic, but it isn't.
- Patents and MP3 - Fraunhofer has protected MP3 technolgies by several patents.
- Philips MPEG audio compression technology - introductory documents to MPEG audio technology, software, components, links
- MP3: The Definitive Guide Sample Chapter 2: How MP3 Works: Inside the Codec
- MPEG Audio FAQ
- mp3/mp3PRO Patent and Software Licensing Information
- The Internet-audio (r)evolution - MP3 has launched digital audio into the mainstream, revolutionizing the way people collect and play music

Computer audio

General

- Computer speech compression techniques - comp.speech.FAQ, speech data, software
- Digital audio experiments - pages devoted to DSP and digital audio
- PC soundcard hardware pages - information about PC soundcard hardware, software and many related topics

File formats

- About Sound Formats - a brief description of the various sound formats
- Audio file formats FAQ

Soundcards

- Interfacing Microphones to Computer Sound Cards
- Tomi Engdahl's PC soundcard pages

Software

- RealAudio - real time audio on Internet
- Sound Processing Kit - object-oriented class library for audio signal processing implemented in C++

Contents rights management

Tecording industry is worried about copying of their products (CDs). The recording industry has tried putting down the first critical pieces for a system it hopes will keep songs on the Net from being pirated.

The general term 'piracy" refers to the illegal duplication and distribution of sound recordings and takes three specific forms: counterfeit, pirate and bootleg. Counterfeit recordings are the unauthorized recording of the prerecorded sounds, as well as the unauthorized duplication of original artwork, label, trademark and packaging of prerecorded music. Pirate recordings are the unauthorized duplication of only the sounds of one or more legitimate recordings. Bootleg recordings are the unauthorized recording of a musical broadcast on radio or television or of a live concert.

The most talked about technology has been Secure Digital Music Initiative (SDMI). The first phase of the SDMI system requires that portable digital music player manufacturers implement several security components, foremost among them a digital rights management system (DRM). This will allow record labels to securely distribute and track files as they are transmitted over the Net and on to portable players.

General information

- Digital Rights Management - information portal
- The Federal Anti-Piracy and Bootleg FAQ Bootleg recordings are the unauthorized recording of a musical broadcast on radio or television or of a live concert.

SDMI

- SDMI: Divide or Conquer? - basic information on SDMI system
- SDMI Introduction by RIAA
- Why Secure Digital Music Initiative is falling apart

CD copy protection

Ever since Napster came to prominence, the music CD publishers have been looking for a way to stop people sharing MP3s extracted from their CDs, and now they think they've found it -- by copy-protecting the CD releases. They hope that by making CDs unplayable on computers that this will reduce the number of MP3s getting onto the internet.

CD copy protection is mainly done to make it harder to make copied of CD or converting them to MP3 files on PC. There have been meny technologies tied, but none of them is completely safe or problem free (and there will not be any such for CD). Generally CD copy protection systems more or less try put some type of errors to CD disk content (tale or contensts or audio itself) so that it does not either play on PC or sound gets corrupted on "CD ripping" process. No CD protection which still allows playing the music is bulletproof, even though protection makers say "We can stop all kinds of copying, even on domestic CD recorders." Sensible industry experts are skeptical on CD protection systems promises of them being any long term solution, ebecause hackers have defeated earlier security (promised to be secure) technologies very quickly. The fact for all CD protections is that if you can play the CD with CD-player, analogue copies can be made to any analogue devices (recorded to tape) or current digital devices with analogue inputs (for example PC soundcard) with some loss of quality.

Cactus Data Shield protection for CD is designed to prevent MP3-ripping by corrupting the table of contents so it won't play on a PC. Cactus inserts modifications to original CDs in a way that confuses CD-ROM devices during the copying process. If the CD is copied, however, the copier machine (a PC or disc-to-disc copier) sees some fake control data as music. So when the copied disc is played, there are bursts of distortion as the player tries in vain to decode the garbage.

SafeAudio is an Audio CD protection developed by Macrovision. SAFEAUDIO add samples that sound like bursts of static, and scramble the ECC data around to make it look like an uncorrectable error. Audio CD players will interpolate the samples during playback, but CD-ROM drives doing digital audio extraction generally won't. The result is a disc that plays back correctly on a CD player, but won't "rip" or copy correctly on a CD-ROM drive. SAFEAUDIO v3 changes the music data at the bit level, flipping a fraction of a disc's billions of 1s and to 0s ("very subtle" degree of data corruption). y MusicGuard introduces selective alternative alterations to the EFM data stream (can cause problems for CD-to-CD copying programs).

CDs with the current version of Sony DADC Key2Audio cannot be recognised by standard CD/DVD-ROM, CD-R and CD-RW drives, thus they do not play on PC.

Suncomm MediaClòQ v1.0 protected CDs are designed such that they can not be read by any CD-ROM player or CD burner. The CD has a heavy visible band at about the place where the audio ends. Those CDs generally say "This CD is designed to play in standard Audio CD players only and is not intended for use in DVD players."

- Anti-piracy system could damage loudspeakers - Some CD protection system can be potentially dangerous to your HiFi system.
- CD Protections - This page has lots of information on game CD-ROM and audio CD protection systems. This page contains a list of all known CD Protections which are currently on the market, how they can be detected and how to by-pass many of them.
- Compromise for CD copying is in the works - The record industry is experimenting with a new strategy for protecting CDs from being copied in CD burners or on computers. Unlike previous anti-copying measures, this plan will place two versions of an album on a single disc: one in standard CD form, modified so that it can't be transferred to a computer hard drive, and another in Microsoft's Windows Media Audio digital format, rigged so that files can be copied to a PC, but with some restrictions on how they can be used.
- Copy-protected CDs - Buying a new CD? Watch out for inferior imitations
- Could CD protection end music swaps?
- Lawmaker: Is CD copy-protection illegal? - Record companies' efforts to protect CDs against digital copying are beginning to draw scrutiny from lawmakers concerned that the plans might violate the law. A 1992 law allows music listeners to make some personal digital copies of their music. In return, recording companies collect royalties on the blank media used for this purpose.
- Natalie Imbruglia wins CD protection race - Aussie pop star (and former Neighbours actress) Natalie Imbruglia has won the race to produce the first copyright-protected CD on general release to the UK public.
- SafeAudio explained and should we fear it ? - How does it work ? - SafeAudio is an Audio CD protection developed by Macrovision.

Other protection systems

Related pages

- Audio pages
- Digital Signal Processing page
- PC Soundcard page
