AudioControl®

Technical Paper
Number 105

How Recordings are Made

by

Richard Chinn

Rick Chinn is a working sound engineer who has been involved in all facets of recording, sound reinforcement and acoustical design for the past fifteen years. He has written articles for Modern Recording and Music, Music and Sound Output and A/V Video and has worked for TAPCO and Kaye-Smith studios.

AudioControl®

Price $1.00

22410 70th Avenue West, Mountlake Terrace, WA 98043 (425) 775-8461

How Recordings are Made

The Myth

For many audiophiles or music lovers, the search for flat audio response is closely related to the search for the Holy Grail. To many of these people, the mere presence of tone controls on a preamplifier is enough to send shivers down their spines. The thought of using an equalizer or other signal processor to enhance their listening experience is as foreign as water to the desert. Unfortunately, or fortunately (depending upon your viewpoint) the majority of audio source material available to most consumers is more processed than a Chicken McNugget.

The recording process is fraught with compromise necessitated by limitations in the recording medium, and the consumer's playback system. Overcoming these compromises requires skillful manipulation of the audio signal during the recording, mixing, and mastering phases of a record project. Every aspect of the signal must be scrutinized: dynamic range, tonal balance, level, and signal-to-noise ratio. The final musical and tonal balance achieved in the mixdown session is a product of technical considerations and personal taste. There is absolutely no guarantee that your taste will agree with that of the producer, engineer and artist. The truth is that sensible and tasteful use of signal processing can actually enhance your listening pleasure.

The reality of the situation

Most of us have been convinced by the media and equipment manufacturers that flat frequency response is a desirable (and attainable) goal.

What exactly is flat frequency response? It is the ability of a sound system to reproduce a range of musical pitches at exactly the same volume regardless of the actual pitch. If you graph the volume level versus frequency, the resulting graph would resemble a straight horizontal line. We call this type of response flat. For amplifiers and other purely electronic devices, this is easy. For signal storage devices, like tape machines and turntables, it's more difficult. For electro-acoustic devices, like loudspeakers, it's very difficult. Even digital devices have their own set of problems. It's pretty easy to achieve flat frequency response, but a host of other problems that we are just beginning to understand are present. So, even though it's possible to achieve zero frequency response error via digital storage, the problems of dynamic range and low-level signal dynamics present new challenges to the digital mastering engineer.

The truth is that over the years, we have been conditioned to accept non-flat frequency response as being flat!

Consider: Almost without exception, every bit of source material available to the audio consumer has seen the likes of an equalizer at least once (and probably more like three times) on its way from the minds and souls of the performers to our ears.

Compounding this are loudspeakers that are voiced to sound "pleasing" to the average consumer in average situations (read: living rooms). There is no set criteria for what is "pleasing" since it is human judgment making that decision. If you've ever shopped for speakers before, then you know that the loudspeaker is the primary factor in determining the "sound" of a sound system. The primary reason that loudspeakers sound different is because their frequency responses, in real rooms, are all different.

The bottom line is simple: We've lost sight of what is truly flat. Perhaps this will become clearer with a brief look at the recording process.

Anatomy of the Most Popular Recording Process

The majority of music sold and listened to today is what has become known as popular music. This includes country, folk, jazz and other non-classical styles. For the most part, this material is recorded using anywhere from 16 to 48 track analog recording equipment. (48 track requires two 24 track machines that are synchronized with each other.) For digital recording, it's either 2 track or multiples of 32 track.

The (supposed) big advantage to today's multitrack technique is the ability to go back and enhance or modify individual elements without effecting the rest. This process is known as overdubbing or sweetening. For the most part, almost everything we hear is recorded in piecemeal fashion. Today's song usually starts out having the rhythm track (rhythm section and rough vocals) re-corded. Then the vocals get added, sometimes phrase-by-phrase. String parts are overdubbed. Last, certain parts of the tune are embellished by careful use of percussion instruments, handclaps, or other musical instruments. In the final mixdown session, the prerecorded tracks are mixed down to 2 tracks and special effects added to make a product that we can enjoy at home.

Keep in mind that recording equipment has become standardized enough to allow the different elements to be recorded in wildly diverse geographic areas. This geographic leapfrogging may be done to accommodate an artist's touring schedule or to take advantage of unique musical or recording situations. Many recording communities have become renowned through the years for the session musicians, arrangers and engineers that choose to live and work there.

Let's take a look at a typical song from conception to our ears. First, there is the matter of what microphones to use since the microphones add their own brand of equalization in the form of coloration.

Microphone Selection

If you've ever visited a recording studio, you probably noticed the wide variety of microphones used to pick up different instruments. Each microphone is selected for the way that its frequency response and directional characteristics enhances the way that it picks up certain instruments. There are certain microphones that are used more than others for more things, but there is definitely no such thing as a universal microphone. Consider the engineer's choice of microphones to be the first step in equalizing the sound.

The Neumann U-87 is a condenser microphone used for a variety of purposes: drums, vocals, strings, piano. It has a full-bodied sound with a warm sounding high end. The Neumann U-47 is another condenser microphone. Unlike it's more recent brother, the U-87, the 47 uses a vacuum tube. Many engineers prefer the 47's warm sound on vocals to that of the 87, which can sound cold by comparison. Another factor is the overload characteristic of the tube preamplifier in the mike. When a 47 approaches overload, it soft-clips. This probably accounts for it's warm sound. The U-47 hasn't been manufactured for at least 20 years and used ones that still work are worth 10 times what they sold for new.

The AKG C414 is another condenser microphone that is quite similar to the U-87. It too is used for a wide variety of purposes, but unlike the 87, it has a much brighter high end. It is a favorite vocal microphone for many engineers. The AKG C451 is a small condenser microphone with a characteristically bright sound. If you look at its frequency response curve, it is fairly smooth, but it is definitely not flat. From 50Hz to 15KHz, it's response has a rising characteristic that amounts to a 6dB difference between 50Hz and 15KHz. The 451 is my personal favorite for acoustic guitar.

The Shure SM 57 is a medium-cost dynamic microphone with a highly tailored response. The bass region is deliberately rolled off and there is an intentional peak in the 3 to 8KHz region for presence. It is an almost universal standard for snare drums and electric guitars.

Each of the five microphones just mentioned has a distinct sound. Each is used for a different purpose. None of them have flat frequency response. Figure 1 shows the response curves for some typical studio microphones. Figure 2 lists some common microphones found in the recording studio and their usual usage.

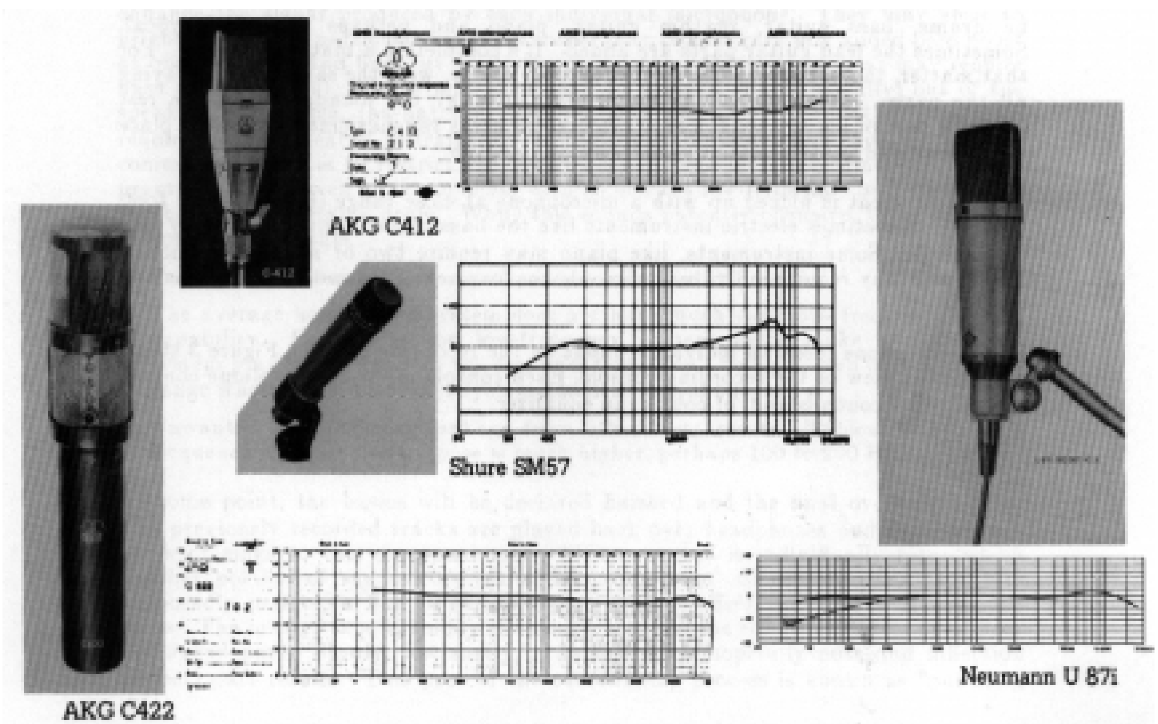

Figure 1. Some typical studio microphones.

Model	Type	Usage
AKG D12	Dynamic	kick drum
AKG C414	Condenser	vocals, drums, piano percussion, string bass, strings
AKG C451	Condenser	acoustic guitar, hi-hat, cymbals, percussion, piano
Neumann M-49	Tube Condenser	vocals
Neumann U-47	Tube Condenser	vocals, drums, piano
Neumann U-67	Tube Condenser	vocals, drums, piano, strings
Neumann U-87	Condenser	vocals, drums, piano, percussion, acoustic guitar, strings
Neumann KM-84	Condenser	piano,, drums, cymbals, percussion, acoustic guitar, strings
Electro-Voice RE-20	Dynamic	vocals, drums (especially kick), electric guitar, sax, brass
Electro-Voice RE-15	Dynamic	drums, electric guitar
RCA 44DX	Ribbon	brass, voice (spoken)
RCA 77DX	Ribbon	brass, voice (spoken), vocal
Sennheiser MD-421	Dynamic	drums, electric guitar, sax
Shure SM 57	Dynamic	snare drum, drums, electric guitar
Shure SM 7	Dynamic	vocals, drums, percussion, electric guitar
Shure SM54	Dynamic	drums, percussion, electric guitar

Figure 2. Common Recording Microphones

Recording the Basic Tracks

It all starts in the studio where the basic tracks are recorded. Typically, this might be drums, bass guitar, rhythm guitar, piano and perhaps a scratch vocal. Sometimes the lead guitar parts are added. It's completely a matter of choice. For that matter, the parts may be recorded one at a time, with the same person playing all the parts. The name of the game here is to lay the foundation for the rest (melodic portion) of the song. The scratch vocal helps the musicians keep their place in the song and usually gets replaced with an overdub.

Each instrument is picked up with a microphone at close range (less than one foot, usually). Sometimes electric instruments like the bass guitar are wired directly into the console. Some instruments, like piano, may require two or more microphones. The drums may require individual microphones for each of the individual drums and cymbals.

Each microphone feeds an individual input on the recording console. Figure 3 shows a simplified view of the recording console. Each console input has a volume control, several other controls, and, of course, an equalizer.

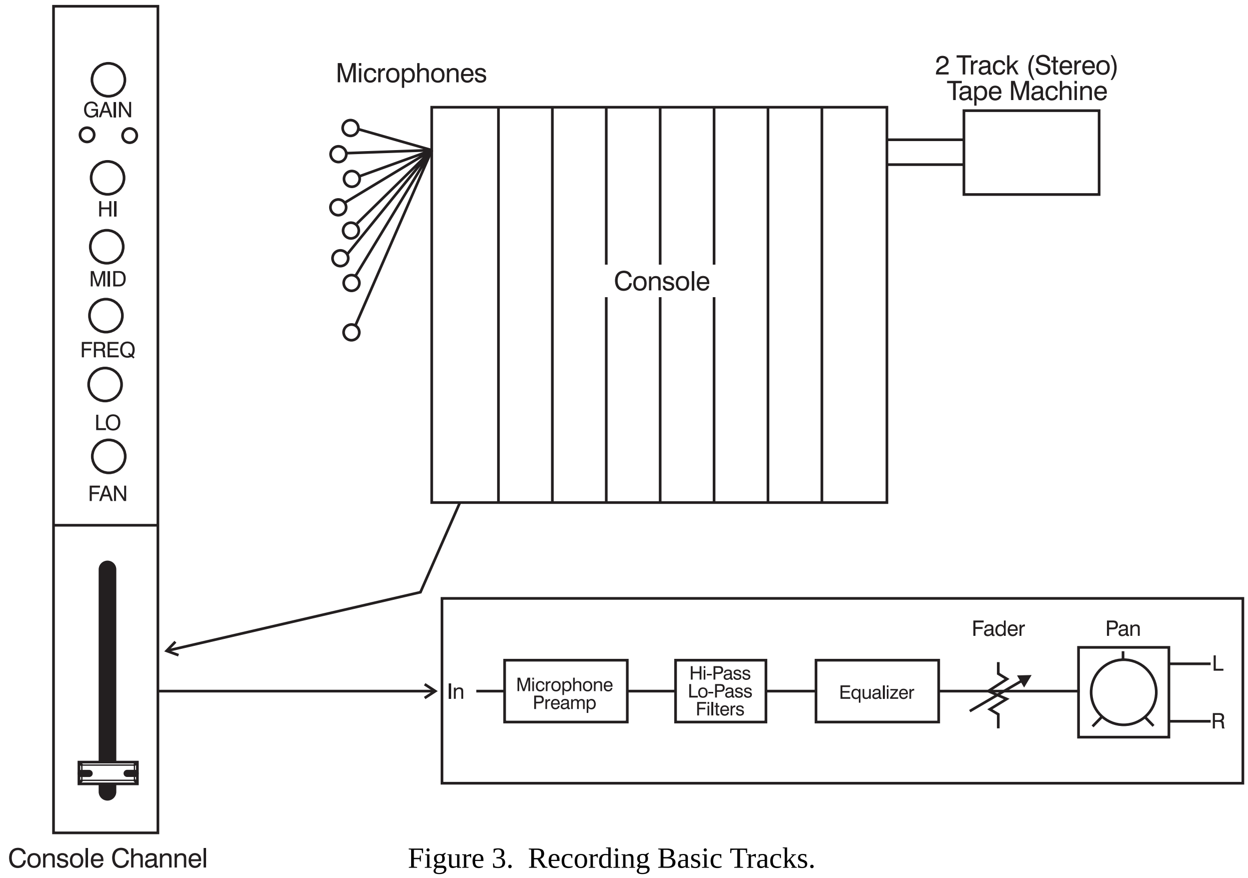

Figure 3. Recording Basic Tracks.

The equalizer allows the recording engineer, producer and artist to selectively enhance the signal produced by each individual microphone. They may elect to brighten or dull a particular sound. Another may require carefully added bass boost at one frequency and bass cut at another frequency. An electric guitar may require bass and treble cut, with a midrange peak to make it climb (literally) out of the speakers. Of course, (ha-ha) the equalizer may not be used at all, but his requires a great deal of restraint on the part of the producer and engineer. A very common practice is to restrict the low frequency response of certain inputs. This practice extends even to instruments with significant low frequency output, such as the bass guitar or kick drum. While this sounds patently offensive, there are good reasons for doing this.

- Excessive low frequency energy requires a wider record groove.
- The average home sound system does not have much very low frequency output capability. Restricting the low frequency content helps make the listener's system play louder because the amplifier power isn't being wasted on a frequency range where the loudspeaker system has little useful output.
- Unwanted low frequency leakage from other instruments. Typically the cutoff frequency used for this purpose is much higher, perhaps 100 to 200Hz.

At some point, the basics will be declared finished and the final overdubs begin. The previously recorded tracks are played back over headphones and the musician plays or sings to this accompaniment. This new part is individually recorded on another portion of the same piece of recording tape. Since the other parts are supposedly perfect, a fluff in the new part doesn't effect the previously recorded parts. The master tape is simply rewound a bit, and the recording starts at or near the fluffed part. Slowly but surely, a finished (and hopefully note and inflection perfect) part results. This part of the overdubbing process is known as "punching in".

Once the overdubbing is finished, the song is ready to be mixed.

The Mixdown Session

Mixing is the art (process?) of combining the many individual elements that were previously recorded into one or two (mono or stereo) channels. At this point, the material could be played on a home system (assuming that you have the equipment to play the tape). Since the mixing process occurs after-the-fact, the artist can participate directly in the final part of the creative process. Figure 4 shows how everything goes together.

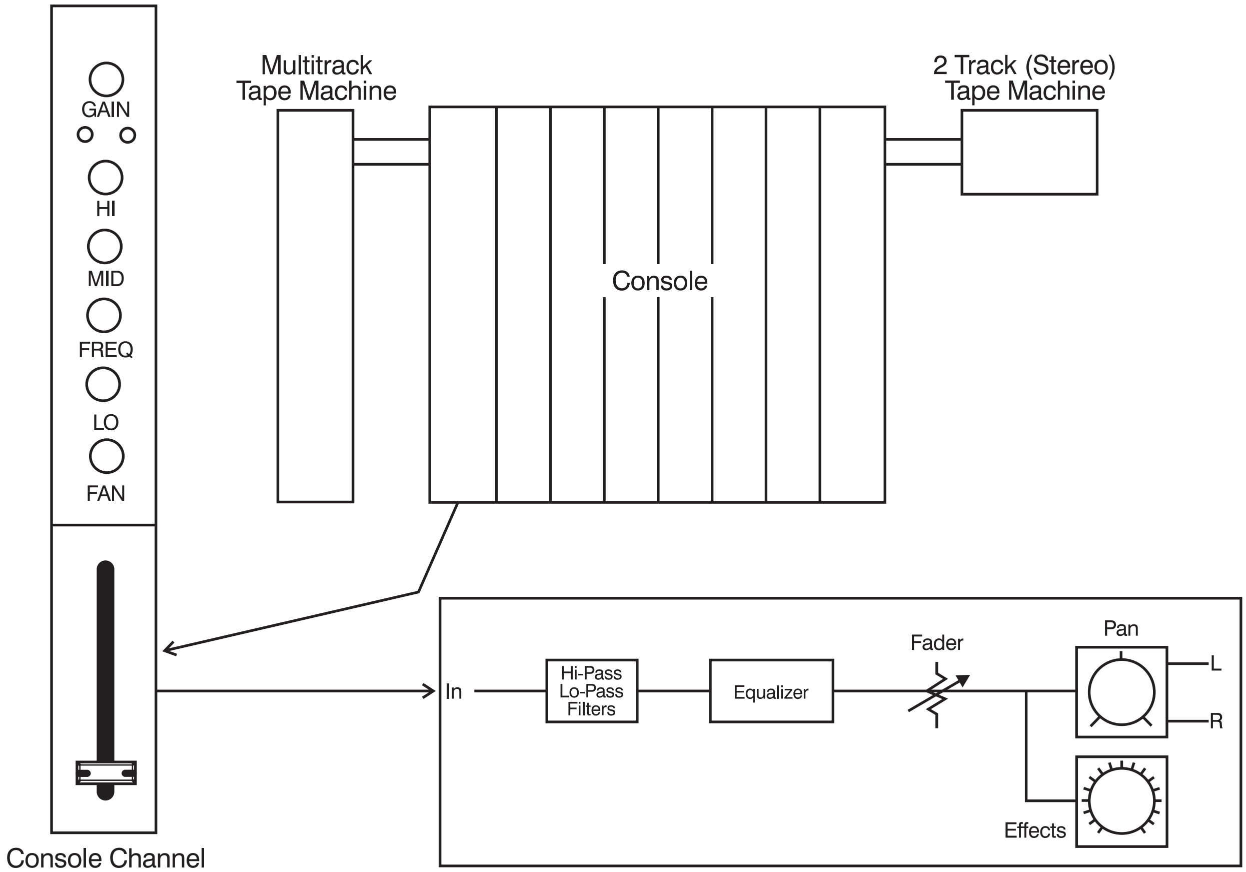

Figure 4. Mixing Down.

Fundamentally, mixing is quite simple. The multitrack tape is rewound. The tape machine is connected to the recording console (much like the microphones were) and the tape played back. Again, the engineer has control of level, equalization and spatial characteristics. After roughly setting levels to establish a semblance of balance, each track is scrutinized for tonal characteristics and re-equalized, if necessary. Dynamic variations are worked out and rehearsed. Echo, delay, reverberation and other effects are added as necessary.

While the basic tracks were probably monitored on a pair of fairly large, high-output monitor speakers, the mixdown session is not. For basic tracks, your monitoring requirements are: brutal accuracy and extremely high output capability. (Sometimes high volume levels are used to mask the absence of talent in the studio, other times they are used to be sure that “everything” gets heard.) For mixdown, the opposite applies: you want a speaker that is representative of the speaker that the end-user listens to. Smaller speakers, mounted nearby are the norm here. These speakers are boringly ordinary: Yamaha NS-10’s, JBL 4406’s, and the ubiquitous Auratones.

Finally, the tape is rewound and a try at the real thing happens. Punch the go button on the multitrack and on the two track mixdown machine, listen to the music, fiddle with all the knobs, adjust the faders, do the final fade. If it was all done right, you’ve got a master.

Again, it’s definitely not flat.

Mastering

Disk mastering is the art (black or otherwise) of transferring the finished master tape to a laquer master disk. Aside from the more or less mechanical part of pressing the finished records, this is the last step in the creative process. Don't get me wrong, it's as easy to blow it here as in any other part of the process.

The finished, mixed master tape is played on a special tape machine with two playback heads located one after another. The first head, known as the preview head, gives the lathe a sneak preview of what lies ahead in the song. From this information, the lathe decides how close together or far apart to put adjacent grooves and how deep to cut the groove. The second head is the normal playback head. It's output is the one that is actually recorded on the disk. Figure 5 shows the signal processing used for the tape-to-disk transfer.

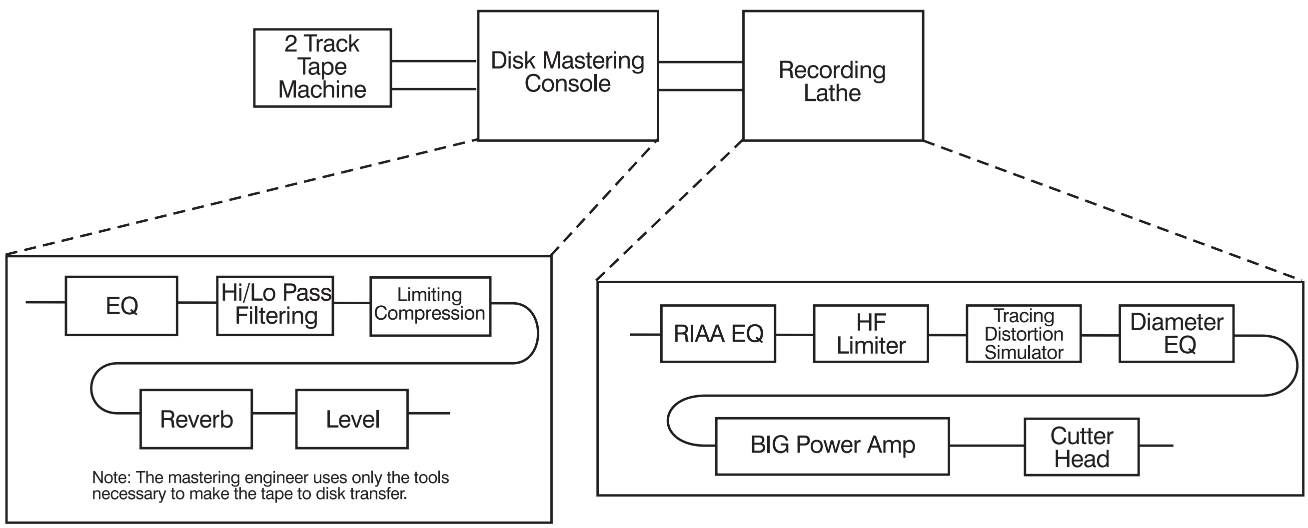

Figure 5. Disk Mastering.

The playback output of this machine is connected to a special type of console, the mastering console. This console usually has two identical channels of level control, equalization, and dynamic range control (compression or limiting). It is the mastering engineer's art (or job) to make the material on tape fit within the confines of a record groove. To do this, he (or she) must sometimes use every trick in the book. Additionally, this is the last chance for the artist and/or producer to make equalization changes.

Disc mastering has it's own set of compromises. Additionally, signal processing is used to help overcome some of the limitations of the disc recording and playback process, or to protect the expensive (over $10,000) cutter-head from burnout.

Basically, the scenario faced by the mastering engineer is that of a pie. The pie is divided into pieces: level, recording time, sequence, and dynamic range. The pieces can be divided as needed, but they always equal 100%. Excess time means reduced level. If the producer wants the record to be 3dB louder than anything else, it's going to cost time. If the last cut on the side is louder than the rest, the overall level of the cut may have to come down because high levels in the inner diameters of the record cause distortion in playback. Diameter losses in the inner grooves may necessitate high frequency boost for those cuts. Excessive dynamic range could result in a signal-to-noise ratio problem on playback. Excessive stereo low frequency material may cause the cutterhead to actually left completely clear of the lacquer master, or (worse) cut through the lacquer film into the aluminum substrate, which ruins the cutting tool and the lacquer master. For this reason, many records are mono in the low frequencies, which guarantees a relatively stable groove depth.

While it's quite possible to record a wide range signal, the extreme quiet could cause manufacturing problems because the pressing plant may not be able to maintain the required standards. Furthermore, it's entirely possible to cut a record where the groove velocities are so high that the record is unplayable.

Flat? What's that?

Compact Disk Mastering

From an audio standpoint, the compact disk (CD) presents an almost perfect audio storage medium. For the audio consumer, a CD is the closest thing that they will ever hear outside of the studio where the mixdown session occurred.

The CD presents its own set of problems. While the mechanical problems of phonograph record mastering or playback aren't present, the CD opens new cans of audio worms.

- The CD is ruthless. It will expose every flaw and imperfection in the master tape. Tape hiss, sloppy punch-ins, noise chairs, human-based noise pollution are clearly audible. Listen to a CD of a symphony orchestra on headphones; listen for squeaky chairs and body noises emanating from the musicians...after all, they're only human.

- The CD is capable of around 90dB dynamic range. While this seems like a lot, the problem is that this number is absolute. In analog recording, the upper and lower limits are somewhat poorly defined because most analog mediums gradually distort at the upper limit, and noise becomes the limitation at the lower limit. In digital recording, you can't record a signal that is below the lower signal-level limit of the equipment. This means that extremely low level signals like ambiance and reverberation will not be handled on an exact one-for-one basis.

- Record companies exist for profit. Most recordings today are still produced with the limitations of analog storage and transmission systems in mind. It's unreasonable to expect record companies to produce separate masters for different storage media.

A Brief Anatomy of the Classical Recording Process

Two wildly divergent techniques are usually used in recording classical music. These are: two (sometimes slightly more) microphones into a two channel tape machine or lots of microphones into a multitrack tape machine. Figure 6 shows two different minimalist techniques.

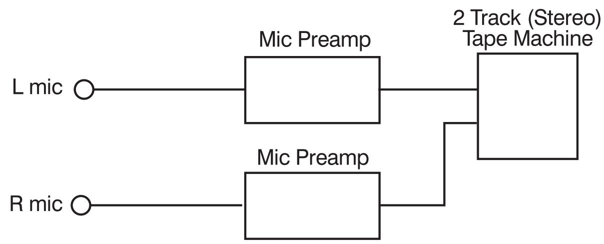

Purist, 2 Mic Approach

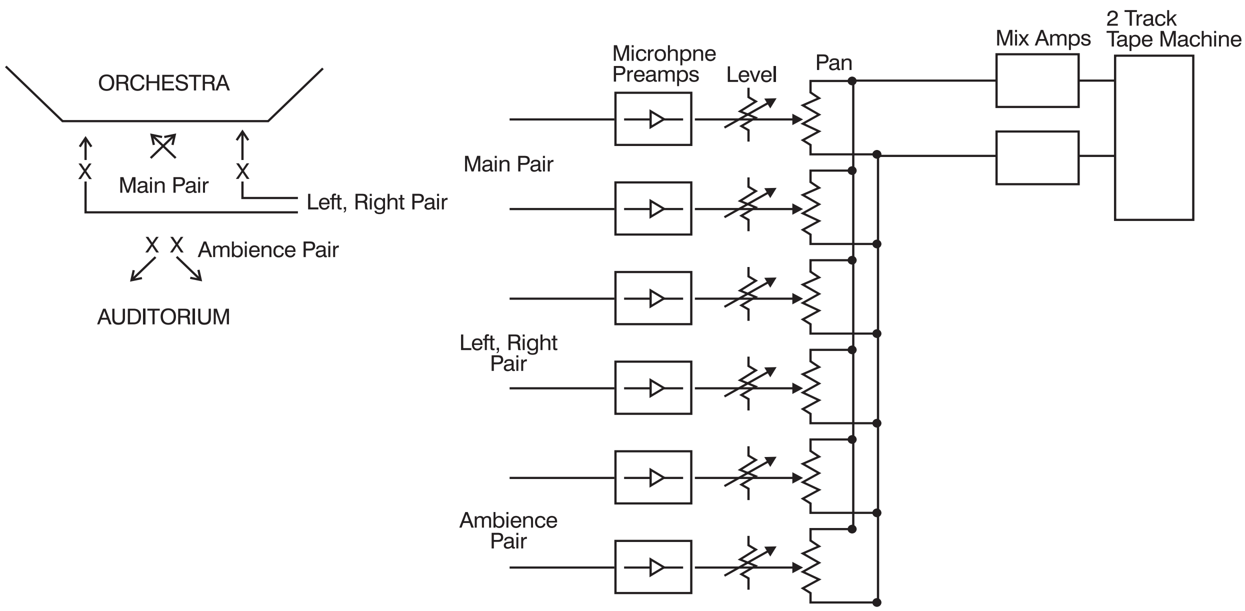

Classical Setup
(Eargle for Seattle Symphony Recordings)

Figure 6. Two Approaches to Classical Recording.

If you use two microphones, aside from placing them at the right place in the hall and setting the controls on the tape machine properly, it's totally up to the musicians and the conductor.

The second setup shown in Figure 6 is one used by John Eargle for some of his classical music recordings on the Delos label. Here he uses a primary stereo pair, supplemented by two additional microphones set up to the left and right of the primary microphones. An additional stereo pair set up towards the rear of the auditorium, and aimed towards the rear of the auditorium captures the natural reverberation of the room. The microphones are all mixed directly to a two-channel digital recorder, without equalization.

If you use lots of microphones and a multitrack, it's still up to the musicians and conductor, but everyone gets a second chance at the mixdown session. The purists in the crowd may also retch.

The ubiquitous equalizer once again presents the opportunities for sonic enhancement. Push up the 3K region to give the strings a little bite, pull back the 5K region to make them shimmer. Remove some 10K from the brass to remove the edge. Boost the 5K region to make the percussion crash better. Add some 50Hz to the basses to make them sound fuller.

Even if you use two microphones, it still probably isn't flat. Like a loudspeaker, a microphone is a complex electro-mechanical system. Each different microphone has its own individual sound. Engineers like to pick different microphones for different instruments, so the characteristics of the microphone can complement the sound of the instrument. Thus, microphone selection and choice can be viewed as the first step in the equalization process.

Maybe it's flat. Sorta.

Why You Can't Really Expect All Recordings to Sound Good On All Loudspeakers

At each step in the recording process, the sound is evaluated on loudspeakers. When basic tracks are recorded, the goal of the control room monitor speakers is brutal, unerring accuracy. The monitor speakers must ruthlessly expose every shortcoming in the quality of the reproduced material.

These speakers are the studio's reference standard for "flat" response. It's important to note that the term "flat" does not mean unchanging frequency response from 20Hz to 20KHz. Even if the monitors have been equalized, they still aren't flat. What they are is flat to about 8KHz with a slow, controlled rolloff beyond. If they were equalized to be ruler flat to 20KHz material that was judged on them would sound dull and lifeless on "average" speakers. Needless to say, the average studio monitor is a very distant cousin of the average home loudspeaker.

During mixdown, the speakers used are more typical of what the finished product will be listened to on. Typically many pairs of different speakers are used during the mixdown process to be as sure as possible that the finished product will sound acceptable on the widest possible selection of equipment. A popular extreme is a medium quality four-inch loudspeaker in a small enclosure, the Auratone 5C. These can be found in over 80% of all US studios. They do a pretty good job of simulating a small radio or automobile speaker. They don't have much low frequency output below 150Hz.

One studio has a small AM radio transmitter. They play your tape over it while you drove around the block listening to your next hit on your car radio! Of course, the speakers in your car are anything but flat.

Finally, during mastering, the product is listened to again. Again, loudspeakers are used for listening. Probably different ones at that.

The loudspeaker's tonal characteristics have a reverse effect on the finished product. If the speaker is boomy, the engineer may tend to mix a trifle shy on the bass. If the speaker yells (has a mid-range peak), the finished product may be deficient in the midrange (coincidentally at the same place where the speaker had its peak). At any rate, if the monitor speaker is not flat, it will be reflected in the finished product in reverse fashion. You ask, "Why don't they just use flat speakers to monitor on, and we'll do the same for listening?"

To make a long story short, the product must be mixed on non-flat loudspeakers so that the production team (artist, engineer and producer) does not lose sight (or sound) of what the fruits of their labors are really going to be listened to on. It takes careful balance and equalization if they want the listener to even start to hear the bass guitar when he's listening to it on a clock radio. If you try to listen to this highly processed mix on flat (inherently flat or equalized flat) speakers, it will sound bright and tinny, with no bass response.

It's really hard to be flat.

What you can do about it

With all this tomfoolery going on in the recording process, it's truly a wonder that this stuff sounds good at all when it's played at home.

Because of the differences in the characteristics of the monitor speakers used when the record was recorded, the recording must be played back on an identical system to that which it was mixed on, preferably in an identical environment. If not, you have no idea of what the engineer heard when he (or she) performed the final mix. Considering that there are maybe a half dozen "favorite" small mixdown speakers, this is hardly practical. As you can see, mixing for records is a study in calculated guesswork. Fortunately, a good mix sounds good on everything, and a good mix on great equipment sounds awesome.

No-Cost Solutions

There aren't many no-cost solutions. The only ones are judicious use of the tone controls provided on the preamp or receiver. Careful use of any filters (hi or lo filters) may help tame a boomy low end or peaky high end. Adjusting the speaker placement relative to the walls and floors may help even out the low frequency response. Last, some careful readjustment of furniture may help to reduce wall-to-wall or floor-to-ceiling reflections.

If you are dissatisfied with the tonal quality of your present speakers, try adjusting the high and/or mid frequency level controls on the speakers. Use a record that you are very familiar with that has good extended lows and highs. Pick another that has vocals on it.

With the tone controls on your system set flat, adjust the controls on the speakers for the most pleasing sound. The highs should be present, without being shrill, the midrange should make voices sound natural, rather than sounding shrill or nasal. Many people adjust the tweeter controls until they can hear the tweeters working, which is usually too loud. My trick is to place my hand in front, blocking the tweeter and listening for the difference.

Cost Solutions

If you are willing to spend some money, there are several solutions on the horizon. Here are a few.

New Speakers

Naturally, new speakers are worth considering. The loudspeaker/room interface is a critical part of any music reproduction system. It's entirely possible for a speaker to sound great in the dealer's showroom and to be a disappointment when you get it home.

Room Acoustics

If your listening room is an acoustical disaster, you can try to alter the absorption characteristics somewhat by adding or removing sound-absorbing objects (curtains, sofas, rugs, tapestries, etc.) or by moving the loudspeakers to take advantage of a particular acoustical aspect of your listening room. Even a good room has a large effect on the sound of your system.

Corner placement will help a speaker that is bass shy. Moving a speaker that is bass heavy out of the corner and/or off the floor is another possibility. Careful equalization with an equalizer may help to minimize a bad standing wave. Equalization will not help to minimize room reflections (like those caused by parallel walls and/or ceilings). If this isn't successful, the only real solution is to dynamite the room and start over. Short of dynamite, planing with the room by rearranging the speakers and furnishings can improve things a great deal.

Signal Processing

If you're willing to use electronics, there's a whole world of possibilities.

As explained earlier, there isn't a lot of low frequency energy on most commercial recordings. This is even true for many compact discs. If you really want to hear the low frequencies, consider bass restoration, a technique that recovers the lost bass information. Oddly enough, speakers without extended bass response can still benefit from this technique.

Last, but not least, is equalization. If your tapes, records and CDs don't sound right, for whatever reason, chances are that careful equalization with a more sophisticated equalizer than the standard bass and treble controls on your receiver can go a long way toward enhancing your listening pleasure.

Careful equalization can help take a bad peak out of a marginal speaker. It can also help you to "voice" your speakers so that they sound more like you want them to, rather than what the designer thought they should. You can also use the equalizer to alter the instrumental balance achieved in the mixdown session to a degree. Equalizing can help if you can't quite hear the bass, the drums sound remotely like oatmeal boxes, or if the singer seems a little too far back in the mix.

For most applications both audio channels should be adjusted in equal amounts. An equalizer with the controls for each channel interleaved (left and right channel controls for each frequency are adjacent to each other) make this much easier. Equal adjustment for each channel helps to preserve the channel-to-channel balance and helps to keep the stereo image stable. Keeping the two channels separate, however, still allows dissimilar channel equalization when made necessary by a marginal recording or extreme circumstances.

Conclusions

For most people, a truly flat audio system would probably not sound good on the majority of source material available to us today. Considering the pre-equalization that occurs before you, the consumer, actually receive the product, it is a shame to limit your options by refusing to admit that perhaps some carefully applied corrective equalization or signal processing could enhance your listening experience.

If you consider the inverse effect of the voicing of the final mixdown monitor speakers on the finished product and the number of different speakers used for that purpose, the final tonal balance is at best a calculated guess. It's a credit to everyone involved that the finished product sounds as good as it does on such a wide range of equipment. For these reasons, an equalizer may be able to improve your "listening quotient" from very good to excellent.

While an equalizer can assist in taming a particularly offensive room resonance, it is no substitute for good acoustical design of the room and sensible placement of absorbing objects within the listening environment. It is, however, a viable means of modifying the voicing of a particular loudspeaker system and improving the interface between the speaker and the acoustical environment.

As you can see, the concept of "flat response" is somewhat altruistic. Very few commercially available recordings are actually made with flat microphones through unequalized recording channels. While some audiophiles will continue their quest for sonic nirvana and flat response, the rest of us will learn to use the tools at our disposal to enjoy our music more. Why not? I do!