

THE SECRET OF THE MASTERING ENGINEER

BY BOB KATZ

ADVANCED USE

TABLE OF CONTENTS

INTRODUCTION

<i>Table of contents</i>	3
<i>Welcome</i>	5

TUTORIAL

<i>Getting Started</i>	6
<i>Your Room Your Monitors</i>	7
<i>Metering</i>	8
<i>Dynamics Processing</i>	9
<i>Sequencing:</i>	
<i>Relative Levels, Loudness</i> <i>and Normalization</i>	10
<i>Recipe for Radio Success</i>	11
<i>Dither</i>	13
<i>Equalization</i>	13
<i>Sibilance Control</i>	15
<i>Noise reduction</i>	15
<i>Monitors</i>	16
<i>Advanced Mastering Techniques</i>	17

APPENDIX

<i>Reference</i>	19
<i>Glossary</i>	19

GETTING STARTED

Mastering vs. Mixing

Mastering requires an entirely different "head" than mixing. I once had an assistant who was a great mix engineer and who wanted to get into mastering. So I left her alone to equalize a rock album. After three hours, she was still working on the snare drum, which didn't have enough "crack"! But as soon as I walked in the room, I could hear something was wrong with the vocal. Which brings us to the first principle of mastering: Every action affects everything. Even touching the low bass affects the perception of the extreme highs.

Mastering is the art of compromise; knowing what's possible and impossible, and making decisions about what's most important in the music. When you work on the bass drum, you'll affect the bass for sure, sometimes for the better, sometimes worse. If the bass drum is light, you may be able to fix it by "getting under the bass" at somewhere under 60 Hz, with careful, selective equalization. You may be able to counteract a problem in the bass instrument by dipping around 80, 90, 100; but this can affect the low end of the vocal or the piano or the guitar - be on the lookout for such interactions. Sometimes you can't tell if a problem can be fixed until you try; don't promise your client miracles. Experience is the best teacher.

Think Holistically

Before mastering, listen carefully to the performance, the message of the music. In many music genres, the vocals message is most important. In other styles, it's the rhythm, in some it's intended distortion, and so on. With rhythmic music, ask yourself, "what can I do to make this music more exciting?" With ballads, ask "is this music about intimacy, space, depth, emotion, charisma, or all of the above"? Ask, "How can I help this music to communicate better to the audience?" Always start by learning the emotion and the message of the client's music. After that, you can break it down into details such as the high frequencies, or the low frequencies, but relate your decisions to the intended message of the music.

Some clients send a "pseudo-mastered" demonstration CD illustrating their goals. Even if you don't like the sound on their reference, or you think you can do better, carefully study the virtues of what they've been listening to. During your mastering, refer back to the original mix; make sure you haven't "fixed" what wasn't broken in the first place. There is no "one-size-fits-all" setting, and each song should be approached from scratch. In other words, when switching to a new song, bypass all processors, and listen to the new song in its naked glory to confirm it needs to be taken in the same or different direction than the previous number. Likewise, as you gain experience, you may want to "tweak" the "presets" in your equipment. Presets are designed to make suggestions and provide good starting points, but they are not one-size-fits-all and should be adjusted according to the program material and your personal taste. The Wizard function of the TC dynamics processors may be a better way of establishing a starting point than a static preset. Listen carefully to the results of the Wizard and study what it has done and how, then, let your ears be the final musical judge.

When you ask your spouse to turn the TV down a touch - do you mean only 1 dB?

YOUR ROOM - YOUR MONITORS

Very few recording studios are suitable for mastering. For optimal mastering, use a different room than your recording studio or room. The typical recording control room has noisy fans, a large console and acoustic obstacles that interfere with evaluation of sound.

With few exceptions, you won't find near-field monitors in a professional mastering room. No little speakers, cheap speakers, alternative monitors. Instead, there's a single set of high quality loudspeakers. The loudspeaker-room interface in a mastering room is highly refined, and the mastering engineer tuned into their sound, so that he/she knows how the sound will translate over a wide variety of systems.

What's wrong with near-field monitors?

Near-field monitoring was devised to overcome the interference of poor control-room acoustics, but it's far from perfect. In many control rooms, with large consoles and rack gear, the sound from the "ideal" big speakers bounces from these surfaces, producing inferior quality. Reflections from the back of the console are often neglected. Even with absorptive treatment, you can't defeat the laws of physics; some wavelengths are going to reflect. But near-field monitors mounted on console meter bridges are not necessarily cures. Nearby surfaces, especially the console itself, cause comb filtering, peaks and dips in frequency response.

The mix engineer may try to compensate for problems which are really caused by monitoring acoustics; resulting in recordings with boomy or weak bass, peaks or dips (suckouts) in the midrange, thumpy bass drums, and so on.

Sound travels over more than one path from the loudspeaker to your ears - the direct path, and one or more reflected paths, especially the bounce off the console. That reflected path is so problematic that it's almost impossible to locate nearfield monitors without breaking a fundamental acoustic rule: The length of the reflected signal path to the ears should be at least 2 to 3 times the direct signal path.

Very few near-field monitors pass the "bandwidth and compression test". Almost none have sufficient low frequency response to judge bass and subsonic problems, and very few can tolerate the instantaneous transients and power levels of music without monitor compression. If your monitors are already compressing, how can you judge your own use of compression? Near-field monitoring also exaggerates the amount of reverberation and left-right separation in a recording. Clients are often surprised to learn their singer has far less reverb than they had thought, and the sound less stereophonic when they hear the recording played with more-normal monitoring. Yes, the best mix engineers have learned how to work with near-field monitors and mentally compensate for their weaknesses, but these same mix engineers know better than to master in that environment. There's no excuse for monitor weakness in a mastering room.

Subwoofers

Subwoofers, or prime loudspeakers with infrasonic response, are essential for a good mastering studio. Vocal P pops, subway rumble, microphone vibrations, and other distortions will be missed without subwoofers, not just the lowest notes of the bass. Proper subwoofer setup requires knowledge and specialized equipment. I've been in too many studios where the subs are inaccurately adjusted, usually "too hot", in a vain attempt to impress the client. But the results won't translate when the subs are incorrectly adjusted.

Room Acoustics

Whether your loudspeakers are mounted in soffits, or in "free space", a properly designed room must have no interfering surfaces between the loudspeakers and your ears. Secondary reflections will be carefully controlled, and the dimensions of the room and solidity of the walls defined. A good mastering room should be at least 20 feet long, preferably 30 feet, and the monitors, if not in soffits, anchored to the floor, and placed several feet from walls and corners. There's obviously a lot more to this part of the story, but the bottom line is to get an acoustic consultant unless you really know what you're doing.

Monitor Translation

Mastering engineers have long ago learned that the widest-range, most accurate loudspeakers translate to the widest variety of alternate playback systems. If you follow all of the above in your mastering room, your masters will translate to the majority of systems out there. Good mastering engineers hit the mark the first time, better than 7 times out of 10.

Monitoring Levels and Fletcher-Munson

There is a scientific reason for not monitoring too loudly. The Fletcher-Munson equal loudness contours reveal that the human ear does not have a linear response to bass energy. The louder you monitor, you can be fooled into thinking a program has more bass energy. Thus it is extremely important to monitor at approximately the same level as the ultimate listener to your recording. No matter how good your monitors, if you turn them up too far, then you will put too little bass into the program, and vice versa.

When you go to a concert, do you identify an 80 Hz resonance under the third balcony?

METERING

Truth in Metering

1999 marks the 60th anniversary of the VU Meter standard, yet many people still don't know how to read a VU! Despite all its defects, the VU meter has survived because it works. The VU meter, with its 300 millisecond averaging time constant, is closer to the loudness-sensing of the human ear, while sample accurate peak-reading meters tell you nothing except whether the capabilities of the digital medium are being exceeded.

Two different programs, both reaching 0 dBFS on the peak meter, can sound 10 dB (or more) apart in loudness! This makes an averaging meter an essential supplement to the mastering engineer's ears. Some meters have dual scales, displaying both average and peak.

While mixing or mastering, use the average meter and glance at the peak meter. For popular music mastering, here's a conservative calibration setting that will help to produce masters in a similar ballpark to the best-sounding CDs ever made: With sine wave tone at -14 dBFS, adjust the averaging meter to read ZERO. If the averaging meter reaches 0 on typical musical peaks, and occasionally +3 or +4 on extreme sustained peaks, you're probably right in the ballpark. Every decibel of increased average level means that considerably more than 1 dB additional compression has been applied; which might or might not be the perfect thing for your kind of music. Listen and decide.

The Ear Is The Final Judge

Wide dynamic range material, such as classical music, folk music, some jazz and other styles are often mastered without any dynamics processing at all. In such cases, you may find the averaging meter reading well below 0. This is probably not a problem as long as the music sounds proper to the ears. Some mastering engineers working with wide-range music recalibrate their averaging meters to -20 dBFS = 0 VU, or else recognize that the averaging meter may read well below 0 VU with such music. Also realize that meters are generally not frequency-sensitive, but the human ear judges loudness by frequency distribution as well as level. Thus, two different programs reaching 0 VU (average) may have different loudness.

Quasi-Peak Meters and Judgment of Quality

The ear is the final arbiter of quality, but meters can help. The VU helps demonstrate if average levels are too hot, but as I've described, it requires interpretation. An objective measure of quality is to measure transient loss - to see if audible peaks are reduced. The ear has a certain "rise time"; we probably can't hear the difference between a 10 millisecond transient and a 10 microsecond transient.

The digital **Peak Program Meter** is too fast; it measures inaudible (short duration) peaks as well as audible ones.

A popular meter for detecting audible peaks is a quasi-peak meter, or analog PPM, defined by an EBU standard. It's usually made with analog circuitry, but can also be constructed with digital circuits. This meter's 10 millisecond integration time is much slower than the 22 microseconds of the sample-accurate digital PPM. Short overloads, or short bursts of limiting can be inaudible, as long as the level on the quasi-peak meter does not drop.

Peaks shorter than about 10 ms can usually be limited without audible penalty. Wide range program material with a true peak to average ratio of 18 to 20 dB can be transparently reduced to about 14 dB. That's one of the reasons 30 IPS analog tape is desirable, as it performs this job very well.

The Finalizer can also do this job, with the aid of a quasi-peak meter to verify the audible peak level is not coming down, and/or the VU meter to see if a 14 dB peak/average level is obtained. A rule of thumb is that short duration transients of unprocessed digital sources can be transparently reduced by 4 to 6 dB; however, this cannot be done with analog tape sources, which already have removed the short duration transients. Any further transient reduction (e.g., compression/limiting) will not be transparent but may still be esthetically acceptable or even desirable.

Over Counters and Increased Level

0 dBFS (FS=full scale) is the highest level that can be encoded. Most mastering engineers have discovered that you can often hit 0 dBFS on a digital PPM without hearing any distortion. In fact, a single peak to 0 dBFS is not defined as an over level. Over levels are measured with over counters. Conventional wisdom says that if three samples in a row reach 0 dBFS, then an overload must have occurred somewhere between the first and third sample. In an A/D converter, even if the source analog voltage exceeds 0 dBFS, the end result is a straight line at 0 dBFS. However, the ear forgives certain overloads. Note that a 3 to 6 sample over will often be inaudible with drums or percussion, but the ears may hear distortion with only a 1-sample over with piano material. The original Sony digital meter established the standard of 3 contiguous samples equals an over, but has a dip switch to indicate 1-sample overs. Some engineers conservatively use the 1-sample standard, but I've had no problems with a set of good ears and a 3-sample over counter. You can often raise gain by 2 or more dB without having to limit or compress, when you trust the over counter and your ears, instead of a digital PPM.

DYNAMICS PROCESSING

Dynamics Processing

Both compression and limiting change the peak to average ratio of music, and both tools reduce dynamic range.

Compression

Compression changes sound much more than limiting. Think of compression as a tool to change the inner dynamics of music. While reducing dynamic range, it can "beef up" or "punch" low level and mid-level passages to make a stronger musical message.

Limiting

Limiting is an interesting tool. With fast enough attack time (1 or 2 samples), and fairly fast release (1 to 3 milliseconds), even several dB of limiting can be transparent to the ear.

- Consider limiting when you want to raise the apparent loudness of material without severely affecting its sound;
- consider compression when the material seems to lack punch or strength.

Remember, the position of your monitor volume control has a tremendous effect on these matters of judgment. If it sounds properly "punchy" when you turn up the monitor, then maybe all you need is to turn up the volume rather than add another DSP process!

If the music sounds adequately punchy, yet high levels are not approaching ZERO (reference -14 dBFS) on a VU meter, then consider limiting to raise the average level without significantly changing the sound.

Equal-Loudness Comparisons

Since loudness has such an effect on judgment, it is very important to make comparisons at equal apparent loudness. The processed version may seem to sound better only because it is louder. That's what makes the Finalizer's unique "matched compare" system so important. Adjust the gain so that there is no apparent change in loudness when the processing is bypassed. This puts everything on a "level" playing field. You may be surprised to discover that the processing is making the sound worse, and it was all an illusion of loudness. If the sound quality is about the same, then you have to decide if you really need the loudness gain. Don't join the loudness "race" (which has no winners); make an informed, not arbitrary, decision.

To judge the absolute loudness of the Finalizer, you need average metering, and a calibrated monitor. See the appendix for references on calibrated monitor and metering systems.

Manipulating Dynamics: Creating the Impact of Music

Consider this rhythmic passage, representing a piece of modern pop music:

shooby dooby doo WOP...shooby dooby doo WOP...shooby dooby doo WOP

The accent point in this rhythm comes on the backbeat (WOP), often a snare drum hit. If you strongly compress this music piece, it might change to:

SHOOBY DOOBY DOO WOP...SHOOBY DOOBY DOO WOP...SHOOBY DOOBY DOO WOP

This completely removes the accent feel from the music, which is probably counterproductive.

A light amount of compression might accomplish this...

shooby dooby doo WOP...**shooby** dooby doo WOP...**shooby** dooby doo WOP

...which could be just what the doctor ordered for this music.

Strengthening the subaccents may give the music even more interest. But just like the TV weatherperson who puts an accent on the wrong syllable because they've been taught to "punch" every sentence ("The weather FOR tomorrow will be cloudy"), it's wrong to go against the natural dynamics of music. Unless you're trying for a special effect, and purposely creating an abstract composition. Much of hip hop music, for example, is intentionally abstract....anything goes, including any resemblance to the natural attacks and decays of musical instruments.

- Back to "Shooby doo". This kind of manipulation can only be accomplished with careful adjustment of thresholds, compressor attack and release times.

If the attack time is too short, the snare drum's initial transient could be softened, losing the main accent and defeating the whole purpose of the compression.

If the release time is too long, then the compressor won't recover fast enough from the gain reduction of the main accent to bring up the subaccent. If the release time is too fast, the sound will begin to distort.

If the combination of attack and release time is not ideal for the rhythm of the music, the sound will be "squashed", louder than the source, but "wimpy loud" instead of "punchy loud". It's a delicate process, requiring time, experience, skill, and an excellent monitor system.

DYNAMICS PROCESSING

Here's a trick for compressor adjustment:

Find the approximate threshold first, with a fairly high ratio and fast release time. Make sure the gain reduction meter bounces as the "syllables" you want to affect pass by. Then reduce the ratio to very low and put the release time to about 250 ms to start. From then on, it's a matter of fine tuning attack, release and ratio, with possibly a readjustment of the threshold. The object is to put the threshold in between the lower and higher dynamics, so there is a constant alternation between high and low (or no) compression with the music. Too low a threshold will defeat the purpose, which is to differentiate the "syllables" of the music.

Note! With too low a threshold and too high a ratio EVERYTHING WILL BE BROUGHT UP TO A CONSTANT LEVEL.

Multiband processing can help in this process

Transients (percussive sounds) contain more high frequency energy than continuous sounds. By using gentler compression or no compression at high frequencies (e.g., higher threshold, lower ratio), you can let the transients through while still "punching" the sustain of the subaccents or the continuous sounds. Practice by listening to the impact of the percussion as you change compressor attack times. With care, you can have punch and impact, too. But with overcompression, or improperly-adjusted compression, you may get the punch, but lose the transient impact. Most music needs a little of both.

Multiband compression also permits you to bring out certain elements that appear to be weak in the mix, such as the bass or bass drum, the vocal or guitars, or the snare, literally changing the mix. Learn to identify the frequency ranges of music so you can choose the best crossover frequencies.

Compression, Stereo Image, and Depth

Compression brings up the inner voices in musical material. Instruments that were in the back of the ensemble are brought forward, and the ambience, depth, width, and space are degraded. But not every instrument should be "up front". Pay attention to these effects when you compare processed vs. unprocessed. Variety is the spice of life. Make sure your "cure" isn't worse than the "disease".

At the ballet, do you notice the music first, before the dancers?

SEQUENCING

Relative Levels, Loudness, and Normalization

Sequencing an album requires adjustment of the levels of each tune. We've seen that the ear judges loudness by the average, not peak levels of the music. We've also seen that compression and limiting change the loudness of the music by changing the peak to average ratio. Normalization is the process of finding the highest peak, and raising the gain until it reaches 0 dBFS. But do not use normalization to adjust the relative loudness of tunes, or you will end up with nonsense. The ear is the final arbiter of relative loudness of tunes. But the ear can be fooled, it's better at making relative than absolute judgments. We've all had the experience of mixing at night, and returning in the morning to find everything sounds much louder! So don't make your judgments by "needle drops"...play the end of each tune going into the beginning of the next. It's the only way.

Do you know when the musicians are out of tune ?

RECIPE FOR RADIO SUCCESS

The Myth of Radio-Ready

Advertisements are created by marketing people, whose goal is to sell products, and often use ambiguous terms. The most ambiguous of those terms is **radio ready**. Be an aware consumer. Radio is the great leveler. It will take songs that sound very soft and "unpunchy", and bring them up to compete with the hottest recordings; it will take songs that are extremely hot and processed, and squash them down in a very unpleasant manner. In other words, mastering with overzealous dynamics processing can actually make a record sound bad on the radio, or at least, not as good as properly-prepared competition.

I discovered this fact at least 12 years ago, when I found that my audiophile recordings made with absolutely no compression or limiting were competing very well on the radio against heavily-processed recordings! Radio engineers will confirm this fact: Almost no special preparation is required to make a recording radio ready.

The Music Always Comes First

1 Write a great original song, use fabulous singers, and wonderful arrangements. Be innovative, not too imitative (if you can get past the "format" censors, your innovative music will attract attention).

2 Sparse, light arrangements often work better than dense, complex ones, because the dynamics processing on the radio can turn dense arrangements into "mush". When you examine the apparent exceptions (e.g., Phil Spector's wall of sound), the main vocal element is always mixed well above "the wall".

The Sound Comes Second

3 Radio Ready does not mean "make it sound like it's already on the radio".

4 Make sure your music sounds good, clean and dynamic at home and in the studio. That will guarantee it will sound good on the radio.

5 Many people are not familiar with good sound production and reproduction. First you must have a background, an ear education. Don't imitate the sound that you hear on the radio speaker. Compare your music to good recordings, auditioned on the best possible sound system. And don't forget the ultimate reference: the dynamic sound of live, unamplified music performed in a concert hall. There's also evidence that prolonged exposure to loud music is causing hearing loss in an entire generation of our children. This leads to a preference for

compressed sound, because dynamics bother the hearing-impaired. This, in turn, leads to a vicious cycle of even more loudness and further hearing loss. Do you hear me?

Preparing For The Radio

6 Peak to average ratio is the difference between the level on an averaging meter, such as a VU meter, and the peak level of the music as read on a PPM. A meter which displays both peak and average on the same scale is most desirable, otherwise, you have to do some arithmetic, and look at two meters at once. If the dual-function meter reads -17 dBFS average, and -6 dBFS peak during some short musical passage, then your music has approximately an 11 dB peak to average ratio. Choose high peak to average ratio (14 dB or more) or low peak to average ratio (less than 14 dB) according to the sound you are trying to create at home-in general, without fear of how that will translate on the radio. If lowered peak to average ratio is part of your creative sound, it will translate on the radio, unless your processing was so severe that the average level becomes high enough to cause the radio processors to bring your music down (squash it). Avoid the danger zone, anything less than 6 dB peak to average ratio is dangerous, since radio processors are designed to try to maintain an average level, and they literally clamp material with too high an average level (material that would "pin" an ordinary VU meter). That material will probably sound worse on the radio than your competition with a larger peak to average ratio.

Think of your dynamics processor as a tool to help create your sound, not to be used for radio ready. The more compressed your material, the less the transient impact of the drums, clarity of the vocal syllables, and percussion. Sometimes that's esthetically desirable, but often it is displeasing, depending on the type of music. Use a wide range, uncompressed monitoring system, to help decide which choice is best for your music. Compressors have always been used for effect in music production, and sometimes misused, from the 50's through the 90's. The newly-invented digital compressors are far more powerful than the old analog versions. Entirely new effects can be created, and some of today's hit records are even based on those effects. But watch out when you step on the gas of that Formula One Race Car! I feel that many rock CDs made in 1991 (before the popularity of powerful digital processors) sound better than major releases made in 1998. Only you have control over your sound; there's no official speed limit, and no policeman to revoke your driver's license, even though engineers are crashing all over the place.

RECIPE FOR RADIO SUCCESS

7 Subsonics

Excessive subsonics can drain unnecessary energy away from the total loudness. In addition, excessive subsonic material can cause radio compressors to pump or be exercised unnecessarily. Check for excess subsonic energy in several ways: By looking - with a real time analyzer, by listening - with a pair of subwoofers, and by testing: If you are confident in the calibration of your subwoofers, test if the subsonics are musically meaningful by comparing the sound with and without a high-pass filter. If the sound gets clearer with the filter in, and you hear no losses in musical information, then use the filter on the program. Ironically, bass instruments (especially direct boxes) sometimes sound clearer when filtered below 40 Hz. But use your ears; don't extend this advice to the general case, and don't make this critical judgment with inferior monitors.

8 Excessive Sibilance

The combination of FM Radio's 75 microsecond pre-emphasis and poor sibilance controllers at the radio station can make a bad broadcast. Better to control excessive sibilance in the mastering. My definition of excessive sibilance is that which would be annoying on a bright playback system.

9 Excessive peak percussion

This problem is rare. Be aware of how radio processing reacts to percussive music. Watch out for a repetitive rhythmic transient that's many dB above the average level of the rest of the music, e.g. very sharp timbale hits with peaks at least 8 dB above the average vocal level. Radio processing, with its slow release times, can bring the vocal down severely with each timbale hit, and render the vocal (and all the background) inaudible for seconds at a time.

Ideally, fix this problem in the mix, not in the mastering. Proper mix techniques, with selective processing, can keep this situation under control. Of course, if you can no longer fix it in the mix, then careful application of the Finalizer's multi-band dynamics module will cure the problem without destroying the music's percussive nature. Just remember, this is a very rare situation that should be repaired with conservative, experienced ears, or your music will be ruined. Overcompression can ruin that beautiful percussive sound.

Loudness and the radio

Subtle multi-band compression and soft clipping can make you appear louder on the radio. If you feel this compromises the sound of the CD when played on the home system, why not make a special compressed single just for radio release. This gives you the best of both worlds. But remember, if you make the average level too high, it may trigger the radio processors to drop the level of your precious song.

Do you know what comb filtering is?

DITHER

Wordlengths and Dithering

Dither is perhaps the most difficult concept for audio engineers to grasp. If this were a 24-bit world, with perfect 24-bit converters and 24-bit storage devices, there would be much less need for dither, and most of the dithering would go on behind the scenes. But until then (and the audio world is heading in that direction), you must apply dither whenever wordlength is reduced. The fine details of dithering are beyond the scope of this booklet; learn more by consulting the references in the appendix.

Here are some basic rules and examples:

- 1 When reducing wordlength you must add dither.
Example: From a 24-bit processor to a 16-bit DAT.
- 2 Avoid dithering to 16 bits more than once on any project.
Example: Use 24-bit intermediate storage, do not store intermediate work on 16-bit recorders.
- 3 Wordlength increases with almost any DSP calculation.
Example: The outputs of digital recording consoles and processors like the Finalizer will be 24-bit even if you start with a 16-bit DAT or 16-bit multitrack.
- 4 Every "flavor" of dither and noise-shaping type sounds different. It is necessary to audition any "flavor" of dither to determine which is more appropriate for a given type of music.
- 5 When bouncing tracks with a digital console to a digital multitrack, dither the mix bus to the wordlength of the multitrack. If the multitrack is 16-bit digital, then you're violating rule #2 above, so try to avoid bounces unless you have a 20-bit (or better) digital multitrack.
Example: You have four tracks of guitars on tracks 5 through 8, which you want to bounce in stereo to tracks 9 and 10. You have a 20-bit digital multitrack. You must dither the console outputs 9/10 to 20 bits. If you want to insert a processor (like the Finalizer) directly patched to tracks 9 and 10, don't dither the console, just dither the Finalizer to 20 bits. The Finalizer's ADAT interface makes this process relatively painless. One complication: The ADAT chips on certain console interface cards are limited to only 20 bits. Consult your console manufacturer. Although the Finalizer's ADAT interface carries a true 24 bits, if the console's ADAT interface is limited to 20, you need to dither the console feed to the Finalizer to 20 bits and once again dither the Finalizer output to 20 bits to feed the multitrack!

EQUALIZATION

What Is An Accurate Tone Balance

Perhaps the prime reason clients come to us is to verify and obtain an accurate tonal balance. The output of the major mastering studios is remarkably consistent, pointing to very accurate studio monitoring.

As I've pointed out, the goals of equalization in mastering are generally different from equalization in mixing. It is possible to help certain instruments (particularly the bass, bass drum and cymbals), but most of the time the goals are to produce a good spectral balance. What is a "good" tonal balance? The ear fancies the tonality of a symphony orchestra. On a 1/3 octave analyzer, the symphony always shows a gradual high frequency roll-off, as will most good pop music masters.

Everything starts with the midrange. If the mid-frequency range is lacking in a rock recording, it's just like leaving the violas or the woodwinds out of the symphony. The fundamentals of the vocal, guitar, piano and other instruments must be correct, or nothing else can be made right.

Specialized Music Genres

It's one thing to understand the symphony, and another to properly balance all the different music genres. The bass plays very different roles in each popular music genre. You could think of Reggae as a symphony with lots more bass instruments, but let's not get too hung up on the symphony analogy. Just remember to keep the symphony balance in your head as a basic reference, especially in the mid to high frequency balance.

EQ Tricks

Remember the yin and the yang: Contrasting ranges have an interactive effect. For example, a slight dip in the lower midrange (~250 Hz) can have a similar effect as a boost in the presence range (~5 kHz). Harshness in the upper midrange/lower highs can be combatted in several ways. For example, a harsh-sounding trumpet-section can be improved by dipping around 6-8 kHz, and/or by boosting circa 250 Hz. Either way produces a warmer presentation. The next trick is to restore the sense of air which can be lost by even a 1/2 dB cut at 7 kHz, and this can often be accomplished by raising the 15 to 20 kHz range, often only 1/4 dB can do the trick. Remember the interactivity of the frequency ranges; if you make a change in any of them, you must reevaluate your choices on them all.

When you go to concerts, do you sometimes think you hear edits?

EQUALIZATION


High Q or low?

Gentle equalizer slopes almost always sound more natural than sharp. Q's of 0.6 to 0.8 are therefore very popular. Use the higher (sharper) Q's (greater than 2) when you need to be surgical, such as dealing with narrow-band bass resonances or high-frequency noises. The classic technique for finding a resonance is to start with a large boost (instead of a cut) to exaggerate the unwanted resonance, and fairly wide Q, then sweep through the frequencies until the resonance is most exaggerated, then narrow the Q to be surgical, and finally, dip the EQ the amount desired.

Equalizer types

Most of you are familiar with the difference between parametric and shelving equalizers. Parametric is the most popular equalizer type in recording and mixing, because we're working with individual instruments. In mastering, shelving equalizers take on an increased role, because we're dealing with overall program material. But the parametric is still most popular as it is surgical with defects, such as bass instruments that have resonances. Very few people know of a third and important curve that's extremely useful in mastering: the Baxandall curve (see ill.). Hi-Fi tone controls are usually modeled around the Baxandall curve. Like shelving equalizers, a Baxandall curve is applied to low or high frequency boost/cuts. With a boost, instead of reaching a plateau (shelf), the Baxandall continues to rise. Think of the spread wings of a butterfly, but with a gentle curve applied. You can simulate a Baxandall high frequency boost by placing a parametric equalizer (Q= approximately 1) at the high-frequency limit (approximately 20 kHz). The portion of the bell curve above 20 k is ignored, and the result is a gradual rise starting at about 10 k and reaching its extreme at 20 k. This shape often corresponds better to the ear's desires than any standard shelf.

Baxandall Curve (grey) vs. Shelf (black)



Most times the same EQ adjustment in both channels is best, as it maintains the stereo balance and relative phase between channels. But sometimes it is essential to be able to alter only one channel's EQ. With a too-bright high-hat on the right side, a good sounding vocal in the middle and proper crash cymbal on the left, the best solution is to work on the right channel's high frequencies. The Finalizer does not currently operate on separate channels, but other TC products provide this flexibility.

Sometimes important instruments need help, though they should have been fixed in the mix. The best repair approach is to start subtly and advance to severity only if subtlety doesn't work. Remember: With a 2-track, every change affects everything!

If the piano solo is weak, we try to make the changes surgically:

- only during the solo.
- only on the channel where the piano is primarily located, if that sounds less obtrusive.
- only in the fundamental frequencies if possible.
- as a last resort by raising the entire level, because a keen ear may notice a change when the gain is brought up.

Instant A/B's?

With good monitoring, equalization changes of less than 1/2 dB are audible, so subtlety counts. You probably won't hear these changes in an instant A/B comparison, but you will notice them over time. I will take an equalizer in and out to confirm initial settings, but I never make instant EQ judgments. Music is so fluid from moment to moment that changes in the music will be confused with EQ changes. I usually play a passage for a reasonable time with setting "A" (sometimes 30 seconds, sometimes several minutes), then play it again with setting "B". Or, play a continuous passage, listening to "A" for a reasonable time before switching to "B". For example, over time it will become clear whether a subtle high frequency boost is helping or hurting the music.

Equalization or Multiband Compression?

Many people have complained that digital recording is harsh and bright. This is partly accurate: low-resolution recording (e.g., 16 bit) doesn't sound as warm to the ear as high-resolution. In addition, digital recording is extremely unforgiving; distortion in preamplifiers, A/Ds, errors in mike placement are mercilessly revealed. The mastering engineer recognizes these defects and struggles to make a pleasant-sounding result. Use equalization when instruments at all levels need alteration, or one of the best tools to deal with these problems -multiband compression, which

EQUALIZATION

provides spectral balancing at different levels. It is possible to simulate the often-desirable high-frequency saturation characteristics of analog tape with a gentle high-frequency compressor. Use increased high-frequency compression at high levels when the sound gets harsh or bright. Or, vice versa, if you find that at low levels, the sound is losing its definition (which can happen due to poor microphone techniques, noise in the recording, or low-resolution recording)...then apply gentle high frequency upward compression, engaged at lower levels. This function, often called AGC, is not available in the Finalizer, but can be found in the DBMAX by TC Electronic.

EQ Interaction with the compressor

If you're using split dynamics, make your first pass at equalization with the outputs (makeup gains) of the three bands. 3-band compression and equalization work hand-in-hand. If you're splitting dynamics processing, then tonal balance will be affected by the crossover frequencies, the amount of compression, and the makeup gain of each band. Before engaging an equalizer, first try to correct overall tonal balance with the makeup (output) gains of each compressor band. In general, the more compression, the "duller" the sound, because of the loss of transients. I first try to solve this problem by using less compression, or altering the attack time of the high-frequency compressor, but you may prefer to use the makeup gain or an equalizer to restore the high-frequency balance.

SIBILANCE CONTROL

Sibilance Control (exaggerated s sounds) is a natural artifact of compressors. This occurs because the compressor doesn't recognize the continuous s sound as over threshold, but the ear is extremely sensitive in that frequency region. In other words, the compressor doesn't correspond with how the ear works.

The solution is a very fast, narrowband compressor working only in the sibilance region (anywhere from 2.5 kHz to about 9 kHz).

At concerts do you try to identify the microphones that are used?

NOISE REDUCTION

Compression tends to amplify the noise in a source, because when the signal is below threshold, the compressor raises the gain. A possible antidote for noise is gentle low level expansion, especially at selective frequencies. Tape hiss, preamp hiss, noisy guitar and synth amplifiers can be perceived as problems or just "part of the sound." But when you think the noise is a problem, don't be overzealous in its removal. I often refer to the sound of poorly-applied noise reduction as "losing the baby with the bathwater." The key to good-sounding noise reduction is not to remove all the noise, but to accept a small improvement as a victory. Remember that louder signals mask the hiss, and also remember that the general public does not zero in on the noise as a problem. They're paying attention to the music, and you should, too!

1 to 4 dB of reduction in a narrow band centered around 3-5 kHz can be very effective and if done right, invisible to the ear. Do this with the Finalizer's multiband expansion. Start by finding a threshold, with initially a high expansion ratio, fast attack and release time. Zero in on a threshold that is just above the noise level. You'll hear ugly "chatter" and bouncing of the noise floor. Now, reduce the ratio to very small, below 1:2, perhaps even 1:1.1, and slow the attack and release until there is little or no perceived modulation of the noise floor. The attack will usually have to be much faster than the release so that fast crescendos will not be affected. This gives gentle, almost imperceptible noise reduction. Use the Finalizer's compare button to see how successful you've been. Hiss can be dramatically reduced, but make sure you haven't damaged the music along with it. The thresholds in the other two bands may have to be set very high (expansion off). The Finalizer's look-ahead delay actually allows the Expander to open before it's hit by the signal, thereby conserving transients.

Know Your Limits

Noise reduction through simple expansion has its limits. If you're not satisfied, you may have to put the recording through specialized dedicated noise-reduction units, which employ algorithms that took years to perfect. In Noise-Reduction, you do get what you pay for, and if it's inexpensive, it's either ineffective, or probably no good.

MONITORS

Monitors and Equalization

An inaccurate or unrefined monitor system not only causes incorrect equalization, but also results in too much equalization. The more accurate and linear your monitors, the less equalization you will apply, so it pays to talk a bit about monitor adjustment. The ear/brain must be used in conjunction with test instruments to determine monitor accuracy. For example, some degree of measured high-frequency roll-off usually sounds best (due to losses in the air), but there is no objective measurement that says "this roll-off measures right", only an approximation. Thus, for the high frequencies, the ultimate monitor tweak must be done by ear. Which leads us to the chicken and egg problem: "If you use recordings to judge monitors, how do you know that the recording was done right?" The answer is to use the finest reference recordings (at least 25 to 50) to judge the monitors, and take an average. The highs will vary from a touch dull to a touch bright, but the majority will be right on if your monitor system is accurate. Try to avoid adding monitor correction equalizers; better to fix the room or replace the loudspeakers; my techniques include tweaks on speaker crossover components until the monitors fall precisely in the middle of the "acceptance curve" of all 50 reference recordings. Even with monitor brands that sound perfect elsewhere, your room, interconnect cable capacitance, power amplifiers, D/A converters, and preamplifiers affect high frequency response especially, so if you make any changes, you must reevaluate your monitors with the 25 best recordings!

Monitors and Stereo Imaging

The Finalizer provides powerful techniques for adjusting stereo imaging. But first, your monitors and acoustics must be up to the task. Separate your monitors to approximately a 60 degree angle. There is a test record that objectively evaluates stereo imaging, and detects comb-filtering caused by nearby surfaces, as well as defects in speaker crossovers. It's called the LEDR test, short for Listening Environment Diagnostic Recording, and is available from Chesky Records, (<http://www.chesky.com>) on JD37. First play the announce track and confirm that the announcer's positions are correct. If not, then adjust speaker separation and angle.

Then play the LEDR test. The **beyond** signal should extend about 1 foot to the left and right of the speakers. If not, then look for side wall reflections. Similarly, the **up** signal should rise straight up, 3 to 6 feet, and the **over** signal should be a rainbow rising at least as high as the **up**. If not, look for interfering objects above and between the speakers, or defective drivers or crossovers.

Adjusting stereo balance

Stereo balance must not be judged by comparing channel meters. The only way to accurately adjust stereo balance is by ear. Confirm your monitors are balanced by playing pink noise at an exact matched channel level. Sit in the "sweet spot". All frequencies of the pink noise must image in a narrow spot in the center of the loudspeakers.

If a film or TV actor wearing a microphone crosses his arms, do you immediately notice the change in sound quality?

ADVANCED MASTERING TECHNIQUES

Mastering benefits from the digital audio workstation. The powerful Digital Audio Workstation (DAW) lets you make edits, smooth fades, emphasize or de-emphasize the loudness of sections. A client brought a DAT with 10 songs. On one of the songs, the bass was not mixed loudly enough (this can happen to even the best producer). We were able to bring up the bass with a narrow-band equalizer that had little effect on the vocal. But when the producer took the ref home, he was dissatisfied. "You've done a wonderful job on the bass, but the delicacy of the vocal is affected too much for my goals. Do you think I can bring you a DAT of the bass part so we can raise it there? I can't possibly duplicate this mix."

I told him we could handle that, asking for a DAT with a full mix reference on one channel, and the isolated bass on the other. I was able to load the DAT into my workstation, synchronize the isolated bass, and raise the bass in the mastering environment, without affecting the vocal. It was an unequivocal success. Another client doing the album of a "new age pianist" brought a four-track Exabyte archive in our workstation's format. Tracks 1 & 2 contained the full mix minus the piano, and tracks 3 & 4 contained only the piano. If all four tracks were mixed at unity gain we would end up with the full mix, but if necessary, we could level, compress, or equalize the piano separately in the mastering.

Alternate Mixes

Another approach is to ask the client to send separate "vocal up", "vocal correct", and "vocal down" mixes because the mastering environment is ideal for making those decisions, and mastering processing may affect that delicate balance. But often it's a luxury to make separate mixes, and we dream of ways of tweaking the mix on an existing two-track. A recent client had mixed in a bass-light room and his bass was very boomy, right up to about 180 Hz. The vocal came down slightly when I corrected the boomy bass, but through special M-S processing techniques, I was able to produce a perfectly-balanced master...which leads us to...

MS Mastering Techniques

Prior to the advent of digital processors like the Finalizer, mastering engineers were fairly limited in what we could accomplish; today, we still tell a few clients to "go back and fix it in the mix", but we have tricks up our sleeves that can accomplish wonders with a two-track mix. One ancient technique is incredibly powerful - MS Mastering.

MS stands for "Mid-Side", or "Mono-Stereo". In MS microphone technique, a cardioid, front-facing microphone is fed to the M, or

mono channel, and a figure 8, side-facing microphone is fed to the S, or stereo channel. A simple decoder (just an audio mixer) combines these two channels to produce L(Left) and R(Right) outputs. Here's the decoder formula: M plus S equals L, M minus S equals R. Here's how to decode in the mixer: feed M to fader 1, S to fader 2, pan both to the left. Feed M to fader 3, S to fader 4, invert the polarity of fader 4 ("minus S"), pan both to the right. The more M in the mix, the more monophonic (centered) the material, the more S, the more wide-spread, or diffuse the material. If you mute the M channel, you will hear out of phase sound, containing largely the reverberation and the instruments at the extreme sides. Mute the S channel, and you will largely hear the vocalist; the sound collapses, missing richness and space.

There's not perfect separation between M and S channels, but enough to accomplish a lot of control on a simple 2-track. It's great for film work - the apparent distance and position of an actor can be changed by simple manipulation of two faders.

M-S technique doesn't have to be reserved to a specialized miking technique. By using MS, we can separate an ordinary stereo recording into its center and side elements, and then separately process those elements. I tell my clients I'm making three tracks from two.

Do you try to identify the frequency every time you hear system feedback?

The MS adventure begins

MS is another tool that reduces compromises and increases the possibilities of mastering. The possibilities are only limited by your imagination. The Finalizer, and especially the Finalizer 96K, allows you to manipulate stereo separation using MS technique. Let's take a stereo recording with a weak, center-channel vocalist. First we put in our MS encoder, which separates the signal into M and S. Then we decrease the S level or increase the M level. We then decode that signal back into L and R. Presto, the vocal level comes up, as does the bass (usually) and every other center instrument. In addition, the stereo width narrows, which often isn't desirable. But at least we raised the vocalist and saved the day! The Finalizer's built-in width control does this job by changing the ratio of M to S.

ADVANCED MASTERING TECHNIQUES

But we can accomplish a lot more, often with no audible compromise to the presentation, and make clients very happy. Let's take our stereo recording, encode it into MS, and apply separate equalization to the M and S channels. Here's the traditional (pre-Finalizer) method: Feed the output of the MS encoder to a dual-channel equalizer. Channel one of the equalizer contains the M channel, which has most of the vocal. Channel two contains the S channel, which has most of the ambience and side instruments. With the M channel EQ, we can raise the vocal slightly by raising (for example) the 250 Hz range, and perhaps also the presence range (5 kHz, for example). This brings up the center vocal with little effect on the other instruments, and lowers the stereo separation almost imperceptibly.

The Finalizer 96K's Spectral Stereo Imager can also "remix" this material, with a slightly different user interface. By raising the M level (reducing the width) of the 250 Hz and/or 5 kHz range, we bring up the center vocal very similarly to the traditional method, and without seriously deteriorating the imaging of the other instruments. In addition to this "remix" facility, the spectral stereo imager has very creative width control, limited only by your imagination. Spread the cymbals without losing the focus of the snare, tighten the bass image without losing stereo separation of other instruments, and so on.

Even More Advanced M-S Technique

Currently the Finalizer has a single threshold for both channels, but other TC Electronics products can accomplish even more sophisticated M-S mastering. You've all heard the mix that sounds great, but the vocal is sometimes a bit buried when the instruments get loud. We try compressing the overall mix, or even narrow band compression of the vocal range, but it worsens the great sound of the instruments. MS compression can help us isolate the compression to the center channel.... By only compressing the M channel, we delicately bring up the center channel level when signals get loud. Or, better yet, use multiband MS compression, so, the bass (for example) is unaffected by our compression. In other words, compress only the midrange frequencies of only the M channel... A very selective and powerful process, only available in today's digital world.

Make an UnFinalized Safety

Now you have a master, ready to send to the plant. The most professional mastering engineers make unprocessed safety copies of the music for future release on high-resolution media. If you're creating a submaster to take to a mastering house, also make an unprocessed version, as the mastering house may have a different idea of how to make your music shine.

If you answered yes to most of the questions in this tutorial, then you'll make a great mastering engineer.

APPENDIX - REFERENCE

References

The Digital Domain website, <http://www.digido.com>, contains further references to topics mentioned in this booklet, with information on dither, compression, metering, monitor calibration, "good-sounding" commercial CDs to listen to, scientific references, and more.

APPENDIX - GLOSSARY

dBFS - dB reference full scale

Full scale digital is defined as 0 dBFS, the maximum numeric level which can be encoded.

Gain, Loudness, Volume and Level

Loudness is the subjective judgment of level, by ear. Loudness is an approximate quantity, while **level** can be repeatable measured, if the averaging time of the measuring instrument is specified. In a professional context, don't use the term **volume** to avoid confusion with quarts and liters - use loudness instead. Use the professional term gain control rather than **volume control**. Gain is often confused with Level. **Gain** is the property of an amplifier or attenuator, while level refers to the amount of signal going through that amplifier. For example, a signal may measure a level of -10 dBFS before going through an amplifier. If the amplifier has a gain of 6 dB, then the level of the signal on the output of the amplifier will be -4 dBFS. Absolute values are applied to level (e.g., 6 volts, or -12 dBFS), while only relative values can be used for gain (e.g. a gain of 2X, or +6 dB). It is incorrect to say that a device has a gain of +7 dBm or dBv...dBm or dBv or dBFS are absolute terms reserved for measured level, not for gain.

Sample Rate

The number of samples per second. The preferred abbreviation is S/s, or kS/s, e.g., 44.1 kS/s. Common usage is "44.1 kHz", but this approach is often confusing when the same sentence also refers to bandwidth or frequencies of interest.

ACKNOWLEDGMENTS

A big thank you to Bob Ludwig, of Gateway Mastering, Portland Maine, and Glenn Meadows, of Masterfonics, Nashville, Tennessee. Bob and Glenn reviewed the manuscript and added helpful suggestions that made an even better booklet.