

THE MASTERING ENGINEER'S HANDBOOK

Fourth Edition



Bobby Owsinski

The
Mastering
Engineer's
Handbook

Fourth Edition

Bobby Owsinski



The Mastering Engineer's Handbook 4th edition

by Bobby Owsinski

Published by:
Bobby Owsinski Media Group
4109 West Burbank, Blvd.
Burbank, CA 91505

© Bobby Owsinski 2016

ISBN 13: 978-0-985033-3-2

ALL RIGHTS RESERVED. No part of this work covered by the copyright herein may be reproduced, transmitted, stored, or used in any form and by any means graphic, electronic or mechanical, including but not limited to photocopying, scanning, digitizing, taping, Web distribution, information networks or information storage and retrieval systems, except as permitted in Sections 107 or 108 of the 1976 Copyright Act, without the prior written permission of the publisher.

For permission to use text or information from this product, submit requests to requests@bobbyowsinski.com.

Please note that much of this publication is based on personal experience and anecdotal evidence. Although the author and publisher have made every reasonable attempt to achieve complete accuracy of the content in this Guide, they assume no responsibility for errors or omissions. Also, you should use this information as you see fit, and at your own risk. Your particular situation may not be exactly suited to the examples illustrated herein; in fact, it's likely that they won't be the same, and you should adjust your use of the information and recommendations accordingly.

Any trademarks, service marks, product names or named features are assumed to be the property of their respective owners, and are used only for reference. There is no implied endorsement if we use one of these terms.

Finally, nothing in this book is intended to replace common sense, legal, medical or other **professional** advice, and is meant to inform and entertain the reader.

To buy books in quantity for corporate use or incentives, call 818.588.6606 or
email office@bobbyowsinski.com.

Introduction

It's already been 16 years since the first edition of *The Mastering Engineer's Handbook* came out, and boy, have things changed. It's safe to say that there's been a mighty revolution in the mastering world, with old technologies replaced and new ones continually evolving. Gone are the days of tape machines (for the most part), and soon even the CD will be a thing of the past.

Gone—again, for the most part—are the days of “heavy iron” customized outboard gear that was necessary for a high-quality mastering job. Even though the basic mastering tools are still the same, they've mostly moved into the world of the DAW, so even someone with the most entry-level system now has the use of powerful tools that only the top mastering pros had access to in the past. And maybe best of all, it's now possible to totally prep just about any kind of audio for any kind of distribution (which is what mastering really is) at home in your personal studio.

Just like everything else in music and recording, some really excellent mastering tools are available to just about anyone with a DAW (which is most of us that are into music these days). That makes the process of mastering very inexpensive compared to previous generations of musicians and audio engineers, but just because you own a hammer doesn't mean that you know how to swing it. A lot of harm can come from misuse of the tools of mastering if the process and concepts are not thoroughly understood.

And that's what this book is about.

In it, we'll take a look at how the top mastering pros perform their magic as some of the top mastering engineers describe their processes in the interviews. Through this, we'll develop a good, strong reference point so we can either do our own mastering (and hopefully do no harm to the material, just like a doctor) or know when it's time to call a pro and then properly prep the program for them to get the best results possible.

More so than any other process in audio, mastering is more than just knowing the procedure and owning the equipment. Yes, more than any other job in audio, mastering done at its highest level is about the long, hard grind of

experience. It's about the cumulative knowledge gained from 12-hour days of listening to both great and terrible mixes; from working on all types of music, not just the type you like; from saving the client's butt without him ever knowing it; from doing 10 times more work than the client ever sees.

Among the many things this book will provide is an insider's look at the process, not so much from my eyes, but from that of the legends and greats of the business.

My goal with this book is a simple one: to help the guy who wants to do his own mastering do a better job of it, and to show that there's a lot more to a professional mastering job than meets the eye.

For those of you who have read my previous books like *The Mixing Engineer's Handbook* and *The Recording Engineer's Handbook*, you'll notice that the format for this book is similar. It's divided into two sections:

- **Part I: The Mechanics of Mastering** provides an overview of the history, tools, philosophy, background, and tips and tricks used by the best mastering engineers in the business.
- **Part II: The Interviews** is a behind-the-scenes look at the mastering world through the eyes of some of the finest (and in some cases, legendary) mastering engineers in the world.

Along with this book, you might also want to take a look at my *Mastering Audio Techniques* course at lynda.com for a more visual approach to how mastering is done.

Meet The Mastering Engineers

Here's a list of the mastering engineers who have contributed to this book, along with some of their credits. I've tried to include not only the most notable names in the business from the main media centers, but also engineers who deal with specialty clients. I'll be quoting them from time to time, so I wanted to introduce them early on so you have some idea of their background when they pop up.

- **Doug Sax.** Perhaps the Godfather of all mastering engineers, Doug

was the first independent when he starting his famous Mastering Lab in Los Angeles in 1967. Since then, he has worked his magic with such diverse talents as The Who, Pink Floyd, The Rolling Stones, the Eagles, Kenny Rogers, Barbra Streisand, Neil Diamond, Earth, Wind & Fire, Diana Krall, the Dixie Chicks, Rod Stewart, Jackson Browne, and many, many more.

- **Bernie Grundman.** One of the most widely respected names in the recording industry, Bernie has mastered literally hundreds of platinum and gold albums, including some of the most successful landmark recordings of all time, such as Michael Jackson's Thriller, Steely Dan's Aja, and Carole King's Tapestry.
- **Bob Ludwig.** Bob certainly stands among the giants in the mastering business. After leaving New York City to open his own Gateway Mastering in Portland, Maine, in 1993, Bob has worked on literally hundreds of platinum and gold records, and mastered projects that have been nominated for scores of Grammys.
- **Greg Calbi.** One of the owners of Sterling Sound in New York City, Greg's credits include Bob Dylan, John Lennon, U2, David Bowie, Paul Simon, Paul McCartney, Blues Traveler, and Sarah McLachlan, among many, many others.
- **Glenn Meadows.** Glenn is a Nashville-based two-time Grammy winner and a multi-TEC award nominee who has worked on scores of gold and platinum records for a diverse array of artists, including Shania Twain, LeAnn Rimes, Randy Travis, Delbert McClinton, and Reba McEntire, as well as for multi-platinum producers such as Tony Brown, Jimmy Bowen, and Mutt Lange.
- **Gene Grimaldi.** Gene is the chief engineer at Oasis Mastering in Los Angeles, and has a list of blockbuster clients that include Lady Gaga, Jennifer Lopez, Carly Rae Jepsen, Ellie Goulding, Nicki Minaj, and many more.
- **David Glasser.** David is the founder and chief engineer of Airshow Mastering in Boulder, Colorado, and Takoma Park, Maryland, and has worked for some 80 Grammy nominees. He's also an expert in catalog restoration, having worked on releases by Smithsonian Folkways Recordings and the Grateful Dead, among many others.

- **Dave Collins.** Operating out of his own Dave Collins Mastering studios in Hollywood, Dave has mastered projects for Sting, Madonna, Bruce Springsteen, and Soundgarden, among many others.
- **Colin Leonard.** With credits like Justin Bieber, Jay-Z, Echosmith, Leona Lewis, Al Di Meola, John Legend and many more, plus a dedicated following of A-list mixers, Colin uses some proprietary analog gear at his Atlanta-based SING Mastering to take a different approach to mastering. Colin is also the creator of Aria automated online mastering, the latest trend in convenient and inexpensive mastering

While you probably won't have access to the gear, playback systems, and rooms that the above engineers have, that's okay because a great mastering job can be at your fingertips if you follow their advice and examples and use the greatest tool you have available—your ears.

PART I

The

Mechanics Of Mastering

Chapter 1

The Essence Of Mastering

The term “mastering” is either completely misunderstood or shrouded in mystery, but the process is really pretty simple. Technically speaking, mastering is the intermediate step between mixing the audio and having it replicated or distributed. Up until recently, we would define it as follows:

Mastering is the process of turning a collection of songs into an album by making them sound like they belong together in tone, volume, and timing (spacing between songs).

That was the old way to explain mastering when the album was king. Since we live in a singles world today, the definition has to be tweaked for our current production flow. Let’s use this definition instead.

Mastering is the process fine-tuning the level, frequency balance, and metadata of a track in preparation for distribution.

That first definition isn’t obsolete though, since albums are still around (and probably always will be), but the fact of the matter is that individual songs are always played in a collection. The collection can be an album or, more usually, a playlist where the song is played before or after someone else’s track on the radio, on an online distribution service, or on someone’s playback device. Of course, you want all your songs to sound at least as good as the others that you listen to or the one’s they’re played around.

I think that mastering is a way of maximizing music to make it more effective for the listener as well as maybe maximizing it in a competitive way for the industry. It's the final creative step and the last chance to do any modifications that might take the song to the next level.

—Bernie Grundman

So loosely speaking, that’s what mastering is. Here’s what mastering is not—it’s not a tool or a plugin that automatically masters a song with little or no effort from the operator. All too often people have the misconception that

mastering is only about EQing the track to make it sound bigger, but it's really more of an art form that relies on an individual's skill, experience with various genres of music, and good taste. In fact, it's been said that 95 percent of all mastering is in the ears, and not the tools.

I think that mastering is, and always has been, the real bridge between the pro audio industry and the hi-fi industry. We're the ones that have to take this stuff that sounds hopefully good or great on a big professional monitor system and make sure it also translates well to the home systems. We're the last link to get it right or the last chance to really screw it up and make it bad, and I think we're all guilty at times of doing both.

—Glenn Meadows

While the tools for audio mastering do require more precision than in other audio operations, the bottom line is that this is an area of audio where experience really does matter.

Why Master Anyway?

Mastering should be considered the final step in the creative process, since this is the last chance to polish and fix a project. Not all projects need mastering, especially if they're not destined to be heard by the public, but here are a few instances when mastering can help:

- If you have a song that sounds pretty good by itself but plays at a lower volume when played after another song.
- If you have a song that sounds pretty good by itself but sounds too bright or dull next to another song.
- If you have a song that sounds pretty good by itself but sounds too bottom heavy or bottom light against another song.

A project that has been mastered simply sounds better if done well (that's the key phrase, of course.) It sounds complete, polished, and finished. The project that might have sounded like a demo before now sounds like a “record”

because:

- Judicious amounts of EQ and compression were added to make the project sound bigger, fatter, richer, and louder.
- The levels for each song of the album (if there is one) are adjusted so they all have the same apparent level or have the same level as other professionally mastered songs in the same genre.
- The fades have been fixed so that they're smooth, if needed.
- The distorted parts or glitches have been edited out.
- All the songs of an album blend together into a cohesive unit.
- In the case of mastering for CD or vinyl, the spreads (the time between each song) have been inserted so the songs flow seamlessly together.
- The songs destined for a CD or vinyl record are sequenced so they fall in the correct order.
- ISRC codes and the proper metadata has been inserted into each track.
- A backup clone is created and stored in case anything should happen to the master.
- Any shipping or uploading to the desired replication facility is taken care of.

As you can see, there's a lot more to mastering than it seems when you really get into it. To begin to understand mastering, let's see how it has evolved over the years.

From Vinyl, To CDs, To MP3s, And Beyond

Until 1948, there was no distinction between different types of audio engineers, because everything was recorded directly onto 10-inch vinyl records that

played at 78 rpm. In 1948, however, the age of the mastering engineer began when Ampex introduced its first commercial magnetic tape recorder. Since most recording of the time began using magnetic tape, a transfer had to be made to a vinyl master for delivery to the pressing plant to make records, hence the first incarnation of the “mastering engineer” was born. There was no concept of what we now consider mastering at the time though, so he was called “transfer engineer.” (see Figure 1.1).



Figure 1.1: A disc-cutting lathe
© 2017 Bobby Owsinski

There was a high degree of difficulty in this transfer process because the level applied to the master vinyl lacquer when cutting the grooves was so crucial. Too low a level and you get a noisy disc, but hit it too hard and you destroy the disc and maybe the expensive (\$15,000 in '50s and '60s dollars) cutting stylus of the lathe too (see Figure 1.2).




Figure 1.2: A disc-cutting stylus

© 2017 Bobby Owsinski

In 1955, Ampex released tape machines that had a new feature called Selective Synchronous Recording, or Sel Sync, which gave the multitrack recorder the ability to overdub, thus changing the recording industry forever. At this point there became a real distinction between the recording and mastering engineer, since the jobs now differed so greatly, although many were trained at both jobs (the EMI training program at Abbey Road made mastering the last job before you became a full engineer).

In 1957, the stereo vinyl record became commercially available and really pushed the industry to what many say was the best-sounding audio ever. Mastering engineers, who were now known as “cutters,” found ways to make the discs louder (and as a result less noisy) by applying equalization and compression. Producers and artists began to take notice that certain records would actually sound louder on the radio, and if it played louder, then the listeners usually thought it sounded better (although they were speculating instead of using any scientific data), and maybe the disc sold better as a result. Hence, a new breed of mastering engineer was born—this one with some

creative control and ability to influence the final sound of a record, rather than just being a transfer jock from medium to medium.

An interesting distinction between American and British mastering engineers developed though. In the U.S., mastering was and still is considered the final step in the creation of an album, while in the UK they look at it as the first step in manufacturing. As a result, American mastering engineers tend to have much more creative leeway in what they're allowed to do to the audio than British engineers.

With the introduction of the CD in 1982, the cutting engineer, who was now finally known as a “mastering engineer,” was forced into the digital age, using a modified video tape recorder called a Sony 1630 (see Figure 1.3) to deliver the digital CD master to the replicator, but still utilizing many of the analog tools from the vinyl past for EQ and compression. The 1989 introduction of the Sonic Solutions digital audio workstation with “pre-mastering software” provided a CD master instead of a bulky 1630 tape cartridge (see Figure 1.4). Now mastering began to evolve into the digital state as we know it today.



Figure 1.3: A Sony 1630
© 2017 Bobby Owsinski




Figure 1.4: A tape cartridge used in a 1630
© 2017 Bobby Owsinski

In the first half of 1995, MP3s began to spread on the Internet, and their small file size set about a revolution in the music industry that continues to this day. This meant that the mastering engineer had to become well versed in how to get the most from this format, something it took many mastering engineers years to get the hang of.

In 1999, 5.1 surround sound and high-resolution audio took the mastering engineer into new, uncharted but highly creative territory. And by 2002, almost all mastering engineers were well acquainted with the computer, since virtually every single project was edited and manipulated with digital audio workstation software. Nowadays a majority of engineers are firmly in the box unless given a 1/2- or 1/4-inch tape to master.

Today's mastering engineer doesn't practice the black art of disc cutting as much as was once required, but he's no less the wizard as he continues to shape and mold a project like never before.

The Difference Between You And A Pro

There are a lot of reasons why a commercial mastering facility usually

produces a better product than mastering at home. If we really break it down, a mastering pro usually has three things over the home studio.

The Gear: A real pro mastering house has many things available that you probably won't find in a simple home or small studio DAW room, such as high-end A/D and D/A converters and signal path, a great-sounding listening environment, and an exceptional monitoring system (see Figure 1.5).



Figure 1.5: The Tannoy monitoring system at Oasis Mastering
© 2017 Bobby Owsinski

The reason people come to a mastering engineer is to gain that mastering engineer's anchor into what they hear and how they hear it and the ability to get that stuff sounding right to the outside world.

—Glenn Meadows

The monitor system of these facilities sometimes costs far more than many

entire home studios (and even more than entire homes, for that matter). Cost isn't the point here, but quality is, since you can rarely hear what you need to in order to make the fine adjustments that you need to make on the nearfield monitors that most recording studios use. The vast majority of monitors, and the rooms in which they reside, are just not precise enough.

The Ears: The mastering engineer is the real key to the process. This is all he does day in and day out. He has “big ears” because he masters at least eight hours every day and knows his monitors better than you know your favorite pair of sneakers. Plus, his reference point of what constitutes a good-sounding mix is finely honed thanks to working hours and hours on the best- and worst-sounding mixes of each genre of music.

Most people need a mastering engineer to bring a certain amount of objectivity to their mix, plus a certain amount of experience. If you (the mastering engineer) have been in the business a while, you've listened to a lot of material, and you've probably heard what really great recordings of any type of music sound like, so in your mind you immediately compare it to the best ones you've ever heard. You know, the ones that really got you excited and created the kind of effect that producers are looking for. If it doesn't meet that ideal, you try to manipulate the sound in such a way as to make it as exciting and effective a musical experience as you've ever had with that kind of music.

—Bernie Grundman

I personally think experience is as valuable as equipment in a large sense because after you've done it for 10 or 20 years, you've heard almost everything that can possibly go wrong and go right on a mix, so you can, in one respect, quickly address people's problems. When a guy writes a book, he doesn't edit the book himself. He sends it off to an editor, and the editor reads it with a fresh set of eyes, just like a mastering engineer hears it with a fresh set of ears.

—Dave Collins

A Backup: I don't know who said it, but this phrase rings true: "The difference between a pro and an amateur is that a pro always has a backup." Good advice for any part of recording, but especially for mastering. You wouldn't believe the number of times masters get lost, even when major record labels are involved. This is the one thing that you can do just as well as a pro can with no trouble at all!

Finally, if mastering was so easy, don't you think that every big-time engineer or producer (or record label, for that matter) would do it themselves? They don't, and mastering houses are busier than ever, which tells you something.

It is the impartial ear that you get from your mastering engineer that is valuable. All this equipment and new technology that we've got is a great thing, but you're really asking for someone who has never heard the record before to hear it for the first time fresh.

—Dave Collins

Mastering is more than just knowing how to manipulate the sound to get it to where somebody wants it to go. I think that a lot of it is this willingness to enter into another person's world, and get to know it and actually help that person express what he is trying to express, only better.

—Bernie Grundman

There's Always Room For DIY

While the above section seems like I'm trying to discourage you from doing your own mastering, that's really not the case. In fact, what I'm trying to do is give you a reference point of how the pros operate and why they're so successful. From there you can determine whether you're better served by doing it yourself or using a pro.

But the reason that you're reading this book is because you want to learn about all the tricks, techniques, and nuances of a major mastering facility, right? For one thing, there are mastering situations that don't need a professional's touch, and for another, sometimes budgets are so tight that there's just no money left over for a mastering pro no matter how much you'd like to use him.

As far as the person who might be trying to learn how to do his own mastering, or understand mastering in general, the main thing is that all you need is one experience of hearing somebody else master something. Your one experience at having it sound so incredibly different makes you then realize just how intricate mastering can be and just how much you could add to or subtract from a final mix.

—*Greg Calbi*

Read on, and you'll discover the hows and whys of mastering in detail.

Chapter 2

Digital Audio Basics

Now is a good time for a brief review of some of the basics of digital audio. While you may be familiar with the sample rate and word length already, there always seems to be a lot of questions about the differences between file formats, such as AIFF and WAV, so we'll try to take care of them straight away.

Sample rate and word length determine the quality of a digital audio signal. In order to understand how that happens, a brief discussion is in order. Remember, this is just an overview and only gives you the general concepts of digital audio. If you really want to get under the hood of digital audio, refer to a book like *Principles of Digital Audio* (McGraw-Hill, 2010) by Ken Pohlmann that thoroughly covers the subject.

Sample Rate

Sample rate is one of the determining factors when it comes to the quality of a digital audio signal. The analog audio waveform amplitude is measured by the analog-to-digital converter (more on this device in Chapter 4, “Monitoring for Mastering”) at discrete points in time, and this is called sampling. The more samples that are taken of the analog waveform per second, the better the digital representation of the waveform is, which results in a greater frequency response of the signal (see Figure 2.1).




Figure 2.1: Sample rate

© 2017 Bobby Owsinski

For example, if we were to use a sampling rate of 48,000 times a second (or 48kHz), that would present us with a frequency response of 24kHz, or half the sample rate. That's because of a law of digital audio called the Nyquist Theorem, which states that your sample rate has to be twice as high as the highest frequency you wish to record, otherwise, digital artifacts called aliasing will be added to the signal. A low-pass filter is used to limit the bandwidth to half the sampling rate.

A sample rate of 96kHz provides a better digital representation of the waveform because it uses more samples, and yields a usable audio bandwidth of about 48kHz. A 192kHz sample rate provides a bandwidth of 96kHz.

While it's true that we can't hear above 20kHz on even a good day, the frequency response that's available as a result of the high sampling rate means that a less intrusive filter can be used, so the digital signal sounds better, which is why we strive to use higher sample rates if possible.

TIP: *The higher the sampling rate, the better the representation of the analog signal and the greater the audio bandwidth will be, which means it sounds better!*

Although a higher sample rate yields a better representation of the analog signal, some

people can't always tell the difference due to the speakers, the listening environment, the signal path, the type of music, or how it was mixed. Couple that with the fact that higher sample rates require a higher powered computer, fewer tracks and plugins are available, and some plugins won't work at some of the very high sample rates, and you can see that sometimes a lower sample rate can be a better decision when it comes to recording. That said, 96kHz has become the new standard for music recording, especially since iTunes now encourages delivery of high-resolution files at that rate.

The downside of a higher sample rate is that it takes up more digital storage space, with 96kHz taking up twice as much as 48k, and 192k taking up twice as much again as 96k. That's no longer much of a problem though, as hard-drive disk or even flash-drive storage is massive compared to the needs of a typical song.

TIP: *It's always best to mix to the highest resolution possible both for archival purposes and because a high-resolution master makes for a better-sounding lower-resolution file. This applies even if the ultimate delivery medium is to be a lower-resolution CD or MP3.*

That said, a mastering engineer must work at certain sampling rates to deliver a product for a particular distribution medium.

Table 2.1: Sample Rates For Various Distribution Mediums

Typical Sample Rates	Comments	Caveats
44.1kHz	The CD sample rate	Fewer CDs are being made, so the minor advantage of recording using the similar sample rate is lost.
48kHz	Standard for film and TV	Lowest recommended sample rate.
96kHz	High-resolution standard	Most pro records are recorded at 96kHz. The recommended master delivery rate for iTunes. However, it takes up twice the storage space of 48kHz.
192kHz	Audiophile standard	Only half the channels and plugins of 96kHz on some DAWs. However, many plugins don't operate, and it takes up twice the storage space of 96kHz.

Bit Depth

Bit depth is the length of a digital word, and digital word length is the other factor involved in audio quality. It's somewhat the same as sample rate in that more is better. The more bits in a digital word, the better the dynamic range, which once again means the final audio sounds better. Every extra bit that's used means there is 6dB more dynamic range available. Therefore, 16 bits yields a maximum dynamic range of 96dB, 20 bits equals 120dB, and 24 bits provides a theoretical maximum of 144dB. From this you can see that a high-resolution 96kHz/24-bit (usually just abbreviated 96/24) format is far closer to sonic realism than the current CD standard of 44.1kHz/16-bit.

Today most recording is done at 24 bits, as there's virtually no advantage to using less. While once upon a time hard drive space or bandwidth was at a premium, neither applies now. That said, both CD and MP3 formats require 16 bits, but iTunes now encourages 24-bit delivery regardless of the sample rate.

TIP: *The longer the word length (the more bits), the greater the dynamic range and therefore the closer to real the sound can be.*

Even though bit depths like 32-bit, 32-bit float, and 64-bit float potentially provide much higher quality, chances are that you've not recorded in the other resolutions, so consider them there for a future delivery application.

TIP: *Never export to a higher than the resolution at which your project started, since you gain nothing in quality, and your file will be a lot larger. For instance, if your project started at 16-bit, selecting 24-bit or higher buys you nothing.*

Standard Audio File Formats

There are several audio file formats used today on most digital audio workstations. A file format specifies how the digital word is encoded in a digital storage medium. Some formats are universal and some are proprietary, while some, like JPEG and PNG files, are very specific as to what type of information they store. There are also specific types of file formats used for audio.

LPCM (Linear Pulse Code Modulation) is the process of sampling an analog waveform and converting it to digital bits that are represented by binary digits (1s and 0s) of the sample values. When LPCM audio is transmitted, each 1 is represented by a positive voltage pulse, and each 0 is represented by the absence of a pulse (see Figure 2.2). LPCM is the most common method of storing and transmitting uncompressed digital audio. Since it's a generic format, it can be read by most audio applications—similar to the way a plain text file can be read by any word-processing program. LPCM is used by audio CDs and is represented in a file format on a DAW by AIFF, BWF, WAV, or SD2 files.




Figure 2.2 : Linear PCM

© 2017 Bobby Owsinski

AIFF (Audio Interchange File Format) is a file format for storing LPCM digital audio data. It supports a variety of bit resolutions, sample rates, and channels of audio. The format was developed by [Apple Computer](#) and is the standard audio format for Macintosh computers, although it can be read by any type of computer workstation these days. AIFF files generally end with .aif.

WAV (Waveform Audio) is another file format for storing LPCM digital audio data. Created by Microsoft and IBM, WAV was one of the first audio file types developed for the PC. Wave files are indicated by a .wav suffix in the file name and are often spelled wav (instead of wave) in writing. The .wav file format supports a variety of bit resolutions, sample rates, and channels of audio.

BWF (Broadcast Wave) is a special version of the standard WAV audio file format developed by the European Broadcast Union in 1996. BWFs contain an extra chunk of data, known as the broadcast extension chunk, that contains information on the author, title, origination, date, time, and so on of the audio content. Perhaps the most significant aspect of BWFs is the feature of time stamping, which allows files to be moved from one DAW application to another and easily aligned to their proper point on a timeline or

an edit decision list.

SDII (Sound Designer II, sometimes seen abbreviated as SD2) is a mono or stereo audio file format for storing LPCM, originally developed by Digidesign for their DAW software applications. When used on a PC, the file must use the extension of .sd2. SD2 files are fast losing favor to the AIFF and WAV formats and should be considered obsolete.

There's really no operational difference between AIFF and WAV files these days. Once upon a time you'd use an AIFF audio file if you were on a Mac and a WAV file if you were on a PC, but both platforms now happily read either one without any difficulty.

SDII files are a different story, though. This is a file format that Digidesign (now Avid) introduced in their early days for their Sound Designer 2 application, which was the precursor to the now widely used Pro Tools. Although the format has the advantage of storing a limited amount of metadata (information about the data), its use has diminished over the years, and it's not entirely compatible with all playback software and DAWs. The only time SDIIs are completely safe to use is if your export is expressly intended for Pro Tools, but even then it's best to stay with a WAV or AIFF file format in case you ever decide to use another DAW in the future.

CAF (Core Audio Format). Some DAWs can export a new file format developed by Apple around its Core Audio technology for use with operating systems 10.4 and higher. CAFs are designed to overcome some of the limitations of the older WAV and AIFF file containers, such as the limit on file size. A CAF file doesn't have the 4GB limit of the other formats, and it can theoretically hold a file that is hundreds of years long (that's a big file!).

The format is also able to hold practically any type of audio data and metadata, any number of audio channels, and auxiliary information such as text annotations, markers, channel layouts, and other DAW data. One of the more interesting features of CAF as a file format is that you can append new audio data on the end of the file, making it ideal as an archive format.

All Apple software products, including Logic Pro, GarageBand, and QuickTime Player, as well as T-RackS, now support CAF files and can open them directly. If you want CAF files to be played on other systems, convert them to a WAV or MP3 file with a utility such as Factory Audio Converter.

Data Compression

Linear PCM files are large and, as a result, can sometimes be painfully slow to upload and download, even with a dedicated high-speed connection. As a result, data compression was introduced to keep a certain amount of sonic integrity (how much is in the ear of the beholder) while making an audio file imminently transportable.

Data compression isn't at all like the audio compression that we're going to be talking about later in the book. Data compression reduces the amount of physical storage space and memory required to store an audio file, and therefore reduces the time required to transfer a file. Files using data compression include MP3, AAC, FLAC, Dolby Digital, DTS, and many more. Check out Chapter 9 for more on the different types of data compression.

Chapter 3

Prepping For Mastering

In order for the mastering session to go smoothly, sound great, and save you money, some prep work is required beforehand. Even if you're doing your own mastering, these tips can really help improve your end product.

Mixing For Mastering

Regardless of whether you master your final mixes yourself or take them to a mastering engineer, things will go a lot faster if you prepare for mastering ahead of time. Nothing is as exasperating to all involved as not knowing which mix is the correct one or forgetting the file name. Here are some tips to get your tracks mastering-ready.

- **Don't over-EQ when mixing.** A mix is over-EQ'd when it has big spikes in its frequency response as a result of trying to make one or more instruments sit better in the mix. This can make your mix tear your head off because it's too bright, or have a huge and unnatural-sounding bottom. In general, mastering engineers can do a better job for you if your mix is on the dull side rather than too bright. Likewise, it's better to be light on the bottom end than to have too much.
- **Don't over-compress when mixing.** Over-compression means that you've added so much mix bus compression that the mix is robbed of all its life. You can tell that a mix has been over-compressed not only by its sound, but by the way its waveform is flat-lined on the DAW timeline. You might as well not even master if you've squashed it too much already. Hyper-compression (see Chapter 6, "Mastering Techniques") deprives the mastering engineer of one of his major abilities to help your project. Squash it for your friends, squash it for your clients, but leave some dynamics in the song so the mastering engineer is better able to do his thing. In general, it's best to compress and control levels on an individual-track basis and not as much on the stereo bus, except to prevent digital overs.
- **Having the levels match between songs is not important.** Just make your mixes sound great, because matching levels between songs is one of the reasons you master in the first place.
- **Getting hot mix levels is not important.** You still have plenty of headroom even if you print your mix with peaks reaching -10dB or so. Leave it to the mastering engineer to get those hot levels. It's another reason why you master.
- **Watch your fades and trims.** If you trim the heads and tails of your track too

tightly, you might discover that you've trimmed a reverb trail or an essential attack or breath. Leave a little room and perfect it in mastering, where you will probably hear things better.

- **Have the right documentation.** See the next section.
- **Make sure to print the highest-resolution mixes you can.** Lossy formats such as MP3s, Windows Media, or Real Audio and even audio CDs won't cut it and will give you an inferior product in the end. Print the highest-resolution mixes possible by staying at the same resolution as the tracks were recorded at. In other words, if the tracks were cut at a sample rate of 96kHz/24-bit, that's the resolution your mix should be. If it's at 44.1kHz/24-bit, that's the resolution the mix should be.
- **Don't add dither.** Adding dither to your mix actually reduces the resolution. Leave that for the mastering engineer (see more about dither in Chapter 7).
- **Alternate mixes can be your friend.** A vocal up/down or instrument-only mix can be a life-saver when mastering. Things that aren't apparent while mixing sometimes jump right out during mastering, and having an alternative mix around can sometimes provide a quick fix and keep you from having to remix. Make sure you document them properly though.
- **Reference your mixes in mono when mixing.** It can be a real shock when you get to the mastering studio, the engineer begins to check for mono compatibility, and the lead singer or guitar solo disappears from the mix because something in the track is out of phase. Even though this was more of a problem in the days of vinyl and AM radio, it's still an important point because many so-called stereo sources (such as television) are either pseudo-stereo or only heard in stereo some of the time. Check it and fix it before you get there.
- **Know your song sequence.** Song sequencing takes a lot of thought in order to make an album flow, so you really don't want to leave that until the mastering session. If you're cutting vinyl, remember that you need two sequences—one for each side of the disc. Remember, the masters can't be completed without the sequence. Also, cutting vinyl is a one-shot deal with no chance to undo, like on a workstation. It'll cost you money every time you change your mind.
- **Have your songs timed out.** This is important if you're going to be making a CD or a vinyl record. First, you want to make sure that your project can easily fit on a CD, if that's your release format. Most CDs have a total time of just under 80 minutes, so that time shouldn't be much of a problem unless it's a concert or a double album. When mastering for vinyl, cumulative time is important because the mastering engineer must know the total time per side before he starts cutting. Due to the physical limitations of the disc, you're limited to a maximum of about 25 minutes per side if you want the record to be nice and loud.

Mastering Session Documentation

You'll make it easier on yourself and your mastering person (even if that's you) if everything is well documented, and you'll save yourself some money, too. Here's what to include:

- **The title of the album and songs.** This should include the final titles, not shortened working titles, using the exact spelling that will appear on the final product.
- **The metadata, including the artist information,** especially if you want the mastering facility to make MP3 files for you.
- **Any flaws, digital errors, distortion, bad edits, fades** or anything out of the ordinary that a file might have.
- **Any FTP or shipping instructions** to send your master to a replicator.
- **Any ISRC and UPC codes.** We'll go over both in depth in Chapter 6.
- **Properly ID'd files.** Make sure that all files are properly titled for easy identification (especially if you're not there during the session), including alternative mixes.
- **A mastering reference.** Providing the mastering engineer with a commercially released CD that has a sound you really like may give him an idea on how to approach the job.

Having the right documentation can make your mastering session go a lot faster, which can be important especially when you're trying to make a release date. All it takes is a little bit of forethought and preparation.

Why Alternative Mixes Can Be Essential During Mastering

Even though mixes in a DAW can be almost instantly recalled and changed, most mixers still print alternate mixes to make ultra-quick fixes during mastering possible.

While alternate mixes with a vocal up and down a dB used to be the norm, today's mixers find that three types of alternate mixes can accomplish most fixes:

- **The instrumental mix.** This is often used to clean up objectionable lyrics on a song by editing in a small piece over the final mix. That way, the mix sounds a lot better than if a word is bleeped out with an audio tone. It's also sometimes used for licensing to television shows.

- **The acappella mix.** By using a combination of the instrumental mix with the acappella mix, it's possible to raise or lower a word that might be too loud or masked.
- **The TV mix.** The TV mix has everything but the lead vocal, so the artist or band can appear on television and sing live against a prerecorded background. Sometimes it's used instead of an instrumental mix.

While editing may be an overlooked skill of the mastering engineer, it can come in handy when alternate mixes are available. Even though mix fixes in a DAW can be fast, sometimes using the alternate mixes to make a fix can be even faster. If you're on the fence about the level or EQ of a certain instrument in the mix, print a couple of options.

Chapter 4

Monitoring For Mastering

The heart and soul of the mastering signal chain are the loudspeakers. More than any one device, these are the main link of the mastering engineer to both the reference point of the outside world and the possible deficiencies of the source material. More great pains go into choosing the monitoring system than just about any other piece of gear in the mastering studio.

Probably the one biggest and most important piece of equipment that a mastering engineer can have is his monitor, and he has to understand that monitor and really know where it should be. If you know the monitor and you've lived with it for a long time, then you're probably going to be able to make good recordings.

—Bernie Grundman

The Acoustic Environment

Having the finest reproduction equipment is all for naught unless the acoustic environment in which they're placed is optimized. Because of this, more time, attention, and expense is initially spent on the acoustic space than on virtually any other aspect in a high-end mastering facility.

I think a lot of people have heard about the effort we've gone through to make our room as acoustically perfect as possible. Many times people come into the room and go, "Oh my God!" or something like that.

—Bob Ludwig

That said, when it comes to mastering your own material in your personal studio, the single greatest impediment to doing an acceptable job can be your monitoring environment, which is why it's so important to try to improve it.

I know, you can't afford George Augspurger to trick out your garage (he's considered the father of modern studio acoustic design), and those Oceanway monitors are still about \$40,000 out of reach. Don't worry, there's still hope.

Let's Fix Your Listening Area

The listening environment is probably the most overlooked part of most home studios. While it's easy to spend a lot of money trying to improve your listening area, here are a

few zero-cost placement tips that in some cases can really make a difference.

- **Avoid placing speakers up against a wall.** The farther away you can get from the wall, the smoother the monitor speaker response will be, especially with the frequencies below 100Hz.
- **Avoid the corners of the room.** A corner reinforces the low end even more than when placed against a wall. The worst case is if only one speaker is in the corner, which will cause the low-end response of your system to be lopsided.
- **Avoid being closer to one wall of the room than the other.** If one speaker is closer to a side wall than the other, you'll get a totally different frequency response between the two because the reflections from the wall are different than on the other side. It's best to set up directly in the center of the room if possible.
- **Avoid different types of wall absorption.** If one side of the room uses a wall material that's soft and absorbent while the other is hard and reflective, you'll have an unbalanced stereo image because one side will be brighter than the other. Try to make the walls on each side of the speakers the same in terms of the amount of absorption that is used.

The above tips aren't a cure-all for big acoustic problems, and they won't do anything for your isolation (it still takes big bucks for that), but you'd be surprised how much better things can sound by just moving the gear around the room.

You can find out additional tips about how to improve your listening environment in my book, *The Studio Builder's Handbook* (Alfred Music) or by watching my *Music Studio Setup and Acoustics* video series on [lynda.com](https://www.lynda.com).

The Monitors

Pro mastering facilities will always choose a mastering monitor with a wide and flat frequency response. Wide frequency response is especially important on the bottom end of the frequency spectrum, which means that a rather large monitor is required, perhaps with an additional subwoofer as well. This means that many of the common monitors used in recording and mixing, especially nearfields, will not provide the frequency response required for a typical mastering scenario.

Smooth frequency response is important for a number of reasons. First, an inaccurate response will result in inaccurate equalization in order to compensate. It will also probably mean you'll overuse the EQ in an unconscious attempt to overcome the deficiencies of the monitors themselves.

Let's say that your monitors have a bit of a dip at 2kHz (not uncommon, since that's about the crossover point of most two-way nearfield monitors). While recording and mixing, you boost 2k to compensate for what you're not hearing. Now, it might sound okay on these monitors during mastering, but if you play it back on another set of speakers, you might find that the midrange is tearing your head off.

Now let's bring the environment into the equation, which compounds the problem. Let's say that between the monitors you're listening to (like a typical two-way system with a 6- or 8-inch woofer) and your room, you're not hearing anything below 60Hz or so. To compensate, you add +8dB of 60Hz so it sounds the way you think it should sound. If you master on the same monitors in the same environment, you'll never realize that when you get the song outside of your studio, it will be a big booming mess.

Large monitors with a lot of power behind them are not for loud playback, but for clean and detailed, distortion-free level. These monitors never sound loud, they just get bigger and bigger sounding and yet reveal every nuance of the music.

Although the selection of monitors is a very subjective and personal issue (just like in recording), some brand names repeatedly pop up in major mastering houses. These include Tannoy, Dunlavy, B&W, Lipinski, and Duntech.

One reason I've always tried to get the very best speaker I can is I've found that when something sounds really right on an accurate speaker, it tends to sound right on a wide variety of speakers.

—Bob Ludwig

It's not that we're going for the biggest or the most powerful sound; we're going for neutral because we really want to hear how one tune compares to the other in an album. We want to hear what we're doing when we add just a half dB at 5k or 10k. A lot of speakers nowadays have a lot of coloration and they're kind of fun to listen to, but boy, it's hard to hear those subtle little differences.

—Bernie Grundman

Basic Monitor Setup

Too often, musicians and engineers haphazardly set up their monitors, and this is a leading cause of mix and mastering problems later on down the line. How the monitors are placed can make an enormous difference in the frequency balance and stereo field and should be addressed before you get into any serious listening. Here are a few things to experiment with before you settle on the exact placement.

- **Check the distance between the monitors.** If the monitors are too close together, the stereo field will be smeared with no clear spatial definition. If the monitors are too far apart, the focal point or “sweet spot” will be too far behind you, and you’ll

hear the left or the right side distinctly, but not both together as one. A rule of thumb is that the speakers should be as far apart as the distance from the listening position. That is, if you're 4 feet away from the monitors, then start by moving them 4 feet apart so that you make an equilateral triangle between you and the two monitors (see Figure 4.1). A simple tape measure will work fine to get it close. You can adjust them either in or out from there.




Figure 4.1: Set the monitors in an equilateral triangle

© 2017 Bobby Owsinski

That being said, it's been found that 67-1/2 inches from tweeter to tweeter at the distance of a console meter bridge seems to be an optimum distance between speakers, and focuses the speakers just behind your head (which is exactly what you want).

- **A really quick setup that we used to use in the days of consoles is to open your arms as wide as possible to each side and place the monitors at the tips of the fingers of each hand.** This seemed to work well because of the built-in depth that the console would provide, but it doesn't really apply in these days of workstations, where the monitors are a lot closer to you than ever. If that's the case, go back to the equilateral triangle outlined above.
- **Check the angle of the monitors.** Improper angling will also cause smearing of the stereo field, which could mean that you'll have a lack of instrument definition as a result. The correct angle is determined strictly by taste, with some mixers preferring the monitors to be angled directly at their mixing position and others

preferring the focal point (the point where the sound from the tweeters converges) anywhere from a foot to about 3 feet behind them to eliminate some of the “hype” of the speakers.

TIP: *A great trick for getting excellent left/right imaging is to mount a mirror over each tweeter and adjust the speakers so that you can see your face clearly in both mirrors at the same time when you are in your listening position.*

- **Check how the monitors are mounted.** Monitors that are mounted directly on top of a desk or console meter bridge without any decoupling are subject to comb-filter effects, especially in the low end. That is, the sound from the monitor causes the desk or console to resonate, causing both the desk and the speaker to interact at certain frequencies either add or subtract. This is what’s known as phase cancellation, and it causes a subtle yet very real blurring of the sound. As a result, it will be just a little harder to hear your low end distinctly, which makes it more difficult to EQ. Phase cancellation can be more or less severe depending on whether the speakers are mounted directly on the desk or metal meter bridge or they are mounted on a piece of carpet or similar material (which is very popular).
- One of the quickest ways to improve the sound of your monitor system is to **decouple your speakers** from whatever they’re sitting on. This can be done with a commercial product, such as Primacoustic’s Recoil Stabilizers (see Figure 4.2), or you can make something similar relatively cheaply with some open-cell (closed-cell will work, too) neoprene or even some mouse pads.



Figure 4.2: Equator D5 with decoupler
© 2017 Bobby Owsinski

Decoupling your subwoofers (if you're using them) from the floor can really help too. Although sometimes the coupling with the floor can make your low end feel bigger, it will be a lot clearer and distinct if decoupled. Auralex even has a product for this called the SubDude HD, although you can probably put together a DIY setup that can work just as well.

Regardless of the brand, model, and type of speakers you use, decoupling is a cheap and easy way to improve your sound right away.

TIP: *The best solution is to mount your monitors on stands just directly behind the desk or meter bridge. Not only will this improve the low-frequency decoupling, but it can greatly decrease the unwanted reflections off the desk or console.*

- **Check how the monitor parameters are set.** Many monitors are meant to be

used in an upright position, yet users frequently will lay them down on their sides. This results in a variety of acoustic anomalies that deteriorate the sound. Also, with powered monitors, be sure that the parameter controls of both monitors are set correctly for the application and are the same on each (see Figure 4.3).



Figure 4.3: Monitor speaker parameter controls
© 2017 Bobby Owsinski

- **Check the position of the tweeters.** Most engineers prefer that the tweeters of a two- or three-way speaker system be on the outside, thereby widening the stereo field. Occasionally, tweeters to the inside work, but this usually results in a smearing of the stereo image. Experiment with both, however, because you never know exactly what will work until you try it (see Figure 4.4).




Figure 4.4: Tweeter position
© 2017 Bobby Owsinski

On The Bottom

Getting a project to have enough low end so that it translates well to speaker systems of all sizes is one of the things that mastering engineers pride themselves on, and it's one of the reasons why nearfield or even popular soffit-mounted large monitors are inadequate for mastering. The only way that you can properly tune the low end of a track is if you can hear it; therefore, a monitor with a frequency response that goes down to at least 40Hz is definitely required if you want to hear the low end well enough to work on it.

To hear that last octave on the bottom, many mastering engineers are now adding subwoofers to their playback systems. A great debate rages as to whether a single subwoofer or stereo subwoofers are required for this purpose. Those that say stereo subs are a must insist that enough directional response occurs at lower frequencies to require a stereo pair. There is also a sense of envelopment that better approximates the realism of a live event with stereo subs. Either way, the placement of the subwoofer(s) is of vital importance due to the standing waves of the control room at low frequencies.

Three Steps To Adding A Subwoofer

It's not unusual for musicians and engineers to crave more bottom end in the speakers that they're using in their personal studios. As a result, the first thing they think about is adding a subwoofer to their monitor system. That's all well and good, but there are a few steps you can follow that might make your venture into low-frequency territory a lot easier.

1. Do you really need a subwoofer? Before you make that purchase, it's a good idea to be sure that a sub is actually necessary. Here are a couple of things to check out first:

- **Are you monitoring at a loud enough level?** This is a trap that people with home studios fall into - they don't listen loudly enough, at least for a short period of time. First of all, if you monitor too quietly, your ears begin to emphasize the mid-frequencies. This is great for balance but bad for judging the low end of a song. Crank up your monitors to a moderately loud level, at least when you're working on the low-frequency end of the spectrum. If you still don't have enough low end, go on to the next point.
- **Do you have an acoustic problem in your room?** Chances are that either your monitors are too close to the wall or they're placed at a point of the room length where standing waves cause some of the low end to cancel out. This is more likely to be the cause of just one area of the low-frequency spectrum rather than the entire low end though. Just to be safe, move your speakers a foot or so backward and forward to see whether you get some of the low end back. If not, move on to Step 2.

2. Purchase a subwoofer from the same manufacturer as your main monitors. The easiest way to get a smooth-sounding low end that doesn't cause you more grief than it's worth is to buy a sub to match the monitors that you use most of the time. That means if you're using JBLs, choose a JBL sub that's made specifically for that system; if you're using Genelecs, do the same; KRKs, the same, and so on. This will make a huge difference, especially at the crossover frequency point where the mains cross over to the sub. It's usually extremely difficult to get that area to sound natural if you mix brands.

3. Calibrate your sub correctly. Most musicians and engineers that choose to use a sub just randomly dial in the level. You might get lucky and get it right, but it's more than likely that your level will be off, causing a number of unbalanced-sounding mixes until you finally figure it out.

Calibrating Your Sub To Your System

1. With the sub bypassed, send pink noise to your main monitors. At the listening position and while listening to *one monitor only*, use an SPL meter (just about any of them will do to get you in the ballpark, even an iPhone app) and adjust the level of the monitor until it reads 85dB. The SPL meter should be set on *C Weight* and *Slow*. Repeat on the other channel and set that so it also reads 85dB.

2. Turn off the main monitors. Send pink noise just to the subwoofer. Set the level of the SPL meter so it reads 79dB. Although it may seem like it will be lower in level, 79dB works because there are fewer bands of low frequencies than high (three for the low and eight for the high), so this number takes that into account. You might have to tweak the level up or down a dB, but this will get you into the ballpark.
3. If there's a polarity switch on the sub, try both positions and see which one has the most bass or sounds the smoothest in the crossover area. That's the one to select.

If you follow these steps, you'll find that integrating a subwoofer into your system (if you decide you need one) will be as painless as possible.

Placing The Subwoofer

Here's a method that will get you in the ballpark, although you'll have to do a bit of experimenting. Keep in mind that this method is for single subwoofer use.

1. Place the subwoofer in the engineer's listening position behind the console.
2. Feed pink noise only into the subwoofer at the desired reference level (85dB SPL should do it, but the level isn't critical).
3. Walk around the room near your main monitor speakers until you find the spot where the bass is the loudest. That's the spot to place the sub. For more level, move it toward the back wall or corner, but be careful because this could provide a peak at only one frequency. We're looking for the smoothest response possible (which may not be possible without the aid of a qualified acoustic consultant).

Amplifiers

While the trend for most recording-style monitors is toward self-powered units, many of the preferred monitor speakers used in a pro mastering environment still require an outboard amplifier, and a rather large one at that. It's not uncommon to see amplifiers of well over 1,000 watts per channel in a mastering situation. This is not for level (since most mastering engineers don't listen all that loudly), but more for headroom, so that the peaks of the music induce nary a hint of distortion. Since many speakers used in

mastering are rather inefficient as well, this extra amount of power can compensate for the difference.

Although many power amps that are standard in professional recording, such as Manley, Bryston, and Hafler, are frequently used, it's not uncommon to see audiophile units such as Cello, Threshold, Krell, and Chevin.

When I started Gateway, I got another pair of Duntech Sovereigns and a new pair of Cello Performance Mark II amplifiers this time. These are the amps that will put out like 6,000-watt peaks. One never listens that loudly, but when you listen, it sounds as though there's an unlimited source of power attached to the speakers. You're never straining the amp, ever.

—Bob Ludwig

Listening Techniques For Mastering

Regardless of what kind of monitors or room you have to work with, there are some proven techniques that will yield reasonable results even under the worst conditions. These all depend upon your ears, which are still the primary ingredient in mastering, and not the gear.

- **Listen to some CDs that you love first.** You want to listen to the highest quality program that you can get, so this is one time when the CD beats an MP3 (although a FLAC file or a high-res WAV or AIFF file will work, too). Listen to a favorite recording or two that you know really well and understand how it sounds on your system so that you can establish a reference point. This will help you from over-EQing or compressing too much. If you do nothing else, this one trick will help you more than anything else.
- **Establish two different listening levels.** You need one level that you consider fairly loud, where you can easily hear how the lower-frequency instruments (especially bass and drums) sit with each other, and another that's at a much lower listening level, somewhere near the point where you can hold a conversation while the music is playing.
- **Use these two listening levels only.** Mark them down on your volume control, make a note where the level is in the software, and do whatever you have to do to make these two levels repeatable. The levels are somewhat arbitrary in that they depend on your monitors and your environment, but the idea is that you want one level that's loud enough for you to gauge the low end and another that's quiet enough that you can hear the tonal balance. If you listen at varying levels, your reference point will be thrown off, and you'll never be sure exactly what you're listening to, which is why you keep it to two levels only.

- **Use two sets of speakers: a large set and a small set.** The only way you can ever be sure of how things really sound is if you have two different sets to reference against in the event that you don't have a large super-high-res reference monitor. Even if the largest speaker system that you can afford is a two-way bookshelf speaker with a 6-inch woofer, you should have an even smaller set to reference against. Although not the best, even a pair of computer speakers will do as long as you can feed them from the same source as your larger monitors. Mastering pros usually use a huge set of monitors with double 15-inch woofers plus a subwoofer, and an average two-way bookshelf speaker or something even smaller. Even if you have more than two sets of monitors available, limit your listening choices during mastering so you don't confuse yourself and end up chasing your tail.
- **If you attempt to master your own mix, use a different set of speakers than what you mixed on.** It doesn't matter whether you're mastering in your bedroom or in a million-dollar SSL room with George Augspurger acoustics, you're at a huge disadvantage if you master on the same monitors that you mixed on. Why? Because all monitors have flaws, and if you use the same monitors for mastering, you're either overlooking the problem or just making it worse. This is really important because if you use the same monitors, you'll only be compounding any frequency response problems that the speakers might have in the first place.

TIP: *There are some projects that you're better off remixing instead of trying to fix during mastering. Don't be afraid to send the project back (or to redo it yourself, if you were the mixer) rather than spending a lot of time making it sound different, not better. Believe it or not, pro mastering engineers make this suggestion all the time, even to some of the top mixers.*

Monitors Versus Headphones

Sometimes it's just not possible to listen to your monitors when you're working on music at home. When it's late at night and your kids, significant other, or neighbors are in the next room, separated only by paper-thin walls, you have no choice but to try to listen on headphones.

Mastering (or mixing, for that matter) on headphones does have four significant downsides, though:

- **Your ears get tired.** You can't wear them for as long as you need to (8, 10, 12

hours) before your head and ears get tired from the extra weight.

- **It's easy to get ear fatigue.** You have a tendency to turn them up, which can lead to some quick ear fatigue, again limiting your ability to listen for long periods.
- **You get a false sense of what the song sounds like.** Because most of the more expensive professional headphones really sound great, you get a false sense of what you're listening to (especially on the low end), and it causes you not to work as hard at getting the frequency balance right.
- **It might not translate to speakers.** If you master something only on headphones, it might not work when played back on normal monitors.

Although it's really helpful to know what your master might sound like on headphones, you still need to do most of your work on speakers to be sure that it will translate to a playback medium of any type.

Chapter 5

Mastering Tools

All tools created for mastering, regardless of whether they're analog or digital, have two major features in common: extremely high sonic quality and repeatability. The sonic quality is a must in that any device in either the monitor or the processing chain should have the least possible effect on the signal. Repeatability is important (although less so now than in the days of vinyl) in that the exact settings must be repeated in the event that a project is redone (as in the case when additional masters or changes are called for weeks later).

While this feature isn't much of a problem in the digital domain because the settings can be memorized, many analog mastering devices are still used in pro facilities. As a result, these hardware devices require special mastering versions that have 1dB or less increment selections on the controls, which can seriously add to the cost of the device.

That said, the mastering that you'll most likely do will probably all be "in the box," and there are a variety of much lower-cost options to choose from. In this chapter, we'll look at each tool as well as their ideal placement in the audio signal path.

The Mastering Compressor

In mastering, the compressor is the primary way of raising the relative level of the program and giving the master both punch and strength. Relative level is how loud we perceive the volume rather than the absolute level that's on the meter.

Compressor Overview

A compressor is a dynamic level control that uses the input signal to determine the output level. The *Ratio* parameter controls the amount that the output level from the compressor will increase compared to the input level (see Figure 5.1). For instance, if the compression ratio is 4:1 (four to one), for every 4dB of level that goes into the compressor, only 1dB will come out once the signal reaches the threshold level (the point at which the compressor begins to work). If a gain ratio is set at 8:1, then for every 8dB that goes into the unit, only 1dB will come out its output. Some compressors have a fixed ratio, but the parameter is variable on most units from 1:1 (no compression) to as much as 100:1 (which makes it a limiter, a process that we'll look at later in this chapter). A *Threshold* control sets the input-level point where the compression will kick in. Under that point, no compression occurs.



Typical Compressor Controls

Figure 5.1: Compressor Attack and Release controls

© 2017 Bobby Owsinski

Most compressors have *Attack* and *Release* parameters. These controls determine how fast or slow the compressor reacts to the beginning (attack) and end (release) of the signal envelope. Many compressors have an *Auto* mode that automatically sets the attack and release according to the dynamics of the signal. Although Auto works relatively well, it still doesn't allow for the precise settings required by certain source material. Some compressors (such as the revered Teletronix LA-2A or dbx 160) have a fixed attack and release, which helps give the compressor a distinctive sound.

When a compressor actually compresses the signal, the level is decreased, so there's another control called *Make-Up Gain* or *Output* that allows the signal to be boosted back up to its original level or beyond.

Using The Compressor In Mastering

For mastering, the compression ratio of the mastering compressor is usually set very low, from about 1.5:1 to 3:1, in order to keep the compression fairly gentle-sounding. The higher the ratio, the more likely you'll hear the compressor work, which can cause the program to sound unnatural. That might be a good thing sometimes in recording and mixing where a certain color is desired, but mastering is trying to keep the sound of the original program as intact as possible.

The keys to getting the most out of a compressor in mastering are the Attack and Release controls, which have a tremendous overall effect on a mix and therefore are important to understand. Generally speaking, transient response and percussive sounds are affected

by the Attack control setting. Release is the time it takes for the gain to return to normal after compression occurs.

In a typical pop-style mix, a fast Attack setting will react to the drums and reduce the overall gain on each beat. If the Release is set very fast, then the gain will return to normal quickly but can have an audible effect by reducing some of the overall program level and attack of the drums in the mix.

As the Release is set faster the gain changes, which might cause the drums to be heard as “pumping,” which means that the level of the mix will increase and then decrease noticeably. Each time the dominant instrument starts or stops, it “pumps” the level of the mix up and down.

Mastering compressors that work best on a wide range of full program material generally have very smooth release curves and slow release times to minimize this pumping effect. Chapter 6, “Mastering Techniques,” will further detail the effects that Attack and Release settings can have on the program.

Widely used mastering compressor plugins include the Shadow Hills Mastering Compressor, PSP Zenon and Vintage Warmer, and the various versions of the Fairchild 670.

Hardware compressors include the Manley SLAM Master, the Shadow Hills Mastering Compressor, and the Maselec MLA-2.

Multi-Band Compression

Multi-band compression splits the input audio signal into multiple frequency bands, each with its own compressor. The main advantage of a multi-band is that a loud event in one frequency band won’t affect the gain reduction in the other bands. That means that something like a loud kick drum will cause the low frequencies to be compressed, but the mid and high frequencies are not affected. This allows you to get a more controlled, hotter signal with far less compression than with a typical single-band compressor.

The multi-band compressor is unique in that it can shape the timbre of a mix in ways that an EQ just can’t. By raising or lowering the *Level* control of each band and then adjusting the Crossover controls, the tonal quality of the mix will change in a way similar to an EQ, because only that particular band will be compressed.

For instance, if you wanted to just tighten up the low end, you’d increase the level of the low-frequency band while adjusting the low/mid frequency Crossover control to focus in on the exact frequencies you want to affect.

Frequently used multi-band software compressors include the Waves C6, the Universal

Audio UAD Precision Multi-Band (see Figure 5.2), or the multi-band compressor in iZotope Ozone 5. Hardware multi-band compressors include the Maselec MLA-3 and the Tube-Tech SMC 2BM.



Figure 5.2 : The Universal Audio UAD Precision Multi-Band Compressor plugin
© 2017 Bobby Owsinski (Source: Avid, Universal Audio)

The Mastering Limiter

One of the most essential tools for a mastering engineer is the limiter, since it's the principle way that high levels are achieved without digital overs.

Limiter Overview

A limiter is a compressor with a very high compression ratio and a very fast attack time, so it's able to catch the fast peaks of an audio signal. Any time the compression ratio is set to 10:1 or more, the result is considered limiting. While it's true that a compressor can be adjusted to work as a limiter, in mastering the limiter is usually a dedicated unit created specifically for the task.

A limiter can be thought of as a brick wall for level, allowing the signal to get only to a certain point and little more. Think of it like a governor that's sometimes used on trucks to make sure they don't go over the speed limit. After you hit 65 mph (or whatever the speed limit in your state is), no matter how much more you depress the gas pedal, you won't go any faster. It's the same with a limiter. When you hit the predetermined level, no matter how much more signal is supplied to the limiter, the level pretty much stays the same.

Using The Limiter In Mastering

To understand how a limiter works in mastering, you have to understand the composition

of a typical music program first. In general, the highest peak of the source program (the song in this case) determines the maximum level that can be achieved in a digital signal. Because many of these upper peaks have a very short duration, they can usually be reduced in level by several dB with minimal audible side effects. By controlling these peaks, the entire level of the program can then be raised several dB, resulting in a higher average signal level.

Most digital limiters used in mastering are what's known as *brick-wall* limiters. This means that no matter what happens, the signal will not exceed a certain predetermined level, and there will be no digital overs after its Ceiling level has been set. A brick-wall limiter is usually set anywhere from -0.1dB to -1.0dB , and once set, the level will never go beyond that point.

Many popular mastering limiter plugins are commonly used, from the Waves L1 and L2, to the Universal Audio Precision Limiter (see Figure 5.3), to the T-RackS Brickwall Limiter, to Massey L2007 and many others.



Figure 5.3: The Universal Audio UAD Precision Limiter
© 2017 Bobby Owsinski (Source: Universal Audio)

By setting a digital limiter correctly, the mastering engineer can gain at least several dB of apparent level just by the simple fact that the peaks in the program are now controlled.

Multi-Band Limiter

Multi-band limiters work in the same manner as multi-band compressors, splitting the audio spectrum into separate frequency bands that can be individually limited. This can provide more level without sounding as compressed as when a single-band limiter is used. The operations are identical to that of a multi-band compressor, except that the compression ratio is higher, making it a limiter instead of a compressor.

Typical mastering multi-band limiter plugins are the Waves L3, iZotope Ozone Multi-Band Limiter, and McDSP MC2000 and ML4000, among others.

The Mastering Equalizer

One of the most important duties of the mastering engineer is fixing the frequency balance of a project if it's required. Of course this is done with an equalizer, but the type used and the way it's driven are generally far different than during recording or mixing.

Most precision work in mastering is done using a parametric equalizer, which allows the engineer to select a frequency, the amount of boost and cut, and the width of bandwidth around the selected frequency that will be affected (known as Q).

Using The EQ in Mastering

Where in recording you might boost or cut large amounts of EQ anywhere from 3 to 15dB at a certain frequency, mastering is almost always done in very small increments, usually in tenths of a dB to 2 or 3dB at the very most. What you might see are a lot of small shots of EQ along the audio frequency band, but again in very small amounts.

For example, these might be something like +1dB at 30Hz, +0.5 at 60Hz, +0.2 at 120Hz, -0.5 at 800Hz, -0.7 at 2500Hz, +0.6 at 8kHz, and +1 at 12kHz. Notice that there's a little happening at a lot of places across the frequency spectrum.

Another technique that's used frequently is known as *feathering*. This means that rather than applying a large amount of EQ at a single frequency, small amounts are added at the frequencies adjoining the one being focused on. An example of this would be instead of adding +2dB at 100Hz, you would add +1.5dB at 100Hz and +0.5dB at 80 and 120Hz (Figure 5.4). This generally results in a smoother sound by not stressing any one area of the equalizer.



Figure 5.4: EQ feathering

© 2017 Bobby Owsinski (Source: Avid, Universal Audio)

Mastering is one area where large amounts of EQ are an indication that there's something wrong with the mix. Top mastering engineers will frequently send a mixer back to redo a mix since a fixed mix will sound better than one where the mastering engineer has to do major EQ surgery, and that's something you should consider as well. In mastering equalization, less is definitely more.

Hardware mastering equalizers differ from their recording counterparts in that they usually feature stepped rather than continuously variable controls in order to be able to repeat the settings, with the steps being in increments as little as 0.5dB, although 1dB is the increment that's mostly seen. Examples are the Manley Massive Passive (see the mastering version in Figure 5.5) and the Avalon AD2077.



Figure 5.5: The Manley Massive Passive mastering version
Courtesy of Manley Labs

Most equalizer plugins are inherently capable of 0.1dB steps, so the overall audio quality imparted to the program is more a factor in which one is chosen. EQs that impart their own “color” to the audio are usually passed over in favor of more transparent versions that change the sonic quality of the program very little. Examples are the Massenburg DesignWorks MDWEQ5 the Sonnox Oxford EQ, the Brainworx bx_digital V2, the PSP Neon HR, and many more.

TIP: Many equalizers have a high-resolution mode for increasing the smoothness of the equalizer, called Linear Phase. If available, this setting reduces the sonic coloration of the equalizer.

The Mastering De-Esser

Sibilance is a short burst of high-frequency energy where the S’s of a vocal are over-emphasized, which comes from a combination of mic technique by the vocalist, the type of mic used, and heavy compression on the vocal track and mix buss. Sibilance is generally felt to be highly undesirable, so a special type of compressor called a de-esser is used to suppress it.

Most de-essers have two main controls, *Threshold* and *Frequency*, which are used to compress only a very narrow band of frequencies anywhere between 3k and 10kHz to eliminate sibilance (see Figure 5.6).




Figure 5.6: Sibilance band between 3k and 10kHz

© 2017 Bobby Owsinski (Source: Avid)

Modern software de-essers are much more sophisticated than analog hardware de-essers, but the bulk of the setup still revolves around those two parameters. One frequently included feature is a Listen button that allows you to solo only the frequencies that are being compressed, which can be helpful in finding the exact band that's offensive (see Figure 5.7).



Figure 5.7: A de-esser with the Listen function

© 2017 Bobby Owsinski (Source: Avid)

While sibilant vocals are the usual reason for de-essing, sometimes a de-esser might be used to control an excessive high frequency from other instruments. Cymbals, guitars, and even the snare drum can occasionally benefit from this unique tool.

Typical de-esser plugins include the Massey De:Esser, McDSP DE555, and Waves DeEsser, among others.

Metering

Metering is extremely important in mastering, much more so than mixing, especially when you're trying to achieve hot levels. There are more metering tools available to the mastering engineer than the simple metering that we're used to during recording, because the mastering process requires a lot more visual input to tell you the things you need to know.

While you don't need all of the following meters to do a proper mastering job, they all do serve a purpose. You'll find that they'll all contribute a great deal to making a great master. Let's look at some of the meters that are frequently used.

The Peak Meter

The peak meter was created by the BBC when they realized that the common VU meter (see Figure 5.8) wasn't precisely telling an engineer exactly what the program signal was doing, which is especially important in broadcast, where over-modulation of the signal can bring a fine from the Federal Communications Commission.

The standard analog VU (which stands for Volume Units) meter, which was common on all professional audio gear until the late '90s, only shows the average level of a signal and has a very slow response time. As a result, you'd have to guess at the signal peaks, since they were too fast for the meter to accurately read them. A good example of this effect takes place during recording of a high-pitched percussion instrument, such as a triangle or a tambourine, where the signal is almost all peaks. The experienced engineer would record the instrument at a barely visible -20 on the VU meter to keep the recording from distorting. Couple the slow response of the VU meter with the fact that it's an analog, mechanical device that could easily be knocked out of calibration, and you can see the need for a new metering system.




Figure 5.8: A typical VU meter

© 2017 Bobby Owsinski

The peak meter, on the other hand, has an extremely fast response, which is almost fast enough to catch most peaks (more on this in a bit), and could be simulated on a digital display instead of using an actual meter (a hardware peak meter used to be very expensive before the digital age; see Figure 5.9). The peak meter also became a necessity for digital recording because any signal beyond 0dB could cause anything from a harsh edginess to a very nasty-sounding distortion. As a result, all peak meters have a red Over indicator that lets you know you've exceeded the zone of audibly clean level.




Figure 5.9: A typical digital peak meter

© 2017 Bobby Owsinski (Source: IK Multimedia)

There's also a not-too-well-known phenomenon called inter-sample distortion, where the signal peaks exceed 0dB between the samples of really hot signals and are never indicated by the Over indicator as a result. This can cause trouble when the song or program is played back later on a CD or MP3 player and the digital-to-analog convertor (D/A) is overloaded even though the Overload indicator never lights (which is why some mastered programs sound so harsh).

Peak meters can also be either somewhat accurate or very accurate, depending upon the resolution of the meter. A peak meter used in mastering would normally be calibrated in tenths of a dB, while some inexpensive implementations might have a resolution of 1dB or more (see Figure 5.9).

The fact is that many standard peak meters can't accurately tell the difference between a level of 0dBFS (FS = Full Scale) and an over, which can mean that an overload is occurring without you knowing. That's why mastering engineers rely on super-accurate peak metering that counts the number of samples in a row that have hit 0dB. This number can usually be set in a preference window from three to six samples. Three overload samples equals distortion that will last for only 33 microseconds at 44.1kHz, so it will probably be inaudible, and even six samples is difficult to hear. That said, this type of peak meter is much preferred during mastering.

The RMS Meter

Even when your peak meter is tickling 0dB, an RMS meter will settle at a point much lower since it's measuring the signal differently. We don't use RMS meters much these days, since a peak meter is much more precise, but in the pre-digital days, that's all that was available (since a VU meter is an RMS meter), and it's still what many older engineers are used to.

Today there are digital versions of the RMS meter (see Figure 5.10). RMS stands for the "root mean square" measurement of the voltage of the electronic signal, which roughly means the average. In an app like T-RackS CS, the RMS meter combines both left and right channels into a single display that measures the power of a signal.




Figure 5.10: The T-RackS digital RMS meter (on the bottom)

© 2017 Bobby Owsinski (Source: IK Multimedia)

The frequency response of an RMS meter is flat, which can give you a false sense of level if the song has a lot of low end. For that reason, it's best to read the meter in conjunction with the other meters to give you an idea of where you're at in terms of level and loudness (they're different, as you'll soon read).

TIP: *One thing that the RMS meter is very good at is telling you if two (or more) songs are approximately the same level. Don't rely solely on the meter for that, though, since you have to use your ears to make the final determination.*

The K-System Metering

Mastering engineer Bob Katz has developed a metering system with simultaneous peak and RMS displays that is supported by many metering packages (see Figure 5.11). The K-System has three different meter scales with the zero point at either -20, -14, or -12dB, depending upon the program that is being worked on. These three scales are named K-20, K-14, and K-12, and they are intended to help the engineer maintain sufficient headroom in the program to eliminate the possibility of clipping.




Figure 5.11: A K-System meter

© 2017 Bobby Owsinski

The K-20 shows 20dB of headroom above 0dB and is intended for theatrical mixes. K-14 shows 14dB headroom and is intended for music mixing and mastering. K-12 shows 12dB headroom and is intended for broadcast.

The Perceived Loudness Meter

The perceived-loudness meter is a newer meter now found in the mastering chain after being used in broadcasting for some years (see Figure 5.12). It determines how “loud” a program is by measuring all the frequencies that make up the program and then applying a weighting measurement to the frequency bands in the same proportions as our ear perceives them. For example, the ear is most sensitive in the 2k to 4kHz range, so frequencies at 100Hz or 10kHz will measure much higher in level than the 4kHz tone, yet they’ll seem equally loud.




Figure 5.12: A loudness meter (at the top)

© 2017 Bobby Owsinski (Source: IK Multimedia)

As said before, level and loudness are two different things. In the case of mastering, level is the signal voltages you read on a meter, and loudness is what you hear. Two different programs can have identical peak and RMS levels, yet one can still sound louder than the other.

Among the dedicated loudness meters are the Waves WLM Plus, the TC Electronic LM5D, the Nugen Audio MasterCheck, and on the analog side, the Dorrough 40-A.

The Phase Scope

The phase “scope” gets its name from the fact that in the early days of recording, the phase between the left and right channels of a program was checked by using an old-fashioned oscilloscope (see Figure 5.13), which was just called a “scope” for short.




Figure 5.13: A classic oscilloscope

© 2017 Bobby Owsinski

Phase is extremely important in a stereo signal, because if the left and right channels are not in phase, not only will the program sound odd, but instruments panned to the center (such as lead vocals and solos) may disappear if the stereo signal should ever be combined into mono.

While you may think that mono isn't used much these days, you'd be surprised. If your song is ever played on AM radio, it's in mono on 99 percent of the stations. On FM radio, if a station is far enough away from where you're listening, the stereo signal may collapse into mono because the signal strength is weak. On television, it's not uncommon for the stereo mix of the show to be automatically converted to mono on some networks.

Sometimes the settings in the iTunes player can be switched to mono or a stereo song can be ripped in mono, so they'll play back in mono on your playback device. Mono is everywhere, so it's a good thing to pay attention to the phase of your program.

Using The Phase Scope

The phase scope isn't very good for measuring absolute levels (that's why you have the other meters), but it does provide a wealth of information about stereo source positioning, and the relative phase and level between the two channels. The signals are displayed in a two-dimensional pattern along the X and Y axis called a Lissajous figure (see Figure 5.14). An identical signal on both channels results in a 180-degree vertical line representing a central mono signal (see Figure 5.15), while a true stereo signal will give you a more or less random figure that's always moving (see Figure 5.16).






Figure 5.14: A Lissajous figure on an oscilloscope
© 2017 Bobby Owsinski



A Mono Guitar Track On The Phase Scope

Figure 5.15: A mono signal on the phase scope
© 2017 Bobby Owsinski (Source: IK Multimedia)



A Stereo Signal On The Phase Scope

Figure 5.16: A stereo signal on the phase scope
© 2017 Bobby Owsinski (Source: IK Multimedia)

After you watch the phase scope for a while, you'll find that you can instantly tell a lot about the signal as you recognize the different shapes it can take on with different signals. Simpler and nearer sounds, such as mono one-shots or notes and chords, are illustrated by

thick, bold-looking, solid lines. Widen the stereo image, and you'll see a relatively wider and stringier image. Heavily reverberated or delayed sounds form shapeless images with lots of small dots. The complex arrangements normally found on most records will show all these and everything in between. The more defined the borders are, the more of the signal is above 0dB. As you can see, the phase scope shows everything from the width, phase, panning, amplitude, and even clipping info in the signal.

Many metering packages that feature a scope function are available, include PSP Stereo Analyser and the Flux Stereo Tool V3 (which is free), among others.

The Phase Correlation Meter

While the phase scope takes some time to get the hang of, the phase correlation meter is dead simple. Anything drawn toward the +1 side of the meter is in phase, while anything drawn toward the left-hand -1 side of the meter is out of phase (see Figure 5.17).



Figure 5.17: A phase correlation meter

© 2017 Bobby Owsinski (Source: Avid)


In general, any meter readings above 0 in the positive side of the scale have acceptable mono compatibility. A brief readout toward the negative side of the scale isn't necessarily a problem, but if the meter consistently sits in the negative side, it could represent a mono-compatibility issue. Keep in mind that the wider the stereo mix is, either because of panning or wide stereo reverbs, the more the phase correlation meter will tend to indicate toward the negative side. But as long as the signal stays mostly on the positive, your compatibility should be good to go.

If the phase correlation meter or phase scope indicates that there might be a mono-compatibility problem, it's important to immediately listen in mono to verify whether it's an issue and whether the track is acceptable. In the event that the out-of-phase condition is verified, sometimes flipping the phase of one channel can fix it, but usually a remix may be the only answer.

Examples are the phase correlation sections of both the Flux Stereo Tool and T-RackS Meters, among others.

The Spectrum Analyzer

The spectrum analyzer, sometimes known as the real-time analyzer, is an excellent tool for determining the frequency balance of your program by looking at it in octave or sub-octave portions (see Figure 5.18). It's especially effective for singling out particular frequencies that are too hot and for dialing in the low end.



Typical Spectrum Analyzer Response

Figure 5.18: A spectrum analyzer
© 2017 Bobby Owsinski (Source: IK Multimedia)

Contrary to what you might think, when you look at the analyzer, the object is not to aim for a totally flat response. The deep bass (below 40Hz) and the ultra-highs (above 10kHz) almost always have less energy compared to the other frequencies. It's very useful to look at songs, CDs, mixes, or any program that you think sounds really good and get a feel for what it looks like on the analyzer. Keep in mind that your mastering job will probably not look like your chosen example because each song is unique, but if the song is in the same genre, it might not be that far off by the time you've finished working your mastering magic.

The precision of a spectrum analyzer is determined by how many bands the audio spectrum is split into. One-octave analyzers can provide an overall picture of the frequency response, but 1/3- and 1/6-octave versions provide a higher resolution into what's happening frequency-wise within the program and are normally used in mastering.


TIP: *Most spectrum analyzers also allow the response to be adjusted, usually from very fast to a slow average of as many as 30 seconds to get a better picture of the response over time.*

Examples include the analyzers included in T-RackS and iZotope Ozone, and the True Audio TrueRTA, among others.

The Dynamic Range Meter

The dynamic range meter is very similar to a peak meter, but it adds the additional function of measuring the dynamic range of an audio signal. While it's easy to think that all music must be mastered the same way, different genres of music have different dynamic ranges that require a different mastering approach.

Dynamic range is a term for the degree of variation in volume level within a piece of music. Music with a low value, like a DR3, means that there's only a 3dB variation in level, so there's a lot of compression being used and not much variation in the dynamics (see Figure 5.19). Something that's more natural-sounding might have a value of about DR12 or more, meaning that there's at least a 12dB difference from the lowest to the highest peak in the song (see Figure 5.20).






Figure 5.19 (left): Dynamic range meter showing a value of DR3
© 2017 Bobby Owsinski (Source: KVR Audio, Avid)

Figure 5.20 (right): Dynamic range meter showing a value of DR12
© 2017 Bobby Owsinski (Source: KVR Audio, Avid)

Dynamic Ranges Of Different Genres Of Music

Different genres of music sound different at different DR levels. For instance, most acoustic music would be considered unpleasant-sounding at DR6, but that range might be perfectly acceptable for electronic music. With most pop, rock, R&B, and hip-hop, a

DR of 8 might be quite comfortable, but that just won't work for jazz, folk, country, or classical music, which sounds a lot better with a DR of at least 12 and probably a lot more.

Here's a list of different averages for different genres of music. As you can see, some genres, such as jazz and classical, have a large dynamic range, while others, such as hip-hop and rock, have a very narrow one.

Music Genre	Average Dynamic Range
Hip-hop	8.38
Rock	8.50
Latin	9.08
Electronic	9.33
Pop	9.60
Reggae	9.64
Funk	9.83
Blues	9.86
Jazz	11.20
Folk, world, country	11.32
Stage and screen	14.29
Classical	16.63
Children's	17.03

Dynamic range is one of the most important aspects of mastering, but it's all too often overlooked. As you go forward in the book, keep it in mind as a major tool in the sound of your finished project.

Among the examples of dynamic range meters are the Brainworx bx_meter or the Blouder Dynamic Range meter.

Convertors

Digital audio requires a device to encode the analog signal into a digital stream of 1s and 0s, and then after the digital audio has been recorded and processed, to turn the 1s and 0s back into an analog signal that we can listen to. These are called analog-to-digital (A/D for short) and digital-to-analog (D/A) convertors.

Most mastering studios are especially concerned with the quality of the D/A convertor (sometimes also known as a DAC), since the highest quality signal path to the monitors is a priority. As a result, the built-in convertors of most commonly used audio interfaces

are not sufficient, and most facilities opt for stand-alone outboard convertors.

One criterion for a mastering DAC is that it must operate at a wide range of sample rates, from 44.1kHz to 192kHz, which most devices on the market currently do. In the future, 384kHz may also be a required option.

Because each brand has a slightly different sound (just like most other pieces of gear), major mastering facilities may have numerous versions of each type available for a particular type of music. Among the current popular converters are Prism Sound, Lavry Engineering (see Figure 5.21), Mytek, Apogee, and Benchmark Media, among others.



*Figure 5.21: Lavry 3000S analog-to-digital converter
Courtesy of Lavry Engineering*

Unless you're mastering from an analog source like tape, or doing processing outside of the DAW with an outboard analog device, external analog-to-digital convertors are not necessary for mastering.

Consoles/Monitor Control

Although mastering consoles (sometimes referred to as transfer consoles) at one time were much more sophisticated and the centerpiece of the mastering studio, these days mastering consoles are more control devices that switch between different input and speaker sources and control the monitor level. That said, once again the emphasis is on a very high-quality signal path that degrades the signal as little as possible.

Back in the analog days when vinyl records ruled, a mastering console was a sophisticated device with two sets of equalizers and, in many cases, built-in compression. As mastering moved into the digital age, mastering consoles were some of the first pieces of gear to become totally digital, with all EQ, compression, and limiting done on-board.

Since the vast majority of processing is done in the DAW these days, all that's required is the monitor section to control the volume level of the monitors, switch between different monitors, and control input sources.

As with the digital-to-analog convertor, a major criterion for a mastering monitor controller is that it must operate at a wide range of sample rates, from 44.1kHz to 192kHz, which most devices on the market currently do. In the future, 384kHz may also be a required option.

Due to the unique nature and relatively small size of the mastering market, not many companies currently manufacture dedicated mastering consoles/monitor controllers. Among the manufacturers are Crookwood, Maselec, Dangerous Music, and Manley Labs.



Figure 5.22: The SPL DMC
Courtesy of Sound Performance Lab

The Digital Audio Workstation

Although not always the case, the digital audio workstation (DAW) has now become the heart and soul of the mastering studio, allowing the engineer to complete tasks such as editing and sequencing with far greater ease than was ever thought possible. Plus, the DAW allows new tasks to be carried out in ways that couldn't even be conceived of only 10 years ago.

Mastering DAWs

Although in a pinch just about any DAW can be used for mastering, a few manufacturers

have established themselves as the mastering engineer's favorites, primarily because dedicated mastering features are included. These features are mostly CD-specific and include DDP export, dither, and PQ code insertion and editing, all of which will be covered in Chapter 7, "Mastering for CD."

Among the DAWs that provide these functions are Steinberg WaveLab, Sonic soundBlade, Sonoris DDP Creator, DSP-Quattro, and SADiE, among others.

TIP: *Many plugins raise the sample rate in order to more precisely perform the digital processing, which is called upsampling. This results in a cleaner sound.*

There are also a number of dedicated software mastering suites available that provide precision processing for mastering, including iZotope Ozone and Waves, among others. Standalone mastering apps on the computer desktop include Roxio Toast, CD Architect, and T-RackS.

Other Tools

There are also other tools sometimes used in the mastering environment that don't fall under one of the previous categories.

Stereo Enhancement

Many times the mix seems like it's either too wide or too narrow, so a stereo enhancer is called for to adjust the width of the stereo field. This is more of a fix-it tool than something that a mastering engineer would use every day, but it does come in handy during those rare times when the stereo field needs manipulating.

Examples include the PSP StereoEnhancer and StereoController, iZotope Ozone, Waves Center, and S1 Stereo Imager.

M-S Processing

Some processors feature an M-S mode, which uses a mid-side matrix that allows you to select where the process works within the stereo image by manipulating the in and out-of-phase signals.

The Mid assigns the processing to the center of the stereo image as usual, but the Side makes it seem as if the processing is occurring at the outside edges of the stereo spectrum. Most of the time you'll find that the Mid mode works best, but occasionally the Side mode will allow you to add processing in a way that works better, such as brightening up a track without affecting a lead vocal.

M-S processing can be found on many T-RackS processors as well as Brainworx Control V2, among others.

Mono

As previously stated earlier in this chapter, it's always a good idea to check your work in mono. If you have an outboard monitor controller, chances are that it's equipped with a mono switch, but if that's not something yet in your gear arsenal, then you have to switch to mono in software. If the processors in your mastering chain don't provide this ability, a great little utility is Brainworx bx_solo (and it's free), which can also be used for stereo width adjustment.

Chapter 6

Mastering Techniques

Now that you've seen the basic philosophy of mastering, let's tackle the creative aspects. The actual mechanics of mastering can be broken down into a number of functions, namely maximizing the level of a song or songs, adjusting the frequency balance if necessary, performing any editing, adding fades and spreads, and inserting PQ codes, ISRC codes, and metadata.

What really separates the upper-echelon mastering engineers from the rest is the ability to make the music (any kind of music) as big and loud and tonally balanced as possible, but with the taste to know how far to take those operations. The mastering related DAW functions that they use, on the other hand, are somewhat mechanical, and usually don't get the same amount of attention as the more creative functions. We'll look at all of those techniques in this chapter, but first let's look at the basic approach used by most pro mastering engineers.

The Basic Mastering Technique

If you were to ask a number of the best mastering engineers what their general approach to mastering was, you'd get mostly the same answer.

1. Listen to all the tracks. If you're listening to a collection of tracks, such as an album, the first thing to do is listen to brief durations of each song (10 to 20 seconds should be enough) to find out which tracks are louder than the others, which ones are mixed better, and which ones have better frequency balances. By doing this, you can tell which songs sound similar and which ones stick out. Inevitably, you'll find that unless you're working on a compilation album where all the songs were done by different production teams, the majority of the songs will have a similar feel to them, and these are the ones to begin with. After you feel pretty good about how these feel, you'll find it easier to get the outliers to sound like the majority than the other way around.

2. Listen to the mix as a whole, instead of hearing the individual parts. Don't listen like a mixer, don't listen like an arranger, and don't listen like a songwriter. Good mastering engineers have the ability to divorce themselves from the inner workings of the song and hear it as a whole, just like the listening public does.

3. Find the most important element. On most modern radio-oriented songs, the vocal is the most important element, unless the song is an instrumental. That means one of your

jobs is trying to make sure that the vocal can be distinguished clearly.

4. Have an idea of where you want to go. Before you go twisting parameter controls, try to have an idea of what you'd like the track to sound like when you're finished. Ask yourself the following questions:

- Is there a frequency that seems to be sticking out?
- Are there frequencies that seem to be missing?
- Is the track punchy enough?
- Is the track loud enough?
- Can you hear the lead element distinctly?

5. Raise the level of the track first. Unless you're extremely confident that you can hear a wide frequency spectrum on your monitors (especially the low end), concentrate on raising the volume instead EQing. You'll keep yourself out of trouble that way. If you feel that you must EQ, refer to the section of the EQing later in the chapter.

6. Adjust the song levels so they match. One of the most important jobs in mastering is to take a collection of songs, like an album, and make sure each has the same relative level. Remember that you want to be sure that all the songs sound about the same level at their loudest. Do this by listening back and forth to all the songs and making small adjustments in level as necessary.

Making A Loud Master

The amount of perceived audio volume, or level, without distortion (on either an audio file, a CD, a vinyl record, or any other audio-delivery method yet to be created) is one of the things that many top mastering engineers pride themselves on. Notice the qualifying words “without distortion,” since that is indeed the trick: to make the music as loud as possible (and thereby competitive with other products on the market) while still sounding natural. Be aware that this generally applies to modern pop/rock/R&B/urban genres and not as often to classical or jazz, whose listeners much prefer a wider dynamic range where maximum level is not a factor.

Competitive Level

The volume/level wars that we experience today really began way back in the vinyl era of the '50s, when it was discovered that if a record played louder than the others on the radio, the listeners would perceive it to be better-sounding and therefore make it a hit. Since then, it has been the charge of the mastering engineer to make any song intended for

radio as loud as possible in whatever way he can.

This also applies to situations other than the radio. Take the MP3 player, CD changer, or streaming music playlist, for instance. Most artists, producers, and labels certainly don't want one of their releases to play at a quieter level than their competitor's because of the perception (and not necessarily the truth) that it won't sound as good if it's not as loud.

The limitation of how loud a "record" (we'll use this term generically) can actually sound is determined by the delivery medium to the consumer. In the days of vinyl records, if a mix was too loud, the stylus would vibrate so much that it would lift right out of the grooves and the record would skip. When mixing too hot to analog tape, the sound would begin to softly distort and the high frequencies would disappear (although many engineers and artists actually like this effect). When digital audio and CDs came along, any attempt to mix beyond 0dBFS would result in terrible distortion as a result of digital overs (nobody likes this effect).

As you can see, trying to squeeze every ounce of level out of the track is a lot harder than it seems, and that's where the art of mastering comes in.

Level Technique #1: The Compressor-Limiter Tandem

The bulk of the audio-level work today is done by a combination of two of the mastering engineer's primary tools: the compressor and the limiter (see Figure 6.1). The compressor is used to control and increase the level of the source audio, while the limiter controls the instantaneous peaks. Remember that the sound of both the compressor and the limiter will have an effect on the final audio quality, especially if you push them hard. Here's how you raise the level:

4. Set the master level on the limiter to -0.1dB to contain the peaks and avoid digital overs.
5. Set a compressor at a ratio of around 2:1 or 3:1 to gain apparent level. Generally speaking, the trick with compression in mastering is to use either a slow release time or one that's timed to the drums, and less (usually way less) than 3dB of compression.
6. Adjust the attack time to let the desired amount of transients through. The slower the attack time, the punchier the sound (generally speaking).
7. Adjust the release time to avoid hearing any pumping. Time it to the track to keep it punchy-sounding. Set it to slow to keep it smooth-sounding.
8. Increase the level of the program to the desired level by increasing the compressor's

Output control.

Remember that the less limiting you add, the better it will usually sound. Most of the gain and punch comes from the compressor.




Figure 6.1: The compressor-limiter tandem
© 2017 Bobby Owsinski (Source: IK Multimedia, Universal Audio)

TIP: Use a multi-band compressor and/or limiter to increase the level without hearing as much of the side effects from compression or limiting.

Level Technique #2: Multi-Compressor Packages

Some mastering engineers dislike the sound of a limiter so much that they'll do anything not to use one. One of the ways this is possible without resulting in the red overload LED continually lighting is by using multiple compressors instead. In this case, each compressor would be different and would use slightly different compression ratios in order to exert the same kind of control over the signal while increasing the level (see Figure 6.2).




Figure 6.2: Multi-compressor packages

© 2017 Bobby Owsinski (Source: IK Multimedia, Universal Audio)

Advanced Level Techniques

Another way to achieve a high signal level with a minimum of limiting is to insert the limiter only on the sections of the songs with peaks. That means finding the brief periods within the song with the peaks are strong enough to trigger the limiter, and either automate the limiter so it's inserted into the signal path at that point, or editing those sections and processing them separately.

I try not to use a limiter if I'm not hearing any distortion. When I run into situations when I need to use a limiter, I'll just use it only on the portions of the song that need it. What I do is go in and slice up a song and just smooth out only the rough edges. I can put up to 300 tiny edits in one song if need be. If it's an open track that has to be loud, I'll just cut all the tiny pieces that need limiting and limit only those. That way those sections go by so fast that your ear can't hear it as it flies by. It gets rid of any overload crackles and keeps the kick hitting hard. It's time-consuming, but I don't mind doing it if it comes out better. It actually goes a lot faster than you think once you have an ear for what to listen for.

—Gene Grimaldi

The Effects Of Hypercompression

Over the years it's become easier and easier to get a record that's hotter and hotter in perceived level, mostly because of new digital technology that has resulted in better and more effective limiters. Today's digital "look ahead" limiters make it easy to set a maximum level (usually at -0.1dBFS) and never worry about digital overs and distortion again, but this can come at a great cost in audio quality, depending on the situation.

Too much bus compression or over-limiting when either mixing or mastering results in what's become known as *hypercompression*. Hypercompression is to be avoided at all costs because:

- For the most part, it can't be undone later.
- It can suck the life out of a song, making it weaker-sounding instead of punchier.
- Lossy codecs like MP3 have a hard time encoding hypercompressed material and insert unwanted side effects as a result.
- It's known to cause listener fatigue, so the consumer won't want to listen to your record for as long or as many times.
- A hypercompressed track can actually sound worse over the radio because of the way it interacts with the broadcast processors at the station.

A hypercompressed track has little or no dynamics, leaving it loud but lifeless and unexciting. On a DAW, it's a constant waveform that fills up the DAW region. Figure 6.3 shows how the levels have changed on recordings over the years.




Figure 6.3: From very little compression to hypercompression
© 2017 Bobby Owsinski (Source: Audacity)

This practice has come under fire since we've just about hit the loudness limit, thanks to the digital environment we're now in. Still, both mixing and mastering engineers try to cram more and more level onto the file, only to find that they end up with either a distorted or an overcompressed product. While this might be the sound that the producer or artist is looking

for, it does violate the mastering engineer's unwritten code of keeping things as natural-sounding as possible while performing his level magic.

When digital first came out, people knew that every time the light went into the red that you were clipping, and that hasn't changed. We're all afraid of the "over" levels, so people started inventing these digital domain compressors where you could just start cranking the level up. I always tell people, "Thank God these things weren't invented when the Beatles were around, because for sure they would've put it on their music and would've destroyed its longevity." I'm totally convinced that over-compression destroys the longevity of a piece. Now when someone's insisting on hot levels where it's not really appropriate, I find I can barely make it through the mastering session. I suppose that's well and good when it's a single for radio, but when you give that treatment to an entire album's worth of material, it's just exhausting. It's a very unnatural situation. Never in the history of mankind has man listened to such compressed music as we listen to now.

—Bob Ludwig

Competitive Level Isn't What It Used To Be

Although everyone wants a hot mix, making it too hot might soon become a thing of the past in this online world that we now live. Many online music services like Apple Music and Spotify now normalize the uploaded content so it all plays at the same exactly the same level regardless of how hot or quiet the original uploaded master played.

As a result, there's no longer a good reason to make the master levels extremely hot, and in fact, it's proven to actually be counterproductive. A song with more dynamic range and less level will actually end up sounding better than the louder one after the service's normalization process. Mastering engineers everywhere (as well as mixing engineers) will jump for joy as sanity returns to mixing levels and the "volume wars" finally end. We're not quite to that point yet, but for the first time in many years, it looks like some progress is being made.

Setting The Compressor

The key to getting the most out of a compressor are the Attack and Release (sometimes called Recovery) parameter controls, which have a tremendous overall effect on a mix and therefore are important to understand. Generally speaking, transient response and percussive sounds are affected by the Attack control setting. Release is the time it takes for the gain to return to normal or zero gain reduction.

In a typical pop-style mix, a fast Attack setting will react to the drums and reduce the overall gain. If the Release is set very fast, then the gain will return to normal quickly but can have an audible side effect of reducing some of the overall program level and attack of the drums in the mix. As the Release is set slower, the gain changes so that the drums might cause an effect called pumping, which means that the level of the mix will increase, then decrease noticeably. Each time the dominant instrument starts or stops, it

“pumps” the level of the mix up and down. Compressors that work best on full program material generally have very smooth release curves and slow release times to minimize this pumping effect.

Compression Tips And Tricks

Adjusting the Attack and Release controls on the compressor and/or limiter can have a surprising effect on the program sound.

- Slower Release settings will usually make the gain changes less audible but will also lower the perceived volume.
- A slow Attack setting will tend to ignore drums and other fast signals but will still react to the vocals and bass.
- A slow Attack setting might also allow a transient to overload the next piece of equipment in the chain.
- Gain changes on the compressor caused by the drum hits can pull down the level of the vocals and bass and cause overall volume changes in the program.
- Usually only the fastest Attack and Release settings can make the sound “pump.”
- The more bouncy the level meter, the more likely that the compression will be audible.
- Quiet passages that are too loud and noisy are usually a giveaway that you are seriously over-compressing.

Don’t just set those Attack and Release controls to the default setting and forget about them. They can make a big difference on your final mastered sound.

Setting The Limiter

Most digital limiters used in mastering are set as brick-wall limiters. This means that no matter what happens, the signal will not exceed a certain predetermined level, and there will be no digital overs.

Thanks to the latest generation of digital limiters, louder levels are easier to achieve than ever because of more efficient peak control. This is thanks to the look-ahead function that just about all digital limiters now employ. Look-ahead delays the signal a small amount (about 2 milliseconds or so) so that the limiter can anticipate the peaks and process them before they get by. Since there is no possibility of overshooting the threshold, the limiter then becomes known as a brick-wall limiter. Analog limiters don’t work nearly as well because they can’t predict the input signal like a digital limiter with look-ahead can.

Most limiters intended for mastering will have an Output control that will set the maximum output level (see Figure 6.4). This is usually set to around -0.1dB or even -0.01dB for the absolute loudest CD level. For online delivery, this level is usually set lower, at -1dB or even less, because most encoders produce a slightly higher output.

For a vinyl record, the level may be set considerably lower still. A Threshold parameter then controls the amount of limiting that will take place, which is then measured by the gain reduction meter. Ideally, no more than 2 to 3dB of limiting should take place in order to keep the limiting less audible, although this amount might be even less for a vinyl master.



Figure 6.4: A typical digital limiter's parameter controls

© 2017 Bobby Owsinski (Source: Universal Audio)

Most digital limiters have only a Release control, since the Attack control is no longer needed because it's superseded by the look-ahead function. Almost all digital limiters have an auto-release function that automatically determines the release time based on the audio program it's processing. This is generally a safe selection, although manually setting the Release can be effective in keeping the audio punchy.

TIP: *A release time set too long can cause the sound to dull and cause the mix to lose its punch.*

Using Multi-Band Compressors And Limiters

As stated previously, multi-band compressors and limiters are extremely effective at increasing the level with far fewer audible side effects than single-band units, since they operate mainly on only a small area of the frequency spectrum at a time, leaving the others somewhat untouched (see Figure 6.5). That said, in order to keep the side effects to a minimum, a number of precautions must be taken:



Figure 6.5: The UAD multiband compressor
 © 2017 Bobby Owsinski (Source: Universal Audio)

- Use the same ratio in all bands, since using different ratios can sound unnatural.
- Use roughly the same amount of gain reduction in all bands to avoid changing the balance of the mix too much.
- Too much compression or limiting in a single band usually means that band has been excessively EQed. Either reduce the EQ if you've added it before the compressor or limiter, try cutting the offending frequency, or suggest that the mix be redone.

Limiting the program comes at a cost because the limiter does change its sound, softening transients on the drums and other percussive instruments, and taking away the punch of a track. That's why most mastering engineers do their best to use avoid using great amounts of limiting or try to use none at all, if possible.

Reducing Sibilance With A De-Esser

Sibilance is a short burst of high-frequency energy where the S's are over-emphasized; suppressing it requires a special type of compressor called a de-esser (see Figure 6.6). Sibilance often occurs in a mix when there's a lot of compression applied to either the vocal or the mix buss, or EQ is added to make the vocal rise above the mix. To use a de-esser, do the following:



Figure 6.6: The DigiRack De-esser
 © 2017 Bobby Owsinski (Source: Avid)

1. Insert the de-esser.
2. Lower the Threshold control level until the sibilance is decreased but you can still hear the S's. If you can't hear them, then you've lowered the Threshold too far.
3. Span the available frequencies with the Frequency control until you find the exact spot where it's most offensive, then adjust the Threshold control until the S's sound more natural.
4. Use the Listen feature to determine the exact sibilance frequency.

TIP: When using the Listen feature, remember that the audio you're hearing isn't in the signal path, just the sidechain. Don't forget to disengage Listen when you've found the correct frequencies.

Frequency Balance Techniques

EQing is usually the place that gets engineers mastering their own mixes into trouble. There's a tendency to over-compensate with the EQ, adding huge amounts (usually of bottom end) that wreck the frequency balance completely. Luckily, there are some rules you can follow to avoid this.

Rule 1. Listen to other CDs (not MP3s) that you like first, *before you touch an EQ parameter*. The more CDs you listen to, the better. You need a reference point to compare to, or

you'll surely over-compensate.

Rule 2. A little EQ goes a long way. If you feel that you need to add more than 2 or 3dB, you're better off mixing the song again!

EQing during mastering is almost always in small increments of anywhere from tenths of a dB to 2 or 3 at the very most.

Seriously, though, if you have to add a lot of EQ, go back and remix. That's what the pros do. It's not uncommon at all for a pro mastering engineer to call up a mixer, tell him where he's off, and suggest that he do it again.

Rule 3. Keep comparing the EQ'd version to the original version. The idea of mastering, first of all, is to make the song or program sound better with EQ, not worse. Don't fall into the trap where you think it sounds better just because it sounds louder. The only way to understand what you're listening to is to have the levels pretty much the same between the EQ'd and pre-EQ'd tracks. That's why an app like T-RackS can be really useful for mastering. It has an A/B function that allows you to compensate for the increased levels so that you can really tell if you're making it sound better.

Rule 4. Keep comparing the song you're currently working on to all the other songs in the project that you've worked on. The idea is to get them all to sound the same. It's pretty common for mixes to sound different from song to song even if they're done by the same mixer with the same gear in the same studio, but it's your job to make the listener think that the songs were all done on the same day in the same way. They've got to sound as close as possible to each other as you can get them, or at least reasonably close so they don't stand out.

TIP: *Even if you can't get the songs to sound just like your best-sounding CD, your mastering job will still be considered "pro" if you can get all the songs to sound the same in tone and volume!*

The Mastering Signal Path

The way the various processors are inserted in the signal path can make a big difference in how the mastering of a song takes place. Here are a couple of different possibilities.

The Basic Mastering Signal Chain

In its most basic form, the mastering signal chain has three elements: an equalizer followed by a compressor, which is then followed by a limiter (see Figure 6.7).




Figure 6.7: A basic mastering signal chain
© 2017 Bobby Owsinski (Source: IK Multimedia)

This order is a leftover from the analog days, when this processor order provided the best sonics when loading the program into a workstation from an analog source like tape. In many cases, it simply wasn't possible to get enough gain from the output stage of the EQ without introducing some distortion (especially if one band was driven hard), while the compressor usually is designed to add plenty of clean gain if needed. That's not as much of an issue in the digital domain, especially if there are only small 1 and 2dB increments of boost, but the EQ first in the chain still remains popular.

The downside is that whatever frequency is being boosted by the EQ can be the first to trigger the compressor, which can give you some unexpected and unpleasant results. That's why it's OK to reverse the order of the processors, with the compressor in the first spot and the EQ in the second, especially if large increments of EQ need to be applied.

As a result, there is a general rule of thumb for compressor/EQ order that goes like this:

- If you're going to use a large amount of EQ, place the EQ after the compressor.
- If you're going to use a large amount of compression, place the compressor after the EQ.

The limiter is always the last in the chain, no matter how many other devices you add and in which order, because that's what stops any overs from occurring.

An Advanced Signal Chain

Many times the simple setup outlined above isn't sufficient for a particular mastering job. Just about any time that you feel you have to use a processor to an extreme (such as

10dB of EQ or 6dB of gain reduction), then you’re probably better off using a little bit from multiple processors instead to keep the signal clean, smooth, and punchy.

Here’s the approach that many pro mastering engineers take, opting to use a little processing from a number of processors (see Figure 6.8). Once again, the limiter is at the end of the chain.





Figure 6.8: An advanced signal chain using multiple processors
© 2017 Bobby Owsinski (Source: IK Multimedia, Universal Audio)

Parallel Processing

While it’s possible to keep adding processors in the signal path, each doing a small bit of processing instead of one processor doing a lot, remember that each processor is influenced by the previous one in the chain. Sometimes a better solution to the problem is parallel processing. In this case, the main program signal would be split into two separate processing chains and then recombined before it enters the limiter. T-RackS 3 is a mastering application that’s perfectly configured for this, with four additional parallel slots in its first four processor positions (see Figure 6.9).



A Parallel Processing Mastering Signal Chain

Figure 6.9: A parallel-processing signal chain in T-RackS 3

© 2017 Bobby Owsinski (Source: IK Multimedia)

The signal path is critical to mastering success. Whether it's a simple three-processor chain or something much more complex, be sure that the limiter is the last processor in the path.

Just as a reference point, most major mastering facilities have both analog and digital signal paths, since so many of the tools and source materials exist in both domains. That being said, the overall signal path is kept as short as possible, with any unneeded gear removed, so it can't inadvertently affect the audio quality.

Adding Effects

Although mastering engineers have occasionally been asked to add effects through the years, it has now become far more commonplace than ever before. This is partly due to the proliferation of the digital audio workstation, where a poorly chosen fade is used prior to mastering or the tail of the song is cut off when prepping the file. And then there's still the fact that many artists and producers are sometimes horrified to find that the amount of reverb is suddenly less than they remembered during the mix.

Most mastering engineers prefer to add any effects in the digital domain, from both an ease-of-use and a sonic standpoint, so a reverb plugin like the Audio Ease Altiverb is chosen.

Sometimes this is done by sending the output of the workstation into the effects device and then recording the result back into the workstation on two different tracks. The

resultant effects tracks are then mixed in the proper proportions in the workstation. Because this processing is done in the digital domain, an outboard effects device with digital I/O is essential.

A lot of people assemble mixes on Pro Tools, and they don't listen to it carefully enough when they're compiling their mix, so they actually cut off the tails of their own mixes. You can't believe how often that happens. A lot of times we'll use a little reverb to just fade out their chopped-off endings and extend it naturally. I do a fair amount of classical-music mastering, and very often a little bit of reverb is needed on those projects, too. Sometimes if there's an edit that for some reason just won't work, you can smear it with a bit of echo at the right point and get past it. Sometimes mixes come in that are just dry as a bone, and a small amount of judicious reverb can really help that out.

—Bob Ludwig

Editing Techniques For Mastering

Today's mastering engineer is called on to do several types of editing that, while similar to what you might do during production, are quite specialized. Here are a few examples.

Inserting Fades

Sometimes a default fade that's added to the beginning or end of a track just doesn't sound natural. Either the fade is too tight and cuts off the attack or release of the part, or the fade itself just isn't smooth-sounding enough. Now is the time to fix any fades that don't work in the track by adjusting the fade timings.

In the case of a fade that seems unnatural (especially a fade-out), try an exponential power fade instead of the default fade (see Figure 6.10).




Figure 6.10: A linear versus an exponential fade

© 2017 Bobby Owsinski (Source: Avid)

TIP: When it comes to song fade-outs, many times the default fade just won't sound smooth enough. Be prepared to help that fade by trying some of the other types available.

Eliminating Intro Noise And Count-Offs

Leaving noise or count-offs, such as drumstick clicks on a song intro, is a sure sign of a demo recording, and is something that usually no one wants to listen to. The trick here is to use a fade-up, and don't cut off the attack of the downbeat of the song.

Clean intros are a sign of a professional mastering job. It only takes a minute and can make a big difference in how the song is perceived.

Making A “Clean” Master

In the case where a mix contains lyrics that some might find objectionable, the mastering engineer may be called upon by the record label to create a “clean” version suitable for radio airplay. This can be accomplished in multiple ways:

- If a TV track or an instrumental mix is available, edit in a piece everywhere that the objectionable lyric takes place. The vocal will drop out for the duration of the lyric,

but the song will continue and it will fly by so fast that most casual listeners may not notice if it's short enough.

- Find a similar instrumental section during the song where you can copy and paste a part that will work similarly to #1.
- If an alternate track is not available, cut a piece of a 1kHz sine wave in place of where the objectionable lyric occurs. This method really stands out, which may be just what the artist wants in order to signify that he's been censored. Then again, maybe not if it happens a lot and for long durations during the song, so try to get an instrumental mix from the producer or label if at all possible.

Parts Production

Although the more high-profile and documented part of mastering lies in the studio with the mastering engineer, the real bread and butter of the business happens after the fact, during what's known as production. Production is the time when the various masters are made, verified, and sent to the replicator. While not a very glamorous portion of the business, it's one of the most important nonetheless, because a problem there can negate a perfect job done beforehand.

Once upon a time, production was a lot more extensive than it is today. For instance, in the days of vinyl, many masters had to be made because a pair (one for each side of the disc) had to be sent to a pressing plant in each area of the country, and overseas if it was a major international release. When you consider that every master had to be cut separately with exactly the same process, you can see that the bulk of the mastering work was not in the original rundown, but in the actual making of the masters (which was very lucrative for the mastering house). In fact, many large mastering facilities employed a production engineer just for this purpose, sometimes running a night shift at the studio (a few of the larger mastering houses still do this). Over the years, parts production has dwindled to the point that we're at today, where digital copies are so easy to make that the mastering engineer does them himself.

Multiple Masters

A project may have a number of different masters cut, depending upon the marketing plans and the label's policy. This usually breaks down as follows:

- **The CD master.** This is the master from which the glass master at the pressing plant will be cut, which in turn will ultimately make the replicated CDs.
- **The vinyl master.** If a vinyl record is desired, once again a separate master is required due to the song sequence of the two-sided format. Many times the vinyl

master will also be made at a lower level and with less compression and limiting.

- **The MFiT master.** Because online is now such a large part of the overall distribution and sales picture, a separate master intended for MP3 and/or AAC for iTunes may be made. This master may be specially tweaked and the level lowered to provide the best fidelity with the least amount of bandwidth.
 - **Alternate takes.** Many artists signed to major and large indie labels also supply alternate takes, such as TV or instrumental mixes for mastering. A TV mix contains the entire mix minus the lead vocal and is used when the artist appears on television. An instrumental mix contains the entire track minus any lead and background vocals.
 - **Backup masters.** Most major labels will ask for a backup master that they will store in the company vault. Many times the mastering facility will make a “house” backup as well, to save time should a new updated master be required at a later date.
-

Chapter 7

Mastering For CD

Mastering for CD requires the mastering engineer to know far more than the basics of EQ, dynamics and editing. In fact, a proper and efficient job entails awareness of many additional processes, from inserting start/stop and track identification codes, to making choices for the master delivery medium, to checking that master medium for errors. In this chapter we'll look at all those things and more that are involved in modern CD mastering.

CD Basics

The compact disc (the proper name for the CD) was developed in the mid-'70s as a joint venture between Dutch technology giant Philips and Japanese tech powerhouse Sony, which established the standard for both the disc and players. The first commercial CD was released in 1982 by ABBA and was called *The Visitors*, although Billy Joel's *52nd Street* received the first worldwide release. More than 200 billion have been sold since then.

The digital audio resolution of the CD is 44.1kHz and 16 bits. While some of today's DAWs are capable of working with sample rates as high as 384kHz at bit depths of 32 bits, when the CD was first invented there were major limitations for both storage and transmission of digital information. This led to the 44.1kHz/16-bit standard for the CD.

The standard of 44.1kHz was decided upon because it was easily handled by the only unit capable of recording digital audio at the time, which was a modified video-tape recorder (see "Obsolete Formats" later in the chapter for more details). This frequency was also high enough to provide a 22kHz frequency bandwidth (remember the Nyquist Sample Frequency Theorem back in Chapter 2?), which would accommodate the bandwidth of human hearing.

The standard of 16 bits was chosen because it provides 96dB of potential dynamic range (6dB per bit x 16), which was felt to be substantial enough for a pleasing musical experience. There was also a limitation of A/D and D/A convertor technology at the time, and 16 bits was state of the art.

The original CDs could store 650MB of information, which is the equivalent of just over 74 minutes of music. This length was determined by the then-president of Sony, who wanted to be sure that a CD could hold the entirety of Beethoven's Ninth

Symphony. Later versions of the CD increased the storage space to 700MB, which extended the playing time to nearly 80 minutes.

How CDs Work

A CD is a plastic disc 1.2mm thick and 5 inches in diameter that consists of several layers. First, to protect the microscopically small pits (more than 8 trillion of them) against dirt and damage, the CD has a plastic protective layer on which the label is printed. Then there's an aluminum coating that contains the ridges that represent the digital data and reflects laser light. Finally, the disc has a transparent carrier through which the actual reading of the disc takes place. This plastic forms a part of the optical system (see Figure 7.1).




Figure 7.1: The CD has several layers. Notice how the ridges contain binary information
© 2017 Bobby Owsinski

Mechanically, the CD is less vulnerable than the record, but that doesn't mean that it can be treated carelessly. Since the protective layer on the label side is very thin (only one ten-thousandth of an inch), rough treatment or granular dust can cause small scratches or hairline cracks, enabling air to penetrate the evaporated aluminum coating. If this occurs, the coating will begin to oxidize, which can cause drop-outs, glitches, or distortion in the audio.

The reflective side of the CD is the side that's read by the laser in the CD player. People tend to set the CD down with the reflective side up, but it's actually the label side that's

the most vulnerable, since it's not as well protected as the reflective side. That's why it's best to store the CD in the jewel case, where it can be safely held by its inside edge.

TIP: CDs are easily scratched and should only be cleaned with a soft cloth wiping from the center to the outside edge, not radially. If a smear, however small, should remain on the CD, running along the direction of the grooves, information could be lost when being read by the player, which could cause the audio to drop out.

The area of the disc that contains data is divided into three areas (see Figure 7.2):

- **The lead-in** contains the table of contents and allows the laser pickup head to follow the pits and synchronize to the audio before the start of the program area. The length of the lead-in is determined by the number of tracks stored in the table of contents.
- **The program area** that contains the audio data and is divided into a maximum of 99 tracks.
- **The lead-out** contains digital silence or zero data and defines the end of the CD program area.




Figure 7.2: The CD layout
© 2017 Bobby Owsinski

Scanning The Disc

Like vinyl records, the information on optical discs is recorded on a spiral track in the form of minute indentations called pits (see Figure 7.3). These pits are scanned from the reflective side of the disc (this makes them appear as ridges to the laser) by a microscopically thin red laser beam during playback. The scanning begins at the inside of the back of the disc and proceeds outward. During playback, the number of revolutions of the disc decreases from 500 to 200rpm (revolutions per minute) in order to maintain a constant scanning speed. The disc data is converted into electrical pulses (the bit stream) by reflections of the laser beam from a photoelectric cell. When the beam strikes a land, the beam is reflected onto a photoelectric cell. When it strikes a ridge, the photocell will receive only a weak reflection. A D/A converter converts these series of pulses to binary coding and then back to an analog waveform (see Figure 7.4).




Figure 7.3: An electron-microscope look at CD pits and land
© 2017 Bobby Owsinski




Figure 7.4: The disc data is converted into electrical pulses (the bit stream) by reflections of the laser beam off a photoelectric cell

© 2017 Bobby Owsinski

It should be noted that the ends of the ridges seen by the laser are 1's, and all lands and pits are 0's, so turning on and off the reflection is a 1, while a steady state is a string of 0's.

Thanks to this optical scanning system, there's no friction between the laser beam and the disc, so the discs don't wear regardless of how often they're played. Discs must be treated carefully, however, since scratches, grease stains, and dust could diffract the light and cause some data to be skipped or distorted. This problem is solved by a fairly robust error-correction system that automatically inserts any lost or damaged information. Without this error-correction system CD players would not have existed, as even the slightest vibration would cause audio dropouts, glitches, and distortion.

Many CD players use three-beam scanning for correct tracking. The three beams come from one laser, as a prism projects three spots of light on the track. It shines the middle one exactly on the track, and the two other "control" beams generate a signal to correct the laser beam immediately, should it deflect from the middle track.

Mastering For CD

Mastering for CD requires a number of extra steps beyond what's required for music destined for online distribution. That's because there are a number of technical and creative processes that apply only to CD mastering.

Editing PQ Subcodes

If you've ever tried to export an album's worth of songs from your normal DAW timeline, you know that what you get is a single large file that plays like a single song instead of individual songs like we're used to on a CD. That's because you need a special editing workstation for making CD masters that allows what's known as PQ subcode editing to separate the songs out.

PQ subcodes control the track location and running-time aspects of a CD and enable the CD player to know how many tracks are present, where they are, how long they are, and when to change from one track to another. Editing software applications such as Sony CD Architect, Steinberg WaveLab, Audiofile Triumph, DSP-Quattro, and Magix Sound Forge, among others, have the ability to place these codes as needed.

When the CD was first developed, it had eight subcodes (labeled P to W), and there were a lot of uses intended for them that never came to pass. Today, the P and Q subcodes are mostly used, although the others can contain other information like CD-Text, which we'll cover shortly.

Most PQ editors also allow a PQ logsheet to be printed out, which is then sent with the master to the replicator as a check to ensure that all the correct data and information has been provided (see Figure 7.5).

CD Subcodes

- P Channel indicates the start and end of each track and was intended for simple audio players that did not have full Q-channel decoding.
- Q Channel contains the timecode (minutes, seconds, and frames), the table of contents or TOC (in the lead-in), the track type, and the catalog number.
- Channels R to W were intended for digital graphics known as CD+G, b

Client:	Test Records			
Project:	Sample			
Title:	Sample			
Date:	March 6, 2014			
Studio:	Test Mastering			
Disc Type:	Audio			
Time Format:	30/NDF			
PQ Track 1 Offset:	00:00:00:10	PQ StartOffset:	00:00:00:10	
PQ SpliceOffset:	00:00:00:06	PQ EndOffset:	00:00:00:02	
PQ MinIndex0Width:	00:00:01:00	UPC/EAN CODE:	00000000000000	
PQ Track/Index Information:				
T-X TITLE/ISRC	COPY EMPH D/A	NO OFFSET TIME hh:mm:ss:ff	OFFSET TIME hh:mm:ss:ff	CD TIME hh:mm:ss:ff
1 GBCNP7880010 OFF OFF A				
0 Pause		00:01:58:00	00:01:57:20	00:00:02:00
1 Suffering & Smiling- Part 1&2		00:02:00:00	00:01:59:20	00:21:31:02
			TOTAL:	00:21:33:02
2 GBCNP7780130 OFF OFF A				
0 Pause		00:23:30:20	00:23:30:22	00:00:02:08
1 No Agreement- Album		00:23:33:10	00:23:33:00	00:15:30:00
			TOTAL:	00:15:32:08
3 GBCNP7780140 OFF OFF A				
0 Pause		00:39:02:28	00:39:03:00	00:00:01:18
1 Dog Eat Dog- Album		00:39:04:28	00:39:04:18	00:15:32:18
			TOTAL:	00:15:34:06
LeadOut		00:54:37:04	00:54:37:06	52:39:40
Total				00:52:39:16

Figure 7.5: A PQ logsheet

© 2017 Bobby Owsinski

Inserting ISRC Codes

Most songs that are commercially released have what's called an *ISRC code*, which is short for International Standard Recording Code. It's a unique identifier for each track that lists the country of origin, the registrant (releasing entity, usually the label), the year, and a designation code (a unique identifier created by the record label). This code stays with the audio recording for its entire lifetime. Even if it later appears on a compilation, the same ISRC will accompany it.

If a recording is changed in any way, it requires a new ISRC, but otherwise it will always retain the same number regardless of the company or format it's in. An ISRC code also may not be reused.

In the U.S., the codes are administered by the Recording Industry Association of America (RIAA), which is a trade organization for music labels and distributors. ISRC codes can help with anti-piracy and royalty collection, though U.S. radio isn't very diligent about using the codes. There is better support for them in Europe.

The ISRC is contained in the Q-channel subcode of a CD and is unique to each track. Each ISRC is composed of 12 characters. Figure 7.6 shows what an ISRC looks like and what all the characters mean.




Figure 7.6: ISRC code character description

© 2017 Bobby Owsinski

Certain circumstances can cause confusion about when to apply a new ISRC code, but most of them are covered in the following list:

- Multiple recordings or takes of the same song produced even in the same recording session and even without any change in orchestration, arrangement, or artist require a new ISRC per recording or take.
- A remix, a different mix version, or an edited mix of a song requires a new ISRC.
- If the playing time changes due to an edit, the song requires a new ISRC.
- Processing of historical recordings requires a new ISRC.
- Re-release as a catalog item requires a new ISRC.
- A recording sold or distributed under license by another label can use the same ISRC.
- A compilation album where the track isn't edited or changed in any way may use the same ISRC.

So how do you get an ISRC code? If you digitally distribute your music through TuneCore or CD Baby, they'll automatically assign one for you. Many CD replicators will also assign ISRCs for you, but they'll charge you a fee for doing so. That being said, it's easy enough to register yourself. Go to usisrc.org to register (it will cost a one-time fee of \$80), and they'll assign you a three-digit registration number. You can then assign ISRC codes to all your new or previously recorded music that doesn't have an ISRC assigned to it already. Just be sure to keep a good list of the numbers and follow the rules, which are provided on the site.

Inserting UPC Codes

Another code used in the release of most albums is a UPC code. The UPC stands for *Universal Product Code*, which is the number represented by the barcode on the back of the packaging for just about any item you buy in a store these days (see Figure 7.7).



Figure 7.7: A typical UPC code

© 2017 Bobby Owsinski

While an ISRC refers to a single track, the UPC code is for the entire album, and each unique physical product that is put on a store shelf has this unique code. In addition to the barcode that you find on the back of the CD package, you can actually encode this into the PQ information on a CD.

If you have any intention of selling your CD at retail and having it recorded by SoundScan for inclusion on the Billboard music charts, you need a UPC. Most retailers only stock products with barcodes so they can easily keep track of them in their inventory, and SoundScan doesn't know you exist until you have a barcode to identify your CD.

UPCs are administered by the Uniform Code Council. If you want to obtain a manufacturer's number from this organization so you can issue your own barcodes, it will cost \$750 for the registration fee, but you can get a single UPC from CD Baby for \$20 if you're already a member, or from Nationwide Barcode (nationwidebarcode.com) for \$7.50.

Inserting CD-Text

CD-Text information includes the album title, song title, artist, and song-time information that's displayed on a playback device when the CD is playing, if the device is CD-Text enabled (some older players are not). This data is stored in the R through W subcodes, as is karaoke info, graphics, and other extended features not standard to the original CD spec. Most applications that allow you to insert PQ codes will also allow CD-Text info to be inserted, but it's not automatic and must be selected in a menu. Once again, only specialized mastering applications allow you to insert CD-Text information with your CD master.

Song Order

You don't have to think about the song order much if you're planning to release your songs online, but the song order (or sequence) becomes important as soon as an online album, a CD, or a vinyl record comes into play.

The sequence of an album has become an important creative decision all its own. In the case of a CD, a sequence that grabs the listener right at the beginning and orders the songs in such a way as to keep listener attention continually high is the goal, but it's a creative decision, so really anything goes. Because there are two sides to a vinyl record, a whole new set of decisions arises because you now have two sequences to determine —one for each side.

Selecting the sequence is not normally the domain of the mastering engineer, since it's a creative decision that should be made by the artist and producer well before mastering even begins. That said, it's something that the mastering engineer needs to know before the job can be completed and the CD master delivered.

Adjusting The Spreads

When mastering for CD or a vinyl record, the time between the songs is called the *spread*, and it can be used as a creative tool just as much as the sequence of the songs. The spreads determine the pace of the album. If the songs are close together, then the pace feels fast, and if they're further apart, it feels slower. Sometimes a combination of the two feels about right. Many times the spread is timed so that the next song will correspond to the tempo of the previous song. In other words, if the tempo of the first song was 123 beats per minute, the mastering engineer times the downbeat of the next song to stay in tempo with the previous one. The number of beats in between depends upon the flow of the album. Occasionally a cross-fade is used between songs so there's no real spread, but that's still a decision usually left for mastering as well.

Many disc-burning utilities, such as Roxio Easy Media Creator and Toast, have only limited spread selections, usually in 0.5-second intervals. That should be enough for most situations, but if you need more precision, you'll need a dedicated PQ editor, as discussed previously.

Using Dither

Dither is a low-level noise signal that's added to the program in order to trim a large digital word into a smaller one. Because the word length for an audio CD must be 16 bits, a program with a longer word length (like the usual 24 bits normally used in a DAW) must eventually be decreased. Just lopping off the last 8 bits (called truncation) degrades the sound of the audio, so the dither signal is used to gently accomplish this task. A truncated and undithered master will have decay trails stop abruptly or will have a buzzing type of distortion at the end of a fade-out, and generally will not sound as good as one that's dithered.

All dither is not created equal. There are currently many different algorithms to accomplish this task, with each DAW manufacturer having either their own version or one supplied by a third party. Generally speaking, dither comes in two major flavors—flat and noise-shaped, with the difference being that flat sounds like white noise and therefore makes what it's applied to a tiny bit noisier, while noise-shaped sometimes moves much of this injected noise to an audio band beyond where we can hear.

Although it seems like using noise-shaped dither would be a no-brainer, many mastering engineers continue to use flat dither because they claim that it tends to “pull together” mixes. Plus, if it's a loud track, you'll be hard-pressed to hear it anyway.

There are many excellent dither algorithms by Waves, iZotope, and POW-r, as well as dither plugins like PSPaudioware's X-Dither (which has a great feature that allows you to actually hear only the dither). Each provides different flavors of dither for different types of music. It's worth trying all the types before settling on one, because each can have a different effect from song to song, even in the same genre of music.

TIP: Remember when you're mixing to turn off the dither in your DAW before exporting your master mix file. Dither should only be applied once at the very end of the signal chain during mastering for it to be effective.

Rules For Using Dither

- Dither the signal once and only once. Because dither is a noise

signal, it will have a cumulative effect if applied more than once. Plus, dither introduced too early in the signal chain can have a very detrimental effect on any subsequent DSP operations that occur afterward.

- **Insert dither only at the end of the signal chain.** The time to dither is just before exporting your final master.
- **Try different types of dither.** All dither sounds different, and one may be better for a certain type of music than others. That said, the differences are usually pretty subtle.

Delivery Formats

There are two ways to deliver your master to a replication facility: audio CD or DDP file. While audio CDs work for this purpose, they are far from ideal because no matter how good the media and the burner are, there will be a number of errors in the data simply as a byproduct of burning the disc. Disc Description Protocol (DDP) files, however, are delivered as data on a CD-ROM, DVD-ROM, Exabyte tape, or FTP file transmission, and are the industry-standard method for audio delivery files for CD replication. The error correction employed by DDP is designed to be more robust than that of a CD and ensures that the audio master received by the replicator will have as few errors as possible in the data.

The DDP Master

DDP has quickly become a master medium of choice, and there are many reasons why:

- **DDP has far fewer errors than any master medium, thanks to computer data error correction.** CD-Rs and PMCDs (see the next section) have a lot less robust error correction and will output data whether it's bad or not. It's therefore possible to get different data each time you play the disc, which requires a diligent replicator to get an error-free transfer from a CD-R. Audio CDs don't protect the audio data from errors, since they assume that the CD player will hide or conceal any errors during playback. This situation leads to errors in replication when recordable CDs are used as replication masters.
- **It's easier and safer to go past the 74-minute boundary with DDP.** Long CD-Rs are less reliable, although that doesn't mean they won't work.
- **DDPs are safer.** It's impossible to play back a DDP without the right equipment

or software, which isn't readily available. This means there's less chance for accidental playback of the master, which may damage the medium. A CD-R can get smudged and scratched, but the DDP will stay in its baggie until it hits the plant.

Once again, a DDP file can only be generated from a DAW app that has true mastering capabilities.

FTP Transmission

Almost all replicators will now accept master files via FTP (File Transfer Protocol). In fact, many prefer to receive your master that way. When using FTP, the best thing to send is a DDP file, since it already contains the necessary error correction to protect against transmission errors.

All replicators have a secure portion of their website dedicated to FTP transfers. After placing your order, they'll send you the host name, user ID, and password.

Obsolete Formats

Although pressing plants will routinely accept common recordable CDs as masters, this wasn't always the case, nor is it the best way. For background's sake, here are a couple of formats that have since been made obsolete by DDP.

The Sony PCM-1630

Time for a bit of history. A longtime staple of the mastering scene was the Sony 1630 (see Figure 7.8), which is a digital processor connected to a Sony DMR 4000 or BVU-800 3/4" U-matic video machine. Once the standard format for the mastering facility to deliver to the replicator, the 1630's 3/4" U-matic tape was noted for its low error count.



Figure 7.8: A Sony 1630 digital processor with a BVU-800 3/4" machine
© 2017 Bobby Owsinski

The PCM-1630 (and its predecessor the 1610) is a modulation format recorded to a 3/4" videotape cartridge (see Figure 7.19). In the early days of the CD, it was the only way one could deliver a digital program and the ancillary PQ information to the replicator for pressing. If a replicator still accepts the format, the glass mastering from the U-matic tape can only be done at single speed, so the audio data is usually transferred to DDP for higher-speed cutting (which is not necessarily a good thing to do from an audio standpoint).




Figure 7.9: A 3/4" U-matic video cartridge

© 2017 Bobby Owsinski

When directly mastering from U-matic tape, the audio must be recorded at 44.1kHz to the Sony 1610/1630 format and the PQ code recorded on Channel 1 so that the title can be mastered directly from the U-matic tape. This PQ burst (which sounds similar to SMPTE timecode or the sound of an old dial-up modem) is basically just a data file placed on the tape before starting audio.

The PMCD

Another relic from the early days of the CD age is the PMCD. PMCD stands for Pre-Mastered CD and is a proprietary format jointly owned by Sonic Solutions and Sony. It originally was an effort to replace the Sony PCM-1630 as the standard media delivered to the replicator. It differs from a normal CD-R in that a PQ log is written into the lead-out of the disc (see “How CDs Work” earlier in this chapter). At read-back, this log is used to generate the PQ data during glass mastering, which eliminates a step during replication.

Although a great idea at the time, PMCD was relatively short-lived due to the fact that the devices used to play the discs back went out of production, and much of the other equipment used in the process was only available from Sony, which limited the replicators’ purchase options. The more modern DDP has proven to be an able replacement thanks to its robust nature and superior error correction.

How CDs Are Made

CD replication is very similar to the process of making vinyl records (which is outlined in Chapter 8, “Mastering for Vinyl”) in that it takes multiple steps to make the components that actually stamp out the disc. Here’s an overview of how that works.

1. Replication is composed of a number of stages that are required to create a master from which CD stampers are produced. All of the processes are carried out in a Class 100 clean room that’s 10 times cleaner than a hospital operating room and has a very low amount of dust particles, chemical vapors, and airborne microbes. To keep the room in this state, the mastering technicians must wear special clothing, such as facemasks and footwear, to minimize any stray particles.

The replication master begins with 8-inch diameter, 6mm thick glass blanks that are recycled from a previous replication, so glass master preparation begins by stripping the old photo-resist from the surface, which is then followed by a washing with de-ionized water and then a careful drying. The surface of the clean glass master is then coated with a photo-resist layer a scant 150 microns thick with the uniformity of the layer measured with an infrared laser. The photo-resist coated glass master is then baked at 176 degrees for 30 minutes, which hardens the photo-resist layer and makes it ready for exposing by laser light.

Laser-beam recording is where the photo-resist layer is exposed to a blue gas laser fed directly from the source audio of a DDP master tape or file. The photo-resist is exposed where pits are to be pressed in the final disc. The photo-resist surface is then chemically developed to remove the photo-resist exposed by the laser and therefore create pits in the surface. These pits then extend right through the photo-resist to the glass underneath to achieve the right pit geometry. The glass itself is unaffected by this process (see Figure 7.10).




Figure 7.10: Making the CD master

© 2017 Bobby Owsinski

2. The surface of this glass master, which is called the *metal master* or *father*, is then coated with either a silver or a nickel metal layer. The glass master is then played on a disc master player to check for any errors. Audio masters are actually listened to at this stage.
3. The next stage is to make the reverse image stamper or *mother* (a positive image of the final disc pit and land orientation).
4. Stampers (a negative image) are then made from the mother and secured into the molding machines that actually stamp the CD discs.
5. After a CD has been molded from clear polycarbonate, a thin layer of reflective metal is bonded onto the pit and land surface, and a clear protective coating is applied.
6. The disc label is printed on the non-read surface of the disc, and the CD is inserted into a package such as a jewel case, with a tray, booklet, and backliner.

A single unit called a *Monoliner* is used to replicate CDs after the stamper has been created. The Monoliner consists of a complete replication line made up of a molding machine, a metalizer, a lacquer unit, a printer (normally three-color),

and inspection. Good and bad discs are automatically transferred to different spindles. Finished discs are removed on spindles for packing. It's also possible for the Monoliner to not include a printer, so a new job can continue without being stopped while the printer is being set up.

A *Duoline* is a replication machine made up of two molding machines, a metalizer, a lacquer unit, and inspection. Each molding machine can run different titles, with the discs being separated after inspection and placed on different spindles.

Duplication Versus Replication: What's The Difference?

- Some people use the terms duplication and replication interchangeably but there is a difference. Duplication means to make a copy of something, so in the case of a CD it would mean burning a copy of disc. A computer extracts the digital data from a master disc and prints it onto a recordable CD, DVD or Blu-ray disc, thus the new disc is a copy of the master. The recordable disc was manufactured first, then the information was burned into it later.
- Replication means that the new disc is made from the ground up complete with the information from a digital master during the manufacturing process. The replicated disc was just a lump of polycarbonate before it was stamped. It's then manufactured already containing the data.
- There's a quality issue between duplicated and replicated discs, although it's fairly small these days. A duplicated disc may have errors as a result of the burning process that could degrade the sound. A replicated disc theoretically has none of these errors.
- Some golden-eared mastering engineers claim they can hear the difference between the two every time, although the average person probably won't perceive one.

Chapter 8

Mastering For Vinyl

Although it seems like almost an ancient technology in these days of 1's and 0's, the vinyl record is making a resurgence in the marketplace, increasing in sales dramatically. That said, at a sales level of around 9 million a year or so (the ones that are counted, at least), vinyl is no threat to other music distribution methods and probably never will be again. Still, the format is in no danger of dying either, and there are still many requests for vinyl masters every year.

While the vast majority of engineers won't be purchasing the gear to cut vinyl anytime soon, it's still important to know what makes the format tick in order to get the best performance if you decide to make some records to complement a CD or an online project. Before we get into the mastering requirements for vinyl, let's take a look at the system itself and the physics required to make a record. While this is by no means a complete description of the entire process of cutting a vinyl record, it is a pretty good overview.

A Brief History Of Vinyl

It's important to look at the history of the record because in some ways it represents the history of mastering itself. Until 1948, all vinyl records were 10 inches in diameter and played at 78 revolutions per minute (RPM). When Columbia Records introduced the 12 inch 33 1/3rd RPM disk in 1948, the age of hi-fidelity actually began since the sonic quality took a quantum leap over the previous generation. However, records of that time had a severe limitation in that they only held about 10 minutes of playing time per side, since the grooves were all relatively wide in order to fit all the lower audio frequencies of the music on the record.

To overcome this time limitation, two refinements occurred. First, in 1953 the Recording Industry Association of America (the RIAA) instituted an equalization curve that narrowed the grooves, thereby allowing more of them to be cut on the record, which increased the playing time and decreased the noise at the same time. This was done by boosting the high frequencies by about 17dB at 15kHz and cutting the lows by 17dB at 50Hz when the record was cut (see Figure 8.1). The opposite curve is then applied during playback. This is what's known as the RIAA curve. It's also the reason why it sounds so bad when you plug your turntable directly into a mic or line input of a console. Without the RIAA curve applied, the resulting sound is thin and tinny due to the overemphasized high frequencies and attenuated low frequencies.




Figure 8.1: The RIAA equalization curve

© 2017 Bobby Owsinski

The second refinement was the implementation of variable pitch, which allowed the mastering engineer to change the number of grooves per inch according to the program material. In cutting parlance, pitch is the rate at which the cutter head and stylus travel across the disk. By varying this velocity, the number of grooves can be varied as well by changing the spacing between them. These two advances increased the playing time to the current 25 minutes or so (more on this later) per side.

In 1957, the stereo record became commercially available and really pushed the industry to the sonic heights that it has reached today.

How A Vinyl Record Works

To understand how a record works, you really must understand what happens within a groove. If we were to cut a mono 1kHz tone, the cutting stylus would swing side to side in the groove 1,000 times per second (see the groove pictures 8.2 through 8.15). The louder the signal, the deeper the groove that must be cut.

While this works great in mono, it doesn't work for stereo, which was a major problem for many years. As stated before, stereo records were introduced in 1957, but the fact of the matter is that the stereo record-cutting technique was actually proposed in 1931 by famed audio scientist Alan Blumlein. His technique, called the 45/45 system, was revisited some 25 years later by the Westrex Corporation (the leader in record-equipment manufacturing at the time) and resulted in the eventual introduction of the stereo disk.

Essentially, a stereo disk combines the side-to-side (lateral) motion of the stylus with an

up-and-down (vertical) motion. The 45/45 system rotated the axis 45 degrees to the plane of the cut. This method actually has several advantages. First, mono and stereo disks and players become totally compatible, and second, the rumble (low-frequency noise from the turntable) is decreased by 3dB.

Let's take a look at some groove pictures provided by Clete Baker of Studio B in Omaha, Nebraska to illustrate how the system works.



Figure 8.2: This is a silent groove with no audio information. The groove width across the top of the "v" from land to land is 0.002 inches, and the groove depth is approximately the same as the width for this particular stylus (made by Capps).

Courtesy of Clete Baker




Figure 8.3: From the outside of the disk (right-hand side) going inward, you can see a low-frequency sine wave, a mid-frequency sine wave, and a high-frequency sine wave, all in mono (lateral excursion). All frequencies are at the same level at the head end of the system, which is prior to application of the RIAA curve. This demonstrates that for any given level, a lower frequency will create a greater excursion than a high frequency, and as a result will require greater pitch to avoid intercuts between grooves.

Courtesy of Clete Baker




Figure 8.4: This is a sine wave applied to the left channel only toward the outer part of the record, summed to mono in the center of the view, and applied to the right channel only toward the inner part of the record. You can easily see the difference between the purely lateral modulation of the mono signal and the vertical of the left- and right-channel signals.

Courtesy of Clete Baker



*Figure 8.5: A human hair laid across the groove offers a point of comparison.
Courtesy of Clete Baker*



Figure 8.6: Again, lower-frequency and higher-frequency sine waves demonstrate that more area of the disk is required to accommodate the excursion of lows than of highs.

Courtesy of Clete Baker



Figure 8.7: This figure shows variable pitch in action on program audio. To accommodate the low-frequency excursions without wasting vast amounts of disk real estate, variable pitch is employed to spread the groove in anticipation of large excursions. This narrows the groove if the material doesn't contain any low frequencies that require it.

Courtesy of Clete Baker




Figure 8.8: This is what happens when variable pitch goes bad, which is a lateral intercut caused by insufficient variable pitch for a wide lateral excursion. Toward the bottom center of the slide, the outside wall of the loud low frequency has touched the adjacent wall of the previous revolution, but the wall has not broken down, and a safe margin still exists so it won't cause a skip. However, on the next revolution an excursion toward the outside of the disk has all but overwritten its earlier neighbor, which is certain to cause mistracking of the playback stylus later.

Courtesy of Clete Baker




Figure 8.9: Lateral excursions aren't the only source of intercuts. This shows a large low-frequency vertical excursion caused by out-of-phase information, which is not severe enough to cause mistracking, but will probably cause distortion.

Courtesy of Clete Baker

TIP: This is exactly why low frequencies are panned to the center, or an elliptical equalizer (see later in the chapter) is used to send them to the center, during disc mastering.




Figure 8.10: Large vertical excursions can cause problems not only by carving out deep and wide grooves that result in intercuts, but by causing the cutting stylus to literally lift right off the disk surface for the other half of the waveform. This would cause a record to skip.

Courtesy of Clete Baker




Figure 8.11: Here a near lift is accompanied on the following revolutions by lateral intercut, which will result in audible distortion. To solve this problem, the mastering engineer can increase groove pitch and/or depth, lower the overall level at which the record is cut, reduce the low-frequency information, sum the low frequencies at a higher crossover point, or add external processing, such as a peak limiter. Each of these can be used alone or in combination to cut the lacquer, but none can be employed without exacting a price to the sound quality.

Courtesy of Clete Baker



Figure 8.12: Here is the same audio viewed in Figure 8.11 only after processing. In this case a limiter was employed to reduce dynamic range (the surrounding material is noticeably louder as well) and rein in the peaks that were causing intercuts and lifts. This section is cut more deeply in order to give vertical excursions plenty of breathing room. Pitch too has had to be increased overall to accommodate the slightly wider groove, despite the reduced need for dramatic dynamic increases in pitch due to the reduction of peaks by the limiter.

Courtesy of Clete Baker




Figure 8.13: Because of the physics of a round disk, high-frequency information suffers terribly as the groove winds closer to the inner diameter. Here what high-frequency-rich program material near the outer diameter of the disk looks like.

Courtesy of Clete Baker



Figure 8.14: Here is the same audio information as in Figure 8.13 only nearer the inside diameter of the disk.

Courtesy of Clete Baker



*Figure 8.15: The ideal: normal, healthy-looking program audio.
Courtesy of Clete Baker*

The Vinyl Signal Chain

While the signal chain for vinyl is similar to that of a CD, there are some important distinctions and unique pieces involved. Let's look at the chain from the master lacquer (the record that we cut to send to the pressing plant) on back.

The Master Lacquer

The master lacquer is the record that is cut to send to the pressing plant. It consists of a mirror-smooth substrate of aluminum coated with cellulose nitrate (a distant cousin to nitroglycerine) along with some resins and pigments to keep it soft and help with visual inspection (see Figure 8.16). The lacquer is extremely soft compared to the finished record, and the master can never be played after it's cut. To audition the mastering job before a lacquer is cut, a reference disk called a "ref" or an acetate is created. Since this is made of the same soft material as the master lacquer, it can only be played five or six times at most before the quality is significantly degraded. There's a separate lacquer master created for each side of the record. The lacquer is always larger than the final record (a 12-inch record has a 14-inch lacquer), so repeated handling does not damage the

grooves.



Figure 8.16: A master lacquer
© 2017 Bobby Owsinski

The Cutting Stylus And Cutter Head

The cutting stylus, which is made of sapphire, sits inside the cutter head, which consists of several large drive coils (see Figure 8.17). The drive coils are powered by a set of very high-powered (typically 1,000 to 3,500 watts or higher) amplifiers. The cutting stylus is heated for an easier and quieter cut.

The Lathe

The lathe contains a precision turntable and the carriage that holds the cutter head assembly, as well as a microscope to inspect the grooves and adjustments that determine the number of grooves and the depth of cut. No lathes are currently being manufactured (although there's been a rumor of a new one to be introduced soon), but models by Scully and Neumann were once among the most desirable (see Figure 8.17).



Figure 8.17: Neumann VMS-80 with SX 84 cutter head from 1984
© 2017 Bobby Owsinski

Way back in the '50s, the first cutting systems weren't very powerful. They only had maybe 10 or 12 watts of power. Then, in the '60s, Neumann developed a system that brought it up to about 75 watts per channel, which was considered pretty cool at the time. In the '70s, the high-powered cutting systems came into being, which were about 500 watts. That was pretty much it for a while, since it made no sense to go beyond that because the cutter heads really weren't designed to handle that kind of power anyway. Even the last cutting system that came off the line in about 1990 at Neumann in Berlin hadn't really changed other than it had newer panels and prettier electronics.

—David Cheppa of Better Quality Sound

In the physical world of analog sound, all the energy is in the low end. In disk-cutting, however, it's the exact opposite, with all the energy in the upper spectrum. As a result, everything from about 5kHz up begins to require a great amount of power to cut, and this is why disk-cutting systems need to be so powerful. The problem is that it can be devastating if something goes wrong at that power, with at least a lacquer, and maybe a stylus or even a cutter head, at the risk of possibly being destroyed.

The Mastering Console

The mastering console for a disk system is equal to that used today for mastering in sound quality and short signal path, but that's where the similarity ends. Because of the

unique requirements of cutting a disk and the manual nature of the task (thanks to the lack of computerized gear at the time), there are several features found on this type of disk that have fallen by the wayside in the modern era of mastering (Figure 8.18).




Figure 8.18: A Neumann SP-75 vinyl mastering console

© 2017 Bobby Owsinski

The Preview System

Chief among those is the preview system, which is an additional monitor path made necessary by the volatile nature of cutting a disk. Here's the problem—disk-cutting is essentially a non-stop operation. Once you start to cut, you must make all your changes on the fly without stopping until the end of the side. If a portion of the program has excessive bass information, a loud peak, or something out of phase, the cutter head will cut right through the lacquer to the aluminum substrate. This would destroy not only the lacquer, but maybe an expensive stylus as well. Hence the need for the mastering engineer to hear the problem and make the necessary adjustments before any harm comes to the disk.

Enter the preview system. Essentially, the program going to the disk is delayed. Before digital delays were invented, an ingenious dedicated mastering tape machine with two separate head stacks (program and preview) and an extended tape path (see Figure 8.19) was used. This gave the mastering engineer enough time to make the necessary adjustments before any damage was done to the disk or system.



Figure 8.19: A Studer A80 tape machine with preview head
© 2017 Bobby Owsinski

Equalization

Since a disk had to be cut on the fly and computer automation was still years away, a system had to be created in order to make EQ and compression adjustments from song to song quickly, easily, and, most necessarily, manually. This was accomplished by having two of each processor with their controls stepped so that adjustments could be repeatable.

The mastering engineer would then run down all the songs for one side of the LP and mark down the EQ settings required for each song on a template. Then, as the first song was being cut through the “A” equalizer, he would preset the “B” equalizer for the next song. As song 2 was playing through the B equalizer, he would preset equalizer A for song 3, and so on (Figure 8.18).

Although this method was crude, it was effective. Naturally, today it’s much easier, because the master is a file in a workstation where the EQ and compression have already been added.

The Elliptical Equalizer

One of the more interesting relics of the record days is the elliptical equalizer or low-frequency crossover. This unit moves all low frequencies below a preset frequency (usually 250, 150, 70, and 30Hz) to the center. This is done to stop excessive lateral

movement of the cutting stylus because of excessive low-frequency energy on one side only or excessive out-of-phase material. Obviously, this device could negatively affect the sound of a record, so it had to be used judiciously.

How Records Are Pressed

Pressing records is such a primitive process by today's standards that it's pretty amazing that they sound as good as they do. This is a multi-step operation that's entirely mechanical and manual, with a host of areas that could influence the end product in a mostly negative way.

1. **The master lacquer** is used as the first of several metal molds from which the vinyl records are pressed. The lacquer is first coated with a layer of tin and silver nitrate and then dropped in a nickel sulfamate bath and electroplated. The lacquer is removed from the bath, and the nickel coating is peeled away. The separated nickel is what's known as the *metal master* and is a negative of the lacquer.
2. **The metal master** is dropped back into the nickel bath and electroplated again. The resultant separated metal part is known as the *mother* and is a positive copy that can be played since it has grooves (although that would destroy it).
3. **The mother** is dropped back into the nickel bath and electroplated again. The resultant separated metal part is known as the *stamper* and is a negative copy that is bolted into the record presser to actually stamp out the plastic records.
4. It should be noted that, just like tape, each resultant copy is a generation down and will result in 6dB worse signal-to-noise ratio. Also, great care must be used when peeling off the electroplating, since any material left behind will result in a pop or click on the finished product.
5. **The vinyl** used to make records actually comes in a granulated form called *vinylite* and isn't black, but honey-colored. Before being pressed, it's heated into the form of modeling clay and colored with pigment. At this point it's known as a *biscuit*.
6. **The biscuit** is then placed in the press, which resembles a large waffle iron, and is heated to about 300 degrees. Temperature is important because if the press is too hot, then the record will warp; if it's too cold, the noise will increase. After pressing, excess vinyl is trimmed with a hot knife, and the records are put on a spindle to cool at room temperature.

All of these metal parts wear out, and a stamper will go dull after about 70,000 pressings. Because of that, several sets of metal parts have to be made for a

large order, and in the case of the big-selling records of yesterday, even several lacquers. For some nice lathe pictures, go to aardvarkmastering.com/history.htm.

New Advances In Vinyl Technology

The technology of record pressing has been virtually stagnant for more than 40 years, with lathes and record pressing machines relying on cannibalized parts from old worn-out gear to keep them going. Thanks to the resurgence in vinyl sales that put more money in this side of the business, we're now seeing some new gear being manufactured as well as some major improvements in the technology about to be implemented. Let's take a look at some of these new developments.

New Record Presses

Until now, if a pressing plant wanted to put add another record press they had to scour the world looking for forgotten presses in a warehouse. While that actually did happen from time to time, most stampers found this way weren't in good enough condition to immediately go online as they required a good deal of refurbishing, which took a considerable amount of time and expertise.

Now for the first time not one, but two companies are releasing brand new record stamping machines, which pressing plants can't seem to get enough of (the backlog is about 5 months are most facilities). Toronto-based Viryl Technologies Warm Tone press is even computerized and fully automated, allowing a record to be stamped every 25 seconds, about half the time of a normal record press. This is something that the vinyl record community has longed for and it has finally come to pass.

A New Way Of Pressing

As you've seen, making a vinyl record is a messy, time consuming business. It involves toxic chemical baths, huge mechanical presses, stampers that wear out easily, and maybe worst of all, the final product is made from an offshoot of petroleum. Record pressing has shown small improvements over the years, but other than Viryl Technologies Warm Tone mentioned above, it's still done the way it was 40+ years ago.

That could change soon however. A new injection moulding process invented by the Dutch company Symcon, promises not only to cut production costs, but to improve the sound quality, and reduce the environmental impact of conventional record pressing as well.

In a conventional record press, a PVC puck is heated with steam until it's soft, then placed between the two stampers that press the puck for about 8 seconds. It takes about

20 seconds to press, then another 16 seconds is then required for the record to cool off before the process can begin again.

In the new process, the plastic mixture is heated in advance, injected between the two stampers, then pressed for a few seconds and cooled for another 20 seconds to make sure the mixture reaches the outer edges of the stampers.

There are several big advantages with injection moulding. First of all, the amount of energy used is cut by up to 65%. There's no excess vinyl around the record that needs to be cut off, and the stampers last much longer before they degrade. Currently, a stamper may last for as little as 2,000 records before it must be replaced (although that figure is normally higher). Yet another happy byproduct is that the record's surface noise is reduced by up to 10dB over conventionally pressed records.

This seems like a slam dunk, but there are still a few challenges to overcome. So far, injection moulded records are less durable, as they show signs of wear after 35 plays compared to 50 times for a vinyl record. The price is also about 25% higher, although that should come down over time. It also takes more time to actually press the record, which is a serious disadvantage. This new system holds a lot of promise, but it's too early to tell whether it's revolutionary or not.

HD Vinyl

An Austrian company called Rebeat Digital has filed a patent to develop what would be they call the world's first "high definition vinyl" technology. This new technology is capable of producing records with longer playing times, 30% more volume, and twice the fidelity of today's vinyl records, plus could potentially cut the wait times at pressing plants in the process.

The beauty of HD Vinyl is that it would eliminate the toxic process of electroplating to make the metal parts to create the stampers. Instead it would use a laser to directly cut the stamper, which could reduce not only the groove wear but shave off about 60% of the time it takes to actually make a vinyl record.

This all sounds good on paper, but the process has yet to be attempted in a real-world pressing plant. That might be coming soon however, so don't be surprised if the next time you purchase a disc it suddenly sounds a lot better than what you're used to.

Chapter 9

Mastering For Online Distribution

It wasn't that long ago that much care was needed when creating files to be shared online. Because of restricted storage and bandwidth, file sizes had to be kept small, and getting the highest-quality audio from them was much trickier. Today it still requires some expertise to make a great-sounding streaming or download file, but the margin for error is a lot larger, thanks to increased bandwidth and the resulting file sizes that go with them. Let's take a look at how we can tweak those online audio files to get the most out of them.

File Compression Encoding

Although the MP3 is the file compression that most people are familiar with, there are actually a lot of alternatives that are regularly used online today, especially by streaming networks. That said, many of the principles of MP3 file compression are the same for the other formats as well, and the techniques that we've learned to make audio files that shine online come directly from basic MP3 encoding.

First of all, most of the widely used online file formats like the MP3 are encoded using what's known as lossy data compression. Data compression is not at all like the audio compression that we've been talking about so far in this book. Data compression means decreasing the number of bits in a digital word to make the file smaller. File formats like the MP3 do this in a lossy manner, which means that they literally throw away certain audio information that the encoder thinks isn't important and won't be missed. Of course, if we compare an data compressed file to its original non-data-compressed source file, we might hear a difference as a result. That's why the following information and parameter settings are so important—so you can get the best-sounding MP3 file that sounds as close to the uncompressed source file as possible.

The encoder that does the data compression also does the decompression as well, which is why it's known as a codec, which is short for compression/decompression. Table 9.1 shows a list of the popular codec currently in use online. Keep in mind that there are dozens more besides these, but most haven't risen to widespread use for music distribution.

Table 9.1: Commonly Used Audio Compression Formats

Compression	Means	Type	Comments
AAC	Advanced Audio Coding	Lossy	Generally thought to be the best sounding of the lossy codecs
ALAC	Apple Lossless Audio Codec	Lossless	Supports up to 8 channels at 384kHz but mostly used for archiving
FLAC	Free Lossless Audio Codec	Lossless	The most widely used lossless codec
MP3	MPEG 2 - Audio Layer III	Lossy	Once the standard for file compression. Quality varies with different encoders.
OGG	Derived from the jargon of the computer game Netrek	Lossy	Better sounding than MP3 but as widely adopted. Used in many video games.

Lossy Versus Lossless Codecs

Lossy compression utilizes a perceptual algorithm that removes signal data that is being masked or covered up by other signal data that is louder. Because this data is thrown away and never retrieved, it's known as *lossy*. This is done not only to make the audio file smaller, but also to fit more data through a small data pipe, like an Internet service with limited bandwidth.

To illustrate what lossy compression does, think of the inner tube for a bicycle tire that's filled up with air. When you let the air out of the tube, it takes up less space, yet the same amount of rubber remains and it can fit into a small box that sits on the store shelf. This is the same idea behind lossy data compression.

Depending upon the source material, lossy compression can be either completely inaudible or somewhat noticeable. It should be noted that even when it is audible, lossy compression still does a remarkable job of recovering the audio signal and still can sound quite good with the proper preparation, like what we'll soon discuss.

The opposite of lossy compression is *lossless* compression, which never discards any data and recovers it completely during decoding and playback. FLAC is the most widely used lossless codec, although ALAC (Apple Lossless) is also used. The lossless compressed files are usually 40 to 60 percent as large as unencoded PCM files, but that's still quite a bit larger than the typical MP3 or AAC file, which are about a tenth of the original size.

File Compression Encoder Parameters Explained

Regardless of the encoder, there's one parameter that matters the most in determining the quality of the encode, and that's bit rate, which is the number of bits of encoded data that are used to represent each second of audio. Lossy encoders such as MP3 provide a number of different options for bit rate. Typically, the rate chosen is between 128 and

320 kilobits per second (kbs), although it could go as low as 96kbs for a mono podcast. By contrast, uncompressed stereo audio as stored on a compact disc has a bit rate of about 1400kbs (see Table 9.2).

Table 9.2: Bit Rate Quality Comparison

Bit Rate	Quality	Comments
64kbps	Poor	For voice and mobile
128kbps	Poor	Minimum bit rate for music
192kbps	Good	Acceptable quality
256kbps	Very Good	iTunes bit rate
320kbps	Excellent	Premium streaming rate
1411.2kbps (44.1kHz/16 bit)	CD Quality	Too large for streaming
2304kbps (48kHz/24 bit)	Very High Quality	Television and movie standard
4608kbps (96kHz/24 bit)	Audiophile Quality	Hi-resolution audio - MFIT suggested rate

Data compressed files encoded with a lower bit rate will result in a smaller file and therefore will download faster, but they generally play back at a lower quality. With a bit rate too low, compression *artifacts* (sounds that weren't present in the original recording) may appear during playback. A good demonstration of compression artifacts is provided by the sound of applause, which is difficult to data compress because it's so random. As a result, the failings of an encoder are more obvious and become audible as a slight ringing.

Conversely, a high bit-rate encode will almost always produce a better-sounding file, but also a larger file, which may require an unacceptable amount of storage space, be too large to send via email, or take too much time to download (although in these days of seemingly unlimited storage and widespread high-speed Internet, these are becoming less and less of a factor).

The norm for acceptable-quality compressed files like MP3s has mostly become 160kbs. Here are the pros and cons of the different bit rates used for music.

- **128kbps.** Lowest acceptable bit rate, but may have marginal quality depending upon the encoder. Results in some artifacts but small file size.
- **160kbps.** Lowest bit rate considered usable for a high-quality file.
- **320kbps.** The highest quality; it has a larger file size but the sound can be indistinguishable from CD under certain circumstances.

Three modes are coupled with bit rate and have a bearing on the final sound quality of the encode.

- **Variable Bit Rate (VBR)** mode maintains a constant quality while raising and

lowering the bit rate depending upon the program's complexity. Size is less predictable than with ABR (see below), but the quality is usually better.

- **Average Bit Rate (ABR)** mode varies the bit rate around a specified target bit rate.
- **Constant Bit Rate (CBR)** mode maintains a steady bit rate regardless of the program's complexity. CBR mode usually provides the lowest-quality encode, but the file size is very predictable.

At a given bit-rate range, VBR will provide higher quality than ABR, which will provide higher quality than CBR. The exception to this is when you choose the highest possible bit rate of 320kbps, where, depending upon the encoder, the mode may have little bearing on the final sound quality.

Some additional parameter settings may be available that can have a huge influence on the quality of the final encode. These include:

- **Mid-Side Joint Stereo** (sometimes called MS Joint Stereo) encodes all of the common audio on one channel and the difference audio (stereo minus the mono information) on the other channel. This is intended for low bit-rate material to retain surround information from a surround mix source, and it is not needed or desired for stereo source files. Do not select this under normal circumstances.
- **Intensity Joint Stereo** is again intended for lower bit rates. It combines the left and right channels by saving some frequencies as mono and placing them in the stereo field based on the intensity of the sound. This should not be used if the stereo audio contains surround-encoded material.
- **Stereo Narrowing** is again intended for lower bit rates, narrowing the stereo signal to increase overall sound quality.

TIP: *It's better not to check any of the above parameters when encoding stereo files that originate at 16-bit or above. With these disabled, the encoding will remain in true stereo, with all of the information from the original left channel going to the left side and the same for the right channel.*

Encoding a compressed audio file like an MP3 may seem easy, but making it sound great requires a bit of thought, some knowledge, and a little experimentation. The idea is to encode the smallest file with the highest quality, which is, of course, the tricky part. Here are some tips to get you started in the right direction so you won't have to try every possible parameter combination. Remember, though, that the settings that might work on one particular song or type of music might not work on another.

It's All About The Source File

Lossy coding like MP3 makes the quality of the master mix more of an issue because high-quality audio may be less damaged by this type of encoding than low-quality audio will. Therefore, it's vitally important that you start with the best audio quality (highest sample rate and most bits) possible.

It's also important to listen to your encode and perhaps even try a number of different parameter settings before settling on the final product. Listen to the encode, A/B it to the original, and make any additional changes you feel are necessary. Sometimes a big, thick wall of sound encodes terribly, and you need to ease back on the compression and limiting of the source track. Other times, heavy compression can make it through better than with a mix with more dynamics. There are a few predictions one can make after doing it for a while, but you can never be certain, so listening and adjusting is the only sure way.

Choosing An Encoder

Unfortunately, all encoders, especially for the MP3, are not created equally, and therefore they don't provide the same quality output, so using a good encoder is the biggest advantage you can give yourself.

When it comes to an MP3 encoder, one to consider is LAME, which is an open-source application. LAME is an acronym for "LAME Ain't an MP3 Encoder," although the current version really is a stand-alone encoder. The consensus seems to be that LAME produces the highest-quality MP3 files for average bit rates of 128kbit/s and higher, although the differences between other encoders are minimal. Another good MP3 encoder is the one found within iTunes.

The Ins And Outs Of File Metadata

An digital music file contains not only the actual audio, but also information about that song called metadata. You can think of metadata as a small database associated with each song, and within that database there are tags that identify the song name, artist, album, music genre, release year, and more. Obviously, those tags tell more about the file than a filename ever could. You could have an MP3 called Jjbr#\$.mp3, but as long as it has accurate tags, it will be identified as "Electrolux" by SNEW off the *What's It to Ya* album.

The most common metadata fields added to MP3 files are:

- **Title.** The track title.

- **Artist.** The artist that recorded the track.
- **Album.** Which album the track belongs to (if applicable).
- **Track.** The track number from the album (if applicable).
- **Year.** The year that the track was published.
- **Genre.** The type of track—for example, speech, rock, pop.
- **Comment.** Additional notes about the track.
- **Copyright.** Copyright notice by the copyright holder.
- **Album Art.** Thumbnail of the album art or artist.

In addition to these common fields, other data can be included, such as ISRC codes, web addresses, composer, conductor, and orchestra.

Metadata is supported by MP3s, Ogg Vorbis, FLAC, AAC, Windows Media Audio, and a few other file formats that aren't used that often.

It's more critical than ever that the metadata be accurate. It's the best way for a song to be identified when it's streamed so that the artist, label, songwriter and publisher get paid. Make sure that you take the time to fill in all the metadata fields before you release your digital music file to the world. Even though there are plenty of editors that allow listeners to insert the data after the fact, wouldn't you prefer that it comes from you?

Creating Great-Sounding Online Files

Encoding a digital music file intended for online distribution may seem easy, but it requires a bit of thought to make it sound great, as well as some knowledge and occasionally some experimentation. Here are some tips to get you started in the right direction so you won't have to try every possible parameter combination. Remember, though, that the settings that might work on one particular song or type of music might not work on another.

If you want the best-sounding digital music files possible, follow these tips:

- **Start with the highest-quality audio file possible.** High-quality audio will be damaged much less during encoding than low-quality source audio. Therefore, it's vitally important that you start with the best audio quality (highest sample rate and most bits) possible. That means that sometimes it's better to start with the 24-bit master or make the encode (usually an MP3) while you're exporting your mix, rather than using something like a 16-bit CD master as the source for your encodes.
- **Lower the level.** Files that are lower in level than your CD master usually result in a

better-sounding encode. A decrease by a dB or two can make a big difference in how the encode sounds.

I did a panel at the Nashville Recording Workshop where I took some stuff that was pretty heavily compressed and was able to run it at different levels into the Mastered for iTunes software to A/B things. I had a DAC that I could vary the level in 1dB steps so that I could lower the level into the encoder, then raise it by the same amount during playback so the reference level was the same. As I pulled it down, the crowd was amazed at how much better it sounded the lower the level to the encoder was, to the point that you could even hear it over the sound system in this big room. When we got down 4dB in level, they could hear that it was virtually identical to the original 192kHz/24-bit master, so it does make a difference, and people are starting to realize that.

—Glenn Meadows

TIP: *MP3 encoding almost always results in the encoded material being slightly hotter than the original source. Limit the output of the material intended for MP3 to -1.1dB or less, instead of the commonly used -0.1 or -0.2dB , so you don't get digital overs.*

- **Filter out some high frequencies.** Filter out the top end at whatever frequency works best (judge by ear). Most codecs (especially for MP3) have the most difficulty with high frequencies, and rolling them off liberates a lot of processing for encoding the lower and mid frequencies. You trade some top end for better quality in the rest of the spectrum.
- **A busy mix can lose punch after encoding.** Sparse mixes, like acoustic jazz trios, seem to retain more of the original audio punch.
- **Use Variable Bit Rate mode** when possible.
- **Turn off Mid-Side Joint Stereo, Intensity Joint Stereo, and Stereo Narrowing** if it's available.
- **Don't use a bit rate below 160kbs** (higher is better).
- **Don't hyper-compress.** Leave some dynamic range so the encoding algorithm has some dynamics to look at.
- **Set your encoder for maximum quality**, which allows it to process for best results. The encoding time is negligible anyway.

TIP: *The above suggestions apply to both MP3 and other online codecs, including those used for the various streaming services.*

Creating Files For Streaming Services

Submitting song files to the various streaming services is usually done either through an aggregator such as TuneCore or CD Baby, or directly from a record label. Regardless of who submits the files, the requirements are the same in most cases – at least 44.1kHz/16-bit audio, the same as with a CD master (although Apple Music is the exception - more on that in Chapter 10). In some cases a high-quality MP3 at 320kbps is also acceptable for submission, but obviously this doesn't represent the best that the music can sound.

After the file is submitted, the streaming service then encodes it to their specifications, which vary considerably from service to service. It's good to know how the music will be encoded in order to provide the best-sounding source file, so here's a chart with the current streaming specs of the most popular services.

Table 9.3: Streaming Specs for the Most Popular Services

Service	Audio Format	Bit Rate
Amazon Music Unlimited	AAC	256kbps
Spotify	Ogg Vorbis	96kbps mobile 160kbps standard quality 320kbps premium quality
Pandora	MP3	128kbps free 192kbps premium
YouTube	AAC	240P video - 64kbps 360P video - 128kbps 480P video - 128kbps 720P video - 192kbps 1080P video - 192kbps
Apple Music	AAC	256kbps
Beats 1 Radio	AAC	256kbps
iHeartRadio	AAC	64kbps
Slacker	MP3	320kbps
Deezer	AAC	128kbps free 320kbps standard 1411kbps elite
Rhapsody	MP3	64kbps mobile 128kbps desktop 192kbps premium
Google Play Music	MP3	320kbps
SoundCloud	MP3	128kbps
Tidal	FLAC	320kbps standard 1411kbps HiFi

TIP: Remember to treat files intended for streaming services as you would for an MP3 file and lower the level 1 to 2dB to ensure distortion-free encoding.

Be Aware Of Sound Check

In order to make sure that all music is played back at the same level, several streaming services have now implemented loudness normalization. For example Apple Music uses

what it calls "Sound Check" to adjust all songs to a set target gain level of -16LUFS.

Spotify has implemented its own normalization process that aligns everything to -11LUFS and it uses a limiter when necessary to ensure peaks don't go above 0dBFS, which means that the process can be more destructive to the audio. Other services have other target levels, which means that there's currently no standard.

What this means is that a song that's absolutely crushed in level will play back at the same level as one with loads of dynamic range. The big difference is that the one with lots of dynamic range will sound better. This is something to keep in mind while mastering and there's the temptation to get every last tenth of a dB in level. Thanks to loudness normalization processes like Sound Check, the loudness wars will soon be over and our audio will sound better than it has in a long time.

Creating A FLAC File

A format that has recently gotten a lot of attention is the lossless FLAC format, which stands for Free Lossless Audio Codec. It works somewhat the same as a standard MP3 file, only it's lossless, like a zip file, and designed specifically for audio. Unlike other lossless codecs by DTS and Dolby, FLAC is non-proprietary, is unencumbered by patents, and has open-source implementation. What's more, FLAC has been adopted as a release format of choice by some of the world's biggest recording artists, from Pearl Jam to Nine Inch Nails to the Eagles, and even reissues from The Beatles.

FLAC supports a bit depth from 4 to 32, and up to 8 channels. Even though it can support any sampling rate from 1Hz to 655,350Hz, you don't need to specify a bit rate because it automatically determines it from the source file. Plus it has a "cue sheet" metadata block for storing CD tables of contents and track and index points. It's an excellent way to deliver the highest-fidelity music file with a reasonably small file size, but it's not yet supported by all applications or players.

Although many DAWs don't have a FLAC encoder built in, there are a number of players and encoders that can be downloaded for free, as well as QuickTime playback components and iTunes scripts.

Submitting To Online Stores and Services

If you want to distribute your work via online stores, you might want to consider one of

the many distribution services. While you may be able to submit your songs to some online stores, others (such as iTunes) require a large record-label account, meaning that you need a digital distributor to get your songs placed in the store. Plus, each online store has different file format requirements, which can cause you to spend a lot of time with file preparation, when submission is just a single click away with CD Baby, TuneCore, DistroKid or ReverbNation (among others).

Table 9.4 gives a quick overview of some of the available digital music distributors. Keep in mind that this data was accurate as of the writing of this book, but things change rapidly in this digital world we live in so they may be completely different by the time you read this.

Table 9.4: Digital Music Distributors

	CD Baby	TuneCore	DistroKid	ReverbNation
Single fee	\$9.95	\$9.99	\$19.99 for unlimited songs	\$34.95 (annual fee)
Album fee	\$49	\$29.99 (first year); \$49.99 (each year thereafter)	\$19.99 for unlimited albums	\$34.95 (annual fee)
Sales commission	9%	0%	0%	0%
Number of digital partners (iTunes, Amazon, etc.)	95	74	150	40
CD sales	Yes	No	No	Yes
Commission on physical sales	\$4	n/a	n/a	\$5.49
CD and vinyl distribution	Yes	No	n/a	No

In a nutshell, TuneCore, DistroKid, and ReverbNation charge annual fees but don't take a percentage of your sales. CD Baby takes a 9 percent cut but doesn't charge an annual fee. If you want physical distribution, only CD Baby and ReverbNation offer that service.

Submitting To Online Song Databases

A number of online databases store album information that is accessed by such programs as iTunes, Windows Media Player, and Winamp to display album and song info on computers. Perhaps the best known is CDDB, or Compact Disc Database, which is a database that allows a music-player application to look up audio CD information online, which then displays the artist name, CD title, track list, and any additional information available.

The information in a database like CDDB is linked to CDs that have had their information submitted to the database service. The artist is totally responsible for all information in the CD text, but it's still possible for the information to be corrected after the fact if an error was made in the original submission.

There are other online databases besides CDDB, including Muze, freedb and MusicBrainz. Although the CD-identification process used by these databases may differ from the original CDDB process, the concept is the same.

Since the CDDB database was just purchased by Nielsen, the best way to submit your album data at the moment is through the iTunes app. This is done by naming the CD tracks and then using the Submit Track Names option under the Advanced heading in the toolbar to submit track information. Here's how it's done:

1. Open iTunes and place your CD into your CD drive. (Click No when asked to import.)
2. Click into a track name and select *Get Info*.
3. On the *Info* tab, type in track name, artist name, album name, select a genre and year of release. The next button will take you to the next track. Continue until all tracks are titled.
4. Go to Advanced and click Submit CD Track Names.

Wait two or three days, call up the album, go to *Advanced*, and click *Get CD Track Names*. This is a re-query button that clears your local cache and shows that your CD's information now comes from CDDB.

TIP: *Because identification of albums is based on the length and order of the tracks, CDDB can't identify playlists in which the order of tracks has been changed, or compilations of tracks from different CDs. CDDB also can't distinguish between different CDs and digital albums that have the same number of tracks and the same track lengths.*

Mastering For iTunes

With the iTunes Store and Apple Music now such a huge part of online music distribution, knowing the latest on how they treat your audio can be beneficial to the final audio quality of your file. iTunes is actually composed of two different categories: standard iTunes audio and the new Mastered for iTunes high-resolution audio. Let's look at both.

A Look At AAC, The iTunes File Format

iTunes uses the AAC file format (which stands for Advanced Audio Coding) as a standard for the music in its store and for the Apple Music streaming platform. Contrary to popular belief, it's not a proprietary format owned by Apple. In fact, it's part of the MP4 specification and generally delivers excellent-quality files that are about 30 percent smaller than a standard MP3 of the same data rate. All new music destined for the iTunes Store or Apple Music is encoded at 256kbs at a Constant Bit Rate with a 44.1kHz sample rate.

While Apple does the encoding for you (there's no way around it if you want your music available on its platforms), here are some of the parameters of the AAC encoder that are available if you do your own encoding.

- **Stereo Bit Rate.** This allows you to select the bit rate. The standard setting is now 256kbs, and the highest-quality setting for this format is 320kbps.
- **Sample Rate.** This enables you to select the sample rate.

TIP: Never use a higher sample rate than the rate used for the source. In

other words, don't use 48kHz if your source was 44.1kHz. Doing so will make the file larger without gaining anything in terms of quality.

- **Variable Bit Rate Encoding (VBR).** This option helps keep the file size down, but the audio quality might be affected. VBR varies the number of bits used to store the music depending on the complexity of the sound. If you select the Highest setting from the Quality pop-up menu for VBR, iTunes encodes up to the maximum bit rate of 320 kbps in sections of songs where the sound is complex enough to require a high bit rate. Meanwhile, iTunes keeps the rest of the song at a lower bit rate to save file space. The lower limit is set by the rate you select in the Stereo Bit Rate pop-up menu.
- **Channels.** This pop-up menu enables you to choose how you want the music to play through speakers—in stereo or mono. Select the Auto setting to have iTunes detect the number of channels in the file.
- **Optimize for Voice.** This option is meant for podcasters, and it filters the audio to favor the human voice, which is obviously not something you'd want for music.

It's best to select the highest bit rate in the Stereo Bit Rate pop-up menu and leave the other two pop-up menus set to Auto.

Mastering Tips For iTunes And Apple Music

There are a number of tips to follow to get the best sound quality for an iTunes or Apple Music encode. As it turns out, the considerations are about the same as with MP3 encoding:

- **Turn it down a bit.** A song that's flat-lined at –0.1dBFS isn't going to encode as well as a song with some headroom. This is because the iTunes AAC encoder sometimes outputs a tad hotter than the source, so there are some inter-sample overloads that happen at that level that aren't detected on a typical peak meter, since all DACs respond differently to it. As a result, a level that doesn't trigger an over on your DAW's DAC may actually be an over on another playback unit.

If you back it down to –0.5 or even –1dB, the encode will sound a lot better and your listener probably won't be able to tell much of a difference in level anyway.

- **Don't squash the master too hard.** Masters with some dynamic range encode better. Masters that are squeezed to within an inch of their life don't; it's as simple as that. Listeners like it better, too.
- Although the latest AAC encoder has a fantastic frequency response, sometimes **rolling off a little of the extreme top end** (16kHz and above) can help the encode as well.

Any type of data compression requires the same commonsense considerations. If you back off on the level, the mix buss compression, and the high frequencies, you'll be surprised by just how good your AAC encode can sound.

iTunes Sound Check

As stated in the previous chapter, iTunes utilizes a feature called Sound Check, which scans the songs in your library and normalizes the volume information so they all play back at the same level. This volume data is stored either in the "Normalization Information" ID3 tag or in the iTunes Music Library database. If you encode or rip a song from a CD with iTunes, the Sound Check level is stored with the song's ID3 tags. Sound Check is designed to work with MP3, AAC, WAV, and AIFF files and won't work with a file that iTunes can't play. The audio data of the song is never changed.

Sound Check is always on in Apple Music, but it defaults to Off for regular iTunes playback. This does cause a dilemma for artists, producers, record labels, and mastering engineers though.

Sound Check will automatically lower a very loud song and boost a quiet one so they play at the same level. Because the loud track has few dynamics, it appears to sound lifeless and even quieter than the one that was less compressed, because it has fewer peaks. This is causing people to rethink the value of a loud master, as it could end up being counterproductive when played back in a level-controlled environment.

Spotify is another service that uses a form of Sound Check, and other services may institute something similar as well.

TIP: *Keep Sound Check in mind when mastering for online distribution. Lowering the level a few dB can actually be advantageous, resulting in a better sounding and even louder master when played back in this environment.*

The Mastered For iTunes Format

Mastered for iTunes is an alternative to the standard iTunes quality, where the iTunes Store accepts high-resolution master files and provides higher-quality AAC encodes as a result. Music files that are supplied at 96kHz/24-bit will have a Mastered for iTunes icon placed beside them to identify them as such, although any sample rate that's a 24-bit file is also considered high-resolution (see Figure 10.1).

Songs	Ratings and Reviews	Related	Mastered for iTunes	
Name	Artist	Time	Popularity	Price
1. I Blame You	Ledisi	4:14		\$1.29
2. Rock With You	Ledisi	3:24		\$1.29
3. That Good Good	Ledisi	3:17		\$1.29
4. Lose Control	Ledisi	4:39		\$1.29
5. Like This	Ledisi	3:48		\$1.29

Figure 10.1: Songs indicated as Mastered for iTunes
(Source: Apple Inc)

Mastered for iTunes (or MFiT, for short) doesn't necessarily mean that the mixer, producer, or mastering facility does anything special to the master except to check what it will sound like before it's submitted to iTunes, and perhaps check it later again before it's posted in the iTunes Store. That said, this is not something that any engineer or mixer can do. In order to have your master qualify for MFiT submission, it must be mastered by a certified master studio. All of the aggregators have the ability to submit to the MFiT program as well.

All encoding for MFiT is still done by Apple, not by the mastering engineer, record labels, or artists (see Figure 10.2). According to Apple, the reason for this is to keep the encodes consistent and to prevent anyone from gaming the system by hacking the encoder, but also to avoid any potential legal problems that might occur when a mixer, producer, or mastering house sends the files directly to iTunes without the label's permission, or uses different submission specs.




Figure 10.2: A block diagram of the MFIT production chain
© 2017 Bobby Owsinski

Mastered for iTunes is only an indication that a high-res master was supplied; it's not a separate product. There will always be only one version of the song on iTunes, and it will be available at the same price regardless of whether it's Mastered for iTunes or a normal submission. Mastered for iTunes doesn't mean you get to charge more or that iTunes charges you more. Everything is like it was before, you just supply a high-res master so it sounds better.

The Mastered For iTunes Tools Package

Even though the mixer or mastering engineer doesn't do any encoding directly, Apple has provided a set of tools that can be used to hear what the final product will sound like when it's encoded. That way, any adjustments can be made to the master before it's submitted to iTunes to ensure that the encode will provide the highest quality. You can find these tools, as well as a PDF explaining Mastered for iTunes in depth, at apple.com/itunes/mastered-for-itunes.

Along with the mastering tools, be sure to also download the AU Lab tool, because it acts as a host for one of the important tools in the set, AURoundTrip. Included in the mastering tools are two utilities, afconvert and afclip, that can be accessed only via the Terminal program in Mac OS X and that require some Unix command-line knowledge.

There is a PDF called Mastered For iTunes that describes each tool and the Terminal command lines. We'll cover the various MFiT apps below, but understand that not many mastering engineers use them all that much.

Using The afconvert Tool

The afconvert utility is command-line tool that will let you encode your masters using exactly the same technology used to encode files for the iTunes Store. The afconvert utility is actually built into Mac OS X and can be accessed using the Terminal application.

Using The Droplet Tool

The Mastered for iTunes Droplet is a stand-alone drag-and-drop tool that's a much quicker and easier way to encode your masters to the AAC format than using the afcovert tool. All you have to do is drag and drop AIFF or WAV format source audio files, or folders containing those files, onto the Droplet. Keep in mind that regardless of the sample rate, the Droplet will automatically convert the file to 44.1kHz.

Once again, the only reason that you'll be converting to an AAC file format is to hear what the file will sound like after it's posted on the iTunes Store. iTunes does not accept AAC files, as it's done by Apple online from a converted AIFF or WAV file.

TIP: Remember that to be Mastered for iTunes, you have to supply the highest-resolution AIFF or WAV file available.

Using The Audio To WAVE Droplet

Just as the iTunes Droplet tool made encoding to AAC easier, the Audio to WAVE Droplet makes converting an AAC file back to WAV easy as well. It works with the WAV file format, as well as any audio file that is natively supported on Mac OS X, such as MPEG or CAF files.

To use the Audio to WAVE Droplet, drag and drop source audio files, or folders containing those files, onto the Droplet. The Droplet will then convert those files to WAV-format files. The newly created WAV files will be titled using the names of their corresponding source files, and they will be placed in the same folder as the source files.

Using The afclip Tool

The afclip command-line utility can be used to check any audio file for clipping. This tool works by examining an audio file and identifying areas where clipping has occurred (see Figure 10.3). It accepts audio files as input and outputs a stereo sound file containing the left channel of the original file and a right channel with graphically represented impulses corresponding to each clipped sample in the original. This sound file can then be loaded into a DAW so that you can see a visual map to locate any clipping that may have occurred.

297.713568	14290251.25	1	1.001909	0.016570
297.713573	14290251.50	1	1.006547	0.056684
297.713573	14290251.50	2	1.002930	0.025410
297.801276	14294461.25	2	1.021100	0.181363
297.801281	14294461.50	2	1.032531	0.278058
297.801286	14294461.75	2	1.024644	0.211463
297.805089	14294644.25	2	1.022340	0.191905
297.805094	14294644.50	2	1.023801	0.204315
297.805099	14294644.75	2	1.004619	0.040030
297.810130	14294886.25	2	1.020037	0.172319
297.810135	14294886.50	2	1.015371	0.132492
297.812109	14294981.25	2	-1.016063	0.138417
297.812115	14294981.50	2	-1.028146	0.241094
297.812120	14294981.75	2	-1.014752	0.127202
297.939589	14301100.25	2	1.020224	0.173911
297.939594	14301100.50	2	1.017918	0.154256
297.997297	14303870.25	2	-1.010133	0.087573
297.997302	14303870.50	2	-1.034622	0.295637
297.997307	14303870.75	2	-1.029110	0.249239
298.013411	14304643.75	2	-1.004668	0.040450
298.024568	14305179.25	2	-1.016847	0.145116
298.024573	14305179.50	2	-1.017927	0.154329
total clipped samples		Left	on-sample:	0 inter-sample: 177
total clipped samples		Right	on-sample:	0 inter-sample: 1837

Figure 10.3: A readout of the points in a song where clipping occurred
© 2017 Bobby Owsinski

To check an audio file for clipping with afclip:

1. Open a Terminal window.
2. In the Terminal window, type the following on one line, *followed by a space (this is important)*: afclip
3. Drag and drop the audio file you wish to check onto the Terminal window.
4. Press Return to run afclip.

The readout will contain the following information for each instance of clipping:

- **Seconds.** The time, in seconds, where the clipping occurs.
- **Sample.** The sample number that was clipped.
- **Channel.** The channel of the clipped sample. A value of 1 means the clipping occurred on the left channel, while a value of 2 means the clipping occurred on the right channel.
- **Value.** The raw value of the clipped sample. Since clipping happens when a value exceeds the range of –1 through 1, these values will be below –1 or above 1.
- **Decibels.** The number of decibels by which the sample exceeds the clipping point.

The readout will end with a summary of how many total clipped samples the audio file contains for the left channel and the right channel.

By default, afclip will give a readout of any clipping found. A sample that indicates a decimal other than .00 (such as 63224.25 or 2821.50) means that the distortion is occurring between the samples. This is known as *intersample distortion*.

iTunes won't reject a master file based on the number of clips the file contains. This tool is there just so you can make an informed decision about whether to submit an audio file or go back to the drawing board and make adjustments, which is a creative decision that's entirely up to you.

TIP: *afclip will find many instances of clipping that are inaudible. Unless you're hearing clipping that you're trying to track down, using this app can be a lesson in futility, especially with a client who just wants the loudest product in spite of any clipping.*

Using The AURoundTrip AAC Audio Unit Tool

The AURoundTrip AAC is another tool that can be used to compare an AAC file to the original source audio file to check for clipping. It includes clip and peak detection, as well as a simple listening test environment. The audio unit can be used in any audio-unit host application, such as Logic or AU Lab. AU Lab is available as a free download at apple.com/iTunes/Mastered-for-iTunes. It's mainly designed as a platform reference Audio Unit host, and is designed for Audio Unit auditioning and live applications.

AURoundTrip AAC outputs only the distortion and clipping so you know exactly where it's occurring. Once again, iTunes won't reject a file with distortion and clipping, but if you know it's there you can fix it before it's submitted so your listeners get the best-sounding songs available.

Using The Test-Pressing Feature

One unique aspect of Mastered for iTunes is something that hasn't been publicized much, called a test pressing. To those in the Mastered for iTunes program, iTunes will send an AAC file back to the label/engineer/artist to check before it's posted. If they sign off on it, the song then goes on sale in the iTunes Store (see Figure 10.4).




Figure 10.4: Test-pressing block diagram

© 2017 Bobby Owsinski

This has proven to be a great tool not so much for catching bad encodes, but for finding more egregious errors, such as the wrong master or even entirely incorrectly labeled songs. Hopefully, the test-pressing feature will be used more and more in the future.

Submitting To The iTunes Store Or Apple

Music

It's not possible to submit directly to the iTunes Store or Apple Music if you're an indie artist, a band, a producer, or even a small label. Apple reserves that feature for larger labels with hundreds of titles already in their catalog, and who release titles on a regular basis. However, you can easily get your music on the iTunes Store by using a digital distributor, such as TuneCore, CD Baby, DistroKid, ReverbNation, Nimbit, or others.

To submit directly to the iTunes Store, labels must pass the requirements for submission, then they use a program known as iTunes Producer to submit individual songs and albums. This free program not only allows uploading product to the store and Apple Music, but also allows the label to set the price and input all the metadata associated with the project.

This is one of the reasons why mastering engineers can't directly submit songs to iTunes on behalf of their clients. The use of the program is exclusive to the label for their product only, and any payments from sales will only return to the bank account associated with that particular iTunes Producer account.

Chapter 11

Other Types Of Mastering

Although music is more often mastered than other kinds of audio projects, mastering is often called for when music is to be used in other mediums as well. In this chapter we'll examine some of those situations as we look at mastering for high-resolution audio, television, and film, as well as some of the new online mastering solutions.

Online Mastering

Mastering costs vary a great deal and can vary anywhere from about \$350 an hour to about \$500 for an A list mastering engineer. As a result, you can expect the overall price for mastering an album at a top facility to come in anywhere from around \$1,000 to around \$5,000.

There now are some very adequate alternatives to spending a lot of dough on mastering however. By far the cheapest is to use one of the automated online mastering services like Landr, eMastering or Aria. The prices of these services vary anywhere from as low as \$4 a song to a monthly fee of around \$100, where you can upload an unlimited number. Depending upon the mix, the results can be surprisingly good, and you're usually allowed to try different mastering settings for just a single price so you're able to choose been a more aggressive or more dynamic mix.

Another alternative is to submit your finished tracks online for mastering at some of the top mastering houses like Abbey Road or Sterling Sound. Although you won't get one of the top engineers nor will be able to attend the session, you'll get a great job by a junior engineer at a price that's much reduced from the facility's normal rates.

High-Resolution Mastering

Thanks to high-capacity discs such as Blu-ray and even some online sites, such as HDTrax.com, the demand for high-resolution audio files is growing rapidly. While the exact specification of "high resolution" can vary depending upon the situation, to most audiophiles it means a sample rate of at least 96kHz at 24 bits. While the mastering process is the same as with standard-def material, other considerations apply when it comes to high resolution.

The first is storage. We know that a 44.1kHz/16-bit stereo minute on a CD needs approximately 10-1/2MB of storage (actually 10.58 MB). Most of us record on our DAWs at 48/24, and that takes about 17.28 for each minute that we record (see Table

11.1).

Table 11.1: High-Resolution Storage Requirements

Resolution	Storage Required (per Minute)
44.1kHz/16-bit CD file	10.58MB
48kHz/24-bit DAW file	17.28MB
96kHz/24-bit high-res	34.56MB
192kHz/24-bit high-res	69.12MB
5.1 surround at 192kHz/24-bit	207.36MB

When it comes to high-res, we go to a different league though, with a minute of true 96kHz/24-bit stereo needing 34.56MB, and a minute of 192k/24 needing 69.12MB. Again, this is just for stereo. If we were dealing with a six-channel 5.1 file at 192k, the storage required would be a staggering 207MB per minute.

In these days of cheap storage where a terabyte doesn't cost very much, this might not sound like a lot of storage space, but it can really get you in a bind if you haven't planned for it.

96/24 operation doesn't stop just at storage though. All equipment in the digital signal chain, including A/D and D/A converters, plugins, and workstations, must now be able to process at least 96/24 as well. And keep in mind that the higher the resolution, the more processing power that's required from your computer's processor and RAM.

Most modern digital audio workstations are up to the task of working with high-resolution files, but be aware that everything is not always as easy as in the standard-res world.

***TIP:** While many delivery formats can now support sample rates up to 192kHz, it does no good to export a file at a higher sampling rate than the source. There's no audio advantage, and the file will be substantially larger as a result. For instance, upsampling a 48kHz source file to 96kHz will not result in an increase in quality. Once a file is exported with a selected sample rate, it should only be exported at that same sample rate or less when mastered.*

Direct Stream Digital (DSD)

When the Super-Audio CD was first introduced in 1999, it was heralded as a new age in audio reproduction, mostly thanks to the encoding process known as Direct Stream Digital (DSD). Despite the massive marketing efforts by its creators Sony and Philips

(who also created the CD), SACD never achieved the market penetration that was expected, and today it has gone the way of CDs, as fewer consumers want to invest in a shiny plastic disc. That being said, DSD files are still desirable, at least among many audiophiles, and can be downloaded from a number of online sites.

DSD differs from other analog-to-digital recording processes in that a single bit measures whether a waveform is rising or falling rather than measuring an analog waveform at discrete points in time (see Table 11.1). In current systems, this one bit is then decimated into LPCM, causing a varying amount (depending upon the system) of unwanted audio side effects (such as quantization errors and ringing from the required brick-wall filter). DSD simplifies the recording chain by recording the one bit directly, thereby reducing the unwanted side effects. Look at DSD-Guide.com for more information.

Indeed, on paper DSD looks impressive. A sampling rate of 2.8224MHz (which is 64 times 44.1k, in case you're wondering) yields a frequency response from DC to 100kHz with a dynamic range of 120dB. Most of the quantization errors are moved out of the audio bandwidth, and the brick-wall filter, which haunts current LPCM systems, is removed. To enable a full 74 minutes of multi-channel recording, Philips also developed a lossless coding method called Direct Stream Transfer that provides a 50 percent data reduction. The original DSD sampling rate has since given way to a higher rate of 5.6448MHz, which is now offered by some recorders, including those in the Korg range of products.

Even though many agree that DSD sounds superior to the LPCM technology used in most of the audio world today, there are severe limiting factors when it comes to using it. First of all, processing and editing DSD natively in a workstation isn't easy. There are only two DAWs being made today, Merging Technologies' Pyramix and the Sonoma (which was developed by Sony). Other systems (such as Korg's AudioGate) transcode the DSD stream to LPCM at 96/24 or 192/24 for editing and processing, then back to DSD again. Of course, separate DSD D/A convertors are also required. These systems can be quite a bit more expensive than a normal PCM mastering system, and considering the limited market, usually aren't worth the investment for most mastering engineers. That said, many "golden ears" claim that it's the next best thing to analog.

Blu-Ray Disc

Blu-ray is the name of the optical disc format initially developed by Sony and Philips (inventor of the compact disc, cassette, and laserdisc) as a next-generation data and video storage format alternative to DVD.

The format was developed to enable recording, rewriting, and playback of high-definition audio and video, as well as storing large amounts of data. It offers more than five times the storage capacity of traditional DVDs and can hold up to 25GB on a single-layer disc and 50GB on a dual-layer disc.

The name Blu-ray is derived from the underlying technology, which utilizes a blue-violet laser to read and write data. The name is a combination of “Blue” (blue-violet laser) and “Ray” (optical ray). According to the Blu-ray Disc Association, the spelling of Blu-ray is not a mistake; the character “e” was intentionally left out so the term could be registered as a trademark.

Blu-ray has an extremely high data rate compared to DVD, which means that both high-definition audio and video are possible at the same time. It’s also possible that Blu-ray could allow the storage capacity to be increased to 100GB–200GB in the future simply by adding more 25GB layers to the discs.

Blu-Ray Audio Specs

Blu-ray supports just about every audio codec format, but because of the high data rate, it’s now possible to store up to 8 channels of 96/24 LPCM audio and 6 channels of 192/24 without data compression and hi-def picture. That being said, codecs, lossy and lossless, are normally used if video is present on the disc, depending upon the definition.

Table 11.2 presents a list of the mandatory and optional audio formats. Mandatory means that all Blu-ray players are required to decode the format. A secondary audio track, if present, can use any of the mandatory formats or one of the optional codecs.

Table 11.2: Blu-Ray List Of Codecs

Codec	Channels	Sample Rate	State	Type
Linear PCM (LPCM)	Up to 8	96/24	Mandatory	Lossless
Linear PCM	Up to 6	192/24	Mandatory	Lossless
Dolby Digital	5.1	48/24	Mandatory	Lossy
Dolby Digital EX	6.1	48/24	Optional	Lossy
Dolby Digital Plus (DD+)	7.1	48/24	Optional	Lossy
Dolby True HD	Up to 8	96/24	Optional	Lossless
Dolby True HD	Up to 6	192/24	Optional	Lossless
DTS Digital Surround	5.1	48/24	Mandatory	Lossy
DTS Digital Surround ES	6.1	48/24	Optional	Lossy
DTS Digital Surround 96/24	5.1	96/24	Optional	Lossy
DTS-HD Hi-Res Audio	Up to 8	96/24	Optional	Lossy
DTS-HD Master Audio	Up to 8	96/24	Optional	Lossless
DTS-HD Master Audio	Up to 6	192/24	Optional	Lossless

Mastering Music For Film

Except on rare occasions, the only audio that gets mastered for film is for the songs intended for the movie, as the film studio or production company usually takes care of the underscore, dialogue, and effects. In fact, most of the time the studio does the music as well, but occasionally a recording artist is asked to record the score or songs specifically for a movie, and since the artist feels comfortable continuing his or her normal way of working, the score or songs get mastered.

In that case, the music is mastered as normal and delivered to the dubbing stage, where the dubbing mixer lays it into the movie at the required level. The need for the hottest level doesn't really exist because it will always get adjusted to fit the film anyway.

On a side note, one of the reasons that the music score for a movie is not normally mastered is that the movie powers-that-be (producer, director, music editor, dubbing mixer) usually ask for the score to be delivered as a 5.1 surround mix with stems. Stems are individual submixes of the final mix that allow the dubbing mixer to weave the music around the effects and dialogue so all can be distinctly heard. Stems are usually delivered as a 5.0 (no LFE channel) mix of the music bed minus the bass, any lead instrument or vocals, and any instruments with a lot of high-frequency information. The bass is then delivered on a separate track, and the lead instrument or vocal, and instruments with high-frequency info, are each delivered as separate 5.0 mixes (which include all reverbs and ambience). The dubbing mixer then completes the music mix with the rest of the movie.

Mastering Music For Television

The majority of the time, mastered music intended for television is delivered to the post-production facility or video editor editing the program, where it's mixed in against the video. The video editor then determines the correct level against the effects and dialogue, just like with film.

On the rare occasion when the television audio is coming from the mastering engineer, the first thing to do is obtain a technical specification from the engineering department of the network on which it's going to be shown. This will tell you exactly what they want and how they want it.

The Effect Of The CALM Act

In December of 2013, a new law called the Commercial Advertisement Loudness Mitigation (CALM) Act went into full effect. The rule requires TV stations, cable

operators, and satellite-television providers to control the audio loudness of all programs and commercials that are broadcast. The law is in response to years of complaints from television viewers about commercials being much louder than the programs surrounding them because they had been more heavily compressed.

The law has some serious teeth in it for violations, including heavy fines and loss of a violator's broadcast license for continued offenses. Because of that, all networks and broadcasters are now especially concerned about the level of any program or commercial supplied, and have a zero-tolerance policy for volume that strays outside of the spec (officially called the Advanced Television Systems Committee A/85RP).

As anyone who has mixed or mastered knows, the relative volumes of two different songs can be very different even though their levels can look the same on a variety of meters, thanks to the amount of compression added. That's why a new metering system had to be developed to measure the loudness of a program as our ears hear it.

Meet The LKFS Scale

The new measurement is called LKFS, which stands for Loudness, K-weighted, relative to Full Scale, which distinguishes itself from the normal dBFS peak meters found on all digital gear in that it measures the loudness not instant by instant, but over a period of time. In Europe, the measurement is called LUFS, which stands for Loudness Units relative to Full Scale. At one point there was a difference between the two, but today they are identical, so most loudness meters indicate both. LUFS is easier to say, so that seems to take precedence when engineers talk loudness.

The new federally mandated loudness specification is $-24\text{LKFS} \pm 2\text{dB}$, which means that the loudness of your program better be between -26LKFS and -22LKFS , or the program is getting kicked back to be redone at the mandated level.

It's even trickier than that, though, since the measurement must be made around an "anchor element," which for television means dialogue. That means that the dialogue must always be around the -24LFS level, while music and effects can momentarily peak above, maybe as high as -16LKFS for brief periods.

While the spec calls for -24LKFS at $\pm 2\text{dB}$, many broadcast networks have even tighter specs, holding their clients to $\pm 1\text{dB}$. That means there's little room for error when mixing a program intended for television with dialogue.

TIP: *If the master is all music that will be supplied to a video editor, this spec does not apply and you can master as always.*

There's really no way to estimate LKFS from a VU, peak, or PPM meter, so it requires a

specialized metering tool made just for this application. The Dolby Media Meter 2, TC Electronic LM2 or LM6, and the Waves WLM meter are the most widely used products on the market at the moment (see Figure 11.1).



Figure 11.1: LUFS Readout

Source: Waves Audio LTD

Keep in mind that television networks are very strict with their specs, and a violation will result in the project being kicked back to you to do it again. So on those times that you're asked for television delivery, paying close attention to all the details will ultimately result in a lot less hassle.

Chapter 12

Archiving The Master

After the creative part of mastering is complete, the job still isn't finished. Many mastering engineers are asked to deliver the product for replication, while in other cases the client prefers to do it himself. That said, the final task at hand is archiving the project in case changes are required at a later date.

Delivering The Master To The Replicator

When sending a master to a CD replicator, most now prefer the file be sent via FTP. The best way to do that is with a FTP app, such as Fetch on a Mac or FileZilla on a PC. Most FTP apps are either free or very low cost, and are the very best way to quickly and safely send large files to a client, distributor, or replicator. Some DAWs specifically designed for mastering even have FTP delivery built in.

Your replicator will provide all the info needed to make the transfer, as well as any help if needed.

We like to send it directly to the factory. When someone asks for the CD master direct, we explain to them that it may or may not be a good idea, and they usually agree. We charge a premium for a master, and that includes a QC check by my assistant. It's time-consuming, but it's necessary. We send mostly DDPs these days. Now QC includes spell-checking of the metadata as well.

—David Glasser

Archiving The Project

Someone once said that the difference between an amateur and a pro is that a pro has a backup, and nothing could be truer. Even though you may have given the client a final signed-off master, it's always a good idea to archive the project in case you're called on to do some fixes later, or as a backup in case the master given to the client is lost (it happens more often than you think, especially after some time has passed).

While normal backups are performed on hard drives, flash drives, or in the cloud, long-term archiving requires a different strategy. Essentially, two different backups are required:

- **The master.** This is a copy of all the master formats that were requested for the project, including CD, MFIT, MP3, DDP, vinyl, or high-res files.
- **The master session.** This contains the DAW session as well as all of the original source files.

The archive can live locally on a hard drive, an archival-grade disc, or a tape backup system or in the cloud for the long-term storage.

Most mastering engineers will keep a mastering session available on a data drive for as long as a year, since clients frequently request changes or even new masters after the project is released. Most facilities will usually make a backup of a project directly after the session, and a more permanent version well after the project is approved and completed. This backup is then put into the facility vault onsite. The record company (if there is one) or client may be charged for this service, but not always. Many mastering engineers will always keep a backup on file anyway, because it's not uncommon for masters to be lost by the client.

Just to show how backups can come in handy, with the advent of Mastered for iTunes, record labels are asking for high-res versions of older albums. For the mastering facilities that made it a point to begin archiving high-res versions years ago, this is now a simple case of retrieving the session and exporting the file for MFIT, instead of searching for the source files and mastering the session all over again.

Archiving your sessions may seem like a pain, but it'll make you look like a hero when a client calls in a panic. Do it often.

***TIP:** It's a good idea to have more than one backup, one onsite and the other offsite. The offsite backup might be an additional drive or be stored by a cloud service, such as CrashPlan, Mozy, Carbonite, or others.*

PART II

THE INTERVIEWS

As always, the interview portion of the book is the most enjoyable from a personal standpoint. It's a wonderful thing to finally meet (at least over the phone) the people whose work I've been listening to for many years. Not only were the contributors most willing to share their working methods and techniques, but they were most gracious in taking time from their busy schedules to do so. For this I am most grateful and extend to them my heartfelt appreciation.

Since this book is about mastering as an entire profession, I've included a cross section of the industry. Not only are the legends and greats represented, but also some engineers that deal in the specialty areas of mastering (the near greats?). Regardless of their perceived industry stature, they all toil in the everyday trenches of mastering, and much can be learned from their perspective.

Chapter 13

Greg Calbi - Sterling Sound

Greg Calbi started his career as a mastering engineer at the Record Plant in New York in 1973 before moving over to Sterling Sound in 1976. After a brief stint at Masterdisk from 1994 to 1998, Greg returned to Sterling as an owner, where he remains today. Greg's credits are numerous, ranging from Bob Dylan and John Lennon to David Bowie, Paul Simon, Paul McCartney, Norah Jones, Branford Marsalis, and Bon Iver, among many, many others.

Do you have a philosophy on mastering?

I do. My philosophy in general is to try to figure out how to improve what the person brings me and then try to figure out what his intent was. In other words, I don't just plug in my own idea without first really communicating with the client. It's a little tricky because it really is different for every project. You have to get a good communication flow going, which sometimes is actually one of the most difficult parts of the job.

One time somebody said something to me that I thought was the best compliment I ever got in mastering. He said, "The reason I like your work is because it sounds like what I did, only better." That's what I've always tried to do. I try not to change the mix; I just try to enhance it. I go with the spirit of what was given to me, unless I really feel that it's totally missing the mark.

Is there a difference between mastering from coast to coast or city to city?

There's really more of a difference from person to person. I've listened many years to all the different sounds that different guys have, and they really all do something different, and I respect every one of them for it. I could be blown away by something that any of 10 guys do, it's so recognizable.

We once hosted a great symposium that NARAS ran for their members. They had about 90 people come up, and the four of us from Sterling—George Marino, Tom Coyne, Ted Jensen, and I—had the same mix to work on. We had 10 people in the room at a time, and we had a make-believe producer who asked producer-type questions so people could see how a session went. We all EQ'd the same song, and after it was over, we all went out to the main room and listened to it with everybody there. All four sounded like four different mixes, and they all had their own thing about them. None of them sounded bad, but it was amazing how different they all were.

Can you hear the final product in your head as you're running something down?

Yeah, I can hear where I want it to go. I use kind of an A/B method most of the time, so

I'm always referring to other mixes on the album. What I try to do is get a listen to everything on the album before I start to work on it. I really want to know what the producer and the engineer are capable of doing at their best before I start to force it in a direction.

In other words, if the first song goes a certain way, all of a sudden you're trying things and going back and forth and just going crazy. Then all of a sudden, about an hour or two later, you find that you might not have done your best work because you were moving a mix in a certain direction. Whereas if I go to the stuff that I really like hearing in the beginning, it gives me more of a realistic expectation of what I'm going to be able to get from this stuff later on. It's just a good way to give your ears something to compare to.

Do you listen to the whole album before you start?

I'll listen to snatches of everything, with maybe a minute or two of a few songs. I'll ask the client, "What's your favorite mix on the album? What's the one that everybody seems to really like?" because that'll give me an indication of their expectations. If they point me to something that I think is horrible and that they think is great, then I know I have to use a combination of engineering and psychology because I need to bring them to where I know it might have to be.

The funny thing is that as the years have gone on, they now seem to throw it into my hands almost totally, and I have to drag them back into it. I find I work better when the client gets involved because when they take some responsibility for the project in the room, they'll also take that same responsibility when they're listening out of the room. A lot of mastering guys kick the people out and are really secretive about what they're doing, but I'm completely the opposite. The black-magic thing is really totally overrated. It's kind of a fallback for a certain amount of not taking responsibility.

What do you think makes the difference between a really great mastering engineer and someone who's just competent?

A great set of ears, but communication skills is another thing that makes somebody great, as well as a willingness to try different things. It's kind of a combination of creativity and tenaciousness.

What do you think is the hardest thing for you to do?

Hard rock and metal have always been the hardest thing for me to make sound good because the density of the music requires a lot of aggressiveness. If the aggressiveness goes just that one step too far, it diminishes the music. You reach a point where all of a sudden it starts to reverse itself, where big becomes small and exciting becomes overbearing, and it works against the rhythms of the music. If it's just one step past the point, it loses impact.

Another thing that's hard is when the low end is thin and light, because it's really hard to create low end when there is none. If you have a really muddy project, you can always clear stuff away, but it's really tough when the bottom end isn't there. Most of the problem projects have to do with the bass being recorded poorly. If you made book of excuses, the chapter on bass would be eight times bigger than the chapter on everything else.

The fact of the matter is that you never have a great-sounding mix and master if you don't have a great bass sound. It can't be great unless the bass is great. It could be good, but bass is what takes it to the level where it's really something special. It's the thing that engineers are the most frustrated about.

What would you tell someone who's trying to learn how to master?

The main thing is that all you need is one experience of hearing somebody else master something where it sounds so incredibly different that it makes you then realize just how intricate mastering can be, and just how much you could add to or subtract from a final mix.

Also realize that there's a hidden element where the more flexibility and patience you have, the more likely you are to come up with something that's going to be better. There's no shortcut to it. You just have to keep A/Bing back and forth and back and forth. It's pretty amazing how far off you can be sometimes even when you think you're doing everything right.

The satisfaction of knowing that you've really got something great is just an amazing feeling. I really don't want to give something back and have them say, "What the heck did you do?" I just want them to listen to it and go, "Wow, it sounds better."

Do you think cutting vinyl helped you in the way you work now?

There's nothing like cutting vinyl because of the attention that you have to pay to dynamics, because it's so critical to whether you're actually going to have a successful cut or not. You train yourself to see the VU meters and the music in one continuum.

I think that it probably helped to focus me on how to concentrate on listening to music. Somebody today could say to me, "Did you like the way the song took off in the second bridge?" and I'd say, "I wasn't even listening to the structure of the song at that point. I was listening to the whole." There's a whole other thing that's going on.

There are guys that know how to make things sound really loud and big, but over-compression will keep the rhythm from working right. Once you take away the beat, then you just don't have the same intensity anymore. Maybe from cutting lacquer all those years, I started listening to drums a lot.

What's your signal path?

My console [designed by Chris Muth, former chief tech and now with Dangerous Audio] is set up with a patchbay so that any piece of analog gear can be placed in any position and in any order. That's a great luxury, and it gives me a little more creative control.

Generally, I keep everything pretty standard, and if there's an exception, I'll repatch. For instance, I have some producers who don't like the sound of certain pieces of gear, so I'll patch around those. Some producers are so fastidious that they want the signal path as short as possible, so I make a shorter signal path. To many of my clients, getting that sort of sonic purity is not as important as it once was. More often than not, it's more getting it somewhere that it didn't get rather than preserving what's there. That was always the case, but now far more projects are in that direction.

Are you doing all of your processing in the analog domain?

I do 90 percent of my processing in the analog domain. I have only one digital equalizer and a digital limiter for making loud CDs. I have three different analog compressors, including one I just got from Dangerous Music that complements my Pendulum OCL-2 very well. I have a Z Sys digital limiter that I've had forever that I use for shaping on the low end and corrective work when I have to do a revision of something that I might have already mastered.

Are you using any plugins?

I don't use many plugins. I'm in the process of evaluating the multi-band and maximizer Cube-Tec plugins that the other guys here use. The only time I use a plugin is when I get an instruction that the client wants something really loud.

What are you playing out of?

I'm playing out of Pro Tools and recording into Pyramix 8, where we do all the editing. We're clocking with the Antelope Atomic Clock and coming out of a D/A convertor to the amplifiers and speakers that one of the guys here built. I also use the AudioGate for a sample-rate convertor, which is fabulous.

I have two A/D convertors: an Ayre and a Burl B2. They're completely different, but I really like to have choices between warmth as opposed to clarity and combine the two of them in any direction. Almost all of my gear balances between those two elements. My equalizers are the same thing. I have a Prism, which is really clear, and a Buzz Audio, which is really warm, plus a Focusrite Blue, which is super precise.

What I do is approach a mix with what I think is the best combination of gear for it, then I immediately switch to the other gear to see whether there's something in the electronics that might bring out something musical that I missed.

Once I figure out what kind of EQ and compression I need, then it's a matter of figuring out which of these boxes are going to give me the best image. I used to listen to a song top to bottom and try to dial in an EQ, but now I try to get a quick impression and dial something in. I'll listen for it for like 30 seconds and get another quick impression. With a sound file, you can look at the peaks and find where the dynamics are going to be in the song and save a tremendous amount of time. If you have an analog tape, all that rewinding is a much different process and a lot more time-consuming.

What are you using for monitors?

I'm using the same ProAc Response 4's that I've had since 1993 with Pass Labs mono-block amplifiers. I have a couple of different interconnect cables. One of them is from Harmonic Technology, and the other is WireWorld. Again, it's like two different sounds. The WireWorld has a little more midrange and is a bit more meaty, while the Harmonic Technology one is very wide and a little splashy on the top end, and the bottom goes really low. That's the starting point.

The beginning of any project is listening to the mix through the different interconnects and convertors before even getting into processing. You do it for a couple of songs and see if there's a pattern to be had, although sometimes you have to do it for every song. It sounds time-consuming, but it really doesn't add that much to the overall mastering time. Clients really enjoy it when they're here. They're always astounded that there are such big differences in the wires. Once you get the hang of it, you really can see how you can add a 3D quality to the mixes independent of the EQ.

You're getting a lot more indie work and more from out of the country, I assume.

Yes. Almost everything is indie. The jobs that come from major labels are few and far between anymore. It's a byproduct of the fact that it's so difficult to make a living in this business, so the careers are a shorter and everybody with a band can scrape together some money to record an album without having a major label. It's all been fantastic for us, as it's increased business tremendously because this is the one entry point into the business that they can afford.

How much are you doing for vinyl?

Most people don't have the money for a separate master for vinyl today, but because I don't cut things massively loud, most of my stuff can be cut from the CD master. A separate master also opens up a creative can of worms because once you get away from the CD level approved by the client, it changes the relationship between the instruments.

I've attempted to do non-compressed files intended just for vinyl, but it requires another approval process, which no one has the time or money for after the project is over. I'm

not talking about an artist like Coldplay or a high-budget major label album, where that's not a problem.

I take into account during mastering if someone is also going to vinyl, but the real thing is to make the mixes substantially more pleasing than where the mixes started. That's the way it's always been from when I started in 1973. The goal is for the client to walk out of our place with something that was better than when they walked in. If they get a feeling that they've got something that's substantially better, then you've got a happy client.

Chapter 14

Dave Collins - Dave Collins

Mastering

A mainstay at Hollywood's A&M Mastering for many years, Dave Collins has brought his unique approach to a host of clients, such as Sting, Madonna, Bruce Springsteen, and Soundgarden. He now operates out of his own Dave Collins Mastering studios in Hollywood. While Dave has an extremely deep technical knowledge base, it's the creative part of mastering that lights him up, as you'll see in the following interview.

What is your philosophy on mastering?

The first philosophy is like the Hippocratic oath to "do no harm." The client is investing a tremendous amount of trust in the mastering engineer when he gives you the mix and expects it to sound better than it did when he brought it to you. I personally think experience is as valuable as equipment in a large sense, because after you've done it for 10 or 20 years, you've heard almost everything that can possibly go wrong and go right on a mix, so you can, in one respect, quickly address people's problems.

When I listen to a record I've never heard before, I don't know that the guitar player was fighting with the singer through the whole session or any politics that entered into the equation. I just listen to the sound that comes out of the speakers and take it from there.

One of the hardest things—and it took me forever to get this—is knowing when to *not* do anything and leave the mix alone. As I have gained more experience, I am more likely to not EQ a mix or to just do tiny, tiny amounts of equalization. I think some people feel like they really have to get in there and do something, and put their stamp on the mix somehow.

I don't really care about that. I only care that the client is happy and he comes back. I don't really feel that I need to put any particular personality on it. And hey, if the mix sounds good, let it sound good.

What distinguishes a great mastering engineer from someone who is just merely good or competent?

It's probably two things. I think the best mastering engineers understand a wide range of music. Believe me, I buy tons of music and listen to everything so I can stay current with what is going on, because I have got to get what the fans are hearing and understand that, so having aesthetics for a wide range of music is probably a fundamental skill.

Secondly, I would say that having a technical background, especially these days, certainly doesn't hurt, because both recording and mastering now are far more complicated than ever before. The palette of signal processing that you have today is enormous, both in analog and digital, and it is growing all the time.

How important is mono to you, and do you listen that way often?

One thing that happens after you've listened for a long time, I can tell by how phase-y it sounds to me in stereo if it's going to sum to mono. Once I get a certain amount of that crossed-eyed feeling, I can pretty much tell that it's not going to sum to mono, so yes, I always check for compatibility. I've certainly had mixers come in with stuff and I'd say, "Man, that is some wide stereo you've got going there. How does it sound in mono?"

And the guy goes, "I don't know. How does it sound in mono?" Of course, you put it in mono and now one of the guitars has disappeared, so it's an issue, but perhaps less important as time goes on.

Can you hear the final product in your head when you first run through a song?

No, not always, and in fact I frequently go down a dead end EQ or processing wise. There are some styles of music that I will intrinsically hear faster because the sonic presentation is pretty standardized in a lot of ways, but there are times when I can hear 90 percent of what it's ultimately going to sound like immediately when I put it up. There are other times when you go around in a big circle.

What is the hardest thing that you have to do? Is there one type of operation or music that is particularly difficult for you?

The hardest thing to do is a compilation album where you have 13 songs with 13 producers and 13 engineers and in some cases 10 different mix formats. Those are the hardest to try to get any consistency to it, just from a strictly sonic point of view.

Second to that is working on projects that have a "too many cooks and not enough chefs" condition, where you've got a lot of people kind of breathing down your neck and a lot of people with different, usually contradictory, opinions. Some of those projects—and usually they are your major-name artists—can be a little problematic because you have so much input and everyone is trying to pull you in a different direction at once, so that can be a little nerve-wracking, but it's all in a day's work.

What do you enjoy the most?

The day after the session, when the client calls and tells you everything sounds great and "I can't believe how good it sounds. I had no idea my mixes sounded that good."

That's the best, when I have someone who really got what I was doing and really got

what my room is able to produce. It's not every project, of course, but those are a good call to get.

How has mastering changed?

I think in a broad sense mastering hasn't changed at all and the concept of a specialist engineer and a room built for the purpose will never change. The tools change, and to some degree the client requirements change, but the general job description hasn't changed, and I don't think it will ever change.

Now from the business point of view, what's changed is the fact that there's an enormous amount of international business right now. I would say about a third of my business is from outside the U.S. With the Internet, it's now a completely electronic world. The downside to that is that you never meet the client. You communicate via email, he sends the files and I send the mastered tracks back, you get paid by PayPal or whatever, and you never even talk to anyone on the phone. At least in my world, that's one thing that's been a major change.

Do you find that a lot of indie projects bring a pro in at mastering to finish it off?

It's interesting that this is the one area where they bring in someone with a lot of experience. Sometimes that's misguided because a lot of people have been led to believe that mastering can fix everything. Of course this is nonsense. I can do a lot of things, and I can process the music in ways that were unavailable 10 years ago, but it's no substitute for taking care every step of the way of production.

For some reason mastering has been considered this black art, which is something that I don't like and don't subscribe to and try to explain to anyone who will listen, because I really don't like the idea of some mysterious "man behind the curtain" thing. I don't think that helps anyone.

It may sound cynical, but there are more people doing mastering than ever, and there are more people doing bad mastering than ever. The guys who have done it for a long time stay in business by keeping their customers happy for a long time. It's that simple.

A lot of engineers have added mastering to their repertoire, so you end up with strange results because a mixer starting to master is not hearing it like a mastering engineer. A mastering engineer hears in a very strange way, because I only listen to everything at once. I don't care what the snare drum sound is, or that there was a U 67 used on the guitar cabinet. I have no interest in that. The mixer has agonized over those details, which can be counterproductive to doing a good mastering job because they're hearing it in a different way.

I know of some very successful mix engineers that try to do their own mastering for

budgetary reasons, and they call you sometimes on the verge of tears saying, “I cannot figure this out. How come when you send me the master back it sounds like my mix, only better? When I do it, I’m making it worse.” Sometimes I’ll say, “How about you don’t do anything and turn it up a couple of dB and see how that sounds?” They’ll say, “You know, I think you’re on to something,” because they’re approaching it from a different aesthetic than I would.

What gear are you using?

I’m using mostly custom gear right now. I’m using the Davelizer, which is a low-Q broad-peaking EQ, a custom parametric EQ designed by Barry Porter [who was a seminal audio circuit designer from the UK], a custom VCA compressor, a Pendulum OCL-2 with some minor modifications, a Pendulum ES-8 compressor, and a TC Electronic System 6000. The A/D and D/A conversion is custom-designed along with the clock. I have a pair of K&H 0300D small monitors and a pair of Quested 2108 large monitors powered by Douglas Self Blameless amplifiers, which were built from a kit. Then, for workstations, I’m using WaveLab for playback and a Sonic Studio to record. My monitor controller is a Manley.

For tape I have a highly customized Ampex ATR-100 tape machine for 1/2- and 1/4-inch playback that has vacuum-tube electronics as well as something similar to Ampex electronics that I can use with it. I’ve also got a Korg MR-2000 DSD recorder, and I get some masters in on that format. It has an analog output only, so I treat it just like any other analog source.

Do you do all of your processing in the analog domain?

Not always. I probably do 80 percent in analog. There are some projects that sound so good there’s no benefit bringing them back to analog. While I pride myself in having an extremely transparent and euphonic signal path, some projects come in so “done” that there’s really no benefit to it.

I was a little late to the game with plugins. I played with them for years and never really thought they were better than what I was already doing with outboard gear until I found a company from the UK called DMG that I really like a lot. I think their EQ plug sounds great, and I prefer it to any of the other digital EQs that I’ve used.

When I need an all-digital signal path, I come out of WaveLab into the TC 6000 into the DMG equalizer, and then maybe into their Compassion compressor, which I really love.

I can take any analog compressor and make it do what I want in 8 seconds, and I can decide if it’s right for the song in 9 seconds, but I could never get a digital compressor to do the thing that analog just did automatically. The DMG Compassion actually works very much like analog does. I also use a peak limiter called the FabFilter Pro-L, and I have an iZotope package for click and noise removal. That said, I’m really at heart an

analog guy, so that's what I run that most of the time.

Do you ever have to add any effects?

Oh, sure. We used to do a lot of soundtrack mastering at A&M, and it was very common to add a touch of reverb at the final stage. Generally, you won't want to add reverb to a whole pop mix because it gets too washy. Sometimes you'll have to add a little reverb at the end just to give you something to fade over if the tail has been cut off. I generally try to caution people, don't trim it too tight because it's a lot easier to take it off than it is to put it back.

I've done things like overdubbed vocals in the mastering room before, and guitar solos, too. Live, right to the master. I remember the last time we were doing vocals, the guy was like, "So, what kind of cue mix are you gonna send me?" I said, "I'm gonna turn the level down low on these speakers, and you can listen to it and you're gonna sing. How's that?" It does happen, but fortunately not often.

Do you get a lot of projects that are already crushed?

Sure, but I'd have to say a lot less than five years ago. One thing I have noticed is that people are not operating their workstations at the same high levels that they used to. In fact, I get a lot of mixes where the peak level might be -6 to -10 or so, and that's a really great trend. I think a lot of people have realized that the gain structure in a workstation has to be paid attention to in the same way as analog. Running everything in the red all the time has negative side effects. Also, I'd have to say that the things that are crushed are crushed better than they were five years ago, because the stereo buss processing is better than it once was.

I can make it loud if that's what the client wants. In fact, anyone who's done mastering for any length of time has most likely come into the studio on a Saturday and monkeyed around with their system to find out what's the best way to get it loud. Personally, I think that the whole loudness thing might have peaked, and maybe we're on the downward side of it.

How much MFiT do you do?

I do everything at 96k, so I have a 96k version in the computer for every project. I'd say it's more common than not to do MFiT, and if it's a label project, then it's 100 percent. If people are doing vinyl, I will always make a 96/24 version with much less or no peak limiting, because the lacquer cutter is going to set the appropriate level when he cuts it.

Also, things that are flat-topped provide a very unsatisfying sound on vinyl, which is sad because good vinyl is normally extremely satisfying.

You have such an in-depth technical knowledge of audio and great ears, but it

seems like your clients always come first.

I'm not there to preach and wag fingers at the client; I'm there to make the record that the customer wants. I will enable that in any way they want. People come to me for my tastes, but it's their record.

Chapter 15

David Glasser - Airshow

Mastering

David Glasser is the founder and chief engineer of Airshow Mastering in Boulder, Colorado, and Takoma Park, Maryland. With two Grammy awards for his work, he has mastered thousands of records over the course of his 35 years in the business, including those for some 80 Grammy nominees. An expert in catalog restoration, David has worked on culturally significant releases by Smithsonian Folkways Recordings and the Grateful Dead, among many others.

How did you get started in the business?

I started in college radio; then I worked for a classical music station in Boston that had the contract along with WGBH to record and syndicate the Boston Symphony and Boston Pops. Syndication in those days consisted of high-speed reel-to-reel duplicates that were sent out in the mail to the various stations, which is what I did. After a while I started doing the recordings as well. From there I worked at NPR in Washington, DC for eight years as a staff engineer doing all the news stuff as well as location jazz and classical recordings and post-production.

How did you get from there to Boulder?

After 18 years in DC, it was time to make a change. We'd really outgrown the small mastering room that I had there and looked around for some alternatives, only to realize that if we did that, we'd be stuck in the DC area for at least another 10 years. We made a shortlist of cities to check out and got as far as Boulder and decided this was it.

Why didn't you go to any of the bigger media centers where there was more work?

I'm not really a big-city guy, and since our clientele had always been small boutique record labels and independent artists, we figured that we could almost be anywhere. That part worked out really well. It turns out that Boulder is a great music town, although we didn't know that at the time.

How did you get into mastering?

I was doing a lot of location recording and jumped on the Sony F1 [the first inexpensive digital recorder], and like a lot of people got blinded to the fact that it didn't sound that good but was so convenient. You didn't have to haul tape machines and a rack of Dolbys [noise-reduction units to quiet the tape hiss] to every gig, but you couldn't edit any of the tapes.

I got one of the first Sound Tools systems [the precursor to Pro Tools], and I was suddenly able to edit the things, and it grew from there, from doing simple editing to preparing CD masters. All of a sudden, I was in the mastering business. From there, I bought a used 1610 and just started acquiring all the tools for mastering around 1990 or so, which happened to be when a lot of the small boutique labels were just starting to repurpose their catalogs for CD.

What's your philosophy on mastering, and how did you come to it?

A lot of it was by studying records that other mastering engineers were doing. A lot of it came out of my NPR experience, where pretty much all the work we did was direct to two-track, so we were dealing with making stereo recordings sound good on air, mostly because we didn't have the budget for anything else. That made me develop the mastering mindset.

How do you approach a project?

I think anybody that's decent at this can hear what the producer's going for and where it works when you hear a mix. Your job then is to get it to where you think the producer was headed. After a while that becomes intuitive and almost a little automatic. You know what's going to work and what won't.

How have things changed in mastering?

Most budgets have gotten way smaller, even for independent artists who didn't have much of a budget in the really good times. More and more people are doing EPs and singles today.

Are you doing separate masters for online?

Usually not. It's easy enough to check to see what's going to work. Most people don't have the budget for two versions, but often I'll do a CD version and then a less limited version for the downloads.

At NPR, one of our mentors always said that if you can make it sound good in your studio, then it's going to sound good on a TV, AM radio, or anything else. I think that's true, so I don't really see a reason to do separate versions except for overall level and limiting.

Are you asked to make loud masters?

It depends. Some people want it to sound good but don't want it to get lost. Other people will compare it to other records, including ones that aren't very appropriate for comparison, and ask for a louder record. I've found out that my version of really loud is not like other people's idea of loud. I pull my hair out trying to make it sound really good and really loud. That's a challenge.

What are you using for a workstation?

soundBlade. I also have Pro Tools and a Sonoma workstation for DSD. I use Pro Tools for surround and video stuff because it just works so well for that. I also use it for capturing at a different sample rate than the source files are at.

Do you have any favorite plugins?

You know, I don't use that many plugins. I've got the Fraunhofer Pro-Codec, which is really more of a useful tool than a plugin. All the processing that I do is with outboard gear.

On the digital side, I've got a Z Systems EQ and compressor and a Weiss EQ and compressor, a TC Electronic TC 6000, and a Waves L2 limiter. In analog, I have some of the new Pultecs, a Prism EQ, an API EQ, a Fairman compressor, an SSL compressor, and a Maselec console.

For convertors, we have a choice between Pacific Microsonics or Prisms. For monitors, I'm using Dunlavy SC-Vs for left and right and SC-IVs for the center and surrounds, with Paradigm subs and Ayre amplifiers.

We've also got Ampex ATR and Studer 820 tape machines with headblocks for pretty much anything, if we get a project in on tape.

What's your signal chain like?

I usually go out into the analog domain. With good enough convertors it's not totally transparent, but it's a compromise worth making. Usually it's analog EQ, analog compressor, A/D convertor, and then maybe some digital EQ and a final limiter. It's pretty basic and standard.

Is your L2 always at the end of the sign chain?

When I use it, yeah, before dithering. Occasionally I'll use the limiter in the TC 6000. If it's something that doesn't need a lot of limiting, then I'll stick with the Maselec analog limiter. I've just started using the Sonnox Oxford Limiter, and I'm pretty impressed. Very occasionally it will be the limiter in the HCD model 2. I've been starting to use one of the LUFS loudness meters, and I'm going hotter than default Apple setting [-16 LUFS] but only by about 3dB LUFS, so it's around -13 LUFS. Some of the louder mastering jobs done by other engineers come out at about -9 LUFS.

How about your room?

The room's great and was designed by Sam Berkow. It's about 350 square feet, with lots of bass trapping and large diffusors in the back.

I see you do a lot of restoration. How did that come about?

That started when I was back in DC. One of our clients was Smithsonian Folkways, who had tons of great old recordings. We bought one of the first Sonic Solutions NoNOISE

systems, and we've been doing that sort of stuff ever since.

I love doing oddball things like that. We've always jumped on esoteric things like SACD. We don't do SACDs much anymore, but we do get DSD mixes in, so we have the tools for that. We've also invested in the Plangent Processes replay electronics for one of our tape machines. The process removes wow and flutter from analog tapes by extracting the bias signal that's still on the tape and using it as a reference for the software to eliminate all the speed problems. We used it on all the Grateful Dead studio records, and the results are amazing.

Do you do much MFiT?

Not a lot. Not that many people ask for it specifically. It's a big extra expense for indie artists, because they have to go through CD Baby, which charges extra for that. Most of the labels usually ask for high-res masters, though.

Are you delivering the project back to the client or directly to the distributor or replicator?

We like to send it directly to the factory. When someone asks for the CD master direct, we explain to them that it may or may not be a good idea, and they usually agree. We charge a premium for a master, and that includes a QC check by my assistant. It's time-consuming, but it's necessary. We send mostly DDPs these days. Now QC includes spell-checking of the metadata as well.

Is your mastering approach different when you're doing catalog?

Maybe slightly. If it's a well-known album, then you can't stray far from what everybody already knows and loves, so you have to check back with the original a lot. If it's an archival thing that wasn't widely known, then probably not. The goal then is to just make it sound as good as you can.

Chapter 16

Gene Grimaldi - Oasis

Mastering

Gene Grimaldi started his career in mastering at the CBS Records/Sony Music CD manufacturing facility in Pitman, New Jersey, in 1986, learning about the business from a much different perspective than most mastering engineers. Wanting to know more about how masters were created, Gene eventually made his way to Future Disc Systems in Hollywood, where he worked as a production engineer, soaking up all the aspects of mastering and cutting vinyl. When Oasis founder Eddy Schreyer opened up the new studio in 1996, he invited Gene on as a full-time staff engineer. Today, Gene is Oasis's chief engineer, with a list of blockbuster clients that include Lady Gaga, Jennifer Lopez, Ellie Goulding, Carly Rae Jepsen, Lana Del Ray, Nicki Minaj, and many more.

How has mastering changed for you in the last few years?

The process is basically the same. You get your files or tapes in, and you still have to balance out the mix and make it sound good. As far as the way the industry has changed, at times it seems a lot more amateur than it used to be because so much of the mixing is done in home studios. It's harder to get it to sound good, and you're dealing with the client more to help them get their mix in a better place before I can work on it. They appreciate the help, too.

Are you getting more unattended sessions?

I have more unattended than attended, but a lot of that has to do with the number of foreign projects that come in from Asia and Europe.

Can you describe your signal path?

On the analog side, the source gets patched into a custom Manley console, then into my Avalon 2055 and GML 8200 EQs, Tube-Tech SMC 2A multiband compressor, and a Manley Vari-Mu. I can interchange anything by physically moving the patches on my patchbay, but I usually go through the EQs first and then the compressors. From there it's into the Lavry Blue A/D into the Lynx interface board and then WaveLab 7.

The output is through a Lavry convertor into a set of Hot House amps, to a pair of Tannoy System 215 DMT with an external Tannoy supertweeter. I also use Tannoy System 600s for the a small speaker reference.

The Tannoys are a little light from, say, 30Hz down, so to fill in the extreme lows more accurately, we use a pair of dual 15-inch subwoofers from Aria. You're pretty much

locked down positionally if you use just one, but with the two you actually have a little more flexibility.

In the digital domain, I use a handful of plugins from Universal Audio, like the Cambridge EQ and UAD multiband compressor, and AudioCube EQ and De-esser VPI. For a limiter where I have to fix something and cut it in, I'll use the Limiter VPI in the AudioCube. If I need to use a limiter all the way through a song, I'll use the Sonnox Oxford limiter.

You don't use a limiter so much anymore, do you?

I have always tried not to use a limiter if I'm not hearing any distortion. When I run into situations when I need to use a limiter, I'll just use it on the portions of the song that need it. Perhaps one day we'll get back to normal levels.

What I do is go in and slice up a song and just smooth out only the rough edges. If it's an open track that has to be loud, I'll just cut all the tiny pieces that need limiting and limit only those. Those sections go by so fast that your ear can't hear the slight audio differences between the fixed limited sections as they fly by. It gets rid of any overload crackles and keeps the kick hitting hard. It's time-consuming, but I don't mind doing it if it comes out better. It actually goes a lot faster than you think once you have an ear for what to listen for.

I notice that you use a lot of processing plugins, but just a very little of each.

I just tickle them, because I look at it like it's all cumulative, especially when adding EQ. I get most of the impact from when I set up the gain structure, because once I get the loudness to where I want it, the mix starts getting into that window where I know what balance adjustments to make. If it's way off from the get-go, I have to get in there right from the beginning and start balancing the bottom or the top end right away, and then I will increase the gain.

What's your typical plugin signal chain?

I would probably go multiband compressor, EQ, limiter, and sometimes I put the de-esser at the very end, because the limiter can add a brightness to it that the de-esser can catch. If I'm not using the limiter, then I'll put another compressor after the EQs. That said, there are some times where I'll put the de-esser at the front of the chain; it all depends on the song and how much it all needs.

I flip-flop between the different EQs and compressors and drive them all just a little to increase the gain. Sometimes one will color more than another, but that's how I give the client different choices. I'll take the gain away from one and add it to another, and it will sound different—either more transparent or smoother.

Are most of the songs that you get crushed level-wise?

It's really all over the place. If it is crushed, I try to talk to the mixer and ask him to take any mastering plugins off before he sends it back to me. You can put your compressors and EQs on the two buss to get your sound, but give me some headroom to work with.

Some of my clients really insist on having the maximum level possible, though. In that case I give them multiple choices. I give them one that I think is loud but sounds good, another that's pushed a little more, and one where you're getting to where you don't want to go any further. It helps them see the light. When you compare them in real time, you can really hear the difference. I do have a limit to how much I push it, though. My name's on it in the end, so it has to sound good. It can be loud yet still have some dynamic range and sound good.

That's kind of what you get when you don't use a limiter. As soon as you put that limiter in, a lot of it starts sounding really soft and smoothes out. I'd rather have it hit you. It does depend on how you set the limiter and how hard you hit it, but it does soften it up.

Then again, the artists trying to do it at home don't have the advantages of hearing it like I do. The studio here is really like a giant microphone, and you can hear every little thing. There's a big difference between an artist's or a producer's home listening environment and listening in our designed and tuned rooms.

How long does the average mastering job take per song?

About half an hour. If I feel like I'm not getting it, then maybe 45 minutes at the most. Once that's approved, then I'll just drop in any alternate versions, which are usually the TV mix for a single and maybe a 48k version for any video that goes along with it.

How much MFiT or high-res do you do?

The major labels ask for MFiT with every mastering job. The indies aren't quite there yet. If a project comes in at a high-res rate, I'll do it that way for them if they want it. For vinyl, I try to deliver a 96k/24-bit file and drop the level much lower from the CD version.

Do you do any different processing for MFiT?

I try not to deviate too much from my original CD master. The level will get dropped a little, but that's about it. It's pretty much the same as MP3. I drop it through Apple's MFiT droplets to listen to the way it's going to sound after it's encoded, just to be sure it'll sound good.

What kind of masters are you delivering?

We hardly make CD-Rs anymore. It's all DDP that's FTP'd directly to the client, although we do WAM!NET delivery for certain record labels, like Universal. It's a lot

easier than it was, although we used to get paid for all the extra work we had to do. Those old 1630s worked, but they were a mess, so I'm glad we've moved on since then.

Chapter 17

Bernie Grundman - Grundman Mastering

One of the most widely respected names in the recording industry, Bernie Grundman has mastered hundreds of platinum and gold albums, including some of the most successful landmark recordings of all time, such as Michael Jackson's *Thriller*, Steely Dan's *Aja*, and Carole King's *Tapestry*. A mainstay at A&M records for 15 years before starting his own Grundman Mastering in 1984, Bernie is certainly one of the most celebrated mastering engineers of our time.

Do you have a philosophy on mastering?

I think that mastering is a way of maximizing music to make it more effective for the listener, as well as maybe maximizing it in a competitive way for the industry. It's the final creative step and the last chance to do any modifications that might take the song to the next level.

There are a couple of factors that come into play when we're trying to determine how to master a recording. Most people need a mastering engineer to bring a certain amount of objectivity to their mix, plus a certain amount of experience. If you [the mastering engineer] have been in the business a while, you've listened to a lot of material, and you've probably heard what really great recordings of any type of music sound like, so in your mind you immediately compare what you're hearing to the best ones you've ever heard. If what you're hearing doesn't meet that ideal, you try to manipulate the sound in such a way as to make it as exciting and effective a musical experience as you've ever had with that kind of music.

You have to interface with the producer or the artist, too, because they might have a vision that may be slightly different than where you intuitively want to take it. They might want to emphasize some aspect of the music that you may not have noticed. A lot of it is definitely trial and error on your part, but it's also give and take between the producer and the artist because you can't sit there and arrogantly think that you know where this recording ought to go and that they don't.

Can you hear the final product in your head when you first run something down?

Well, you do get ideas. If you've been in it a while and you've heard a lot of things, then you know where to go. Like if you put on a hip-hop record, you know that it's very rhythm-oriented and it has to be really snappy and punchy on the bottom end. You know that some of the elements are really important and that this kind of music seems to feel better if it has them.

Maybe the client had a monitoring system that had a lot of bottom end and the mix comes out bottom-light as a result. That's why probably the single most important piece of equipment that a mastering engineer can have is his monitors. If you know the monitors and you've lived with them for a long time, then you're probably going to be able to make good decisions. The only problem with that is, if the monitor is something that is a little bit esoteric and only you understand it, the producer or artist can become very insecure with the result. That happened to me when I first worked at A&M and I had a monitor system where I knew what it should sound like, but it was really kind of wrong for everyone else. They had to trust me, and they did, but I could see them get really concerned about what they were hearing, so in my studio I've gone to great lengths to make it a very neutral system that everyone can relate to.

What monitors are you using?

We put them together ourselves. We build our own boxes and crossovers using all Tannoy components. It's not that we're going for the biggest or the most powerful sound, we're going for neutral because we really want to hear how one tune compares to the other in an album. We want to hear what we're doing when we add just a half dB at 5k or 10k. A lot of speakers nowadays have a lot of coloration and they're kind of fun to listen to, but it's hard to hear those subtle differences.

We just use a two-way speaker system with just one woofer and one tweeter so it really puts us somewhere between nearfields and big soffited monitors.

Do you use nearfields as well?

We have some NS10s and some little RadioShack cubes that a lot of people around town like to hear what it's going to sound like on. Usually if you can get it sounding good on our main system, it's just that much better on the other ones.

Do you still cut lacquer?

Oh yes, we sure do. We have one room with two lathes where we cut all of our lacquers now that's going all day long. We can't do it fast enough. I don't know if this is going to last, because it's gotten into this area where people think that vinyl is "happening," but the expense that you have to go through to make a vinyl album and the cost of manufacturing is way more than CDs. I don't know how many clients actually make their money back.

I have a love/hate relationship with vinyl because there are so many things that can go wrong and there are so many limitations. It can sound incredible if everything is right and you're careful not to exceed any of those limitations. The problem is that it's analog. Any little thing that goes wrong, you're going to hear it.

Are you doing a separate master for vinyl?

Ideally, but not necessarily. Some of the clients want us to cut from the CD file, but you're using a signal that's been modified to be very aggressive. That's right where you're going to start having trouble with vinyl, because those grooves get radical when they have that much energy in them.

How do you think that having experience cutting vinyl has helped you in the CD age?

It takes a lot more knowledge to cut a good vinyl disc than it does to do a CD. With CDs, except for artifacts and various changes that occur in the digital domain, what you get on the monitors is very close to what you get on the disc, and you don't have all the various distortions that vinyl can come up with. Vinyl has inner groove distortion and tracking distortion because of too much energy in the high frequencies, but this doesn't happen on CDs or digital files. With CDs, of course, the quality is the same from the beginning to the end of the disc, which isn't the case with vinyl. High frequencies might get a little brittle, but they don't distort on a CD, whereas they will on vinyl, so there's this whole grab-bag of problems with vinyl that you have to consider. Part of being a good vinyl cutter is knowing how to compromise the least.

Can you talk about the level wars for a minute?

That's one of the unfortunate things about the industry, and it was that way even way back in the days of vinyl. Everybody was always trying to cut the loudest disc, and then if you got into a new generation of playback cartridges that could track cleaner, they would push it again until those were on the edge of distortion. It didn't matter if you had better cartridges because that just meant that you could go that much louder and get right up to the same amount of distortion you were at before.

Usually anything that sounds louder gets at least some attention. It might not hold up on the long haul, but the main thing that a lot of promotion guys want is to at least attract attention so that the song gets a chance. What happens is everybody is right at that ceiling where the level is as high as it can go, so now guys without a lot of experience try to make things loud and the stuff starts to sound awful. It's smashed and smeared and distorted and pumping. You can hear some pretty bad projects out there.

I try to give the client a number of options. I give them one at full level, with a small amount of clipping because of the compression and limiting, then one a couple of dB down. They usually like that one a lot better, but they'll also take the loud one anyway. As idealistic as we all would like to be, it's the last chance anyone has before it goes out to the public and will be put up against everything else. When a client realizes that, it's really hard for them to be a purist and reduce the level and make it more dynamic and natural.

As much as I hate to say it, I always tell them that you don't want to go too much lower, because you still have to be close, or the public will think there's something wrong with the record. People tend to gravitate to loud because there's a certain excitement in that.

Would you have any words of advice for somebody that's trying to master something themselves to keep them out of trouble?

I don't think that you should do anything that draws attention to itself. Like if you're going to use a compressor or limiter on the bus, you have to realize that you're going to degrade the sound, because compressors and limiters will do that. It's just another process that you're going through no matter if it is in the digital domain or analog.

Analog and digital are very, very much alike when it comes to signal processing. If you put an equalizer in the circuit, even if it's all in the digital domain, you will hear a difference. If you put a compressor in the circuit, not even compressing, you will hear a difference and it will sound worse.

What is the hardest thing that you have to do?

One of the things that is really hard is when the recording isn't uniform. In other words, a whole bunch of elements are dull and then just a couple of elements are bright. That's the hardest thing to EQ because sometimes you'll have just one element, like a hi-hat, that's nice and bright and crisp and clean, and everything else is muffled. That's a terrible situation because it's very hard to do anything with the rest of the recording without affecting the hi-hat. You find yourself dipping and boosting and trying to simulate air and openness and clarity and all the things that high end can give you, so you have to start modifying the bottom a lot. You do the best you can in that situation, but it's usually a pretty big compromise.

If the client just had a bright monitor system and everything in the mix was just a little bit dull, that's easy. It's almost like a tone control because you bring the high end up and everything comes up, but when you have inconsistencies in the mix like that, it's tough.

Then there's something that's been overly processed digitally, where it gets so hard and brittle that you can't do much with it because once you've lost the quality, you can't get it back. If I am starting out with something that is really slammed and distorted and grainy and smeary, I can maybe make it a little better, but the fact that a lot of that quality is already gone is going to handicap that recording. It's never going to be as present as the way something that is really clean can be.

That's part of what gives you presence—when it's clean. The cleaner it is, the more it almost sounds like it is in front of the speakers because it's got good transients, where if it has poor transients, it just stays in the speakers and sounds like it's just coming out of those little holes. It doesn't ever fill up the space between the speakers.

What makes a great mastering engineer as opposed to someone who is just competent?

I think it would be trying to get a certain kind of intimacy with the music. It doesn't even have to be music that you like. The real test is if you can stop yourself from having all kinds of preconceived ideas and just open yourself up to see how the song is affecting you emotionally and try to enhance that. I think that a lot of it is this willingness to enter into another person's world, and get to know it and actually help that person express what he is trying to express, only better.

How long do you think it takes to get to that point?

I think it varies. It depends on the emotional issues that people have, their personal defenses and their sense of self-esteem. Some people have such low self-esteem that it's really hard for them to even admit that there's a better way to do something. If a client suggests something, they're very defensive because they feel that they have to have the answers. A lot of engineers are that way, but mastering is more than just knowing how to manipulate the sound to get it to where somebody wants it to go.

What's your signal flow?

Most of what we do is converted to analog for processing, so we don't much care what format we're sent. Not all of it is done in the analog domain, though, especially if you're trying to get competitive level for something like EDM, which is slammed so hard that it's shocking. That's the thing with mastering; you just have to take what comes. If the client wants it, we're here to help him realize his dream. I'm there to show him different ways that I see it being better, but they have to make the final decision. It's their project.

What are you using for a workstation?

An AudioCube, which we chose because it was the best-sounding. We use Pro Tools 10 for playback because it's such a standard. We actually get full sessions in that are put together with the song spacing they want. In our surround room, we have an AudioCube that goes up to eight channels.

We used to get stems in, but not so much anymore. Even then, I only used them to vary the vocal level, because that's where mixes tend to err. One of the most common problems I find is that the vocal is too buried. That's because the client gets so used to the song and the lyrics that they think they're hearing it when they're really not hearing it well at all. They've been working on getting the maximum support from the track, but you need to do that just to the point where it starts distracting you from the central figure, or the vocal. You want to get maximum support but never distract your listener from the vocal. I hear the vocal, but it's not carved out enough, it's not detailed enough, it's not out front enough. What are we listening to, a track accompanied by the vocal or a vocal accompanied by a track?

Do you use any plugins?

I use some of the AudioCube VPIs, and I'm getting pretty good results with the PSP maximizer. I only use them on certain types of programs, because most of my stuff is outboard. I use a really good compressor/limiter that we built here that uses some special parts that you can't even buy anymore. You almost can't hear it in the circuit, but the problem is that analog just doesn't have a fast enough attack time. Maybe on a signal that's not that complex, but you try to do some rock thing that's really dense, and you don't get good results from it. I don't think you can get great results if you only use analog compressors and limiters to get the kind of level that people want. You have to use digital.

What console are you using?

One that we built. We build most of our own equipment mostly to avoid a lot of extra electronics and isolation devices and so forth. When you buy most pieces of audio equipment, each one has its own isolation transformer or electronically balanced outputs, or however they arrive at a balanced output. When we buy outboard equipment, we completely rebuild it and put all of our own line amps in and take out the transformers or the active transformers. You'd be amazed at how much better they sound as a result. We also use Lavry convertors that are completely hot-rodded with our line amps and power supplies.

We have all separate power to each one of our rooms and a very elaborate grounding setup, and we've proven to ourselves that it helps time and time again. We have all custom wire in the console. We build our own power supplies, as well as the equalizers and compressors and everything else.

There are so many tools now for manipulating the sound that there are so many more possibilities to help you get whatever it is you're looking for. With that number of possibilities, you need some kind of vision, or else you can easily go way out on a limb and be really wrong. That said, you can really do some minute changes on something and get the most out of it, which is the goal. Again, it's the same old thing—it comes down to experience. It's one of those things where it has to come together in your mind first.

Chapter 18

Colin Leonard - SING

Mastering

It takes a lot to convert hit-maker mixers like Phil Tan and Dave Pensado into big fans, especially if you're not in one of the media country's media centers, but Atlanta-based Colin Leonard and his proprietary mastering technology has managed to do just that. With credits that include Justin Bieber, Jay-Z, Echosmith, Leona Lewis, Al Di Meola, John Legend and many more, Colin is proving that there's a new way to look at mastering that's equally effective as the traditional techniques. Along with his custom mastering service at SING Mastering, Colin is also the creator of Aria automated online mastering, the latest trend in convenient and inexpensive mastering.

How did you learn mastering?

The way I learned the most was by pulling up mixes and learning it by ear. I would spend every day doing that for a long while. That's the most important thing because even when you learn other people's techniques you still have to find out what is going to work for you. I don't think that I use many of the techniques that I learned from other people that much.

How long did it take you until you felt that you were good at it?

Every day [laughs]. I feel that it's a lot like playing an instrument. You have hills and plateaus. I always felt the same way playing guitar. You reach a place where you're feeling really confident, then you reach another point where you're working on improving. Mastering is the same thing. There are always those days where you feel like you're doing amazing work and then there are other days where it just doesn't seem like it's happening, so it's a constant work in progress.

The sonic fashion of music constantly changes as well. Masters from 1999 don't sound like the masters from today, and working like that wouldn't fly with artists or labels. It's a constant learning process.

Is there a certain type of music that you find easier or more difficult to work on?

No, I don't think so. I'm lucky in that I get a lot of different genres, but I don't necessarily find certain ones harder than others. The club aspect of some pop or rap music can make it a little more difficult because the DJs have a different loudness perspective, so that can be a bit challenging because they really want everything really loud. Making it loud yet still having the bass stay loud and clean can be a challenge.

I'm told by some of your clients that you manage to get things a lot louder than anyone else. How do you do that?

I have some proprietary processes that I use. I don't really use plugins. Almost everything that I use is analog. I'm always trying to come up with creative ways to get things louder while trying to keep as much apparent dynamics as possible in the masters.

How did your proprietary processes come about?

Out of frustration, really. I would be mastering some club record that needed to be on 11 and I would get it on the verge of distorting yet the client would say they needed it even louder. I was pulling my hair out. I try to look at the loudness as something creative, like "How can we get around this without distorting?" All the good mastering engineers face the same hurdles, I think.

Isn't it a lot harder to get the same hot level in the analog domain that you can get with digital?

Yes, and no. I think that for me with digital limiters you basically have a glass of water that's already full once it's hitting zero, and then you're boosting the level up and then it's just overflowing. With analog, depending on how things are calibrated and how much headroom your equipment has, you can operate it in such a way where the analog signal is much louder and clearer. Getting it back to the digital domain is where everything falls apart [laughs]. That's where you have to work on creative ways to get the same level without causing problems and for me that's the key to making it happen. I prefer analog to digital because I find it a little more pleasing to my ears. There are some guys doing all digital that sound pretty good, but I prefer analog.

What's your signal path like?

It's kind of a classic setup. I have a playback computer based around the Cube-tec platform that just plays back the files. I don't really do any processing in it. Then I convert to analog, and then I have a Dangerous Music Mastering mastering console that tweaked a little bit, but the stock one is fantastic. From there it's mostly EQ and not a lot of compression. I only have a couple of analog compressors that I don't use that much. I have 6 analog EQs that are all good at different things. I have some newer high voltage EQs and then some old Neumann cutting EQs that sound really good in the midrange. I have some Manley stuff with transformers and an SPL PQ that I really like. I then go back into a different Cub-tec to get back into the digital domain.

I find it interesting that you don't use compression that much.

I don't get using analog compression for mastering. Sometimes it can be nice for softer acoustic pieces or maybe rock stuff, but I think it causes more damage than it helps to songs with transient drums and extended bass notes. I'm more about transient energy. I concentrate on not screwing up the really fast high-frequency transients since a lot of the energy from the song is there. A lot of times I'll have a compressor or two in line but it

won't be doing anything. It's just there for the tube sound. It's more like using a compressor as an EQ.

Isn't the stuff that you get in that's pretty crushed already?

Yeah, totally. Why would we need to crush it any more? At that point you're just losing more dynamic range and that's not what we need.

I do get mixes that are all over the map level-wise though. Sometimes engineers send me stuff that's really open, while others send me stuff that's really compressed. I have to learn each different mix engineer's style, then after a while I know what I'm going to get from them, but it's always part of a learning curve. Even within pop music, there might be one great mixer that sends me really compressed stuff, and another great mixer that sends me really dynamic stuff.

I think a lot of it has to do with approvals on mixes. Most mixers don't want to send something in for approval on big projects and not have it be at least close to the level of other commercial releases. Some of the guys that I know that mix quiet and dynamic get a lot more complaints about level. The clients think the energy's not there so they try another mixer. If the other guy makes it louder then they'll lose the client. At the end of the day you have to do what the client wants.

What's odd about competitive levels is if you go to AES and listen to the speaker demos, they'll always use really dynamic material and people will always comment on how great it sounds.

Yeah, I tend to do that to show off my playback system if someone wants to hear it. I'll play something that's less compressed and quieter and I'll just turn it up louder.

There's also something that's going on in playback systems that makes it advantageous to have a louder brighter final product and that's the full range (singer driver) speaker, which is what most people listen to these days. The speakers on your phone are a single full range speaker, the same as headphones and earbuds, and the same as most computer speakers. Since you don't have a dedicated tweeter, it's just a darker sounding speaker, which has caused productions to get brighter and louder to get as much volume out of those speakers as possible.

What do you supply for vinyl master?

I used to have a lathe and would cut the lacquers, but I don't right now. It's a fun thing to do but it's also really time consuming. You'll go through days where you'll have a lot of problems with production. Maybe the lacquers or the plates got damaged and then you'll have to spend another day cutting new lacquers.

I'll just create 24 or 32 bit files that are a little quieter for the vinyl premaster, and then just lay

out each side as a long wave file. If there are 4 songs on one side, it will be all 4 songs connected as one big wave file. I'll also send the CD PQ information along so the cutter knows where the breaks are between the songs.

How do you view spreads these days? It used to be that spreads between songs were really important and you'd time everything out so it felt right between songs. We're in a singles world, so I guess people don't care as much about that now.

I still do. I do a lot of albums and EPs and I spend a lot of time on that. It's kind of an art on its own. Clients like to sit here and do it. For me it's still important, but yeah, if you just have a bunch of singles it doesn't really matter.

I do it in such a way so that the iTunes files use my spacing, so if you buy an album you will get the same spacing that I created. If there's space between the CD markers, I'll use a setting that keeps the that same space on the final file export so that it will carry over into the individual files.

Is there a typical problem that you see with mixes that you get?

I can't name one thing that happens all the time. Most of the guys I work with are really good so I don't have a lot of complaints.

When you think about it, who's to say what a problem is? I try to view it from the aspect of the mixing engineer and the artist making a piece of art. It's their vision, so what can I bring to it to make it better? If there's something that's really obvious that wasn't supposed to be there, maybe I'll ask in a real cautious way, but I try not to get too technical about it.

What monitors are you using?

I use Duntech Sovereign powered by Cello amplifiers.

What's your take on high end cables?

Oh, man [we both laugh]! I think they make a difference for sure. How much of a difference for the amount of investment I think you need to decide for yourself.

I actually see some of the biggest improvements with cables coming from power cables. It effects certain pieces of gear more than others, and it also depends what the power is like at your facility. I was at another facility where some filtering cables made more of a difference than they do here because the power was worse. My power here is transformer balanced so it's really clean and I don't get as much of an advantage as I was getting. There are so many variables and I think it's more of a trial and error kind of thing. It's not about getting all solid silver cables or your masters won't sound good though, but I think that good quality cables matter.

How did Aria come about?

I guess some of it was being loaded up with a lot of work and sometimes being a little overwhelmed about how fast people needed things turned around. I'm sure that a lot of busy mastering engineers feel that same frustration sometimes. The deadlines might not be realistic or even real, but your client sets those deadlines and you have to follow them. A lot of times, if you don't get those projects done on time, you just don't get those projects anymore.

Another part of it was that from mastering so many projects I have a basic setup that I start with, and figured I could automate some of it. It's a unique idea in that we're doing it in the analog domain. This was before I even heard of any of the other online processes.

The tools that I use in the Aria system are all the tools that I use in my normal mastering setup. It comes in digital, we do some scans and snapshots of the audio in the digital domain, then it uses our custom software to play back from one computer. It then goes through a real high-end D/A convertor and through the whole analog chain, which has automation elements built in, and then it gets recorded again into the digital domain after a real high-end D/A convertor. From there we can control the exports into the file format that we want. The whole system is hosted in-house, so the servers are here, which makes it really fast. The file goes directly from you to our server here, so it's immediate. As soon as it's mastered it hits your account for download.

How long did it take you to develop the idea?

I had the idea in 2013. I financed the whole thing from my mastering work, so I've been juggling these two things this entire time. It probably could have gone a little faster if I hadn't done that, but I wanted to maintain control of everything so we wouldn't have to do any silly marketing or anything.

Philosophically, aren't mastering engineers opposed to automated mastering?

Yeah, as I've seen online most of them are, but I think Aria is a different product than what I offer at SING Mastering. For one thing, budgets have really dropped. There are some clients that are willing to pay for high-end professional mastering, but there are also clients that maybe don't have the money to spend on their project, or maybe they need it in like an hour, so I think it fills a void for some clients.

Isn't there some pushback from mixers as well, mostly because they don't believe the final product will be as good as they expect?

I haven't had a lot of that from mixing engineers. Most of my beta testers were really good mixing engineers and they are still using it every single day. If they didn't like it they wouldn't use it.

Part of the idea is that it creates reference level material. Most engineers need to get a level that's competitive with what's commercially available to get approval from the client. I think it fills a void for mixing engineers in that they don't have to focus on the loudness of the track. It's almost like they're handing in a mastered reference.

How loud is it when it comes out the other side?

It depends on the mix. There are 5 settings that go in order of compression, so it depends on what you're going for. A and D and E are very safe settings and they're not pushed hard; B and C have a lot more analog push to them, which can be great on certain types of mixes. Once you play with it a few times you get the hang of it. Also, once you pick a certain setting you're not married to that setting. You can remaster it at a higher or lower level at no extra charge.

Explain the multiple levels for buying the product.

You can get it on a per song basis or a subscription basis. The subscriptions are on a monthly basis. For instance, you can buy a subscription for as many as 100 tracks a month if you have a big album to do with a bunch of different versions. That way you'll be able to do it a lot cheaper than what the single rate is.

Where do you think mastering is headed now that we're in this realm of automation?

There will always be room for "personal mastering," which is what I do at SING Mastering, but I think it's nice to be able to give another good sounding option that's more convenient and less money. We live in an Amazon Prime world where people would rather get something done at home at 3AM rather than waiting around to schedule a mastering session, so there's a big convenience factor involved as well. That said, there's always going to be a market for high level custom mastering.

Chapter 19

Bob Ludwig - Gateway

Mastering

After having worked on literally hundreds of platinum and gold records and mastered projects that have been nominated for scores of Grammys, Bob Ludwig certainly stands among the giants in the mastering business. After leaving New York City to open his own Gateway Mastering in Portland, Maine, in 1993, Bob has proved that you can still be in the center of the media business without being in a media center.

What do you think is the difference between someone who's just merely competent and someone who's really great as a mastering engineer?

I always say that the secret of being a great mastering engineer is being able to hear a raw tape and then in your mind hear what it could sound like, and then knowing what knobs to move to make it sound that way.

You know where you're going right from the beginning then, right?

Pretty much. It's a little bit like the Bob Clearmountain school, where after 45 minutes of mixing he's practically there and then spends most of the rest of the day just fine-tuning that last 10 percent. I think I can get 90 percent of the way there sometimes in a couple of minutes, and just keep hanging with it and keep fine-tuning it from there. It comes very, very fast to me when I hear something, and I immediately can tell what I think it should sound like. The frustration is, sometimes you get what I call a "pristine piece of crud," because it's a bad mix and anything you do to it will make it worse in some other way. Ninety-nine percent of the time, I hear something and I can figure out what it needs, and fortunately I know what all my gear does well enough to make it happen.

How many of your sessions are attended?

When I started my own business after working at Masterdisk and at Sterling Sound before that, our business plan called for a 20 percent reduction in overall business, but the opposite actually happened. We thought that half the people that had attended sessions in New York would attend up here. It turns out more people attend sessions here than in New York, which was a total surprise.

Why do you think that is?

I'm not sure. To tell you the truth, I think a lot of people have heard about the effort we've gone through to make our room as acoustically perfect as possible. They know that we've got speakers that retail for \$100,000 a pair, so a lot of people just want to come and see what it's about. Many times people come into the room and go, "Oh my God!" or something like that. It's

a trip to get that kind of reaction from people.

I felt that if I stayed in New York, I'd never be able to have a room that was acoustically as perfect as we knew how to make it. Sterling and Masterdisk were always in high-rises, so you're always limited to very low-ceiling rooms. In order to get as near-perfect a situation as possible, you actually need a fairly large shell that's at least 30 feet long and accommodates a 17- or 18-foot ceiling.

Do you think that there's a difference between the ways people master from coast to coast?

I don't think there's so much a difference between coast to coast as there is between some of the major personalities in mastering. Some engineers might master almost everything into the analog domain because they love working with analog gear. I certainly do that sometimes, but I would say that I've tried to accumulate what I think is the very best new gear as well as funky old gear that has a certain sound. If a mix comes in sounding really, really good, I have gear that will stay out of the way and do exactly what I need without inflicting any damage on the thing at all. Occasionally we'll get a mix in that's so good that I'm just happy to change the level on it if that's all that's needed.

There are some engineers that just like to slam everything. It seems like their only criterion is how loud they can make it, not how musical they can make it. For me, I'm under pressure from A&R people and clients to have things loud, but I try to keep the music at all costs. I'll think nothing of doing a Foo Fighters record one day, where it's totally appropriate to have it smashed, then the next day doing something that's perhaps even 4dB quieter than that because it suddenly needs the dynamics for it to breathe.

The loudness wars... Where did that come from?

I think it came from the invention of digital-domain compressors. When digital first came out, people knew that every time the overload light went red, you were clipping, and that hasn't changed.

We're all afraid of the over levels, so people started inventing these digital-domain compressors where you could just start cranking the level up. Because it was in the digital domain, you could look ahead in the circuit and have a theoretical zero attack time or even have a negative attack time if you wanted to. It was able to do things that you couldn't do with any piece of analog gear. It will give you that kind of an apparent level increase without audibly destroying the music, up to a point. And of course, once they achieved that, then people started pushing it as far as it would go.

I always tell people, "Thank God these things weren't invented when the Beatles were around, because for sure they would've put it on their music and would've destroyed its

longevity.” I’m totally convinced that over-compression destroys the longevity of a piece. When someone’s insisting on hot levels where it’s not really appropriate, I find I can barely make it through the mastering session.

Another thing that contributed to it was the fact that in Nashville, the top 200 Country stations got serviced with records from the record company, but there was an agreement that the major record companies have for all the other stations to get serviced with a special CD every week that had the different label’s new singles on it.

When they started doing that, the A&R people would go, “Well, how come my record isn’t as loud as this guy’s record?” And that led to level wars, where everyone wanted their song to be the hottest one on the compilation. When the program director of the radio station is going through a stack of CDs, a mediocre song that’s twice as loud as a great song might seem more impressive at first, just because it grabs you by the neck. It has a certain impressiveness about it, so you listen to it before realizing there’s no song there, but at least on first listen it might get the program director’s attention.

I suppose that’s well and good when it’s a single for radio, but when you give that treatment to an entire album’s worth of material, it’s just exhausting. It’s a very unnatural situation. Never in the history of mankind have we listened to such compressed music as we listen to now.

That’s all beginning to change a little, don’t you think?

We’re trying to. Whenever I’m doing an album that starts with a single, I’ll have time to do it at several different levels so they can hear what they’d be missing if they squash it to death. Then there are some artists like Jack White, where we did a version of his *Blunderbuss* record with no compression except what he did in mixing. You actually have to turn up the level on your playback system when you play it back, but when you do it sounds amazing. His fans were all raving about how great it sounded. The Daft Punk record [*Random Access Memories*] is not heavily squished either, compared to other electronica records. We raised the level, but it’s not insane.

I’ve been working with the Core Audio people at Apple about different issues. Album Sound Check in iTunes means that if you download an album, there’ll be one sound-check figure that all the albums will refer to, and all the mastering between the loud and soft tracks will remain correct. Apple did that only as a preference, but there’s no reason why they can’t turn on Sound Check on iTunes as a default, like they did for iTunes Radio. We keep trying to educate producers all the time that you have to check it out on your computer with Sound Check turned on to hear how it will sound when it’s streamed.

Do you often get asked to add effects?

Oh yeah, it happens often enough. A lot of people assemble mixes on Pro Tools, and they don’t

listen to it carefully enough when they're compiling their mix, so they actually cut off the tails of their own mixes. You can't believe how often that happens. A lot of times we'll use a little reverb to just fade out their chopped-off endings and extend it naturally. I do a fair amount of classical music mastering, and very often a little bit of reverb is needed on those projects, too. Sometimes if there's an edit that for some reason just won't work, you can smear it with a bit of echo at the right point and get past it. Sometimes mixes come in that are just dry as a bone, and a small amount of judicious reverb can really help that out.

Tell me about your monitors.

I used to have Duntech Sovereign 2001 monitors when I worked at Masterdisk, and when I started Gateway, I got another pair of Duntechs with a new pair of Cello Performance Mark II amplifiers. These are the amps that will put out like 6,000-watt peaks. One never listens that loudly, but when you listen, it sounds as though there's an unlimited source of power attached to the speakers. You're never straining the amp, ever. I used those Duntechs for quite a while.

Then, when I began doing 5.1 surround music, I really fell in love with EgglestonWorks Andras. I told Bill Eggleston if he ever decided to build a bigger version of the Andras to let me know, and maybe I'd consider changing my Duntechs if I thought they sounded better. He decided to build what he thought was the ultimate speaker, which is called the EgglestonWorks Ivy speaker. [He names all of his speakers after former wives or girlfriends.] These speakers have granite on the sides of them and weigh close to 800 pounds a piece. There are three woofers on the bottom, a couple of mids, a tweeter, and then a couple of more mids on the top. Actually, each cabinet has 23 speakers in it.

They're amazing. Every client that comes in, once they tune in to what they're listening to, starts commenting on how they're hearing things in their mixes that they never heard before, even sometimes after working weeks on them. It's great for mastering because they're just so accurate that there's never much doubt as to what's really on the mix.

One reason I've always tried to get the very best speaker I can is I've found that when something sounds really right on an accurate speaker, it tends to sound right on a wide variety of speakers. I've never been a big fan of trying to get things to sound right only on NS10Ms.

Do you listen only with that one set of monitors, or do you listen to nearfields?

Primarily just the big ones, because they tell you everything, but I do have a set of NS10Ms and some ProAcs and stuff like that. Lower-resolution nearfields have their place. In the case of the NS10Ms, the reason we have them there is just so the client can hear what he thought his mix was like. The NS10M kind of dials in a little bit more reverb than you think you have and more punch than is really there. When I'm teaching people, I make sure that they listen on NS10s and ProAcs and speakers like that a lot, so

they can learn in their head how to translate from one to the other.

Do you think that having experience cutting lacquers helps you now in the digital domain?

It does. I'm certainly more concerned about compatibility issues than a lot of the mixers are, especially as more people are getting into synthetic ways of generating outside-of-the-speaker sound. Some people just get into this and don't realize that their piano solo is gone in mono. People do still listen in mono, but some artists just don't seem to be bothered by the lack of compatibility. Nevertheless, I'm probably more hypersensitive to sibilance problems than I would otherwise be if I hadn't cut a lot of disks.

Does that mean you still listen in mono a lot?

I certainly check in mono. We have correlation meters on our consoles. In my room, if you're sitting in the sweet spot and flip the phase on one of the speakers, the entire bass goes away. It's almost as if you were doing it electronically, so you can hear any phase problems instantly. Plus, I have the ability to monitor L minus R, as well as to hear the difference channel if I need to.

Tell me about your signal path.

First of all, we were doing so much surround work that we stripped the room down to the floating floor and re-laid Transparent Audio cable. We had used it before, but their cable technology has improved to the point where we had to switch. For people who think cable doesn't make a difference, it was quite dramatic, and clients that really knew my room well were marveling at the difference. We also still run our studio off batteries and make our own 60Hz here, so we have very clean power. They're like the size of a refrigerator. [Laughs]

Then there's the eight-channel SPL mastering console and a 256 x 256 Z Sys router. I used to waste like two hours of my time every day switching from 5.1 to stereo and back. It wasn't time that you could bill a client, so it was just wasted from my life. Now a single click completely reroutes my room for 5.1.

What are you using for a workstation?

We've been using the Pyramix workstation, and we have the Horus convertor that will do everything up to 384kHz PCM or DSD or DXD [Digital eXtreme Definition]. It's a really good system. We play off of Pro Tools and we also have a Nuendo system that's part of an AudioCube system that supports a couple of plugins that are only on the AudioCube.

Does that mean you're going from the digital domain to analog and then back to digital?

It completely depends on the project, but normally, yes. And of course we still get tape.

That Daft Punk record that I got three Grammys for was on five different formats, including DSD and PCM, and the tape won out for that particular record.

If it's something like a mix from Bob Clearmountain, where the mix is almost perfect to begin with and you don't want to do too much to it, then it works just to stay in the digital domain.

We do place a lot of attention on analog at our place. We've got six different ways of playing back analog tape. We've got a stock Studer A820 and a Studer that's got Cello class A audiophile electronics. We've got a stock ATR, a tube ATR, and an unbalanced ATR. We also have one of the Tim de Paravicini 1-inch two-track machines with his fantastic tube electronics. When you record with his custom EQ curve at 15 ips, it's basically flat from 8 cycles up to 28kHz. It's unbelievable. You put an MRL test tape on his machine, and it comes back zero VU all the way.

Is all your processing done analog?

No. If I'm in the analog domain, I have six channels of Manley Massive Passive and six channels of GML mastering equalizers, six channels of the SSL compressor, and six channels of a Millennia Media compressor. If we're in the digital domain, then I use plugins, or if it needs nothing, just level.

How much of what you're doing will end up in high-res?

A lot. We've been archiving everything in high-res since around 2004 or so. I've always mastered at the highest resolution that made sense for what the client gave us. Now with MFiT, the record companies are loving us because there's tons of catalog stuff that all we have to do is un-archive it and use it for the MFiT.

How are you archiving?

The old stuff was all AIT, and now it's LTO tape.

Are you doing multiple masters?

Oh, yeah. We usually do separate masters for vinyl, for HD tracks or other high-res downloads, Mastered for iTunes, and normal 44.1/16 for normal downloads and CD, if they order it.

What's the hardest thing that you have to do? Is there a certain type of music or project that's particularly difficult?

I think the most difficult thing is when the artist is going through the period where they just can't let go of the project. You get into the psychological thing where in the same sentence they say, "I want you to make the voice more predominant, but make sure it doesn't stick out." Just contradictory things like that. They'll say, "This mix is too bright," and then you'll dull it up like half a dB and they say, "Oh, it doesn't have any air

anymore.” It’s that kind of thing.

Do you have a specific approach to mastering?

To me, music is a very sacred thing. I believe that music has the power to heal people. A lot of the music that I work on, even some of the heavy-metal stuff, is healing some 13-year-old kid’s angst and making him feel better, no matter what his parents might think about it, so I treat all music very seriously.

I love all kinds of music. I master everything from pop and some jazz to classical and even avant-garde. I used to be principle trumpet player in the Utica, New York Symphony Orchestra, so I always put myself in the artist’s shoes and ask myself, “What if this were my record? What would I do with it?” That’s why I try to get some input from the artist. If they’re not here, at least I try to get them on the phone and just talk about what things they like. I just take it all very seriously.

Chapter 20

Glenn Meadows - Mayfield

Mastering

Glenn Meadows of Mayfield Mastering is a two-time Grammy winner and a multi-TEC award nominee. He has worked on scores of gold and platinum records for a diverse array of artists, including Shania Twain, LeAnn Rimes, Randy Travis, Vince Gill, and Steely Dan, as well as for producers and engineers such as Tony Brown, Jimmy Bowen, and Mutt Lange.

What's your philosophy on mastering?

I think that mastering is, and always has been, the real bridge between the pro audio industry and the hi-fi industry. We're the ones that have to take this stuff that sounds hopefully good or great on a big professional monitor system and make sure it also translates well to the home systems. We're the last link to get it right or the last chance to really screw it up and make it bad, and I think we're all guilty at times of doing both.

What makes a great mastering engineer?

The ability to use discretion. The ability to listen to a piece of product and say, "You know, this really doesn't need much of anything." At this point in my career—I've been doing this for more than 30 years now—if I put a client's mix up and I don't have a pretty good clue by the time I'm at the end of the first run of the first song as to what that song needs, they ought to go back and remix.

It has to do with the experience of the engineer working in his environment. He's in the same room every day for years. I can walk into this room in the morning and know if my monitors are right or wrong just by listening to a track from yesterday. To me, that's the value of a mastering engineer. What they bring to the table is the cross-section of their experience and their ability to say, "No, you really don't want to do that."

When you use your compression technique, are you using the typical radio attack and release settings? Long attack, long release?

No, it varies. It depends on what the tempo of the music is doing. I'll adjust it track by track. Most everything I do is tailored to what the music dictates that it needs. There's no preset standard that I'm aware of that I use, although I have had a producer come in, and he had me master a record, and then he went back and matched it and stored the setting: "Ah, there's the Glenn Meadows setting." He told me he did the same thing for Bob Ludwig, too. He had a couple of things mastered up there and then found a common setting, and now he's got it as his Gateway preset. He does his own mastering now. "Ah,

make it sound like Gateway. There it is.” I told Bob [Ludwig] that, because he and I have been friends for probably 20 years, he just died laughing. He said, “If you can find out what that setting is, send it to me. I’d love to have it, because I don’t know what I do.”

My typical approach is to use like a 1.5:1 compression ratio and stick it down at –20 or –25dB so you get into the compressor real early and don’t notice it going from linear to compressed, and basically just pack it a little bit tighter over that range. I’ll get maybe 3dB of compression, but I’ve brought the average level up 3 or 4dB, and it just makes it bigger and fatter. People think that they have to be heavily compressed to sound loud on the radio, and they don’t.

What has changed from the last time we talked for the book?

For me, 98 percent of what I do now is working all in the box. I know there are people that work 100 percent outside the box, and their workstation is nothing more than a storage medium and an assembly point for making their masters. For me, it’s the other way around.

I’ve found that the quality of in-the-box processing has really improved, to the point where I find the same color of analog in the box if that’s what I want. I find that I can keep the finished product truer to what they originally had by keeping it all in the box.

I’m getting 24-bit files and mixes that actually have dynamic range and headroom, which allows for the ability to do high-fidelity mastering if that’s what we need to do, or a more in-your-face processed master. We can do an expanded and open mastering if they want to take it to vinyl, because a lot of people are learning that vinyl really can’t take that squeezed, squashed master that you’re doing for CD and online. It doesn’t make the mixes sound like those great vinyl records from years ago, because the mixes aren’t the way they were 15 or 20 years ago. If you want it to sound like that, you’ve got to mix it with dynamic range and lots of transient response. A lot of guys doing it today haven’t heard that style of mixing before. There’s a whole generation of people who’ve only mixed on workstations and approved mixes on iPods and earbuds. They’ve never really heard high-res, high-definition audio.

That said, the hardware that’s now coming out to play high-res audio will play every possible format, so we won’t be caught in another format war again. The new decoder chips support everything from MP3 all the way up to DSD, all on one chip, and some of the new consumer hardware will connect directly to the Internet via Wi-Fi so you won’t have to connect to a computer to get files in and out.

Do you get a lot at 192k?

We get about 10 to 15 percent. The interesting thing is that a lot of it has still been

compressed and destroyed, which defeats the purpose.

How often are you called upon to crush something?

Probably more than half the time, and we get into some discussions about it. The one good thing I'm seeing is more people are backing away from that and understanding what happens when you take this crushed stuff and convert it to MP3s. We use Apple's MFiT apps to show the person what happens when you do that. If we pull the level back a little, you can really hear the difference.

I did a panel at the Nashville Recording Workshop, where I took some stuff that was pretty heavily compressed and was able to run it at different levels into the Mastered for iTunes software to A/B things. I had a DAC that I could vary the level in 1dB steps so that I could lower the level into the encoder and then raise it by the same amount during playback so the reference level was the same. As I pulled it down, the crowd was amazed at how much better it sounded the lower the level to the encoder was, to the point that you could even hear it over the sound system in this big room. When we got down 4dB in level, they could hear that it was virtually identical to the original 192kHz/24-bit master, so it does make a difference, and people are starting to realize that.

The other thing that we're doing is showing people the difference in sound quality when Apple's Sound Check is employed. I took some of the old Steely Dan stuff that I mastered that had a lot of dynamic range and some current product, and we listened to both with Sound Check turned on and off. Everyone is blown away how great the Steely Dan stuff sounds when Sound Check is turned on, versus the more modern stuff that initially they thought sounded louder and better. It appears that Spotify is targeting a -20LUFS reference point, so something that had a -8LUFS reading would be turned down 12dB, which means the peaks only go 6 to 8dB above that. This means that the song at the lower LUFS level now has more peak level, which means it's going to sound louder. iTunes Radio appears to be set at -16LUFS, and you can't turn it off. Spotify does something similar as well.

The next question for mastering engineers is, When will Sound Check permanently default to On in iTunes, rather than just On in iTunes Radio? Once they do that, all of a sudden the levels will be different from what they were on many tracks, with what was loudest before now sounding quieter, and they won't be able to change it.

I hope that this is the beginning of the end of the "dark ages" of mastering, so we can go back to making music again, because that's what ultimately connects to the consumer.

What gear are you using?

My mastering workstation is still SADiE. I've been using it since it came out in 1991, and for me it's still one of the cleanest and fastest systems on the market. The monitors are

PMC IB1's powered by Bryston 7B amps. My room is incredibly accurate and easy to work in. One of the downsides is that I work too fast. If I work on an hourly rate, I get done too quickly [laughs], so I switched to a flat rate.

I listen to either the Benchmark DAC2 or the DA convertors that are in my Crookwood C4 monitor controller, which allows me to do both stereo and 5.1. It's one of the cleanest things going.

I've got a bunch of plugins, from iZotope Ozone to some of the Slate stuff, a package called the Dynamic Spectrum Mapper, and some various Waves stuff. It's plugin du jour.

Are there any that you come back to all the time?

I tend to come back to the Ozone 5. It's so amazingly flexible and clean, and built by guys who literally are rocket scientists! I also have the RX3 restoration package and several Cedar plugins built into SADiE for declicking and retouch.

Do you do much restoration?

It varies. There are some months where there's a lot to do, but then it goes away for a while. Sometimes I'll do some forensics for the police, where they need some voices pulled out of some bad surveillance tapes.

Do you make a comparison of what you're doing on a typical home hi-fi system?

No, what I think is really difficult is that if you put up two or three different monitors to get a cross-section, then you don't really know when anything is right because they all sound so different. I used to run little B&W 100s, and I'd also have the requisite NS10s in the room; and during that time when I was switching back and forth, I found my mastering suffered radically because I didn't have an anchor anymore. I didn't have a point where I knew what was right, because the character of the speakers was so different from each other. Once you listened to one for a couple of minutes, you lost your reference point on the others.

The reason people come to a mastering engineer is to gain that mastering engineer's anchor into what they hear and how they hear it, and the ability to get that stuff sounding right to the outside world. If you start putting all this stuff up on small speakers and try this and try that, you've basically created a big, confused image for the mastering engineer.

So you never listen to a smaller pair?

I do at home. I do in the car. I do outside of the mastering room, but in the room itself when I'm working? No, it's the one set of monitors.

If I get a producer that says, "Well, I've gotta listen on...fill in the blank," then we get a pair, and it's like, "Okay, here's the button that turns them on. Here's how you start."

Here's how you put the EQ in and out if you want to listen that way. Call me when you're finished listening." Then I leave the room and let them listen, because it literally rips me away from my anchor. If I start listening on different-sounding monitors, then I'm completely lost; but on the monitors that I've worked on in the same room, I know how they sound. I know what they need to sound like, and the repeat clients go, "Yep, that sounds right. Yep, that sounds good." What you find is typically within a song or two of working with somebody who has been in here, they settle into it and say, "Okay, yeah. I really can hear all that detail. I understand exactly what you are doing." We put other things up for them to listen to that they're familiar with to get a cross-check on what I'm used to hearing.

How have your clients changed?

The clients are a lot of the kids coming up the line that realize the advantage of having an outside set of ears rather than just making it loud themselves. The reality is that they end up with a better product by letting somebody else do it. They realize that if they get their mixes in pretty good shape, they can take it to a mastering facility to get it where it really needs to be.

We get a lot of stuff early on in mixing, and we give them feedback. I've found over the years that corrections are pretty broad curves where you're correcting for room and speaker anomalies. Once you find what that is, you're now pretty much on track, and the rest of the album will fall in place pretty quickly.

How many sessions are attended?

Thirty-five to forty percent. It's a lot less than it used to be. A lot of time we're getting work from people all over the country. It's so convenient to just pop a file up into the cloud, but we also host our own FTP site here in the building. If the client is close to computer literate, they can use an FTP program instead of two-stepping it through a cloud service. We just put a DDP file and a DDP player back in their folder, they log in, and they're done. It cuts out sending CD references by FedEx back and forth like we used to.

How much are you doing that's for vinyl?

Maybe 5 percent. It's not a lot. In most cases we do a second master where it fits more in line with what works for vinyl. I cut vinyl for 30 years, so I know what's going to work and what doesn't.

Chapter 21

Doug Sax - The Mastering Lab

If ever there was a title of “Godfather of Mastering,” Doug Sax has truly earned it, as evidenced by the extremely high regard in which the industry holds him. One of the first independent mastering engineers, Doug literally defined the art when he opened his world-famous Mastering Lab in Hollywood in 1967. Doug recently passed away, but his magic remains a big part of the albums he worked on for major diverse talents as The Who, Pink Floyd, The Rolling Stones, the Eagles, Diana Krall, Kenny Rogers, Barbra Streisand, Neil Diamond, Earth, Wind & Fire, Rod Stewart, Jackson Browne, and many, many more.

Do you have a philosophy about mastering?

Yes. If it needs nothing, don’t do anything. I think that you’re not doing a service by adding something it doesn’t need. I don’t make the stew, I season it. If the stew needs no seasoning, then that’s what you have to do, because if you add salt when it doesn’t need any, you’ve ruined it. I try to maintain what the engineer did. A lot of times they’re not really in the ballpark due to their monitoring, so I EQ for clarity more than anything.

When you first run something down, can you hear the final product in your head?

Oh yes, virtually instantly, because for the most part I’m working with music that I know what it’s supposed to sound like. Once in a while I’ll get an album that’s so strange to me because of either the music or what the engineer did that I have no idea what it’s supposed to sound like, and I often will pass on it. I’ll say, “I just don’t hear this. Maybe you should go somewhere where they’re glued into what you’re doing.”

For the most part, I’m fortunate to usually work on things that sound pretty good. I work on most of the recordings from great engineers like Bill Schnee, George Massenburg, Ed Cherney, and Al Schmitt. These are clients that I’m the one they go to if they have a say in where it’s mastered. Every room has its claim to fame, and mine is that I work on more albums nominated for engineering Grammys than any other room, and probably by a factor of three or four to the next closest room.

How has mastering changed over the years from the time you started until the way it is now?

My answer is maybe different than everyone else’s. It hasn’t changed at all! In other words, what you’re doing is finessing what an engineer and artist has created into its

best possible form. If an engineer says, “I don’t know what it is, but the vocal always seems to be a little cloudy,” I can go in there and keep his mix the same, yet still make the vocal clearer. That’s what I did in 1968, and that’s what I still do. The process is the same, and the goal is the same. I don’t master differently for different formats, because you essentially make it sound as proper as you can, and then you transfer it to the final medium using the best equipment.

One thing that has changed recently is that every client that comes in wants vinyl again. Almost nothing comes into the Lab that doesn’t do vinyl anymore. For one thing, it doesn’t cost that much. For another \$1,500 you can be doing vinyl, and you’re in a young market as the people buying these turntables are 18 to 25, and that’s proven. If you want to get your album to people that are really listening to the music, that’s the way. It’s also where the people that are going to buy high-res downloads are coming from.

Right now we’re mastering a Jackson Browne album and making a CD master, MFiT master, 96k master, 192k master, DSD master, and vinyl. That’s six different formats. Three years ago we made a CD master, and that was it. That’s becoming more and more routine.

I think this is all an offshoot from the phonograph record in the home. The fact that someone has to make a commitment to listening to a record and won’t be listening on earbuds, but real loudspeakers, is a revolution right there.

Do you think that working on vinyl would help a newer mastering engineer who’s never had that experience?

I don’t know if working on vinyl helps. I think having worked on many different types of music over the years helps. In one sense, being from the vinyl days I was used to doing all the moves in real time. I always cut directly from the master tapes, so if you blew a fade on the fourth cut, you started over again. So the concept of being able to do everything in real time instead of going into a computer probably affects the way I master now. I don’t look at things as, “Oh, I can put this in and fine-tune this and move this up and down.” I look at it as to what I can do in real time.

I find that the idea that you have a track for every instrument and you put them all together to have great clarity doesn’t work. I think it works the opposite way. The more you separate it, the harder it is to put together and have clarity, so if you’re EQing for musical clarity to hear what is down there, that’s unchanged today from way back 40 years ago. It’s the same process, and the EQ that would make somebody call up and say, “Wow, I really like it. I can hear everything and yet it’s still full,” is still as valid today as it was then.

Many artists won’t spend the money on recording and mixing, but it seems they’ll

spend it on mastering to get the ears of a pro. Have you experienced that?

Yes, but I think the caveat is how much money they're willing to spend. The amount of money it takes to open up a mastering facility today is minuscule compared to what it used to be. You can do almost everything without a big investment. The question then becomes, "Are you willing to spend the extra money for the expertise of the mastering engineer?" Just owning a Pro Tools system does not make you a mastering engineer.

The fact that you have a finely tuned room and a super high-quality playback system is hard to compete with.

Yes, but if you figure in the jobs that no one attends, they can't experience that, so we have to supply something that they can hear at home that's better than what they could do themselves. There are some engineers that do come to our facility, but the majority is being sent to us now.

You're just out of the way enough in Ojai [a little more than an hour from Los Angeles without traffic] that many might not want to make the trip from LA.

I was always concerned that the distance would affect people that wanted to attend, but that turned out not to be true. They still come and make a day of it. If someone wants to attend a session and they're in LA and don't want to drive, they won't even book us. I don't think it's hurt us in the long run. I think that everyone who comes here really enjoys it.

What's the hardest thing you have to do?

I come from a time when an album had a concept to it. The producer worked with one engineer and one studio, the group recorded everything, and there was cohesiveness as to what was put before you. Once you got what they were doing, you sort of had the album done. The multiple-producer album to me is the biggest challenge, because you might have three mixes from Nashville, a couple from New York, and two that are really dark and muddy, and three are bright and thin. The only good part that I see about this is that you absolutely have to use a mastering engineer in this case, or the mixes don't work together. The hard part for the mastering engineer is to find some middle ground, so that the guy with the bright, thin sound is still happy with what he's done and doesn't drive off the road when the dull, thick one plays after the bright, thin one. That's the biggest challenge in mastering—making what is really a cafeteria sound feel like a planned meal.

I'm very proud of the fact that I've trained a lot of good mastering engineers, and I'll tell them, "You're not going to learn how to master working on a Massenburg mix. It's pretty well done, and if he didn't like it, he wouldn't have sent it. When you get mixes from engineers that are not great, or you get these multiple-engineer things, then you can sort of learn the art of mastering by making these things work using your ears."

Is it true that you were the first independent mastering engineer?

Absolutely. Independent has to be clarified because if you go back to the late '60s and before, everything was done in-house. You were signed to a label, you were given an A&R man, and you stayed within the label. If you recorded at Capitol, then you went down to Capitol's mastering to get your product cut to lacquer. You went to Capitol's art department, and they gave you the artist that designed your cover, and that's the way it was.

It was really at the end of the '60s that certain top producers would say, "I love the security, but I would like to work with an artist that's not on this label. I would like to work with Streisand, but she's on Columbia." So they started to break off from the label and really started the process where nobody is tied to one anymore. The cry became, "If you sign me, I'll use the engineer I want, and I'll record and master where I want." That's 40 years of hard-fought independence, so from the standpoint of an independent that is not aligned with a label, just a specialty room that handles mastering, the answer is yes.

I was one of the pioneers when there was no independent business. We opened up our doors on December 27, 1967, and by '71 or '72, you couldn't get into the place because we were so busy. By '72, we were doing 20 percent of the top 100 chart, and there weren't a lot of competitors. There was Artisan in LA, and Sterling and maybe Masterdisk just starting in New York, and that was it. Now there seems to be a thousand, because the reality is that it's very easy for someone to go into this business now, or for the artist or engineer do it himself. You can get a workstation with all the bells and whistles for a song and a dance. A Neumann lathe setup in 1972 was \$75,000, and that was just the cutting system; you still needed a room and a console, so you had to have a big budget, and there were only a few people doing it as a result. Now you fire it right up.

And don't forget that in the industry, for almost 10 years there were no tones on an analog tape, so you didn't know how to line up to the machine.

There were no tones?

No tones. I'm one of the instigators in railing on these guys to go back and print the tones so I could at least set my machine to where your machine was. There was no such thing as nearfield monitoring, either. It didn't exist. People used to go to these strange studios with big speakers in the wall, most of which were useless as far as relating to the real world, and the engineers never knew that they were out in left field because they had nothing to take home. The cassette was just starting, and only a handful of engineers that I can think of actually had a 15 ips (inches per second) tape machine at home that they could take home a mix and find out where they were.

I started the process in the early '70s just in self-defense. I would say, "Look, before you do anything, come in with your first mix on the house, and find out if you're in trouble. We'll listen to it and get you straight." I just got tired of watching these guy's eyes open the first time they ever heard their mixes outside of the studio. "Oh my God. I couldn't hear any highs in the studio, so I kept adding highs." That absolute horrendous reality is really the reason why nearfields came in.

Did you have a lot of experience with the early digital technology as well?

My partner and I did some of the pioneering work in digital in the late '70s. The classic 3M digital tape machines were designed out in Camarillo [California—near Los Angeles], and my partner lived in Camarillo and did the original piano tests for them in '78. We participated in the very first recordings that were done on the Soundstream machine [the first digital recorder] before it was even up to a 44.1k sampling rate, so when I was critical of digital in the past, it was because I really have heard digital from the beginning and I knew that it was not up to the best of analog. That said, we're talking about 1980, and there's been a lot of development since. I get a lot of 96/24 stuff in now, and when I say that a 96/24 recording done with great converters sounds terrific, that's also true.

The truth of the matter is that the tools are getting so much better. Digital technology is moving so fast, and it's gone from, in my view, absolute garbage to "Hey, this is pretty good." Mastering engineers don't like that because they used to be the ones that made it loud, but the reality is that everyone can make their mix loud. Once that becomes absolutely no trick at all, then the question becomes, "Are there things that maybe we should do besides just make it loud?" I'm hoping that there's still going to be a business for someone that treats the music with love and respect when they're mastering it, and I think there's going to be a small reversion away from, "I want the loudest master."

Did you ever get caught up in the loudness wars?

We never did. We always had healthy levels, but most of the people that are looking for the loudest product don't actually come to us. I'm noted for leaving some dynamic range when I master. We do get people that occasionally say, "I want the loudest CD ever made," and I say, "You're in the wrong place." Once in a while they'll pull out a CD and put it on, and it's absolutely blazing, and I'll say, "Find out where that was mastered and go there and get what you're looking for."

The trend back to vinyl is actually getting rid of some that, though. You don't normally cut the record from the CD master, although you can. You make it a little better for the LP in that you can't make it super-loud because you just don't have room on the grooves.

Can you take us through your signal chain, or is that proprietary?

No, it's not proprietary. My gear hasn't really changed that much over the years. Obviously I have workstations, and everything that comes in is opened up in some

version of Pro Tools. The files are put into Pro Tools for playback, and then we choose a clock, because each one sounds different on it. Ninety-nine percent of the time, I do no processing in Pro Tools, but we might change the level or remove a few S's so we don't have to worry about that again.

We treat it as analog from there, exactly as we did back when I started. All the processing is done in the analog domain using the same EQs and limiters that we've used since 1968. The output from analog can go to DSD, or 96k or 192k or 44.1k, or 48k for surround, since it's connected to video. We use Josh Florian's converters [JCF Audio] and Benchmark converters for both the AD and DA.

What do you use for monitors?

I use ATC 150s in the main room, and I use Mastering Lab Tannoys in the vinyl room.

It's fantastic that what you have has weathered the test of time.

Yes. It's the same concept that I have about mastering. I don't master any differently today than I did in 1968. The speakers allow me to put the right stuff on, and if they steer me wrong, then they're worthless.

Glossary

0dB FS (Full Scale). The highest level that can be recorded in the digital domain. Recording beyond 0dBFS can result in distortion.

5.1. A speaker system that uses three speakers across the front and two stereo speakers in the rear, along with a subwoofer.

1630. A first-generation two-track digital tape machine made by Sony utilizing a separate digital processor and a 3/4-inch U-matic video tape machine for storage. The 1630 was the primary master tape delivered to the pressing plant in the early years of the CD, but they are considered obsolete today. A model 1610 predated this machine.

AAC. Advanced Audio Coding is a standard lossy data compression encoding scheme for digital audio used exclusively on Apple iTunes.

acetate. An acetate is a single-sided vinyl check disc, sometimes called a *ref*. Due to the extreme softness of the vinyl, an acetate has a limited number of plays (five or six) before it wears out. See *ref*.

A/D. Analog-to-digital converter. This device converts the analog waveform into the digital language that can be used by a digital audio workstation.

AIFF. Audio Interchange File Format (also known as *Apple Interchange File Format*) is an audio file format designed for use in the Apple Macintosh operating system but now widely used in PCs as well.

airplay. When a song gets played on the radio.

asset. A multimedia element, either sound, picture, graphic, or text.

attack. The first part of a sound envelope. On a compressor/limiter, a control that affects how that device will respond to the attack of a sound.

attenuation. A decrease in gain or level.

Augspurger. George Augspurger of Perception Inc. in Los Angeles is one of the most revered studio designers. He also designs large studio monitors, each having dual 15-inch woofers and a horn tweeter.

automation. A system that memorizes and then plays back the position of all faders, mutes on a console, and just about every parameter in a DAW.

bandwidth. The number of frequencies that a device will pass before the signal degrades. A human being can supposedly hear from 20Hz to 20kHz, so the bandwidth of the human ear is 20 to 20kHz. Sometimes applies to computer data rate, where a high rate per second represents a wider bandwidth.

barcode. A series of vertical bars of varying widths in which the numbers 0 through 9 are represented by a different pattern of bars that can be read by a laser scanner. Barcodes are commonly found on consumer products and are used for inventory control, and in the case of CDs, to tally sales.

big ears. The ability to be very aware of everything going on within the session and with the music.

bit rate. The transmission rate of a digital system.

Blu-ray. The name of the optical disc format initially developed by Sony and Philips (inventor of the compact disc, cassette, and laserdisc) as a next-generation data and video storage format alternative to DVD. It supports just about every audio codec format and is able to store up to eight channels of 96/24 LPCM audio and six channels of 192/24 without data compression, along with high-def picture.

bottom. Bass frequencies, the lower end of the audio spectrum. See also *low end*.

brick-wall filter. A low-pass filter used in digital audio set to half the sampling rate so the frequency response does not go beyond what is suggested by the Nyquist Theorem.

brick-wall limiter. A limiter employing look-ahead technology that is so efficient that the signal will never exceed a certain predetermined level, and there will be no digital overs.

buss. A signal pathway.

CALM Act. Legislation passed by Congress to ensure that television commercials and programs are broadcast equally loud.

catalog. Older albums or recordings under control of the record label.

clip. To overload and cause distortion.

clipping. When an audio signal begins to distort because a circuit in the signal path is overloaded, the top of the waveform becomes “clipped” off and begins to look square instead of rounded. This usually results in some type of distortion, which can vary from soft and barely noticeable to horribly crunchy-sounding.

codec. Code-decode. A codec is a software algorithm that encodes and decodes a particular file format such as FLAC, AAC or MP3 ..

color. To affect the timbral qualities of a sound.

comb filter. A distortion produced by combining an electronic or acoustic signal with a delayed copy of itself. The result is peaks and dips introduced into the frequency response.

competitive level. A mix level that is as loud as your competitor’s mix.

compression. Signal processing that controls the dynamics of a sound.

compressor. A signal-processing device used to control audio dynamics.

cross-fade. In mastering, the overlap of the end of one song into the beginning of the next.

cut. To decrease, attenuate, or make less.

cutter head. The assembly on a vinyl-cutting lathe that holds the cutting stylus between a set of drive coils powered by very high-powered (typically, 1000 to 3500 watts) amplifiers.

D/A. Digital-to-analog converter, sometimes called a DAC. This device converts the digital audio data stream into analog audio.

data compression. A process that uses a specially designed algorithm to decrease the number of bits in a file for more efficient storage and transmission.

DAW. A digital audio workstation. A system designed for recording, editing and playback of audio. Modern DAW systems are primarily run using software on computers.

dB. Stands for *decibel*, which is a unit of measurement of sound level or loudness.

decay. The time it takes for a signal to fall below audibility.

decoupling. Isolating speakers from a desk or console by using rubber or carpet.

DDP. Disc Description Protocol. A proprietary format that is low in errors and allows high-speed glass-master cutting. It is currently the standard delivery format for CDs and DVDs.

digital domain. When a signal source is converted into a series of electronic pulses represented by 1s and 0s, the signal is then in the digital domain.

digital overs. The point beyond 0 on a digital processor where the red over indicator lights, resulting in a digital overload.

'dither. A low-level noise signal used to limit quantization distortion when lowering the bit resolution of an audio file.

DSP. Digital Signal Processing. Processing within the digital domain, usually by dedicated microprocessors.

dynamic range. A ratio that describes the difference between the highest and lowest signal level. The higher the number, equaling the greater dynamic range, the better.

edgy. A sound with an abundance of midrange frequencies.

element. A component or ingredient of the mix.

EQ. Equalizer, or to adjust the equalizer (tone controls) to affect the timbral balance of a sound.

equalization. Adjustment of the frequency spectrum to even out or alter tonal imbalances.

equalizer. A tone control that can vary in sophistication from very simple to very complex. See *parametric equalizer*.

feather. Rather than applying a large amount of equalization at a single frequency, small amounts are added at the frequencies adjoining the one of principle concern.

FLAC. Free Lossless Audio Codec. A lossless file format used to make digital audio files smaller in size, yet that suffer no degradation of audio quality.

Fletcher-Munson curves. A set of measurements that describes how the frequency response of the ear changes at different sound-pressure levels. For instance, we generally hear very high and very low frequencies much better as the overall sound pressure level is increased.

gain. The amount a sound is boosted.

gain reduction. The amount of signal level attenuation as a result of compression or limiting.

gain staging. Adjusting the gain of each processing stage of the signal chain so the output of one doesn't overload the input of another.

glass master. The first and most important step in the CD replication process involving electroplating a glass block in order to make the CD stampers.

groove. The pulse of the song and how the instruments dynamically breathe with it. Or, the part of a vinyl record that contains the mechanical information that is transferred to electronic info by the stylus.

headroom. The amount of dynamic range between the normal operating level and the maximum output level, which is usually the onset of clipping.

Hertz. A measurement unit of audio frequency, defined by the number of cycles per second. High numbers represent high-pitched sounds, and low numbers represent low-pitched sounds.

high-pass filter. An electronic circuit that allows the high frequencies to pass while attenuating the low frequencies. Used to eliminate low-frequency artifacts, such as hum and rumble. The frequency point where it cuts off is usually either switchable or variable.

hypercompression. Too much buss compression during mixing or mastering in an effort to make the recording louder results in what's known as hypercompression, a condition

that essentially leaves no dynamics and makes the track sound lifeless.

Hz. Short for Hertz. See *Hertz*.

I/O. The input/output of a device.

ISRC code. The International Standard Recording Code is used to uniquely identifying sound recordings and music video recordings. An ISRC code identifies a particular recording, not the song itself; therefore, different recordings, edits, and remixes of the same song will each have their own ISRC codes.

kbs. Kilobits per second. The amount of digital information sent per second. Sometimes referred to as *bandwidth*.

kHz. One-thousand Hertz (example: 4kHz = 4000Hz).

knee. How quickly a compressor will turn on once it reaches the threshold. A *soft knee* turns on gradually and is less audible than a *hard knee*.

lacquer. The vinyl master, which is a single-sided 14-inch disc made of aluminum substrate covered with a soft cellulose nitrate. A separate lacquer is required for each side of a record. Since the lacquer can never be played, a ref or acetate is made to check the disc. See *ref* and *acetate*.

latency. Latency is a measure of the time it takes (in milliseconds) for your audio signal to pass through your system during the recording process. This delay is caused by the time it takes for your computer to receive, understand, process, and send the signal back to your outputs.

limiter. A signal-processing device used to constrict or reduce audio dynamics, reducing the loudest peaks in volume.

LKFS. Stands for Loudness, K-weighted, relative to Full Scale, which distinguishes itself from the normal dBFS peak meters found on all digital gear in that it measures the loudness not instant by instant, but over a period of time.

look-ahead. In a mastering limiter, look-ahead processing delays the audio signal a small amount (about 2 milliseconds or so) so that the limiter can anticipate the peaks in such a way that it catches the peak before it gets by.

lossless compression. A compression format that recovers all the original data from the compressed version and suffers no degradation of audio quality as a result. FLAC and ALAC are lossless compression schemes.

lossy compression. A digital file-compression format that cannot recover all of its original data from the compressed version. Supposedly some of what is normally recorded before compression is imperceptible, with the louder sounds masking the softer ones. As a result, some data can be eliminated since it's not heard anyway. This selective approach, determined by extensive psychoacoustic research, is the basis for lossy compression. MP3

and AAC are lossy compression schemes.

low end. The lower end of the audio spectrum, or bass frequencies usually below 200Hz.

low-pass filter. An electronic frequency filter that allows only the low frequencies to pass while attenuating the high frequencies. The frequency point where it cuts off is usually either switchable or variable.

LPCM. Linear Pulse Code Modulation. This is the most common method of digital encoding of audio used today and is the same digital encoding method used by current audio CDs. In LPCM, the analog waveform is measured at discrete points in time and converted into a digital representation.

LUFS. Stands for Loudness Units relative to Full Scale. This was formerly the European standard for loudness, but is now identical to LKFS.

makeup gain. A control on a compressor/limiter that applies additional gain to the signal. This is required since the signal is automatically decreased when the compressor is working. Makeup gain “makes up” for the lost gain and brings it back to where it was prior to being compressed.

mastering. The process of turning a collection of songs into an album by making them sound like they belong together in tone, volume, and timing (spacing between songs). The process also applies to a single song in terms of making it sound similar in level to songs played before and after it on any distribution medium.

metadata. Data that describes the primary file. For instance, metadata can be information about an audio file that indicates the date recorded, sample rate, resolution, and so on.

midrange. Middle frequencies starting from around 250Hz up to 4000Hz.

modeling. Developing a software algorithm that is an electronic representation of the sound of hardware audio device down to the smallest behaviors and nuances.

monaural. A mix that contains a single channel and usually comes from only one speaker.

mono. Short for monaural, or single audio playback channel.

mother. In either vinyl or CD manufacturing, the intermediate step from which a stamper is made.

MP3. A standard data-compression format used to make audio files smaller in size.

multi-band compression. A compressor that is able to individually compress different frequency bands as a means of having more control over the compression process.

mute. An on/off switch. To mute something would mean to turn it off.

native resolution. The sample rate and bit depth of a distribution container. For example, the native resolution of a CD is 44.1kHz and 16 bits. The native resolution in film work is 48kHz and 24 bits.

noise shaping. Dither that moves much of the injected noise to a part of the audio spectrum beyond where the ear is less likely to hear it.

normalization. A selection on a DAW that looks for the highest peak of an audio file and adjusts all the levels of the file upward to match that level.

Nyquist Sampling Theorem. A basic tenet of digital audio that states that the frequency response of a system cannot go beyond half the sampling rate. If that occurs, artifacts known as aliasing are introduced into the signal, destroying the purity of the audio.

out of phase. The polarity of two channels (it could be the left and right channels of a stereo program) are reversed, thereby causing the center of the program (like the vocal) to diminish in level. Electronically, when one cable is wired backwards from all the others.

overs. Digital overs occur when the level is so high that it tries to go beyond 0dBFS on a typical digital level meter found in just about all equipment. A red overload indicator usually will turn on, possibly accompanied by audible clipping.

pan. Short for *panorama*, pan indicates the left and right position of an instrument within the stereo spectrum.

panning. Moving a sound across the stereo spectrum.

parametric equalizer. A tone control where the gain, frequency, and bandwidth are all variable.

parts. The different masters sent to the pressing plant. A mastering house may make different parts/masters for CD, cassette, and vinyl, or send additional parts to pressing plants around the world.

peaks. A sound that's temporarily much higher than the sound surrounding it.

phantom image. In a stereo system, if the signal is of equal strength in the left and right channels, the resultant sound appears to come from in between them. This is a phantom image.

phase. The relationship between two separate sound signals when combined into one.

phase meter. A dedicated meter that displays the relative phase of a stereo signal.

phase shift. The process during which some frequencies (usually those below 100Hz) are slowed down ever so slightly as they pass through a device. This is usually exaggerated by excessive use of equalization and is highly undesirable.

pitch. On a record, the velocity of the cutter head. Measured in the number of lines (grooves)

per inch.

plug-in. An add-on to a computer application that adds functionality to it. EQ, modulation, and reverb are examples of DAW plug-ins.

PMCD. Pre-Mastered CD, an obsolete format similar to a CD-R except that it has PQ codes written on the lead-out of the disc to expedite replication.

PQ codes. Subcodes included along with the audio data channel as a means of placing control data, such as start IDs and the table of contents, on a CD.

pre-delay. The time between the dry sound and the onset of reverberation. The correct setting of the pre-delay parameter can make a difference in the clarity of the mix.

presence. Accentuated upper-midrange frequencies (anywhere from 4k to 6kHz).

producer. The equivalent of a movie director, the producer has the ability to craft the songs of an artist or band technically, sonically, and musically.

proximity effect. The inherent low-frequency boost that occurs with a directional microphone as it gets closer to the signal source.

'Pultec. An equalizer sold during the '50s and '60s by Western Electric that is highly prized today for its smooth, unique sound.

pumping. When the level of a mix increases and then decreases noticeably. Pumping is caused by the improper setting of the attack and release times on a compressor.

punchy. A description for a quality of sound that infers good reproduction of dynamics with a strong impact. The term sometimes means emphasis in the 200Hz and 5kHz areas.

Q. The bandwidth, or the frequency range of a filter or equalizer.

ratio. A parameter control on a compressor/limiter that determines how much gain reduction will occur when the signal exceeds the threshold.

record. A generic term for the distribution method of a recording. Regardless of whether it's a CD, vinyl, or a digital file, it is still known as a record.

ref. Short for *reference record*, a ref is a single-sided vinyl check disc, sometimes called an *acetate*. Due to the extreme softness of the vinyl, a ref has a limited number of plays (five or six) before it wears out. See *acetate*.

reference level. This is the audio level, either electronic and acoustic, to which a sound system is aligned.

release. The last part of a sound envelope. On a compressor/limiter, a control that affects how that device will respond to the release part of the sound envelope.

resonance. See *resonant frequency*.

resonant frequency. A particular frequency or band of frequencies that is accentuated, usually due to some sympathetic acoustic, electronic, or mechanical factor.

return. Inputs on a recording console especially dedicated for effects devices such as reverbs and delays. The return inputs are usually not as sophisticated as normal channel inputs on a console.

RIAA. Recording Industry Association of America. A trade organization for record labels but dominated by the major labels.

RIAA Curve. An equalization curve instituted by the Recording Industry Association of America (the RIAA) in 1953 that enabled the grooves to be narrowed, thereby allowing more of them to be cut on the record, which increased the playing time and decreased the noise. This was accomplished by boosting the high frequencies by about 17dB at 15kHz and cutting the lows by 17dB at 50Hz when the record was cut. The opposite curve is then applied during playback.

roll-off. Usually another word for high-pass filter, although it can refer to a low-pass filter as well.

sample rate. The rate at which the analog waveform is measured. The more samples per second of the analog waveform that are taken, the better the digital representation of the waveform that occurs, resulting in greater bandwidth for the signal.

sequencing. Setting the order in which the songs will play on a CD or vinyl record.

scope. Short for oscilloscope, an electronic measurement device that produces a picture of the audio waveform.

shelving curve. A type of equalizer circuit used to boost or cut a signal above or below a specified frequency. Usually the high- and low-band equalizers built into many mixing boards are the shelving type.

sibilance. A short burst of high frequencies in a vocal sometimes due to heavy compression, resulting in the S sounds being overemphasized.

soundfield. The direct listening area.

SoundScan. The company (a division of the Nielsen Company) that measures record sales. Whenever a CD or DVD is sold, the barcode on the unit is scanned and recorded by SoundScan.

source. An original master that is not a copy or a clone.

spectrum. The complete audible range of audio signals.

SPL. Sound-pressure level.

spread. The time in between songs on a CD or vinyl record.

SRC. Sample-rate conversion.

stamper. In either vinyl or CD manufacturing, a negative copy bolted into the presser to actually stamp out records or CDs.

stems. A mix that has its major elements broken out separately for individual adjustment at a later time.

sub. Short for subwoofer.

subwoofer. A low-frequency speaker with a frequency response from about 25Hz to 120Hz.

tempo. The rate of speed, usually represented in beats per minute, that a song is played.

test tones. A set of tones used to calibrate a playback system. In the days of tape, they were added to a tape to help calibrate the playback machine.

threshold. The point at which an effect takes place. On a compressor/limiter, for instance, the threshold control adjusts the point at which compression will begin.

timbre. Tonal color.

track. A term sometimes used to mean a song. In recording, a separate musical performance that is recorded.

transformer. An electronic component that either matches or changes the impedance. Transformers are large, heavy, and expensive, but are in part responsible for the desirable sound in vintage audio gear.

trim. A control that sets the gain of a device, or the process of reducing the size or playing time of an audio file.

tube. Short for *vacuum tube*; an electronic component used as the primary amplification device in most vintage audio gear. Equipment utilizing vacuum tubes runs hot, is heavy, and has a short life, but it has a desirable sound.

TV mix. A mix without the vocals so the artist can sing live to the backing tracks during a television appearance.

U-matic. An industrial digital-video machine utilizing a cassette storing 3/4-inch tape. The U-matic is the primary storage device for the 1630 digital processor.

underscore. The instrumental score of a movie that plays below the dialogue and/or sound effects.

unity gain. When the output level of a process or processor exactly matches its input level.

variable pitch. On a record, varying the number of grooves per inch depending upon the program material.

Vinylite. The vinyl used to make records actually comes in a granulated form called *Vinylite*. Before being pressed, it is heated into the form of modeling clay and colored with pigment.

WAV. A WAV file is an audio data file developed by the IBM and Microsoft corporations, and is the PC equivalent of an AIFF file. It is identified by the “.wav” file extension.

word length. The number of bits in a word. Word length is in groups of eight. The longer the word length, the better the dynamic range.

About Bobby Owsinski

Producer/engineer Bobby Owsinski is one of the best selling authors in the music industry with 23 books that are now staples in audio recording, music, and music business programs in colleges around the world, including the *Deconstructed Hits series*, *Music 4.1 Internet Music Guidebook*, *The Mixing Engineer's Handbook* and more. He's also a contributor to Forbes writing on the new music business, his popular blogs have passed 7 million visits, and he's appeared on CNN and ABC News as a music branding and audio expert.

Visit Bobby's music production blog at bobbyowsinskiblog.com, his Music 3.0 music industry blog at music3point0.com, his Forbes blog at forbes.com/sites/bobbyowsinski/, his podcast at bobbyoinnercircle.com, and his website at bobbyowsinski.com.

Bobby Owsinski Bibliography

The Mixing Engineer's Handbook 4th Edition (BOMG Publishing)

The Recording Engineer's Handbook 4th Edition (BOMG Publishing)

The Mastering Engineer's Handbook 4th Edition (BOMG Publishing)

The Drum Recording Handbook 2nd Edition [with Dennis Moody] (Hal Leonard Publishing)

How To Make Your Band Sound Great (Hal Leonard Publishing)

The Studio Musician's Handbook [with Paul ILL] (Hal Leonard Publishing)

Music 4.1 - A Survival Guide To Making Music In The Internet Age 4th Edition (Hal Leonard Publishing)

The Music Producer's Handbook 2nd Edition (Hal Leonard Publishing)

The Musician's Video Handbook (Hal Leonard Publishing)

Mixing And Mastering With T-Racks: The Official Guide (Course Technology PTR)

The Touring Musician's Handbook (ISBN 978-1423492368 Hal Leonard)

The Ultimate Guitar Tone Handbook [with Rich Tozzoli] (Alfred Music Publishing)

The Studio Builder's Handbook [with Dennis Moody] (Alfred Music Publishing)

Abbey Road To Ziggy Stardust [with Ken Scott] (Alfred Music Publishing)

The Audio Mixing Bootcamp (Alfred Music Publishing)

Audio Recording Basic Training (Alfred Music Publishing)

The Music 3.0 Guide To Social Media. *Tips and Tricks for using Facebook, Twitter, YouTube and Google+.* (BOMG Publishing)

Social Media Promotion For Musicians *The Manual For Marketing Yourself, Your Band or your Music Online* (BOMG Publishing)

Social Media Promotion For Small Business and Entrepreneurs *The Manual for Marketing Yourself or your Business Online* (BOMG Publishing)

Deconstructed Hits: Classic Rock Vol. 1 (Alfred Music Publishing)

Deconstructed Hits: Modern Pop & Hip-Hop (Alfred Music Publishing)

Deconstructed Hits: Modern Rock & Country (Alfred Music Publishing)

The PreSonus StudioLive Mixer Official Manual (Alfred Music Publishing)

You can get more info and read excerpts from each book by visiting the excerpts section of bobbyowsinski.com.

Bobby Owsinski Lynda.com Video Courses

The Audio Mixing Bootcamp Video Course
Audio Recording Techniques
Mastering For iTunes
Audio Mastering Techniques
Music Studio Setup and Acoustics
Selling Music Merchandise
Selling Your Music: CDs, Streams and Downloads
Bookmarking Sites for Musicians and Bands
Pinterest for Musicians and Bands
Google+ for Musicians and Bands
YouTube for Musicians and Bands
Blogging Strategies for Musicians and Bands
Twitter for Musicians and Bands
Mailing List Management for Musicians and Bands
Website Management for Musicians and Bands
Facebook for Musicians and Bands
Social Media Basics for Musicians and Bands

[Get 7 free days to check out these and other courses on lynda.com.](#)

Bobby Owsinski's Online Connections

Music Production Blog: [bobbyowsinskiblog.com](#)

Music Industry Blog: [music3point0.com](#)

Forbes Blog: [forbes.com/sites/bobbyowsinski/](#)

Facebook: [facebook.com/bobby.owsinski](#)

YouTube: [youtube.com/polymedia](#)

Pinterest: [pinterest.com/bobbyowsinski/](#)

Linkedin: [linkedin.com/in/bobbyo](#)

Twitter: @bobbyowsinski

Website: [bobbyowsinski.com](#)

Courses: bobbyowsinskicourses.com



Table of Contents

Introduction	6
Meet The Mastering Engineers	7
The Essence Of Mastering	12
Why Master Anyway?	13
From Vinyl, To CDs, To MP3s, And Beyond	14
The Difference Between You And A Pro	18
There's Always Room For DIY	21
Digital Audio Basics	23
Sample Rate	23
Table 2.1: Sample Rates For Various Distribution Mediums	25
Bit Depth	25
Standard Audio File Formats	26
Data Compression	28
Prepping For Mastering	30
Mixing For Mastering	30
Mastering Session Documentation	31
Why Alternative Mixes Can Be Essential During Mastering	32
Monitoring For Mastering	34
The Acoustic Environment	34
Let's Fix Your Listening Area	34
The Monitors	35
Basic Monitor Setup	36
On The Bottom	41
Three Steps To Adding A Subwoofer	41
Calibrating Your Sub To Your System	42
Placing The Subwoofer	43
Amplifiers	43
Listening Techniques For Mastering	44
Monitors Versus Headphones	45
Mastering Tools	47
The Mastering Compressor	47
Compressor Overview	47
Using The Compressor In Mastering	48
Multi-Band Compression	49

The Mastering Limiter	50
Limiter Overview	50
Using The Limiter In Mastering	50
Multi-Band Limiter	51
The Mastering Equalizer	51
Using The EQ in Mastering	52
The Mastering De-Esser	54
Metering	56
The Peak Meter	56
The RMS Meter	58
The K-System Metering	58
The Perceived Loudness Meter	59
The Phase Scope	60
Using The Phase Scope	61
The Phase Correlation Meter	63
The Spectrum Analyzer	63
The Dynamic Range Meter	65
Dynamic Ranges Of Different Genres Of Music	67
Convertors	68
Consoles/Monitor Control	69
The Digital Audio Workstation	70
Mastering DAWs	70
Other Tools	71
Stereo Enhancement	71
M-S Processing	71
Mono	72
Mastering Techniques	73
The Basic Mastering Technique	73
Making A Loud Master	74
Competitive Level	74
Level Technique #1: The Compressor-Limiter Tandem	75
Level Technique #2: Multi-Compressor Packages	76
Advanced Level Techniques	77
The Effects Of Hypercompression	77
Competitive Level Isn't What It Used To Be	79
Setting The Compressor	79
Compression Tips And Tricks	80
Setting The Limiter	80

Using Multi-Band Compressors And Limiters	81
Reducing Sibilance With A De-Esser	82
Frequency Balance Techniques	83
The Mastering Signal Path	84
The Basic Mastering Signal Chain	84
An Advanced Signal Chain	85
Parallel Processing	86
Adding Effects	87
Editing Techniques For Mastering	88
Inserting Fades	88
Eliminating Intro Noise And Count-Offs	89
Making A “Clean” Master	89
Parts Production	90
Multiple Masters	90
Mastering For CD	92
CD Basics	92
How CDs Work	93
Scanning The Disc	95
Mastering For CD	96
Editing PQ Subcodes	96
CD Subcodes	97
Inserting ISRC Codes	98
Inserting UPC Codes	100
Inserting CD-Text	101
Song Order	101
Adjusting The Spreads	101
Using Dither	102
Rules For Using Dither	102
Delivery Formats	103
The DDP Master	103
FTP Transmission	104
Obsolete Formats	104
The Sony PCM-1630	104
The PMCD	106
How CDs Are Made	106
Duplication Versus Replication: What’s The Difference?	109
Mastering For Vinyl	111

A Brief History Of Vinyl	111
How A Vinyl Record Works	112
The Vinyl Signal Chain	126
The Master Lacquer	126
The Cutting Stylus And Cutter Head	127
The Lathe	127
The Mastering Console	128
The Preview System	129
Equalization	130
The Elliptical Equalizer	130
How Records Are Pressed	131
New Advances In Vinyl Technology	132
New Record Presses	132
A New Way Of Pressing	132
HD Vinyl	133
Mastering For Online Distribution	134
File Compression Encoding	134
Table 9.1: Commonly Used Audio Compression Formats	134
Lossy Versus Lossless Codecs	135
File Compression Encoder Parameters Explained	135
Table 9.2: Bit Rate Quality Comparison	136
It's All About The Source File	138
Choosing An Encoder	138
The Ins And Outs Of File Metadata	138
Creating Great-Sounding Online Files	139
Creating Files For Streaming Services	140
Table 9.3: Streaming Specs for the Most Popular Services	141
Be Aware Of Sound Check	141
Creating A FLAC File	142
Submitting To Online Stores and Services	142
Table 9.4: Digital Music Distributors	143
Submitting To Online Song Databases	143
A Look At AAC, The iTunes File Format	145
Mastering Tips For iTunes And Apple Music	146
iTunes Sound Check	147
The Mastered For iTunes Format	147
The Mastered For iTunes Tools Package	149

Using The afconvert Tool	149
Using The Droplet Tool	149
Using The Audio To WAVE Droplet	150
Using The afclip Tool	150
Using The AURoundTrip AAC Audio Unit Tool	151
Using The Test-Pressing Feature	152
Submitting To The iTunes Store Or Apple Music	152
Other Types Of Mastering	154
Online Mastering	154
High-Resolution Mastering	154
Table 11.1: High-Resolution Storage Requirements	155
Direct Stream Digital (DSD)	155
Blu-Ray Disc	156
Blu-Ray Audio Specs	157
Table 11.2: Blu-Ray List Of Codecs	157
Mastering Music For Film	157
Mastering Music For Television	158
The Effect Of The CALM Act	158
Meet The LKFS Scale	159
Archiving The Master	161
Delivering The Master To The Replicator	161
Archiving The Project	161
Greg Calbi - Sterling Sound	163
Dave Collins - Dave Collins Mastering	171
David Glasser - Airshow Mastering	177
Gene Grimaldi - Oasis Mastering	181
Bernie Grundman - Grundman Mastering	185
Colin Leonard - SING Mastering	191
Bob Ludwig - Gateway Mastering	197
Glenn Meadows - Mayfield Mastering	204
Doug Sax - The Mastering Lab	209
Glossary	215
About Bobby Owsinski	226
Bobby Owsinski Bibliography	226
Bobby Owsinski Lynda.com Video Courses	227

