SECOND EDITION

COMPANION WEBSITE

ROUTLEDGE

MIC IT!

MICROPHONES, MICROPHONE TECHNIQUES, AND THEIR IMPACT ON THE FINAL MIX

A Focal Press Book

IAN CORBETT

Mic It!

Capture great sound in the first place and spend less time "fixing it in the mix" with Ian Corbett's *Mic It!* With this updated and expanded second edition, you'll quickly understand essential audio concepts as they relate to microphones and mic techniques and learn how to apply them to your recording situation. *Mic It!* gives you the background to explore, discover, and design your own solutions, enabling you to record great source tracks that can be developed into anything from ultra-clean mixes to massive, organic soundscapes.

Beginning with essential audio theory and a discussion of the desirable characteristics of "good sound," *Mic It!* covers microphones, mono and stereo mic techniques, the effect of the recording space or room, and large classical and jazz ensemble recording. This second edition also features new chapters on immersive audio, immersive recording concepts, drum tuning, and recording techniques for audio for video. *Mic It!* provides in-depth information on how different mic techniques can be used, modified, and fine-tuned to capture not only the best sound, but *the best sound for the mix*, as well as how to approach and set up the recording session, prepare for mixing, and avoid common recording and mixing mistakes.

- Train your ears with practical audio examples on the companion website.
- Develop and test your knowledge as you learn, with concise, applicable exercises and examples that cover the concepts presented.
- Record the best sound possible in any situation with *Mic It!*

Corbett's expert advice ranges from vital knowledge no novice should be without, to advanced techniques that more experienced engineers can explore to benefit and vary the sound of their recordings. Whether you only ever buy one microphone, are equipping a studio on a budget, or have a vast selection of great mics to use, with *Mic It!* you'll learn how to make the most of the tools you have.

Ian Corbett is the Coordinator of Audio Engineering at Kansas City Kansas Community College. He also owns and operates *off-beat-open-hats LLC* recording and sound reinforcement. He is a frequent presenter at conferences and universities around the world, and currently a regional Vice-President of the Audio Engineering Society. He has experience in a wide variety of audio fields, including location recording, sound reinforcement, studio recording, theatre, television and radio.

Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

Mic It!

Microphones, Microphone Techniques, and Their Impact on the Final Mix

Second Edition

Ian Corbett

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

Second edition published 2021
by Routledge
52 Vanderbilt Avenue, New York, NY 10017

and by Routledge
2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

First edition published by Focal Press 2015

Library of Congress Cataloging-in-Publication Data
Names: Corbett, Ian, 1972- author.
Title: Mic it!: microphones, microphone techniques, and their impact on the final mix / Ian Corbett.
Description: Second edition. | New York, NY: Routledge, 2021. | Includes index.
Identifiers: LCCN 2020031529 (print) | LCCN 2020031530 (ebook) | ISBN 9780367470364 (pbk) | ISBN 9780367470449 (hbk) | ISBN 9781003033042 (ebk)
Subjects: LCSH: Microphone. | Sound–Recording and reproducing. | BISAC: MUSIC / Recording & Reproduction.
Classification: LCC TK6478 .C67 2021 (print) | LCC TK6478 (ebook) | DDC 621.389/30284–dc23
LC record available at https://lccn.loc.gov/2020031529
LC ebook record available at https://lccn.loc.gov/2020031530

ISBN: 978-0-367-47044-9 (hbk)
ISBN: 978-0-367-47036-4 (pbk)
ISBN: 978-1-003-03304-2 (ebk)

Typeset in Optima
by Deanta Global Publishing Services, Chennai, India

Visit the companion website: www.routledge.com/cw/corbett

Contents

About the Author

Ian Corbett is the Coordinator and Professor of Audio Engineering and Music Technology at Kansas City Kansas Community College. He also owns and operates "off-beat-open-hats LLC," providing live sound, recording, and consulting services to clients in the Kansas City area. He has experience in a wide variety of audio industry fields, including location concert recording, live sound reinforcement, studio recording, theatre sound, television and radio, and forensic audio evidence recording. As the co-founder and technical director of Electronic Music Midwest he enjoys multichannel sound and presenting immersive concert experiences. A highly active Audio Engineering Society member and officer, Ian has mentored, presented, and served on panels at local, regional, national, and international AES events, and been an invited presenter at many other audio events in the USA and Europe. He frequently serves as an external examiner, program evaluator, and competition judge in the USA and Europe. He has also written articles for *Sound On Sound* magazine. Ian holds a Doctor of Musical Arts degree in composition, and can occasionally be found playing saxophone in Kansas City's jazz clubs and restaurants.

For more information, please see:
www.offbeatopenhats.audio

Acknowledgments

Thank you to everybody who helped throughout the various stages of production of this project, from proposal to final copy edit, including: Bob Beck, Jason Bolte, Mark Drews, Genelec, Dave Greenspan, Dr. Eric Honour, Kerry-Anne Kubisa, Dave Maclaughlin, Wes Maebe, Aki Mäkivirta, Giuliano Mingucci, Shannon Neill, Hannah Rowe, Kyle Snyder, Chris Townsend, Megan Turner, Anaïs Wheeler, and Meagan White. Additionally, Eric Honour and Dave Maclaughlin, thank you for your invaluable assistance as Technical Editors.

Thank you for letting me take and use photographs of you, your equipment, musical instruments, or your studio facilities, or supplying photographs for me to use: Bob Beck, Patrick Conway, Ian Dobyns, Chris Hazelton, Levi Holgerson, Dionne Jeroue, Kansas City Kansas Community College, Jennifer Kim, Margaret Kocher, Ken Lovern, Jim Mair, Armond McDonald, Erin McGrane, Chandler Meierarend, Giuliano Mingucci, B. Scott Nelson, Neumann-Berlin, Jim Schrader, Sky Recording (Kansas City), Tuesday Stewart, Travis Stewart, Andrew Tremain, and Elizabeth Vallé.

Thank you to the professionals who graciously let me interview them so their valuable insights could fill the last chapter: Lenise Bent, David V. R. Bowles, Joel Hamilton, Kerry-Anne Kubisa, Wes Maebe, Matt Ross-Spang, and Mark Rubel, Catherine Vericolli, Paul Womack.

Thank you to the musicians and engineers who generously provided their talents for the accompanying web site audio examples:

- Carswell & Hope: Nick Carswell (guitar and vocals), Dan Hines (bass), Austin Quick (keys), Jason Slote (drums). *Hunger* from the album *Carswell & Hope*. Written, recorded and produced by Carswell and Hope. Audio Example 2.21 mixed by Jason Slote and Nick Carswell (http://carswellandhope.com). Audio Examples 1.5, 1.6, 1.7, 1.8, 2.20, 6.12, 6.13, 6.15, and 6.16, remixed by Ian Corbett.
- Bryan Corbett (trumpets, flugelhorn, flute, recording and mixing) and Chris Dodd (bass, drums, piano, keyboards, recording and mixing): *Funk 75*, written by Chris Dodd and Bryan Corbett, from the album *Green*. Audio Examples 2.7, 2.8, and 2.9. (http://bryancorbett.co.uk)
- Ian Corbett (saxophones), Audio Examples 2.10, 6.9–6.11.

- Ian Corbett (all instruments and drum programming), David Garibaldi (Tower of Funk drum samples), and Sarah Williams (vocals): *Smile!*, Audio Examples 1.1–1.4, 6.14, and *The Closer I Get To You*, Audio Examples 2.17–2.19, written by Ian Corbett, from the album *If That's You*.
- Aaron Crawford (drums). Audio Examples 2.15, 2.16, 3.7, 3.8, 4.6–4.9, 7.1–7.8, 9.6–9.9, 10.1, 10.2, 12.1–12.8, 12.15–12.19, 16.4–16.7, 16.18–16.21.
- Ian Dobyns (acoustic and electric guitars). Audio Examples 2.11, 2.12, 4.10, 4.11, 7.9–7.10, 9.1–9.5, 14.1–14.4, 14.11–14.14, 16.8–16.15.
- Eboni Fondren (vocals and narrations). Audio Examples 1.1–1.8, 1.13, 2.13, 2.14, 2.22–2.26, 3.3–3.6, 10.3, 10.4, 16.1–16.3, 16.13.
- Kelley Gant (vocals). Audio Examples 4.12–4.14, 11.1–11.4, 16.16, 16.17.
- Levi Holgerson (acoustic and electric bass). Audio Examples 14.5–14.10.
- Dr. Nahyun Kim (piano). Audio Examples 6.1–6.8, 14.15–14.21.
- Giuliano Mingucci (drums). Audio Examples 12.9–12.14, 13.1–13.12.
- Now Now Sleepyhead: Aaron Crawford (drums, percussion, programming, additional vocals), Phil Park (vocals, guitar, bass, and additional programming), Megan Zander (additional vocals). *Influenza* from the album *The Violator*. Written by Phil Park and Aaron Crawford. Produced, engineered, mixed, and mastered by Aaron Crawford. Audio Examples 2.1–2.6. (www.facebook.com/nownowsleepyhead)
- John Stafford (voice). Audio Examples 20.1–20.5.
- Steven Sterner (vocals and narration). Audio Examples 1.13, 2.10, 4.1–4.5.
- "The Standard" Vocal Jazz Ensemble: Natalie Bennett, Stormy Borsella, Deandre Clark, Laquita English, Allison Her, Lauren Irving, Andrew Roberson, Andrea Rodriguez, Shelby Stephenson, Morgan J. Tinsley. John Stafford, Director, Kansas City Kansas Community College. Audio Examples 11.5 and 11.6.
- Mike "Topspin" Tosspon (vocals). Audio Examples 3.1 and 3.2.

Last but not least, lots of love to my family in the UK, and friends and "family" around the world.

To anybody I've omitted, I apologize!

1 Audio Basics

In This Chapter:

1.1 It's Not Always About the Gear!
1.2 What is Sound?
1.3 The Decibel (dB)
1.4 Power Relationships
1.5 Decibel Scales
1.6 Dynamic Range
1.7 Signal-to-Noise Ratio
1.8 Frequency vs Pitch
1.9 Frequency Response
1.10 Waveforms, Fundamentals, and Harmonics
1.11 Wavelength, Velocity, Phase
1.12 Amplitude Summation and Comb Filtering
1.13 Human Hearing
1.14 Signal Flow and Audio Level Standards
1.15 Gain Structure and Recording Levels
1.16 Analog Audio Connectors
1.17 Digital Audio Connectors
1.18 Digital Audio Basics

An engineer who knows how to use a mic is more important than the mic.
–Wes Maebe, Engineer/Producer: Sonic Cuisine, RAK Studios, London, UK

The best microphone I've used in my life was the one that was up when the artist did something great.
–Joel Hamilton, Producer, Engineer, Co-Owner: Studio G Brooklyn, New York, USA

1.1 It's Not Always About the Gear!

Gear, gear, gear! We all want to get our hands on new toys, plug that gear in and make sound… There's nothing wrong with that! But even the best audio equipment available won't produce great sound unless you understand how to use it properly. Lower quality equipment used well, will always sound better than great quality equipment used poorly.

Sound is our artistic medium. As sound engineers, we create sonic artwork from it. The more we understand it, the more we can predict its behavior, and more easily produce better recordings. You'll be much better prepared to get the most from the concepts discussed later in this book, and to produce great recordings, if you have a good understanding of fundamental audio theory and studio basics. *Don't skip ahead!* There might be something in this chapter that will change the way you think about sound, use your equipment, or hear what you're listening to – improving the audio you record and mix.

1.2 What is Sound?

Examples of objects that produce sound include the strings on a guitar or violin, the reed in a wind instrument mouthpiece, a trumpet player's lips, the head on a drum, and a loudspeaker cone or headphone driver. All of these sources have one thing in common – they *vibrate*, creating variations in air pressure, which become *sound waves*. Sound does also travel through other mediums, such as water and solid objects – but seeing as air is the medium that usually surrounds us, we'll concentrate on that!

Figure 1.1 shows a simplified, illustrative picture of a vibrating guitar string. The string is anchored at both ends, and stretched so it is taut. When it is plucked, bowed, or struck, it is set in motion and vibrates. During this motion it moves from its point of rest (the position it naturally returns to when not in motion), labeled **A**, and out to an extreme, labeled **B**. As it approaches **B**, the tension in the string increases until it is not able to move any further, and it rebounds back in the opposite direction, through **A** to the opposite extreme, **C**. Tension builds as it moves towards **C**, and causes the string to reverse its direction again, so it moves back through **A** towards **B**. A little energy is lost with each consecutive change in direction, so the string gradually moves less and less (getting quieter and quieter) until it is stationary and silent back at its point of rest, **A**.

As the string moves from **A** to **B**, it squashes the air molecules to the left of the string closer together. This increases the air pressure in that spot, causing a *compression*. This compression then travels outwards from this source at the speed of sound. The air molecules themselves do not move far – the compressed group of molecules bump into the adjacent molecules, which bump into the next set, passing the compression on.

As the compression travels outwards, the string moves back through **A**, where normal atmospheric air pressure is restored. During the subsequent motion from **A** to **C**, the air molecules to the left of the string are drawn further apart, to fill the space where the string

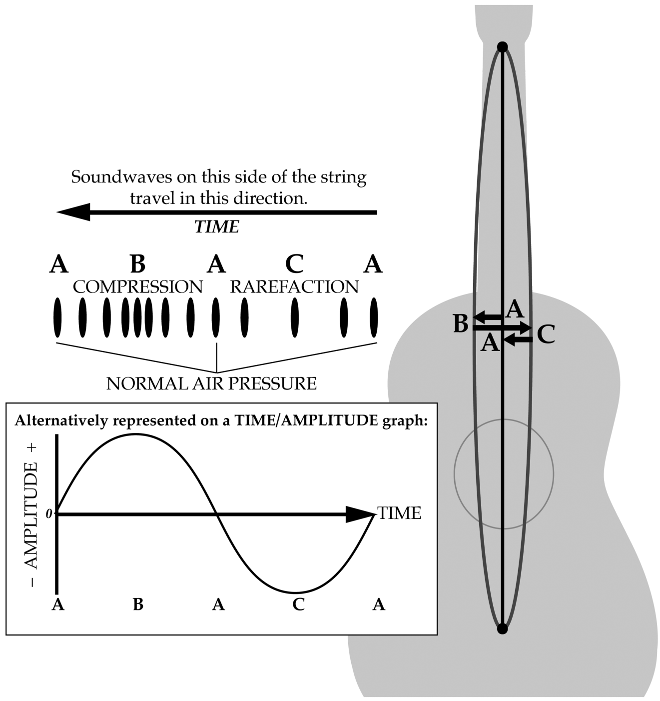

Figure 1.1 The motion of a vibrating string, and a resulting sound wave.

used to be. This causes a decrease in air pressure to the left of the string, called a *rarefaction*. This rarefaction travels outwards, following the previous compression. As the string returns to **A**, normal air pressure is once again restored behind the rarefaction. This continuing motion propagates the alternating compressions and rarefactions of *sound waves* behind each other. Greater motion and displacements from the point of rest create greater variations in air pressure and louder sounds.

THE SPEED OF SOUND

The speed of sound is commonly quoted as being around 344 meters (1130 ft) *per second* – for dry air, at 20°C (68°F), at sea level. As general principles, the speed of sound:

- Increases as air temperature rises, and decreases as air temperature falls.
- Decreases slightly as altitude increases (though this has more to do with the altitude-associated temperature decrease than the change in air pressure).
- Increases a little as humidity rises.

Figure 1.1 also shows two ways of graphically representing these sound waves. The first is a series of blobs representing the relative spacing of the air molecules – the air pressure.

The second is a time/amplitude graph. This graph may look familiar. It is the shape of a *sine wave* – a pure tone made up of one single frequency. An actual vibrating string produces a much more complex harmonic waveform than a sine wave, as shown in **Figure 1.2**.

Time/amplitude graphs are based on measurements taken at a single point in space, over a duration of time. This is similar in principle to what a microphone does – a mic is positioned at a single point in space and takes measurements of air pressure over time, capturing frequency and amplitude information.

1.3 The Decibel (dB)

The *amplitude*, or "amount of" a sound's energy is measured in decibels. Meters, faders, and the scales on many equipment knobs are in dBs, so it is vitally important to understand the concept in order to use audio equipment correctly. The dB is a unit of convenience. By itself it *only* implies that we have reduced a much larger range of numbers into a smaller, more manageable range of numbers in the form of a dB scale of some sort.

For example, a more familiar unit might be the *watt* – used for measuring *power*. Power is a measure of the energy consumption or energy transfer of a device. The human ear is capable of detecting sounds with powers from 0.000000000001 W to almost 100 W. *Do not confuse these figures with amplifier power!* These numbers represent the energy of the sound waves in the air (that you're actually listening to) and *not* the electrical amplifier power required to propagate sound from a loudspeaker. This huge range of numbers is unusable and impractical!

The dB is a logarithmic unit that folds this large power range down into a smaller range of numbers that represent power ratios. *The dB unit always requires clarifying with a second unit – in the form "dB something."*

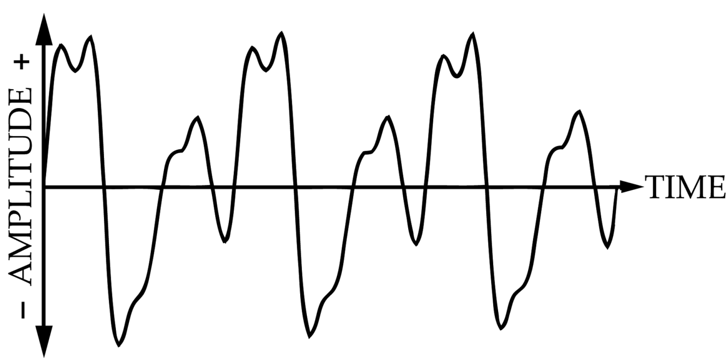

Figure 1.2 The waveform produced by an actual guitar string, showing the overall motion of the string produced by the interaction of all the frequencies present.

1.4 Power Relationships

Sound Pressure Level, or *SPL*, is a measure of how loud a sound actually is, where you're listening to it or measuring it.

- Doubling the power of a sound causes a 3 dB increase in SPL, which we perceive as "slightly louder."
- Halving the power of a sound causes a 3 dB decrease in SPL, which we perceive as "slightly quieter."
- Doubling the electrical voltage representing a sound causes a 6 dB increase – however, this does not double the perceived loudness of the sound.
- In order to perceive something as being "twice as loud," a 9 or 10 dB increase is necessary.
- In order to create a 9 or 10 dB increase, the power has to be doubled (= +3 dB), doubled again (+3 dB = 6 dB total increase), and then doubled again (+3 dB = 9 dB total increase).

These same calculations and observations apply to *any control scaled in dB* on *any piece of audio equipment*.

KNOW YOUR POWER RELATIONSHIPS!

- Turning something up by 3 dB makes it "slightly louder."
- Turning something down by 3 dB makes it "slightly quieter."
- Turning something up by 9 dB makes it "twice as loud."
- Turning something down by 9 dB makes it "half as loud."

Knowing just these simple facts, you're already a more informed, educated, and better audio engineer! If you're mixing and thinking "the vocal is *slightly* too quiet," you now know what you need to do in order to achieve the right correction, and anticipate the results of that correction, before moving any controls:

- If the vocal is "slightly" too quiet, increasing the vocal fader level by +3 dB will make it "slightly louder" – and should come close to correcting the problem.
- If you're thinking, "the sax needs to be *half* as loud," then you should be able to make an informed judgment, reduce the sax fader level by about –9 dB, and come close to correcting the problem.

This process is much more professional than grabbing the fader and randomly moving it until you stumble across the correct level!

AUDIO EXAMPLES
Can be found on the companion website
(The companion website address can be found on the back cover or page iv).

Amplitude Changes

Example 1.1: A music excerpt, and then the same excerpt +3 dB (slightly louder).
Example 1.2: A music excerpt, and then the same excerpt –3 dB (slightly quieter).
Example 1.3: A music excerpt, and then the same excerpt +9 dB (twice as loud).
Example 1.4: A music excerpt, and then the same excerpt –9 dB (half as loud).

Contextual Mix Changes

Example 1.5: A mix example with the lead vocal +3 dB (slightly louder) in the second excerpt.
Example 1.6: A mix example with the lead vocal –3 dB (slightly quieter) in the second excerpt.
Example 1.7: A mix example with the bass +9 dB (twice as loud) in the second excerpt.
Example 1.8: A mix example with the bass –9 dB (half as loud) in the second excerpt.

1.5 Decibel Scales

dB SPL

dB SPL is a measure of how loud a sound actually is in an environment, at a specific point in space. It is measured with an SPL meter. The most commonly encountered SPL scale is 0 to 140 dB SPL – which is based on the human ear's threshold of hearing (0 dB SPL) and threshold of pain (130 dB SPL). The SPL scale does in fact go up to about 194 dB SPL – which is the point at which the air cannot handle the amount of sonic energy, and distorts the sound.

dBV

dBV (*decibel volt*) scales are used on analog audio equipment, or digital devices and software trying to emulate an analog experience. They are a measure of the electrical voltage representing the sound wave. Bargraph meters commonly use dBV scales. Their range is usually –∞ dBV (or some negative value) through 0 dBV, to approximately +15 or +18 dBV. **Figure 1.3** shows some typical dBV meters.

On a dBV meter, "0" is considered *nominal* – the optimum level at which the equipment is designed to operate most linearly. However, analog equipment can be pushed and operated above this quite happily, hence the typical +15 or +18 dBV range. dBV meters are

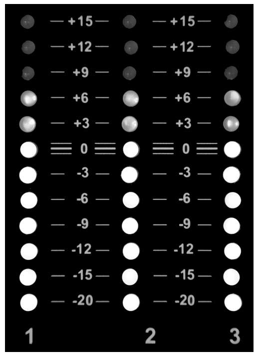

Figure 1.3 Some dBV meters, as found on many analog devices.

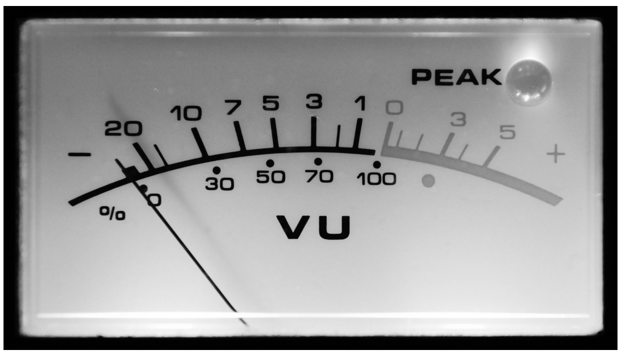

Figure 1.4 A VU meter with a dBu scale, as found on many analog devices.

peak meters – they show the actual maximum voltage peaks of the electricity representing the sound waves. It is normal and desirable to have peak levels in the "+" ranges – but *do not* light any "overload" lights.

Some meters can be switched between peak and RMS (root mean squared) and other non-peak averaging modes.

dBu

Volume Unit meters, as shown in **Figure 1.4**, use the dBu scale. Their mechanical needles (or software simulations) cannot react to instantaneous peak levels, so they do not show any instantaneous or short-term peaks less than about 300 milliseconds long. Instead, they show more of an averaged amplitude measurement – an indication of how we perceive the loudness of the sound. Many VU meters have a peak LED that lights when finite peak levels are exceeded. A VU meter's usual range is from −30 or −20 dBu or so, through 0 dBu, and up to approximately +6 dBu. VU meters are appropriate on analog devices (analog mixing consoles, tape machines, etc.), where "0" can be happily exceeded, and are useful because they indicate the sound's average intensity. VU meters are not found on digital devices because it is imperative to know the actual peak level in order to avoid digital clipping (although software meters can sometimes be switched to averaging modes).

When using a VU meter it is important to leave enough headroom for instantaneous peak levels without lighting the peak LED. Desirable levels on a VU meter are generally lower than on a peak dBV meter, particularly for percussive or high frequency sounds. When recording percussive or plucked instrument sounds using a digital recording system via a preamp or mixer with VU meters, it is essential to make sure the VU meter's peak indicator (if it has one) does not light, and that the dBFS meters in the DAW (Digital Audio Workstation) tracks are not showing digital clipping.

dBFS

Most digital audio devices have *dBFS* meters – *decibels* (relative to) *full scale* – which are instantaneous peak meters. A typical meter is shown in the left of **Figure 1.5**. "Full scale" is the maximum amplitude a digital device can handle without horrible distortion, and is labeled 0 dBFS. A dBFS meter's range is from –∞ dB (or some detectable negative value) up to 0 dBFS. The only thing above 0 dBFS is the "clip" or "over" indicator.

Digital devices cannot be overloaded in the same way as analog devices. There are no potentially desirable distortions or non-linearities above zero, as there might be on analog equipment – only ugly, and potentially equipment damaging, square wave digital nastiness! When using digital devices it is essential to leave enough *headroom* ("spare" level) so the "clip" or "over" indicators do not light if the signal's amplitude increases unexpectedly.

0 dBV on an analog device with a maximum dBV level of +15 correlates to –15 dBFS on a digital device. The right of **Figure 1.5** shows a hybrid meter found on a modern analog mixer that would, today, typically be connected to a digital recording system.

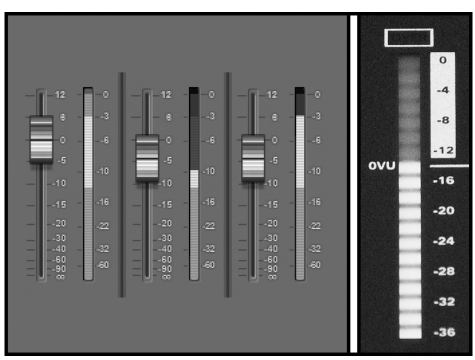

Figure 1.5 **Left:** A software dBFS meter. **Right:** A modern hybrid meter with a dBFS scale on the right, and 0 dBu on the left, correlating to –14 dBFS.

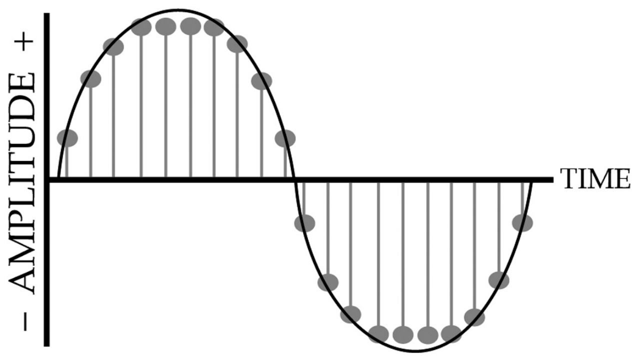

Figure 1.6 The grey dots represent dBFS sample values, and the black curve shows the slightly higher amplitude dBTP waveform created by reconstruction during digital to analog conversion on playback, or when encoding to lossy dissemination formats.

dBTP

Decibels True Peak is an instantaneous peak measurement that accounts for the inter-sample peaks created by digital to analog conversion on playback, or when encoding to lossy dissemination formats (mp3, AAC, Ogg-Vorbis etc.). If adjacent samples are similar in value, the smoothed and reconstructed analog waveform, or lossy encoding, can have a momentary peak level that is higher than indicated on dBFS meters, and higher than 0 dBFS, which can cause clipping and distortion. dBTP is a numerical display that holds the true peak value until reset.

LUFS

Loudness Units Full Scale metering represents perceived loudness in a way that correlates to human perception. LUFS measurements de-emphasize the effect of low frequencies in the level shown. Actual numbers, LUFS meters are easier to accurately read than needles or bargraphs. One LU equates to one dB. LUFS meters show measurements for various time durations:

- *Momentary Loudness* – a measure of the loudness averaged over 400 ms.
- *Short Term Loudness* – a measure averaged over one to several seconds.
- *Integrated Loudness* – a long term measurement averaged over an entire song or file, or all the material played through a real-time meter since it was last reset.
- *True Peak* – an instantaneous measurement, as previously discussed.
- *Loudness Range* – a measure of the loudness variation (the difference between the louder and quieter sections) calculated for an entire song, file, or since a real-time meter was reset.

LUFS levels are important to music mixing and mastering engineers (and anybody producing audio assets for the game industry). Most music streaming services penalize any material submitted to them with *integrated LUFS levels* that are too high – they turn them down to play back at a lower level. *There is no benefit to mixing and mastering a single or album so it is hyper-compressed and super-loud!*

TARGET LEVELS AND LOUDNESS PENALTIES

Because integrated LUFS values now impact music streaming, it is important that the mixing engineer knows whether or not they are not exceeding the target LUFS playback level for streaming services. If an unmastered mix is *at* or *exceeding* that target level, the mastering engineer is not left much room (if any) to do their dynamic magic without the mix suffering a loudness penalty (being turned down) on playback.

The target level of different streaming services ranges between –13 and –16 LUFS, with –14 LUFS being the most common. So –14 LUFS can be a good potential target in so much as if the track is mastered to that level, it will not be turned down when streamed. But ultimately, the mix and mastering engineers should do what the song or material needs, remembering that if it is louder than –14 LUFS it will be turned down by the streaming service.

1.6 Dynamic Range

Measured in dB, audio dynamic range is the difference between the lowest and highest signal levels an audio device (or system) can record, store, or reproduce. Larger (or wider) dynamic ranges are better. The lowest extreme of this range, considered 0 dB, is the level below which the system is unable to resolve details. This limit is usually set by the naturally occurring hiss and noise in a system (its *noise floor*) overpowering and masking low level details, or in a microphone it could be a lack of low-level sensitivity and inability to respond to low levels. The high level limit of a system's dynamic range is the maximum level, relative to 0 dB, that the system can handle before a certain amount of distortion of the audio signal occurs.

Analog devices (mixers, tape and tape machines, outboard gear, etc.) and the analog parts of digital devices (preamps and convertors in interfaces for example) all exhibit hiss and noise that dictate their lower performance limit. Their highest level is limited by the distortions created when their electronic components are gently overloaded with voltage, or the magnetic particles on tape are over-saturated with magnetism. Gentle, slight overloading of analog devices often causes desirable sounding, "pleasing" distortions to occur.

Dynamic range is usually quoted as being "*for x%* THD" (*Total Harmonic Distortion*). THD is a measure of the amplitude of the additional harmonics of the overload distortion, compared to the original waveform.

Digital devices do not exhibit gentle and "pleasing" distortions if their maximum limit is exceeded. Digital distortion is a square wave type of distortion, adding sharp corners to the waveform. It is neither pleasing to the ear or good for loudspeakers! **Figure 1.7** shows an input waveform, and how analog and digital systems would change it if it exceeded their maximums slightly.

1.7 Signal-To-Noise Ratio

SNR, or *S/N*, and dynamic range are quite similar. Whereas dynamic range is a measure of the difference between the loudest and quietest perceivable signals in a sound (or that a system can resolve), *signal-to-noise ratio* is a measure of the difference between a specific reference level (not the maximum level) and the noise floor. As with dynamic range, the higher the quoted SNR specification, the technically better it is.

1.8 Frequency vs Pitch

In addition to amplitude, another way we categorize or identify musical sound is by its pitch or frequency. But frequency and pitch are not the same thing.

So what is *frequency*? Earlier in this chapter the motion of a vibrating string was analyzed. A vibrating sound source's motion from a starting point, through both extremes, and back to the same starting point and direction, is known as one *cycle*. The typical analysis of one cycle is shown in **Figure 1.8** – this is the motion **A-B-A-C-A** of the vibrating string in **Figure 1.1**. Additionally, this figure shows cycles that start from other points in the waveform.

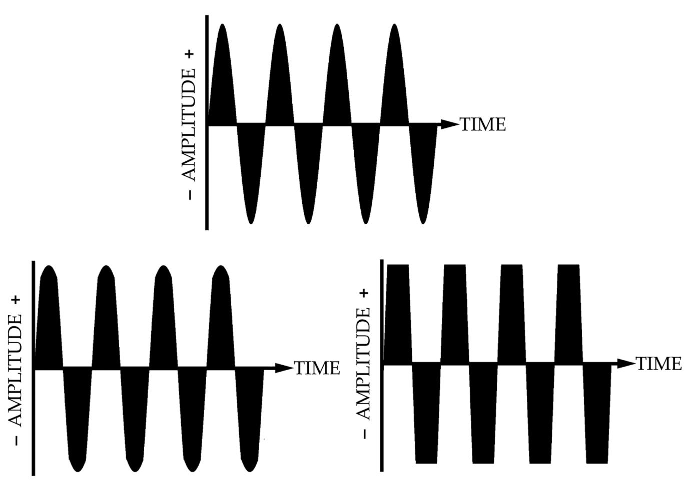

Figure 1.7 **Top:** An undistorted input waveform. **Bottom left:** Gentle saturation distortion characteristics of an analog device. **Bottom right:** Harsh, undesirable, and potentially damaging digital distortion (clipping).

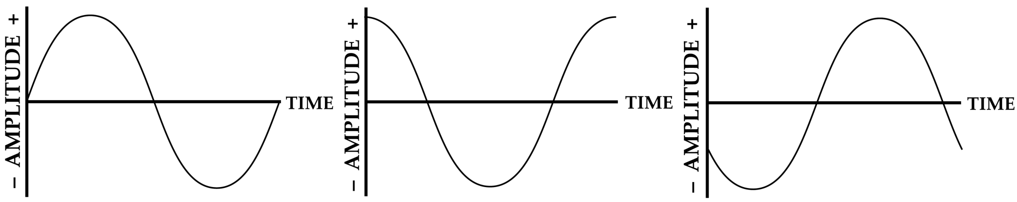

Figure 1.8 **Left:** A single cycle identified starting from the vibrating source's point of rest. **Center:** A single cycle analyzed starting at one extreme. **Right:** A single cycle analyzed starting mid-cycle.

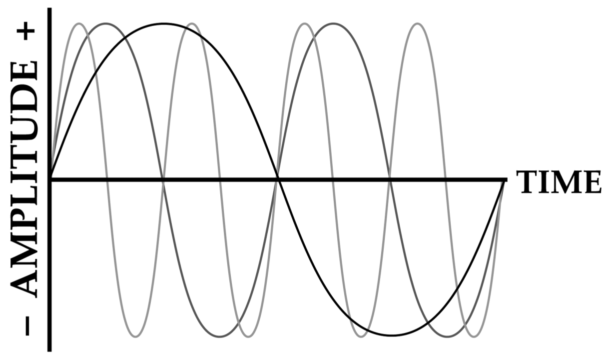

Figure 1.9 Three octave frequencies superimposed on each other show how zero crossing points regularly line up at the beginning, middle, and end of the lowest frequency's cycle, making the combined sound pleasingly powerful and solid.

Frequency is measured in *Hertz*. The abbreviation is *Hz*. The definition of 1 Hz is one cycle per second – so imagine the string pictured earlier taking one second to move **A-B-A-C-A**.

In music, the interval of an octave is very special – octaves are "the same note," but in a higher or lower register. Octaves are created by doubling or halving frequency. The standard Western tuning note **A** is 440 Hz, the octave above is 880 Hz, and the octave below is 220 Hz. **Figure 1.9** shows the integer relationships between octave frequencies that cause their waveforms to interact with each other in a way we find complementary and pleasing.

Frequency is an *exponential* scale – it is non-linear. The difference between tuning note **A** and the octave below it is 220 Hz. The difference between tuning note **A** and the octave above it is 440 Hz. The difference between, and the frequency of each higher octave, increases exponentially – continuously doubling. The frequency, and difference between each lower octave is continuously halving – or decreasing exponentially, never reaching zero.

Pitch is our perception of frequency, and that perception changes with amplitude:

- A loud tone sounds flatter in pitch than a quiet tone of the same frequency.
- A quiet tone sounds sharper in pitch than a loud tone of the same frequency.

1.9 Frequency Response

The frequency response of an audio device or system is the range of frequencies it can efficiently capture, store, or reproduce. The standard range of audio frequencies, based on the limitations of human hearing, is from 20 Hz to 20 kHz (kilohertz, 20,000 Hz) – but many audio devices have frequency responses that extend beyond that range.

1.10 Waveforms, Fundamentals, and Harmonics

The amplitude of the compressions and rarefactions of a sound wave plotted along a time axis produce a graphical *waveform*. **Figure 1.10** shows several waveforms:

- The **left** graph shows a loud, low frequency sine wave.
- The **center** graph shows a quieter, higher frequency sine wave. The positive and negative deflections are smaller, representing a lower amplitude. There are more cycles in the same amount of time, representing its higher frequency.
- The **right** graph shows the waveform of a musical instrument playing a note. The larger (and louder) overall cyclic shape represents the *fundamental* frequency – the lowest frequency in the sound, which gives the note its pitch. (The fundamental is also called the *first harmonic*.) The multiple faster cycling, (quieter) higher frequencies superimposed on the fundamental are the *harmonics*, which give the sound its *timbre* or *tone color*.

The frequencies at which harmonics occur above a given fundamental frequency are dictated by the *harmonic series* – a naturally occurring phenomenon that applies to just about all sounds except metallic bells, and synthesizers and other electronically generated sounds. Harmonics in the harmonic series occur at integer multiples of the fundamental frequency (f):

$$\text{Fundamental or first harmonic } (f) = 1 \times f \left(\text{for example, 100 Hz}\right)$$
$$\text{Second harmonic} = 2 \times f \left(\text{for example, 200 Hz}\right)$$
$$\text{Third harmonic} = 3 \times f \left(\text{for example, 300 Hz}\right)$$
$$\text{Fourth harmonic} = 4 \times f \left(\text{for example, 400 Hz}\right)$$

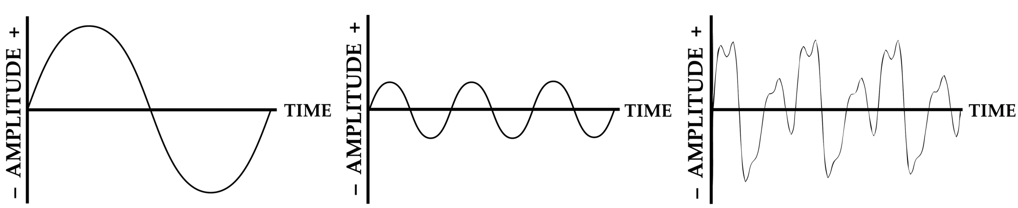

Figure 1.10 The different waveforms discussed in the main text.

Figure 1.11 shows the oscillations of a 100 Hz fundamental, and the second, third, and fourth harmonics. In most naturally occurring sounds, the higher the harmonic, the lower its relative amplitude. Not all sounds contain all harmonics – some do have odd and even harmonics present, others might have more odd harmonics and less even harmonics.

The musical interval between the fundamental and the second harmonic is an octave – the frequency is doubled. The musical interval between the second and third harmonics is approximately a perfect fifth. The musical interval between the third and fourth harmonics is approximately a perfect fourth, creating an octave from the second harmonic. From there the harmonics progress in smaller musical intervals – approximately major third, minor third, then seconds, then microtonal intervals. The word "approximate" is used because the mathematical relationships and resulting harmonic frequencies do not exactly match the musical note tunings of the equal temperament tuning system used by most Western instruments.

Partial is a term used to describe any harmonic in a sound, whether it conforms to the harmonic series or not. An *inharmonic partial* is a harmonic that is not an integer multiple of the fundamental, and therefore is not part of the harmonic series.

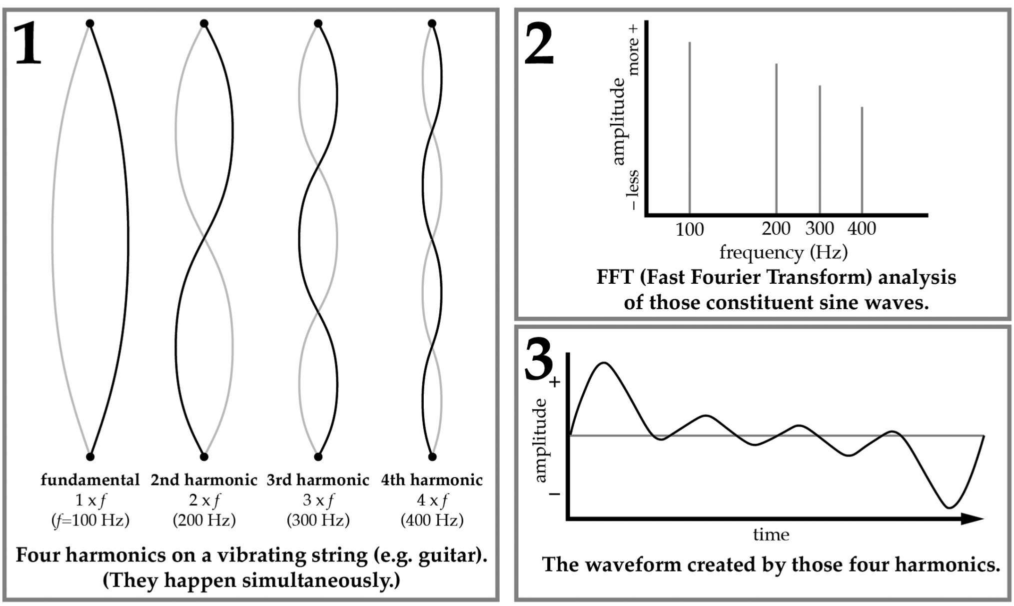

Figure 1.11 Some harmonic structure of a hypothetical sound shown as **1:** Constituent harmonic vibrations on a string. **2:** An RTA (Real Time Analyzer) or FFT (fast Fourier transform) analysis of the constituent sine wave harmonic structure. **3:** The waveform created by that harmonic structure.

AUDIO EXAMPLES
Can be found on the companion website

Simple Waveforms

Example 1.9: A pure (single frequency) sine wave.
Example 1.10: A square wave, produced by the addition of odd harmonics.
Example 1.11: A triangle wave, produced by the addition of odd and even harmonics.

1.11 Wavelength, Velocity, Phase

Wavelength and Velocity

Sound waves travel at a finite speed, which for air is quoted as being 344 m/s or 1130 ft/s. A high frequency cycle takes a shorter time to be propagated than a low frequency cycle – so the start of its cycle doesn't travel as far from the sound source before the next cycle is propagated behind it. Therefore a low frequency cycle takes up more physical space in the air than a high frequency cycle.

Figure 1.12 shows a loudspeaker cone generating two frequencies. The top one is a lower frequency; the middle one is a higher frequency. The wavelengths of one cycle of each can be calculated using the following equations:

$$\lambda = \frac{v}{f}$$

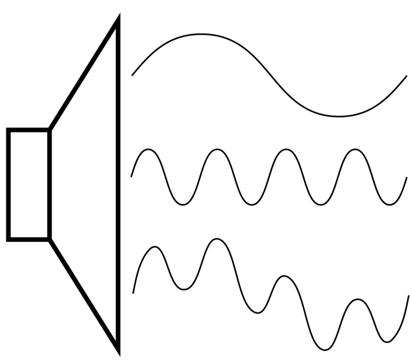

Figure 1.12 A loudspeaker cone generating different frequencies. **Top:** The low frequency. **Middle:** The high frequency. **Bottom:** The actual waveform produced by summing the two frequencies together.

$$\text{Wavelength in meters} = \frac{\text{speed of sound in meters per second}}{\text{frequency in Hertz}}$$

$$\text{Wavelength in feet} = \frac{\text{speed of sound in feet per second}}{\text{frequency in Hertz}}$$

If the low frequency is 100 Hz, its wavelength is 344 ÷ 100 = 3.44 meters (344 cm). In feet, 1130 ÷ 100 = 11.3 ft (11 ft 3.6 in).

If the high frequency is 500 Hz, its wavelength is 344 ÷ 500 = 0.688 meters (68.8 cm). In feet, 1130 ÷ 500 = 2.23 ft (2 ft 3.1in)

The bottom waveform in **Figure 1.12** is the result of the loudspeaker generating both frequencies simultaneously – the shorter wavelengths of the higher harmonic frequency are superimposed on the longer wavelength of the lower fundamental frequency.

Phase

When looking at a waveform, the sine wave in **Figure 1.13** for example, it is impossible to describe any non-peak or non-zero-crossing position along the waveform accurately using non-technical language. Imagine trying to describe the position of the dot on the waveform – "It's almost at the first peak, but not quite…" Anything other than an approximation of its position is impossible to describe. An understanding of phase is necessary to be able to overcome this issue.

Waveform cycles can be divided into 360 degrees of phase:

- 0° and 360° are the same – the beginning or end of the cycle and at the zero-crossing point.
- 90° and 270° are the positive and negative peaks, respectively.
- 180° is the mid-cycle zero-crossing point.

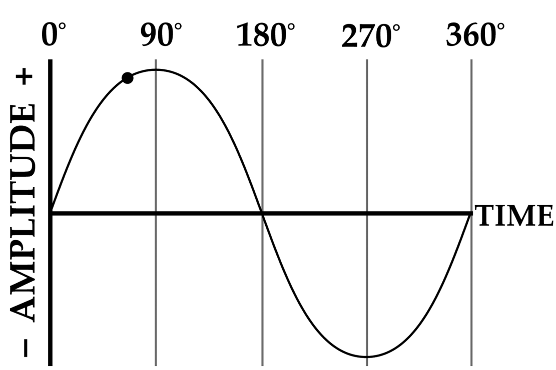

Figure 1.13 A sine wave, with degrees of "phase" marked.

Using degrees of phase, it is possible to accurately calculate and express the position of the dot on the waveform in the previous diagram (66°). It is also possible to describe the phase difference between two identical waveforms that differ in time. The difference between the waveforms **A** and **B** in **Figure 1.14** is 90° – they are 90° out of phase. **A** and **C** are 180° out of phase, and **A** and **D** are 270° out of phase.

Phase is frequency dependent. If the waveforms in the figure were replaced with higher or lower frequencies, displaced by the same amount of *time*, the relative phase relationships would not be the same.

1.12 Amplitude Summation and Comb Filtering

Summing together two waveforms can cause:

- Constructive interference – increasing the amplitude of the positive and negative peaks.
- Destructive interference – decreasing the amplitude of the positive and negative peaks.

If two identical and "in phase" waveforms (or the in phase frequency components of a more complex waveform) are acoustically summed together in the air, they will constructively interfere, and become 3 dB louder because the power is doubled.

If two identical and "in phase" waveforms (or the individual in phase frequencies of a complex waveform) are electrically summed together in a mixer or DAW, they will constructively interfere, and become 6 dB louder because the voltage is doubled.

This type of constructive amplitude summation is shown in **Figure 1.15**.

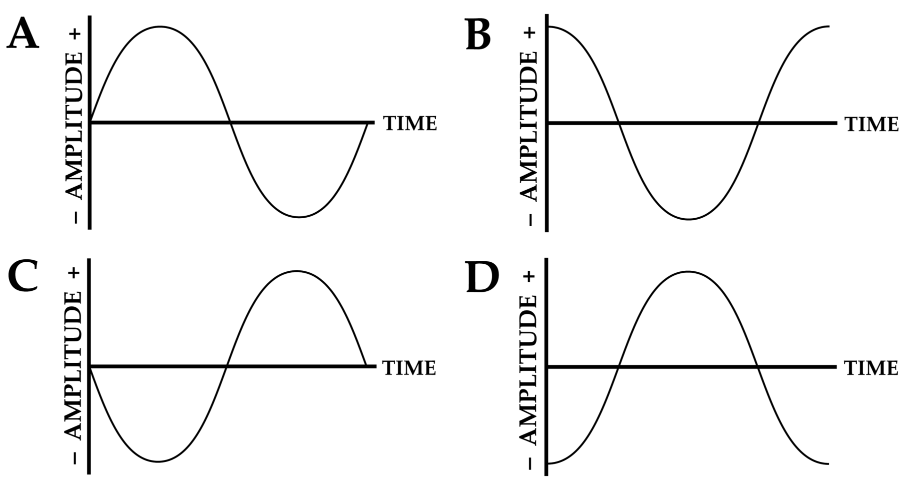

Figure 1.14 The relative phases of two identical waveforms which are displaced in time. **A** is 0°, "in phase." **B** is 90° out of phase compared to **A. C** is 180° out of phase compared to **A. D** is 270° out of phase compared to **A**.

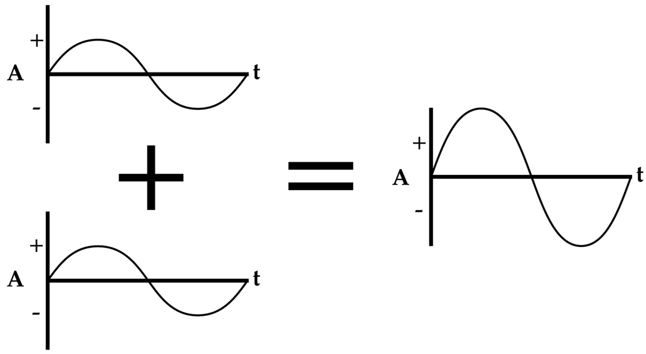

Figure 1.15 Two "in phase" sine waves acoustically sum together constructively, and become 3 dB louder. Two loudspeakers playing the same material would be 3 dB louder than a single loudspeaker.

If two summed waveforms (or individual frequency components of a complex waveform) are 180° out of phase, the positives and negatives happening simultaneously cancel, causing those frequencies to disappear, as in **Figure 1.16**.

If two frequency components have a more "in-between" phase relationship, their behavior is more complex – there will be amplitude increases and decreases alternating throughout the summed waveform, as shown in **Figure 1.17**.

COMB FILTERING

In the real world, we do not record just one or a few sine waves at a time – we record complex waveforms that contain a myriad of frequencies. When two complex waveforms are displaced in time and summed together, different frequency ranges combine with different phase relationships, causing narrow bands of constructive and destructive summation, and the strange, funky, usually undesirable sound of *comb filtering*:

- Some frequencies will be one or more complete cycles out of phase with each other (360° out of phase, or in phase but delayed by an integer number of cycles), and will combine constructively to become louder.
- Other frequencies will be a half cycle out of phase (180° out of phase, or 0.5, 1.5, 2.5 etc. cycles out of phase) and combine destructively to cancel.
- Frequencies with phase relationships between these two extremes partially sum or cancel.

Changing the time displacement between the signals changes the frequencies at which comb filtering occurs.

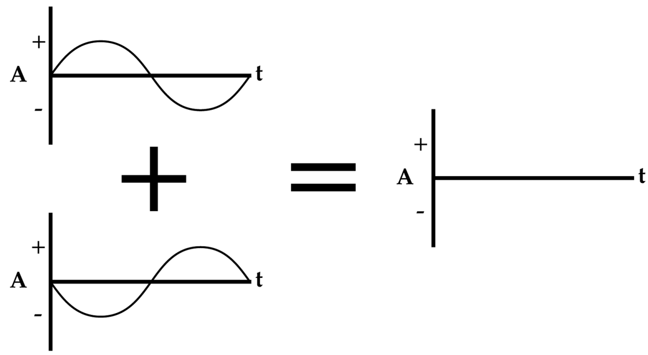

Figure 1.16 Two 180° out of phase sine waves destructively cancel completely.

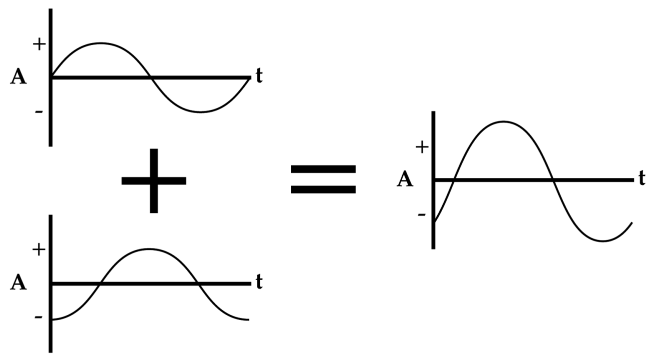

Figure 1.17 Two slightly out of phase waveforms interact partially constructively.

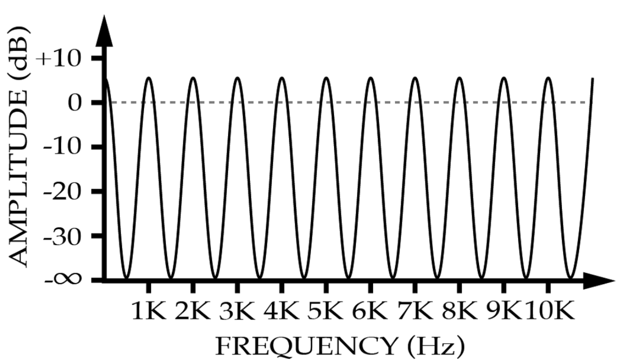

Figure 1.18 Comb filtering – alternating bands of cancellation and summation.

AUDIO EXAMPLES
Can be found on the companion website

Phase Relationships

Example 1.12: Two slightly detuned sine waves cycle in and out of phase as they interact, at times constructively, and other times destructively.

Example 1.13: A vocal excerpt. In the second part of the example, two slightly time delayed versions of the same material produce the frequency dependent on constructive and destructive interference of comb filtering.

1.13 Human Hearing

Figure 1.19 shows a sound wave's compressions and rarefactions travelling through the air to a listener's ear. They are reflected by the pinna of the ear into the ear canal and to the eardrum, a thin membrane of skin-like material.

- If a compression travels towards the eardrum, the pressure in the ear canal becomes greater than normal, causing the eardrum to stretch inwards to equalize the air pressure difference on both sides of the eardrum.
- A rarefaction lowers the air pressure in the ear canal, causing the eardrum to flex outwards to equalize the air pressure difference on both sides of the eardrum.

The eardrum's motion is an analog of the sound wave's air pressure differences, which are an analog of the motion of the original vibrating or oscillating sound source.

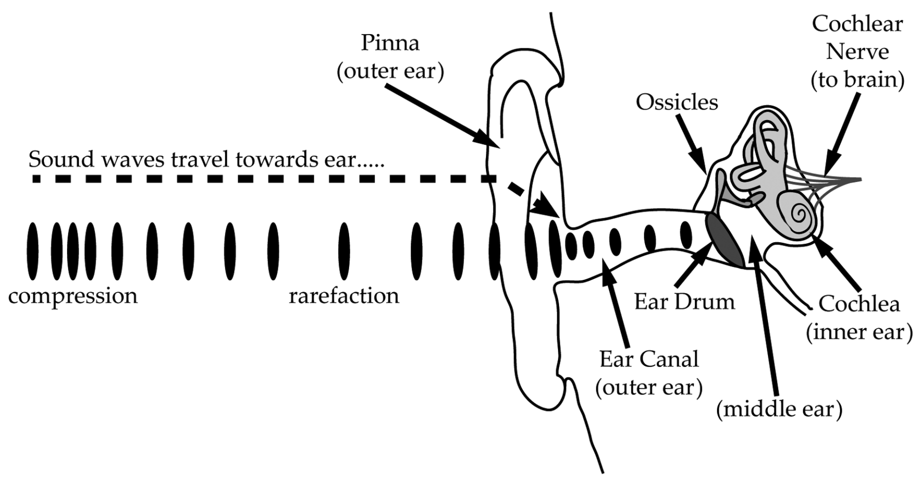

Figure 1.19 Sound waves and the human ear.

The eardrum is connected to the auditory ossicles – an "acoustical gearbox" made up of three tiny bones. The eardrum and auditory ossicles change acoustical sound waves into mechanical motion, and then into pressure waves in the fluid of the cochlea. The cochlea is a very small helix-shaped organ lined with hair-like cells that transmit signals to the auditory cochlear nerve when stimulated. Different regions of hair-like cells in the cochlea are stimulated by different frequencies. The brain then processes these signals, and the listener perceives sound.

The commonly quoted 20 Hz to 20 kHz frequency range of human hearing is a best case scenario, and different from person to person. Factors such as age and exposure to loud sound reduce this range.

- High frequency sensitivity deteriorates as we get older. By the age of thirty, 20 Hz to 16 kHz would be a more typical audible frequency range.
- Extended exposure to loud sounds can impair upper midrange (and high frequency) sensitivity.

Our ears are not equally sensitive to all frequencies. They are most sensitive to the middle of the frequency spectrum (where essential speech intelligibility frequencies are), and less sensitive to extreme high and low frequencies. The linearity (or non-linearity) of our hearings frequency response also changes with amplitude:

- At quiet levels, our ears are very insensitive to extreme low and high frequencies.
- At higher SPLs we perceive frequencies more equally across the spectrum.

THE PHON

The *phon* is a measurement used to describe the relative sound pressure levels of frequencies above and below a 1 kHz reference tone, required in order for us to perceive *all* frequencies at the same loudness as that 1 kHz tone. The phon equates to the actual SPL we are listening to at 1 kHz, so 85 phons = 85 dB SPL at 1 kHz.

- At low SPLs, 40 phons for example, extreme low and high frequencies need to be at significantly higher SPLs than the 1 kHz range, and at around 4 kHz the SPL must be a little less than the 1 kHz range for our ears to perceive the entire frequency spectrum as being the same loudness as 1kHz at 40 dB SPL.
- To perceive frequencies around 20 Hz at the same level as 1 kHz at 40 phons, an actual SPL of almost 80 dB is required at 20 Hz.
- To perceive frequencies around 10 kHz at equal loudness to 1 kHz, an actual SPL of almost 60 dB is necessary at 10 kHz.
- To perceive frequencies around 4 kHz at equal loudness to 1 kHz, an actual SPL of approximately 35 dB SPL is necessary at 4 kHz.

At high SPLs, 100 phons for example, (meaning that 1 kHz is perceived at 100 dB SPL) our ears are more linear:

- Almost no increase in the low frequencies is required to perceive them at the same loudness as 1 kHz.
- To perceive frequencies around 4 kHz at equal loudness to 1 kHz, an actual SPL of approximately 95 dB SPL is necessary at 4 kHz.
- Extreme high frequencies need to be slightly louder than the 1 kHz range to be perceived the same as 1 kHz.

Look up the Fletcher-Munson Curves. They show the relative amplitudes of each frequency necessary in order for us to perceive all frequencies at the same loudness at different Phon levels.

AUDIO EXAMPLES
Can be found on the companion website

Frequencies

All these tones are presented at identical amplitudes – ***try not to adjust your volume knob between examples!*** *The mid frequencies (around 1 kHz to 2 kHz) should appear much louder than the lower and higher extremes.*

Example 1.14: A 32 Hz tone.
Example 1.15: A 63 Hz tone.
Example 1.16: A 125 Hz tone.
Example 1.17: A 250 Hz tone.
Example 1.18: A 500 Hz tone.
Example 1.19: A 1 kHz tone.
Example 1.20: A 2 kHz tone.
Example 1.21: A 4 kHz tone.
Example 1.22: An 8 kHz tone.
Example 1.23: A 16 kHz tone.
Example 1.24: A 20 kHz tone.
Example 1.25: A smooth sweep from 20 Hz to 20 kHz, and back down.

Humans hear amplitudes between the *threshold of hearing* and *threshold of pain:*

- The threshold of hearing is quoted as 0 dB SPL – any sounds below this level are theoretically inaudible, and any sounds at or above this level are theoretically audible.

- The threshold of pain is around 130 dB SPL. This is the level at which the sensation of sound changes from being a sonic one, to one of physical pain.

These levels are theoretical "best case" scenarios, and vary not only from person to person, but also depending upon frequency. Hearing damage and age affect each person's threshold of hearing, and individual tolerance and age change the threshold of pain for each individual.

A leading cause of hearing damage is exposure to loud sound. Various governments around the world require hearing protection and regular hearing screenings to be made available to employees if noise in an industrial workplace exceeds 85 dBA SPL. Needless to say, sitting behind a drum set, in front of a guitar amp, going to a rock concert, or working in a recording studio or audio production environment often exceeds these levels – but the entertainment, music, and recording industries are not covered by the same government legislation in most countries. A drummer's ears can be exposed to over 115 dB SPL while playing!

SAFE SOUND!

If you wish to have a long career as an audio professional or musician it is important to practice safe sound:

- Monitor loudly for brief periods only.
- If you are in a situation where you are exposed to loud sounds (above 85 or 90 dBA for extended periods, or above 95 dBA for any time) wear earplugs.

Protect your hearing. Loud noise exposure eventually destroys the sensitive hair cells in the cochlea – modern medicine cannot yet make them grow back!

Inexpensive SPL meters are available, and most are good enough for approximate readings. A calibrated professional device is of course more accurate. Phone apps are also available very inexpensively, however they really need calibrating (to a known accurate meter) in order to give accurate readings.

Weighting

SPL measurements are usually *weighted*, meaning that they are biased, to better reflect how our ears perceive amplitude at different frequencies:

- *dBA* weighting significantly reduces the measurement tool's sensitivity to the low and high frequency extremes (much like we hear at lower SPLs).
- *dBC* weighting is flatter, only slightly reducing the measurement tool's sensitivity to lows and highs (much like we hear at high SPLs).

dBA weighting is generally used for workplace noise and safety measurements.

Time Weightings

Sound levels fluctuate so quickly in most music and speech that it would be impossible to read an SPL meter trying to continually show those rapidly changing levels. So SPL meters slow and smooth their reaction to sudden changes, producing a more readable display. *Slow* weighting takes one second for the meter to fully register a change in level of a continuous sound. *Fast* weighting takes 125 ms for the meter to fully register a change in level of a continuous sound. Slow measurements do not fully report the short, instantaneous level changes found in music, but are a lot easier to read.

For measurements for personal use, use the Fast setting to try to get a more accurate reading of rapidly changing levels. But if the Fast reading is jumping around by more than a few dB, it becomes difficult to infer a meaningful value, so use the Slow time weighting. Most industrial measurements use the Slow setting, but for measurements for professional or legal use, the time weighting selection is determined by professional standards or laws.

So, some essential theory out of the way, let's discuss some gear-related topics!

1.14 Signal Flow and Audio Level Standards

Signal flow refers to how an audio signal flows through a chain of devices while it is being recorded or played back. If you plug a mic in and there's no sound, or the sound is distorted, a thorough understanding of signal flow makes it possible to quickly troubleshoot and fix the problem – and makes you a better and more employable audio professional.

Signal flow is case-specific, and differs depending upon the type of facility and specific gear in use – so it's impossible to describe every scenario in this book. However, typical hardware mixing console and computer/digital audio workstation recording chains are shown in **Figure 1.20**.

Microphones and Mic Levels

Microphones change acoustical sound waves (variations in air pressure) into electricity – specifically, variations in voltage. This voltage is very small, measured in millivolts (mV), and is called *mic level*. When picking up loud sounds, low output dynamic mics output just a few mV, while high output condenser mics can output 20 to 50 mV, or occasionally more.

A mic level signal needs amplifying before it can go through the main mixing console circuits, or the analog to digital converters in most audio interfaces. A *preamplifier*, *preamp*, or *mic pre* is the first circuit in the signal flow of a mixer (or audio interface) that the mic

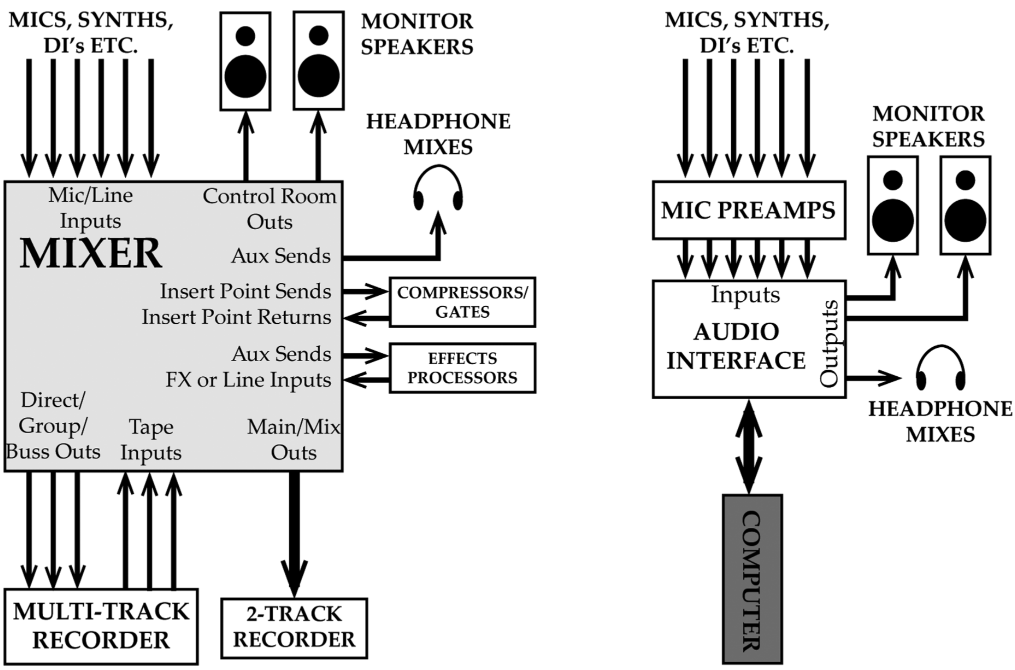

Figure 1.20 **Left:** The signal flow of a typical hardware recording chain, including a mixing console. **Right:** An example DAW recording chain.

is plugged into. A preamp can also be a separate device, external to the mixer or audio interface. The preamp is a critical circuit – it needs to amplify the tiny voltage coming from a microphone, without increasing noise, undesirable distortions, or negatively affecting the frequency response or sound of the microphone.

Line Levels

A preamp applies *gain* (amplification) to the mic level signal to raise it to *line level*, which is the level the rest of the mixing console (or audio interface) operates at. There are two standards for line level signals:

- "*+4 dBu*" or *professional level* represents the audio waveform as voltages of up to about ±1.7 V.
- "*–10 dBV*" or *consumer level* represents the audio waveform as voltages of up to about ±0.4 V.

Consumer level "–10" interconnects are commonly found on domestic devices such as CD, DVD, and Blu-ray players, turntables, and gaming systems.

- If you connect a –10 dBV device to a +4 dBu input you will see low levels (about 12 dB lower than they should be) and have to apply a lot of gain (which can introduce additional noise).
- If you connect a +4 dBu output to a –10 dBV input you will probably overload and distort that input unless there is an input gain stage you can turn down on the receiving device.

Connections between studio hardware are usually “+4” professional line level.

- For a “direct to stereo” recording, the main outputs of the mixing console are connected to a stereo hardware recording device or an audio interface attached to a computer-based recording system.
- For multi-track recording, the subgroups, multi-track busses, or direct outs from the console can be used to send individual microphones to individual tracks in a DAW, hardware multi-track recorder, or an analog tape machine – yes, analog tape is still used!

The playback outputs from a multi-track recording device will return into individual inputs on a mixing console. If mixing is taking place “in the box” (inside the DAW system), only a stereo left/right (or a set of surround or immersive outputs if mixing in those formats) may come out of the DAW and be routed to the monitor speakers.

Loudspeaker Level

If passive monitor speakers are being used, the line level output of the mixer or DAW will need amplifying to *loudspeaker level* – which is the *many* volts required to move a loudspeaker cone (± 10 to ±50 V or more). The line outputs are connected to a power amplifier, which will handle that task. This amplification is built into *powered* or *active* monitors – hence the input is line level, and the speaker requires an AC electrical power connection for its amplifiers and crossover circuitry.

Instrument Level

Electric guitars and basses have *instrument level* connections. In terms of signal level, they are not dissimilar to mic level – so plugging a guitar or bass directly into a mic input will usually result in decent levels. However the impedances (the electrical load each device presents to the other) of the two devices are not designed to work together – and the sound will be muddy and dull. Some audio interfaces and preamps have an instrument input or two on them. If yours does not, you will need to use a *direct box* or *DI*, or a dedicated instrument preamp to convert the instrument signal level and impedance, if you wish to plug a guitar or bass directly into your recording chain.

IMPEDANCE

An analog device's *impedance* is its opposition to electrical current flow. In DC (direct current) steady-state continuous voltage electronics, the term "resistance" is used instead. In AC (alternating current) electronics, impedance takes into account not only resistance, but also the capacitance, and inductance factors caused by the rapidly fluctuating + and – AC voltages. Impedance is measured in Ohms, Ω.

An analog audio *output* creates AC voltage, and presents an output impedance to the device it plugs into. The circuitry in an analog *input* has an input impedance that it presents to the device connected to it.

Audio devices are not designed for both the output and input impedances of a connection to be the same – this can actually cause level drops and losses at certain frequency ranges in some situations. Instead, outputs are usually designed to have low impedances (microphone impedances are about 150 to 200 Ω), and inputs to have much higher impedances, usually at least ten times the expected source device's output impedance (mic preamp impedances are usually between 1.5 KΩ and 3 KΩ).

Some preamps even have variable and adjustable impedances because different input impedances can change a mic's sound.

1.15 Gain Structure and Recording Levels

Levels, levels, levels! A beautifully clean recording (as opposed to one that is full of hiss, noise, or distortion) is created in part by setting appropriate levels as signals leave one device and enter the next:

- If too low a signal enters a device, that device will add a greater relative percentage of its own hiss and noise to the signal. That extra noise will end up being turned up to compensate for the low input level somewhere later in the signal flow.
- If too high a level enters a device, then subtle or not so subtle distortion will result.

Both these types of noise and distortion are undesirable.

With the exception of a *pad* control (which attenuates a microphone's output by a fixed amount) mics do not have gain controls on their outputs. They are connected to a preamp, and the gain control on the preamp is used to set the amount of amplification that the mic signal receives in order to bring it up to an appropriate line level.

On a digital dBFS scale, the golden rule is "*do not clip*." To avoid potentially nasty distortion the "clip" or "over" indicators must not light.

- In the days of 16 and 20 bit recording, the approach was to record with the levels as high as possible, without going "over," while making sure enough headroom was left to avoid clipping unexpected louder moments.

- Given the large dynamic ranges offered by modern 24 and 32 bit (or greater) recording systems, it is much less necessary to "max the meters." Peaking a little lower, maybe around −12 dB or so, leaving plenty of headroom means that the digital convertors are not constantly running at or close to their upper limit. Many engineers claim they sound better that way.
- You do not want levels to be too low, using only the lowest bits of resolution. The track will probably need to be turned up during the mix process. What will be turned up is a grainier low level signal, with a significant amount of noise and distortion created by the recording equipment.

CLIP INDICATORS

In addition to the generally preferred sound of analog to digital converters running 6 to 12 dB below maximum levels, other reasons to keep levels below maximum are:

- Often, clip indicators don't indicate clipping unless several consecutive samples are "over" – by which time the problem is more serious than a single sample overload.
- This also means that if a single sample is clipped, but multiple samples in a row do not, the clip will not be indicated. Just because it doesn't show up, doesn't mean it hasn't happened and distorted the sound.

A big difference between particular digital recording devices and systems is how they distort when overloaded. Some are more sonically forgiving than others. You can expect a cheap device to produce ugly clicks and chirps immediately on clipping, while you might be able to get away with the occasional and light clip on a more expensive, better quality device.

On an analog device, with a dBV peak meter, the level can happily be pushed past "0," and into the "+" values. On many analog devices, more of the unit's real character will be picked up by doing that. Care must still be taken to not light any "over" indicator though!

A VU needle type meter cannot show you peak level, so good levels are when it is peaking at, or just above "0" for non-percussive sounds, and below "0" for percussive sounds. Again – the "peak" light should not come on.

With good input levels established, it should be easy to get good levels to the recording device.

From the Mixer to the Recording Device

Mixer outputs to the recording system may be pre-fader or post-fader, and pre-panner or post-panner. If they are post-fader and/or post-panner, faders in-between the preamps and outputs to the recorder should be set at unity (which is marked "0" near the top of the fader's travel), and pan controls should be set correctly if odd/even panning rules apply to the outputs in use.

- In its unity position, a fader does not boost or cut the level of the signal on that channel. What leaves the channel is the same level as was set at the input gain/preamp stage – possibly minus a few dB depending upon the output being used, because of the stereo panning laws used to maintain perceived equal amplitudes when a sound is panned to different positions. (When a sound is panned centrally for example, two loudspeakers reproduce it – that's twice as much power than if it was hard panned and only reproduced by one loudspeaker. Panning laws attenuate the sound's amplitude slightly as it is panned towards center so that it is perceived at equal amplitude regardless of where it is panned.)
- Turning a fader above unity also turns up mixer channel noise. If a channel needs turning up above unity, it usually implies that the input gain is too low, so the input level to the channel should be checked.
- Turning a fader down below unity is fine during mixing, but if post-fader sends to the recording device are used during tracking, turning the fader down will send lower levels to the recorder.

When mixing, good main fader output levels can be achieved by setting the faders for the most important, loudest channels of the mix at around unity, and other faders below unity. If too many faders are set above unity, the main output levels can be too high, clipped, or overloaded; or less clean, due to the additional amplification and noise added by faders that are above unity. Conversely, running too many faders too low will result in low output levels and an increased relative percentage of noise. A compensating gain increase will be required in the next recording device that while turning up the desired signal, will also turn up the undesired extra noise.

1.16 Analog Audio Connectors

Microphones and all of the equipment in the previous recording chain diagram need to be plugged in and connected together correctly. You need specific cables to do this, so you need to know the names of the connectors used, and what each type of connector is commonly used for.

XLR Connectors

XLR connectors are the preferred connectors for single channel analog connections. They are robust and lock into place – a little release button needs to be pushed to disconnect them. They can be found with varying numbers of pins in the connector. Three-pin XLR connectors are usually used for standard mic and line connections. XLR connectors can be

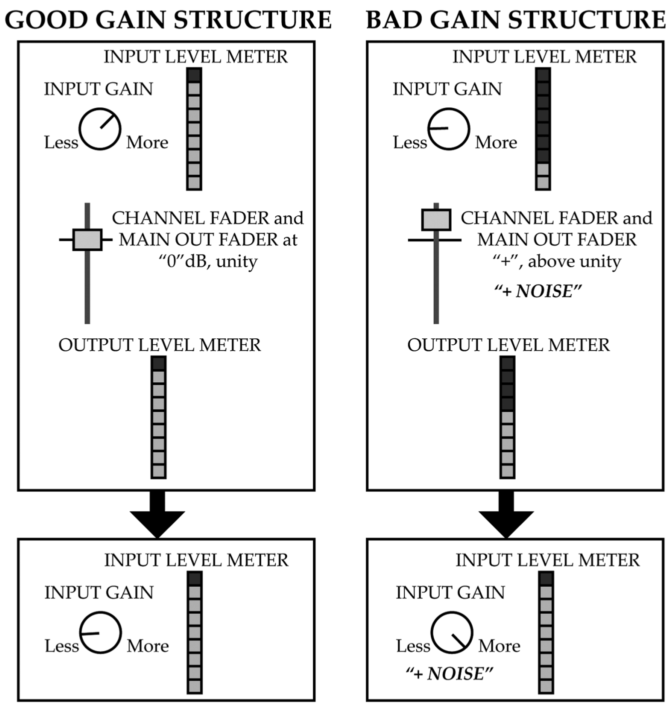

Figure 1.21 **Left:** Good input levels and the fader at unity create good output levels on the top device, and good input levels on the bottom device. **Right:** Low input levels and undesirable fader gain produce a low level output on the top device, requiring increased input gain on the bottom device. This results in a noisier signal, due to the two additional stages of amplification after the top device's input gain.

male or female – the male plug is the one with the pins, the female (jack) the one with little holes. Signal always flows out of the male and into the female.

Quarter-Inch Phone Connectors

Quarter-inch phone connectors are also used on a lot of audio equipment. They come in two types:

- *Quarter-inch TRS* connectors have three segments to the connector – the tip, ring, and sleeve.
- *Quarter-inch TS* connectors have just two segments to the connector – the tip and sleeve.

Quarter-inch connectors come in male and female sexes – male plugs and female jacks. The genders of quarter-inch connectors are not dedicated to inputs or outputs like XLR connectors – sound can go in or out of a male plug or female jack.

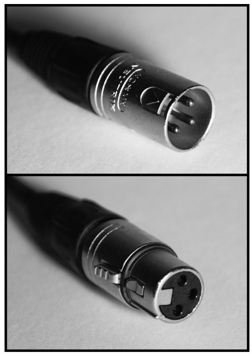

Figure 1.22 XLR Connectors. **Top:** Male. **Bottom:** Female.

Line level connections commonly use quarter-inch TRS connectors. Instrument level, standard "instrument cables," and some speaker level connections use quarter-inch TS connectors. From the outside, all female quarter-inch jack sockets (instrument, line or speaker level, TRS or TS) look identical – so it's really important to know what you're hooking up and use a cable with the correct connectors and wire for the job! *You don't want to fry a mixer channel by connecting a speaker level output to a line level input!*

BALANCED AND UNBALANCED CONNECTIONS

Have you ever heard stray taxicab or radio station interference leaking into your sound system or a guitar amplifier? This is usually due to a two-conductor *unbalanced* audio connection between two devices. A *balanced* audio connection requires three-conductor cable and connectors such as XLR and TRS. Balanced connections offer superior rejection of external electrical and electro-magnetic noise and interference, and allow much greater cable lengths to be used. For studio and sound system interconnects, balanced connections are recommended and preferred – however, both devices that are connected together have to feature balancing circuits on their respective inputs and outputs. Connecting a TRS cable to unbalanced inputs and outputs provides no advantage. Neither does connecting a balanced output to an unbalanced input or vice versa – the connection will still be unbalanced and susceptible to noise and interference.

RCA/Phono Connectors

RCA or *phono* connectors are found on some consumer and prosumer equipment. They are two-conductor, so unbalanced. They are relatively fragile, and do not lock into place – so they are not a preferred method of connection for professional audio systems. CD, DVD

Figure 1.23 **Top:** A quarter-inch TRS male plug. **Middle:** A quarter-inch TS male plug. **Bottom:** A quarter-inch female socket/jack (it could be TS or TRS – there's no way to tell visually, but "Phones" here indicates that it is TRS).

Figure 1.24 **Top:** A male RCA (or phono) connector. **Bottom:** Female RCA connectors.

and Blu-ray players, turntables, DJ equipment, game systems, and other consumer devices usually feature RCA connectors.

Bantam/TT and Quarter-Inch Military Connectors

Patchbays are used in studios to place the inputs and outputs of the audio equipment in a centralized location – so that they can be hooked up as needed without rewiring the back of the equipment racks. Prosumer patchbays often use quarter-inch TRS connectors. Professional patchbays use either *MIL* or *bantam/TT* (*Tiny Telephone*) connections. Both are three-conductor TRS connectors, so can carry balanced signals. The military connector is about the same size as a quarter-inch connector, but is much more robust and secure. Military and quarter-inch TRS patchbay units feature up to 48 connection points per rack-space unit. The

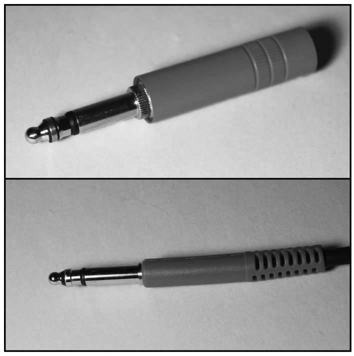

Figure 1.25 **Top:** A MIL/military connector. **Bottom:** A smaller bantam/TT connector. Compare their shapes to a quarter-inch TRS connector.

Figure 1.26 25-pin D-Sub connectors. **Top:** Male. **Bottom:** Female.

bantam connector is not as robust or secure, but has the advantage of a smaller size which allows up to 96 connection points to be put in a single rack-space patchbay unit.

25-PIN D-Sub Connectors

Behind the scenes, in the back of equipment racks, some audio hardware and modular patchbay systems use *25-pin D-sub* connectors to transport eight channels of analog or digital information between devices. The Tascam pin-out is the most common wiring format for this type of connection.

Eighth-Inch TRS Connectors

Small, fragile, and generally unreliable, the eighth-inch, or 3.5mm TRS plug, as found on some mobile phones and earbuds, has no place in professional audio!

1.17 Digital Audio Connectors

SPDIF

S/PDIF (*Sony/Philips Digital Interconnect Format*) is primarily a consumer and prosumer format that can transmit two-channel stereo digital audio data, at up to 24 bit/192 kHz resolution (originally it was 24/48). There is a copper-wire version that uses a single 75 ohm cable and RCA connectors, and an optical fiber version that uses a single optical fiber with *TOSLINK* connectors. On consumer audio-video equipment, the SPDIF connections can also transport some lossy compressed surround sound audio formats.

AES/EBU

The original *AES/EBU* format transmitted two-channel stereo data, at resolutions up to 24 bit/48 kHz, over a single XLR terminated cable. AES3 now exists, which supports stereo data at resolutions of up to 24 bit/192 kHz, over balanced XLR, unbalanced RCA, unbalanced BNC, and TOSLINK optical fiber connections. In addition to offering sturdy locking connectors, another advantage the XLR version has over SPDIF, is that it is a balanced connection, allowing much longer cable lengths.

ADAT Lightpipe

The *ADAT Lightpipe* format uses a single TOSLINK optical fiber to transmit eight channels of digital audio at up to 24 bit/48 kHz resolution. There are also higher sample rate versions:

- SMUX-2 splits 88.2 and 96 kHz data over two fibers; the first carries channels 1 through 4, and the second carries channels 5 through 8.
- SMUX-4 transmits two channels of 176.4 and 192 kHz data down a single fiber.

Figure 1.27 A TOSLINK optical fiber connector as used for optical SPDIF connections, and ADAT Lightpipe connections.

Both the transmitting and receiving devices must be equipped with the appropriate SMUX protocol for high resolution use.

TDIF

TDIF (Tascam Digital Interface Format) is a bidirectional, unbalanced copper-wire format that transmits 8 channels of audio via a 25-pin D-sub cable. Its original specification was for up to 24 bit/48 kHz resolution, but there is now a version that supports up to 96 kHz.

MADI

MADI (*Multichannel Audio Digital Interface*) connections, commonly found on higher-end equipment, but making their way onto less expensive equipment, transmit up to 64 channels of standard resolution audio (up to 24 bit/48 kHz), or a reduced number of channels at higher resolutions. MADI connections can be either copper-wire via 75 ohm BNC terminated cables, or optical fiber terminated with SC type plugs. Cable lengths can be a lot longer than other formats – up to 100 m of copper-wire, or 2000 m of fiber.

Audio Over IP

Audio over IP (Internet Protocol) systems transmit multichannel digital audio data over inexpensive Ethernet networks and cables. While most systems can use shared and pre-existing internet infrastructure, they perform best on their own dedicated network.

DANTE is one of the most popular formats, and is found on a wide variety of live sound and recording equipment, at a wide variety of price points. Using 10 Gigabit networks, up to 1024 channels of audio can be transported per link – 512 in each direction, in and out.

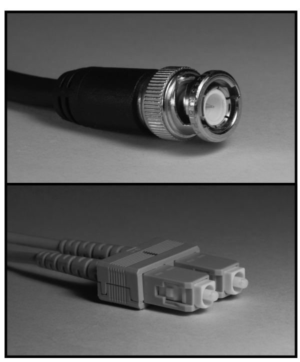

Figure 1.28 MADI connectors. **Top:** A BNC terminated cable. **Bottom:** A pair of SC plugs provide bidirectional communication.

Sample rates of up to 192 kHz, and bit depths of up to 32 bit are supported. When using Gigabit or 100 Mbps networks the number of channels may be reduced depending upon the sample rate and bit depth in use.

DANTE hardware devices place limits on the number of channels that may be input to or output from that device – for example a DANTE interface may support “32 in x 16 out” operation – but that does not limit the whole network to that lower number of channels. Different devices can be set to use different channels of data. Controller software is used to set up the network and the input and output assignments of each device. Once programmed, the equipment retains those settings, so they can operate without the computer attached to the network.

Ravenna is another audio over IP format. Typically implemented on high-end recording equipment including products for super-high resolution recording, it supports sample rates of up to 384 kHz, and up to 32 bits (both fixed and floating point formats).

CobraNet and *EtherSound* are other Ethernet-based technologies, developed before DANTE and Ravenna.

1.18 Digital Audio Basics

PCM

Most digital recording systems (hardware, software, and plug-ins) use a method known as *PCM* (*Pulse Code Modulation*) to encode and represent audio waveforms. An understanding of PCM will enable you to make more educated decisions as you plan and record projects.

Acoustically, audio is an analog format – its amplitude varies smoothly, with infinite resolution between minimum (silence) and maximum (the point at which the air distorts the waveform); and smoothly with infinite resolution over time. In an analog audio device, a voltage that similarly varies smoothly, with infinite resolution between the minimum and maximum amplitudes it is capable of, and with infinite resolution over time, represents the audio waveform. *A PCM digital device is more like a camera – taking snapshots of the audio amplitude at specific intervals of time.*

Sample Rate and Frequency Response

Figure 1.29 shows an analog audio waveform being sampled by an *Analog to Digital Convertor* (*ADC*, or *A to D convertor*) in a PCM system. The A to D convertor changes the continuously varying analog waveform into discrete *samples* – measurements of the waveform’s amplitude taken at regular intervals of time. The number of samples taken per second is known as the *sample rate* or *sampling frequency*, abbreviated to *fs*.

The sample rate dictates the highest frequency a system can record. The *Nyquist Theory* states that the sample rate must be at least double the highest frequency to be recorded.

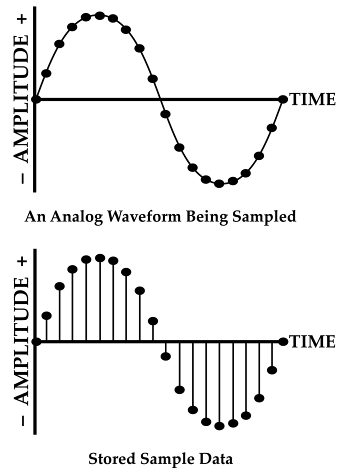

Figure 1.29 **Top:** The smooth line is the analog input waveform. The dots represent samples – amplitude measurements taken at regular intervals of time. **Bottom:** The digitally stored information.

This potentially allows for a sample measurement to be taken during both the positive and negative parts of that highest frequency's waveform. If both a positive and negative portion of the waveform could not be stored, the reconstructed waveform, when played back, would contain added erroneous lower frequencies – an artifact known as *aliasing*.

Standard CD sample rate is 44.1 kHz. This means that 44,100 measurements of the waveform's amplitude are taken per second, on both the left and right channels of a stereo signal. The frequency response of this sample rate is up to about 20 kHz – it approaches the capabilities of human hearing. Lossless streaming services stream standard resolution audio at 44.1 or 48 kHz.

Given that the best human hearing only extends up to 20 kHz anyway, why would we want to record at commonly available sample rates of up to 192 kHz (and beyond), with frequency responses of up to and beyond 90 kHz? Higher sample rates offer two major benefits:

- They raise the highest frequency that can be recorded.
- They increase the stored resolution and accuracy of the waveform being recorded.
- There is also research that shows that although we can't directly hear frequencies above 20 kHz, we can detect their interactions with, and effect on lower frequencies.

According to the Nyquist Theory, higher sample rates should not be necessary to accurately record and reproduce the highest frequency possible for a given sample rate, however technological limitations mean that good quality, higher sample rate converters produce high frequencies that sound smoother, more natural, and less brittle and harsh than lower resolution converters. But – high quality low sample rate converters may still sound better than mediocre quality high sample rate converters!

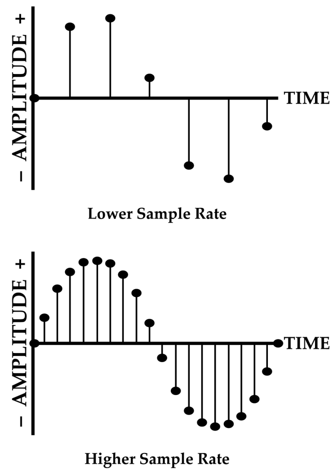

Figure 1.30 **Top:** At a lower sample rate, less detailed data, and a more approximate version of the sound wave is stored. **Bottom:** At a higher sample rate, much more detailed data, and a more accurate version of the sound wave is stored.

Bit Resolution and Dynamic Range

CD quality audio is still a benchmark (CDs still exist in some markets!) and uses 16 binary digits to represent the value of each sample, using fixed point numbers. This is known as its *bit resolution*. An approximate way of calculating the dynamic range of a fixed point PCM system is to multiply the bit resolution by six. Each bit provides 6 dB of dynamic range. Therefore, a 16 bit system offers a theoretical dynamic range of approximately 96 dB. This falls far short of our ear's 130 to 140 dB dynamic range. Professional recording systems are capable of at least 24 bit resolution (and many offer 32 bit, 32 bit floating point, and even 64 bit floating point resolution). Twenty-four bit resolution produces a theoretical dynamic range of about 144 dB, which equals the capability of the human ear. However this is not its primary benefit, as we will hopefully never listen to a playback system so loud that it transfers this range into actual SPL!

Being digital, and non-continuous, the amplitude scale is "stepped" into integer values. A 16 bit (fixed point) system allows 65,536 steps to exist on the amplitude scale. If a sampled amplitude falls between two of those values, it will be rounded to one of the adjacent integer step values. This is a form of distortion known as *quantization error* or *rounding error*.

For professional recording and mixing, a minimum of 24 bit resolution is used. At 24 bit resolution (fixed point) the number of steps on the amplitude scale is 16,777,216. This means that less rounding will occur, and low level details are no longer so close to being masked by the noise-floor of the system. Amplitude values are stored more faithfully, low-level details sound less grainy and metallic, the waveform reconstructed is smoother and more transparent, and a fuller, more precise stereo image is produced.

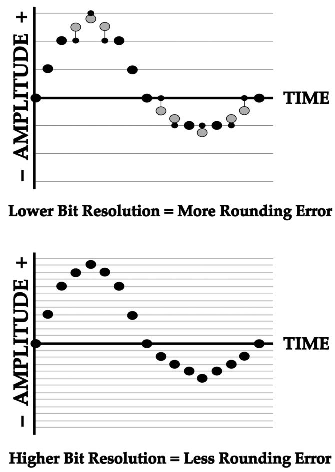

Figure 1.31 **Top:** At lower bit resolutions there is a greater probability that intended sample values cannot be accurately stored. The grey dots representing the actual input waveform cannot exist at their correct amplitudes and will be rounded up or down to the nearest possible step on the amplitude scale (the small black dots). This is particularly noticeable in low level sounds (the right part of the graph) and low level details contained in louder sounds (the left part of the graph), where the relative percentage change is greater. **Bottom:** At higher bit resolutions there is less rounding, and improved accuracy in both the subtle details of louder sounds (the left part of the graph), and quiet sounds themselves (the right part of the graph).

24, 32, OR 64 BITS? FIXED OR FLOATING POINT?

Thirty-two bit *fixed point* recording offers more dynamic range than 24 bit recording – up to 192 dB.

Floating point formats don't use the same straight forward bitmapping as fixed point formats. They use scientific notation such as:

$$3.4 \times 10^{38}$$

$$\text{MANTISSA} \times 10^{+/-\text{EXPONENT}}$$

In a 32 bit floating point system, 23 bits are dedicated to storing the mantissa, and the other nine bits to storing the exponent value. This allows the amplitude to be "floated" above or below the limits of 24 (or even 32) bit fixed point resolution. The potential dynamic range is increased to 1528 dB (!) – so it is theoretically impossible to clip, and very quiet sounds still get represented with good resolution!

However, the A to D convertors used to record (and on any equipment used during mixing), do still have finite amplitude limits – if the input level is too loud the signal will still clip and distort as it overloads the convertor, and if the sound gets too quiet it will still disappear below and into the convertor's noise floor. So it is still important to maintain good levels and not to clip or overload convertors!

D to A convertors cannot simply play back a floating point file – software needs to scale the levels for correct conversion. Most D to A convertors do not have 32 or 64 bit capabilities. Even if a D to A convertor offers 32 bit conversion it is doubtful that it's dynamic range and noise specifications will allow a 32 bit file to actually sound any better than a 24 bit file.

Weighing up the increased file sizes (32 bit files are 33 percent larger than 24 bit files) and actual use of the potentially increased dynamic range, 32 bit formats offer little advantage over 24 bits in most recording and playback situations. Rarely will sound sources, or the technical capabilities of the preamps, convertors, and other equipment in the recording chain exceed the capabilities of 24 bits of dynamic range.

32 and 64 bit floating point resolution is important and does produce tangible benefits when digital audio data is being processed and manipulated – when mixing in a DAW, or when a file will be sent to another piece of software for further processing, for example. The levels within the software can be greatly reduced or increased without the noise, quantization errors, or distortion that might occur with fixed point resolutions. But the levels should be made sensible prior to bouncing or any digital to analog conversion during mixing or on playback.

If a project is mixed in a DAW there can be benefits to bouncing a 32 bit (usually floating point) file to be sent for mastering. But 24 bit files are still standard, particularly for projects mixed using analog consoles and hardware.

File Size

Recording at 16/44.1 resolution produces a file size of approximately 5 MB per minute for a mono track. Mono 24/96 files are over 17 MB per minute, and 24/192 files over 34 MB per minute. Higher resolution files take up significantly more space than lower resolution files, particularly when you multiply by the number of tracks recorded. Doubling the sample rate potentially reduces the available track count by half due to the storage media's maximum read/write speeds, and also reduces the number of plug-ins a DAW can run before the CPU overloads or the DSP system is maxed out. *But storage media is getting larger and cheaper per TB every day, and computers getting faster – so it's definitely worth recording at least 24/88.2 or 24/96 files if your system can handle the necessary track count and processing – particularly for major label projects, or if you know a high resolution release is planned.* That said, many independent projects are still recorded at 24/44.1 or 24/48.

It's advisable to record at a high sample rate so your work can take advantage of future high quality dissemination formats, and not have a new format reveal the limitations of the initial recording resolution. If you start a project at a high resolution, and run into track count or CPU horsepower problems, then most DAWs allow you to convert and export the entire session and all its associated files at a lower sample rate.

It used to be advisable to record at integer multiples of the intended dissemination format's sample rate – to avoid potentially degrading non-integer sample rate conversion. But modern sample rate conversion algorithms are very good, so the sample rate used to record, mix and master should be based on the highest quality intended release format – and converted down for standard resolution for other delivery formats after mastering.

2 "Good Sound"

In This Chapter:

2.1 Recognizing Good Sound

Before even thinking about plugging in and setting up a microphone, it is important to understand what the desirable (and undesirable) characteristics of sound and a mix are. That's why there are some chapters in this book that discuss things you need to know in order to use microphones to capture the best, most suitable sound possible.

There's no simple answer to the question "what is good sound?" The best answer might be along the lines of "whatever is stylistically and artistically appropriate." Good sound is subjective. One person's ideal guitar sound may be another person's worst nightmare

– however that is often related to whether a sound is appropriate for its context, rather than the sound being simply "good" or "bad." But bad sound certainly does exist! Poor quality sound sources, poor quality equipment, bad recording techniques, and poor mixing skills can all result in sound that is inappropriate, questionable, or just plain "wrong"!

How do you learn to record and mix well? There are basic concepts and skills that should be mastered before developing your own style. Musicians develop their skill sets and individual musical sound and style by listening to other musicians, emulating them, and eventually synthesizing many influences into their own unique characteristics. As a sound engineer or producer, you should similarly find good quality recordings, listen to them, analyze them, and try to emulate them – learning how to capture and control sound before eventually developing your own style.

Listening to, and becoming intimately familiar with a wide variety of musical and production styles will make you more marketable in the industry – if you only listen to electronica, good luck when an acoustic folk band shows up for a show or session you're working! Hopefully, like many audio professionals, you can combine your love of great sounding recordings with artists you enjoy musically. But don't just listen to music because you like the style or the artists – when starting to listen critically, it's often easier to concentrate on the sound (and not be distracted by the music) when you're listening to artists and musical styles you're not a fan of. It's very easy for your love of an artist or band to persuade you that the recording is better than it actually is!

2.2 Sound Reproduction Formats

Mono

A *mono*, or *monophonic* playback system has only one loudspeaker, as in **Figure 2.1**. Mono playback systems include some TVs, bedside alarm clock radios, ceiling type loudspeaker

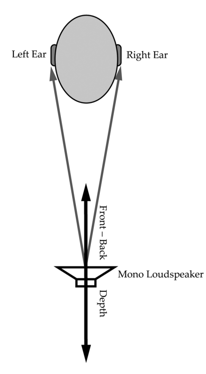

Figure 2.1 A mono reproduction system with a single loudspeaker. Sound from the loudspeaker arrives at both ears at the same time (if the loudspeaker is centered).

systems in retail outlets (which have many distributed loudspeakers, that are all sent the same signal), and speakers on many tablets and mobile phones. Mono systems are one-dimensional – they can offer a sense of front/back depth through the use of creative recording and mixing techniques, but there is no sense of width. It is potentially difficult for a listener to clearly hear everything that might be going on in a mono mix, which makes achieving a good mono mix very challenging.

A small, nasty sounding mono *Auratone* type loudspeaker is common in many professional recording studios – listening on this enables engineers to anticipate the effects of poor quality mono sound systems, and make sure their product translates acceptably to them.

Stereo

Stereo, or *stereophonic* systems feature two playback channels – left and right. *Different* signals are sent to each, as shown in **Figure 2.2**. In a stereo loudspeaker system, sound from the left loudspeaker travels to the left ear, *and* the right ear – where it arrives slightly later, changed in timbre due to the extra distance and the filtering/equalization effects of wrapping around the face and head. These delay and EQ effects are known as *HRTFs*, or *head related transfer functions*. Similarly, sound from the right loudspeaker travels to the right ear, and also the left ear – again, slightly delayed and changed in timbre. The sound transfer to each "opposite" ear is known as *inter-aural crosstalk*, and is an essential and desirable component of a stereo playback or monitoring system.

The reason a sound coming from just the left loudspeaker sounds like it's coming from just the left loudspeaker (even though the sound travels to both of our ears) is because our brain uses the HRTF sub-millisecond-level time-delay and the EQ differences of the wavefront's arrival at each ear to determine directionality. We perceive the sound as coming from the direction of the wavefront that arrives at our ears first. This is related to the *Law of First Wavefront*, *Haas Effect*, or *Precedence Effect* (which are all the same thing).

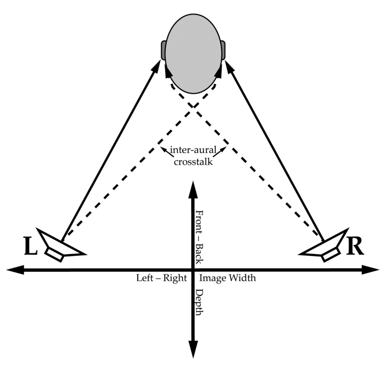

Figure 2.2 A stereo reproduction system.

If both loudspeakers reproduce an identical sound wave simultaneously, both ears receive the same wavefront at the same time, followed by identical slightly delayed inter-aural crosstalk. The listener perceives this sound as coming from a central location, directly between the loudspeakers – where there is no physical loudspeaker. The illusion of a sound located where there is no loudspeaker is called a *phantom image*.

The stereo format is a two-dimensional format when creative recording and mixing techniques are used:

- There is a clear a sense of left/right directionality, because it is possible to position sounds anywhere between the loudspeakers, and to create the illusion that sound is coming from just beyond them.
- It is also possible to create the illusion of front/back depth.
- To a lesser degree, it is possible to create the limited illusion of a third dimension – height.

Surround and Immersive Audio

Older *surround sound* formats, such as 5.1 and 7.1, have been around for decades. They utilize additional loudspeakers positioned around the listener on the horizontal plane. They offer the advantage of being able to envelop the listener with sound coming from behind and to the side of the listening position. Additionally, a dedicated center channel loud-speaker can take the place of center phantom images, and provide alternate sonic imaging qualities. The ".1" refers to a subwoofer channel.

The term *immersive audio* is now used to describe systems with more than two independent playback channels. Fully immersive formats with added height channels (such as 7.1.4, the last number being the number of height channels) are now established and used in the movie, TV, and music industries, and even in concerts. These formats create a truly three-dimensional listening experience.

It is possible to listen to immersive audio on regular headphones if content is encoded for binaural dissemination. So, with a suitably equipped playback system and multi-channel source material you can listen to movies, games, and even music releases in three-dimensional immersive audio using headphones – although the results are not as accurate or predictable as using a multi-channel loudspeaker system.

STEREO PRODUCTION FIRST!

Surround sound and immersive audio formats and recording techniques are introduced later in this book, but stereo production should be mastered before attempting any type of surround or immersive production.

2.3 Monitoring Options – Loudspeakers, Headphones, and Earbuds

Stereo is the most common consumer listening format, and a pair of studio loudspeakers, called *monitors*, are the preferred listening system for stereo recording, mixing, and mastering. Boom boxes, hi-fi systems, stereo televisions, car stereos, mobile phones, and portable media players all offer two-channel "stereo" playback. But are all of these devices actually stereo?

Figure 2.2 shows that a stereo system creates inter-aural crosstalk that our hearing system uses to determine the directionality of sound. **Figure 2.3** shows a pair of headphones (or earbuds) on a listener. There is no inter-aural crosstalk in this system – the sound from the left driver goes only into the left ear, and the sound from the right driver goes only into the right ear. The only directional information presented to and processed by our brain is the amplitude difference of the sound between each ear. For this reason, headphones and earbuds should not really be called stereo, despite the label which says "stereo" on the retail packaging! They are in fact *binaural* – featuring two channels, without the inter-aural crosstalk necessary to make them a true stereo system.

There are several reasons that mixing using high quality, professional loudspeakers is preferable to mixing using headphones:

- Mixes created on loudspeakers translate very well to headphones.
- Mixes created on headphones do not translate as well to loudspeaker systems.
- "But most of the listeners will be using earbuds anyway." It is perfectly OK, and good, to check a mix on headphones or earbuds. But to create a stereo mix that will best translate to the greatest variety of playback systems, it should be crafted on a stereo loudspeaker system. A mix that doesn't translate well to many reproduction systems, or future systems, will have a more limited audience, and a more limited lifespan.
- Headphones do not present the same stereo image or front/back depth that stereo loudspeaker systems are capable of. While a conventional stereo headphone image may be correctly described as more precise and surgical, it is also smaller and more compact. If you are mixing for headphones exclusively, then knowledge of some of the miking techniques discussed later in this book can allow you to incorporate delay-based image steering into your recordings, to add some delay and HRTF-like artifacts into the binaural reproduction system.

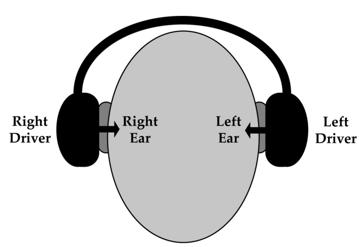

Figure 2.3 A binaural headphone or earbud system.

- Professional monitor loudspeakers are built to tell the truth, and not to hype certain frequency ranges, or to sound instantly "pleasing" (like most lower cost headphones, and consumer loudspeakers). Professional monitors reveal more problems in your mix or recording than most headphones or consumer loudspeakers. This makes it more challenging to get the sound right on good loudspeakers – but when you do, the results should translate well to a larger variety of playback systems. Most headphones and consumer loudspeakers don't reveal so many details or potential problems – which can lead to unfortunate surprises when a mix is played on a better sound system.
- Many headphones and consumer loudspeakers hype the bass for "instant impact." Professional monitors have a flatter, more linear frequency response in order to more correctly present the sound, and not mask any potential mix issues. This means that a first impression of a professional monitor speaker might be that it lacks bass, and is not as full sounding. If you are new to professional monitors, you need to learn how *your* monitors sound – intimately. Listen extensively to industry-respected, good sounding reference recordings on your new monitors. Recognize the characteristics of good sound on those monitors, and get used to the balance of the low, mid, and high frequencies.

If you do not know and trust how your speakers sound, you cannot reliably use them to craft a mix that will sound good elsewhere.

There *are* some great sounding headphones on the market, and these *can* be used for monitoring, subject to the caveats above. Earbuds however, particularly the cheap stock items that the majority of the public happily use, sound horrible. They either lack real bass, have a highly inaccurate "bass" boost that isn't in a really low frequency range, and their high frequency performance is usually not very good.

EARBUDS – A HEALTH HAZARD?

Earbuds connected to phones and portable media players can produce peak SPLs of up to 127 dBA in the ear canal, depending on the device! The average sustained level is a little below the peak level, but still well in excess of the recommended 85 dBA maximum sound exposure level without hearing protection.

Stock, and most cheap earbuds are non-sealing "open" designs that let in extraneous sound so the listener is aware of their surroundings. In noisy environments listeners usually listen at least 6 dB louder than the environmental noise. City traffic is about 85 dB SPL, the noise on a subway train, or in a loud restaurant or bar can be over 90 dB SPL! Listening to music 6 dB louder than this background noise immediately puts the listener at risk of permanent hearing damage in a short time. Sealing, or "closed" earbuds should promote listening at quieter levels because they block some of the background noise – but many people still end up turning them up excessively loud.

KNOW HOW YOUR MONITOR SPEAKERS SOUND

- If cheap headphones or earbuds are what you are used to listening to, then you will need to retrain your ears by listening extensively on good stereo monitors.
- If you go straight from earbuds to good studio monitors without learning how the monitors sound, you'll end up trying to recreate the earbud sound on the monitors – and the result will be exponentially worse when your mix is played back on earbuds!
- If you're transitioning to new monitors, it's a good idea to have some respected commercial mixes spinning continuously while you mix, so you can A/B compare the characteristics of your mix to those commercial mixes as you are working.

2.4 Mono Compatibility

Checking your stereo (or immersive) project sounds good in mono – its *mono compatibility* – is vitally important. You never know what type of system your project will end up being heard on, or when the LR (or immersive) channels will be summed to mono for playback. Cheap or old TVs, mobile phone speakers, elevator music systems, bedside alarm clock radios, and ceiling speaker systems in shops and stores are potentially mono systems. The last thing you want is for your mix to sound bad, for comb filtering to occur on some elements, or in extreme cases for some parts to completely disappear when it is summed to mono.

Most mixing consoles and DAWs (Digital Audio Workstations) have a *mono* button in their monitor section. Use it every so often! In DAWs without a monitor section, it is necessary to insert a plug-in with a mono feature on the master fader to do this. When you switch from stereo to mono the mix will lose its width and become a narrow center phantom image – that's expected. What you're listening for are significant tonal or amplitude balance changes – if there are any, find the cause, and reduce them.

2.5 Compressed Audio Formats

Dissemination formats such as MP3, AAC, and Ogg-Vorbis, used by many music streaming and download services, are *lossy* compression codecs. They reduce (compress) the audio file size compared to PCM, by removing audio that the encoder doesn't think will be perceived. As a result, songs can be downloaded faster, or more songs fit onto your phone. The process is not transparent though, and more and more detail is removed as the bit rate is reduced. Lossy compression does negatively affect the quality of the sound.

LOSSY COMPRESSION ARTIFACTS

Some of the many undesirable byproducts of lossy audio compression codecs include:

- ► Frequency content discarded throughout the spectrum.
- ► Drastic high frequency loss at low bit-rates.
- ► MP3's reduction of bass frequencies.
- ► Time and phase smearing.
- ► Roughness and chattering.
- ► Ringing frequencies.
- ► Swirling and unstable higher frequencies.
- ► Loss of transient detail.
- ► Flattening of dynamics.
- ► Loss of reverberation tails, and reduction of other time-based effects.
- ► Stereo image narrowing and blurring.
- ► The addition of low-level noise.

Lossy compressed audio formats are not formats an audio professional, or aspiring audio professional should be listening to when learning what good sound might be, and how to achieve it.

Lossless compression codecs, such as FLAC and Apple Lossless (ALAC), are used by some higher quality streaming and download services. They do not reduce sound quality, however, they do not offer such large file size reductions.

While most modern portable media players (PMPs) can store lossless files of one type or another, or even uncompressed PCM files, their cheap convertors, amplifiers, and earbuds mean they're still not an acceptable solution. So, when listening, try to use good quality audio interfaces and good monitor loudspeakers (or headphones if loudspeakers are not available, or the listening room acoustics are bad).

Lossy formats should not be used for any system evaluation, or to analyze mixes – they do not accurately portray the mix as the engineer intended.

AUDIO EXAMPLES
Can be found on the companion website

Common Lossy Formats

These excerpts all feature the same musical example, encoded to different formats.

Example 2.1: Uncompressed, raw PCM.
Example 2.2: 256 kbps MP3. (amazon.com download standard prior to 2019.)

Example 2.3: The audio *removed*, and *distortions* added by the MP3 encoding in Example 2.2.
Example 2.4: 64 kbps MP3. (This quality and lower are experienced when streaming many services to mobile devices.)
Example 2.5: 256 kbps AAC+. (2020 Apple Music standard.)
Example 2.6: The audio *removed*, and *distortions* added by the AAC+ encoding in Example 2.5.

STREAMING AUDIO QUALITY IS IMPROVING, BUT NOT FOR VIDEO

There are now several online music streaming services offering lossless downloads, and even high resolution streaming at up to 24 bit/96 kHz and 24 bit/192 kHz. The situation is improving, and the number of vendors and amount of material offered in lossless and high resolution streaming and download formats is increasing every year. Some services even stream 3D immersive audio.

Unfortunately, most popular video streaming services tend to give whatever bandwidth increases they implement to the video picture – and the audio content is still supplied as very low quality lossy file types. If a video service says it is "HD" it is usually referring to the picture quality, *not* the audio quality.

2.6 Dynamic Range

If you compare a pop or rock recording from the 1970s to one from the late 2000s, one thing will probably be really obvious – the latter one is *louder!* The compact disc became the most popular dissemination format in the late 1980s, and the availability of the digital limiter in the mid-1990s prompted competition to have the "loudest" most in-your-face mix – the *Loudness War*. By the late 2000s, the average levels on CDs had become so high that distortion was apparent on many releases.

Some dynamic range compression and limiting are a desirable and essential part of the mastering process – particularly for commercial music. However, *hyper-compression* – when the peak levels are reduced too much, and end up very close to the more continuous average levels – produces a squashed, gritty, and often distorted sound, lacking true detail and punch. The transients – the initial attack portions of loud punchy sounds such as the kick and snare drums – are particularly prone to this type of sonic damage.

Forget about radio! AM is mono, and its frequency response and dynamic range are so limited, the sound is terrible! FM radio sounds better (but still not good), but radio stations employ equalizers, multiband compressors, stereo enhancers, aggressive limiters, and automatic gain control prior to their transmitter. So in addition to potential hyper-compression during mixing and mastering, the mix is squashed and processed even more by these often poor-sounding devices and processes. Most radio stations store their broadcast content as lossy compressed files nowadays, and digital and hybrid digital (HD) radio (both terrestrial transmitter and satellite stations) use further lossy transmission codecs.

Luckily there is a growing anti-hyper-compression movement, and more and more artists are bucking the trend. As discussed in the previous chapter, there is no longer an advantage to mixing and mastering a song to be as loud as possible, because streaming services and some media playback software will turn down albums or singles with LUFS values of higher than about -14. Legally mandated, average level standards have also been imposed on broadcasters.

Why does this matter? Casualties of the loudness war, or current hyper-compressed products are not considered desirable goals.

To learn, evaluate, and analyze good sound, you need to listen to industry-respected recordings, from uncompressed or lossless source formats.

AUDIO EXAMPLES
Can be found on the companion website

Dynamic Range Processing

These excerpts are all the same musical example, with different amounts of dynamic limiting applied during mastering.

Example 2.7: No limiting applied. The dynamics are as they came off the mixer output.

Example 2.8: Moderate limiting. About 6 dB of gain reduction on the loudest peaks, allowing the gain and average levels to be turned up by 6 dB. The mix is much louder and more powerful, yet maintains most of its punch and clarity.

Example 2.9: Loudness war hyper-limiting. Over 12 dB of gain reduction on the loudest peaks, allowing the average levels to be made over 12 dB louder. The mix sounds distorted, and the impact and power of the kick and snare drum transients are destroyed.

2.7 What About Distortion?

Is distortion bad? Well…yes. And no.

Distortion as a result of digital recording levels being too high, digital mix bus levels being over-hot, clipping processors or plug-ins – yes, that's bad. But as discussed earlier, gentle overloading of analog circuits and magnetic tape can produce pleasing effects. What would the sound of an electric guitar be without gentle, moderate, and severe forms of analog distortion, and digital pedal or processor emulations of those characteristics? What would bands such as Nine Inch Nails sound like without deliberate bit-crushing digital distortion?

Distortion is an essential part of some instrument sounds! A clean DI'd guitar track is usually out of place, inappropriate, and just does not work in a heavy rock or blues context – a stylistically appropriate guitar sound has either gentle distortion on the attacks of the notes, or more aggressive and continuous self-compressing distortion throughout the notes. The loudspeaker cone of a guitar or bass amp distorts the electrical audio signal fed into it, even if all the distortion knobs are set to "off." Non-linearities in the cone's behavior are technically distortion, and produce a unique character. The rotating horn in a Leslie speaker distorts an organ's output waveform, and again, this distortion *is* the instrument's sound and character – without it, a B3 just wouldn't be the same!

Subtle distortions are why one piece of gear sounds different to another, and can be used to give sounds a little extra character. Moderate or severe forms of distortion can be used for creative effect, giving elements of a mix hyper-character, grit and grunt when needed.

Once you have learned how to create clean, technically perfect mixes, carefully applied distortion can be a very powerful and effective tool. As part of the recording chain, it is fair to say that any accidental digital distortion, or much more than slight analog distortion, is undesirable. Software is available to undo distortion – but it's not free, and there's a limit to what it can do without sounding unnatural.

AUDIO EXAMPLES
Can be found on the companion website

Distortion – Friend and Foe

Example 2.10: Inappropriate distortion on an instrument track. The first section of this example is analog distortion caused by overloading the mic preamp. The second half is digital distortion caused by clipping the inputs (or outputs) of a digital device, DAW, or plug-in.

Example 2.11: A clean, undistorted guitar track. It is bland, and lacks character.

Example 2.12: The slight distortions of the guitar amp cabinet loudspeaker color the sound and make it more characterful.

2.8 What Is a Good Recording?

Musicians practice scales as a prerequisite to finding their own style. A similar prerequisite for an aspiring audio professional would be learning to make technically and artistically correct recordings and mixes prior to exploiting more creative stylistic techniques. "Artistically correct" means that the recording and mixing styles are appropriate to the project, and musical style. "Technically correct" means free of any undesirable technical artifacts. So what makes a technically good recording?

- Good sounding sources and musicians are essential.
- Appropriate microphone choice.
- Good mic placement and mic techniques.
- No noise or distortion problems created by incorrect or inappropriate use of any of the equipment in the recording and mixing chain.
- Good balances and use of the stereo soundstage.

By understanding how to capture, process, and mix sound using equipment technically correctly, you learn to:

- Use the equipment to control the sound.
- Really hear the effect of more creative, artistic use of the equipment.
- Anticipate how creative processing may benefit a project you're working on.

BUT IT'S ABOUT THE MUSIC, NOT THE RECORDING!

Let's not forget one important thing – the music creates a hit song, *not* the recording! Many hit records are not technically perfect – there may be minor engineering mistakes and errors because a great musical performance transcends a little distortion on a killer vocal take!

2.9 Accuracy

When learning to record and mix, one characteristic to aim for is *accuracy*. Is the recorded sound a faithful reproduction of the instrument or singer? If acoustic musicians are professional and are used to playing together, they know how to blend themselves. You, the recording engineer simply need to capture their performance appropriately. The recording room, mic choice, and mic placement are huge factors that impact the characteristics of a recording. Accuracy is unachievable if you place the wrong mic in the wrong position on an instrument in a bad room – all you can do is wrestle the sound somewhat into shape as part of the mixing process.

Great sound sources, a good sounding room, the right microphones, and good mic technique will capture sounds that mix themselves more. Getting the initial recording right results in a better, quicker, and easier mix. Capturing a sound accurately is difficult – after you know *you* can control the equipment, room, and situation to do this reliably (and not let the equipment dictate the goals) you can more deliberately develop your own personal style.

2.10 Non-Natural Sounds and Balances

Having discussed the importance of accuracy in a recording, it has to be admitted that most pop and rock music is not about overall accuracy – it's about achieving a sound that is stylistically and artistically appropriate. A good example of this is the sound of a modern rock drum set – the recorded kick, snare, and tom tom sounds are quite different to how they sound naturally from a normal listening position. Microphone and production techniques are exploited to make those sounds larger than life, phat, and in-your-face. They are *not* necessarily accurate, but they *are* stylistically and artistically desirable.

In real life, most singers would not be heard above the naturally much louder drums, guitar, and bass amps of a typical rock band. Recording equipment allows the engineer to create non-natural balances so the singer is heard clearly. Effects such as compression and reverb are used creatively during mixing, to improve the way the sounds work as a recorded piece of art. The balance and mix of sounds and effects we're used to hearing in pop music has little to do with natural balance, but everything to do with style.

"Accuracy" in pop music involves capturing the actual source material in a way that allows construction of stylistically appropriate sounds. If you are not familiar with either the natural sound of the instruments being recorded, or the stylistic goals of the type of project you're working on, how can you record appropriate tracks and create a good mix from them?

It is important to listen, listen, and listen to industry-respected recordings of styles you might one day be called to work on.

AUDIO EXAMPLES
Can be found on the companion website

Inappropriate Recording Techniques

Example 2.13: A vocal recorded in a room that is too live and reverberant for the project's style.
Example 2.14: A vocal recorded in a more suitable, drier vocal booth.
Example 2.15: Drums, recorded in a small dry room. The sound is small, compact, and lifeless.
Example 2.16: Drums, recorded in a live drum room. The sound is bright, punchy, and exciting.

2.11 What Are the Elements of a Good Mix?

Regardless of musical or production style, there are some fundamental characteristics that are essential and common to any good mix. They include:

- Appropriate frequency balance.
- Clarity and intelligibility.
- Effective use of the stereo image, and stereo imaging concepts.
- Effective use of soundstage depth, and front/back imaging concepts.
- Appropriate focus and amplitude balance.
- Good use of processing and effects.

2.12 Frequency Balance

A mix should have appropriate amounts of low, mid, and high frequency content. A *real time analyzer* (*RTA*) is a tool that can be used to visually show frequency content. **Figure 2.4** shows an RTA plot of a great sounding mix. The RTA is set to respond fairly slowly and average out the frequency content over fairly long intervals of time. If the RTA showed instantaneous readings, it would jump around too much, be difficult to interpret, and be dissimilar to human hearing (which tends to determine frequency balance based on longer-term averages).

An "ideal" frequency curve, based on current mixing trends, is for a slight hype in the low end, relatively "flat" content into the midrange, and a gentle roll off up into the highest frequencies.

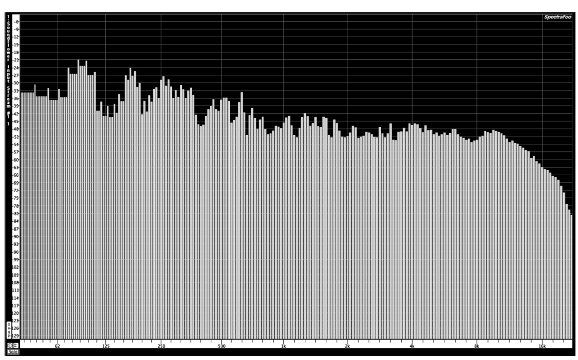

Figure 2.4 An RTA plot of a great mix. Note the relative "flatness" of the middle of the plot, with a gentle rise in the bass (on the left), and a rolling off of the high frequencies above about 6 kHz.

- If a mix doesn't have enough bass it will sound thin and lack power.
- Not enough midrange, and it will sound distant, thumpy, sizzly, and lack diction and clarity.
- Not enough high frequency content, and it will sound dull.

Visual tools *are just tools* to help train or confirm what the ear is hearing – your trained ear should *always* be the final judge.

A good way to confirm that your mix is on the right track is to have a respected, good-sounding mix available for playback, and to frequently and quickly switch to – A/B it with your mix in progress. The commercial mix will probably be louder and more "in-your-face" because it has been mastered, but you should listen to frequency content, balance and clarity, not overall loudness.

Your choice of microphone, the position you put it in, and the room you record in are just a few of many variables that impact the frequency balance of the sounds you record. It is essential that you record a sound that will give you what you need in the mix.

A mix with too much low and high frequency content, and not enough midrange will be boomy, bright, and sizzly, and lack definition. This is what is referred to as the *smile curve* or *loudness curve* sound, and is commonly dialed in by many consumers on their equalizers. It can be instantly pleasing because it hypes the extreme frequency ranges our ears are less sensitive to, particularly when listening at lower volumes. But at higher volumes, it creates an inappropriate frequency balance.

A mix lacking bass and high frequencies will sound thin, AM radio-like, and similar to cheap earbuds! If you only listen on earbuds, this is the type of sound you might think is

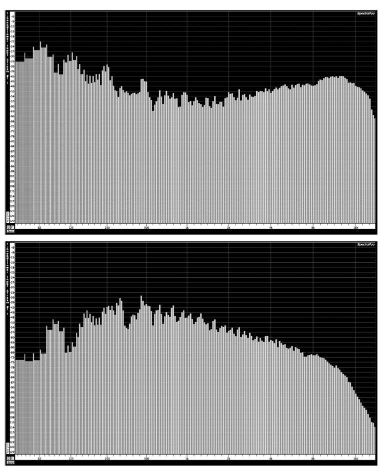

Figure 2.5 **Top:** An RTA plot of a mix with too much low and high frequency content. Note the areas of "smile curve" hype in the low (left) and high (right) extremes. **Bottom:** An RTA plot of a thin mix that is lacking low and high frequencies.

desirable – because it's the way you're used to hearing things. If you do think cheap earbuds sound good, invest in some good loudspeakers, throw away the earbuds, and re-train your ears to recognize the characteristics of a good mix on good loudspeakers – before you record or mix anything else!

SOME ESSENTIAL FREQUENCY CHARACTERISTICS TO LEARN TO IDENTIFY

- 150 Hz and below: Is the low frequency content of the kick and bass guitar/bass line boomy and undefined, tight and full, or small and compact sounding?
- 150 to 400 Hz: Does too much of the low-mid frequency range cloud the mix and make it too thick and muffled? Or does the mix lack fullness and body because there is too little of this frequency range?
- 400 to 500 Hz: Too much of this frequency range can give the mix a confused, boxy sound. Try banging a cardboard box with a stick, or talking into a cardboard box to get an idea of this quality!
- 600 to 800 Hz: Not quite a nasal sound, but like talking into a toilet roll tube. Too little of this frequency range, and many sounds lack definition. Too much of this range results in a honky sound.
- 1 to 1.6 kHz: A lot of nasal, diction, and definition characteristics are contained in this frequency range. Too much of this range can result in a thin, nasal, AM radio, megaphone-like sound. Not enough, and a "smile curve" type of EQ can result, where the mix lacks power and sounds distant.
- 2 to 3 kHz: Probably the ugliest sounding frequency range! Too little of it and the smile curve effects described above can result. Too much, and your mix could have a very cheap sounding, thin, tinny timbre.
- Around 4 kHz: "Not quite high frequencies" but "not quite mid frequencies"! Too much can give a harsh edge to instruments such as horns, guitars, drums, and cymbals. Too little of this frequency range can reduce clarity, particularly of vocals.
- Around 8 kHz: "Proper" brightness and high frequencies. Too little content in this range makes mixes dull and lifeless. Too much, and the mix can become over splashy, too bright, or too sibilant.
- Above 12 kHz: This range contains the "sprinkles" – the magic dust that can give your mix "air," and a sparkly sheen. Too little of this frequency range and the mix can sound flat and unexciting. Too much, and the mix will be too sizzly.

AUDIO EXAMPLES
Can be found on the companion website

Mix Frequency Balances

Example 2.17: A mix with a good, frequency response.
Example 2.18: A mix with "smile curve" EQ. Too many lows and too many highs create a mix that lacks definition.
Example 2.19: A mix with too much midrange lacks highs and lows. The mix sounds thin, like many cheap earbuds.

2.13 Clarity and Intelligibility

Every element in a mix should be able to be heard clearly, or appropriately to its artistic function. Clarity and intelligibility are products of many things, including:

- Overall frequency balance: If the overall frequency balance of the mix is incorrect, then that will negatively impact the listener's ability to hear everything appropriately in the mix.
- Frequency content of the individual elements in the mix: By nature of the format, a stereo mix places many constituent sounds in a relatively small space – the space between the loudspeakers, or between the headphone drivers. When sounds are positioned on top of each other, or in close proximity to each other, they become more difficult to accurately interpret, and their frequency contents sum together to form a new, combined frequency balance.
- Spatial positioning and panning: By panning sounds to different positions we physically separate them, making them clearer and improving the clarity of a mix – as well as making it more interesting to listen to.

The individual sounds or tracks in a mix might sound great when soloed, but when put together the sound can become muddled and unclear. Equalization (EQ) can be used to fix this. The unnecessary, or less important frequency components of a sound can be attenuated (de-emphasized) so that more important components of *other* sounds can occupy that frequency space. This enables both sounds to be heard more easily because they are no longer "fighting" with each other in that same overlapping frequency range.

Clarity is one of the hardest things to achieve in a mix. The ability to achieve clarity shows *you* can fully control everything from room to performer to mic to mix – rather than have the equipment and situation control you and dictate your objectives. It's an essential prerequisite before intentionally developing your own production style.

EQ-ING FOR CLARITY – GUITARS AND VOCALS

PROBLEM: A distorted electric guitar and a vocal may both have a lot of frequency content around 2 kHz. They get in the way of each other in this essential vocal diction and intelligibility frequency range. Neither is clear.

WRONG SOLUTION: Turning one up just further obscures the other, and/or makes the track too loud.

CORRECT SOLUTION: The guitar is less important musically, and also has a greater amount of other beneficially usable frequency content than the vocal – it has its body at lower frequencies, and brightness at around 5 kHz. An EQ attenuation around 2 to 3 kHz on the guitar can set it back in the mix a little, and make that frequency space available to more essential vocal diction and intelligibility frequencies. The problem is solved without turning anything up.

EQ-ING FOR CLARITY – KICK DRUM AND BASS GUITAR

PROBLEM: Kick drums and bass guitars frequently "fight" due to overlapping low frequency content. If a frequency range is congested, then all elements competing in that range lose, and none are clear.

SOLUTION: Attenuating the kick drum in a frequency range above its fundamental pitch (its "boom") can de-clutter frequency space that the bass can then occupy and therefore be more clearly heard. De-emphasizing the bass, around the frequency of the kick drum's fundamental boom will allow the beef and boom of the kick to be heard more clearly.

AUDIO EXAMPLES
Can be found on the companion website

Clarity Issues

Example 2.20: A mix lacking clarity. A product of poor mixing skills, this mix does not utilize the stereo image, and EQ is not effectively used to separate the elements of the mix.

Example 2.21: The clarity of the previous mix is improved through effective EQ, and positioning sounds more spatially throughout the entire stereo soundstage.

2.14 The Stereo Image

A good mix really takes advantage of the space between the loudspeakers. Leaving all the sound sources panned centrally, parked on top of one another, does not effectively use this space – it's boring, and things are difficult to hear and comprehend. Having sounds come from different locations in the stereo image not only makes the mix more interesting to listen to, but it also means that sounds come from their own spaces – and we can hear each element of the mix more clearly, and the clarity and intelligibility of the mix is improved.

2.15 Focus and Amplitude Balance

Poor amplitude balance between the elements in a mix will negatively affect frequency balance and clarity. For example, if the bass is turned up too loud it will obscure the kick drum, guitars, and keyboard sounds. It will also make the mix generally too boomy, and lacking in mid frequency punch and high frequency brightness.

Amplitude balance is dependent on frequency balance – a good amplitude balance cannot be obtained until a good frequency balance has been worked out between all the tracks. But a good frequency balance is reliant on the amplitude balance of the tracks! Everything is dependent on what you've done to it, what you've done to everything else, and what you haven't yet done to everything else! Every time you change an amplitude or EQ setting, you have to go back and re-evaluate its effect on every other track or channel you've already worked on. You must be prepared to rework every other track or channel every time you change something!

MIX EQ AND AMPLITUDE BALANCE

EQ and amplitude balancing sounds like a daunting combination of processes, but a helpful approach might be to:

1. EQ a sound so it sounds good by itself. Or better yet, EQ and balance (channel fader) a small group of related tracks so they sound good together. Don't get too hung up on how an individual soloed track sounds – even though something sounds amazing by itself, it probably won't fit in the complete mix, and may cause other elements to disappear!
2. Add and balance another track, or small group of related tracks – not all of the mix elements, but just more. EQ and balance them so they work with the track(s) worked on previously.
3. Listen carefully to ensure the previously worked on tracks still sound good since the addition of the new tracks. Also, if a new track just isn't clear, or "popping"

the way it should, the cause could be some of the previously worked on tracks getting in its way. Go back and further refine the EQ and amplitude balance of pre-existing or less important tracks with frequency content in similar ranges to the new track, that need it so the new combination of tracks all work together.

4 Return to step 2 and add more tracks into the mix.

But what is a good amplitude balance?

You need to be familiar with the typical amplitude balances of whatever style of music you're working on (or may unexpectedly end up working on). The desired sound and amplitude of the kick drum and bass in a jazz setting are very different to their respective sounds and levels in a rock project. The only way to learn what is desirable in a particular musical or production style is to listen to industry-respected recordings.

Additionally, a mix needs focus – a focal point. What is the most important element of the mix? In vocally driven music it is of course the lead vocal. If the lead vocal cannot be heard and clearly understood, either because of poor frequency balance, clarity, or amplitude balance problems, the mix has failed.

Trends in vocal mixing change over time – from "very loud and on top of the rest of the mix" pre-1970, to "in the mix" in the later 1980s, to "slightly on top" currently. The source recording needs to be made with the production aesthetic in mind – careful selection of the sounds recorded for the mix, and an effective musical arrangement have more impact on the mix than mix EQ (which can only help so much).

In pop and rock music, the drums drive the rhythm of most songs – particularly the kick drum and snare drum. Usually they are mixed at about equal levels, and with an equal focus to (or just slightly below) the lead vocal. For dance music they are often the loudest elements in the mix. If the drums are too quiet, the mix will lack rhythmic drive. If the bass is too quiet the mix will lack a solid foundation and "bottom."

There are different ways to develop a mix:

- Some engineers start with the drums, add the bass, add the other rhythm section instruments, and then park the focal point, usually the vocals, on top. This might be good for rhythm focused music.
- Other engineers start with the focal point, and then mix the other instrumentation in around it – perhaps the drums, then bass, then other rhythm section. This might be good for vocal focused music.

No method is right or wrong. *You* have to develop a process that works for you, and the type of project you are producing. Novice engineers should master the drums first method before moving on to the vocal first method. One thing is guaranteed though – you will end up with very different results using each of these methods. Try mixing the same song twice

– using a different method each time, make note of what you find yourself doing differently on the second version, and then compare the mixes to hear how different the end results are. Different projects, music, or production styles will benefit from different approaches.

2.16 Processing and Effects

Compression and reverb are like audio "glue" and "makeup." Compression can be used to tighten individual sounds, or to give them power and punch. It can also be used gently on the master output bus to apply a bit of glue that gels the whole mix together. Reverb can smooth over slight blemishes on individual tracks, and generally make things more pleasing.

We usually hear acoustic musical sounds in a room or hall – an enclosed space of some kind. That space imposes its reverberant characteristics on the sound. Some recording studio environments are acoustically dead and dry – particularly iso-booths and home studios treated with acoustical absorption products. When a microphone is positioned relatively close to a sound source in an acoustically dry environment, the recording lacks the acoustic reverberation that we are used to hearing – so artificial reverb can be used to put the sound source into a more characterful space.

The choice of reverb character is very style and tempo dependent – and it's true that some pop music mixing styles can be relatively dry sounding these days. But that doesn't necessarily mean that no effects are used – short room characteristics as opposed to long swishy halls are often used, and early reflections are exploited instead of reverb tails.

Drums sound dull and lifeless without the sound of a bright reflective room around them. This is why the best professional recording studios have dedicated drum rooms that are relatively reflective – so excitement and energy does not have to be added artificially. Vocals and solo instruments usually sound smoother, and more professional and polished after the addition of reverb. A reverb applied to an entire mix, either during the mixing or mastering stage of production, can act like "acoustical glue," gelling the entire band together in a similar environment.

Delays, choruses, flanging, distortion, and a multitude of other creative effects are powerful production tools that should also be used when making an artistically creative recording. Mixes without processing and effects are not as interesting as they otherwise could be!

AUDIO EXAMPLES
Can be found on the companion website

Processing and Effects

Example 2.22: A dry vocal, no processing.
Example 2.23: Compression makes the vocal more even, a little fuller, and bigger sounding.

Example 2.24: The addition of a short, early reflection-based reverb, increases the size and power of the vocal.
Example 2.25: A longer reverb tail acts like makeup, putting a professional sheen on the performance.
Example 2.26: The addition of a stereo delay creates a much more interesting image and effect.

2.17 Musical Arrangement and Song Structure

Like many things discussed in these early chapters, this has nothing to do with microphones or mic techniques! But the musical arrangement of a song can be the difference between a mix and song that works and sounds great, or one that doesn't reflect either the band or you as an engineer favorably! What note each instrument plays and with what tone, directly impacts the mix – and a mix can often be improved by reducing the number of instruments playing simultaneously, or by changing an instrument's part so it is not playing in the same range as another instrument. Musical arrangements should ideally be cleaned up in pre-production meetings well before any mics are set up in the studio. The best mics in the world, good-sounding instruments, and amazing musicians won't help you produce a great mix if the musical arrangement is poor.

In addition to affecting clarity, rarely does having all the musicians thrash away on the same riff or rhythm incessantly for an entire song produce an interesting record that a listener wants to listen to repeatedly. A great record is all about drawing the listener in – to do that there must be flow, development, tension, and release:

- *Boring*: A verse of vocal and full instrumentation followed by chorus of the same full instrumentation. In terms of intensity, the song hasn't gone anywhere and has nowhere to go.
- *Interesting intensity and textural change*: Having only the drums, bass, and keyboard in the mix for the verse, before adding the guitars for the chorus.

A great band, made up of experienced musicians, usually arranges itself. Different musicians and musical parts will come in and out effectively, and they will naturally pitch their parts so that they aren't in similar octaves or frequency ranges. Mixing is a lot easier when a band does this for you!

Less experienced bands might need some production assistance in order to get their songs work well as *recordings*. Less interesting songs and arrangements can sustain the audience's attention when performed on stage, because of the visual distractions of the performance – the musicians to look at, the lights etc. Songs and arrangements need to be expertly crafted in order to succeed in an audio recording where the sole focus is the music:

- If the guitar and keyboards are playing material that is too similar, it may be beneficial to suggest that one of the musicians tries something a little different. This could be as subtle as changing the synth sound, slightly varying the rhythm of a part, changing the octave, or even not playing that section.
- If a band insists on tracking everything (the same way they thrash through a song at a local bar gig), the mixer or DAW has *faders* and *mute* buttons! Just because something is recorded, doesn't mean it has to be used. Taking tracks out of the mix, or turning unimportant tracks down, can improve the clarity and focus of an otherwise cluttered, muddled, and over-busy mix.

How much you, the engineer, can suggest to the band musically depends on your relationship with the performers – and whether there is an actual producer present. If there is a producer, you should keep your mouth closed, and not have an opinion even when asked! If you are hired because of your relationship with the musicians, and you are respected for your opinions, or because there is no producer and you end up functioning as both engineer and producer, then it would be appropriate to educate the band, and shape their performance into a better recording – by suggesting performance, instrumentation, and arrangement changes.

If the band or producer is adamant that everything has to remain in, try to find time to do an extra mix, as you would prefer it to be – even if it's on your own (unpaid) time. Both can be presented to the client if you think they'd be receptive to your input. You never know, they might like the stripped down version – and if not, you have a better mix for your portfolio, or professional satisfaction! Regardless, good musicians and good people should be impressed that you cared enough about their project to go the extra mile.

2.18 Making a Great Record

A Great Performance vs the Best Technology

Time for a reality check!

The best equipment, and recording and mixing techniques can't turn a poor performance, or bad sounding instruments, into anything other than a "slightly improved" version. And pitch correction software can only do so much before the vocal starts sounding unnatural.

It is very frustrating trying to record and mix bad sounding instruments or singers! Corrective EQ can only do so much, and only fix certain problems – for example, a poorly tuned drum set, drum set rattles and buzzes, guitar or bass cabinet loudspeaker cone noises, or wrong inflections in the vocal performance *cannot* be fixed in the mix. If you find yourself spending a long time struggling to fix one specific problem, it could be caused by issues not related to recording or mixing. Solutions?

- Locate better sounding instruments.
- Exercise producer skills and coax, guide, and nurture the performer(s) to a confident performance.
- (Or get better performers!)

WHAT DO YOU REMEMBER ABOUT YOUR FAVORITE SONG?

Is it the pristine vocal sound recorded with a very expensive high end condenser microphone? Or is it the catchy hook and the words – which were captured with a cheaper dynamic mic, because that's what was plugged in when artist showed up and the engineer decided to hit the record button "just in case"?

For audio people it's probably (and hopefully!) a combination of both. But for most consumers, it's usually the latter. Of course, better equipment means better sound, and that has to be desirable – but the fact of the matter is that many top selling records have been recorded on less than the best equipment, and do sometimes exhibit technical problems that are easily overlooked in the context of a great musical performance.

While "we'll fix it in the mix" is certainly a myth, the goal of 100 percent technical perfection does have to be balanced with knowing when a magical and unrepeatable musical performance has been captured.

3 About Microphones, Part 1…

In This Chapter:

3.1 The Microphone

NOT NAMING NAMES…

This book isn't going to tell you which specific brand or model of mic to use for any given situation. There are too many options to list, and there's never only one correct way to record something. It should give you the knowledge to make good, educated mic decisions for yourself. However, there are a few applications for which available products are lesser-known or uncommon – so in those cases it seems sensible to name them.

There are *many* different microphones on the market for good reason – each model makes the source sound different. Selecting the right microphone is part of the art of recording – a mic's characteristics should complement the sound source being recorded, and capture the most appropriate sound for the mix.

A microphone changes sound waves (variations in air pressure) into electrical waveforms (variations in voltage). A mic's *capsule* contains the mechanism that does this. The capsule has either a *diaphragm* or *ribbon* in it – both are flexible membranes that move in response to the sound waves to which they are exposed.

Mics are the critical first stage of the recording chain. The sound quality of a recording is limited by the weakest link in the recording chain. There is no way that subsequent electronic manipulation can fully make up for inappropriate mic choice or poor mic technique – so it is essential to understand the different types of microphones available, and the characteristics of each.

3.2 End Address or Side Address?

Most microphones fall into one of two address types:

- *End address* mics, such as the one shown in the top of **Figure 3.1**. "On-axis," "the front," or 0° is the front "pointy" end of the mic.
- *Side address* mics, such as the one shown in the bottom of **Figure 3.1**, are addressed from the side, as the name suggests. Usually the front is the side with the company logo on it – but always read the product manual to determine this!

Neither of these types are better than the other – they just lend themselves to different orientations, which affects function, ease of set up, and visual aesthetics.

3.3 Directionality and Pick-Up Patterns

The *directionality*, *pick-up pattern*, *polar pattern* or *polar response* of a microphone determines how well it picks up sound coming from different directions. Two-dimensional diagrams are usually used to illustrate pick-up patterns – but remember that mics are placed

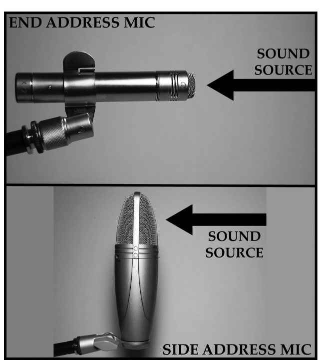

Figure 3.1 **Top:** An end address mic. **Bottom:** A side address mic.

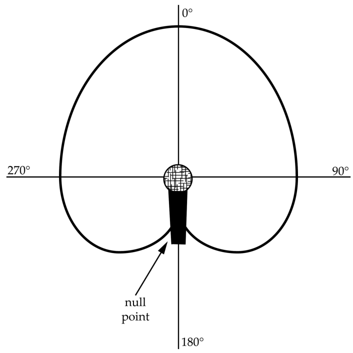

Figure 3.2 The pick-up pattern of a cardioid mic.

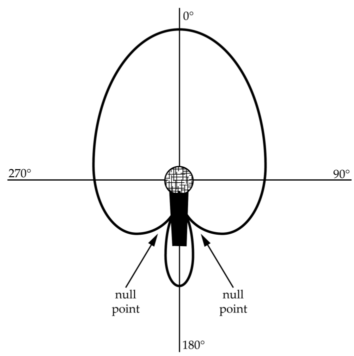

Figure 3.3 The pick-up pattern of a hyper-cardioid mic.

in three-dimensional environments, so the diagrams here should be interpreted as representing both the horizontal and vertical planes. The further the pick-up pattern plots are from the center of the crosshairs of the graph, the more sensitive the mic is to sound coming from that direction.

Cardioid Microphones

Cardioid mics are most sensitive to sound arriving *on-axis*, from directly in front of the mic. They are still relatively sensitive to their sides (and above and below) at 90° and 270°, and have a *null point* (the angle of least sensitivity, or most rejection) at 180° directly behind the mic, as shown in **Figure 3.2**.

Hyper-Cardioid Microphones

Hyper-cardioid mics, as shown in **Figure 3.3**, are less sensitive to sound coming from the sides (and above and below) than cardioid mics – although they do still exhibit some pick-up from those axes. This makes them more directional – they focus more on sound directly in front, and reject *off-axis* sound more than a cardioid mic. An unavoidable trade-off of the hyper-cardioid pattern is an area of sensitivity at 180°, directly behind the mic. This means the null points are 35° to 45° displaced from the rear of the mic.

SINGERS! DO NOT CUP THE MIC WITH YOUR HAND!

All *directional* mics achieve their directionality by allowing sound waves to enter the rear of the capsule. Covering the back, or rear vents of a directional end address mic (or the rear portion of a side address mic) prevents sound from entering the rear of the capsule – changing the mic's frequency response drastically. It will become very non-linear and peaky above 1 kHz, sounding thin and nasally. It will also become more omnidirectional.

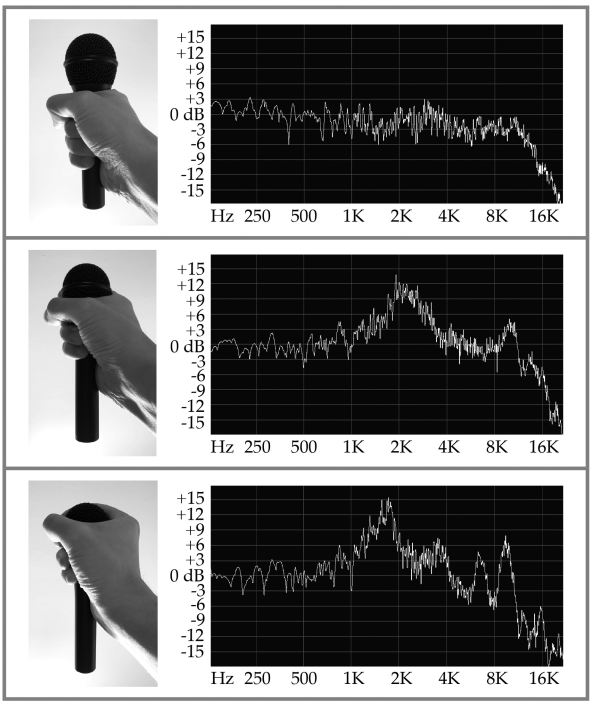

Figure 3.4 **Top:** Holding a mic by the handle produces the mic's true frequency response. **Middle:** Covering the rear vents produces peaks in the frequency response. **Bottom:** Fully cupping the mic produces an even more non-linear response – with larger and additional peaks, and some dips.

AUDIO EXAMPLES
Can be found on the companion website

Cupping the Mic

Example 3.1: A vocal, recorded with a dynamic microphone held properly by the handle.
Example 3.2: The same singer, same mic, the singer's hand cups the capsule.

Super-Cardioid Microphones

A *super-cardioid* pick-up pattern is more directional than a cardioid, but less directional than a hyper-cardioid. A super-cardioid mic has a lobe of sensitivity at 180º, but it is less sensitive on that axis than a hyper-cardioid mic.

Omnidirectional Microphones

Omnidirectional mics pick up sound from all directions more evenly, and are typically represented by a graphic resembling the shape in **Figure 3.5**. However, *all mics become more directional at higher frequencies* – so it is still important to point an omnidirectional mic on-axis to the sound source. An omnidirectional mic will pick up more *spill* (sound from adjacent sources, not really intended to go into the mic) and off-axis room reflections than a directional mic. This additional spill may or may not be desirable – it depends on the sound source, its role in the mix, the room, the type of project, and the desired production style and mix aesthetic. Omni mics have a more open and transparent sound than directional mics. Because they do not have proximity effect, they can generally be positioned closer to sound sources than directional mics, without sounding boomy or muddy. They also pick up lower frequencies than directional mics.

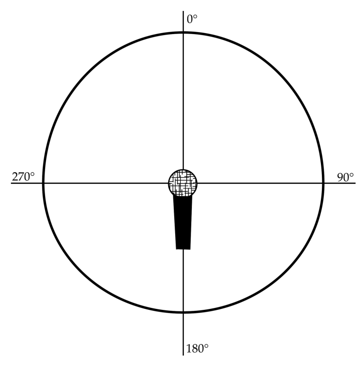

Figure 3.5 The pick-up pattern of an omnidirectional mic.

Wide-Cardioid Microphones

Wide-cardioid mics are less directional than cardioids, but more directional than omnidirectional mics. They favor sound coming from the hemisphere in front of the microphone, and are more sensitive to sounds coming from the sides (and above and below) than cardioid mics. Their null point is directly behind the mic. **Figure 3.6** shows the polar pattern of a wide-cardioid mic. Their sound tends to be a balance of the openness of an omnidirectional mic, with some of the directional control of a cardioid mic.

Bidirectional Microphones

Bidirectional or *figure-8* mics favor sound sources in front of *and behind* the mic capsule. They strongly reject sounds coming from the sides, above, and below – as shown in **Figure 3.7**. The extreme rejection of sound coming from the sides, above, and below the mic is a major benefit of this polar pattern. If there are desired sound sources behind the mic, the rear pick-up can be advantageous, but if the sound sources or spill coming from behind the mic are not

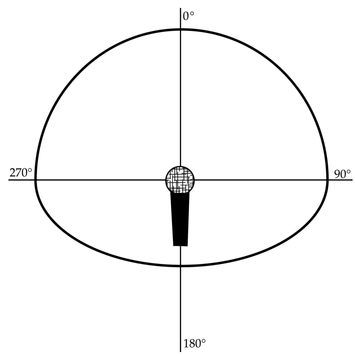

Figure 3.6 The pick-up pattern of a wide-cardioid mic.

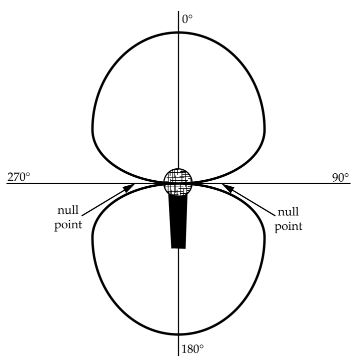

Figure 3.7 The pick-up pattern of a bidirectional mic.

desirable, the rear pick-up can be problematic. Some bidirectional ribbon mics sound slightly different on the front and back sides, providing different tonalities to explore.

Shotgun Microphones

Extremely directional, *shotgun* mics are rarely used in the recording studio, but are certainly used in television, film, broadcasting, live sound, and location/field sound recording – and they are discussed in a later chapter.

3.4 Dynamic Microphones

A *dynamic* or *moving coil microphone* features a thin, lightweight, usually circular, flexible diaphragm – which moves backwards and forwards when excited by the sound waves that travel towards it, as shown in **Figure 3.8**.

Attached to the rear of the diaphragm is a coil of wire. Surrounding the coil is a magnet, which is fixed in place. As the diaphragm moves backwards and forwards, the coil moves in the magnetic field. This causes a small amount of electricity to be induced in the coil of wire – the polarity is determined by whether the coil is moving backwards or forwards. This small electrical voltage is an analog of the sound wave that was picked up by the diaphragm – it is the same waveform, but as an electrical voltage rather than air pressure differences. A head-amplifier in the mic raises this voltage to a mic level signal, which then goes down the mic cable, and is amplified further by either a mixing console, outboard, or interface preamp.

Dynamic mics are relatively rugged – they withstand both sonic and physical abuse better than other technologies. For this reason, they tend to be the go-to mics for loud sound

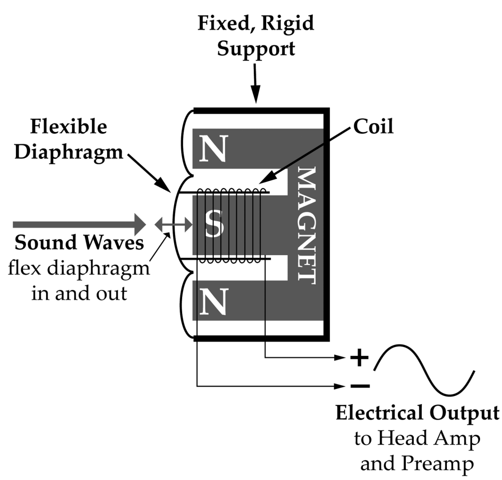

Figure 3.8 A cross section of a dynamic mic capsule.

sources and for situations where the mic may be at risk of being physically damaged. For example:

- Inside a kick drum.
- A couple of inches away from a guitar or bass speaker cabinet.
- In danger of being hit by a fast moving drum stick.
- In danger of being dropped by a singer on a live performance stage.

Generally, dynamic mics have relatively poor *high frequency response*, *sensitivity*, and *reach*. They struggle to pick up the top octave of sound, and often impose a somewhat distorted or fuzzy character on higher frequencies. They are not the best choice for sound sources with essential high frequency content (cymbals, hi-hats, acoustic guitars, etc.), nor if you want the brightest, crispest, smoothest, or most detailed high frequencies in your recording. The reason for this is that the diaphragm has to be made thick and strong enough to support the weight of the coil of wire, and it also has the added mass of the coil attached to it. This means that the diaphragm is not able to move as fast as it otherwise might, and it has more inertia – it is less able to stop and reverse direction to accurately represent the sound waves hitting it. It also takes more energy to set this heavier, less flexible diaphragm in motion, meaning that a very quiet sonic impulse (or the low level details within a louder sound) may not be strong enough to make the diaphragm respond – so those details will not be captured in the recording.

A dynamic mic would be a poor choice when recording quiet low-level sounds, such as acoustic string instruments or the sound of a pin dropping, or when a highly detailed recording containing all the subtle nuances of the sound is desired. Dynamic mics are like "low megapixel cameras" – they give an overall impression, but lack details and fine textural information.

Reach refers to a microphone's ability to pull sounds in from a distance. Because of their poor sensitivity, dynamic mics do not have good reach – they need to be close to a sound source. This makes them good for use in environments where there is a lot of undesirable spill – such as on a live performance stage, or within a compact miking environment such as a drum set. The lack of sensitivity and reach can be beneficial if a vocal (or other loud instrument) has to be recorded in a bad sounding room. Just make sure the mic is used close to the sound source to eliminate as much spill or room sound as possible – but do be aware of the negative effects of close miking discussed later.

3.5 Condenser (Capacitor) Microphones

Condenser microphones, also called *capacitor* mics, offer better high frequency response, and improved sensitivity and reach compared to dynamic mics. A diagram of a condenser mic capsule is shown in **Figure 3.9**.

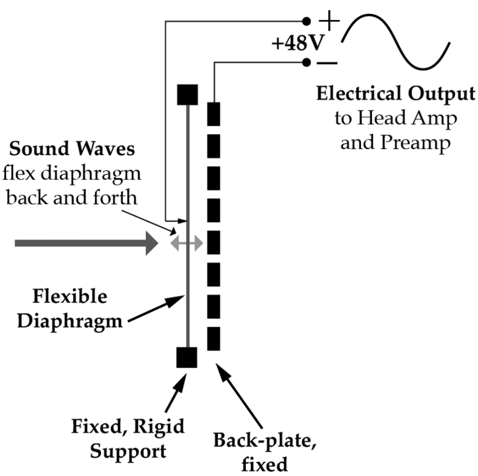

Figure 3.9 A cross section of a condenser mic capsule.

The diaphragm of a condenser mic does not have a coil attached to it, allowing it to be made thinner, so it has much less mass than a dynamic mic's diaphragm. Most condenser mic diaphragms are circular, but some mics have rectangular, and even multiple diaphragms. A condenser mic's diaphragm is coated with a very thin layer of electrically conductive metal particles – just a few microns thick – a process known as *sputtering*. Because it is so thin, the sputtered layer does not impede the diaphragm's ability to move, but it does allow a small amount of electricity to flow across the diaphragm.

ABOUT SPUTTERING

Gold is usually used for sputtering diaphragms – it oxidizes very slowly, giving the mic a long lifespan. A "gold sputtered diaphragm" is nothing unusual in condenser mics – more than 99 percent of condenser mics have gold sputtered diaphragms. Despite the hype some manufacturers give the process, gold sputtering is the norm.

A few microphones feature titanium sputtered diaphragms. Titanium is lighter, stronger, and oxidizes even slower than gold – giving the mic a slightly different character, and an even longer service life.

Behind a condenser mic's diaphragm is a *back-plate* – a rigid metal plate, with many holes drilled through it. The holes allow the air between the diaphragm and back-plate to displace more readily as sound waves cause the diaphragm to move backwards and forwards, as well as allowing sound from behind the back-plate to reach the diaphragm in some designs.

A steady DC (direct current) polarizing voltage known as *phantom power* is applied to the back-plate – by engaging a button labeled "phantom" or "+48V" on the mixing console or preamp. As the distance between the fixed back-plate and moving diaphragm changes, the capacitance changes, creating a modulating AC (alternating current) voltage at the capsule output. This electrical representation of the incoming sound waves is then amplified to a mic level signal by the head-amp in the mic, transmitted down the mic cable, and amplified to line level by the preamp the mic is plugged into.

The thinner, lighter, and more flexible diaphragm in a condenser mic can change directions and move faster and more easily that the thicker, heavier, “diaphragm plus coil” in a dynamic mic. This gives condenser mics the ability to pick up high frequencies more accurately – most condenser mics can efficiently pick up all the way to 20 kHz, and some to 30 kHz, 40 kHz, and even 50 kHz and beyond. It also means that very quiet sounds, and the low-level details contained within louder sounds, are picked up by a condenser mic. Condenser mics also have much better reach than dynamic mics – they pick up sound (including spill and room sound) from much further away.

Good condenser mics are generally more expensive than dynamic mics. They are also more fragile. Wind and air gusts can perforate diaphragms, and physical shock can damage capsules. *Windscreens* should be employed when condenser mics are used outside or in the presence of moving air currents, and *pop filters* should always be used between the mic and a singer close to the microphone.

Condenser mics are the “high-megapixel camera” of microphones – they pick up all the details, which is often a good thing. But some sounds can benefit from the less truthful characteristics of dynamic mics.

DIAPHRAGM SIZE

Condenser mics are additionally categorized by the diameter of their diaphragm:

- Small diaphragms are approximately ½ inch or smaller.
- Large diaphragms are approximately 1 inch or greater.
- Medium diaphragms are between these.

Small diaphragm condenser mics are generally more “accurate,” with less “character” than large diaphragm condenser mics. Small diaphragm condenser mics commonly have a more extended high frequency response than large diaphragm condenser mics – although there are many larger diaphragm mics that are very capable in this regard.

Large diaphragm condenser mics commonly have “bigger characters” and a tendency to “hype” the sound, making them less “accurate.” They make things larger than life – which is often desirable.

Medium diaphragm condenser mics can be generically described as combining these two traits – they retain some of the faithfulness of small diaphragm mics, but with some of the hype of large diaphragm mics.

Large diaphragm mics are generally quieter than small diaphragm mics – they make less hiss and background noise. Additionally, more acoustical energy is captured by the diaphragm and converted into electricity, therefore less (potentially noisy) gain in the mic's head amp and the preamp is required.

There are several manufacturers making microphones with smaller quarter inch diaphragms that have frequency responses up to 50 kHz and beyond. The trade-off is that they are noisier – the mic itself produces more hiss and noise. They often also require more preamp gain compared to a large diaphragm microphone. So, higher quality, quieter preamps are necessary in order to get enough clean gain, without adding the noise of poor quality preamplification to the mic's own noise.

Electret Microphones

Electret microphones are similar to condenser mics, but instead of requiring phantom power, they use a permanently charged back-plate, diaphragm, or capsule surface. The head-amps in the mic do require either phantom power or battery power, depending upon the mic. Because this voltage does not have to be 48 V, compact 9 V or 1.5 V batteries are commonly used. The current drain is very low, so the batteries last a long time.

In terms of performance, some electret mics come close to condenser mics, but most have an inferior frequency response, and are noisier. Electret mics can be used wherever a condenser mic is a good choice, with the exception of very quiet sound sources.

3.6 Single vs Dual Diaphragm Microphones

The pick-up pattern of a single diaphragm microphone is dictated by the capsule design. Most have a single, fixed pick-up pattern. Single diaphragm mics are generally considered truer to their ideal pick-up pattern than multi-pattern or dual diaphragm mics.

- Single pattern omnidirectional mics are more consistent and offer smoother off-axis frequency response compared to multi-pattern mics.
- Single pattern cardioid mics are more directional/less wide than multi-pattern mics.

So, single pattern mics can produce arguably "better" sound. But if the user needs a variety of different pick-up patterns available, the need to purchase multiple microphones makes them more expensive.

There are some single diaphragm microphones that offer a selection of cardioid and omnidirectional pick-up patterns. They achieve this through mechanical means – rear vents are open in cardioid mode, and a mechanical mechanism closes those vents to produce the omnidirectional pattern.

Dual diaphragm microphones have two diaphragms spaced closely behind each other in the same capsule. They can have either a single shared back-plate, or two separate back-plates, depending on the design. The outputs of the diaphragms are summed together electrically in different combinations to produce different pick-up patterns. Most multi-pattern condenser microphones do this summing in the microphone – there is a polar pattern switch on the mic, and a single standard 3-pin XLR connector. Some high-end multi-pattern condenser mics, and most multi-pattern tube microphones use an external box for this summing/pattern switching. Because two channels of information need to be transferred from the mic to the box, a proprietary XLR cable with more than three pins is used, and then a standard 3-pin XLR cable connects the box to the preamp.

There are also some single pattern mics that use dual diaphragms, and combine them in one fixed way to achieve the mic's specific pick-up pattern.

The most popular dual diaphragm multi-pattern design uses two cardioid capsules (one front-facing, the other rear-facing) in the same housing, positioned closely behind each other. By summing the capsules at different levels and with different polarities, a variety of polar patterns can be achieved, as shown in **Figure 3.10**.

- A cardioid pattern is produced by using the output of the front capsule only.
- An omnidirectional pattern is produced by summing the front and rear capsules at equal amplitude, with the same polarity. The less sensitive but overlapping rear half of each capsule's pattern sum constructively to boost the 90º/270º region and make the mic more equally sensitive even all around.
- A bidirectional pattern is produced by summing the front and rear capsules at equal amplitude, with the rear capsule polarity reversed. The overlapping rear half of each capsule's pattern cancel to produce the 90º and 270º nulls.

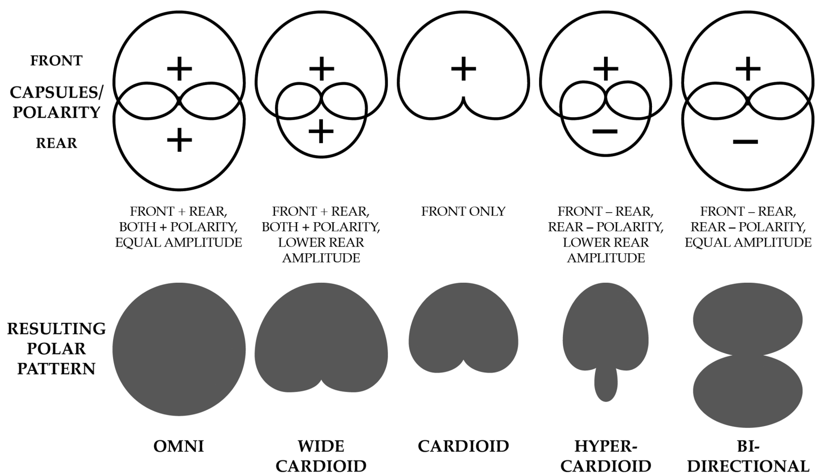

Figure 3.10 Dual diaphragm capsule summations and the pick-up patterns they produce.

- Wide-cardioid patterns are achieved by summing the rear capsule at a reduced level, with the same polarity as the front capsule. The rear portions of each capsule's pattern combine constructively to increase the side and rear sensitivity compared to a cardioid capsule.
- Hyper-cardioid and super-cardioid patterns occur when the rear capsule is summed at a reduced level, with the rear capsule polarity reversed. The overlapping rear half of each capsule's pattern partially cancel to reduce the 90º/270º sensitivity, but in the process the mic is made a little sensitive at 180º.

Some microphones have three-position switches to select cardioid, omnidirectional, or bidirectional patterns. Others have five-position switches, adding wide-cardioid and hyper-cardioid options. Some external boxes even have stepped or non-stepped rotary knobs which allow additional intermediate patterns to be selected, as shown in **Figure 3.11**.

3.7 Pressure and Pressure Gradient Transducers

A microphone capsule is a transducer that changes variations in air pressure (sound waves) into electrical voltage. *Pressure microphones* or *transducers*, and *pressure gradient microphones* or *transducers* are different types of capsules.

A *pressure microphone* is a single diaphragm omnidirectional microphone (**Figure 3.12 Left**). The diaphragm is only sensitive to sound on its front side, because the air behind the diaphragm is in a sealed capsule. Sounds from the sides and rear of the mic are also picked up because they wrap around the capsule to also reach the front of the diaphragm. The front of the diaphragm picks up a constructive sum of the air *pressure* variations from all around the mic. The capsule and microphone body block and attenuate high frequencies from the sides and behind, which is what makes omnidirectional mics more directional at higher frequencies.

Figure 3.11 **Left:** A three-pattern selection switch on a multi-pattern microphone. **Right:** A rotary switch on an external (tube mic power supply) box allows selection of a variety of intermediate pick-up patterns.

A *pressure gradient* microphone (**Figure 3.12 Right**) is a single diaphragm/ribbon *bidirectional* microphone. Both sides of the diaphragm are open, and it moves in reaction to pressure changes on both sides. The diaphragm outputs the sum of these pressures. Positive pressure in front of the mic causes the diaphragm to move inwards. Positive pressure behind the mic causes the same diaphragm to move outwards when viewed from the front – electrically at opposite polarity to sounds from the front. The difference, or the *gradient* between the two sides is what the single capsule outputs. Frontal source sound is usually significantly different to rear source sound, so the summation is constructive. Sounds from the sides of the capsule are picked up equally by both the sides of the diaphragm – but with opposite polarity on each side, which causes less or no diaphragm movement, significantly attenuating the amplitude of sounds coming from the sides.

Single diaphragm *cardioid* microphones are a combination of *pressure* characteristics at lower frequencies (they become less directional at lower frequencies) and *pressure gradient* characteristics at higher frequencies (they become more directional at higher frequencies). The rear of a single diaphragm cardioid mic's capsule is partially open – there are acoustic resistance materials in the rear opening, which cause high frequency attenuation and phase shift of frequencies below that. Single diaphragm cardioid, wide-cardioid, and hyper-cardioid microphones have small openings or "vents" in the mic body behind the capsule so that sound from in front, to the sides, and behind the mic can pass through them to the acoustically resistive capsule opening, through the small holes in the back-plate if it's a condenser mic, and reach the back side of the diaphragm. The microphone's directionality is created by frontal sounds summing constructively, and rear sounds summing destructively between both sides of the diaphragm. Front and rear sound must get to both sides of the diaphragm in order for the mic's directional pattern to be created. See **Figure 3.13**.

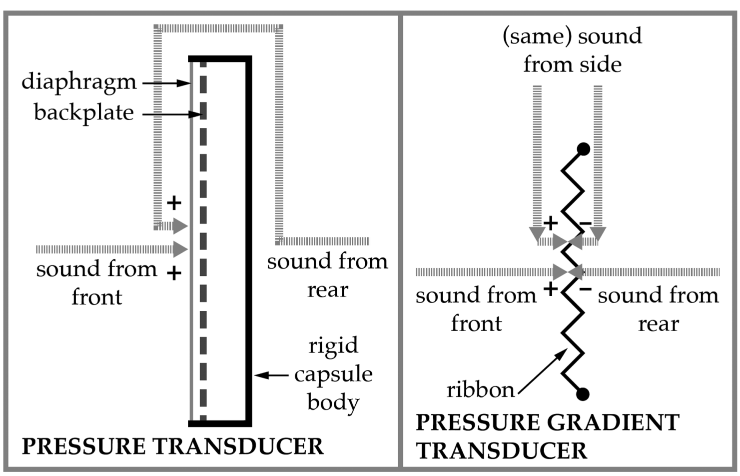

Figure 3.12 **Left:** An omnidirectional pressure transducer. **Right:** A bidirectional pressure gradient transducer.

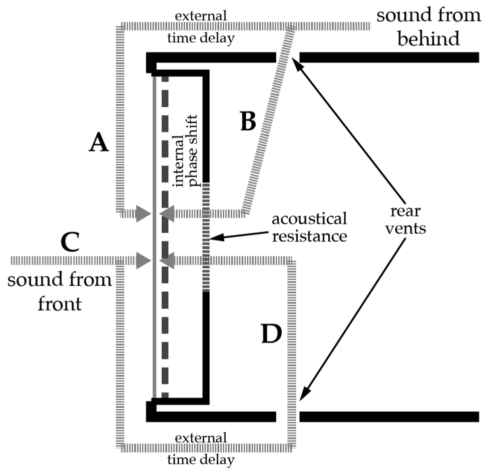

Figure 3.13 **A**: Sound coming from behind a cardioid capsule travels around the capsule to the front of the diaphragm, delayed/phase shifted by the extra distance it travels around the capsule. **B**: The same sound also enters the rear of the capsule but is phase (time) shifted by the acoustic resistance material it has to pass through, as well as the structure and geometry of the capsule. This causes the sound waves to arrive at both sides of the diaphragm, with the most phase difference – and similar positive pressures on both sides of the diaphragm, resulting in little or no pressure difference, so little or no actual diaphragm movement. *Sound coming from the rear is attenuated.* **C**: Sound arrives at the front of the diaphragm. **D**: It also travels to the rear of the diaphragm, taking additional time to travel the extra distance, and being further phase shifted by the acoustic resistance material before arriving at the rear of the diaphragm. The degree of phase shift causes low and mid frequency wavefronts to end up as positive pressure on the front side of the diaphragm, and at the same time, negative pressure on the back side of the diaphragm (or vice versa). Therefore they cause increased diaphragm movement and output. *Mid and low frequency sound coming from the front of the mic does not cancel.*

3.8 Ribbon Microphones

A *ribbon microphone* capsule features a folded "ribbon" of aluminum, anchored at the top and bottom, surrounded by a magnet, as shown in **Figure 3.14**. The ribbon moves backwards and forwards as it is excited by sound waves, and its motion within the magnetic field causes electricity to be induced into it. *Velocity microphone* is another term for a ribbon microphone – the speed of the ribbon's movement is responsible for the voltage generated, whereas in condenser mics, the capsule voltage is dictated by the amount of diaphragm displacement and not its speed.

In order to allow the ribbon to move quickly (as required for high frequencies), and easily when slight pressure changes excite it (for subtle details), the ribbon is very thin. Because of this, most ribbon mics are very fragile. Wind gusts, plosives (the rush of air from the mouth associated with "p" "b" and "t" type sounds), physical shock, and even being in close proximity to very loud sounds can cause the ribbon to stretch, perforate or tear

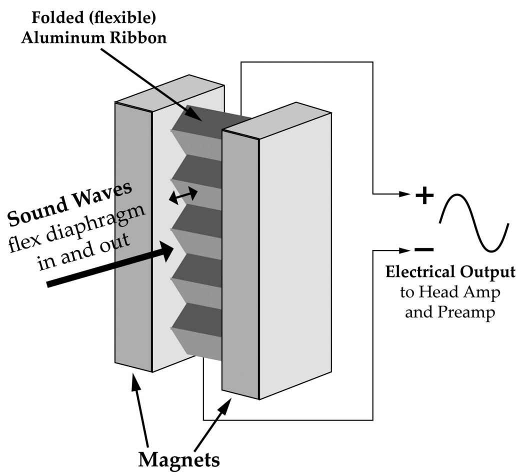

Figure 3.14 A cross section of a ribbon mic capsule. The ribbon is fixed and anchored at the top and bottom (not shown in this diagram). The magnet does not touch the flexible ribbon.

– requiring an expensive re-ribboning. Some modern ribbon microphones are more durable and, taking some precautions, can even be used on stage and in front of loud amp cabinets.

Phantom power can damage or destroy the ribbons in many *vintage* ribbon mics – but most *modern* ribbon mics will not usually be damaged by it, and some even *require* it in order for their internal amplifiers to function. However, ribbon elements *can be damaged* if phantom power is short-circuited to the mic – which *will happen* if quarter-inch, bantam/TT, or MIL patchbay connections between the mic preamp and the mic are plugged or unplugged with phantom power turned on, or there is a short in a mic or patchbay cable. If in a home or project studio there is no patchbay in use and all mic connections are XLR, phantom power typically presents no danger to modern ribbon mics – but do make sure by reading the manual first!

RIBBON MIC TLC…

If in any doubt – keep ribbon mics away from phantom power, away from singers unless pop filters are used, and away from live concert stages unless the environment is controlled and predictable.

You should be very cautious when using ribbon mics on loud sound sources, or even very close to sound sources you might not think are too loud – an SPL louder than the recommended maximum can destroy the ribbon by stretching or tearing it.

Increasing the distance from the sound source to the mic, and/or angling the mic up to 45° up or down from the source (which will of course change the sound), are ways of slightly reducing the SPL to which the mic is exposed.

Some condenser mics have a slightly harsh and brittle high frequency character. In the past, this would be softened by recording to analog tape. Digital recording doesn't do that. Ribbon mics have a smooth and silky high frequency character. Some ribbon mics, especially those with large ribbons, are exploited for their "vintage" character – they have a "dark and round" sound. Others have a more extended and flatter frequency response, but would be considered "warm."

IS THERE SIGNIFICANT SPILL FROM BEHIND THE MIC?

Most ribbon mics are bidirectional. It is important to consider this when positioning the mic.

Front or Back?

Some ribbon mics have slight sonic differences between their front and rear pick-ups. So, do not assume the "front" has to be the front – the characteristics of the rear of the mic may suit some sound sources. In a few designs, the rear portion of a ribbon mic's capsule is closed up, making it an omnidirectional mic, and several have dual ribbons to derive other directional patterns.

Vintage ribbon microphones, and modern recreations of vintage designs, have very low electrical output. This means that they need a lot of gain from the mic preamp. The cheap preamps found in many budget to midrange mixing consoles and interfaces don't have enough clean gain to get good signal levels from these mics – and if they do, they become very hissy and noisy. High quality, clean, quiet, and high-gain mic preamps are usually needed to get the most out of traditional ribbon mics.

In contrast, some modern ribbon mics feature high gain amplifiers in the mics themselves. This means they do not need such a high-gain preamp – so the mic *can* be used with less than top-of-the-line preamps. These active mics are more expensive than their classic counterparts, but you're paying part of the expense of an otherwise necessary and costly preamp when you buy the mic. These active designs *do* typically need phantom power to power the internal amplifier.

Several companies make in-line mic level boosters specifically to provide about 20 to 30 dB of clean gain to a mic signal, and enable low output ribbon and dynamic mics to be used with cheaper noisier preamps. Cloud Microphone's "Cloudlifter," and the sE Electronic "Dynamite" are two examples. They require phantom power to operate, but do not pass phantom on to the mic. These small XLR terminated devices should be plugged directly into the mic, or as close to it as possible, and the longer cable run made from the booster to the preamp.

3.9 Tube (Valve) Microphones

Guitarists have long favored the subtle harmonic distortions and warmth of *tube* or *valve* electronics in their amplifiers. Tube, or valve microphones, use sonically colorful tube amplification circuits instead of more modern, reliable, and stable solid-state head-amplification stages. This tube amplification colors the sound of the condenser capsule with fullness, warmth, and smoother high frequencies.

Original classic vintage tube mics, built using very high quality materials and tubes, were expensive. Faithful modern recreations of these classic designs often use "new old stock" tubes, and can be *very* expensive! But there are also many cheaper tube mics on the market. To keep the price down, some cheap mics use poor quality modern tubes and components – so assuming a tube mic is always better than a regular condenser mic is not correct. A classic tube mic will cost $5,000 to $10,000 (US$). There's just no way a $200 copy/clone/rip-off is going to have the same qualities, or automatically be the right choice just because it's a tube mic!

Additionally, due to the amount of character a tube mic produces, it is not necessarily the best choice for all sound sources or production styles. If a sound source is already very thick and warm sounding, the additional tube coloration may be too much. But if a sound source is a little thin, and needs some warmth, body, or "phatness," a tube mic could provide useful "sonic makeup."

As a tube ages, its characteristics can change – affecting a mic's performance and sound. Tubes do need replacing. They are fragile, and physical abuse will damage them – so care must be taken when setting up and tearing down. Tube circuits also tend to be noisier than solid-state electronics – making tube mics less suitable for recording very quiet sound sources.

For rock music, the warmth, "phatness," and "bigness" characteristics of analog tape have long been favored over digital recording systems by many engineers. With the current inaccessibility and expense of analog tape recording, and the prevalence of digital recording systems, some often desirable sonic colorations are missing from the recording chain. Tube mics, tube preamps, and tube and tape saturation plug-ins have been developed to add analog coloration to digital recording systems.

FETS, TUBES, AND TRANSFORMERS

All mics have electronics in them to amplify the capsule signal, and to change its output impedance. These amplifiers can have solid-state (transistor) circuits, or be tube based. The impedance matching circuits can be solid-state or transformer based.

- A *FET* (*field effect transistor*) is a specific type of transistor that behaves similarly to the amplifier *tubes* it started to replace in the 1960s. FETs are cheaper, more rugged, and quieter than tubes. Tubes produce a warmer sound, while FETs are less characterful (more truthful) and have better transient response.
- *Transformerless* solid state mics are "modern" designs – genericized as sounding "transparent" or neutral, with accurate and extended high and low frequency response. They are more forgiving of overloads than mics with transformers.

- Mics with output *transformers* pre-date solid-state designs, and can be considered "vintage" designs – transformerless outputs didn't start to become popular until the late 1970s. Not only old mics have transformers – many new mics available today are based on "vintage" designs, and have transformers. Transformers introduce frequency and time-based non-linearities to the sound of the mic – which with a good transformer we perceive as silkier or more airy high frequencies, and bigger, fuller low frequencies. They might also have other unique and often desirable characteristics, although their extreme low and high frequency performance may not be as good as transformerless designs. Transformer electronics overload at lower levels than transformerless designs, and can sound characterful when pushed a little too far, but are uglier in excessive overload situations. High quality (expensive) transformers are necessary for good sound, particularly in the high and low frequency extremes – cheap transformers degrade sound quality.

3.10 Stereo Microphones

Stereo microphones have two capsules housed in a single microphone body.

A format often found on hand-held portable recorders and mobile phone accessory microphones is an array of two cardioid microphones in an XY coincident pair or near-coincident pair arrangement. Some devices feature a middle-side array of a forwards facing cardioid microphone, and a sideways facing bidirectional microphone.

For professional recording, better quality full-size stereo condenser microphones should be used. Small diaphragm XY coincident, near-coincident, and MS stereo condenser microphones are available. Large diaphragm stereo microphones have two coincident capsules, some with independently switchable pick-up patterns, which can be rotated relative to each other to allow a variety of coincident arrays to be created. Many of these microphones have external power supply boxes that can be positioned conveniently in the control room so that patterns and array characteristics can be adjusted remotely.

Stereo microphones are of course more expensive than their mono counterparts, but having a single microphone body makes set up easy and less obtrusive (requiring less stand and mounting hardware), and guarantees well matched capsules. A stereo microphone is less versatile than two mono microphones however, because the stereo microphone is limited to coincident arrays.

3.11 Virtual Microphones

A *virtual microphone* system allows the user to change the timbre and tonal qualities of the recorded sound – to model the characteristics of other usually more expensive, vintage, or desirable microphones. Some systems also offer modelling of different polar patterns. All of this can be done *after the sound is recorded*, before or during mixing.

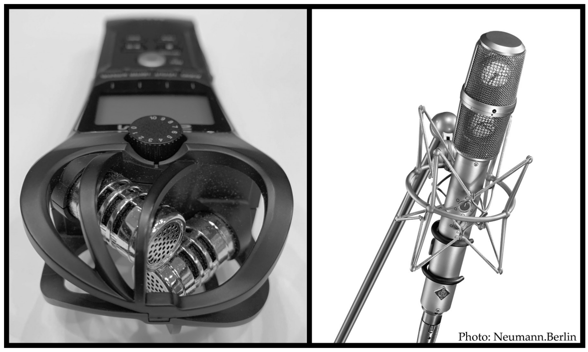

Figure 3.15 **Left:** An XY coincident pair on a hand-held recorder. **Right:** A large diaphragm stereo condenser mic, with rotatable capsules and switchable polar patterns.

Figure 3.16 An example of a virtual mic system plug-in, used to select the model, polar pattern, orientation, and other characteristics of the mic being modelled. This screenshot shows the plug-in in "dual" mode, where the single (dual diaphragm) virtual mic simultaneously models two different mics.

The best virtual microphone systems allow the user to do the following, *after recording*:

- Change the type of microphone.
- Change the polar pattern.
- Change the angle of the microphone relative to the sound source.
- Change the amount of proximity effect.
- Change the impression of the distance between the sound source and microphone.
- Create multiple virtual microphones from the single recorded channel.
- Increase or reduce room reflections (either reverb or undesirable colorations).
- Decrease spill in the recorded track.
- Create coincident stereo (and even surround) recordings from a single microphone.

Software-Only Systems

At the lowest cost, there are several software-only products that attempt to model different microphones from a selection of possible source mics. The limitations of these "software plus miscellaneous source mic" systems include:

- They model on-axis frequency response (and maybe phase relationships and other characteristics) but not the effects of a sound's directionality to the microphone.
- The results are unpredictable because of the large variations in the sound of "identical" mics produced by some manufacturers.
- The effect of proximity (distance cues or actual proximity effect) is rarely modelled, and if it is, its implementation may be very elementary.

Single Diaphragm Systems

Slightly better, but still lower priced virtual microphone systems use a custom single diaphragm microphone to capture the sound. The same company is in control of the hardware and software design, so the mic is perfectly paired with the software. The included microphone is built to tight tolerances, and designed and calibrated specifically for the system. This type of system allows the user to change the tonal qualities of the microphone after recording, but does not allow manipulation of the polar pattern.

Limitations of single channel modelling include:

- They model on-axis frequency response (and maybe phase relationships and other characteristics) but not the effects of a sound's directionality to the microphone.
- They have one fixed polar pattern.
- The system only works with the specific microphone sold as part of the system.

Dual-Diaphragm Systems

More expensive, but usually better sounding and more flexible large-diaphragm dual-diaphragm virtual microphone systems allow more complete emulation of the modelled microphone's characteristics, including:

- Different polar patterns (sometimes even including theoretical emulations of patterns the modelled microphone didn't have).
- The full 360° sound field of the modelled microphone.
- Proximity effect (increasing and decreasing), and the impression of the source sound's distance to the microphone.
- The direction of the sound source to the microphone can be changed, or optimized.

- Accurate and predictable tonal and polar response characteristics because the included microphone is built to very tight tolerances, and designed and calibrated specifically for the hardware/software system combination.
- Some systems even allow coincident stereo mic arrays to be modelled from a single microphone.

VIRTUAL MICROPHONES AND PREAMPS

The only variable unaccounted for with any of these technologies is the sonic character of the mic preamps – so it is best to use less characterful, transparent preamps for the most predictable results.

In order for a dual-diaphragm system to model correctly, both preamps should be the same, and the gains of both the front and back capsules need to be identical – *not* the signal levels on the meters when recording, but the amount of gain provided by each preamp channel. To make this easier to achieve:

- A preamp with stepped gains controls can be used.
- On a digitally controlled preamp, gains can be linked so they always match.
- A good virtual microphone system will have a calibration mode, which sends the output of the front capsule to both of the mic outputs temporarily, so that the gain settings of both preamp channels can be matched.

Virtual Microphones or the Real Thing?

Virtual microphones provide many tonal options at a great price – *much* less expensive than the tens-of-thousands-of-dollars of microphones they emulate. Are they exactly the same as the original mic? Well, which one? One problem with vintage mics is that many of the exact same model sound so different. A recording studio might have two of an "identical" microphone that sound quite different – so use each for different jobs. While a software recreation is never going to be truly faithful to the original microphone, depending on the system in use, a pretty convincing emulation can be achieved – for a much lower price than the usually very expensive originals.

Keeping an actual vintage mic in working condition (if you can even find a good one for sale) is a bit like maintaining an old car – work will be required as parts of it age and deteriorate, and some of those parts are very expensive and difficult to obtain. A new virtual microphone should give you at least a couple of decades of reliable service.

Software plug-ins are another issue… They are only compatible with specific host software that runs on specific operating systems. So there is no guarantee that the software will remain compatible and operational for as long as the microphone itself

– particularly if a smaller company disappears and the software is no longer updated and maintained.

Virtual microphones provide a great learning experience. They allow the user to explore the sound of many desirable microphones, compare them back to back, and to hear the effect of off-axis pick-up and coloration. Stereo pairs can also be used to examine the effect different polar patterns have on stereo mic arrays. Additionally, stereo pairs of virtual mics are more likely to be well matched, unlike pairs of vintage mics.

PRACTICAL EXERCISE

At the time of publication, Townsend Labs offers their Sphere plug-in and demonstration DAW sessions at no cost, after registering on their website. There are "mono microphone" tracks (which are actually stereo tracks containing the sound picked up by the front and rear mic diaphragms), true stereo pairs of tracks, and single mic tracks from which coincident stereo content can be derived.

www.townsendlabs.com/downloads

- From the *Vocals and other solo instruments* session, use the *Sphere* plug-in to experiment with mic choices and polar patterns on the soloed vocal track (which was recorded in a dead vocal booth). Notice the room acoustic doesn't change too much (because of the dead room), but the proximity effect and tonality does.
- From the *Drum Kit – OH and Rooms*, use the *Sphere* plug-in on both the L and R overhead (spaced) pair tracks. Change polar patterns and hear how the stereo image changes, and how the room sound/reverb affects the drum sound.
- From the *Grand Piano – Stereo* or *Acoustic Guitar – Stereo* sessions, solo the Sphere 180 track and use the *Sphere 180* plug-in to hear the effect of different polar patterns (and mic choices) on the single virtual microphone coincident array.

3.12 Other Microphone Technologies

Loudspeaker Cone Microphones

For decades, recording engineers have experimented with reverse wiring loudspeaker cones into mic inputs, and exploiting the exaggerated sound that such a large and heavy loudspeaker cone "diaphragm" produces. Commercial products are also available.

Loudspeaker cone mics lack high frequencies, and usually sound boomy, bassy, and mushy. However, their beefy bottom end sounds great blended with the detailed sound of a traditional mic. Is the resulting sound accurate? No, not really – but a desirably large, beefy, modern recorded kick drum sound is a non-natural, engineered sound anyway!

Figure 3.17 A commercial loudspeaker cone mic on a kick drum.

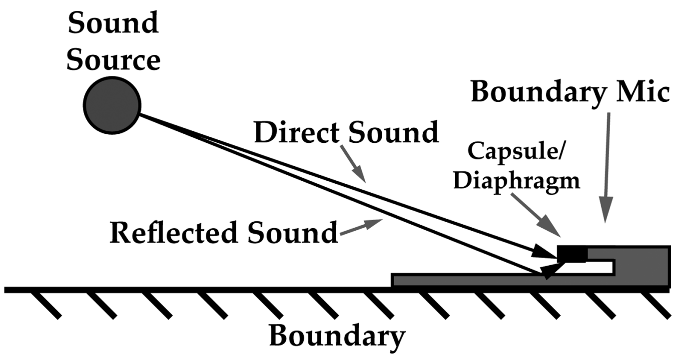

Figure 3.18 A cross section of a boundary mic capsule.

Boundary Microphones

A boundary microphone (trademarked “PZM” by Crown) features a very small condenser mic capsule mounted above, and facing down towards a sound-reflective boundary plate, as shown in **Figure 3.18**. Sound takes two paths to the mic capsule:

1 Directly into it.
2 Reflecting off the boundary plate up into it.

If a regular mic picks up both direct sound and the same sound reflected off the floor or a wall, the reflected sound arrives at the capsule at least a millisecond or so after the direct sound due to the extra distance it travels, and “phasey” comb filtering can be audible. Because of the extremely small distance between a boundary mic’s capsule and the boundary plate – usually around 1 mm (0.05 in) – the sound waves remain in phase throughout the audio frequency spectrum – so there is no comb filtering.

Boundary mics should be mounted flat, on large, hard surfaces (boundaries). The larger the surface area the mic is mounted on, the better the mic's low frequency response. A surface of at least several feet square is necessary for good bass response. Most boundary mics have a hemispherical pick-up pattern – like the "front half" of an omnidirectional pick-up pattern. Directional boundary mics (cardioid, or hyper cardioid) also exist, which reject sound coming from behind the mic to varying degrees. Hemispherical boundary mics do not color off-axis sounds in the way that other types of mic do. This means that sound sources in the entire front hemisphere of the microphone will be picked evenly, and that off-axis reflected sound (such as reverb) will be picked up more neutrally and accurately than by other types of mic.

Less commonly used in music recording, some applications include:

- Mounting inside pianos – on the lid.
- Mounting on walls for room and drum mic applications.
- Inside kick drum miking.
- Mounting on the front of theater stages, where the movement of performers and stage floor reflections would cause phase problems with other mic technologies.
- On desks and tables at conferences, meetings, and for phone and video conferencing.

Lavalier and Headset Microphones

Lavalier, *lav*, *lapel*, or *tie-clip* mics are very small mics that usually clip on to clothing, or inconspicuously mount on an actor's forehead, cheek, or in their hair. They are used to capture voices in the broadcasting, motion picture, and corporate sound industries, and for theater productions. *Headset* mics position the capsule directly in front of the mouth, so generally sound better than lav mics – but they are much more conspicuous.

Lavalier and headset mics are not the best sounding or quietest mics available, so they are not generally used for studio recording.

Clip-On Drum and Instrument Microphones

Many manufacturers now make dedicated drum and instrument mics featuring integrated clip systems to mount them directly to drums and instruments. These mics are primarily designed for the live sound industry, and not serious studio recording. While easy to position, they do not sound as good as stand mounted studio models, are generally noisier, and vibrations transmitted through the mounting clips can affect the sound. In some cases the clip can affect the sound of the drum or instrument. Some manufacturers make instrument mounting clips for their conventional mics – but for the best sound, with maximum isolation from undesirable vibrations, and maximum freedom to adjust the mic's position, a proper mic stand should be used with a good studio mic.

Parabolic Microphones

Not used in the recording studio, but frequently seen capturing distant sounds at televised sporting events, *parabolic* mics feature a condenser mic capsule mounted in a plastic "saucer" about 60 cm (2 ft) wide (similar to a small satellite dish). Parabolic mics are very directional, the dish giving them a very narrow "beam" pick-up.

USB Microphones

A USB microphone is an easy way to record audio into a computer, without needing a preamp or interface. The microphones have gain control and digital and USB conversion built in. Some even have headphone outputs. Most are mono, but there are also stereo models available. Most are large diaphragm condensers, and there are some dynamic models available too. They are plug-and-play, and very portable and convenient – however, there are limitations to their usefulness:

- They are low budget mics, suitable for tasks such as spoken word, podcasting, amateur voice-over, and home studio instrument recording.
- They are not going to produce the depth of sound or details that better mics will.
- It is difficult or impossible to use more than one mic at a time.
- If there is any chance that more than one mic will be needed for any work you do in the near future, it is advised to invest in a system that allows for growth – usually regular mics and an external interface.

For a podcaster, singer, or musician wanting to record amateur or demo tracks at home they are an easy convenient solution. For the audio professional or aspiring audio professional it would be better to invest in higher quality, more versatile systems.

AUDIO EXAMPLES
Can be found on the companion website

Different Microphone Technologies

Example 3.3: A vocal, recorded with a dynamic microphone.
Example 3.4: The same singer, recorded with a small diaphragm condenser microphone.
Example 3.5: The same singer, recorded with a large diaphragm condenser microphone.
Example 3.6: The same singer, recorded with a tube microphone.
Example 3.7: A kick drum recorded with a single kick mic.
Example 3.8: The same kick drum recorded with the same kick mic *and* a loudspeaker cone mic.

4 About Microphones, Part 2…

In This Chapter:

4.1 Phantom Power

Phantom power is a steady DC current, usually sent by mixers and preamps down the mic cable to power condenser microphones, active DI boxes, and other mics requiring power for their head amplifier. True phantom power is +48 V DC. If a piece of equipment is labeled "phantom power" rather than "+48 V," it can be an indication that it does not send the full 48 V – and possibly as little as 12 V or less! Most high quality mics do require the full 48 V to perform close to specification. Some mics can run on these lower voltages, but others will become noisier, and their high frequency response will be impaired.

Not all mixers have phantom power switches on all channels – they might be switchable in groups of four, six, eight, or globally. Phantom power presents no danger to dynamic mics or passive DI boxes – it can be turned on, and will have no effect on them. As previously discussed, some ribbon mics will be damaged by phantom power, while others can tolerate it, and some actually need it. Tube mic power supplies may or may not be damaged by phantom power – so it's best to check with the manufacturer. If your mixer switches phantom in blocks, you may be able to group your mics together in "phantom required" and "phantom not required" blocks.

4.2 Proximity Effect

Proximity effect is a boost of low and low-mid frequencies that occurs when a directional microphone is close to a sound source. As a general rule, the more directional the mic is, the more pronounced its proximity effect. This means that bidirectional mics have the most, hyper-cardioids a little less, cardioids a little less again, and wide cardioids the least. True omnidirectional mics should not exhibit proximity effect. Proximity effect can boost frequencies as high as 400 to 500 Hz, with the amount of boost increasing as the frequency gets lower – the amount of boost usually peaks somewhere between 100 to 200 Hz.

Figure 4.1 contains a frequency response chart that shows the *near-field* (close) characteristics of a directional mic with its associated proximity effect, and the *free-field* (at a distance) characteristics which are free of proximity effect.

Radio DJs, many announcers and emcees, and some singers use proximity effect deliberately to make themselves sound bigger, beefier, and boomier than they do naturally. Generally however, *proximity effect is a bad thing in the context of a busy music mix* – a frequent cause of muddy, confused mixes is low and low-mid frequency congestion, which

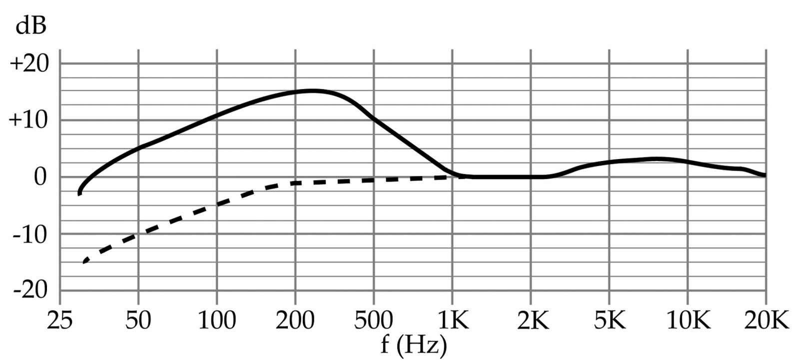

Figure 4.1 The dotted line below 1 kHz represents the flatter low frequency response of a directional mic when it is used in the "free-field" (at a distance). The solid line represents the low frequency boost of *proximity effect*, when the mic is used in the "near field" (close).

can build up cumulatively from just a little too much on multiple tracks. In many recording situations, particularly smaller rooms or home/project studios, this frequency congestion quickly becomes a problem because of the predominant and close-up use of cardioid and hyper-cardioid microphones. Moving a singer or instrument away from the mic by a few more inches (7 cm or so) can really reduce proximity effect problems.

Mics designed for close use, including hand held vocal mics, are built to mitigate proximity effect. This means that they sound great a few inches from a sound source – but sound thin and tinny if they are too far away.

AUDIO EXAMPLES
Can be found on the companion website

Proximity Effect

Example 4.1: A voice recorded accurately, or "flat," from 45 cm (18 in) away.
Example 4.2: The same voice recorded from 8 cm (3 in) away. Boomy proximity effect can be heard.

4.3 Frequency Response

Frequency response is a measure of how effectively a microphone is able to pick up different frequency ranges. There are several different ways of expressing frequency response, and the level of detail and usefulness of the quoted specification usually corresponds to the type, price, and quality of the mic.

"20 to 20 Hz"

Many cheaper mics simply state a range of frequencies. Unfortunately this supplies little useful information. Microphones pick up extreme low and high frequencies less well. This simple specification doesn't tell us how much less sensitive the mic has become at either extreme. Yes, the mic may pick up 50 Hz and 15 kHz, but it is probably doing so much less effectively than it picks up 2 kHz – you just don't know from this uninformative specification.

"20 to 20 kHZ ±3 dB"

Some mics quote an effective frequency range *and* a tolerance. This is more useful. It tells us that at 20 Hz and 20 kHz, the mic is 3 dB ("slightly") less sensitive than its average response between those extremes, and that within the quoted frequency range its sensitivity to a particular frequency could vary by up to 6 dB – it could pick up *frequency* **X** 3 dB hotter than average, and it could pick up *frequency* **Y** 3 dB less well than average. When looking at this type of specification, the smaller the tolerance, the flatter and less colored the mic's sound is – which can be a good thing. However this still does not tell us where those areas of higher and lower than average sensitivities are.

Generic Frequency Response Graphs

Frequency response graphs, such as those found in product literature and manuals, are much more useful. As shown in **Figure 4.2**, these are plots of frequency (x-axis) against a dB scale of relative sensitivity (y-axis), showing how a mic picks up the entire frequency range of an on-axis sound source. While a flatter mic might pick up a much more accurate picture of a sound, it does not mean it's necessarily the most suitable or flattering mic for every application.

Serial Number Specific Frequency Response Plots

The electronic components in a microphone are all built to specific tolerances. Tighter tolerances mean more identical behavior – and that multiple mics of the same model will sound more similar. But the fact of the matter is that multiple mics of the same model and vintages can sound a little, or even quite different, depending upon the tolerances they are built to. So, a generic plot printed in the product manual is not necessarily a completely accurate indication of the performance of the mic in your possession.

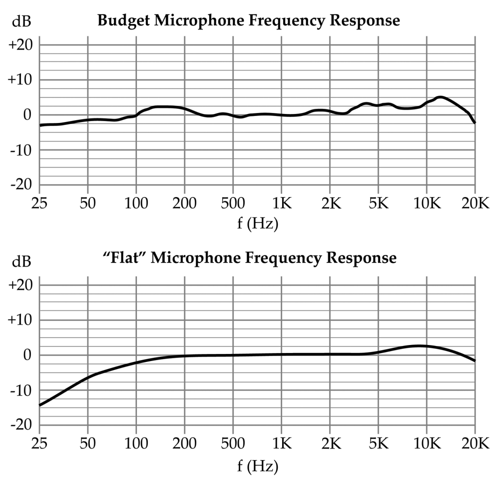

Figure 4.2 **Top:** A frequency response graph for a low budget microphone. The plot is not "flat." It has some peaks in its response where the mic is more sensitive – hyping that particular frequency range. **Bottom:** A "flat" microphone has a much more even, linear response to all frequencies in its effective pick-up range.

More expensive, higher quality microphones often include custom, individual serial number specific frequency response plots for each mic. For professional recording engineers this is valuable information. If you are buying multiple mics for stereo or multichannel arrays, then the mics should be closely matched. *Matched pairs* can be purchased from many manufacturers. They usually cost a little more than buying two unmatched mics, because of the extra testing and analysis required to match up the closest pairs – but for serious stereo pair work, this is money well spent.

4.4 Off-Axis Response

All microphones, regardless of their pick-up pattern, become more directional at higher frequencies, and more omnidirectional at lower frequencies. The plots laid out earlier in this chapter only show how a microphone responds to sound coming from directly in front of it, on-axis. In a multi-mic recording situation, or any live sound situation, there is also *spill* – unintended sound from adjacent sound sources reaching the mic off-axis from the sides or behind. Off-axis, a directional mic is much more sensitive to low frequencies than high frequencies. An omnidirectional mic is slightly less sensitive to high frequencies coming from the sides or behind. This means that off-axis spill can be muddy or thick sounding, due to the lack of higher frequencies that give the sound intelligibility and clarity. This colored spill causes a cumulative buildup of "mud" in projects recorded with predominantly cardioid and hyper-cardioid mics.

A major difference between cheaper, and more expensive higher quality mics, is the quality of the off-axis pick-up. A higher price point usually correlates to a more linear, better sounding off-axis response, which produces better sounding spill – and less off-axis mud and mush build up in a recording.

High quality microphones come with frequency response charts that show the mic's directional sensitivity at different frequencies, as shown in **Figure 4.3**. These charts are not usually serial number specific.

Different line styles represent different frequencies. A frequency's plot is only shown on half the diagram in order to allow additional frequencies to be shown on the other side – but remember, the mic exhibits symmetrical pick-up on both sides, and above and below the mic.

AUDIO EXAMPLES
Can be found on the companion website

Off-Axis Response of a Cardioid Microphone

Example 4.3: A voice recorded at 0°, on-axis.
Example 4.4: The same voice recorded at 90°, off-axis.
Example 4.5: The same voice recorded at 180°, in the mic's null point.

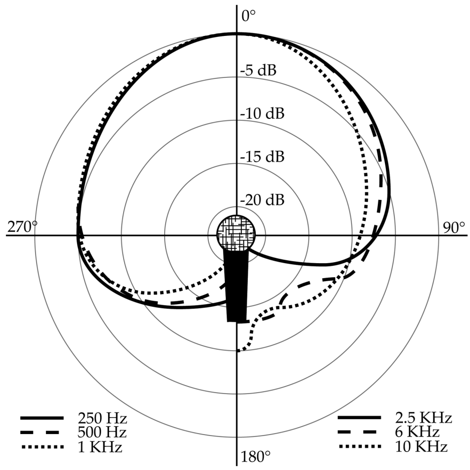

Figure 4.3 A mic's directional sensitivity is shown for different frequencies, each plotted hemispherically.

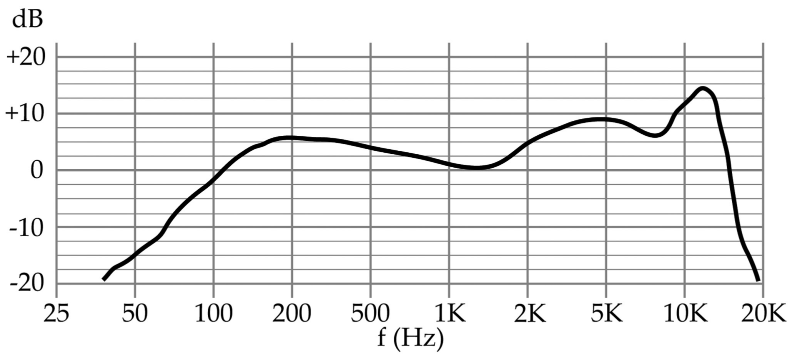

Figure 4.4 The frequency response graph of a mic with two presence peaks: A wide one at about 4 kHz, and a narrower one at about 11 kHz.

4.5 Flat Microphones vs Vocal Microphones

Microphones with a "flatter," more accurate frequency response are generally preferred for instrument, classical orchestra, and choir recording. *Diffuse field* mics, built for very distant use, are not completely flat, but have a slight rise in the high frequencies to compensate for natural high frequency attenuation over distance – this also avoids distant reverberant pick-up in a large concert hall sounding too dark or muddy.

Vocal mics are typically quite "characterful." This means they are not "flat," but that they hype and boost certain frequency ranges. There are usually one or two *presence peaks* in the 3 kHz to 10 kHz range, shown in **Figure 4.4**, which are areas of increased sensitivity which bring the intelligibility, diction, or brightness of a voice "forwards" in the mix.

Depending on the voice being recorded, microphones with a relatively low frequency presence peak (3 kHz to 5 kHz) can sometimes help a voice cut through a busy mix, or they can sound a little harsh – they are not suitable for all vocalists. A higher frequency presence peak (around 8 kHz) can produce a brighter, more shimmery and polished "studio" type sound. Presence peaks can conflict with *sibilance* – the sizzle of "S" and "T" sounds. If a voice is very sibilant, a presence peak in the same frequency range as the sibilance can unpleasantly over-exaggerate it.

No two voices or instruments sound the same, and ultimately the specific voice or instrument being recorded determines the character of the mic necessary. A mic that makes one artist sound great, may not work well on another. The only way to determine which mic brings out the most pleasing characteristics of any sound source is to experiment and try different mics and placements.

4.6 Low Frequency Response

Directional mics do not pick up very low frequencies as well as omnidirectional mics. This may seem counter-intuitive considering the low frequency boost characteristics of proximity effect, however it is important to remember that proximity effect only occurs *when the mic is close to the sound source*, and that this unnatural boost is inaccurate and can cause frequency balance issues that make the sound muddy and confused.

Even when used up close, with a large resulting proximity effect boost, the response of many directional microphones falls off dramatically below about 50 Hz. There's another very powerful octave of sound below 50 Hz! Additionally, when used at a distance from a sound source, the low frequency response of directional proximity effect compensated mics falls off dramatically below 200 Hz.

Omnidirectional mics do not exhibit proximity effect. This means that their low and low-mid frequency response will be the same regardless of the distance they are from the sound source. Also, and importantly, their low frequency response extends down to lower frequencies than directional mics. A good omnidirectional mic will have a flatter, more even frequency response down to 30 Hz, 20 Hz, or below! This means that it will pick up the low frequencies of orchestral bass drums, string basses, and bass guitar speaker cabinets really well. Most directional mics will not pick up the same weight, and if they are used up close that weight will be masked by proximity effect.

AUDIO EXAMPLES
Can be found on the companion website

Low Frequency Response

Example 4.6: Distant cardioid drum room mics.
Example 4.7: Distant omnidirectional drum room mics in the same room.

Figure 4.5 A low cut filter switch on a mic.

4.7 Low Cut Filters

Many microphones have *low cut filter* switches. They can be simple on/off switches (flat vs a fixed filter slope at a fixed cut-off frequency) or multi-position switches offering a variety of filter slope and cut-off frequency combinations, as shown in **Figure 4.5**.

Low cut filters can compensate for proximity effect, reducing the excess low frequency energy before the signal eats headroom in the preamp and recording system – improving the gain structure and level of useful sound recorded. They can also remove low and low-mid frequencies that commonly clutter a mix and are frequently attenuated when mixing. But be sure to research the cut-off frequencies and filter slopes of each setting offered:

- Some settings are designed with a gentle roll off starting in the low-mid frequencies – which will definitely reduce proximity effect muddiness.
- Others roll off more aggressively, with steeper slopes, but below a much lower frequency – reducing only the super-low frequency rumbles below the frequency content of the actual sound source.

The low cut filter can also be used to reduce (but not usually eliminate) mechanical rumbles, booms, and vibrations transferred into the mic through the mic stand, and rumbles produced by wind and moving air currents – but it is always best to prevent these from entering the mic by using supports that mechanically isolate the mic more, or positioning the mic so it is not in moving air currents. Not only will the recorded results be better, but the diaphragms in the mics will be safer from potential damage too!

4.8 Low Frequency Instrument Microphones

Bass and *low frequency instrument* mics are typically dynamic mics, with frequency responses heavily contoured to make sound sources like basses and kick drums more instantly appealing. Generally this means that the low and high frequencies are hyped, and there's a big cut in the mids. These mics can produce excellent results, however they have such exaggerated characters that they either really work or really *do not* work with a particular instrument or drum – there's little middle ground.

It is always preferable to get a great sound by selecting the right mic, and put it in the correct position on a great sounding instrument – rather than to use console or DAW equalization to compensate for preventable and avoidable deficiencies. It can be beneficial to use a low frequency instrument mic rather than equalize a flatter mic into shape if you're using low quality EQ circuits or plug-ins. But if the EQ curves built into this type of mic are not exactly what is needed for a specific situation, you don't want to have to use EQ to undo the effects of the mic, while trying to apply desirable corrections. A more generic, flatter mic may not instantly sound as "big" as a low frequency instrument mic, but you may be able to get better results using some good EQ to tweak the flat mic's sound into shape.

AUDIO EXAMPLES
Can be found on the companion website

Low Frequency Instrument Microphones

Example 4.8: A kick drum recorded with a general-purpose dynamic mic.
Example 4.9: The same kick drum recorded with a low frequency instrument dynamic mic.

4.9 Sensitivity

Microphone *sensitivity* or *transfer factor* is a measure of a mic's electrical output, when it is placed in a sound field of a specific SPL, usually 94 dB SPL. Most manufacturers quote sensitivity as the number of millivolts the mic produces for this reference level (which can be expressed as *mV/Pa*, or *mV per 10 µbars* – the two units are interchangeable). When sensitivity is expressed in one of these formats, higher numbers are usually more desirable – the higher the quoted number, the less preamp gain will be required to bring the mic's signal up to line level, minimizing any noise potentially introduced by lower quality preamps. So, a mic with a sensitivity of 29.8 mV/Pa is a higher output mic than one with a spec of 5.6 mV per 10 µbars.

Some manufacturers express sensitivity in a "–dB" format, and others in the older level standard of a 74 dB SPL input level, 1 µbar, or 0.1 Pascal – those last three unit values are identical. The 94 dB equivalent is expressed as *dB re 1V/Pa*, and the 74 dB equivalent as *dB re 1V/µbar*. Expressed these ways, microphones with greater electrical outputs will have smaller negative numbers – so for recording very quiet sound sources, a mic with a rating of −31 dB re 1V/Pa would be preferable to one with a rating of −45 dB re 1V/Pa.

To compare two mics, when one quotes sensitivity as "–dB" to a 94 dB equivalent, and the other as "–dB" to a 74 dB equivalent reference level:

94 dB reference equivalent = 74 dB equivalent – 20 dB

74 dB reference equivalent = 94 dB equivalent + 20 dB

Greater sensitivity means less preamp gain and noise, but does not necessarily mean that the mic sounds "better." If you expect to record quiet sound sources, or sources with wide dynamic ranges, and you don't have clean, quiet, high gain preamps, higher sensitivity specs can help you produce technically better recordings. For mics which do not need a lot of gain from the preamp (drum and amp mics for example) sensitivity specs are less important.

4.10 Self-Noise and Equivalent Noise Rating (ENR)

Self-noise and *ENR* are the same thing – a measure of the background noise (hiss and rumble) that a mic makes itself. Specifically, it is a measure of the sound pressure level required to create an output level the same as the noise the mic outputs when it is surrounded by silence. Self-noise is usually quoted as a dBA figure. Good quality modern mic preamps and digital recording systems can be very clean and quiet, and are able to resolve the smallest sonic details, so self-noise can be very important – particularly if quiet sounds are being recorded, because the self-noise correlates to the lowest extreme of a mic's effective dynamic range.

Self-noise is only quoted for condenser microphones, where it is a measure of the noise created by the head-amp in the mic itself. This noise is usually greater than the noise introduced by high quality preamp, and can be audible when the gain is turned up high when recording quiet sound sources. Most of the noise in dynamic and ribbon mics is created by the electrons moving within the coil or ribbon. Because the output voltage of these mics is much lower and more gain is needed on the preamp or console, the noise generated by the preamp is louder than the self-noise of the mic itself. This makes the self-noise irrelevant, but makes having good quality, quiet, high gain preamps very important.

Self-noise figures can be lower than 5 dBA for modern large diaphragm condenser mics. Small diaphragm mics tend to be a little noisier, at 12 to 18 dBA. Very small diaphragm mics are usually noisier still, with specs between 22 and 27 dBA. Self-noise in the upper 20s can definitely be noticeable when recording quiet sources, or recording from large distances. If you know that you are going to record quiet or wide dynamic range sources, or do classical instrumental or choral recording, low self-noise figures are important, and mics with lower self-noise figures are definitely preferred. But for most rock, pop, and loud instrument recording, self-noise figures are less important than the actual sound of the mic. Using noisier mics as drum overheads for example, would present no problem – not a lot of gain is required due to the high SPL of the sound source, so the self-noise would be insignificant.

4.11 Signal-to-Noise Ratio

Directly related to self-noise, the *S/N ratio*, or *SNR* of a microphone is a measure of the amount of signal output, compared to the level of mic-induced noise present, when the mic is in a 94 dB SPL sound field. Higher numbers indicate better performance.

The following equations are useful to convert between self-noise and SNR, if you are trying to compare a mic that has only a self-noise spec with a mic that has only an SNR spec:

$$\text{SELF} - \text{NOISE} = 94 - \text{SNR}$$

$$\text{SNR} = 94 - \text{SELF} - \text{NOISE}$$

4.12 Pads

Condenser mics often have *pad switches* or *pads* (labeled "–10 dB," "–20 dB," etc.), which attenuate the mic's capsule output so loud sounds do not overload the other internal electronics of the mic, or the preamp the mic is plugged into.

Some preamps and mixing consoles have pad buttons. These may prevent the preamp from overloading, but it is still possible for a loud sound to overload the mic's electronics. If you are recording a very loud singer (or other sound source) with a condenser mic and are hearing distortion, but all your preamp, input and output levels look good, and a mixer or preamp pad is engaged – it's probably the mic distorting. Engaging the mic's pad, and not the mixer/preamp pad should fix this situation – as long as the sound source is not in excess of the mic's maximum SPL level.

4.13 Maximum SPL

If a mic is overloaded by being put in front of a sound source that is too loud, it will distort. *Maximum SPL* is a measure of the SPL that produces a slightly distorted output from the mic – usually quoted for 0.5 percent THD or 1 percent THD (total harmonic distortion, a measure of the amplitude of the additional harmonics of the overload distortion, compared to the original waveform).

A drummer may sit in a sound field of greater than 110 dB SPL when playing – even though it is dangerous to listen to sounds above 100 dB SPL for any length of time. The actual SPL at a mic a couple inches from a drum can be in excess of 130 dB SPL – so high SPL capable mics are definitely necessary. Maximum SPL figures are quoted for all pads on the mic engaged.

4.14 Dynamic Range

A microphone's dynamic range is the range of sound pressure levels over which the mic will effectively pick up sound. It is the difference between sound being buried and masked by the mic's self-noise, or the diaphragm's inability to resolve low level sonic impulses, and the onset of distortion at the mic's maximum SPL. Greater dynamic range specs are usually desirable.

4.15 Transient Response

Transients are the first few positive and negative peaks of the attack of a sound. They contain a lot of important data our brain uses to identify the sound. Sounds typically described as having important transient details include percussion instruments (drums and pianos) and plucked string instruments (particularly acoustic guitars). To record these sounds with the most accuracy, transparency and "zing," a mic with good *transient response* is desirable. The diaphragm of the mic needs to be able to move fast and change direction easily to respond to positive and negative transient peaks, and the subtle nuances contained within them.

Dynamic mics, with their heavier, more sluggish diaphragms, cannot respond fast enough to accurately capture initial transients – so the attacks of sounds become compressed and subtly distorted. In terms of accuracy this is undesirable, but compressing the amplitude of the initial and loudest peaks of a sound means that preamp gain is turned up more to achieve good signal levels – increasing the sound's average level. This results in a less transparent, but bigger, phatter, and more powerful sound that can be desirable for rock drums, electric guitars, and basses.

Manufacturers making condenser mics with ultra-small diaphragms usually do so because their design philosophy emphasizes the extended high frequency response *and* improved transient response that is only possible with ultra-small diaphragms. Many small diaphragm mics are also respected for their transient response. There is no specification for transient response though. It's not as simple as "the smaller the diaphragm, the better the transient response" – the head-amp electronics also have an effect. Some manufacturers purposefully design and promote the "fast response" of their mic's electronics. If you are unable to audition a mic before purchasing it, the best way to judge its transient response is to read non-biased, reputable professional reviews.

AUDIO EXAMPLES
Can be found on the companion website

Transient Response

Example 4.10: Acoustic guitar, recorded with a small diaphragm condenser mic known for its good transient response.
Example 4.11: The same acoustic guitar, recorded with a small diaphragm condenser mic with less good transient response.

4.16 Pop Filters, Windscreens, and Dead Cats

Pop Filters

Pop filters, *shields*, or *screens* are essential when recording the human voice (singing or speaking) from within 60 cm (2 ft), with the talent aiming at the microphone, as shown in **Figure 4.6**.

Plosives are created when a burst of air associated with B, D, G, K, P, and T sounds leaves the mouth. If this air current hits a mic's diaphragm, it can cause loud pops and distortions, and potentially damage the diaphragm or ribbon. For better sound, studio mics do not usually have built in protection from these stray air currents.

A *pop filter* is an acoustically transparent nylon material or perforated metal screen which is positioned between the talent and the mic. Nylon pop filters have two layers of material which diffuse the plosive air currents randomly so they do not reach the mic. Metal pop filters redirect the air current in a specific direction – above, below, or to the side of the mic, depending on the orientation of the filter.

In order to work properly, there should be about 6 cm (2 in) between the mic and the pop filter, and the same distance between the pop filter and singer's mouth. Goosenecks wear out quickly, particularly if they are aggressively bent in order to position the pop filter. It is often easier to position a pop filter when it is attached to a separate stand, and not the same stand as the mic.

Most musical instruments do not propagate plosive air currents, so pop filters are unnecessary. If a studio vocal mic is positioned at above or below the singer's mouth, and at least 30 cm (1 ft) away, making sure the singer does not sing "into" it but "across the top" of it, the plosives are directed away from the diaphragm and a pop filter may not be necessary.

Figure 4.6 A correctly positioned pop filter, with about "a clenched fist" distance between the singer and the filter, and the filter and the microphone.

Commercial pop filters can be purchased relatively inexpensively. Alternatively, an almost as effective pop filter can be made for minimum expense using pantyhose and a wire coat hanger.

Windscreens

A *windscreen* is a foam cover that is put directly around a microphone's capsule, as shown in **Figure 4.7**. In addition to the boom of plosives, loud low frequency rumbles are produced when a wind current flows over or into a mic's diaphragm. Depending upon their severity, a windscreen can reduce or eliminate these artifacts. Windscreens should be used when a mic is unavoidably positioned in the air stream from an HVAC duct, or when using mics outdoors where wind is unavoidable.

Windscreens do affect a mic's high frequency response, so should only be used when necessary in the studio. When recording singers in a recording studio, pop filters are preferred over windscreens, because they are more acoustically transparent.

Wind Muffs/Dead Cats

Hairy *wind muffs*, or *dead cats*, are used to protect mics from outdoor wind noise in the motion picture, television, and broadcasting industries. In the music industry, about the only time wind muffs are used is on exposed condenser mics at outdoor events in windy conditions – not in studio recording.

4.17 Shock Mounts

Vibrations traveling through the floor (a singer's foot movement, mechanical HVAC vibrations, trucks driving by, low frequency sound waves, etc.) can travel up microphone stands and into the mic, where they manifest themselves as low frequency rumbles and booms.

Figure 4.7 A mic with and without a windscreen.

Figure 4.8 Mics in their shock mounts.

Mic stands can get knocked or bumped, causing loud booms. Shock mounts, as shown in **Figure 4.8**, are elastic suspension mounts that mechanically isolate mics from stands. Most vibrations traveling through the mic stand are absorbed by the elastic material, and do not reach the microphone. Shock mounts are essential for serious studio and location recording.

ABOUT MIC STANDS

- Always use a good, sturdy mic stand.
- Always loosen clutches and joints before adjusting the stand – they won't wear out so quickly, meaning the stand will be less likely to "adjust itself down" in the middle of a recording.
- Always set the legs up properly. Center columns of tripod stands should *not* be touching the floor – the rubber feet on the ends of the legs isolate the stand (and microphone) from vibrations in the floor.
- Always point a leg in the same or very similar direction to the boom arm (particularly if it is extended a lot, or has a heavy mic on it). The stand will be less likely to tip over.
- Do not over-extend the boom segments. The further a mic is boomed out, the less stable it is, and the more strain it puts on the clutch – wearing it out quicker, and making it more likely to droop in the middle of magical recording.
- Mic stands should not touch other mic stands, and mic stands should not touch any part of any instrument. Any physical contact can cause vibrations, buzzes, and booms to get into the mics and your recording.
- Make sure all joints and clutches are secure and tight so the mic won't move.
- Ensure the stand is not going to be knocked or bumped by anyone – even with the best shock mount, this will cause some noise.

4.18 Mic Preamps

If selecting microphones and mic techniques to best match the characteristics of the sound sources and the desired production aesthetic wasn't enough, different preamps also add their unique characters to the sound. This is why recording is an art form!

Cheap preamps found on budget mixers and interfaces tend to sound unrefined and noisy compared to the higher quality circuits in professional consoles and outboard preamps. There are three common categories of desirable preamp sound:

- In-your-face, aggressive, or assertive: These characteristics can really bring a sound forward, and help it stand out in the mix.
- Warmer, fuller, and phatter: Common characteristics of tube-based circuits, these preamps make things sound bigger and smoother.
- Transparent: A preamp with this character is designed to impart little of its own flavor, and simply report the characteristics of the microphone and what it's picking up.

These are generalizations, and you should read product literature and unbiased reviews to help you select preamps to meet your needs. In a small home or project studio, a few high quality outboard preamps as front-end to a DAW recording system are a great investment.

PREAMP INPUT IMPEDANCE

Some preamps have variable input impedances. Different input impedances can subtly change the frequency response and timbre the mic produces. Additionally, you can expect slight changes in the transient punch or fullness, and in the detail of off-axis room sound. These differences are more apparent when using dynamic and ribbon mics, although condenser mics with a transformer output can exhibit tonal changes (some more than others).

Lower impedances might be generalized as producing fatter, fuller, more "vintage" sounds, while higher impedances increase the focus, punch, and tightness of the sound, and can be a little brighter. On ribbon mics, higher impedances can give the sound increased low frequency weight, rather than simply making it brighter.

Drive and Saturation

Some preamps offer adjustable *drive* or *saturation* features. As the levels of these distortions are increased, the more the preamp's tube or solid state electronics are overloaded, so they produce their characteristic distortions. Tube saturation is a warmer, fuller sound. Solid state drive can be anything from tube-like warming to nastier grungier transistor fuzziness. "More" of these colorations is not always "best" though, and it's important to keep the original sound, and what will be most beneficial to the mix in mind. *You can't undo overdone saturation artifacts*.

The traditional location for preamps (whether they are in a console, interface, or are external outboard units) is in the control room. The drawback of this is that very long runs of mic cable are usually needed. How long mic cables can be before the audio quality degrades depends on the quality of the cable, and the circuits each end of the cable plugs into. Analog snakes 60 m (200 ft) long are common in live concert sound systems, but undesirable for critical recording. High frequencies can be attenuated by poor quality cable – and you may not always know what cable, or how much of it is running between rooms in a studio. For critical recording, the best advice is to keep the mic cable runs as short as possible.

If long cable runs are unavoidable there are two other options:

- Remote control preamps put the amplifying electronics in the recording room, near the microphones – so the mic cables can be very short. Balanced analog line level signals (which are less susceptible to interference and degradation than mic signals) then run the distance to the recording equipment. The preamp's gain, phantom power, polarity, and any other adjustable parameters it may have, are operated remotely from the control room, using a control box or computer software. The preamp and remote are usually connected together using some kind of inexpensive data cable.
- Digital snake systems locate the preamp electronics *and* analog to digital converters in the recording room near the mics. The audio is transmitted to the control room as digital data, over copper wire data cables, optical fibers, or Ethernet cables. The cable runs can be well over 100 m (300 ft). The long data cable run (for the audio and control data) is *much* cheaper and lighter than a traditional copper wire snake.

A drawback of both of these alternatives is that all the preamps in the system are the same, and may not have the unique characteristics of other top quality outboard preamps.

YOU'RE ONLY AS GOOD AS YOUR WEAKEST LINK!

There's no point in having top of the line monitor speakers if you have poor quality microphones or preamps! *Your recording chain is only as good as its weakest link*. Purchasing high quality *front-end* first (the input devices) allows you to upgrade your *back-end* (the mixing and playback system) at a later date, but still have great sounding material originally captured. Having poor quality mics or preamps means you are capturing potentially poor sound, which can't be turned into great sound later no matter how good your monitors. In fact, great monitor speakers will just magnify the deficiencies of poor mics or any other devices in the recording chain.

AUDIO EXAMPLES
Can be found on the companion website

Mic Preamps

Example 4.12: A vocal recorded through a budget mixing console preamp.
Example 4.13: The same vocal and mic, recorded through a professional mixing console preamp, known for its warmth and fullness.
Example 4.14: The same vocal and mic recorded through a boutique outboard preamp, known for its punch and aggressive qualities.

4.19 What Mic to Use?

Generic characteristics of microphones and preamps can be anticipated – but that's as far as any science goes! Recording is an *art form* because of the myriad variables in the recording and mixing chain that allow the engineer to craft the sound of the final product.

Good, unbiased product reviews are a great place to start when deciding what mics to use or purchase. Stay away from manufacturer and retailer publications, and catalogs with "reviews" that are really sales pitches. The following summarized information should help you make initial mic choices.

Use Dynamic Mics When:

- The sound source is very loud.
- The sound source contains no essential real high frequency content.
- Amplitude and transient compression characteristics are desired.
- A "smudging over" of details is acceptable or desirable.

Common recording studio applications include: Kick drum, snare drum, tom toms, guitar, and bass amplifiers.

Stage and live sound applications include: Other drum percussion instruments, horns, and vocals.

Use Small Diaphragm Condenser Mics When:

- Accuracy, details, and maximum resolution are desired.
- The sound source contains essential high frequency content.
- The sound source has important initial transients or sharp, percussive characteristics.
- The bigness and "hype" of a large diaphragm condenser mic is not desired.
- The sound source may be quiet.

Common recording studio applications include: Drum overheads, hi-hat, other percussion (particularly transient or bright sizzly instruments), acoustic guitar and other plucked string instruments, horns and other wind instruments (when accuracy is preferred over hype and color), bowed string instruments, and pianos.

Stage and live sound applications are similar, but more care must be taken when using relatively fragile condenser mics on stage. Additionally, the mic's extra sensitivity and reach can increase feedback potential.

Use Large Diaphragm Condenser Mics When:

- A big, full, hyped sound is desired.
- A detailed recording with good high frequencies is desired.
- The slightly warmer and/or smoother sound of many large diaphragm mics will be beneficial to the sound source and mix.
- The sound source may be quiet.

Large diaphragm mics tend to be studio workhorses – they are the generic go-to mic for many applications.

Recording studio applications include almost anything where a big, bold, detailed representation is desired: Vocals, horns and wind instruments, string instruments, acoustic guitars and other plucked string instruments, drum overheads, room mics, and pianos. They are also popular as guitar and bass speaker cabinet mics.

On stage and for live sound they need to be used with care. Their cost, size, and weight make them less popular in some markets, but with knowledge of how to keep feedback at bay they can be used on drums and speaker cabinets, and throughout quieter stages – jazz or folk shows perhaps, for vocals, strings and acoustic guitars, horns and wind instruments, speaker cabinets, and pianos.

Use Ribbon Mics When:

- A smoother, rounder, warmer representation of high frequencies is desired, or when a bright, shrill, or sizzly sound needs "mellowing."
- Their transient character is desired.
- Their vintage character is desired (if they are vintage designs).
- Extreme rejection of sounds coming from the sides (and above or below) is desired, and there is no spill coming from directly behind the bidirectional mic.
- Significant proximity effect is beneficial. Or when you want a warm, full sound from a foot or two away.

- You want the ability to quickly rotate the mic 180°, and exploit the potentially different backside character of the mic.
- High SPLs, moving air currents, and physical shocks *will not* be a threat to the mic!

Typical studio applications include: Vocals (with a good pop filter), drum overheads and room mics, acoustic guitars, electric guitar cabinets, horns, strings and other orchestral instruments, and piano.

They are less commonly used on stage due to their fragility, but professional touring acts do sometimes use them.

Use Tube Mics When:

- A big, phat, warm sound is desired. Maybe the source sound is a little thin or harsh, and needs warming up?
- Warmth and thickness are more important than sizzle and fine detail in the very high frequencies.

Common studio uses include: Vocals, horns, thin sounding guitar and bass cabinets, and other acoustic instruments that need some "sonic makeup."

Directional Mics are Beneficial When:

- Spill from sound sources behind the mic needs to be minimized – use a cardioid mic.
- Spill from sound sources to the sides and towards the rear of the mic (but not at 180°) needs to be reduced – use a hyper-cardioid mic.
- The least pick-up to the sides, above and below the mic is desired (but there is no spill coming from behind the mic) – use a bidirectional mic.
- Room reflections need to be de-emphasized, as is typical in bad sounding, acoustically untreated (or incorrectly treated) small rooms – use cardioid or hyper-cardioid mics.
- Proximity effect is desirable to "beef up" a sound.

Omnidirectional Mics are Beneficial When:

- Their more open, transparent sound is desired.
- Their better, more linear low frequency response is desired.
- Close miking is necessary, and their lack of proximity effect mud is desirable.

- A good sounding room produces desirable reflections that will give a sound excitement and energy. In many typical home and small room recording situations, omnidirectional mics are generally not used, because of undesirable room acoustics.
- A less "point source," and more organic and blended image is desired – and can be achieved through effective use of the better sounding spill omni mics pick up.

4.20 There's More to It Than Specifications!

A lot of this chapter has discussed specifications, and you *should* use mic spec sheets to help identify mics with characteristics beneficial and complementary to the sound sources and projects you're recording. But while specifications *are* an indicator of technical performance, they *are not* an indicator of perceived sonic performance or how a sound will work within the context of a mix. The *actual sound* of a microphone is much less easily quantifiable and is a combination of its technical specifications and many other characteristics that are impossible to quantify numerically.

A mic with stunning technical specs may just not work well on a given sound source – it might exaggerate characteristics in an unflattering way. A noisier, less sensitive mic may make a specific singer sound stunning, even though its technical specs are not as good! What is more important? Having a sound that works aesthetically, or having a more technically clean and less noisy signal to work with? Most engineers would sacrifice technical specs for the *right* sound! A slight technical sacrifice isn't going to stop a record becoming a hit!

PRACTICAL EXERCISES

Record identical sound sources (a voice or an instrument) using the following techniques. Record a minute or two of each to different tracks on your DAW. Adjust the preamp gain to achieve identical levels for each example.

1. A directional dynamic mic from a distance of at least 1 foot.
2. The same directional dynamic mic from a distance of just a couple of inches.
3. A small diaphragm directional condenser mic from a distance of at least 1 foot.
4. The same small diaphragm directional condenser mic from a distance of just a couple of inches (with a pop filter for vocals).
5. A large diaphragm directional condenser mic from a distance of at least 1 foot.
6. The same large diaphragm directional condenser mic from a distance of just a couple of inches (with a pop filter for vocals).
7. An omnidirectional mic from a distance of at least 1 foot.
8. The same omnidirectional mic from a distance of just a couple of inches (with a pop filter for vocals).

9. Using mics of different polar patterns move the sound source around the mic as you record.
10. Then try similar things with tube mics and ribbon mics if you have them.

Listen to each example. How do they sound different?

- How do the near and far positions sound different?
- How do the near and far positions of the omnidirectional mic compare to the directional mic?
- How does the timbre, level of detail, and intimacy of the sound compare between the dynamic, small diaphragm condenser, and large diaphragm condenser (and ribbon and tube mics if you have them)?
- How does the off-axis sound differ between a cardioid and omnidirectional mic? What happens to the sound as you move more off-axis on a cardioid mic?

5 EQ Basics

In This Chapter:

5.1 What Is EQ?

Equalization, or *EQ*, is available on every mixing console and in every DAW. EQ will probably be necessary at some point during most recording and mixing projects. This chapter is included in this book because EQ is such an intrinsic part of the recording process, used to complement the mic techniques discussed later.

EQ changes the frequency balance of a sound. It allows different frequency ranges to be boosted or attenuated (reduced) by variable amounts. Analog EQ works by manipulating the relative phase and timing of frequency bands, causing boosts or attenuations of those frequency ranges. Digital EQ can model this analog behavior, or some digital EQ is more phase linear, and less destructive to the integrity of the sound – but also less characterful!

Different EQs *sound different* – that's the point! But there are certainly some cheap EQs available that are considered to sound bad because of the undesirable artifacts they color the sound with, or to sound generic because they sound similar to many other EQs. Conversely, there are also (usually expensive) EQs available that impose very desirable side-effect characteristics on the sound!

5.2 Last Resort, and Creative Mix Tool

Because the relative phase of a sound's frequency content is changed as part of most EQ processing, it's best to consider EQ a relatively destructive "last resort," and not a go-to tool – to be avoided if there are other solutions. When recording or tracking, you, the engineer, should aim to capture the best, most desirable sound possible. "It sounds OK…. It's not quite right, but we'll EQ later," is *not* a professional way to approach recording.

Corrective EQ i*s not* a substitute for good sound sources, good mics, good mic technique, or a great sounding room to record in. The characteristics and frequency content of a recorded acoustic sound source can be changed significantly by using different microphones or mic techniques, and even other sound sources – and these are preferred acoustic alternatives to EQ. If something doesn't sound right when it is being tracked, the problem should be identified and fixed by changing one or more of those variables. Aside from minor tweaks, it is rarely possible to "fix it in the mix."

EQ *is* a creative tool that should be used to give a mix clarity – assuming great sounding tracks were recorded in the first place. The stereo soundstage is the space-compressed environment between the loudspeakers, into which many sound sources are placed. Most loud or important sounds are panned on top of each other in the center, where they compete for physical and frequency space. While building up a recording layer by layer, it is impossible to anticipate what specific EQ or timbre each element will need in order to sound good *together* – this will only be evident when they are all at appropriate mix balances and pan positions. Minor EQ tweaks are inevitable.

CASE STUDY

Fix It Before You Record

A common problem that should be fixed before pushing the record button is a sound source that sounds too boomy and muffled, and lacks real pitch definition and clarity – a miked bass guitar cabinet for example. Possible solutions include:

- Moving the mic further from the sound source to reduce proximity effect.
- Using a less thick sounding mic.
- Dialing in an alternate sound on the bass amp.

EQ It as You Mix

If the bass sounded great in isolation while tracking, but is obscuring the kick drum in the mix, then EQ *should* be used while mixing, to make those sounds work together.

Sometimes you Just Have to…

Of course, not all recording situations are ideal. Sometimes, you have to deal with a bad snare drum for example. No matter what mic and technique is used, it just doesn't sound as good as it should – but the recording has to happen. In this situation, if EQ helps, then yes, *use it* for corrective reasons. It's not an ideal situation, but "the show must go on!" But don't rush to this conclusion before exploring all other potential solutions.

5.3 Can You EQ Spill?

When evaluating a sound, channel, or track for EQ treatment, make sure you are listening to the source sound itself, and are not distracted by any spill in the background. A tom tom mic will pick up lots of distant tinny sounding snare drum spill. This midrange spill should not be EQ'd out. The tom tom has desired frequency content in that same range, and the tom tom will become dull, and lack definition and punch if the snare spill is EQ'd down. Just about the only kind of spill that can be reduced a little by using EQ is extreme low or high frequency spill in a frequency range that the actual desired sound source does not occupy – for example cymbal splash leaking into a bass mic, or bass guitar boom leaking into an acoustic guitar mic.

In large room recording situations, polar patterns, mic placements, and sound barriers can be explored to minimize spill. If spill is unavoidable yet undesirable and unworkable, tools such as gating, or silence stripping in a DAW, may help.

ABOUT SPILL

Spill is not always the enemy, and for some musical and recording styles it should be embraced as an essential part of the sound and mix. A mix devoid of any spill sounds ultra-clean, and too sterile for some projects. A mix which embraces good sounding spill can sound organic, wholesome, and unified.

5.4 EQ Filters

An EQ circuit is known as a *filter*. Different filter types affect different frequency ranges in different ways. Several types of filters are commonly found on mixing consoles, outboard gear, and in plug-ins.

Shelving Filters

A *low shelf filter*, as shown in **Figure 5.1**, allows frequencies *below* a specified *cut-off frequency* to be boosted or attenuated by up to ±15 to ±18 dB or more. The maximum amount

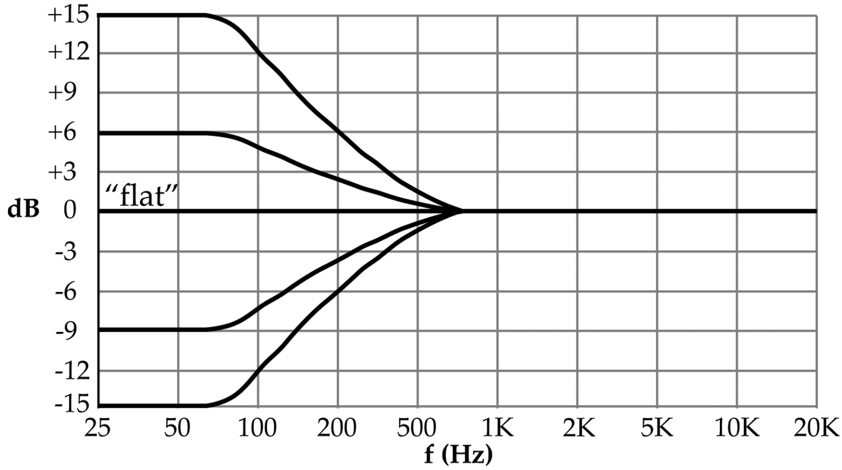

Figure 5.1 The straight line at "0 dB" on the y-axis represents "flat," or "no change" – the filter has no effect on the sound, and nothing is boosted or attenuated. The largest "uphill" plot as you read to the left, represents the maximum boost of +15 dB below the cut-off frequency of 100 Hz in this example. The largest "downhill" plot as you read to the left, represents the maximum attenuation of −15 dB below the cut-off frequency. The other plots represent possible intermediate settings of +6 dB and −9 dB.

of boost or attenuation depends upon the specific equipment. The amount of boost or cut can be anywhere between maximum cut, through 0 dB (no change) to maximum boost.

CUT-OFF POINTS

The *cut-off frequency*, or *cut-off point* of a filter is the frequency at which it is 3 dB less than maximally effective. The filter in **Figure 5.1** is described as operating below the "−3 dB" cut-off point, which is 100 Hz.

Slopes

It is impossible to build a filter that operates like a brick wall – doing nothing to frequencies above 100 Hz, but immediately being fully effective at 99 Hz. That type of filter would sound terrible and be musically unusable anyway! All filters have a *slope* – a range of frequencies over which they become increasingly effective. Even though the filter in **Figure 5.1** is described as having a cut-off point of 100 Hz, its slope means that it *does* affect a broad range of frequencies well above 100 Hz.

A *high shelf filter*, as shown in **Figure 5.2**, allows frequencies *above* a specified cut-off point to be boosted or attenuated by anywhere between minimum and maximum values.

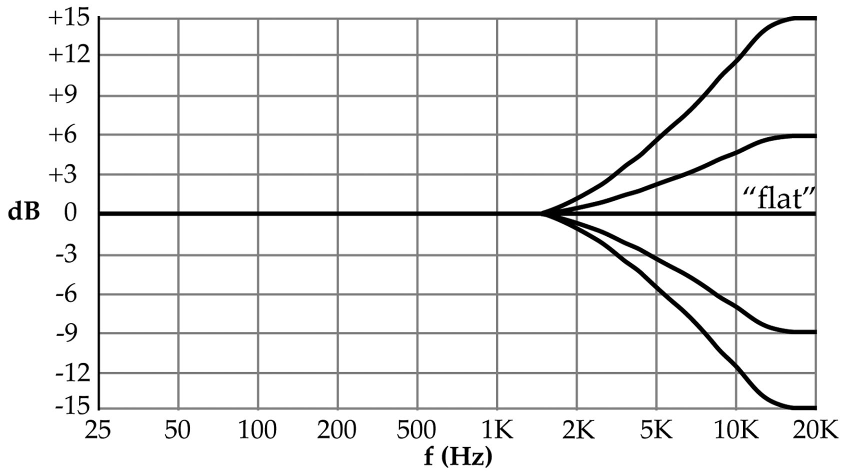

Figure 5.2 A high shelf filter. The maximum boost of +15 dB, above the cut-off point of 10 kHz is represented by the largest "uphill" plot as you read to the right. Note that the –3 dB point, the cut-off point, is at 10 kHz and the filter's slope affects frequencies well below this. The largest "down-hill" plot as you read to the right, represents maximum attenuation.

Shelving filters are often switchable between a couple of different cut-off frequencies, or they are *sweepable* – meaning that the cut-off frequency can be smoothly moved up and down over a range of frequencies. This allows a low shelf filter to be set to boost or cut below a cut-off point of 60 Hz in order to focus on the weight and thump of a kick drum or bass guitar, or swept up to 250 Hz in order to attenuate low and low-mid frequency proximity effect build-up.

If a shelving filter has a switchable or sweepable cut-off frequency, it will usually have two controls:

- A gain knob to dial in the amount of the boost or attenuation.
- The switch to select, or knob to sweep the cut-off frequency.

Low shelf filters are the default low frequency (LF) EQ circuit, and high shelf filters the default high frequency (HF) EQ circuit on most traditional equipment and EQ plug-ins.

Peak/Notch Filters

As pictured in **Figure 5.3**, *peak/notch filters* allow a range of frequencies *around* a center frequency to be boosted or attenuated by any amount up to ±15 dB or ±18 dB or more, depending upon the equipment or plug-in. Mid frequency (MF) equalization filters are always peak/notch filters.

Peak/notch filters may also be sweepable, meaning the center frequency can be moved up and down the spectrum – which allows them to really focus in on an issue. Additionally,

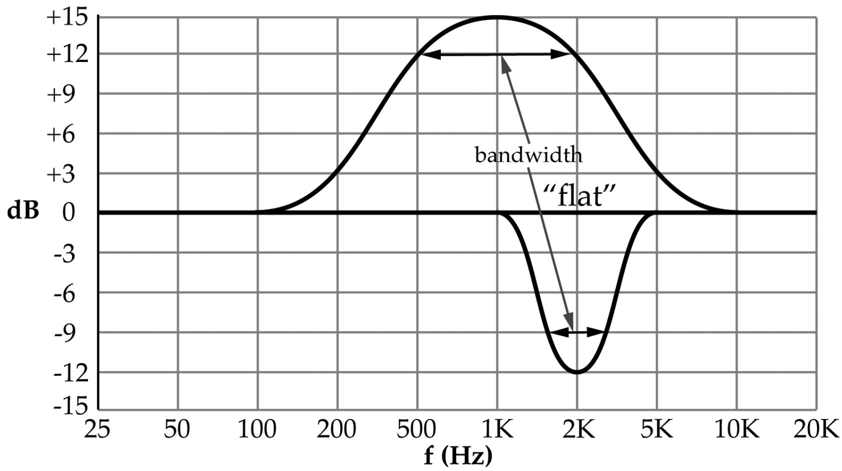

Figure 5.3 A sweepable and bandwidth adjustable (parametric) peak/notch filter, showing two of an infinite number of possible settings – a wide +15 dB boost around 1 kHz, and a narrow –12 dB attenuation around 2 kHz.

they may have a *bandwidth* control, allowing them to affect a wide range of frequencies for broader more "musical" tonal changes, or a narrow range of frequencies for more surgical problem solving. The bandwidth controls on many hardware EQ sections often use icons to represent narrow and wide filter widths. Some plug-ins use *Q* or *Quality Factor* indications of bandwidth.

BANDWIDTH, OR "*Q*"

Q, or Quality Factor, is an indication of a peak/notch filter's musical bandwidth. It is calculated using the following equation:

Q = center frequency/bandwidth

EQ filters are *constant Q* filters, meaning that they affect a similar musical bandwidth (or number of musical octaves) as the filter is swept up and down. A bandwidth of 100 Hz in the low frequencies represents a wide bandwidth – over an octave if it is centered around 100 Hz, giving it a Q factor of 1. Because frequency is a nonlinear exponential scale, higher octaves are made up of much larger numerical ranges. So to retain a Q factor of 1 when centered around 1 kHz, the same filter must have a bandwidth of 1000 Hz.

- Higher Q factor numbers mean that the filter is narrower, and affects a smaller frequency range.
- Lower Q factor numbers mean that the filter is wider, affecting a broader frequency range.

If the bandwidth of the filter described here stayed at a constant 100 Hz, it would have a Q factor of 10 if it was centered around 1000 Hz, and affect a narrow range of frequencies. A filter that behaves like this would be a *constant bandwidth* filter – but they are musically useless, so not found on audio equipment.

Parametric and Semi-Parametric EQ

- *Parametric* or *Fully Parametric* denotes either single filters that have continuously adjustable gain, frequency sweep, and Q controls; or entire EQ sections/plug-ins which have those controls on every band.
- *Semi-parametric* denotes either single filters that are sweepable but do not have a Q control (they have a fixed Q); or entire EQ sections/plug-ins which contain some fully parametric filters (usually the mid peak/notch bands) and some which lack Q controls.
- *Pseudo-parametric*, or *quasi-parametric* can mean the same thing as semi-parametric; or can indicate that only a small selection of preset Q values are offered – low and high, or wide and narrow, for example.

Low Cut Filters

Also known as *high pass filters* or *HPFs*. *Low cut filters* are sometimes found as part of an EQ section or plug-in, or as part of the preamp section of a hardware device. In the recording studio they are more of a technical problem solving filter and less of an artistic or musical filter than shelving and peak/notch filters. As shown in **Figure 5.4**, low cut filters *roll off*, or progressively attenuate frequencies below the cut-off frequency, by a fixed amount.

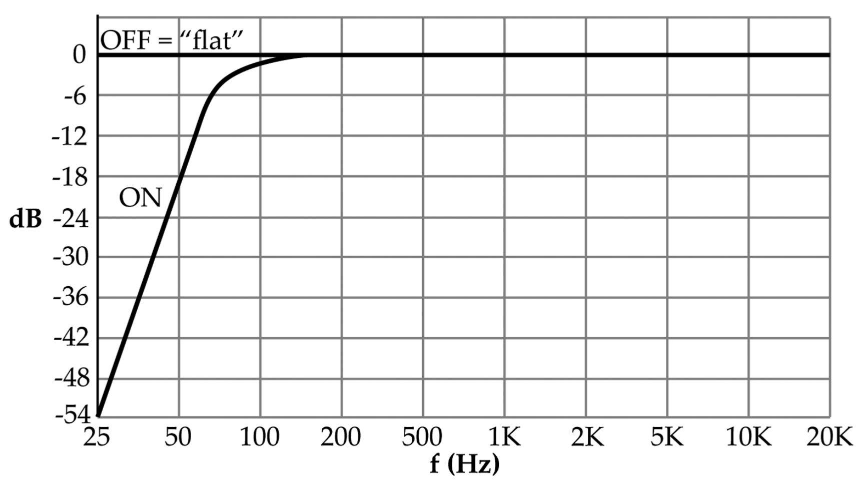

Figure 5.4 This low cut filter severely attenuates below a cut-off point of about 80 Hz when engaged. No boost is possible. The filter is either "in" or "out" – either fully effective, or "off."

Many HPFs only have an *in/out* button to turn them on and off. Some are sweepable, with an additional frequency control. Some, particularly plug-ins, offer adjustable slopes.

MORE ON SLOPES...

- A 6 dB per octave HPF with a cut-off point of 100 Hz will attenuate 100 Hz by –3 dB, then 50 Hz by a further 6 dB (–9 dB), and 25 Hz by another 6 dB (–15 dB).
- An 18 dB/octave HPF also set at 100 Hz will attenuate 100 Hz by –3 dB, then 50 Hz by a further 18 dB (–21 dB), and 25 Hz by an additional 18 dB (–39 dB), making the very low frequency content much less audible.

HPFs can be used to remove the rumbles and booms caused by mechanical or HVAC sources, too much uncontrolled "weight" or "woof," or some proximity effect build-up – although proximity effect usually extends up into higher low-mid frequencies which some HPFs do not affect.

5.5 Analog vs Digital EQ

A couple of frequent observations made by engineers who transitioned from analog EQ to digital EQ are that:

- Many digital EQs require more significant boosts or cuts than they're used to making, in order to hear a difference.
- "Character" is missing from many digital EQs.

Certain analog EQs are popular because they have desirable and pronounced sonic characteristics – making them different to each other, and different to more "transparent" EQs which just "do the job" without imposing much of their own signature on the sound. Regardless of which design philosophy is used, a good analog EQ really has an obvious effect on the sound when a small boost or cut is dialed in. In order to perceive similarly obvious changes on some digital EQs and plug-ins, significantly greater amplitude changes often have to be dialed in.

This means that if you're used to working with plug-in EQ or a digital mixer, but are incorporating analog EQ into your hybrid DAW studio, or mixing on an analog console, you should expect to make smaller EQ adjustments than you're used to making.

There are many plug-in recreations of classic vintage analog EQs, and reviews of these products are generally favorable – however it is an accepted fact that while the plug-in versions come close, they *do not* sound the same as the original devices. Competing plug-in companies make emulations of some of the same classic analog EQs – they don't all sound the same, and some are definitely better than others.

Digital EQs can offer a level of precision, control, and linear behavior that analog devices cannot. In some applications and to some engineers this can be a benefit. In other situations or to other engineers it's a lack of necessary character.

DAW USERS AND EQ

Even though most digital EQs and plug-ins don't sound the same as the best analog devices, it is definitely worth upgrading from the stock EQ plug-ins found in many DAWs, and investing in some with a few specific characteristics beneficial to the musical styles and projects you're working on.

In books and magazines, some interviews with engineers and producers cite the amplitude of EQ changes they have made. Those figures may refer to changes they made on analog EQ, or plug-ins modelling boutique analog EQ hardware. If you're mixing with stock DAW EQ, or less esoteric plug-ins, don't be surprised if you end up using a couple of dB's more than you might read about.

Ever noticed that there seem to be more knobs, buttons, and greater control ranges on plug-ins compared to analog devices? Most hardware devices offer less flexibility because more control or increased control ranges cost more money. So a hardware unit focuses on what you actually need 99 percent of the time! It costs nothing to add additional controls or control ranges to plug-ins, but this can make the best sounding settings harder to find, and you may not need the "extras" most of the time.

5.6 Additive vs Subtractive EQ

It is preferable to use EQ subtractively wherever possible. This means turning down frequency ranges that contain excess, or are causing problems – *not* turning up frequency ranges adjacent to the actual problem.

PROBLEM: A vocal recorded with the singer close to the mic is going to exhibit proximity effect build-up in the lows and low-mids.

POOR SOLUTION: **Figure 5.5 Top** – turning up the high-mid EQ to increase "cut," "presence," "intelligibility," and "bite." In addition to turning up potentially desired characteristics, this additive EQ also turns up whatever noise and spill there is in that frequency range, and eats into available headroom possibly causing distortion or clipping.

BETTER SOLUTION: The problem is *not* that there aren't enough high-mids, but that there is an excess of low-mid frequencies causing congestion that obscures the high-mid details. **Figure 5.5 Bottom** – scooping out the lows and low-mids targets the problem directly. Using this approach, noise, spill, gremlins, amplitude, and undesirable EQ artifacts have not been unnecessarily boosted – so the sound, and mix, will remain cleaner and clearer.

BEST SOLUTION (FOR NEXT TIME): Position the mic more appropriately, and/or use a different mic to capture a less muddy and boomy sound in the first place!

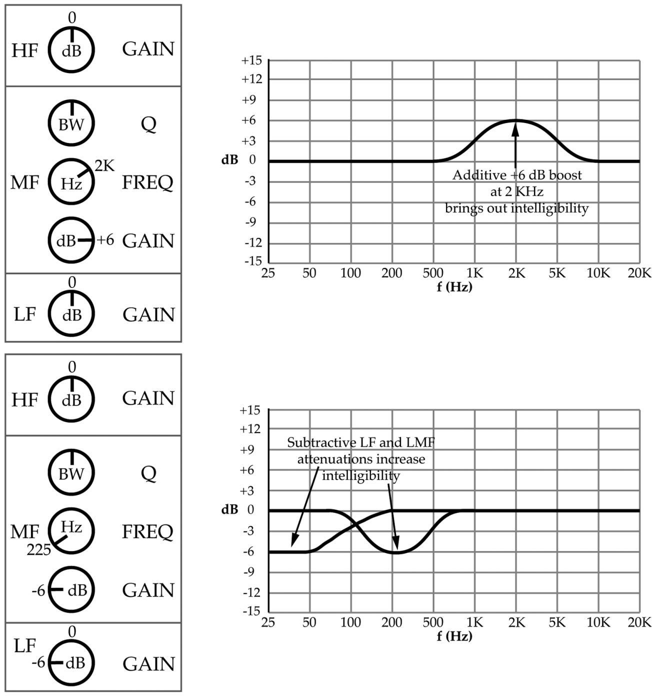

Figure 5.5 **Top:** Additive EQ to boost the "cut," "presence," "intelligibility," and "bite" of a vocal which has too much proximity effect clouding up the lower frequencies is not the best solution to a muddy vocal problem. **Bottom:** Subtractive EQ targets the cause, and attenuates the proximity effect boost of the lows and low-mids without boosting noise or spill, or leaving that excess frequency content present to cloud up the mix in combination with other sounds.

BOOST–SWEEP–CUT

An effective way to identify a potentially problematic frequency range is to:

- **BOOST:** Create a large boost in an EQ filter – at least +10 dB.
- **SWEEP:** Sweep the frequency control around until the nastiness *really* jumps out.
- Return the gain control to "0" for a time to re-familiarize yourself with the original sound.
- **CUT:** Reduce the gain of the problem frequency to correct the problem.

Make sure you are targeting *actual problems identified before any EQ adjustments are made*! Don't fall into the novice trap of boosting and sweeping every filter available, then attenuating the ugliest range each reveals, one after another – those frequency ranges are probably not problems when they are not boosted, and are at a natural balance. Anything turned up loud enough will be unbalanced and ugly!

5.7 The Fewer Filters the Better

Because EQ can negatively impact the technical integrity of the sound, use as few EQ filters as possible. It is usually easier for novice engineers to hear the effects of an EQ boost than it is to hear an attenuation, and a common mistake is to turn up many frequency bands instead of attenuating just one or two. Some of the reasons for this common error are:

- Our ears have a very poor memory – it is difficult to *really* remember how something sounded 30 seconds ago. After making some EQ adjustments it's easy to forget the specific problem that needed fixing.
- Our ears are easily fooled into thinking that something sounds "better" simply because it sounds different.
- Making something louder fools our ears into thinking it sounds "better." Boosting some frequency content certainly makes the sound louder, but "louder" does not always mean "better"!
- Turning a certain frequency range up might make the sound more instantly appealing, and then turning up another frequency range might add another instantly appealing character. But this approach has little to do with locating and fixing an actual problem!

The effect of an EQ curve created by boosting multiple filters, similar to **Figure 5.6 Top**, can be better produced by attenuating the only non-boosted frequency range instead, as shown in **Figure 5.6 Bottom**. The attenuation approach is preferable because:

- It targets the specific problem frequency band.
- It uses fewer filters, adding less potential technical degradation to the sound.
- It is subtractive, meaning less undesired signal is amplified.
- It does not cause the headroom problems additive EQ can.

ADDITIVE EQ

While subtractive EQ is generally preferred, there are certainly times when additive EQ is necessary – after all, there's a reason the gain knobs boost as well as attenuate! Sometimes a sound does need an extra boost to make it pop. If a vocal track needs a little extra "air" in the top end, it's best to do it with only one filter – a couple of dB boost above 10 kHz on the high shelf filter. Using one filter additively is simpler and cleaner than turning down the LF, LMF, and HMF to achieve a similar result subtractively.

Some plug-ins have many more filters than the usual three or four found on a hardware mixer EQ section – so it's very easy to end up with multiple adjacent frequency bands turned similarly up or down. If two or three adjacent bands overlap, and are boosted or cut

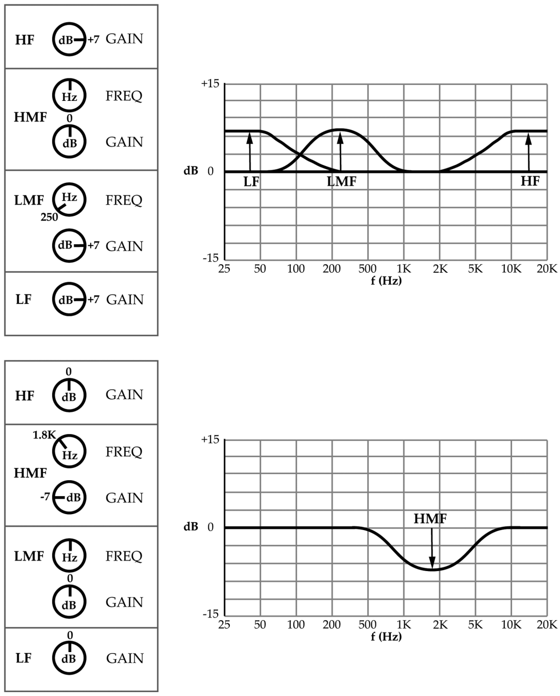

Figure 5.6 The "see-saw" effect of EQ means that similar results are better achieved by attenuating the only non-boosted frequency band in the top diagram, as shown in the bottom diagram.

similarly, it would be simpler and more logical to use only the middle of those filters and to widen its bandwidth.

5.8 How Much to EQ?

Do whatever the sound and mix need. But – if you're working with great sounding source tracks, it's a good idea to question situations where changes greater than –6 or +3 dB are made when using analog EQ, or about –8 or +5 dB when using digital EQ. Sometimes large changes are necessary, but use these suggestions to remind yourself to double-check that such a large change is really necessary.

5.9 When to EQ?

Minor tweaks to fit a sound into the mix are best done as part of the mixing process. It is bad to EQ the sound as it is recorded, and then have to undo and create slightly different EQ during mixing.

But if after trying all the recording and mic technique methods discussed in following chapters and preferred to using EQ, some minor EQ is still necessary in order to capture a great sound, it can be better to EQ into the recording system, rather than during mixing. This is particularly true when the problem is an overabundance of a particular frequency range, commonly low frequencies, that can eat up headroom, resulting in lower average signal levels and gain structure issues when lots of corrective EQ is applied during the mix.

If you record through a mixing console or other hardware EQ, EQ-ing to tape or the DAW is easy. If you do not, then recording an EQ'd sound into a DAW track is not the same as applying EQ during mixing, and requires some extra steps such as, depending on the DAW:

- Routing the mic into an aux track with an EQ in it, and then bussing that sound into an audio track.
- EQ-ing the input channel (if the DAW has these).

Neither of these solutions fixes headroom issues though, because the A to D conversion is before the EQs.

5.10 Golden Rules of EQ

CUTS VS BOOSTS

- Attenuate (cut) *problems*, as gently or severely as the mix needs.
- Gently boost for *character*.

- EQ is a last resort, and not a go-to process in lieu of other solutions. *Don't be lazy while tracking!*
- EQ subtractively whenever possible.
- Use as few filters as possible.
- EQ additively if it means using fewer filters that way.
- Avoid EQ to remove spill – it will probably negatively affect the desired sound source.
- Double-check that any larger EQ changes are really appropriate and necessary.
- Don't accidentally end up with additive and subtractive filters set to the same frequency range.
- Never use EQ to make something louder. A common novice error is turning up the low and high shelf filters to produce a "smile curve" if a sound seems too quiet. If EQ seems to make something "louder," it was probably too quiet in the first place – use the fader to correct balance issues before using EQ to fit the sound in the mix.

PRACTICAL EXERCISE

Record some vocals with a directional mic at different distances, ranging from very close to a meter (3 ft) away. Record each to a different DAW track. Remove proximity effect from the closer recordings using subtractive EQ. Listen to the EQ'd closer tracks and the un-EQ'd more distant tracks, and compare the following:

- The timbral and tonal frequency content differences.
- The reflected room sound/ambience differences.
- The differences in proximity, and sense of distance perceived.

6 Stereo Imaging

In This Chapter:

6.1 The Stereo Soundstage

A stereo loudspeaker system allows sounds to be positioned in the *soundstage* pictured in **Figure 6.1**. A good, clear, interesting music mix usually uses much of this space between, and around the loudspeakers. This chapter will familiarize you with the stereo soundstage and different imaging concepts so that you can more effectively exploit the space in your own projects – and use mics and mic techniques to record the sound the mix needs, rather than trying to "force" inappropriately recorded sounds into a mix using electronic processing.

Remember that this soundstage is specific to stereo *loudspeaker* systems – and *not* headphones, which present a smaller sense of space, and lack front-back depth (as discussed earlier).

6.2 How to Listen

In order to properly perceive the sound of a stereo loudspeaker system, it is essential to sit in the *sweet-spot* – located by creating an equilateral triangle between the loudspeakers and

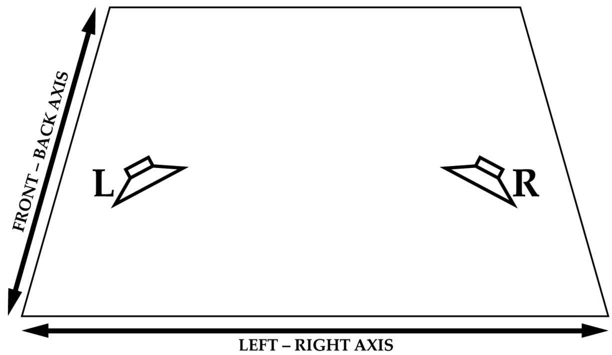

Figure 6.1 A 2-D representation of the effective stereo loudspeaker soundstage. Sounds can be positioned within this space.

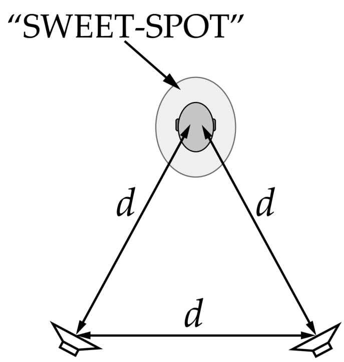

Figure 6.2 The correct listening position, sitting in the "sweet-spot." All *d* distances should be identical.

the listening position, as shown in **Figure 6.2.** The speakers should be positioned so that the tweeters are at ear level – if that is not possible, the speakers should be angled up or down so that the tweeters point towards the engineer's ears.

Where you sit is important!

- If you sit too far forwards, the stereo image will appear too wide and tend to have a hole in the center of some stereo sounds. You will compensate for this by not panning aggressively enough and the result will be a mix that is too narrow.
- If you sit too far away from the loudspeakers the stereo image will appear too narrow to you. You will compensate and create a mix that is too wide with too much extreme panning.
- If you sit off-center, the *law of first wavefront* dictates that the image will steer towards the loudspeaker closest to you and appear asymmetrically balanced and too heavy on that side. Your mix will be too heavy in the opposite loudspeaker when listened to correctly.

It is also important to have symmetrical room geometry on each side of the monitor speakers – the angle of the walls and the materials used. Asymmetry in control rooms can contribute to imaging problems and poor mix decision making. The monitor speakers and sweet-spot should be carefully positioned in the room. Professional studios might have an acoustical study done to ascertain the best position in the room, but if that's not in the budget there are many resources available online from professional monitor manufacturers and reputable magazine websites.

MONITORING LEVELS

Because of our ear's non-linear perception at lower SPLs, it is important to monitor and mix at an average (and safe) overall level – not too quiet, and not too loud. Between 80 to 85 dB C-weighted, slow response is recommended. You should check how your mix sounds at lower and higher levels though, so that you know it will sound good at a variety of playback levels. If you mix too quietly there can be a tendency to turn up the extreme low and high frequencies too much. If you mix too loudly, then hearing loss becomes a concern!

6.3 Phantom and Discrete Images

The pan control on a mixer or in a DAW is used to position sound between the left and right loudspeakers. Modern panners are continuously variable from the left, through the central space between the speakers, to the right.

The perceived location of any single mono channel, panned anywhere except hard left or hard right, is an illusion known as a *phantom image* – there is no physical loudspeaker where the sound appears to be located. Phantom center images have a unique character because of the comb filtering and frequency cancellation (above 5 kHz) produced when each speaker's direct sound combines with the inter-aural crosstalk from the "opposite" loudspeaker at each ear. This is not something you have to consciously think about though while recording and mixing – but it does make creating a mix based on phantom center images challenging!

Discrete, single loudspeaker images are produced when a sound is hard panned exclusively to the left or right loudspeaker. The ear closest to the loudspeaker reproducing the sound receives direct sound, and a fraction of a millisecond later the other ear receives inter-aural crosstalk only – neither ear receives both direct *and* time-delayed inter-aural crosstalk as is the case with a non-hard panned phantom image. Discrete single loudspeaker images have very different sonic characteristics to phantom images. They are:

- Clearer.
- More focused.
- More precise.

- Tighter.
- Timbrally more assertive.
- Arguably “smaller.”

Phantom images are:

- Slightly “larger.”
- Less focused.
- Arguably a little weaker sounding.

6.4 Image Width

Mono Point-Sources

A single (mono) microphone or channel produces a narrow, focused, *point-source* image, which although panned somewhere in the soundstage, has no width or spread. Mono point-source images can be panned by any amount to the left or right:

- Centered.
- A little off center.
- Almost to one side or the other.
- All the way to one side or the other.

Panning many mono point sources throughout the left/right soundstage produces an interesting and wide stereo image, with many discrete details and distinct points of sound spaced throughout it.

Spread Images

Spread images are produced by stereo mic arrays, synthesizers and keyboards with stereo outputs, and stereo reverb and effects units. A spread image takes up left/right space – sound comes from throughout its *width*. A spread image can be:

- Narrow – only a little wider than a point-source mono image.
- Wide – taking up nearly, or all of the space between the loudspeakers.
- Symmetrically panned – the pan controls are equal but opposite from the center position.
- Asymmetrically panned. For example, the L channel of a stereo synth sound panned hard left, and the R channel panned center.

Figure 6.3 shows some possible stereo imaging concepts and the pan positions that achieve them.

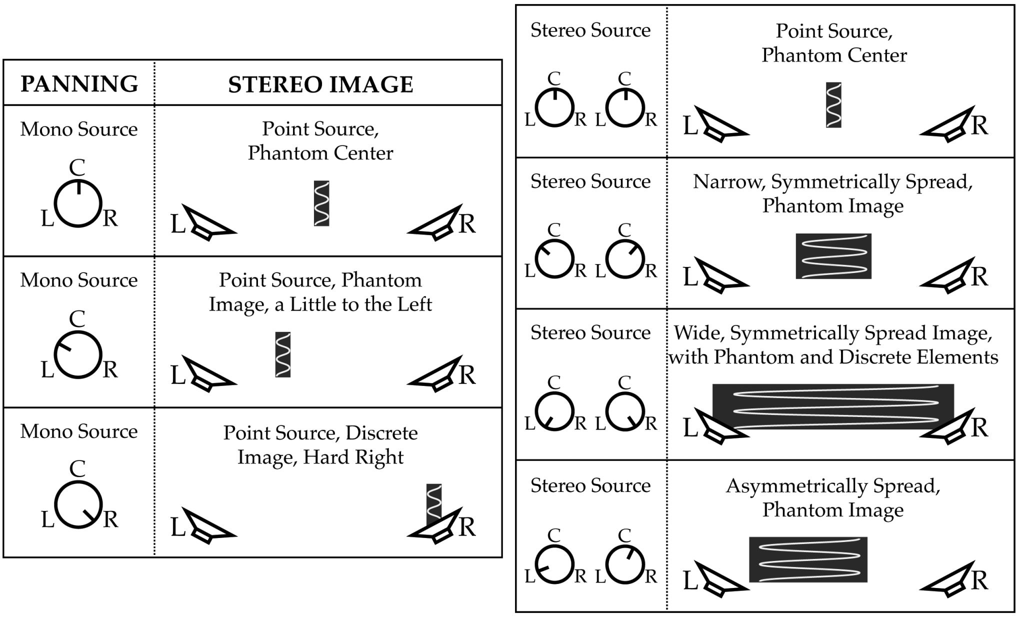

Figure 6.3 Pan positions and the images they create.

AUDIO EXAMPLES
Can be found on the companion website

Stereo Image Width and Panning Concepts

Examples 6.1 through 6.7 correspond to ***Figure 6.3***.

Example 6.1: A mono point source, panned centrally.
Example 6.2: A mono point source, panned slightly left.
Example 6.3: A mono point source, panned hard right.
Example 6.4: A stereo source, both channels panned centrally, produces a mono phantom center. *NOT STEREO!*
Example 6.5: A narrow spread image, panned symmetrically.
Example 6.6: A wide spread image, panned symmetrically.
Example 6.7: An asymmetrically panned spread image.
Example 6.8: An image that is too wide. Sounds can be clearly identified and located in the left and right extremes of the image, but a blurry "hole" is created in the middle where it's difficult to precisely identify and localize sounds.

Listen to these examples on both loudspeakers and headphones to hear how the sense of space sounds different on both reproduction formats.

6.5 Beyond the Loudspeakers

As shown in **Figure 6.1**, it is also possible to create the illusion that sound is coming from just beyond the physical spacing of the loudspeakers. This can be achieved using a variety of methods of varying predictability:

- Context: A widely panned stereo source may appear to be super-wide depending upon what else is going on at the same time, and what other imaging concepts it is being contrasted with.
- Out of phase content: A wide stereo sound source that has a significant amount of decorrelated information in each channel can appear to extend beyond the physical position of the loudspeakers. Some synth pads, and big room/hall reverbs are examples of this.

6.6 Depth Concepts

In addition to the left/right axis, the stereo format allows us to create the illusion of a front/back perspective. Sounds can appear to be positioned from slightly in front of the loudspeaker plane, to well behind the loudspeakers – as shown in **Figure 6.1**. Different techniques and concepts can create the illusion of depth in a mix:

- High frequency attenuation: As sound sources get naturally further away from the listener, the air the sound is traveling through absorbs some high frequency content. Rather than emulating this with EQ, mic it so it sounds that way! Conversely, the brighter and crispier a sound is, the more “up front” and “in-your-face” it usually appears.
- Reverb: In a natural acoustic environment (a room, concert hall, etc.) we hear a sound source close to us as mainly direct, dry sound, with only a small amount of wet, reverberant content. A far away sound source is experienced as predominantly wet, reverberant content. The addition of reverb (with a short pre-delay) can create the illusion that a sound is pushed behind the loudspeakers. The effect can be magnified by adding a little high frequency attenuation to the source sound. Or it can be miked so it sounds that way to begin with.
- Amplitude balance: A contextually quieter sound can also be interpreted as being further away. But, it is possible to have a very quiet sound appear forwards in the mix if it is bright and dry.
- Sound can appear to be forward of the loudspeakers because it is bright and relatively dry. Longer reverb pre-delays can be used to “detach” the reverb from the dry sound, keeping the dry sound up front. The “up-front-and-in-your-face” factor can also be increased with good volume automation and compression, to make the sound more consistent.

Context is also a powerful contributor to front/back imaging:

- If a mix contains sounds which have good frequency balances, and contrasting reverbs and wet/dry balances, the contextual illusion of front/back depth will be much more effective. Be careful though – while differences in reverb times and types can create interesting effects and stylizations in a mix, too much variety can destroy the concept that a band was playing together, and can turn the mix into a blur of disconnected sounds.
- If similar and abundant reverb is added to most of the elements in a mix, reverb becomes a less effective way of pushing sounds back.
- If the entire mix is timbrally dark and mellow, high frequency attenuation will be less effective at pushing sounds back.

In a multi-track production, these characteristics can be adjusted electronically during mixing. The focus of this book though, is adjusting them while recording – through the use of creative mic techniques. The result is very different, and much more organic and natural, and the sounds "mix themselves" more.

AUDIO EXAMPLES
Can be found on the companion website

Depth Concepts

Example 6.9: A bright, dry sound appears forwards in the soundstage.
Example 6.10: Gentle high frequency attenuation, and the addition of reverb, pushes the same source sound back behind the loudspeakers.
Example 6.11: The same instrument, recorded with a more distant mic technique naturally moves the sound to the rear of the soundstage.

6.7 The Illusion of Height

It is also possible to create a limited illusion of height in a stereo mix – but it is the hardest dimension to control and predict. There is no knob for it! As a huge generalization it might be said that brighter, swishy, more sizzly high frequency sounds can "float" above the bulk of the mix – but it's all a matter of context.

6.8 Static and Dynamic Panning

A stereo mix with most elements panned center has two drawbacks:

- A possible lack of clarity due to too many sound sources with overlapping frequency content competing for the same physical space.

- It is boring and uninteresting to listen to! The stereo soundstage exists to be used – it makes the sonic experience more expansive and keeps listeners returning for another listen to see what else they can notice.

Static panning occurs when the pan controls remain unchanged throughout the mix. It is good to pan some sound sources away from the center to improve clarity and interest. In a hypothetical mix, the lead vocal might be panned center, the electric guitar panned to the left, and the acoustic guitar to the right, and they stay in those locations for the entire song.

Dynamic, or *animated panning* involves moving and changing a sound's position, actively, as we hear it:

- *Fly-bys:* Sounds heard moving smoothly through the stereo image. Used too much, or inappropriately, moving pans *can* be distracting, but a swishy sound or a percussion effect heard moving from the left to the right, the measure before a chorus can create a spectacular lead in to the next section of the song!
- *Auto-panning*: An auto-panner is a device or plug-in that automatically moves a sound through the space between a left and right location, at a speed and movement style set by the engineer. It is easy to over-use, but continuous auto-panning has been used very effectively on solos and sound effects on countless rock and pop records.
- *Ping-pong panning:* A musical "question" phrase panned centrally, followed by a musical "answer" that alternates between being panned left and right. Or two parts that fairly rapidly bounce left then right. There is no smooth movement from one pan position to the other, just a bounce, or *ping-pong* between them. If a recording has musical material conducive to this effect, it can really expand the width of a mix and be interesting ear candy to draw listeners into the mix.

Prior to mix automation, engineers had to physically move the pan controls at the correct time, live, during mixdown – usually requiring a few engineers to "perform" big mixes! With automation and DAWs, panning movement has become easy – draw or record the automation, edit it if necessary, and the system will recreate that movement perfectly each and every time.

AUDIO EXAMPLES
Can be found on the companion website

Panning Concepts

Example 6.12: A mix with static panning only.
Example 6.13: A mix in which some elements move around dynamically.
Example 6.14: A mix featuring elements of antiphonal, side-to-side, question/answer panning.

6.9 Image Symmetry

Stereo image symmetry and physical balance are desirable in a mix. If there are too many elements or too much amplitude on one side of the stereo soundstage, the image becomes lopsided and unbalanced. The following principles can help achieve a good, balanced stereo image:

- Divide sound sources of similar role and function equally between the two sides of the stereo soundstage:
  - If a guitar part is panned to the left, it is usually best to balance it with something similar on the right (another guitar part, or perhaps a piano part) – otherwise the image will be left heavy.
  - Percussion sounds are great to pan around – they very effectively create width and stereo interest. A shaker, panned left, balanced with a triangle or clave, panned right, will instantly open up and widen a mix.
  - Stereo image balance does not have to be created by an identical instrument, but the musical functions of the two should be similar.
- The amplitudes of the L and R channels of the stereo image should be evenly balanced. You do not want the mix to end up with hotter levels on one side, due to asymmetrical panning or too many instruments of similar function panned to the same side.
- Both simultaneous symmetry (sounds happening at the same time) and antiphonal symmetry (subsequent sound events alternating sides) can be very effective, as long as the side-to-side alternation is not so slow that each event is perceived as asymmetrical static panning.

A common issue with music that builds up in layers as the verses and choruses progress, as a lot of pop music does, is that the reduced instrumentation of the first verse lacks enough layers of activity to create a balanced image. Subtle double tracking of rhythm parts, and layering slightly different and opposingly panned keyboard sounds are some strategies to avoid this. More obvious double tracking and sound layering can be saved for later sections of the song. Short asymmetrical gestures can be effective at times – but make sure they are artistically justified, convincing, and deliberate.

AUDIO EXAMPLES
Can be found on the companion website

Image Symmetry

Example 6.15: A symmetrically balanced mix.
Example 6.16: An unbalanced, asymmetrical mix.

6.10 Use All of the Soundstage!

Big, interesting, and expansive mixes use all of the available soundstage – there is sound located:

- Front and center.
- Behind the speakers.
- Extreme left and right.
- Between the phantom center and left and right speakers.

A mix that contains only center panned channels, plus additional mono channels panned hard left and hard right, will appear to have holes in the image between the center and extreme left and right.

Great mixes also exploit a variety of imaging concepts simultaneously. For example:

- Narrow point-source images can be surrounded by wide symmetrically panned spread images.
- Narrow point-source images can be framed by different narrowly panned spread images on each side.

There is no single "correct" mix. Wide, expansive, pointillistic mixes can be created by panning many narrow point sources throughout the soundstage. Less precise, but equally if not more expansive and immersive mixes can be created by panning multiple stereo sources throughout different parts of the soundstage. The stereo soundstages of two contrasting mixes are plotted in **Figure 6.4**.

6.11 Reality vs Recording

A recording is not a reflection of reality. Surround sound and immersive audio environments can come closer to recreating reality, but stereo just can't. A recording is always a fold-down or reduction of something happening in a much larger space. Collapsing many microphones, channels, tracks, or sound sources into a two-channel reproduction system creates the challenge of making the mix clear and intelligible. Fully exploiting the stereo soundstage will help the clarity and intelligibility of each component sound source by positioning them in their own unique physical space. The more that panning and the soundstage are exploited and used in a mix, the easier it becomes to achieve clarity in the mix.

Expansive use of the soundstage may not be faithful to an artist's live performance set-up, but that authenticity is less important than making a product that works well as a recorded piece of art.

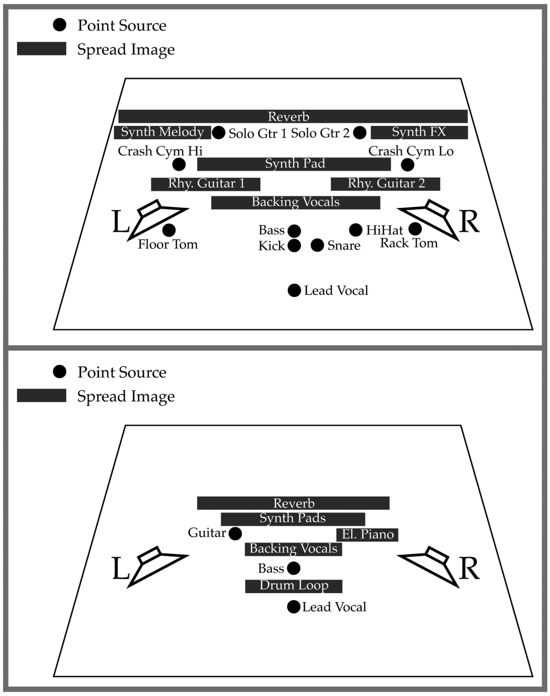

Figure 6.4 **Top:** A bigger, more expansive mix that utilizes the entire soundstage and multiple imaging concepts. **Bottom:** A mix that uses the soundstage more conservatively.

PRACTICAL EXERCISE

1. Sitting in the sweet-spot between a good pair of stereo monitors, analyze some commercial mixes and make soundstage diagrams similar to those above.
2. Identify different imaging techniques used in those mixes, and note:
   - ▶ What makes the mix clear and gives it good clarity?
   - ▶ What makes the mix interesting, wide, and immersive to listen to?
3. Listen to the same mixes on good headphones and see if you get the same sense of space and imaging.
4. Use these observations and techniques in your own mixes!

7 Stereo Microphone Arrays

In This Chapter:

7.1 Microphone Arrays

A microphone *array* is a collection of microphones that function together as a unit to capture the horizontal size of a sound source or environment. *Minimalist* microphone techniques use the same number of mics as loudspeaker channels in the recording or dissemination format – so two mics are the minimum required for a stereo array. More mics can be used, but there are several standard techniques that should be understood and mastered before embellishing them and adding extra mics.

Matched microphones are strongly recommended for stereo techniques. Two mics of the same make, model, and vintage can each sound slightly different due to unavoidable component and manufacturing inconsistencies. *Matched pairs* undergo additional testing by the manufacturer in order to pair together mics with the most similar performance, resulting in more accurate stereo imaging.

Researchers, scientists, and professional recording engineers have tried lots of variations of these stereo array techniques, before deciding which particular geometry produced the results they were after. For optimal results, use a ruler and a protractor to set these arrays up precisely. "Close enough" or "approximate" versions often produce inferior results. Experiment only after you have become familiar with the sound of the correctly set up techniques.

HOW OUR EARS PERCEIVE DIRECTION

Our ears use three methods to determine the perceived location (the *localization* or *directionality*) of a sound source:

- Amplitude differences between the ears: A sound is perceived slightly louder in the ear that it is closest to. This is a more significant indicator of the directionality of frequency content above about 800 Hz.
- Time arrival differences between the ears: The ear that the sound source is closest to hears the sound slightly before the opposite ear, which is further away. Our brains calculate directionality from the sub-millisecond time difference between the wavefront's arrival at each ear. This is a more significant indicator of the directionality of frequency content below about 800 Hz.
- Our anatomy EQ's the sound as it travels around different parts of the face, head, and shoulders to get to each ear.

Put together, these create *head related transfer function* (*HRTF*) effects.

7.2 XY Coincident Pair Techniques

A *coincident pair*, or *XY array*, uses two cardioid or hyper-cardioid microphones positioned so their capsules are as close together as possible. The mics are angled between 90° to 130° apart, with their pick-up patterns crossing in front, as shown in **Figure 7.1**. The array as a whole is aimed at the center of the sound source, putting each mic 45° to 65° off-axis from

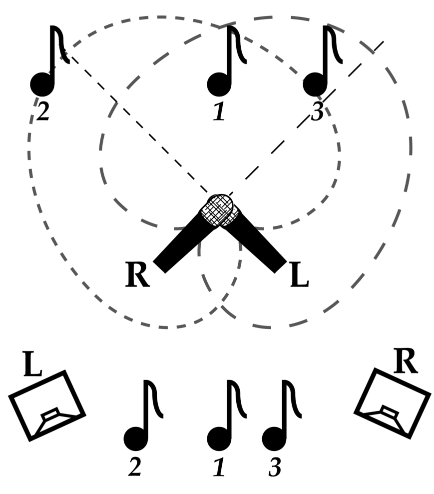

Figure 7.1 An XY coincident pair of microphones. Sound sources are shown at various positions in the array's pick-up and their perceived playback locations are shown between the loudspeakers.

pointing directly forwards. The mics must not touch, acoustically shade, or interfere with each other. If they touch, mechanical vibrations and buzzes can result. If a microphone is physically "in the way" (in line of sight to the sound source) of the other mic, frequency response and phase inaccuracies can result.

Sound source **1** comes from the center of the source, and is equally off-axis and equidistant to both microphones. This means that both mics pick up identical sound at the same time. When the mics are panned hard left and hard right on a stereo loudspeaker system, both loudspeakers reproduce the same sound at the same time. Identical direct sound and inter-aural crosstalk arrive at each of the listener's ears creating a phantom center image directly in-between the loudspeakers.

Sound source **2** is located far to the left of the array's pick-up. It is also equidistant to both mic capsules – so both mics pick up the sound at the same time, and on playback both loudspeakers reproduce the sound at the same time. However, compared to sound **1**, this source is more on-axis to the left mic, and beyond 90° off-axis to the right mic. The left mic picks up a greater amplitude of sound with better frequency balance than the right mic – so the sound is reproduced louder and more clearly by the left loudspeaker than the right loudspeaker. This localizes the sound to the left when we listen from the sweet-spot.

Sound source **3**, located "somewhat" to the right of the array's pick-up area, is again equidistant to both mic capsules – so whatever sound is picked up by each mic will be reproduced at the same time by both loudspeakers. This source is on-axis to the right mic, and less than 90° off-axis to the left mic. A slightly greater amplitude and slightly better frequency balance is picked up by the right mic than the left mic, so the sound is reproduced slightly louder and a little more clearly in the right loudspeaker. This localizes the sound somewhat to the right when we listen from the sweet-spot.

The reproduced locations of the sounds are not as widely spaced as the original sound source. This is because XY techniques capture and present directional information only as amplitude differences between the two mics and loudspeakers. Time arrival differences are powerful indicators of directionality, but this technique doesn't capture any, because the microphone capsules are as close together as possible.

General characteristics of XY coincident techniques include:

- Narrow, compact images due to the lack of time arrival cues, and a lot of overlap of each mic's pick-up.
- Potential for muddy sound due to the significant amount of each mic's off-axis pick-up – no mic is actually pointed towards the center of the sound source. Cheaper directional mics, with their inferior off-axis colorations exaggerate this problem. The increased price of better directional mics is justified by their improved off-axis frequency response in this context.
- Good mono compatibility! With no time arrival differences between the capsules there is no phase cancellation or comb filtering when they are summed together for mono playback.

- The perceived stereo image can be widened by increasing the angle of incidence between the mics – however, this puts each mic further off-axis from the center of the sound source. This makes the center of the image muddier – particularly if the mics have poor off-axis response.
- Using hyper-cardioid mics instead of cardioid mics decreases the overlap of the pick-up patterns and unclutters the center of the image, making it wider. However, sounds coming from the center of the source are now effectively more off-axis to each hyper-cardioid mic's more directional pick-up – so they are picked up with greater off-axis coloration and slightly less amplitude.

This list of mainly negatives might suggest that XY coincident pair techniques are of limited usefulness. Not so! Where mono compatibility is a major concern, or if an expansively wide image is not desired, the technique can work very well. With some sound sources, and in some rooms and concert venues, XY coincident pair techniques produce more pleasing results than other techniques. It is hard to predict exactly how different stereo arrays will sound on different sources or in different rooms, so do try XY coincident pair techniques!

AUDIO EXAMPLES
Can be found on the companion website

Mic Arrays and Stereo Images

The same drum set is recorded throughout this chapter.

Example 7.1: A single mono overhead microphone, centered on the drum set.
Example 7.2: An XY coincident pair of overheads, centered over the drum set.

7.3 Blumlein Pair Technique

A *Blumlein pair* is a coincident array of bidirectional microphones crossed at 90°. Each mic is displaced 45° from the center of the sound source, as shown in **Figure 7.2**. This array works in a similar way to an XY coincident pair – it *is* a coincident pair. There are no time arrival differences between the coincident capsules, so as with XY coincident pair techniques, only amplitude difference information is recorded and reproduced on playback.

Bi-directional microphones *really* reject sound sources positioned to their sides. This results in a lot less overlap or common pick-up between the mics. Blumlein pairs can produce very natural, realistic images that are wider and more spacious than XY coincident pairs. Mono compatibility remains excellent because there are no time arrival differences between the two mics. Again though, no microphone is directly on-axis with the center of the sound source.

Sound from behind the microphones will be picked up due to the bidirectional mic's rear pick-up. This can be useful in a good sounding natural acoustic because room

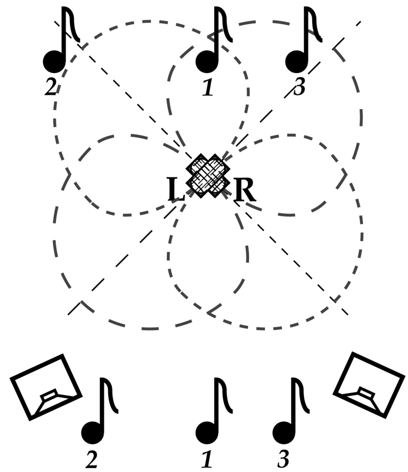

Figure 7.2 A Blumlein pair of microphones. Some sound source and perceived playback positions are shown.

reflections and reverb can really energize the sound. But it can be problematic if the room sound is not desirable, or there are sources of unwanted spill, or walls or ceilings too close behind the mic array.

AUDIO EXAMPLES
Can be found on the companion website

Mic Arrays and Stereo Images

Example 7.3: A Blumlein pair of overheads, centered over the drum set. The image is wider, and more of "the room" can be heard.

7.4 Near-Coincident Pair Techniques

A *near-coincident array* uses two directional microphones with the capsules positioned a small distance apart. The mics face away from each other, resulting in slightly less overlap of their pick-up patterns than an XY coincident array. The angle of incidence between the mics is usually between 90° to 110°, as shown in **Figure 7.3**.

A sound coming from the center of the sound source, labeled **1**, is equally off-axis and equidistant to both microphones. Therefore, both microphones pick up identical sound at the same time. With the mics hard panned, both loudspeakers reproduce identical sound simultaneously. Each of the listener's ears receives identical direct sound and identical

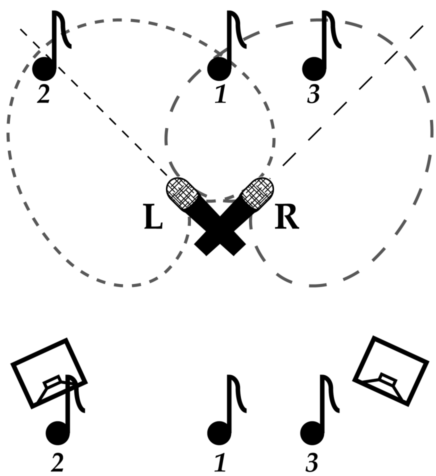

Figure 7.3 A near-coincident pair of microphones. Some sound source and perceived playback positions are shown.

inter-aural crosstalk, causing a phantom image to be perceived in the center, directly between the loudspeakers.

Sound source **2**, located far to the left of the array's pick-up area, is closer to the left mic, which picks up the wavefront a fraction of a millisecond before the right mic. So the left loudspeaker reproduces it first. Additionally, this sound is on-axis to the left mic and beyond 90° off-axis to the right mic – so the left loudspeaker reproduces it louder and with no off-axis coloration.

Sound source **3**, located somewhat to the right of the array's pick-up area, is picked up by the right mic slightly before the left mic – so the right loudspeaker reproduces it slightly before the left loudspeaker. This sound source is on-axis to the right mic, and quite off-axis to the left mic, so the sound is reproduced with slightly more amplitude and less off-axis coloration by the right loudspeaker.

The distance between the mic capsules allows important time arrival information to be recorded and reproduced. Characteristics of near-coincident techniques include:

- A wider, clearer stereo image than XY coincident techniques – often more closely matching the original sound source.
- A less muddy and confused sound.
- Sources in the center of the image might still be slightly muddy due to their off-axis position – no mic is pointed directly at the center of the sound source.
- A slight decrease in mono compatibility because of the small time arrival differences between each capsule.
- The perceived image can be made wider or narrower by increasing or decreasing the angle of incidence between the mics. But as with an XY coincident array, increasing this angle puts each mic further off-axis with the center of the source, further muddying the center image sound, particularly if the mics have poor off-axis response.

- The perceived image width can be increased by increasing the distance between the mic capsules (to a point), or decreased by moving them closer together. Mono compatibility decreases dramatically as the distance between the two capsules increases.
- Using hyper-cardioid mics instead of cardioid mics will widen the image, however any sound coming from the center of the source will be more off-axis to each mic and consequently picked up with increased off-axis coloration and slightly less amplitude.
- It may be possible to decrease the angle of incidence when using hyper-cardioid mics and maintain a similar image width to using cardioid mics – with the advantage of the hyper-cardioid mics being less off-axis to the center of the sound source.

This list of mainly positives make this technique seem like a preferred alternative to XY coincident technique. But when mono compatibility is a major concern, or an expansively wide image is not desired, XY coincident techniques may still be preferable. It's important to use the "mono" monitoring button before recording – to see if much sound disappears, or if any strange phasey comb-filtering is heard when the image is summed to mono. Slight adjustments to the spacing and angle of incidence can reduce these problems and increase mono compatibility.

Specific near-coincident pair techniques include the following (see **Figure 7.4**):

ORTF

Developed by the Office de Radiodiffusion-Télevision Français, the *ORTF* technique uses two cardioid mics set at an angle of 110°, with the capsules spaced 17 cm (7.7 in) apart. With capsule spacing similar to the ears on your head, this technique produces a satisfying, transparent stereo image, with relatively good mono compatibility.

NOS

Developed by Netherlands Radio, the *NOS* technique uses two cardioid mics set at 90° to each other, with the capsules 30 cm (12 in) apart. Due to this wider spacing, mono compatibility is not as assured as with an ORTF array, but due to the narrower angle of incidence between the mics, they are less off-axis to the center of the sound source. The increased phase differences introduced by the wider spacing result in a slightly bigger and more immersive image than a typical ORTF image.

DIN

Standardized by the German Deutches Institut Für Normung organization, the *DIN* technique positions two cardioid mics 20 cm (7.8 in) apart, angled at 90°. The mics are not as

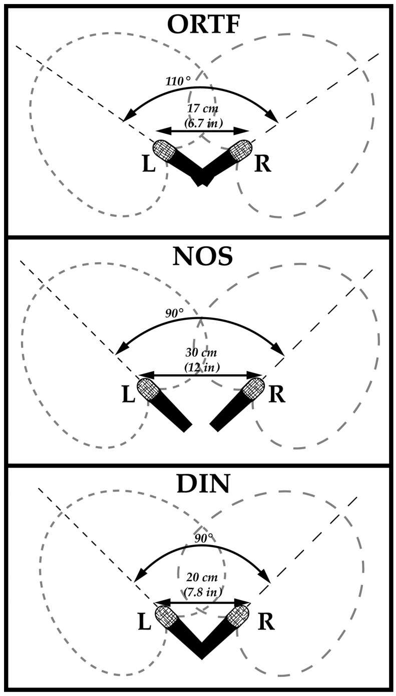

Figure 7.4 ORTF, NOS, and DIN arrays.

off-axis as in an ORTF array, but are spaced slightly further apart. The spacing is significantly less than in a NOS array. This results in a good balance of time arrival and amplitude differences that is particularly effective at shorter distances.

AUDIO EXAMPLES
Can be found on the companion website

Near Coincident Mic Arrays

Example 7.4: An ORTF near-coincident pair of overheads, centered over the drum set.
Example 7.5: A NOS near-coincident pair of overheads, centered over the drum set.
Example 7.6: A DIN near-coincident pair of overheads, centered over the drum set.

7.5 Spaced Pair (AB) Techniques

In *spaced pair* or *AB* techniques, two microphones face directly forwards, with their capsules 40 to 60 cm apart (15 to 24 in), as shown in **Figure 7.5**. *Omnidirectional* microphones are generally used for spaced pair techniques – but there are many applications, such as drum overheads, when spaced directional mics can be used.

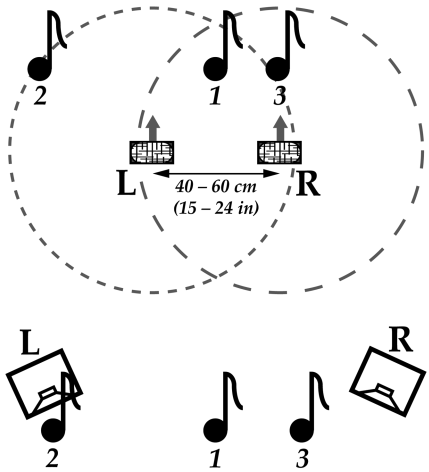

Figure 7.5 A spaced pair of microphones. Some sound source and perceived playback positions are shown.

Sound source **1**, located centrally, travels the same distance to each microphone – the sound arrives at each mic simultaneously, and at the same amplitude. Both loudspeakers then reproduce this sound at the same time and at the same amplitude, creating a phantom image centrally located between the loudspeakers.

Sound source **2**, located to the extreme left, travels a significantly shorter distance to the left mic than to the right mic – so the left mic picks it up a millisecond or two (or even several ms with more extreme mic spacing) before the right mic. This sound source is also slightly quieter when it arrives at the right mic because the right mic is further from it than the left mic – but in most cases this amplitude difference is trivial. On reproduction, the left loudspeaker reproduces the sound the same number of milliseconds before the right loudspeaker, and possibly just slightly louder than the right loudspeaker.

Sound source **3** is picked up by the right mic slightly before the left mic, and is fractionally louder in the right mic. Therefore the right loudspeaker reproduces the sound slightly before, and slightly louder than the left loudspeaker.

Time arrival information is the primary indicator of directionality in spaced mic techniques. Mic spacing of 50 cm (20 in) creates a maximum time arrival difference of approximately 1.5 ms between the mics – which produces a natural listening experience for a listener in the sweet-spot of a stereo pair of loudspeakers.

A suitably spaced AB mic array can capture an image that very closely matches the natural acoustic image, or pleasingly expands it. General characteristics of AB spaced pair techniques include:

- A wide, expansive stereo image created predominantly by time arrival information. The image is not necessarily the most focused or precise, but it is the most *enveloping*.
- If the mic spacing is too wide, sources in the center can lack definition and focus.

- As the spacing between the mics is widened, longer than natural time arrival differences are generated – the image becomes weaker, made up of "ghostly" separated left/right components with a "hole" in the center where nothing can really be localized.
- The relatively large time arrival differences between the mics make this the least mono compatible stereo technique. It is vitally important to check for phasing and comb filtering before recording, and to adjust the array spacing to minimize any problems – although it will probably be impossible to remove all mono compatibility issues.
- The perceived image width can be increased or narrowed by increasing or decreasing the spacing between the mics – but don't go too wide on a close sound source (the image will have a hole in the middle), or too narrow when the mics are further away (the image will become too mono).
- Using directional mics instead of omnidirectional mics decreases common overlap, widening the stereo image – but it also positions centrally located sound sources more off-axis to each mic, subjecting them to off-axis coloration. Directional mics will, however, minimize the pick-up of undesirable room sound and sources of spill to the rear of the mic array.

If a wide sound source such as a symphony orchestra, choir, or very large drum set is being recorded, it is possible to augment a spaced pair with a third, center mic, and position the left and right mics wider than if they were just a pair. This extra center panned information can either aid mono compatibility by adding common sound to both loudspeakers, or further reduce it by adding a third time displaced signal to the mix – it is entirely situational, depending on the spacing and relative balance between the mics.

AUDIO EXAMPLES
Can be found on the companion website

Spaced Mic Arrays

Example 7.7: A spaced pair of omnidirectional overheads, 50 cm (20 in) apart, centered over the drum set.
Example 7.8: A widely spaced pair of cardioid overheads, each mic positioned over the edge of the drum set's width.

7.6 MS (Middle-Side) Techniques

In all of the stereo techniques previously discussed, left and right microphones were present – and each were routed discretely to their respective loudspeaker (for a full width stereo image). An *MS* array, as shown in **Figure 7.6**, is significantly different. It doesn't pick up left/right information. Instead, MS techniques use a forward facing microphone (most commonly a cardioid mic, but it could be any polar pattern) to pick up the center or *middle*

(*M*) information, and a bidirectional microphone turned sideways at 90° to pick up *side* (*S*) information (which is a combination of the sounds coming from the left and right sides of the array). The S mic is a single mic capsule which outputs the sum of the sound waves hitting its front and rear – there is *not* a separate left and right capsule or output.

Sound source **1**, in the center of the array's pick-up, is only picked up by the M mic – it is in the null point of the S microphone.

Sound source **2**, to the left, is picked up by a combination of both the M mic and the S mic. Its polarity in the S mic is normal because it is "in front" (+) of that mic.

Sound source **3**, to the right, is also picked up by both mics, but its polarity in the S mic is reversed because it is "to the rear" (–) of that mic. It is also louder in the M mic and quieter in the S mic than sound source **2** because of its direction relative to each mic.

Converting MS to LR Stereo

In order to be reproduced on stereo loudspeakers or headphones, an MS array needs to be decoded to stereo using a sum and difference matrix – the process is known as *matrixing*.

Raw MS tracks can be recorded and then matrixed to stereo at a later time. Or they can be matrixed prior to the recording device to record LR stereo tracks. If you record raw MS tracks you *should* monitor a matrixed LR stereo form of the signal in order to check the array is working properly and that the matrixed image is desirable. To record LR stereo tracks a hardware matrix is necessary after the mic preamps and before the recording device. Hardware matrixes usually have few controls, and possibly only one – an S level control which varies the level of the S+ and S– signals that are combined with the M signal – to adjust the image width. Matrixing is also possible on a mixing console or in a DAW.

The Sum and Difference Matrix

The left side of the MS mic array's pick-up in **Figure 7.6** is a combination of the M mic and normal polarity sound picked up by the front of the S mic (S+), so:

$$LEFT = M + S+$$

The right side of the mic array's pick-up is a combination of the M mic and polarity reversed sound from the rear of the S mic, (S–), so:

$$RIGHT = M + S-$$

Those equations simplify to:

$$LEFT = M + S$$
$$RIGHT = M - S$$

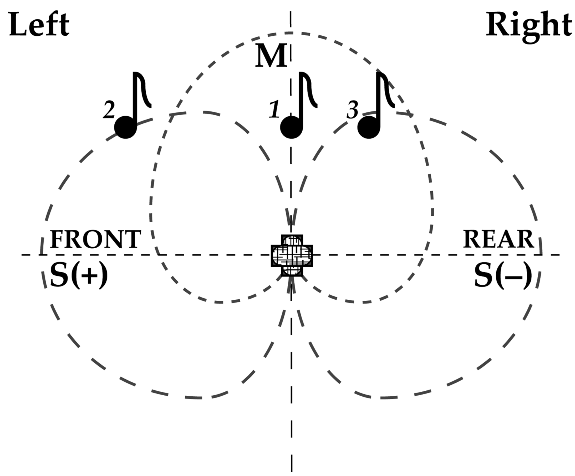

Figure 7.6 An MS array.

LEFT is the *sum* of the M and S mics, and RIGHT is the *difference* between the M and S mics.

MATRIXING IN A DAW

EASY METHOD: Plug-ins are available to matrix MS signals to LR. They are used on a stereo track which contains the MS signals.

MORE DIFFICULT METHOD: Matrixing can be done "manually" if you don't have a plug-in, or you want practical experience of how the matrix works:

- Label the "M" track, pan it center, and put the fader at unity.
- Duplicate the S track.
- Label the original track "S+" and pan it hard left.
- Label the duplicate track "S–," polarity reverse it, and pan it hard right.
- If you need to use a plug-in to polarity reverse, make sure plug-in delay compensation is turned on, or that identical polarity reverse plug-ins are put in the M, S+ and S– tracks, even though they are not actually doing anything in the M and S+ tracks. This ensures all three tracks are delayed by the same plug-in latency and remain perfectly in phase so the matrixing will work properly.
- Put both the S+ and S– faders at –∞, then group them together so that they are always at exactly the same level when you move them.
- With the S+ and S– faders at –∞ the image will be mono because all you are hearing is the center panned M mic. Increase the width of the image by increasing the level of the S+ and S– faders. Decrease the stereo width by decreasing the level of the S+ and S– faders.

- Once the desired stereo image has been achieved, group together the M and both S faders so that their relative balance (and the stereo image produced) stays identical if you adjust their overall level.
- If you adjust any preamp, plug-in, or gain controls related to these three channels you will have to re-balance the M and S channel levels to maintain the stereo image.
- Any insert effects (EQ or compression, etc.) should be applied to the matrixed LR stereo signal and *not* the M, S+ or S– tracks. Route the outputs of the M, S+, and S– tracks to a group or aux track (instead of the main stereo outs), and apply the processing on that new track.

MATRIXING ON AN ANALOG MIXING CONSOLE

Matrixing on a mixing console requires three channels: M, S+, and S–. "Y-cabling" the S signal into two channels is not recommended because analog faders, even supposedly identical ones, do not perform exactly the same throughout their travel, and it is not possible to simply "Y" the output of a mic that requires phantom power. Maintaining an identical level of the S+ and S– channels is essential, and simply visually matching fader positions is not accurate enough.

- Label the "M" track, pan it center, and put the fader at unity. Adjust the gain to achieve good input levels. (The level of the M fader can also be below unity if you are balancing the array's level with other sources in the context of a larger mix.)
- Mute the M channel.
- Label the first S channel "S+," and connect the S signal to its input. Adjust its gain to achieve good input levels. Pan it hard left.
- Connect the S+ channel's post-fader Direct Out (or similar individual post-fader output) to the line input of a third channel, labeled "S–." Polarity reverse this channel and also pan it hard left. (Yes, *left*, the same as the other channel. For now.)
- Put both the S+ and S– faders at unity. You will not move the S– fader from unity, so put some board tape over it so it does not get accidentally knocked.
- Adjust the S– gain control until the two S channels cancel out as much as possible. When maximum cancellation is achieved, both the S+ and S– channels are amplitude matched.
- Pan the S– channel hard right. You should hear a horrible strange image, with extreme weirdness in the center.
- Pull down the S+ fader to –∞.
- Unmute the M channel.

- ► Gradually creep up the S+ fader to increase the S component in the signal and widen the stereo image. Reduce its level if the image gets too wide. Do not adjust the S– channel – it has already been calibrated to the S+ channel and its feed is after the S+ fader, so the S+ fader effectively adjusts it as well.
- ► If you change the level of the M track at any time you will have to adjust the S+ track by exactly the same amount in order to maintain an identical stereo image.
- ► You could route the M and both S faders through a stereo group fader (and not directly to the L/R output fader) and have a simple way to control the level of the matrixed image in your mix.

The level of the S component in the matrix should not be as high as the M signal. The S faders (or preamp level) will commonly be at least 6 dB lower than the M fader to produce a good stereo image. If there is too much S component in the image, a phasey "hole" develops in the center – activity is heard on the left and right extremes, but not throughout the width of the image.

MS technique has a very different sound to the other stereo techniques discussed in this chapter. It produces a very precise image in which it is easy to pinpoint the location of specific sound sources. It is sometimes criticized for being too "surgical" and not as substantial or enveloping as near-coincident or spaced techniques – which trade MS's localization accuracy and focus for slightly less accurate "smeared" but bigger images.

MS technique is completely mono compatible. If the stereo image is summed to mono, the S+ and S– signals cancel out completely (because they are polarity reversed versions of each other) leaving just the M mic signal as the mono output.

MATRIXING LR STEREO TO MS

Any LR stereo signal can be matrixed into its M and S components. This might be done to adjust the width of a stereo image. An LR stereo source can be made narrower by panning the two channels less extremely – but this may cause phase cancellation problems if the signal is not 100 percent mono compatible. It is impossible to widen a stereo signal that is too narrow by using the pan pots – the knobs only turn so far!

LR to MS matrixing to adjust an image width should be a last resort – it is preferable to adjust the mic array to achieve a better image in the first place. After matrixing the LR stereo signal to MS, the relative balances of the M (center) and S (width) components can be adjusted to widen or narrow the image width. The adjusted MS signals are then matrixed back to LR stereo using the techniques described earlier.

Plug-ins are available which input and output LR stereo signals, but will manipulate the M and S components within the plug-in.

Don't expect MS plug-ins to fix every stereo image problem! If a recorded stereo image is far too wide, has a hole in the middle, or lots of out of phase content between the L and R channels, there is little common M signal to be derived. Conversely, if a recorded stereo image is far too narrow and almost mono, there is little S difference signal to be derived. With a shortage of either component the technique is less effective. This technique is best used for subtle changes, not big "rescue" jobs!

AUDIO EXAMPLES
Can be found on the companion website

MS Mic Arrays

Example 7.9: An MS array, matrixed to produce a good wide stereo image. (Compare this example with the stereo image of a spaced pair on an acoustic guitar in a later chapter, **Example 14.16**.)

Example 7.10: An MS array. This example starts with the S component turned all the way down in the matrix, leaving just the mono center. As the S component is turned up you can hear the image widen. By about 20 seconds there is too much S component and the image gets "phasey" with a "hole" in the middle.

7.7 The Decca Tree

The *Decca Tree* was originally developed by the Decca record label for orchestral recording, and can be heard on countless records and film soundtracks. It consists of three omnidirectional microphones, traditionally large diaphragm condensers, which are set up in a triangle as shown in **Figure 7.7**. The L mic is panned left, the R mic panned right, and the

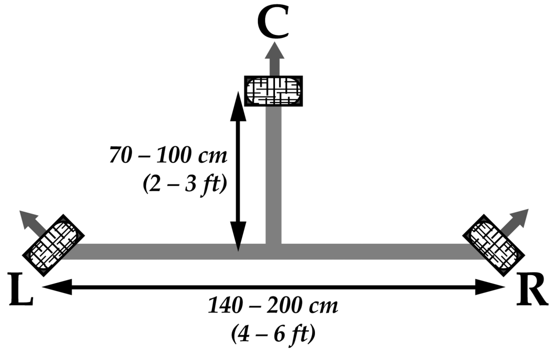

Figure 7.7 A Decca Tree.

C mic panned to the center. The gains of the three mics should be set equally, although adjusting the C mic up or down can narrow or widen the perceived image.

Many variations of the Decca Tree are possible:

- Cardioid or wide-cardioid mics can be used instead of omnidirectional mics.
- Combinations such as a cardioid mic for the C, and hyper-cardioid mics for the L and R can be used.
- The spacing of the C mic can be reduced, and the L and R mics spaced more or less widely in order to achieve the desired stereo image.

This array is famed for its expansive and immersive sound. It is not necessarily the most precise in terms of imaging accuracy – but its image characteristics are definitely pleasing. Neither is it the most mono compatible, with three spaced mics and their associated time arrival differences. Mono compatibility should definitely be checked prior to recording.

7.8 Binaural and Baffle Techniques

Stereo loudspeaker systems can create expansive two-dimensional stereo images, with some sense of depth, because of the inter-aural crosstalk and associated HRTF effects of the sound from each loudspeaker travelling to *both* of the listener's ears. The stereo microphone techniques discussed so far are designed primarily for loudspeaker system reproduction – but they also work well on headphones, although the perceived image is not the same as it would be on a loudspeaker system. Headphones, being binaural and lacking the inter-aural crosstalk and HRTF effects of a true stereo system, produce a more one-dimensional image. The stereo mic techniques discussed so far will not create a sense of depth in headphones.

Binaural microphone techniques capture characteristics similar to natural HRTF time delay and frequency filtering effects, at the time of recording. These pre-recorded HRTF effects (that would otherwise be missing when listening on headphones) are then reproduced by the headphones and interpreted as inter-aural crosstalk and HRTF effects by the listener's brain. This gives the headphone image more dimension, with an increased sense of space and depth. But it also means that binaural mic techniques do not translate well to stereo loudspeaker systems, where an additional dose of HRTF effects are unavoidably imposed by the playback system and the listener's head.

Several manufacturers produce binaural mics in the form of head-shaped, spherical or ovoid baffles which have mic capsules permanently mounted in positions similar to our ears, as shown in **Figure 7.8**.

Flat *baffles* can also be purchased or easily home built, and used between a pair of omnidirectional mics. The presence of the baffle acoustically separates the two mics (opening up the image), and simulates some rudimentary HRTF type effects. The *Jecklin*

Figure 7.8 A binaural "dummy head" microphone. Photo courtesy of Neumann, Berlin.

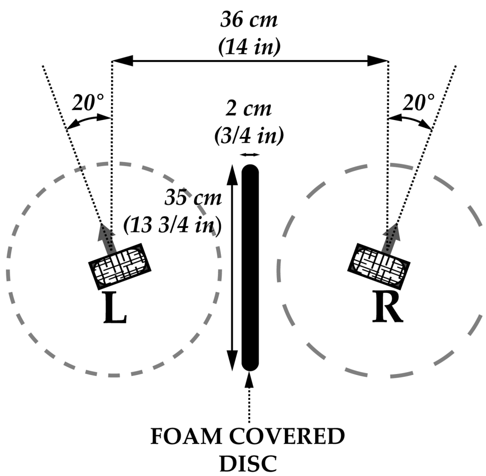

Figure 7.9 Jecklin Disc technique.

Disc baffle technique is shown in **Figure 7.9**. A circular baffle, 25 to 35 cm (10 to 14 in) in diameter, with both sides covered with acoustically absorbent material is positioned between two omnidirectional mics that are 36 cm (14 in) apart. The mics are angled outwards at +20° and –20° from center, although this angle is variable. This technique was developed to produce a pleasing image on headphones, and also to translate usably to loudspeakers.

Baffles are a way of increasing the channel separation and image width of any of the near-coincident techniques previously discussed. They are not true binaural techniques, because full HRTF effects are not simulated. They are also not 100 percent loudspeaker compatible, because of the partially simulated HRTF effects. They are generally considered to reproduce better on headphones.

PRACTICAL EXERCISE

Record a piano, acoustic guitar, drum set (or any other sound source that has width and takes up horizontal space) using a variety of the stereo mic techniques discussed in this chapter. Pan each channel hard left and hard right respectively, and listen carefully to the stereo images produced by each. Answer the following questions:

- When summed to mono how dramatic is the width change? When put back into stereo how much does the image expand and widen?
- When summed to mono do you hear tonal changes or does the amplitude get significantly quieter? (These would indicate mono compatibility issues.)
- How does the overall sense of width and space compare between the techniques? Do some sound wider and more spacious than others? Do some sound smaller and more compact?
- How clear and accurate is the image between the speakers? Can you really pinpoint the precise location of a specific sound? A good test for this is to record some of the mic techniques over the hammers of a grand piano and listen to the transition of the notes from left to right. How much do you hear the notes moving across the soundstage? Do they move slowly and smoothly across and through the center of the soundstage, or do they hang in a certain loudspeaker before suddenly jumping to the other side?
- Tonally or timbrally, how does each technique compare? Are some darker, muddier, or more muffled sounding than others? Are some brighter, more open, and more transparent than others?

8 Immersive Audio

In This Chapter:

8.1 Surround and Immersive Audio

Immersive audio loudspeaker systems have many loudspeakers positioned around the listener, each speaker channel capable of reproducing unique content. This produces a far more enveloping experience than stereo. The term *immersive audio* has replaced the term *surround sound*, although many people use the term "surround" to imply two-dimensional horizontal formats, and "immersive" to imply fully three-dimensional formats.

Digital 5.1 systems ("five dot one" or "five point one") have been around since the early 1990s – they have five main speakers arranged horizontally around the listener. A 7.1 system surrounds the listener with seven main speakers, with both side and back pairs of surround speakers. These systems have one dedicated *LFE* (*low frequency effects*) channel – the bandwidth limited ".1." (Outside the movie and television industry this channel is often referred to as the subwoofer channel.) Systems with up to eleven channels and either one ".1" LFE channel or true stereo ".2" LFE channels have been standardized. With all the main speakers on the same horizontal plane, these surround formats do not produce complete three-dimensional immersion.

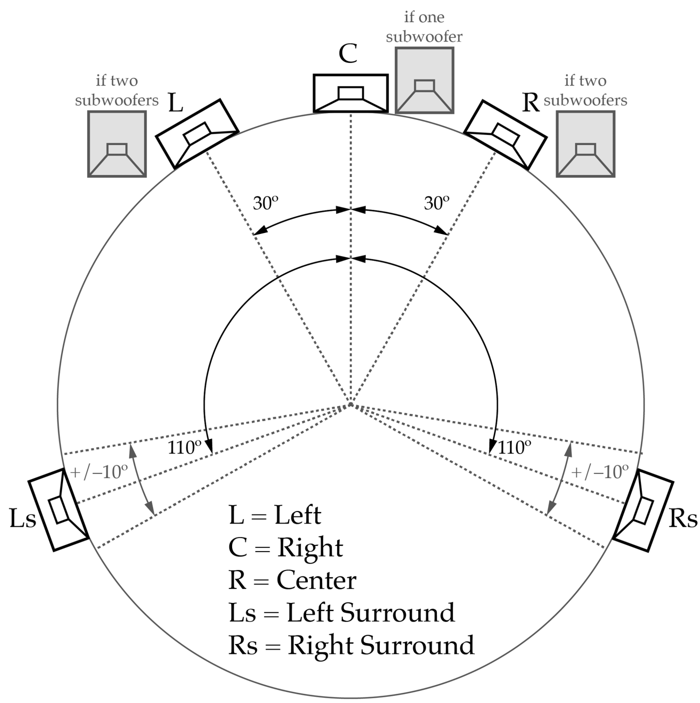

Figure 8.1 The Dolby/ITU recommended 5.1 loudspeaker arrangement for a content creation studio. All speakers should be the same distance from the listener.

CAREFUL WITH THE LFE CHANNEL…

In movie theaters and home theater systems the playback system reproduces the LFE channel 10 dB louder than the main channels – this gives the channel more volume and headroom for floor shaking explosions and sound effects. When subwoofers are used in the music production environment they are usually trimmed so the system presents a flat frequency response to the engineer – for correctly balanced low frequency content.

This means that the LFE channel should be used with caution in music production – either don't use it, or just use it a little bit so that if a low level of low frequency content is played back twice as loud as it should be, it's not going to be noticed.

With speakers on multiple horizontal planes, true *immersive audio* produces a fully three-dimensional experience. The additional "dot something" after the LFE designation in the format description indicates the number of height (top) channels.

Three competing commercial immersive standards in use today are Dolby Atmos, DTS:X, and Auro 3D. They are each different in the number of speaker channels, and suggested placement of those speakers. Their system-specific hardware and software are incompatible with each other. In an immersive music production, a project might be mixed in a studio set up to Dolby Atmos standards, but to be accessible to the largest number of consumers and possible dissemination formats, the mix also needs to play back acceptably on other speaker layouts including legacy 5.1 and 7.1 formats.

The specific technical specifications of each format in this still developing and fast changing field is beyond the scope of this book. Each system developer has a wealth of information on their website – look for the sections geared towards content creation. To summarize:

- Dolby Atmos allows many variations to the number of speaker channels – 5.1.2, 5.1.4, 7.1.2, 7.1.4 etc. up to 11.1.8 for home theater. For content creation studios, a minimum of a 7.1.4 system is recommended, to ensure mixes will translate to smaller or larger formats.

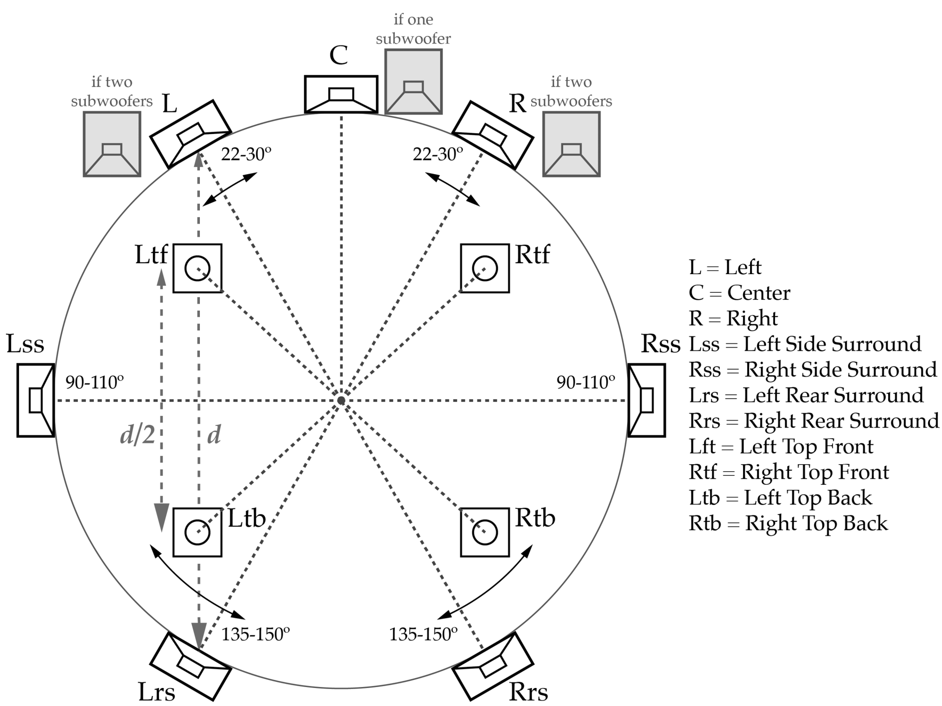

Figure 8.2 A 7.1.4 Dolby Atmos content creation studio speaker set up. The Ltf, Rtf, Ltb, and Rtb height speakers should be two to three times the height of the main speakers, and at an elevation of between 30º to 55º from the listening position. The front/back separation of the height speakers should be about half the length of the distance between the main front and rear speakers. All speakers should be the same distance from the sweet-spot, including the height speakers – all the distances indicated by the smaller dotted lines are the same when in three-dimensional space. If it is not possible to position all speakers equidistant from the listening position, electronic delays (and gain adjustment) should be used to align them.

- DTS:X is more flexible regarding the number and precise positioning of loudspeakers, and relies on the system's calibration and set up procedures in order for it to correctly send (*render*) sounds to the correct location in the room.
- Auro 3D features three layers of loudspeakers – an ear-level surround layer, a mid-level height or upper layer, and a top or overhead layer.

Music releases have been made in Dolby Atmos and Auro 3D formats, and as immersive audio is standard in movie theaters and high-end home theater systems, soundtrack content is commonly mixed in immersive formats. Dolby Atmos speaker layouts have become a common standard in mixing facilities including high-end music facilities.

8.2 Channel Panning and Object Based Audio

Panning in a conventional stereo system is *channel-based panning* – the pan control adjusts the amplitude of the sound between two specific loudspeaker channels, steering it between them. This is also known as *bed mode* panning in immersive audio – directing sound to and between actual loudspeaker channels (5.1.2, 7.1.4 etc.). The mix is stored as an audio file for each loudspeaker channel, and the relative amplitude of each sound in each speaker channel file creates that sound's position in the soundstage when they are all played simultaneously.

Object-based audio stores metadata about an *individual sound's position in the room*, along with that particular sound's audio file – and not which specific loudspeakers it is directed to. The audio file (*asset*) to be panned around the room is known as an *object*, or *sound object*. The playback (*rendering*) system uses the metadata to determine which loudspeaker or loudspeakers to direct that sound object to, and at what amplitude in each speaker. This enables the sound to appear at the correct location in rooms with different speaker configurations in them, and enables a project mixed in a small 7.1.4 Atmos studio to reliably upscale to a large movie theater Atmos system which could have up to 64 loudspeakers surrounding the audience.

There are limits on the number of object-based audio channels that can happen simultaneously. Most content in a mix tends to be channel-based, while special effects and specific moving sounds might be object-based.

It is relatively simple to transfer the location of bed panned channels from one system to another (Dolby Atmos to Auro 3D or DTS:X for example), because they are directed to actual speaker outputs in the DAW. Object panned asset data however, is proprietary to the system used to mix the production – so object panning movement *will not* easily transfer from one system to another.

8.3 The New Challenges of Immersive Audio

Before starting an immersive project, it is important to understand the format, and to record what is needed to use the format effectively. This could mean setting up different or larger mic arrays than for a stereo-only project, or recording more layers of sound to provide immersive panning options. It also means understanding the capabilities of the additional loudspeaker channels, and the imaging options they provide.

WHAT ARE WE LISTENING TO?

A lot of immersive content we are currently able to listen to is not captured using surround, immersive, or binaural microphone techniques exclusively – or at all! Many immersive mixes are created by positioning tens (or hundreds!) of mono or stereo sound sources in the immersive environment. This is the same concept as crafting a stereo mix from mono tracks in a DAW – but now in three dimensions!

Record It!

As with all recording, it is better to record what you think you'll need and end up not using it, than to *wish* you had recorded something. It may not be possible to record that extra layer or track for an immersive mix after the session or live concert – and even if it is, it is almost certainly inconvenient and expensive to do so.

What to Do With the Subwoofer?

In music production, to avoid the subwoofer in a movie theater or on home theater system producing low-frequency content twice as loud as it should, common practice is to send little or no audio to the dedicated LFE channel. This does not mean that the subwoofer will not do anything! Any immersive system using less than full-range main speakers will have a subwoofer and utilize a *bass management* system. The bass management system redirects frequencies below the effective lower limit of the main speakers to the subwoofer, where they are mixed with any content on the dedicated LFE channel. In a movie theater or home theater system, the bass management will control the level of this redirected low-frequency content so that it is reproduced at the correct volume, while the 10 dB of additional headroom in the LFE amplifier will be exploited by explosion and earthquake sound effects.

In a professional facility with good main monitor speakers, a bass management crossover frequency of 80 Hz is common, although 100 Hz and 120 Hz are also standard. Lower frequencies are generally preferred, to avoid higher frequency content localizing to the subwoofer. On home theater systems with small satellite speakers, the crossover frequency can be as high as 200 Hz or more – causing a lot of low and low-mid content to be perceived as coming from the subwoofer, and not the location the content was actually panned to.

What to Do With the Center Channel?

A lot of important content in a conventional stereo mix is panned center and perceived as a phantom image between the L and R speakers – lead vocals, bass, kick drum, snare drum etc. In stereo, a phantom center image is the only option. In immersive formats there are multiple ways to create a center panned image:

- As in stereo, direct a sound to the L and R speakers, and not use the center speaker, creating a phantom center image.
- Send a sound to only the dedicated center speaker.
- Combine both a phantom center image, and the center speaker – either at equal levels, or with the center speaker louder than the LR phantom image, or with the LR phantom image louder than the center speaker.

Phantom and discrete center images have very different qualities and "sizes" – phantom images appear bigger and less focused. Discrete center speaker images are tighter and more compact. A good, expansive immersive mix is often a combination of different types of center images.

Common practice in many commercial mixes is to use the center speaker for specific focal points and highlights, with the left and right channels carrying most of the front mix. Many engineers tend to keep the center channel dry, with little, if any reverb. The center channel master output levels are usually lower, or at most, equal to the left and right master output channel levels – they are not usually higher than the left/right master output levels.

What to Do With the Other Speaker Channels?

When first mixing in a surround or immersive format, it is very easy to think of the environment as stereo pairs of loudspeakers – front LR, side LR, rear LR etc. The danger of this is that the mix can end up being quite different in each stereo pair, with disassociated mixes on each plane – and it does not sound like a unified "whole." While a lot of the source tracks mixed together to create an immersive mix are mono or stereo, it is important to think "360°" and not only in stereo pairs. For example, to unite the front and rear planes of an immersive mix, there should be some common or similar content to bond them together. But there should also be enough different content to open up the space and give the mix size and dimension.

In commercial music a lot of mixes use keyboard pads, rhythm guitars, and reverb as common elements that glue the different planes together. Lead vocal (front) vs backing vocals (surround), contrasting percussion, different synth and sound effects, and different ear candy in each plane can create separation to make the mix more expansive.

Before mixing immersive content, be sure to analyze some immersive mixes to determine what several engineers do similarly (therefore what has been accepted as good practice) and identify what doesn't really work. It's still a developing field, and there's a lot less written about this subject than about stereo panning and mixing techniques. It can be difficult to get hold of immersive music mixes for analysis, but Blu-ray music discs exist, and streaming and downloads are possible. In addition to college and university facilities, conferences and pro-audio trade shows often feature immersive music listening and demo rooms set up by equipment manufacturers, and retail showrooms also exist in major markets.

8.4 Channel Based Microphone Techniques

There are many different immersive microphone techniques in use, and they are still being developed because the format is relatively new. This chapter contains a brief introduction to some different concepts, starting with some two-dimensional horizontal mic arrays, and then three-dimensional immersive mic arrays. Most immersive mic techniques were developed for classical concert hall recording, but similar principles can be applied in large studio live rooms.

IS THERE A ".1" MICROPHONE?

In music production there are no specific ".1" mics. Low frequency content is directed to the LFE channel from the main mics, or any spot mics on sound sources which have important low frequency content.

In creative productions or sound effect recording, anything goes though! A mic may well be set up specifically to pick up content that will only go to the LFE channel.

2D Horizontal Front Microphone Arrays

The stereo mic arrays and large ensemble mic arrays discussed elsewhere in this book can be augmented to provide the front components of surround mic arrays. Omnidirectional mics provide better low frequencies, and a very immersive, enveloping image. Cardioid and hyper-cardioid mics provide better channel separation, more open, and potentially more accurate imaging. For example:

- A near-coincident pair of hyper-cardioid mics can have a forward-facing center mic added to form an LCR (left, center, right) front array.

- LCR spaced array mics can be routed to the LCR loudspeakers respectively. These could be omni or cardioid mics.
- LCR mics in a Decca Tree arrangement can be routed to the LCR loudspeakers respectively.
- A middle-side pair can be matrixed to the left and right loudspeakers, and some of the M signal also sent to the center speaker. (The additional center speaker M signal might necessitate a lower than usual level of M in the LR matrix.)

Front mic arrays might be positioned closer to the sound source than for single array stereo recording, because they are not being used to pick up reverberant content – there will be rear-facing mics dedicated to that.

2D Horizontal Rear Microphone Arrays

Rear facing stereo arrays, provide content for the rear loudspeakers, and can be almost any stereo array except a Blumlein pair (which is avoided because the backside pick-up of the bidirectional mics will pick up the “front” sound, time displaced from the front array). As with front arrays, mics with more directional polar patterns can be used for increased channel separation.

Rear arrays are usually positioned close to the forwards facing array – either on the same front/back plane, or within a few meters (less than 10 ft) of the front array. In a concert hall this puts both off-axis musical content in the rear speakers as well as room sound/reverb – so it is important to use mics with smooth off-axis frequency response so that everything picked up sounds good. In a non-concert recording situation, musicians can be positioned “behind” the rear mic array so that musical content truly surrounds the listener.

2D Horizontal Front and Rear Mic Array Combinations

Omnidirectional arrays produce wide, dramatic, and expansive soundscapes, with beefy low frequencies – but the imaging is not the most precise or accurate. The spacing between the mics in this type of array can be varied, as shown in **Figure 8.3**. Larger spacings, particularly on the front-back axis, can cause image problems because of the significantly delayed time arrival of similar sounds into both the front and back arrays.

Cardioid arrays produce more precise, but less expansive images, due to the increased channel separation. In order to pick up the entire 360° sound field without gaps in the circle of pick-up, the mics are positioned closer together than in omnidirectional arrays. Using hyper-cardioid mics, the set up can become so compact that it becomes a near-coincident or coincident array.

These techniques can be upscaled to seven (or more) horizontal channels. Hybrid approaches employ mics with different polar patterns, for the benefits of each. Double

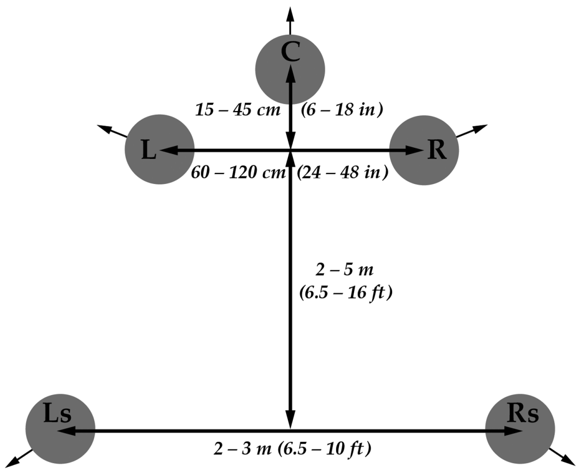

Figure 8.3 DPA Microphones suggest these minimum and maximum omnidirectional mic arrays spacings.

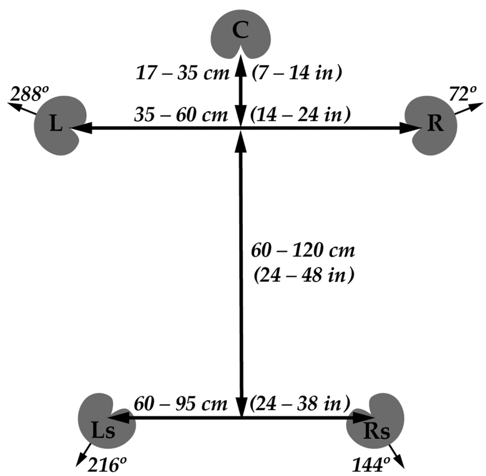

Figure 8.4 A cardioid surround array, which can be upscaled and downscaled between the suggested dimensions to provide more open and expansive, or smaller and more compact images.

MS (Mid-Side) techniques can also be used, forming single-point, mono compatible surround arrays.

OTHER 2D SURROUND TECHNIQUES TO KNOW ABOUT…

The concepts discussed in this section are applied in many surround microphone techniques that have been developed by recording engineers and researchers around the world. Each specific technique has a unique set of spacings and angles, producing different imaging characteristics. Some of these techniques include:

- The INA-5, Ideal Cardioid Arrangement (ICA) Surround array.
- The Optimized Cardioid Triangle (OCT) Surround array.
- The Fukada Tree.
- The Hamasaki Square.
- Double MS (Mid-Side).
- The Corey/Martin Tree.
- The Wide Cardioid Surround Array.
- The Polyhymnia Pentagon.

Combined derivatives also exist, which blend the front array from one technique with the rear array from another.

Immersive Microphone Arrays

To be truly immersive, a height layer of microphones is needed – commonly four, for the L top front, R top front, L top back, and R top back channels. Exact mic placement will depend on the chosen immersive dissemination/speaker format – but it is important to use and develop systems that translate well to multiple reproduction systems.

The *2L Cube*, shown in **Figure 8.5**, uses omnidirectional mics throughout. It features a layer of four height microphones in addition to either a five or seven channel lower layer, for 5.1.4 or 7.1.4 productions. The front, rear, front top, and back top mics are set up in a cube, with the center mic added on the bottom front axis. The dimensions of the cube can vary between 120 cm (4 ft) for large orchestra recording, and 40 cm (16 in) for smaller chamber music recording.

MONO COMPATIBILITY

Mono compatibility is important to surround and immersive mixes, just as it is to stereo mixes. While it's unlikely audio-only immersive content will get played back on a mono loudspeaker system, anything destined for broadcast (or that ends up on a soundtrack) might get played back on a mobile phone speaker or an old TV. The DAW/control room section mono button should be used to check that not too much is lost, and there's no phasey weirdness when the image is summed to mono.

Stereo Compatibility

The stereo version of a surround or immersive project should be a completely separate mix. Although there are plug-ins that will fold-down surround and immersive mixes to stereo, the process rarely works well. Compressing all the elements into the smaller stereo soundstage means increased competition for less physical space, and the mix loses clarity. EQ and amplitude changes are usually necessary in order to maintain clarity and balance, and the reverbs and creative effects may need to be tweaked. Also, time and phase differences between pairs of speakers (the fronts and rears for example) which might give an immersive mix spaciousness, can produce phase problems when the waveforms are electronically summed to the same stereo pair of loudspeakers (or to mono).

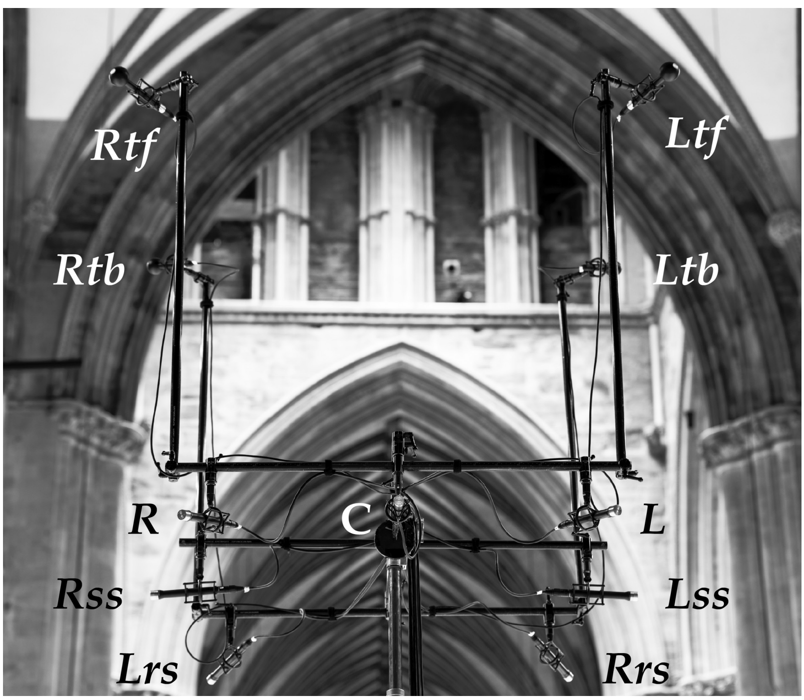

Figure 8.5 Morten Lindberg's 2L Cube array. Omnidirectional DPA 4003 or 4041 mics are used, and the height mics sometimes have *acoustic pressure equalizer* balls attached (which create an upper-mid/high frequency boost, and make the mics more directional at higher frequencies). Photo courtesy of 2L/Morten Lindberg.

Some height mic arrays are positioned a distance above the main horizontal array, providing decorrelated content and slight time differences that can increase the spaciousness of the image. Other mic techniques position the height mics on the same horizontal plane as the main mics for a potentially more precise and focused image. Most techniques position the height mics directly above (or adjacent to) the lower-level front and rear mics.

OTHER 3D IMMERSIVE TECHNIQUES TO KNOW ABOUT…

- The Decca Cuboid.
- ORTF-3D.
- The Hamasaki Cube (Hamasaki Square plus Height).
- OCT-3D.
- PCMA-3D.
- The ESMA-3D, Equal Segment Microphone Array.

8.5 Binaural Techniques

Binaural Recording

Many consumers listen to music and other content on headphones, which are binaural devices – two discrete channels of information, each going to a single ear with no inter-aural crosstalk. Humans have two ears, and can hear three-dimensional sound. This is because the distance between our ears, and our head and shoulder anatomy create unique time arrival, amplitude, and frequency filtering differences between the sounds arriving at each ear – Head Related Transfer Function or HRTF effects. Every individual person's brain has learned to interpret these unique differences and instantly calculate the direction a sound is coming from. Binaural mic techniques emulate this experience by using a mic array containing two microphones, positioned to simulate the ears on an average head, separated by some kind of head-shaped baffle. Binaural techniques only reproduce well on headphones – they do not translate well to loudspeakers.

Some binaural microphone technologies include:

- Dummy head microphones, which use two omnidirectional mics set into ear-shaped forms in the side of a mannequin head or head-sized sphere or ovoid.
- Miniature microphones that are positioned on a person's head. These can be either "in-ear," or "near-ear," and use the person's anatomy to capture real HRTF information.
- Other commercial products that have ear-shaped molds spaced appropriately, but are simply spaced or baffled apart, and do not feature an entire head-shaped barrier between the mics.

The perceived playback accuracy of binaural recording is not guaranteed, because most people's anatomy, and the associated and expected HRFT effects are different to those created by the geometry of the binaural mic used for the recording. In addition to the perceived image size being inconsistent between listeners, image localization and perceived timbre are also unpredictable.

Binaural Reproduction

Headphone reproduction of immersive content that was not originally captured using binaural mic techniques is now commonplace. There are consumer devices and software that can take a surround mix and render it for binaural reproduction. Software and plug-ins which will render fully immersive content for binaural reproduction are also available. Gaming, virtual reality, the music industry, and even broadcasters are using and experimenting with this technology so they can create and easily disseminate immersive content to consumers who do not have multi-channel loudspeaker systems.

Because every listener's head/shoulder anatomy and HRTF effects are different, rendering audio for binaural dissemination using generic HRTF algorithms is not 100 percent reliable – there is a discrepancy between what the renderer creates, and what each listener's brain expects, causing unique imaging and tonal inaccuracies for each listener. Many binaural rendering plug-ins allow the selection of different preset HRTF profiles in the form of SOFA files (Spatially Oriented Format for Acoustics). For content creation it is important to find the one that works best for the engineer/creator. It takes one to two weeks for a listener to learn and adjust to an alien HRTF – and as these are not customized for the actual user, imaging and tonal inaccuracies do not completely disappear.

MORE ON HRTFS

A complete HRTF profile has two components that interact, and are unique to each listener:

- Sound coloration effects, called Common Transfer Function (CTF) components.
- Direction specific effects, called Directional Transfer Function (DTF) components.

In many cases, generic HRTF profiles only contain DTF information, assuming that the headphones will supply adequate CTF effects by design. Lacking necessary listener-specific CTF components, and their interaction with DTF components, this common "standard headphones plus generic HRTF profile" approach does not work well – although after a two week learning period the experience might become "better."

For the best experience it is necessary to custom equalize the user's headphones to remove the (diffuse field) equalization built in to them, and replace that coloration with the listener's unique CTF profile.

The latest advances in binaural reproduction technology include being able to customize HRTF profiles for the specific listener, based on their unique anatomy – including both CTF and DTF components. Genelec's *Aural ID* computers analyze a 360° video of the head and shoulder region to create personal SOFA files. These customized SOFA files then replace the stock, generic algorithms in binaural rendering plug-ins – a DTF-only version for use with standard headphones, and (preferable) a full CTF + DTF version for use with custom de-equalized headphones. Steinberg's *Immerse* analyzes a picture of the user's right ear to generate personalized SOFA files. Aimed at professional content creators, these technologies give the user a much more accurate and compelling binaurally simulated immersive experience than stock SOFA file based systems.

8.6 Introducing Ambisonics…

The science behind Ambisonics is quite heavy, and is not going to be explained in detail here! This section is an introduction to the concepts of more commonplace recording and playback systems.

Developed in the 1970s, Ambisonics had a cult rather than mainstream following due to its complexity, and the very expensive hardware it required. There is now renewed interest and development in the format, because of the use of immersive audio in gaming and virtual reality, and the hardware and software now being available at much lower prices – some software is even free!

Unlike the channel-based microphone and reproduction systems described earlier in this chapter, an Ambisonic microphone, and the data stream transmitted to the consumer, is loudspeaker channel-independent – it is decoded for each listener's unique loudspeaker set up, or headphones. The system encourages the content creator to think of sound sources as coming from directions, and not specific loudspeaker channels.

Ambisonic Microphones and Recording

In *First-Order Ambisonics* (*FOA*), a sound source's position is represented on three axes relative to the microphone:

X = front/back
Y = side, L/R
Z = up/down.

An additional W signal represents the overall amplitude of the soundfield. The W, Y, Z, and X signals are known as *B-format* signals. They do not relate to actual reproduction or loudspeaker channels – instead they represent the three-dimensional soundfield at wherever the mic is positioned.

It is possible to create a first-order B-format Ambisonic mic system using a triple-MS array of four coincident mics, as shown in **Figure 8.6**.

- A forwards-facing omnidirectional microphone picks up the W signal.
- A side-facing bidirectional microphone picks up the Y component.
- An upwards facing bidirectional microphone picks up the Z component.
- A forwards-facing bidirectional microphone picks up the X component.

The Y, Z, and X mics are matrixed with the W signal (usually in software plug-ins) at varying amplitudes, allowing the creation of:

- Many outputs, each representing a mono microphone of any pick-up pattern (from omnidirectional through cardioid to bi-directional), pointing in any direction.
- A coincident stereo array pointing in any direction.
- Multiple coincident stereo arrays pointing in any direction, simultaneously, to create surround and immersive mic arrays.

Because B-format signals are recorded before they are matrixed to create loudspeaker output channel signals, the resulting mic array(s) can be decided upon, manipulated, or changed after recording has taken place.

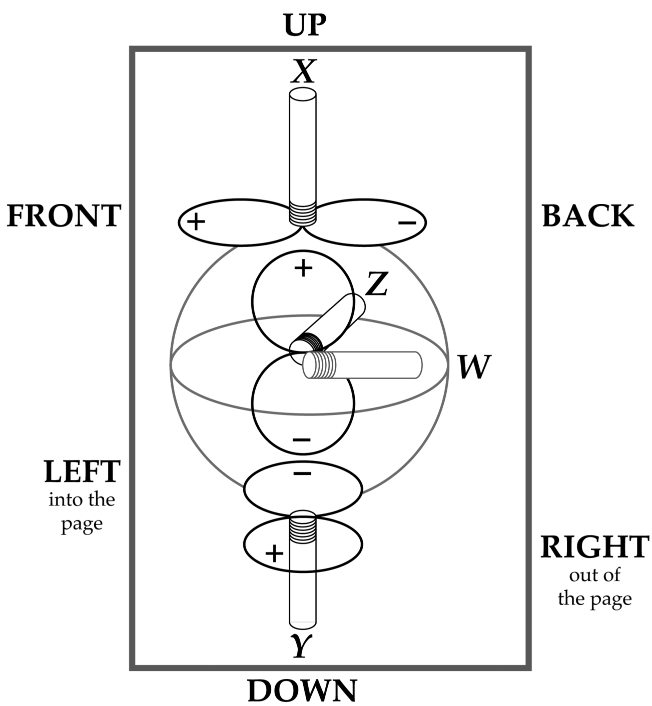

Figure 8.6 A first-order Ambisonic array created using four conventional microphones – an omnidirectional W mic, and bidirectional Y, Z, and X mics.

This array explains the principle of first-order B-format signals – however it is not widely used because it is difficult to set up and get the capsules close enough to truly be coincident, and each mic acoustically shadows the others, blocking the sound trying to get into the other capsules.

A more compact solution is the first-order *A-format* tetrahedral cardioid capsule arrangement found in single body Ambisonic microphones, as shown in **Figure 8.7**. The small diaphragm condenser capsules are positioned extremely close together, so they function almost coincidentally, and do not acoustically shadow each other. A first-order A-format tetrahedral array has four capsules:

FU = front upper
RU = rear upper
LD = left lower
RD = right lower

Before the mic's signals can be manipulated, the A-format signal must be converted to B-format. Some microphones have a matrix in the microphone that does this, others use an external box, and others rely on plug-ins to do this. Once the signal is in B-format it can be manipulated using Ambisonic plug-ins, used in Virtual Reality and gaming, and converted for different listening formats.

Higher-Order Ambisonic (*HOA*) microphones have more capsules, enabling them to capture a sound source's position relative to more axes. This results in greater precision and imaging accuracy on playback. Second-Order microphones have eight capsules, outputting eight channels of A-format signals, which require software to change them into a nine channel B-format signal for processing, transportation, and dissemination. At the time of writing this book a 32 capsule fourth-order microphone exists (producing 25 B-format channels), and a 64 channel seventh-order microphone is in development.

To produce an Ambisonic mix, it is not necessary to record everything with Ambisonic microphones. Non-Ambisonic sound sources can also be used in Ambisonic productions. There are plug-ins available which position traditional mono, stereo, or multi-channel

Figure 8.7 A First-Order Ambisonic mic with a tetrahedral array of cardioid capsules.

surround and immersive signals in the Ambisonic soundfield, and convert the result to B-format signals.

In terms of mixing and production, in order to future-proof Ambisonic projects, it is best to record and mix/process at the highest-order that is realistic. Many consumer applications currently feature first, second, and third-order compatibility, but this will probably increase. Some research has shown that *given current technology*, consumers don't perceive much imaging benefit above third-order. But as technology develops, that could change – so other researchers recommend seventh-order as a future-proof format.

AMBIX AND FUMA

There are two conventions for channel ordering of Ambisonic signals, *AmbiX* (WYZX) and *FuMa* (WXYZ). To complicate things, different software and hardware companies have adopted different formats – so make sure you know which convention your technology is using before inputting and trying to process signals. Software and plug-ins are available to convert the formats.

Ambisonic Playback Systems

B-format Ambisonic signals are completely independent from the playback system and are decoded or transcoded specifically for each unique playback environment – which can range from headphones (binaural) to loudspeaker arrays with as few as three or four speakers (the minimum recommended for horizontal two-dimensional playback) to large three-dimensional multi-speaker arrays (including specific Ambisonic arrangements, other proprietary domes, and also more mainstream systems such as Dolby Atmos). While creating content in a DAW, plug-ins decode the B-format signal for the studio's speaker set up – requiring data about the number of loudspeakers, and their location to be entered. B-format decoding can also happen on some consumer games and virtual reality devices – most commonly decoding for headphone playback. This means that only one B-format mix needs to be produced, and it will be decoded or transcoded as needed for specific playback environments, by each playback system.

One of the simplest three-dimensional Ambisonic loudspeaker layouts is a cube of eight loudspeakers, arranged as if they were in room corners in front and behind the listener (although they may not be in the actual corners). Four loudspeakers are in a square below the listening position, and another four above the listening position, each equidistant from the listening position, as shown in **Figure 8.8**. Larger systems can have tens of speakers surrounding the listener in optimized positions.

Amisonics also provides the ability to change one channel-based recording or mix format to another, or to convert non-Ambisonic channel-based immersive microphone array recordings for different loudspeaker formats. To do this, an Ambisonic encoder is used to

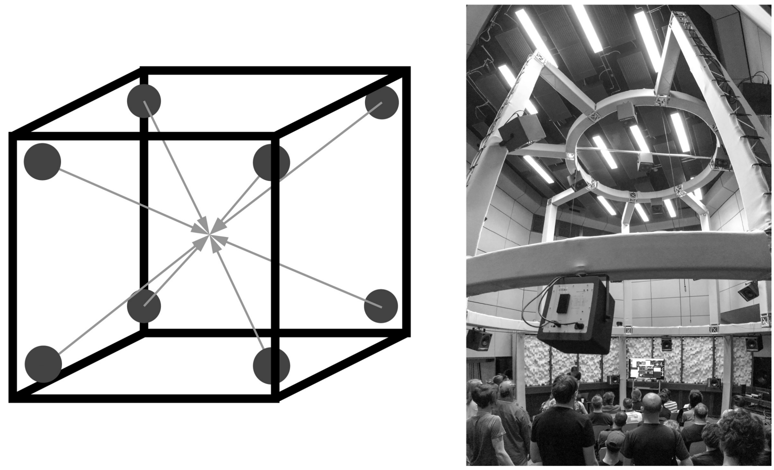

Figure 8.8 **Left:** A simple 3D Ambisonic loudspeaker set up, the grey balls represent loudspeakers pointed towards, and equidistant from, the listening position in the center of the cube. The front center of the soundfield is in the center of one of the "walls." **Right:** A 33.2 dome with five layers (4, 8, 8, 8, 4+1, floor to top), used for higher-order Ambisonic work (Hamburg University of Applied Sciences, Germany).

convert the channel-based mic signals to B-format, and then a decoder is used to change that B-format signal to a different channel-based format.

HEADPHONES TO THE RESCUE?

Research institutions may have facilities with large numbers of loudspeakers correctly positioned around the listener, in dome and inverted dome arrangements, to reproduce higher-order Ambisonic signals accurately – but for the rest of us, that is not practical!

B-format Ambisonic signals can be rendered (transcoded) into binaural formats for headphone listening. Most commercially available software does this by creating the illusion of an array of virtual loudspeakers around the listener. The higher the order of the transcoding, the more virtual loudspeakers surround the listener. Images tend to be precise and focused when a sound source's location is close to one of the virtual loudspeaker positions, but the image can become unfocused, separated, and spread, when the sound's location is between virtual loudspeaker positions. Higher-order transcoding provides more virtual loudspeaker positions. First-order transcoding is not recommended because it significantly reduces the spatial and timbral accuracy of the image. Second-order transcoding is much more reliable. Research has shown that given current transcoding technology, third-order transcoding is optimal, and there is no perceived advantage to transcoding at orders higher than this. There are software packages and plug-ins available to do this transcoding, and they even come as standard in some DAWs.

Another unique feature made possible by Ambisonic processing for headphone immersive audio systems is *head tracking*. The listeners head position is tracked (rotation, inclination, and tilt etc.) and this movement causes the soundfield in the headphones to be adjusted in real time. This makes gaming and virtual reality a more lifelike experience. It could even make watching a recorded concert video a more realistic home experience!

There is a downside to binaural rendering though – the timbre of sound sources can be compromised by the HRTF processing, depending on the sound's position in the soundfield.

3D Headphone Mixing at Home?

There are binaural panning plug-ins available (even free or stock ones in some DAWs!), that when used on DAW channels allow simulated 3D immersive audio mixes to be created. These provide an "out of the head" experience, with an increased soundstage compared to standard stereo – but the perceived results vary from listener to listener, are not as accurate as a multi-channel loudspeaker system.

AUDIO EXAMPLES
Can be found on the companion website

Immersive Audio via Headphones

Example 8.1: A recording made with a binaural pair of mics.
Example 8.2: An immersive mix, created from a mono/stereo multi-track recording, rendered to a binaural format so the consumer can experience immersive audio over headphones.
Example 8.3: An Ambisonic recording rendered to a two-channel binaural format so the consumer can experience immersive audio over headphones.

9 The Effect of Microphone Position

In This Chapter:

9.1 Art and Science

Small changes to a microphone's position can significantly affect the sound it picks up. There is no one "correct" position when miking a particular sound source, because the sound a mic picks up is dependent on many variables, including:

- The specific mic in use.
- Its distance and position relative to the sound source.
- The angle of pick-up relative to the sound source.
- The size and shape of the room.
- The mic's position relative to the room.
- The sound source's position in the room.
- The mic's position relative to any other sound sources in the room.

This short list should explain why in addition to being a little bit of a science, audio recording is an *art*!

9.2 Distance and Tonal Qualities

The distance at which a mic's low frequency proximity effect boost becomes apparent varies significantly between different mic models – even between those of similar pick-up patterns. It certainly varies between mics with different pick-up patterns – the more directional the mic, the more pronounced the proximity effect.

Proximity effect can be exploited in beneficial ways. A thin sound might be warmed up by positioning a directional mic relatively close to the sound source. Conversely, a sound that is quite boomy and full might be made less muddy by moving the mic further away. A few inches difference in mic distance can have a significant effect on the amount of proximity effect.

For most rock and pop recording, particularly in smaller project and home studio rooms, directional mics are commonly used. Proximity effect "mud" is cumulative and somewhat exponential – the more tracks it is apparent on, the bigger the problem. To reduce proximity effect problems, don't automatically position a directional mic super close to each sound source – consider whether it can be backed off a little. The increase in transparency and clarity will make mixing easier because less corrective EQ will be required.

A close mic picks up details we may not be used to hearing at a normal listening distance – powerful indicators of a close proximity to the sound. Those details may not travel to more distant mics.

Air naturally absorbs high frequencies over distance, so if a mic is positioned far away from a sound source, the sound will be darker and less bright, and we will perceive the sound as being further back in the stereo image, particularly when combined with room reflections and reverberation. If you want a sound to naturally appear "back" in the mix, try recording it from a greater distance.

9.3 "Zoom Factor"

We do not generally listen to sound sources from just a few inches away – we usually listen from several or many meters or feet away. Listening to a piano from a normal distance we hear all the constituent elements of the piano which blend together over the distance between the instrument and our ears:

- The impact of the hammers on the strings.
- The warm resonance of the strings and soundboard.
- In the case of the grand piano, the reflected sound coming from underneath the instrument, and bouncing off the open lid.

A directional microphone used close is like a zoom lens on a camera. It focuses in on the sound coming from just a small area of the much larger sound source. For example, a microphone placed very close to a saxophone's bell, as in **Figure 9.1**, focuses on the unflattering honky sound coming from the bell, and not the complete or natural sonic picture of the sounds emanating from all over the instrument.

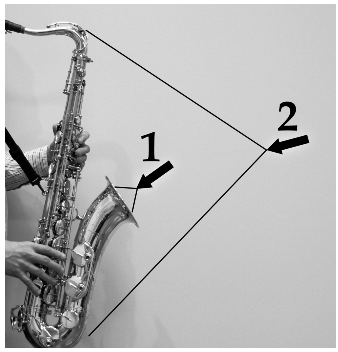

Figure 9.1 Correct and incorrect mic positions on a saxophone. **1:** Put too close to the bell, the mic picks up a honky sound that favors the lowest few notes. **2:** A mic positioned further away will pick up the sound of the entire instrument.

THE LONGEST DIMENSION RULE

A good rule of thumb to pick up a natural sound of most acoustic instruments in the studio is to take the longest dimension of the instrument and position the mic that far away (or perhaps just a little closer if it's more than about a meter or 3 ft). This gives the component sounds radiating from all over the instrument distance to blend together and form the "whole" sound we are used to hearing.

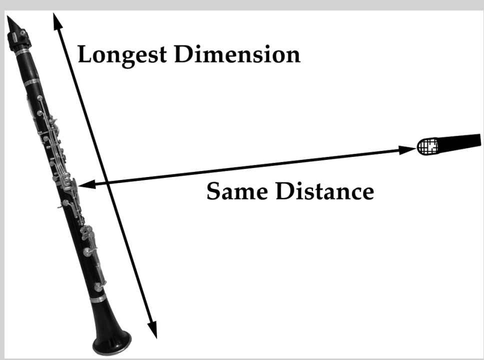

Figure 9.2 "Longest dimension" positioning on a clarinet.

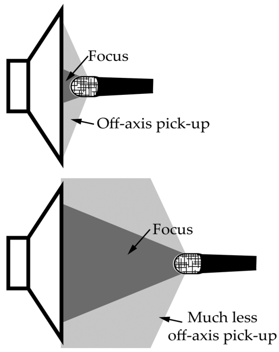

Figure 9.3 Positioning a mic further away allows it to pick up the sound of the entire cone, with less off-axis coloration.

A directional mic positioned right up against the grill of a guitar or bass amplifier cabinet will focus on the sound made by just a small portion of the loudspeaker cone, as shown in **Figure 9.3**. Backing the microphone off by a foot or so will not only make the sound less muddy and boomy, but also pick up the characteristics of the entire speaker cone. Instrument cabinet loudspeakers are very non-linear in their response, exhibiting complex distortion patterns over their entire surface – they are much less accurate than monitor loudspeaker drivers, but much more characterful. Often, a more desirable, detailed, less congested, and more mixable sound can be achieved by introducing a little distance between the speaker cone and the mic.

INCREASING MIC DISTANCE

A very natural acoustic guitar sound can be achieved by using a single mic 45 cm (15 in) or more from the instrument. This extra distance results in:

- Less proximity effect, and a clearer sound.
- A natural balance of all of the instrument's component characteristics.
- Less potentially undesirable off-axis coloration because more of the instrument is on-axis to the mic.

The tradeoff of this approach is that the close perspective is lost and replaced with a more distant sound.

Using an omnidirectional mic instead of a cardioid mic will eliminate proximity effect and make the off-axis pick-up sound better, so omnidirectional mics can be positioned a little closer than directional mics without sounding so isolated. But omnidirectional mics should only be used if the recording room characteristics are desirable – in most small home or project studio rooms this is not the case.

PRACTICAL EXERCISE

Record an acoustic guitar with a single microphone. Keep the mic close, about 12 cm (5 in) away from the guitar.

- Position the mic directly in front of the sound hole. A thick, beefy, body sound should be picked up. The mic is close to, and on-axis with the sound hole, and that character dominates.
- Position the mic a few frets further up the fretboard from where it joins the body of the guitar. Zingy, stringy, sizzly details should be picked up from this mic position. The mic is close to, and on-axis to the neck, and that character dominates.

Neither of these sounds accurately represents the acoustic guitar, because of the mic's "zoom factor" – but each is an important part of the instrument's complete sound character.

- If the mic is put in-between these two positions, perpendicular to the guitar, a blend of both can be picked up.
- Move the mic left or right by 5 or 6 cm (2 in) to change the relative balance of these sound components, and listen to the balance of the characteristics picked up.
- Move the mic a few more times until you find an appropriate blend that sounds like the guitar's "whole" sound – a good balance of "body" and "string."

Increasing the mic's distance will reduce the "zoom factor," and allow a more natural sound to be picked up.

- Position the mic about 30 cm (1 ft) from the guitar and change its position (up, down, left, right) by small amounts until you find the best balanced, natural sounding sweet-spot.

9.4 Off-Axis Response

All microphones (even omnidirectional mics) become more directional at higher frequencies – so it is important that mics are always pointed in exactly the right direction.

- When recording physically small sound sources (point source images such as vocals or brass instruments, for example) the results of positioning the mic on-axis vs off-axis are relatively predictable – a mic sloppily or inaccurately positioned will have its sound colored by the mic's off-axis response, and the sound will become muddy and muffled, or just unpleasant.
- On larger sound sources such as drum sets, acoustic guitars, string basses, and large acoustic ensembles, sound from parts of the sound source the mic isn't specifically aimed at are in the off-axis part of the mic's pick-up pattern.

Consider the acoustic guitar miked in the previous exercise:

- A mic positioned to get great pick-up of the bright, zingy string sound, is off-axis to the sound hole and most of the guitar's body – the sound of body and sound hole will be subject to the mic's off-axis coloration, and will possibly introduce dull, congested, and muddy characteristics, compromising its usability.
- A mic positioned to get great pick-up of the big warm body sound is off-axis to the sound coming from towards the neck – the stringy, fret board sound is subject to the mic's off-axis coloration, and will not be as bright, zingy, or clear as it should, possibly compromising the track's usability.
- Angling the mic a little left or right will change the recording's characteristics by putting different components more on or off-axis. If you want more string sound, but with a darker "zing," then try positioning the mic closer to the neck, angled towards the sound hole.
- An overly bright sound can be darkened by positioning the mic so the sound source is slightly off-axis – if the mic has good, smooth off-axis response characteristics.

The only way to find the best mic position is to experiment – make small changes and listen to the mic in different positions, and/or change its angle slightly to favor the desired characteristics.

AUDIO EXAMPLES
Can be found on the companion website

"Zoom Factor"

Example 9.1: A mic set up close to an acoustic guitar's sound hole primarily picks up the body of the sound colored by proximity effect.
Example 9.2: A mic set up close to an acoustic guitar's neck primarily picks up the lighter stringy quality of the sound colored by proximity effect.
Example 9.3: A mic placed 45 cm (18 in) away picks up a clearer, more accurate, blended picture of the guitar's natural sound, and is not colored by proximity effect.

9.5 Direct vs Reflected Sound

Diagrams explaining the travel and trajectory of sound waves in an environment usually show sound travelling in straight line laser beam patterns from the sound source to the microphone. That's a simplification of what really happens. Yes, a lot of the sound does do that – but not all of it. Most acoustic instruments and sound sources radiate sound in all directions – but it's more complex than that… Different frequency content is radiated in different directions, and not all of that frequency content remains equally audible over

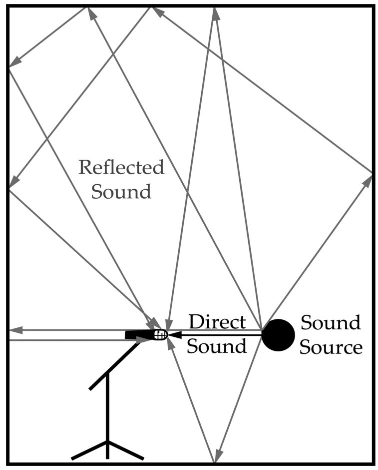

Figure 9.4 Direct and reflected sounds are picked up by a microphone. Complex "room signatures" are quickly generated by the infinite number of reflections in the room.

the same distance as other frequency content! Sound travels outwards to the floor, walls, and ceiling, which reflect it back into the room towards the mic, and the other boundaries, which reflect it again. An infinite number of reflection paths are quickly created, which give the room its character. See **Figure 9.4**.

The closer a mic is to a sound source, the greater the ratio of direct to reflected sound – the sound will be more dry, dead, in-your-face, and up-front. As a mic is moved further away from a sound source, the level of direct sound decreases and the relative amount of reflected sound increases. This means that the sound of the room or environment is heard more. A more distant mic's sound can be anything from "a little less up close and personal, but beneficially more live," to "quite set back, unfocused, reverberant, and distant."

Most of the time we listen to musical instruments and singers in a room – a space with a floor, walls, and a ceiling. The reflections from those boundaries are an essential part of the natural sound we are accustomed to hearing. An uncomfortable sensation of "deadness" is experienced when walking into a recording studio's dry vocal booth for the first time because the reflections and room sound that we are used to hearing are missing.

Unnaturally dry sounds, lacking reflected content certainly have their place in contemporary music production styles – alongside more natural or stylistically wet, reverberant sounds.

9.6 Comb Filtering Problems

Consider a guitar loudspeaker cabinet – comb filtering can be produced when multiple mics are different distances from the speaker cone, shown in **Figure 9.5 Top Left**. Sound takes approximately 2.9 ms to travel each meter (0.9 ms to travel each foot), so the wavefront is picked up slightly later in the more distant mic. At any single moment in time, different frequencies arrive at each mic at different points in their relative phase cycles. When

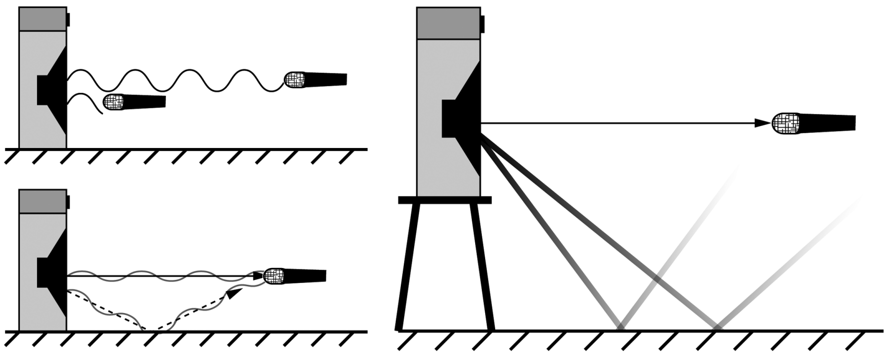

Figure 9.5 **Left:** Causes of comb filtering. **Top:** Between multiple mics, each a different distance from the cone. **Bottom:** A single mic more than 15 cm (6 in) from the loudspeaker cone. **Right:** Raising the cabinet reduces the amplitude of the floor reflections picked up by the microphone.

these time displaced signals are electrically summed in a mixer or DAW, the thinning or phasey hollowness of *comb filtering* occurs. Changing the distance between the two mics will change the frequencies at which comb filtering occurs.

Also, when a single microphone is positioned more than about 15 cm (6 in) from a loudspeaker cone, as shown in **Figure 9.5 Bottom Left**, the sound takes two different length paths to the mic. Each wavefront arrives with different frequencies at different points in their relative phase cycles – so they comb filter when they sum acoustically at the mic. Raising a cabinet off the floor can decrease the amplitude of floor reflections, causing many of them to bounce behind the microphone, in addition to their amplitude decreasing because of the greater distance to the mic, as shown in **Figure 9.5 Right**.

High frequency comb filtering can be reduced by putting carpet, blankets, or acoustical absorption products on the floor between the cabinet and mic – but this may also deaden the sound, making it duller and less exciting. Low frequencies do not suffer comb filtering problems at the miking distances described here, because their wavelengths are very long compared to the miking distances, so the LF arrives at each mic more in phase.

POSITION THE DIAPHRAGM, NOT THE GRILLE

If you are positioning multiple mics at identical distances from a sound source, make sure you know the position of the actual diaphragm inside each mic you use, and make sure the diaphragms line up (and not necessarily the front grille of each mic). The very front of each mic is not where the diaphragm usually is.

In a DAW it is easy to delay a close mic to a distant mic, or advance a distant mic to a close mic. This can resolve some direct sound phase issues, but cause other spill and reflected sound issues, so *it is always best to resolve phase problems acoustically, by changing mic positions*.

AUDIO EXAMPLES
Can be found on the companion website

Comb Filtering – Problems

Example 9.4: An electric guitar cabinet. A close mic and distance mic are combined together. The strange phasey thinning of comb filtering can be heard. It is usually desirable to minimize this type of artifact by adjusting the position of the distance mic.

PRACTICAL EXERCISE

Learn to identify the sound of comb filtering:

- Copy and paste identical solo instrument material to two tracks in your DAW.
- Pan both to the center.
- Use a delay plug-in to delay one of them by 1 to 20 ms.
- Play them both together, and listen to the timbre changes as you change the delay time by a millisecond at a time.
- Each time, bypass or solo one of the tracks so you are just listening to one of them. Listen for how the sound gets fuller and more natural sounding without the effect of comb filtering.
- Polarity reverse one of the tracks to hear the difference between them – this is what causes comb filtering.
- Try this with recordings of different instruments.

9.7 Floor Reflections – the Good, the Bad, and Boundary Mics

When a sound source is close miked in a large room (even a fairly reverberant one), the amplitude of the reflections from the walls and ceiling are of a relatively low level compared to the dry, direct sound reaching the microphone. This is due to the fact that the reflected sound's amplitude drops as it travels a longer path to the walls and ceiling, and then back to the mic. Even at low amplitudes, these reflections can have a big impact on some sound sources. The floor is much closer to both the sound source and microphone, so the reflections from a hard floor are temporally closer to the dry sound, and louder than the reflections from the walls and ceiling – and they have a significant impact on the recorded sound.

Floor reflections are not long swishy reverb tails – they are power and energy, indistinguishable from the dry sound. They can add life and excitement to a sound, and we are used to hearing them as part of a natural listening experience.

- A mic very close to a sound source picks up mainly dry sound which overpowers and masks reflected content – and the result is a dry, dead, in-your-face sound.
- Moving the mic a meter (3 ft) away allows floor reflections to become audible.
- Moving the mic several meters (many feet) away decreases the dry, direct sound and increases the amount of wall and ceiling reflections (room sound or reverb) that the mic picks up.

Hard reflective floors can add beneficial reflections to a recording – but significant floor reflections can also cause phase problems and comb filtering due to the distance and consequential time arrival differences between the direct and reflected sound at the microphone. It is important to really listen for any thinning or strangeness in the timbre of the recorded sound source, in all its pitch and frequency ranges, and at all dynamic levels to confirm that the sound is not being negatively impacted by excess reflections.

If the sound of comb filtering is identified, the mic position should be adjusted to minimize the problem. Moving the mic just a small amount can change the phase relationships of the direct and reflected sound at the mic position, reducing or altering the affected frequency range(s), making things better – or worse! Do not position a mic and assume it is picking up optimal sound. Try a few different positions – they may sound worse, but you'll have verified that the initial placement was best.

Thin, lightweight carpet and rugs can cause poor sounding floor reflections because they only absorb high frequencies – so the reflections are muddy and dull. Many instruments, such as acoustic guitars and drum sets, really benefit from the bright hype of hard floor reflections. Most studios have hardwood floors because the reflections sound great – and if necessary the room can be damped down with thick, heavy carpet or rugs.

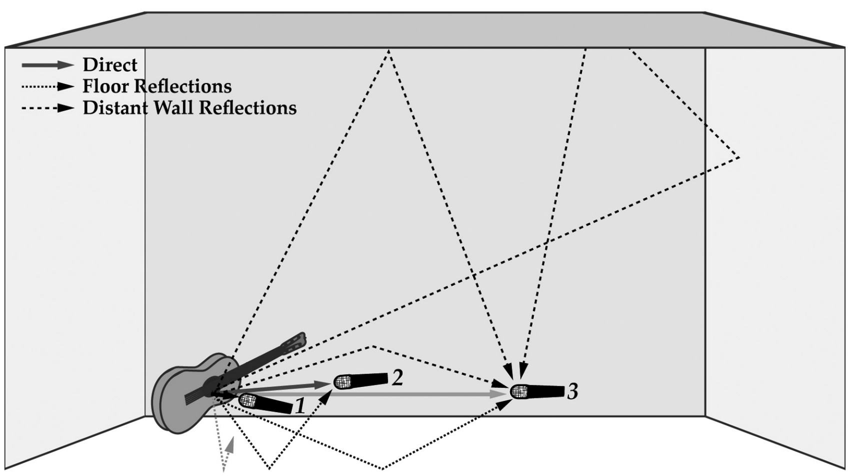

Figure 9.6 **1:** A close mic favors dry, direct sound. **2:** A few feet away, floor reflections become part of the sound. **3:** With a more distant mic position reflections from the entire room (reverb) are picked up by the mic.

Boundary mics positioned on the floor exploit floor reflections as part of their design. They need the floor to give them good low frequency response. Because the diaphragm is only a few millimeters above the floor, the phase and comb filtering problems caused when the direct and reflected sound sum at the mic capsule are in such high frequencies that they are irrelevant. This does not mean that a boundary mic will always pick up good sound though! Many instruments radiate their high frequency content quite directionally, and not necessarily towards a mic positioned below the plane of the instrument. So it is often better to use a conventional mic, up and away from the floor to capture the best blend of direct sound and floor reflections.

AUDIO EXAMPLES
Can be found on the companion website

Floor Reflections

Example 9.5: An acoustic guitar recorded in a carpeted room.
Example 9.6: An acoustic guitar recorded in a room with a reflective wooden floor.

9.8 Stereo Arrays and Distance

The top diagram in **Figure 9.7** shows a near-coincident array positioned relatively close to a sound source. The sound source extends through most of the effective pick-up area of the array. Panned hard left and hard right during mixing this results in an expansive image that takes up the entire width between the loudspeakers.

In the bottom diagram, the same array is moved back from the sound source. The sound source is now more concentrated in the overlapping center pick-up of both mics.

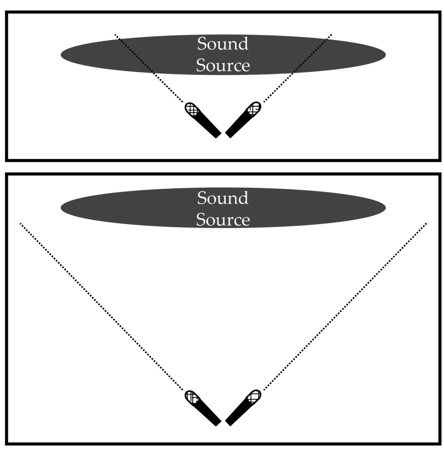

Figure 9.7 **Top:** A near-coincident mic array close to a sound source creates a wide image. **Bottom:** Moving the mic further away from the sound source results in a narrower stereo image.

Even panned as wide as possible during mixing the image will be narrower than when the array was closer to the sound source.

But – image width is not all that changes as the distance between a stereo array and sound source is changed:

- When a mic or array is positioned closer to the sound source, the amplitude of direct sound is increased and the relative amplitude of reflected room sound is decreased – the sound is drier and more up-front.
- Moving the mic(s) further away from the sound source increases the relative amplitude of reflected sound and decreases the amplitude of the direct sound – the sound is more distant and reverberant.

The positioning of stereo arrays is a balancing act between the desired image width and wet/dry balance:

- A stereo array up close = a wider, but drier image.
- A stereo array further away = a narrower, but more reverberant image.

You should not compromise desired image width for the sake of wet/dry balance or vice versa. As previously discussed, different stereo arrays produce different widths and imaging characteristics – you should experiment with different arrays and/or different microphones until the desired combination of image width and reverb is obtained.

CASE STUDY

PROBLEM: An acoustic guitar miked with an XY coincident pair is appropriately wet or dry, but the image is too narrow.

SOLUTION: Try a near-coincident pair at the same distance, in order to widen the image without changing the wet/dry balance too much.

PROBLEM: An acoustic guitar miked with a near-coincident pair is too reverberant or “roomy,” but the image width is good.

SOLUTION: Try a spaced pair of cardioid mics a little closer – for a wider but drier image at a closer distance.

AUDIO EXAMPLES
Can be found on the companion website

Stereo Array Distance, Image Width, and Wet/Dry Balance

Example 9.7: A drum set recorded with a near-coincident pair, above and slightly in front of the drum set. The image of the drum set is wide, and relatively tight and dry.

Example 9.8: The same drum set, with the same stereo mic array positioned 20 feet in front of the drum set. The drum set appears narrower, taking up less of the image width, but an increased amount of reverb takes up the entire width of the image.

9.9 Spill – Enemy or Creative Tool?

Whenever multiple microphones are positioned close to each other, each one will pick up sound that's not really intended for it. This spill is problematic if it does not sound good in itself, or if comb filtering and mono compatibility problems exist between multiple mics in the same room.

THE 3:1 RULE

In any multi-mic situation (including close drum mics, or recording multiple singers or a horn section with individual mics) the 3:1 rule should be used in order to reduce the level of spill relative to the desired sound in each mic. The distance between each mic should be three times the distance between each sound source and its mic. This will certainly not eliminate spill totally, but it will minimize its potentially negative effects.

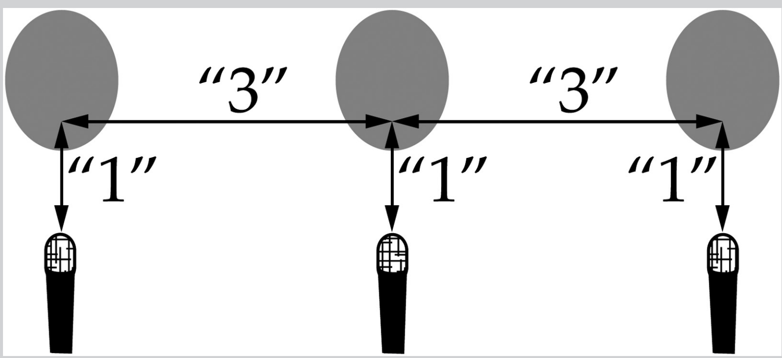

Figure 9.8 Applying the 3:1 rule when miking three singers.

When mixing rock and pop music, *gates* or *strip silence* features can be used to silence mic channels between drum hits. Gating out spill and removing the natural acoustical connection between the multiple instruments can make it become ultra-clean and surgical. This may be stylistically and aesthetically desirable for some rock and pop projects because the drum set tends to be treated as a collection of individual, separate instruments. In jazz and some grungier rock styles though, the drum set is treated more as an organic whole and not as the sum of many individual parts – so a super clean, gated drum sound could be inappropriate and undesirable.

If gating is undesirable, less aggressive *expansion* can be used to reduce the level of the spill rather than eliminate it completely – however, ensuring good sounding spill is captured, or minimizing undesirable or bad sounding spill acoustically through mic technique and instrument location in a room is generally preferred to electronic solutions by most engineers working on acoustic music projects.

Good sounding spill is *not* the enemy. *Bad* sounding spill is. Strategies to "improve" the sound of spill include using mics with good sounding off-axis characteristics, directional polar patterns that put sources of spill in their null(s), and to record in a suitable and good sounding room – because the sound of the room *will* be picked up by the mics. Spill can also be managed by changing the position of the sound sources in the room – both the position of the sound sources relative to each other and their mics, and relative to the room itself. For example, a mic facing an instrument but pointing towards a reflective wall is going to pick up more reflected spill than a mic facing an instrument and pointing towards an absorptive wall surface. With the right mics and techniques, good spill can be used creatively to shape the sound and energy of the mix.

AUDIO EXAMPLES
Can be found on the companion website

Spill

Example 9.9: A multi-miked drum set with all of the mics on all of the time. The spill gives the sound a natural "organic" character.
Example 9.10: A multi-miked drum set with the kick, snare, and tom tom mics gated to remove the significant spill in those mics. The sound is cleaner, tighter, and more focused. For some musical styles it might be considered too "bland."

9.10 Mic Position Practicalities

All of the topics related to microphones and stereo arrays discussed in this and previous chapters have a combined effect on the sound a mic picks up. For example:

- Moving a microphone closer or further away affects proximity effect, distance perspective, zoom factor, details, the direct vs reflected sound ratio, and the mic's phase relationship with any other mics which also pick up the same sound source.
- Moving a microphone left, right, up, or down, affects the on-axis sound picked up, the spill picked up, and the mic's phase relationship with other mics which also pick up the same sound source.
- Changing the angle of a mic affects not only the on-axis sound picked up, but also the off-axis spill picked up.

EQ cannot fix *any* phase related problems, such as the timbral shifts of comb filtering – so it is essential to listen critically to the sound produced by each mic, *and* the combination of all the mics set up in the same room before committing and proceeding with the recording session. If anything isn't right, it should be corrected immediately (through mic choice, and mic and equipment placement).

9.11 Multi-Miking

Move your ears (or a microphone) around a large instrument listening for where different components of the sound are best. Listen (or position the mic) up close. Using multiple microphones, each to pick up a different component of the larger instrument's sound, gives you the ability to blend those constituent elements into the "whole" sound during mixing. Blending multiple mics like this will allow you to position the mics closer than if you were using a single mic, which would have to be positioned more distantly to pick up the entire instrument's sound. These closer mics mean a closer perspective is captured, spill and room sound are reduced, and there is more sound captured meaning less preamp gain is required.

How do you pan multiple close mics?

- Most of the time, when acoustic instruments are multi-miked, with each mic in a different place to pick up discrete component sounds (i.e. *not* a stereo array), the mics should be panned to the same position to form the synthesized point source "whole." Panned widely, these mics will produce an image that is smeared across the stereo soundstage, lacks focus, and has less impact and solidity.
- Two different mics placed next to each other and panned wide, will produce a centered image with very little width – the actual difference in sound between the two mics is relatively small, so it's not an effective way of spreading or creating a wide image. Panned to the same position, two radically different mics can be effectively blended though.
- Multiple different sounding mics, on different guitar cabinet loudspeaker cones, have imaging limitations. Different "identical" loudspeaker cones only usually sound slightly different, and the sonic differences between contrasting mics are also relatively subtle – so when the mics are panned to opposite speakers, a lot of common information produces a smeared image with a limited sense of width.
- Recording multiple "identical" (or slightly different) performances and panning them across the stereo image is a much more effective widening technique.

9.12 Experimentation and Exploration

Mic position *X* might have worked really well on guitar *A*, played by person *B*, in studio *C* when using mic *D*. But if any one of those variables is changed, that mic position may not work as well. Time must be taken to experiment until great sound *is* captured by the microphone(s). In addition to changing mic position, some other things to consider include:

- Changing the instrument: Two different instruments of the same type may sound different, and propagate sound differently, necessitating different mic positions. The new instrument may sound so different that an alternate mic is desirable.
- Changing the recording room or the instrument's position in a recording space will change how the instrument interacts with the room. Finding where the instrument naturally sounds best in the room may necessitate a change in mic position relative to the instrument, and/or a change in the microphone(s) used.
- Changing the mic will change both the on-axis direct sound and off-axis spill picked up, and a slightly different mic position may be necessary to get the best sound from the new mic and instrument combination.

SHOULD I RECORD MANY MICS, WHEN I ONLY NEED ONE OR TWO IN THE MIX?

There are different schools of thought regarding how to choose the best microphone for the sound and the mix. With the high track counts of DAWs, gone are the days of having to experiment with mics and mic positions until you narrow the choice down to one mic and position, and commit to it before pressing the record button.

More Is Just More!

One option is to set up multiple mics, maybe in different positions, audition them to make sure they are each producing potentially useful sound, and to record them all for selection at a later date – deferring decision making. Just because you record them, does not mean you have to use them – and more mics, usually mean more phase problems. Mute and delete features exist for a reason!

However, this approach confronts a novice engineer with a bewildering amount of options when mixing, and it is easy to fall into the trap of wanting to use every track. Even when creatively designing a sound in a mix, rarely are more than two or three mics per sound source needed.

Less Is More!

An alternative option is to commit to just one or two mics. Set up many mics (possibly in many positions), audition them carefully, tweak their position, and select just the best sounding/most appropriate one or two to record. Making these decisions prior to pressing the record button makes mixing easier – you are presented with a less daunting array of better considered options. The more mics you set up, the more physical objects and reflective surfaces are in the way of capturing the sound you're looking for – so do take down any mics you're not using before pressing the record button.

9.13 Practical Tips to Help Set Mic Position

It is time consuming and clumsy for you, the engineer, working alone, to position a mic in the studio, run into the control room to listen to it, then repeatedly run back and forth to make further adjustments. A good way of making this process easier is to route the mic signal to a pair of isolating headphones in the studio. You can then wear these, and monitor the mic's sound as you move it around. Once the mic sounds good in the headphones, it should be double-checked through the monitors in the control room.

Unfortunately for your hearing, to really check the sound a mic is picking up, the performer needs to play at full performance volume. Closed-back isolating headphones used at a sensible level will reduce the risk of hearing damage while you're crawling around a loud drum set or guitar amp, and help you hear the mic's sound while minimizing the natural sound in the room.

If multiple mics are used on a single sound source and the mics are not intended to be a stereo array, it is essential that they be checked for phase and comb filtering problems:

- With each mic panned to the same position, listen to each separately by muting them or soloing them alternately. Make sure they are each contributing good characteristics.
- Then, solo or unmute them all at the same time and listen for any hollowing, thinning, or other phasey negative shifts in tonality.

If tonal or timbral characteristics of the sound disappear when all the mics are on, the mics need repositioning relative to each other.

MULTI-MIKING POLARITY TRICK

When multi-miking a source at different distances (not stereo pairs), polarity reverse one of the mics, and listen for mic positions that produce the least of this ugly difference signal when both mics are on. Then put the polarity back to normal and see if more of the complete sound is retained.

Mono compatibility should also be checked when setting up stereo arrays:

- Listen to the image with the channels panned hard left and right.
- Then listen in mono (either push the mono button, or pan both channels centrally).

The image will collapse to a much narrower mono phantom center – that is expected. But are there are any significant shifts in the timbre of the sound? Do any characteristics of the sound disappear too much? Or does the sound take on any strange, thin characteristics? If so, there are mono compatibility problems that must be addressed by adjusting the mic array or using a different stereo array.

PRACTICAL EXERCISE

- ► Set up a single mic on an instrument. For example, an acoustic guitar, a saxophone, or a kick drum.
- ► Record a minute of sound.
- ► Move the mic a few inches forwards, backwards, up, down, or to the side, and record another example.
- ► Compare the two and identify the specific differences between them. Which one sounds "better," or truer to the original source being recorded?
- ► Move the mic again, and repeat the process.
- ► Try this exercise with different instruments.

10 The Recording Room

In This Chapter:

10.1 Room Sound

A significant and often overlooked weak link in the recording chain is the room in which the recording takes place. A room's *acoustics* (its "sound" and character) are created by the reflections the floor, walls, and ceiling produce:

- Different boundary materials reflect different amounts of sound.
- Different materials reflect and absorb varying amounts of high, mid, and low frequencies.
- Different distances between the boundaries change the character of the room sound.

A room's sound needs to be complementary to the source being recorded, and to the aesthetic goals of the project. Unfortunately, the small square and rectangular rooms commonly found in home and project studios are not ideal for recording or mixing.

Many consumers and gear-heads are excited to buy a new mic or preamp and expect it to radically improve their recordings – but the full value of that new equipment will not be totally realized until it is used in a good sounding room. If a recording room has acoustical problems, great performers with great instruments cannot be recorded well – even by the best mics and preamps. Even with acoustical treatment it isn't possible to turn a small square room into a perfect sounding room – but it is possible to reduce *some* acoustical problems. Ultimately, if you want your recording to sound like it was done in one of the world's best professionally designed and constructed studios, then you have to go and record in one of the world's best professionally designed and constructed studios!

10.2 Live Rooms

A *live* room has *reflective* boundaries – the walls, floor, and ceiling are hard surfaces that reflect sound back into the room. Musicians and performers generally like bright reflective rooms because the reflections provide instant feedback about how they sound. Live rooms:

- Provide natural reflections and reverberant content, giving recordings life, vitality, and excitement.
- Provide a natural reverberant "glue" that gels multiple performers together.

A desirable large room characteristic is *isolation* within the room. This means that the sound from a sound source in one part of the room should not overpower or cause problematic spill issues in the mics set up on another sound source elsewhere in the room. Simple square and rectangular rooms do not have good intra-room isolation. The carefully designed, "irregular" geometry of professional facilities contributes to better intra-room isolation characteristics.

10.3 Dead Rooms

A *dead* room is one in which the boundaries absorb sound rather than reflect it. Absorptive acoustical materials are strategically placed on walls and ceilings. An overly dead room, such as a voice-over booth, is a very unnatural environment and can be uncomfortable to those not familiar with that acoustic. Dead rooms can be challenging to perform in – the room does not naturally supply as much instant feedback as some musicians are used to, so

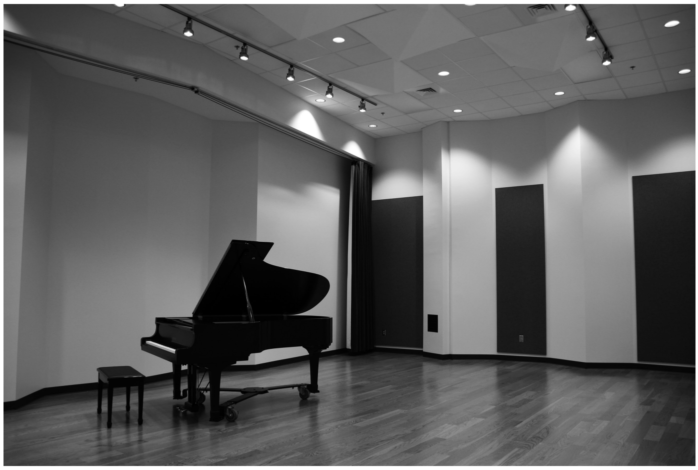

Figure 10.1 A live room with hard, reflective surfaces that reflect the sound back into the room towards the mics and performers (Kansas City Kansas Community College, Kansas City, KS).

they have difficulty relaxing and performing confidently. A good engineer will have a little light reverb dialed into the headphone mix to help the performer and increase their comfort level in a dead space.

Dead rooms:

- Mitigate some of the acoustical problems caused by less than ideal room shapes, sizes, and dimensions.
- Give the engineer more control when adding artificial reverbs and effects electronically.

Spoken word and non-classical vocals are some of the few sources that record better in a dead room.

10.4 Room Size

The size of a room is often an indication of its characteristics and best uses. Generally, the larger the room, the easier it is to control the acoustics and to make it suitable for recording. There are desirable *golden ratios* of length, width, and height for different sized rooms – which are rarely attainable except in new constructions or big-budget remodels.

Small live rooms tend to be the most acoustically problematic. The reflected sound the mics pick up (even when close miking) tends to sound small and compact – because the reflections are temporally very close to the dry, direct sound. Loud sound sources can overload small rooms, resulting in confused and muddy recordings which lack clarity and transparency.

Small rooms are generally treated with acoustical absorption to reduce problematic reflections, however in most cases the treatment needed to control these problems leaves

Figure 10.2 A dead room has absorptive surfaces that minimize reflections (Sky Recording, Kansas City, MO).

the room fairly dull and lifeless – characteristics that then become prevalent in recordings made in the room. Absorption products become less effective at lower frequencies, and in small rooms it's incredibly difficult to control low frequencies – so it's easy to end up with a room that sounds dull because too much high frequency content is absorbed, while boomy low frequency problems are not fixed. On the other hand, a small, dead room works well for vocals and voice recording, as well as other sound sources that are not too loud, and don't have much low frequency content.

Larger live rooms have a more open sound – the reflections add desirable characteristics of depth and space. Depending on the size, geometry, and boundary materials, large live rooms can exhibit either big and bold, but relatively short and tight "room" sounds, *or* more extended reverb tails. A well designed large live room can make instruments like drums sound big and impressive, as well as acoustically glue sound sources together.

Some larger rooms have *adjustable acoustics*, with reversible panels or retractable curtains that can give the room different characteristics for different recording situations.

10.5 Cubic Airspace

Too much sound is problematic in a small room because the sound does not have enough airspace in which to disperse and decay – the frequency balance of the sound is changed, and the reflections in the room become too prevalent in the recording. On the other hand, a room that is large and reflective imposes its longer reverberation characteristics on all the recorded sound sources in the room – even when close miking. Bad, inappropriate, or excess room character is impossible to remove from a recording. Reverb reduction plug-ins certainly exist, but they are far from transparent, and generally leave artifacts that make the processed sound unacceptable for a music project – so it is important to record in rooms that suit the sound sources and the production aesthetic.

The more instruments or voices that are being recorded simultaneously, the larger the room needs to be – not only in terms of square footage, but also height. A room with low ceilings will always sound like a room with low ceilings – there will be a compact character to it due to the timing of the low ceiling reflections, and because the sound has nowhere to go. It may also be impossible to get mics as high as they really should go, or they end up too close to the ceiling. Low ceilings are acceptable for absorbent vocal or amplifier booths, but not for a choir, orchestra, or jazz big-band. Artificial reverb cannot remove the small room or low ceiling reflections already recorded – it can only add to them.

10.6 Standing Waves and Resonant Frequencies

Parallel surfaces are particularly undesirable in recording rooms and control rooms! *Standing waves* occur when a sound wave reflects back over itself. Standing waves can cause varying boosts and dips in the amplitude of specific frequency ranges around the room, as well as booming or ringing noticeably after the actual sound source has silenced. This means that mics pick up incorrect sound, or when mixing, your monitor speakers are unable to tell you the truth so you make inappropriate mix decisions.

Axial modes are one-dimensional standing waves, and the simplest and easiest to predict. They occur between two reflective parallel surfaces – for example, the front and rear walls, or the side walls, or the floor and ceiling. Each of those dimensions is probably different, so in an untreated or poorly treated room there are three sets of potential axial mode problem frequencies.

The lowest frequency axial standing wave for each dimension can be calculated using the following equation where v is the speed of sound, and L the length, width, or height of the room:

$$f = v / 2L$$

This equation calculates the frequency which has a half-cycle wavelength the same as the particular room dimension. So, if the length of a room is 3.6 m (12 ft), the lowest problem frequency along that axis is 47 Hz:

$$344(\mathrm{m/s}) \div 7.2(\mathrm{m}) = 47\ \mathrm{Hz}$$

$$1130(\mathrm{ft/s}) \div 24(\mathrm{ft}) = 47\mathrm{Hz}$$

Other potential standing waves for this dimension will occur at multiples of this frequency – 94 Hz, 141 Hz, 188 Hz, 235 Hz, etc. Similar calculations should also be made for the width

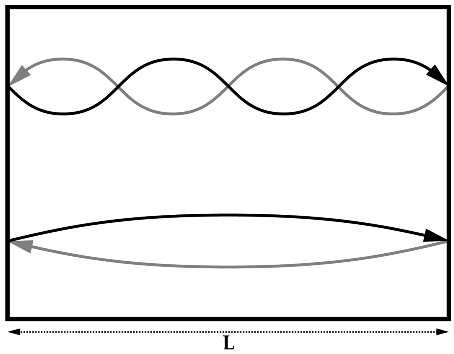

Figure 10.3 Two frequencies a sound source might generate are shown reflecting axially between two parallel walls in a room. The bottom plot represents the lowest frequency that can cause an axial standing wave – a half wavelength. The top plot represents a higher frequency at a multiple of the half wavelength frequency.

and height dimensions if the surfaces are parallel and at all reflective. The biggest axial mode problems are usually the first four or so for each axis.

There will be places in the room where these frequencies either constructively reinforce (boost, creating an *antinode*) or destructively cancel (attenuate, creating a *node*) as they reflect over themselves. A microphone's position in the room will determine whether it is in a node or antinode for a particular frequency. An area of destructive interference for one particular frequency (or set of frequencies) may be an area of constructive interference for a different frequency (or set of frequencies). Moving the microphone to a different position in the room can change the frequency balance dramatically.

Tangential modes occur when a sound wave bounces around a room in a two-dimensional pattern between four surfaces. *Oblique modes* occur in three dimensions between all six (or more) surfaces in a room.

Even though a sound source may be a point source (such as a loudspeaker or a singer's mouth), standing waves are generated throughout the room because sound does not travel like laser beams, and when a sound wave hits a reflective boundary it bounces off over a hemisphere of angles. An infinite number of complex reflection paths is rapidly created, some of which will interact to form standing waves. The higher a standing wave frequency is, the less significant its effect is likely to be. The larger the room, the lower the potential standing wave frequencies are – in some cases approaching the lower limits of audibility.

A room with multiple axes of identical or very similar dimensions is particularly undesirable – square rooms being the worst, because although there are fewer different standing wave frequencies, the problems at those frequencies are greatly increased. A very small cube-shaped room 8 ft (length) × 8 ft (width) × 8 ft (height) will be very difficult to record or mix great sound in, even with extensive acoustical treatment!

10.7 Flutter Echo

Flutter echo is produced when sound reflections between two parallel surfaces are perceived as individual events. A "ringy" or "pingy" character is imposed on all the sound in the room, affecting its frequency content and most noticeably producing an echo after percussive events. Flutter echo can be heard in rooms with dimensions as small as 4.5 m (15 ft) if there are parallel untreated reflective surfaces, *or* if the parallel surfaces at either end of this dimension are not as acoustically treated as the other surfaces in the room. A microphone placed close to a sound source at one end of a 6 m (20 ft) dimension will pick up a reflection approximately 40 ms later, and then at subsequent 40 ms intervals as the sound travels to the opposite wall and back to the microphone repeatedly.

10.8 Microphone Directionality and Room Considerations

If a room is very reflective and live, cardioid and hyper-cardioid microphones used close to the sound source will minimize the room sound picked up. However, directional mics exhibit proximity effect when used close to a sound source and often have muddy sounding off-axis characteristics. Moving the same directional mics further away will reduce proximity effect, but increase the relative level of (potentially poor sounding) off-axis reflected sound picked up.

Omnidirectional mics have similar tonal characteristics regardless of their distance from a sound source, due to their lack of proximity effect. They can produce transparent and spacious recordings at a greater variety of distances than directional mics. They do pick up much more off-axis spill and more room reflections than cardioid or hyper-cardioid mics though. The amount of reflected sound picked up can be decreased by positioning the mic closer to the sound source – but even with an omni mic, the zoom factor of close positioning can impair the sound. A recording that sounds bad because the mic was too close will only sound slightly better even if corrective EQ is possible during mixing.

If a room's sound is not desirable, or bad sounding spill is too prevalent, omnidirectional mics are not a good choice – it is impossible to remove spill and room sound during mixing. In a less than good sounding room, directional mics are usually preferable, even though you have to work harder to find the best mic position and probably end up using corrective EQ.

Bidirectional mics are a great solution when extreme side rejection is required, as long as there is not too much spill or reflected room sound coming from behind the mic.

In order to be able to beneficially use the widest choice of microphones and mic techniques it is imperative to make the recording room suitable for its intended purpose.

10.9 Room Shape

To avoid standing waves and flutter echoes, parallel walls and surfaces should be avoided – but regardless of the room shape or size, most spaces have some acoustical anomalies. The goal of acoustical design or treatment is to reduce these problems so they do not impair the sound in the room. In new construction or an extensive renovation this is easier to do. Even with a low budget it is possible to build additional drywall surfaces into a room to create non-parallel walls, as in **Figure 10.4**. (And with the addition of some doors, any large spaces behind a new angled interior wall can be usefully repurposed as storage!)

If room within room construction or renovation is not possible, various commercially available acoustical products can be applied to existing wall and ceiling surfaces to reduce standing waves, flutter echo, and excess liveness.

10.10 Absorption

Standing waves, flutter echo, and excess reverberation can be controlled by applying *absorption* products to walls and ceilings, as shown previously in **Figure 10.2**. Absorption products are generally foam or mineral fiber based, and applied along a recording room's parallel axes in equal amounts – either symmetrically or in "equal but opposite" patterns. Twenty percent coverage can have a significant effect. *It is not necessary to cover every square inch of the surfaces!* Over-treatment will produce a room that is too dead for most recording purposes, and the recorded sound will be dull, lifeless, and muddy.

Absorption products have an *NRC* rating – *Noise Reduction Coefficient*. This is a measure of how much sound they absorb. Overall NRC is calculated by averaging the absorption coefficients of one-octave bands centered around 250, 500, 1000, and 2000 Hz. Individual NRCs for different octave bands are also usually specified. Values between 0 and 1 (or more) are typical. The higher the NRC, the better the product is at absorbing sound. NRC is not a simple percentage and an NRC of 1 does not indicate complete absorption – values higher than 1 are possible. High frequencies are very easy to absorb and control, while

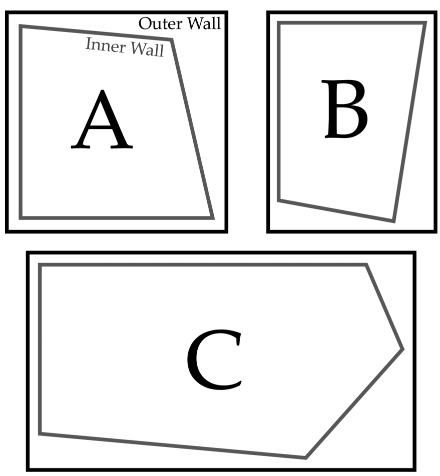

Figure 10.4 Some "room within room" shapes. Rooms **A** and **B** have no parallel inner walls, and only one set of perpendicular walls. Pentagonal shapes can also be used, as shown in room **C**.

lower frequencies are much more difficult. *All absorption products become less effective at lower frequencies*.

To control lower frequencies, the absorption product must be thicker, and be appropriately placed away from a boundary – ideally a quarter wavelength away, which is a realistically impossible distance from a wall for very low frequencies! Inexpensive one inch thick foam, with an overall NRC of about 0.5, is only good at absorbing mid and high frequencies (1 kHz and above) – it has little effect on the low-mid and low frequencies. Two inch thick foam is a commonly purchased broadband absorption product, and has an overall NRC of around 0.8 – but it is only effective above about 500 Hz. This means that a room treated exclusively with this product can still have low and low-mid frequency standing waves and a relative deficiency of high frequency liveness and brightness.

The effective frequency range of broadband absorption products can be extended downwards by positioning them away from the wall (or ceiling) by at least their thickness, using some kind of spacer. However, control of frequencies below 150 Hz is still difficult, because their wavelengths are so large relative to the 1 to 4 inch thick absorption products typically applied for broadband control. *Bass traps* are specific products designed to absorb the longer wavelengths of low frequency energy. They are much bigger and thicker and are usually placed in corners or along axes where multiple boundaries intersect – locations where bass frequencies can be most efficiently absorbed.

Mineral fiber, mineral wool, and modern semi-rigid products offer more effective absorption and better fire retardancy than foam. Some foams do not meet fire code regulations for construction materials. If you intend to use acoustical foam, you should check building codes in your area.

10.11 Diffusion

An alternative to absorption is to scatter reflections so that they don't reflect directly back on themselves to create standing waves or flutter echoes. *Diffusors* randomly disperse reflections. Diffusion retains a room's sense of liveness and open space, and does not drastically change the frequency content of the sound in the room – avoiding the dulling and muddying effects of too much absorption. Commercial diffusion products are readily available, ranging from inexpensive polystyrene products to expensive wooden devices. Or they can be home built.

A mixture of absorption and diffusion is typically applied to many medium and large rooms – absorbing excess reflections while randomly scattering a controlled amount back into the room. In some studio designs, certain walls are treated exclusively with diffusion and others with absorption. In other designs a mixture of absorption and diffusion is applied to each surface. Hiring an acoustical consultant and doing some room-within-room construction is the best way to get great sounding recording rooms. The materials for DIY custom acoustical treatments cost less than buying commercial wall mounted products. But if an acoustician or even simple construction are not realistic options, most reputable acoustical treatment manufacturers offer consultation services to help customers make informed purchasing and treatment decisions.

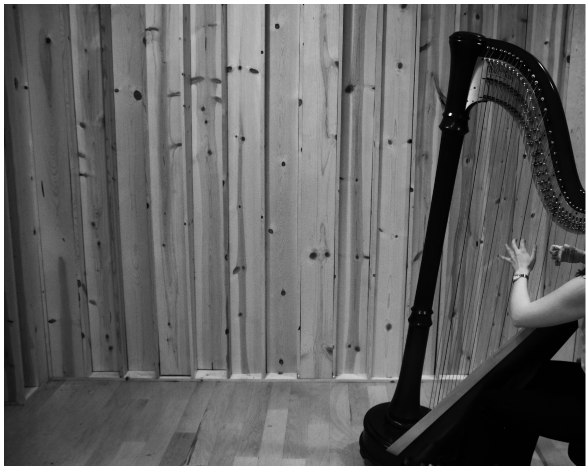

Figure 10.5 A home built diffusive wall randomly scatters reflections, reducing standing waves while opening up the sound of the room, and keeping it relatively bright.

10.12 The Purpose of the Room

A project's sound sources, musical style, and production aesthetic dictate a room's desirable acoustical characteristics. This is why some facilities have only acoustically dead vocal booths, others semi-live large ensemble spaces, and others a mixture of live reverberant drum rooms, semi-live, and dead rooms.

- Drums benefit from being recorded in fairly live, medium to large rooms. Without the reflections of those sized rooms the recorded drum sounds can be dead and lifeless.
- Acoustic guitars and most other acoustic instruments really benefit from being recorded in at least a slightly reflective room – the life is sucked out of them if the room is too dead. Conversely, a recording made in a room that is too reverberant will have too much potentially inappropriate room sound, and lack clarity and intelligibility.
- Pop vocals and spoken word are usually best recorded in very dry rooms to maximize intelligibility and provide a clean, intimate, close perspective. The engineer then has the ability to add artificial acoustics and effects during mixing. It *is* appropriate to cover a high percentage of the surface area of a vocal booth with absorption.
- Rooms can be too live, even though they do not have parallel surfaces and standing wave problems. Minimal absorption can be applied in order to tame the excess liveness. In medium to large rooms sometimes just one or two surfaces will be treated. In smaller rooms, partial but equal treatment of all surfaces usually produces better results.
- A room that is "too bright" generally has a frequency imbalance in its reflection characteristic. Thin walls reflect higher frequencies while letting lower frequencies pass through them and escape. Thinner absorptive treatments can be added to "darken" a room's sound.

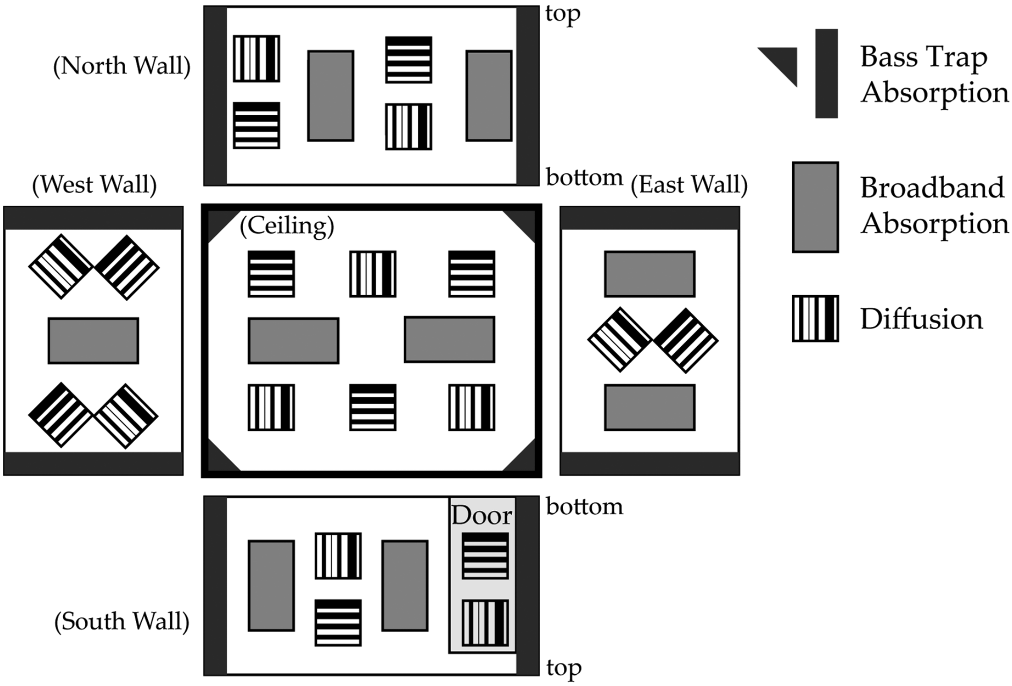

Figure 10.6 A generic treatment of a small to medium sized rectangular room. Standing waves and flutter echoes are reduced using a combination of absorption and diffusion products, and bass traps are used to tighten up the low frequency sound of the room.

► A room that is "too dark" or "too dull" sounding does not have enough high frequency reflections – probably due to too much absorptive treatment in the room. Diffusion can be added to brighten up a dull sounding room without adding standing wave or flutter echo potential.

AUDIO EXAMPLES
Can be found on the companion website

Recording Room Acoustics

Example 10.1: A drum set recorded in a larger live room is exciting!
Example 10.2: The same drum set recorded in a smaller deader room is dull and lifeless.
Example 10.3: A pop vocal recorded in a smaller deader room is controllable, mixable, and in-your-face.
Example 10.4: The same vocal recorded in a larger live room cannot escape the reverb of the room. This "dry" track is less clear, and less punchy.

10.13 The Single Room Home Studio

Single room home studios are more popular than ever. Hopefully the contents of this book have made it clear that for truly world-class sound, there is no substitute for a professionally designed room of a size and shape optimized for the specific sound source and the aesthetic goals of the recording project. But, it is certainly possible to improve on untreated home room acoustics, and make a standard home bedroom or office space acceptable for amateur and semi-professional use.

PRACTICE, PRACTICE, PRACTICE…

If you're learning to record and mix, it is essential to actually record and mix! So it's better to practice as much as possible, even in a less-than-acoustically-perfect space, than to not practice.

A little time and money can drastically improve room acoustics in the middle and upper frequencies. Be realistic – low frequency issues are much more difficult and expensive to fix, and impossible to fix completely in smaller rooms. If all that happens in your home studio is mixing, then the acoustical treatment and layout can be optimized for that. But if recording is also likely, there should be an area of the room optimized for that – versatility needs to be built into the room to make it work for the user.

One end of a room can be treated to optimize the listening position for mixing. This can be done by putting absorptive panels at the mirror/reflection points on the side walls and ceiling, between the loudspeakers and listening position. Bass traps should be placed in the corners. Absorption behind the loudspeakers is also a good idea. The wall treatment should be centered around ear-level – there is no need to treat the lower or upper extremes of the walls, as reflections from those heights are less problematic than direct reflections around loudspeaker and ear level.

The back half of the room can be treated with absorption if a deader acoustic is desirable for vocal recording. For a more versatile room, an alternating mixture of absorption and diffusion will generally produce a less dull and less "small" sound, better for instrument recording.

A possible approach to treating one room with the versatility to be a mixing room, and also offer dead areas for vocal recording, and a less dead area for instrument recording is shown in **Figure 10.7**.

- The front is treated with absorption and bass traps for mixing, as described above.
- The rear corners are treated with a high percentage of absorption, with the treatments going high enough to go above a singer's head. Absorption on the ceiling between the singer and the vocal mic will control ceiling reflections. A vocal mic pointing into the corner of the room will be aimed into a dead acoustic. Reflections from the middle of

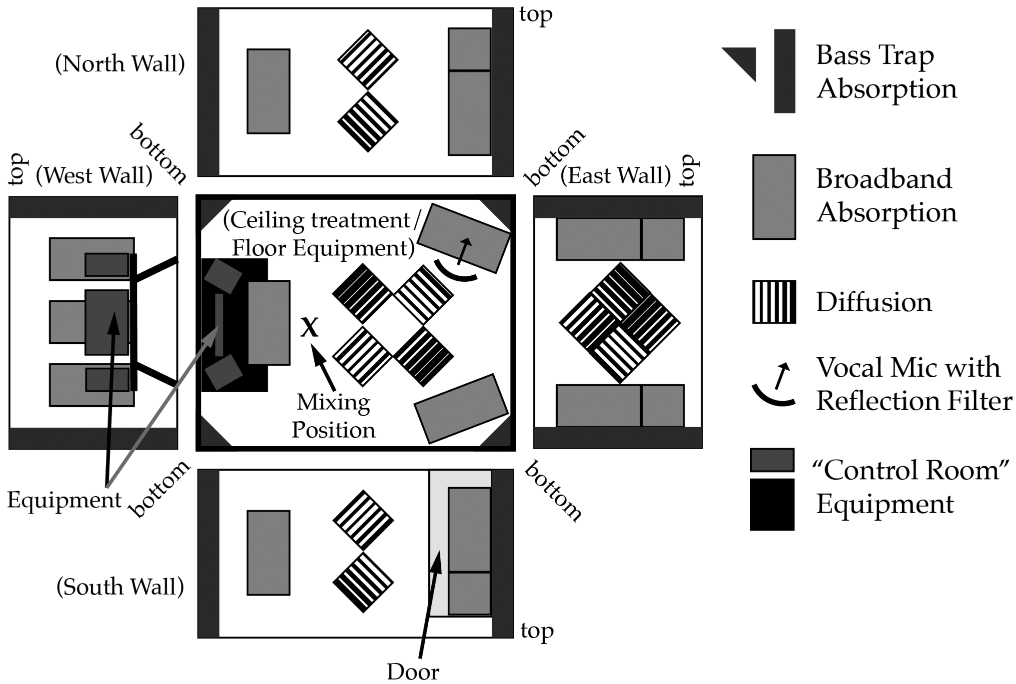

Figure 10.7 An example of a single room home studio, acoustically treated to give three separate areas – in the front for *mixing*, in the rear corners for *vocal recording*, and behind the mix area generally for *generic instrument recording*. Additionally, bass traps are used in each corner.

the room into the rear of the vocal mic can be reduced by using a foam reflection filter (or even a comforter draped over a boom stand) behind the vocal mic.

- The center of the back wall can be treated with diffusion (either custom products or things like shelves with lots of odd shaped things on them) to break up reflections between the front and rear walls. If budget and room size allow, diffusion in the middle of the side walls and ceiling can open up that space and help it sound less "roomy" and small – in combination with the frontal and rear corner absorption it should be neither too dead, or too live.

10.14 Acoustical "Home Remedies"

Let's get a few myths out of the way!

- Egg boxes are not a good absorption product. They are effective over only a very narrow mid frequency range, above which they act as diffusion. Plus, they are not at all fire retardant!

- Domestic curtains and drapes are relatively thin, providing only high frequency absorption, therefore muddying the sound of the room. They are ineffective. If you must use curtains in a recording or mixing room, get the thickest, heaviest ones you can. They will be more effective over a wider frequency range, although no curtains will effectively absorb low frequencies.
- Carpet is relatively thin – so lining walls with it is inappropriate and will darken and muddy the room sound by absorbing only high frequency content.

So, what "on hand," cheaper, temporary home remedies can reduce some of the more significant and common acoustical problems found in smaller rooms?

- Egg crate-style foam mattress pad material can be fixed to walls as absorption – however, commercial absorption products look better, perform better, and reputable products are more fire retardant. It is possible to cut ugly pink or yellow mattress pads into smaller squares, and cover them with nice looking material – but this could make them less fire retardant, so isn't a solution for commercial use. Open-cell foam is preferred for acoustical absorption because of its more even absorption characteristics throughout its effective frequency range. Closed-cell foam should be avoided because its frequency absorption characteristics are uneven. It's definitely advisable to spend the extra money on a good quality commercial absorption product, or make absorbers from wood framed mineral wool wrapped with fire retardant material.
- Cold weather comforters and duvets are relatively thick and heavy, so they can be used as temporary absorption. They can be hung over extended mic stands to form enclosures around instruments or a singer and their microphone. They can also be effective draped over an open grand piano lid to separate the piano sound from the other sources in the room – including a pianist/vocalist who wants to sing at the same time as playing the piano. They can be hung on walls as broadband absorption, reducing standing waves, flutter echoes, and live room sound. Cheap thin, lightweight products will not have much of an effect on low or low-mid frequencies. Low frequency absorption can be increased by spacing the comforter a few inches or more from the wall.
- A sofa or a couple of soft comfortable lounging chairs with thick heavy cushions positioned in the corners or around the edge of the room are effective bass traps.

With some tools and a little know-how, you can make your own acoustical treatments:

- Broadband absorbers can be inexpensively made from construction products like mineral wool insulation panels. One approach is to use 5 cm (2 in) or preferably 10 cm (4 in) thick mineral wool, wood framed around the edges to form 60 cm by 120 cm (2 ft by 4 ft) panels, with the mineral wool material kept in place with chicken wire. (Semi-rigid panels may not need the chicken wire.) The whole thing is then wrapped in an

Figure 10.8 The guitar speaker cabinet and its mic (in the right of the picture) can be placed inside a "cage" of homemade gobos (in the left of the picture) to isolate it from the room and reduce its spill into any mics set up on other sound sources in the room.

acoustically transparent material (such as burlap) in order to keep the mineral wool fibers from escaping and becoming an irritant or dust. Those panels can be wall mounted. Spacing them away from the wall by at least the same distance as their thickness will make them more effective at lower frequencies.

- The thicker and bigger a bass trap is, the more effective it is. One common design involves cutting 60 cm by 60 cm (2 ft by 2 ft) panels of mineral wool diagonally in half to form triangles, lightly gluing these triangles together to form columns, covering them with material so they look nice (and to keep potentially irritating fibers from becoming airborne), and stacking them floor to ceiling in the corners of the room. Less dense, less rigid materials may start to squash under the weight of the column, so lightweight shelves, or a framing system can be built to stop this happening.
- Another common bass trap design involves making cylinders of mineral wool which are covered with material. These bass traps are then placed in the corners of the room.
- A *gobo*, or *sound barrier* can be easily constructed from a wood plank frame and a plywood backing, the plywood covered with absorption, and the whole thing wrapped in material. Put on casters, these moveable barriers can be placed between sound sources to either isolate them from each other, or from the room. Treating one side of a gobo with thick absorption, and the other with diffusion provides various options of sound control.
- Commercial mic stand mounted absorption can be emulated by positioning any homemade absorption behind a microphone to reduce reflections into the back of the mic from a space, or an untreated wall, or control room window behind the mic.
- An internet search for "DIY broadband absorption" or "homemade bass traps" will provide lots of information.

10.15 Monitor Calibration Software?

Several companies now make *DSP* (*Digital Signal Processing*) software that can optimize and calibrate the sound coming from monitor speakers for the room in which they are placed. Some products are proprietary to the loudspeaker company and only work with their own products, while others are more universal and work with any loudspeakers. Most systems analyze the effect the room has on the loudspeaker sound by generating pink noise or sweep tones through the monitor speakers, which is picked up by an analysis microphone positioned in or around the listening position. The simplest systems analyze only which frequency ranges are too loud or too quiet at the listening position. Other more advanced systems also analyze relative phase and time arrival at the measurement position. This allows them to attempt to compensate for things like speaker driver misalignment, and asymmetrical speaker or listening position placement.

For the home user these systems can be beneficial. They can improve and flatten an uneven frequency response caused by monitor speaker and room interactions. They can solidify stereo images, and make them more accurate. However DSP solutions do have limitations, and the nature of the acoustical problems in the room will determine the success of these systems. For example:

- The systems may be able to improve the frequency response of less than “flat” (i.e. bad) speakers.
- Amplitude problems caused by constructive interference of standing waves at the listening position can be attenuated – the DSP system turns down those frequency ranges. However standing waves bouncing between surfaces increase the duration of those problem frequencies, so severe problems can still “ring on” coloring the sound, just at a lower amplitude.
- Amplitude problems caused by destructive standing wave interference at the listening position are less effectively controlled by DSP. Even if the system turns up those frequency ranges, there will still be cancellation at the listening position, and possibly decreased temporal accuracy caused by the now louder standing waves ringing.
- Non-proprietary systems can have a habit of wanting to correct and compensate for normal dips in frequency response around a loudspeakers crossover point(s). Even if the DSP turns this frequency range up, the driver interaction will still cause destructive interference that is impossible to remove.

For these reasons it is essential to acoustically treat your room to fix major problems first, and use DSP systems only to further refine the sound of your speakers in the room. Don’t use a DSP system as your primary means of improvement – they cannot magically make an untreated room perfect, but they can certainly help *reduce* (not eliminate) *some* (not all) problems.

The better the room sounds in the first place, the better the job the DSP system can do. It is always preferable to design and treat a room to be acoustically perfect – however that is not realistic given most home studio budgets!

PRACTICAL EXERCISE

- Set up a vocal or instrument microphone in your usual recording room.
- Record examples of the same sound source using different acoustical treatments behind the sound source, behind the mic, and to the sides of the mic and sound source.
- In a fairly live or untreated room, try absorption products.
- In a well-treated or acoustically dry room, try diffusion products.
- If the room has a hardwood, tile, or stone floor, try recording with some carpet or a rug on the floor, and then with it removed.
- Listen and evaluate the room sound in these recordings. Listen for complex reverberation characteristics (or the lack of them), and phasey, thin sounds produced by flat reflective surfaces if the rear of the mic is close to a wall or window.

11 Recording Vocals

In This Chapter:

11.1 Is it Really About the Mic?

In a book about microphones and mic techniques, I have to start this chapter by saying that mics and recording technology are secondary to *the performance*, and that a good portion of this chapter will not focus on mics or mic techniques! A technically better vocal sound, captured on a super-expensive bling-sputtered large diaphragm condenser mic will *not* sell more records, downloads or streams than a killer vocal performance recorded with a dynamic mic costing the same as two dinners! But as an audio professional, it is important to aim for the highest technical standards.

The vocal is the most important element of many musical styles – particularly pop and commercial music. It is the vocal performance that listeners identify as "the song," and is often the element you spend the most time getting right when tracking and mixing.

IT'S *ALL* ABOUT THE PERFORMANCE!

Having a large selection of legendary (and often expensive) studio vocal mics is irrelevant if the singer is only used to performing live, and has to move and sing like they do on stage to feel "in the zone" and deliver their best performance. *The performance* is more important than the sound quality difference between a handheld mic and large diaphragm condenser mic. If the artist really can't deliver their best performance without a mic in their hand, *give them a handheld mic!* It could be a prop, with a studio vocal mic set up just behind it, or it may end up being the mic you record and use in the mix.

11.2 Getting "the" Performance

Many bands and singers are used to doing live performances where any little mistake is over as soon as it's happened – to hopefully be corrected at the next performance. The recording studio can be an unfamiliar environment – everything about their performance is put under a microscope. Little flaws are magnified (the sound, tone, and timbre; inflections and phrasings; wrong notes, wrong words, etc.) and any blemishes recorded will exist for the life of the recording. This puts the performer under a great deal of pressure, and they have to be prepared to work to a level of precision and perfection that they may not be accustomed to, or initially comfortable with, sometimes for an extended time.

To have a good session and deliver *"the"* performance, a singer needs to be:

- 100% comfortable with the material they're recording. They must know the song top to bottom, inside out, and upside down!
- In good health, and feeling good.
- Having a good day.
- In the mood to focus for an extended period.
- Able to respond to constructive criticism – leaving the ego at home!

To get a singer relaxed and comfortable, have on hand any supplies they might need – drinks, lozenges to suck on between takes, etc. Avoid sugary drinks – focus on lemon, ginger, and herbal teas instead. After doing a few vocal sessions, any engineer, producer, or studio owner should have a list of items to provide for all subsequent vocal sessions.

The recording room also needs to create the right ambience and atmosphere. It should be clean and professional looking – generic, but warm and welcoming. Learn your performer's quirks, likes, and dislikes. If they like to be surrounded by candles and low lighting, make that happen. Anticipating and providing for a client's every need is your job. It makes you look good, assures you future custom, and results in good word of mouth advertising from satisfied clients.

As an engineer or producer, your interactions with the performer can set them at ease and promote a great performance, or conversely, keep the performer from being able to deliver *"the"* performance. With the exception of legendary professionals who can deliver amazing performances regardless of the situation around them, most artists perform better if they feel comfortable with and trust the engineer and producer.

TALKBACK…

You will be communicating with the artist via the talkback system in the studio. Be careful and aware of what is being transmitted through the talkback system – and what is not! You have a performer in the studio who is in an unnatural and uncomfortable place, where everything they do is scrutinized – it can make some people very insecure. If you are telling a joke to someone in the control room, make sure the performer hears it too, otherwise they may look into the control room, see people laughing, and assume that they are the joke.

You should be a coach – the most important use of the talkback system is for you to respond appropriately to what the performer is saying (and what they are not saying), and to get them to deliver their best performance.

KNOW YOUR MUSIC THEORY!

It is essential that music engineers understand basic music theory, and the sound sources being recorded – otherwise it is impossible to communicate musically and coach performers towards delivering a magical performance. Experienced performers have favorite engineers and producers because they have, over the years, developed a bond, trust, and musical understanding, and they work well together.

BE A PROFESSIONAL!

As a home or project studio owner/engineer, or an employee in a commercial studio, you should:

- Have a personality, demeanor, and be presented in a way that is likeable, polite, and positive.
- Be clearly and easily understood.
- Communicate well, and speak the artist's vocabulary – musical and technical.
- Be able to quickly establish professional trust and a close working relationship with the artist.

- Quickly establish what role the client wants you to fulfill. Are you simply the engineer, capturing their performance to the best technical standards? Or is the client receptive to feedback and coaching – *producing* – from you, in order to get them to deliver the performance the song needs?
- Know the sound sources you are recording. How do they sound naturally? How does a good recording of that particular voice or instrument sound? What microphones and techniques can be used to obtain those goals?
- Be able to talk music, instruments, and gear. Know scales, chords, and rhythm values – it will help when discussing the musicality of parts and performances. Know how instruments and amplifiers work, so that if something doesn't sound right, or isn't working properly, you can offer helpful suggestions to remedy the problem.
- Be a psychologist! If things aren't going well, what can you say or do in order to improve the physical, psychological, or emotional state of the artists, and help them deliver the performance of their life?
- Be positive! Ultimately, if something about the performance is not quite right, someone in the control room (the producer or engineer) should be able to identify what's not right and encourage the performer to try something slightly different to make it right. A simple "It sounds great! But let's try the last phrase again, I think you can relax and really go for the end of the line more" promotes a much better performance, vibe, and relationship between both sides of the soundproof glass than, "OK, let's try that again because you screwed up and were flat and late on the last note."

11.3 Vocal Tracking Methods

When using analog tape, the number of tracks is more limited than on modern DAWs. It used to be common practice to record three or four vocal tracks, either in their entirety or with a punch in or two, and then *comp* or *compile* a good vocal track from the best parts of each of those tracks to a spare tape track. The original vocal tracks could then be recorded over. When using a modern DAW, it's possible to use several processes:

- Comp a good vocal track from several complete takes.
- Record a complete take, and then punch in fixes where necessary.
- Work on the vocal track more linearly, top to bottom, concentrating on getting each phrase right before moving on to the next, punching in and re-recording on the same track as many times as necessary.

Punching in on analog tape is a destructive process – there is no Command-Z to undo a punch that is too early and accidentally erased the end of the previous phrase! Tape stretches and wears out the more it is stopped, rewound, stopped, played, and cycled – so re-recording or punching in a particular phrase too many times is not good for the tape or its sound. This type of wear and tear doesn't exist in DAWs, and all have an "undo" feature in case of an incorrect punch or edit.

Some pros and cons of recording a handful of complete vocal takes, and comping the best parts together to form the vocal performance include:

+ Continuity is generally better – phrasing, performance intensity, and style are more consistent throughout the longer takes.
+ The process is more comfortable for a performer who is used to singing the song top to bottom, and not jumping in somewhere in the middle.
– There may still be parts of the performance that are less than perfect, despite recording multiple versions.

Some pros and cons of building up a vocal performance on one track, and punching in, making as many edits as necessary include:

+ A perfect performance is possible on one track. This means that only one set of processing plug-ins, or one mixer channel is necessary during mixing, and the mix process is more streamlined.
– Depending on the singer, there may be less continuity of phrasing, intensity, and performance style.
– There may be more timbre or tonal differences between different punches, created by inconsistencies in the distance between the singer and the mic.
– If the singer does not know the song intimately enough and cannot deliver convincing performances of isolated smaller phrases, punching may be a slow process.

A good approach for most home and project studio situations is a hybrid, somewhere in-between the two tracking methods described above:

- Record multiple takes, top to bottom, or in large sections.
- Identify which take has the best performance of each part of the song.
- Identify any sections or phrases that are still not ideal, and punch in to replace those on an additional track as necessary. Use plenty of pre-roll and post-roll, so that the singer is confident of their entry, and so that you can spot continuity differences, and fix them immediately.
- Comp the final performance together from these tracks.
- Listen for clicks and pops at every edit point, and crossfade to remove them.

Whichever method you use, have a copy of the lyrics on hand so you (or an assistant) can mark off which parts of each take are good, as you are recording. This way, you instantly know which sections you still have to work on. Even if you are building a track a small section at a time, having the singer warm up with a couple of beginning to end run throughs will help them feel comfortable.

IS IT REALLY ABOUT PERFECTION?

QUESTION: What's more important to the success of a song? A) The emotion, energy, and overall performance style? Or B) the fact that one note was a little off here or there?

ANSWER: A! The overall musicality, power, and emotion of a performance are more important than perfection. A minor blemish will go unnoticed by most consumers in the context of an otherwise amazing performance. As a last resort, minor mistakes may be corrected using pitch and time editing tools – but those tools leave telltale signs, and over-correction can remove the groove, style, and soul from the performance.

11.4 Miking Individuals

Large diaphragm cardioid condenser microphones are the usual choice for vocal recording because of their big, hyped, detailed sound. They are often mounted upside down on a large boom stand – this way the mic and stand do not block sight lines to any music stand below the mic. Even though large diaphragm mics are more popular choices, small diaphragm mics, which are usually flatter and more transparent, may suit some singers.

A *pop filter* should usually always be used for voice recording when the mic is within 60 cm (2 ft) of the singer. Make sure there is at least 5 cm (2 in) between the singer and the pop filter, *and* the pop filter and the mic – the pop filter needs this distance on both sides to work effectively and dissipate the plosive air currents that cause pops and booms. If a less-plosive singer is at least 30 cm (1 ft) away, and singing across the top of the mic (not directly into it), a pop filter may not be necessary.

If you have access to multiple vocal mics, you should experiment with them to decide which one suits the singer best and produces the most pleasing sound. Just because a certain mic is another engineer's favorite doesn't mean it will sound great on the singer you're recording. For example, a mic that has a presence peak where a singer's sibilance or nasal resonance lies, will rarely be flattering.

Think about mic distance and proximity effect:

- For a close and in-your-face sound, position the mic closer to the singer, 10 to 15 cm (4 to 6 in) away, with a pop filter between the mic and singer.
- For a more natural, transparent sound with more of a sense of depth, position the mic 20 to 30 cm or more (9 to 12 in) away from the singer.

The further away a directional mic is, the less proximity effect there will be, but the more room sound there will be – so listen for warmth and fullness vs clarity and diction in addition to direct sound vs room sound as you move the mic back and forth to find the ideal position.

Don't forget about the more open, honest, and transparent character of omnidirectional mics! Their lack of proximity effect allows the singer to get closer (for an intimate detailed sound) without muddy coloration – but make sure there is still enough distance to allow the pop filter to work. There will be more room sound picked up by an omnidirectional mic, particularly potentially problematic reflections from any window or wall directly in front of the singer – but if the singer is close to the mic, and the mic is at least a 150 cm (5 ft) away from any wall or window, reflections will be minimized relative to the direct sound.

Some vocal mic techniques to try, shown in **Figure 11.1**, include:

1 The most "in-your-face" and brightest sound is produced by close on-axis miking. The mic is typically placed directly in front of the singer, at a distance of 10 to 15 cm (4 to 6 in), with a pop filter midway between the vocalist and microphone.
2 A slightly less bright sound, less prone to the booms and pops of plosive air currents, is produced when the microphone is positioned just above the singer's mouth, facing directly forward. The mic is not directly on-axis to the mouth, hence the slightly less bright sound which may mellow some excess sibilance. This position could overemphasize nasal characteristics.
3 Positioning the mic at, or slightly above eye height, angled down towards the mouth keeps it out of plosive air blasts, but on-axis to the mouth. This technique restores some of the brightness lost in the previous technique. It also encourages the singer to face straight ahead and slightly upwards, possibly improving their sound and performance. It is also easy to set this mic over and out of the way of a music stand. Angling the mic

Figure 11.1 The different vocal mic positions described in the main text.

towards the nose instead of the mouth will change the characteristics of the sound picked up – whether this is beneficial or not depends on the singer!

4 Positioning the mic below the mouth, angled up towards the mouth can sometimes benefit a singer who sounds thin and nasally, or overly sibilant. The sound will be fuller and warmer, but less punchy than other close placements. If a music stand is being used this technique can be difficult to set up.

If the singer needs lyrics or music, whether on paper or a tablet, use a music stand – otherwise you'll hear the rustling of hand held papers, or the singer is going to get tired holding the tablet. The music stand should be positioned high enough that it does not encourage the singer to look down – if they do occasionally look down, their position relative to the microphone is changed and the timbre and amplitude of the recorded voice will fluctuate. The music stand is also a hard reflective surface that causes reflections that the mic may pick up. Covering the stand with a rug or piece of carpet can help minimize these reflections, as can angling the stand more horizontally so the reflections bounce away from the mic rather than back into it.

KEEPING THE (SAME) DISTANCE

Throughout a vocal recording session, the singer is going to move around the room, leave the room to listen in the control room, go on breaks etc. Then they will return to do more recording. If they don't maintain the same distance from the mic, the recorded vocal timbre can change between takes. In a reflective room the recorded acoustic can get "roomier" if they move away from the mic. To prevent this:

- Put a strip of board tape, gaff tape, or any other visual mark on the floor where their toes should go. Then they will always stand the same distance and maintain the same angle to the microphone.
- Measure the distance from the pop shield to the microphone early in the session, or before starting, and periodically and unobtrusively check it's the same (a good time is when the singer is out of the room, taking a break). Singers with studio experience will often feel comfortable when the pop shield is a certain distance from their face, and return to that position naturally. Pop shields are usually attached using flexible goosenecks, so it's easy for them to get knocked and moved, causing the singer to get closer to (or further away from) the mic without knowing it.

Do encourage good mic technique! If a singer is getting much louder on some notes or phrases, have them lean back from the mic slightly. If there are very quiet passages, they can lean into the mic a little. This will even out their recorded levels, making compression and mixing easier. Do watch out for excessive tonal changes related to proximity effect and reflected sound if the mic distance changes – if there are any distracting timbral changes, have them maintain a more constant distance from the mic, adjust your recording levels accordingly, and be prepared to work a little harder during mixing.

AUDIO EXAMPLES
Can be found on the companion website

Solo Vocal Recording

No EQ is used on these examples.

Example 11.1: A vocal recorded with a cardioid mic just a few inches away. This perspective is *very* close.
Example 11.2: The same singer recorded with a cardioid mic about 9 inches away. The sound is close and powerful, but still fairly natural sounding.
Example 11.3: A vocal recorded with a cardioid mic, about 18 inches away. This sound may be too distant for some aesthetics, but favored for others because of its organic neutrality.
Example 11.4: A vocal recorded with an omnidirectional mic, just a few inches from the mic. Notice the lack of proximity effect compared to the close cardioid mic recording.

11.5 Voice and Acoustic Guitar

Some singer/guitarists need to sing and play at the same time – it is difficult or impossible for them to perform well when they try to record the guitar or vocals separately.

Figure 11.2 Left shows vocal and acoustic guitar being recorded simultaneously. Using a cardioid or hyper-cardioid mic on the voice, the guitar is off-axis, below the vocal mic's main pick-up. The guitar spill is subject to the mic's off-axis coloration, and doesn't usually sound good. A cardioid mic (or two) on the guitar, provide a good guitar sound, but vocal spill is colored by their off-axis characteristics.

Bidirectional mics can be used to minimize spill between the voice and guitar, as long as there are no sound sources or undesirable room reflections behind the mics. Because bidirectional mics reject sound directly above and below them, there will be less voice picked up by the guitar mic, and less guitar picked up by the vocal mic, as shown in **Figure 11.2 Center**.

Aiming the mics by their null points will reduce spill even further. Listening carefully for proximity and perspective changes, front/back distance adjustments can also be made to really get the null points aimed at the sources of spill while keeping the desired sound sources on-axis. With a good mic, even if the voice or guitar do end up just a few degrees off-axis, the sound should still be great. In fact, if the vocals are over sibilant or the guitar is too bright, moving the mic slightly off-axis will gently roll off some high frequencies as an alternative to EQ. A null point optimized set up is shown in **Figure 11.2 Right**.

Figure 11.2 **Left:** A singer/guitarist miked with cardioid mics. Each mic picks up significant off-axis spill. **Center:** Using bidirectional mics on-axis to their desired sound source does not maximize rejection of spill. **Right:** The null points are aimed at the sources of spill, while keeping each mic on-axis with their desired source.

11.6 Small Vocal Groups

If you're recording a small group of singers (backing singers or a harmony vocal group for example) you have to decide whether to record each singer individually or the whole group simultaneously. Singers from performing groups are probably used to working together. They are used to being next to each other and hearing each other at the same time – so it makes sense to set them up in the same room simultaneously. Tracking each singer separately provides more isolation (which is possibly desirable), but it does not give the singers the ensemble sound they are used to, and it may be impossible for them to hear their harmonies correctly when isolated from the full ensemble – so they may struggle and lack confidence.

If you decide to record the whole group simultaneously, you have the choice of giving each singer their own microphone, or setting them up around a single mic or a stereo pair. Individual miking will give the sound a close perspective, a pop music "in-your-faceness," and gives you more control during mixing. When individually miking multiple singers in the same room, do adhere to the 3:1 rule. **Figure 11.3** shows how the more directional pick-up pattern of hyper-cardioid mics can be used to isolate each singer more than cardioid mics. Remember that hyper-cardioid mics are slightly sensitive to their rear, so reflections from a glass window or wall might cause more problems than if cardioid mics were used. Stand or wall mounted absorption products positioned behind the mics can reduce these reflections.

Positioning the singers around a single microphone or a stereo pair are also valid approaches, but they are only effective if the singers are naturally well-balanced. The sound will be more unified – less individual, and more of a blended whole. This may require some

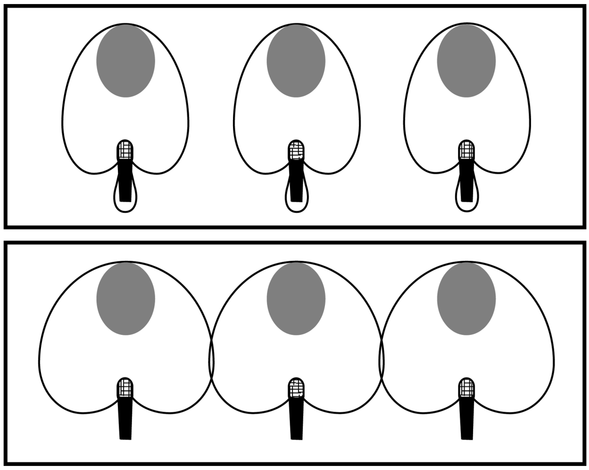

Figure 11.3 The increased directionality of hyper-cardioid mics (**top**) results in less spill from adjacent singers than if cardioid mics were used (**bottom**).

extra time to physically move each singer closer to or further away from the mic until an appropriate balance and blend is achieved. Compression is trickier than on individual mics, so if a punchy, in-your-face, compressed sound is desired, individual mics should be used. Remember that if you mic a group of singers with a single mic, it will only provide a mono point source image, with no width to exploit the stereo soundstage.

JAZZ CHOIRS AND MODERN "A CAPELLA" GROUPS

The solo and small group vocal mic techniques described in the main text work well for most of the vocal parts in jazz choirs and contemporary *a capella* (unaccompanied) groups. Some of these groups have bass vocal parts which imitate bass instruments, and others have vocalists beatboxing. Techniques described as "bad" elsewhere in this book are often used to bring out specific characteristics and help those voices fill their stylized role more effectively:

- Directional mics can be used *very* closely, and even hand held mics can be used, with the bass singer's mouth right at the mic grill for maximum proximity effect and bass boost.
- Dynamic mics traditionally used by radio announcers (because of their extreme plosive protection systems) can be used instead of condenser mics for bass singers and beatboxers, to allow closer use and more proximity effect.
- Hand-held vocal mics can be deliberately cupped and used as close as possible by beat-boxers to bring out the bass and nasal characteristics of the beatboxing.

Super close mic use like this may require additional protection from pops and plosives. Multiple pop filters of different types (metal and nylon) might be used in front of a condenser mic. External foam windscreens can be used on DJ mics and hand-held vocal mics.

To make smaller vocal groups sound bigger, double, triple, and quadruple tracking, and panning each take to different positions in the stereo image is common.

11.7 Larger Contemporary Vocal Groups

Miking options when recording choirs or large vocal groups for non-classical or "pop" projects include:

- Close, individual mics for each singer.
- A stereo pair covering the entire group.
- Multiple, spaced, zone mics.

(An approach to classical choir recording is discussed in a later chapter.)

Individual Mics

Individual miking captures an isolated perspective of each singer, which then needs to be blended into the group sound. Each voice can be processed, shaped, and nuanced individually to form the mix. The engineer is in full control of the soundscape and perspective created. The benefits of individual miking include:

- Each voice can be EQ'd differently to create blend, tonal contrast, or clarity.
- Each voice can be compressed individually for maximum control and impact.
- A close perspective is recorded.
- Individual voices can be panned to create the most effective stereo image for the mix.

Individual, close miking can only be used if:

- There are enough mics, preamps, and inputs!
- There are enough mixer/DAW channels and enough compressors or plug-in DSP available during mixing.
- There is enough space to set up the mics adhering to the 3:1 rule.

A Stereo Pair

Figure 11.4 shows a single stereo pair set up for a group of singers. Any stereo array can be used, for their specific characteristics. The stereo pair should be positioned far enough

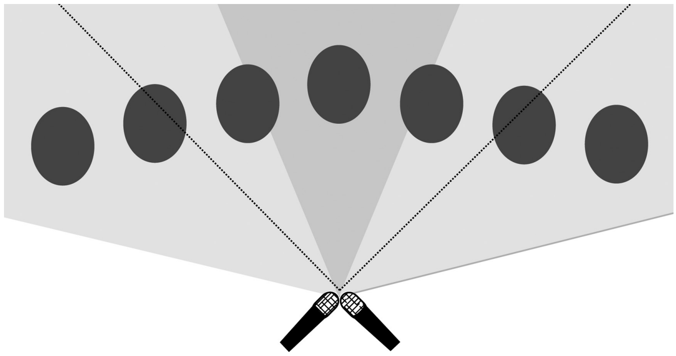

Figure 11.4 A group of singers set up around a stereo pair of mics.

back that its effective pick-up encompasses the entire group. It should be at head height, or slightly above, pointing towards the singer's mouths, or slightly lower if floor reflections are desired.

Characteristics of stereo pair miking include:

- The natural acoustic blend of the singers is captured – they sound like a unified, cohesive group, bonded by the room characteristics.
- EQ can only be applied globally to the whole group. It is not possible to EQ one singer individually (to blend them into the group), without negatively affecting the sound of the other singers.
- Compression can only be applied globally to the entire group. It is not possible to individually compress one singer who is louder than the others. Similar to gentle mix bus compression this global compression can add some "glue" or "gel" to the sound – but it does not allow for the same amount of dynamic control that is typical of pop and rock vocals.
- The stereo image cannot be manipulated much after recording. It should be optimized by altering how the singers are set up and selecting an appropriate stereo array and position prior to recording.
- The balance between direct and reflected/reverberant content is adjusted by positioning the mic array closer to or further away from the singers. This not only affects the amount of reflected sound but also the perception of distance, front/back perspective, and stereo image width.

Zone or Area Mics

For much larger groups of singers, bigger choirs for example, another option is to use multiple mics, each targeting a "zone" of the ensemble, as shown in **Figure 11.5**. The addition of more mics means that they can be positioned closer than a front array, because each mic

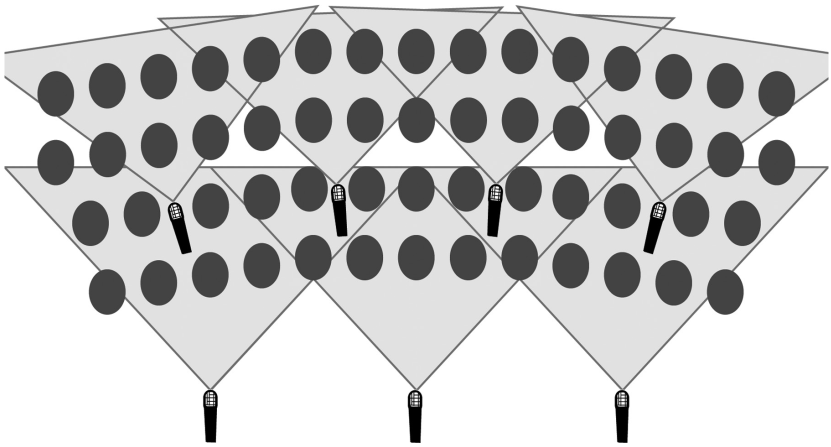

Figure 11.5 Multiple mics on a choir, each focusing on a "zone" or "area" of singers. (The pick-up areas drawn for each mic are just an approximation – it does not mean the mic doesn't pick up beyond those outlines.)

is only required to pick up a small part of the ensemble. This closer multi-miking produces an intimate perspective of a large choir that can be desirable in "pop" or non-classical projects – but is rarely desirable in traditional classical recording styles. The 3:1 rule should be adhered to – you don't want too many mics too close together. You do want the mics far enough away that each picks up a group of singers and doesn't zoom in on just the one or two voices immediately in front of it. The mics should be angled down from at least 30 cm (1 ft) above head height towards the singer's mouths.

AUDIO EXAMPLES
Can be found on the companion website

Vocal Group Recording

Example 11.5: A vocal group recorded with individual mics.
Example 11.6: The same vocal group recorded around a stereo pair of mics.

11.8 Gang Vocals

Frequently used in contemporary pop music styles, gang vocals are different to a choir or chorus because of their unique musical intent. Gang vocals are usually unison, with lots of voices singing or yelling the same musical line to add impact and weight to certain parts of a song. They are not about being precisely in tune or robotically in time – the thickness and character of gang vocals come from tuning and timing imperfections.

The best way to record gang vocals is to get a group of people to perform together. To get a natural gang blend, minimal miking should be used. A stereo array or even a single mono mic can work – just don't put the singers too far to the side of a directional mic, to avoid muddy off-axis sound. The illusion of even bigger gangs can be achieved by recording several takes and panning them throughout the stereo image. If multiple layers of gang vocals are desired, you can:

- Record each take in stereo, panning each one widely.
- Pan three mono takes, L, R, and C.
- Pan two mono takes, L and R, leaving the center of the image for the lead vocal.

Familiarity with different stereo images, panning techniques, and the gang vocal's context in the mix will help you decide what is best for the project.

GANGS AND BVS – STEREO *AND* MONO?

If you have enough tracks, give yourself more panning options by recording group backing vocals and gang vocals using both a stereo array, and a single mono mic. Just because you track something, doesn't mean you have to use it! It's best to have something to discard rather than be wishing you had the extra material during mixing. Just make sure you are disciplined enough to discard tracks that don't bring anything useful to the mix. The mute button or delete feature are frequently responsible for making songs much better and mixes more approachable!

11.9 Vocal Recording Tips

Split Loud and Quiet Sections Over Different Tracks

A vocal performance with a wide dynamic range is a challenge to mix – not only because of the extensive fader movement and automation required to balance the vocals, but also because it is impossible to effectively compress a track that ranges from super-quiet to mega-loud! Vocals are compressed to make the sound more dynamically consistent, with less sudden volume dips and surges, and for the increased fullness, power, phatness, and character that compression imparts on the sound.

- If a compressor is set to compress a loud section of vocals well, it will under-modulate the amplitude of quiet sections of vocal – the character and impact of the sound will be different.

- If a compressor is set to compress a quiet section of vocal well, it will over-modulate the amplitude when the singer gets very loud – squashing it too much, and sucking the life and detail out of it.

With clip-gain features and plug-in automation it is certainly possible to even out sections with very different dynamic levels, automate compressor thresholds, or bypass plug-ins to switch between different compressors to fix these problems. But an easier solution is to record the dynamically different sections of the song to different tracks, each with its own set of EQ and compression plug-ins, or to move the relevant areas to different tracks after recording.

Ride the Gain

With the low noise floor and large dynamic ranges of 24 and 32 bit recording systems, it's tempting to set the vocal mic's gain based on the loudest moment of the singer's entire performance and leave it there. While the quiet parts of the performance will definitely be using enough bit resolution to sound good, it's more difficult to mix and process a performance with widely varying dynamics.

Engineers who work with the smaller dynamic range and limited track count of analog tape are used to having to keep recording levels optimized and more consistent, so that quiet sounds don't get too low and the tape hiss too present. By riding the gain to the recording device while recording you can even out the recorded levels – you are being a manual, finger operated compressor! You can do this regardless of whether you are recording to tape or a DAW:

- As the singer gets quiet for a word, phrase, or section of the song, increase the gain to the recording device.
- As the singer builds up to a loud peak, reduce the gain to the recording device.

Have a copy of the lyrics next to you as you record. Get to know the song and the singer's performance tendencies during a practice run through or two. Be able to anticipate their dynamics. Mark on the lyric sheet when the singer is getting louder than average and when they are getting quieter than average. Then during subsequent takes ride the signal level up during the quiet moments and down during the louder moments. Pay particular attention to "swallowed words" or long sustaining notes, especially at the end of phrases. You should feel involved and connected to the singer and their performance. You will end up with a much more even vocal track, appropriate compression will be easier to achieve throughout the mix, and less volume automation will be necessary.

This is easiest if you're using a mixing console as the front-end into your recording system. It should be possible to feed the recording device post-fader – in which case a fader can be used to push and pull the recording levels. Non-stepped mic gain controls can be used to do this, but they usually offer less precision and it's easy to over-adjust sensitive gain knobs.

Be careful though, particularly before you become experienced at using this technique – inappropriate gain adjustments are not easy to correct and compensate for after the fact. And a compressor or expander won't magically fix them for you!

Double Tracking and Harmony Layers

For thicker, more modern lead vocal sounds in pop and rock music, consider double tracking them:

- Entire vocal lines can be double (or triple etc.) tracked to make the vocal bigger and thicker.
- Double tracking the choruses, but not the verses, will give the chorus emphasis.
- Double tracking the verses, then adding harmony layers on the choruses will make the choruses bigger.
- In rap and hip hop, selective double tracking emphasizes single words or short phrases.

This is easiest when the material is rhythmically regular, and the singer can repeat their performance accurately. Even if the artist thinks they are repeating their performance exactly, there will be natural subtle differences between each take – and some differences are necessary for the thickening effect. All the layers need to be rhythmically together, and close in pitch. But:

- If they are too different in rhythm and pitch, the result is a sloppy mess!
- If they are completely identical in rhythm and pitch, they blend together and the effect is reduced.

In DAWs there are manual and automated ways to edit and align performances to follow a lead track – but the tracks need to be at least "close" in pitch and time. Tight rhythms between tracks increase groove. Identical pitches between tracks can cause the sounds to blend together too much, and the thickening effect is reduced. If you're editing tracks, you don't need 100 percent correction – you can generally tighten the rhythm more than the pitch (make the pitches "close" but not identical), and still retain the thickening effects.

If vocals are double, triple (or more) tracked, in unison or harmony, various stereo imaging options are available:

- Panning all the vocal takes center produces the most powerful sound.
- Panning them extensively across the stereo soundstage produces an expansive and interesting image – but it's not quite as hard-hitting or powerful as the center panned image.
- Panning them around the center, but not too wide, creates an image that retains some power and keeps the vocal image more center focused.

The context of the vocals in the song, and potentially the meaning of the words can dictate where double tracking and harmony layers might work best.

AUDIO EXAMPLES
Can be found on the companion website

Double Tracking Vocals

Example 11.7: Multi-tracked vocals all panned center produce a powerful solid sound.
Example 11.8: Multi-tracked vocals panned widely produce a less powerful but more expansive stereo image.

Compression While Tracking

It used to be common for engineers to gently compress while tracking to analog tape, to avoid over-saturating the tape during loud moments and keep the levels from getting too low relative to the tape hiss during the quiet moments. The dynamic ranges offered by today's high resolution digital recording systems make this technically unnecessary most of the time, but a little gentle analog compression while recording means the A to D converters are less likely to clip unexpectedly, and the track is easier to mix. Just a couple of dBs of gain reduction can also impart a desirable compressor's character (rather than perceived compression) onto the track.

If you do try this, make sure the sound of the compression *is* desirable and use just a little – it is *impossible* to undo bad sounding or inappropriate compression. Until you have thoroughly mastered compression it is advisable to track without it when using 24 bit or greater resolutions.

RECORDING THROUGH A PLUG-IN?

By default, most DAWs do not record through their plug-ins. They just play back through them – so simply inserting a compressor on the audio track does not mean it's being recorded that way, and does not prevent A to D converter clipping. Characterful analog compressors are favored for this application in hybrid studios.

Compress Individual Tracks, Not Groups of Tracks

While sending twenty individual vocal tracks to a stereo group, and putting a stereo compressor on the group might be more DSP efficient (or necessary if 20 hardware compressors are not available), it is usually impossible to effectively compress much without undesirable amplitude modulations. A few dBs of "gel" or "glue" compression is all that is effectively possible even if all the performers are fairly consistent.

In an amateur choir there's often one singer who sings loud high notes much louder than the others! (And way out of tune, but that's another issue…!)

- Imagine this choir recorded with just one mic. If this single mic channel is compressed, the loud singer's amplitude will cause the compressor to turn the gain down – and it will turn down *all* of the other singers inappropriately.
- If all of the singers have their own individual mics and each channel is compressed with its own compressor, the loud singer will be pulled down by their own compressor, while the others will not – blending the louder singer into the group.

More dramatic and better blends can be achieved by first compressing each track or channel individually, and then applying stereo bus or group compression.

Don't Over-EQ as You Record

Fundamental frequency balance problems *should* be fixed as you record – particularly problems that reduce available headroom, such as an over-abundance of proximity effect on a close miked source. EQ is a last resort though, and the best solution is to adjust the mic position or mic choice in order to obtain a sound that does not need corrective EQ to sound great. But if all else fails and reducing proximity effect by moving the mic further away, or using a different mic is not an option, use pre-recorder EQ to tweak the lows and low-mids down a little in order to get higher average levels to the recording medium. Do not be too creative at this stage – you do not yet know how the track will fit into the mix. You don't want to make EQ adjustments while tracking that end up inappropriate and require a second round of EQ to undo as you mix. Characterful analog EQ is used to do this in hybrid studios.

Editing Software vs Performance

As an engineer or producer you should never settle for a substandard performance while recording. As previously discussed, "fix it in the mix" is a myth. Good luck with a singer who exclaims "you can just fix it later!" With the exception of very minor corrections (which *can* be invisible), pitch and timing editors leave behind audible artifacts as evidence of their processing when asked to do too much. After you become familiar with specific editors you'll be able to hear them working! The general public may not hear their telltale signs (with the exception of the super-clichéd Cher/T-Pain effect) – but you, and your professional colleagues (who judge everything you do) will hear them! If pitch and rhythm correction software must be used, target the processing to the small segments that need correcting – do not leave the plug-ins running continuously, because some can impart undesirable artifacts when correction isn't needed.

Don't forget that a memorable song is a *musical performance*. A powerful and emotional performance is *not always about perfection*. In the context of a magical performance, a slight flaw will probably go unnoticed by most listeners. If a phrase is re-recorded in order to correct a slight problem, the emotion and power may just not be the same.

RE-RECORD, OR LET IT GO…?

Learning which mistakes can be fixed invisibly by editors and which ones will require re-recording will allow you to produce better vocal tracks:

- If the singer can do another take and repeat the same great performance but be more in tune or better in time, record another take!
- Only if the emotion and power of the performance is lost in subsequent takes, should editing software be used.

Extensive vocal editing takes time – so a "you can fix it later" artist is great for the engineer or studio billing by the hour! (But of course, more expensive for the client!) If the performer is able, it's usually much quicker to do another take! Being too attached to a performance can also make you overcritical – get a trusted third party set of ears in, someone who can listen with a fresh perspective and give you their opinion on what *really* needs editing or processing to fix.

11.10 Vocal EQ Frequencies

The fundamental pitch frequencies of vocals range from about 80 Hz (for the lowest notes a bass can sing) to just over 1 kHz (for the highest notes a soprano can sing). These are the frequencies that identify the pitch of the note. Harmonics above these frequencies produce the timbre of the voice, the sibilance of the "CH," "S," and "T" and "Z" sounds, and the breath and air of the sound. The following frequency information will help you find the correct frequency area when working on vocal tracks – but the specific frequencies and amount of boost or cut necessary will depend upon the singer, the mic, the room characteristics, and the mix context.

- *Bottom, weight*. Below 120 Hz. Usually unnecessary in the voice track, so it can be cut.
- *Boom, muddiness*. 150 to 300 Hz. Too much thickness in this range can be caused by proximity effect, in which case a low or low-mid attenuation will help vocal clarity. Conversely, some singers may need a little boost here to reduce perceived thinness or scratchiness.
- *Boxiness*. 400 to 600 Hz. This range can often be attenuated in order to give the vocal sound more clarity and intelligibility. This range includes some fundamental frequencies

of higher pitched or female vocals, so attenuating too much can make the vocal sound unnaturally thin.

- *Honkiness*. 600 to 800 Hz. Attenuation here can remove a bull-horn, megaphone, AM radio-like character.
- *Definition*, *nasally*. 1 to 2.5 kHz. Our ears are most sensitive to this frequency range because it contains so much important speech intelligibility and diction information. It can be reduced to make a voice less *thin*, or boosted (when lacking) for more *clarity*.
- *Presence*. Around 5 kHz. Can be boosted to give the singer an edge and allow them to cut through the mix – but make sure it does not make sibilance too harsh or brash.
- *Sibilance*. 5 to 10 kHz (different for each singer). *Brightness* and *sizzle*. If a singer's S's and T's stick out too much, use the boost-sweep-cut technique described in a previous chapter to find the center of their sibilance and apply gentle attenuation.
- *Air*. A wide filter above 12 kHz. A boost of a few dBs can give a vocal a sultry, breathy intimacy. If the frequency is too low, it can boost too much sibilance, or undesirable, harsh characteristics below 8 kHz.

12 Drum Miking

In This Chapter:

12.1 What Kind of Sound Does the Project Call For?

In most pop and rock music, the drums are the second most important instrument – second only to the vocals. They provide the beat, the groove, and the foundation over which the other instruments play. Their sound can be natural, or heavily stylized and influenced by microphone, mixing, and recording techniques. Before deciding how to mic a drum set you need to be familiar with the band's sound, and all the instruments being recorded – if you don't know how specific instruments should sound, or how the musicians are trying to use them in the context of the band's sound, how can you appropriately record them? Preferably this research should be done before the band comes into the studio. Listen to other recordings they've made, or check out a gig or rehearsal. If neither of those are possible, and even if they have already been done, you can't really decide what mics and mic techniques to use until you've listened to the instruments in the studio. Every instrument is different, and the sonic goals of each project are different – "this is how I always mic it" doesn't guarantee the best results and is an inexcusably lazy approach to recording.

A drum set can be miked with anything from a single mono microphone, to 20 or more mics! Natural sounds tend to be preferred for jazz and acoustic styles, whereas huge, studio-crafted, power-house sounds provide the basis of rock music.

- Directional mics zoom into individual drums and instruments, producing a more focused sound.
- Omnidirectional mics produce a more transparent blended sound.
- As more mics are used to individually mic drums and cymbals, the closer and more discrete the recorded sound becomes.
- Using fewer mics promotes a more natural acoustic sound.
- Close miking will reveal and magnify rings and resonances our ears don’t notice when listening from a distance.
- Gating can reduce close miking rings and resonances, or it can produce an ultra-clean, surgical sound.
- Compression can do anything from subtly even out varying dynamics in a performance, to totally reshaping the drum sounds into brighter snappier sounds or big sustaining driving forces.

A modern rock or pop drum sound is a studio crafted sound, influenced by recording and mixing techniques so much that you must have a sonic goal in mind before starting the project.

DRUM MIKING PRACTICALITIES

- Make sure mic stands do not touch any part of the drum set.
- Make sure mic stands do not touch other mic stands.
- Make sure mic stands are set up properly – no center columns touching the floor, only the rubber feet or wheels touch the floor, clutches should be tightened but not over-tightened.
- Make sure mic cables are secured so they will not fall and touch things. Cables should always run horizontally along the floor or vertically up stands – no “loose” cables strung diagonally.
- Stress relieve cables so they do not pull on mics – you don’t want to end up with hi-hat or overhead mics pointing at the ceiling half way through your session! Stand clips or gaff tape can be used for this.
- Coil slack neatly (over-under, or figure-8) near or under the mic stands. Mics and mic stands may move, so you need slack there, *not* at the snake box or preamps.
- Make sure the mics are positioned so that they are not in the way of the player. Your job is to capture a performance – the gear cannot interfere with that.

12.2 How Many Mics Do You Really Need?

On an average drum set you need at least five mics to get a contemporary drum sound and have the ability to adjust balances, EQ, compression, and effects:

- A stereo pair of overheads (usually condenser or ribbon) – in front of and over the drummer's head, or directly above the drum set.
- A dynamic mic would be a traditional choice for the kick. (There are also condenser and boundary mics made for kick drums, and some mics that have both dynamic and condenser elements.)
- A dynamic mic on the snare.
- A small diaphragm condenser on the hi-hat.

These five mics allow discrete processing of the most important and most used elements of the drum set – the kick, snare, and hi-hat.

In this situation, the overhead mics also function as tom tom mics. Remember that EQ applied to the overhead mics will also change the tom toms. The overheads pick up the snare and hi-hat – so the sound of those instruments is a combination of their dedicated mics *plus* the overheads – when EQ-ing, make sure you check how the close mics and overheads sound *together*.

Without their own mics, the tom toms will appear more ambient and distant than the snare and kick – so if they are used frequently and more than five mics are available, give the tom toms and any other additional instruments their own mics.

The less a drum set is miked individually, the less it is possible to adjust the balance, EQ, or process individual instruments, while mixing, without negatively affecting the other instruments. When fewer mics are used it is more important that each instrument in the drum set sounds great and that the natural balance of the drum set is perfect.

FEWER THAN FIVE MICS?

If you have a limited number of mics or input channels, here are some options:

Only one mic? Mono is your only option! The mic must be carefully positioned to capture a perfect recorded balance. Try these positions:

- A meter (3 ft) in front of and 60 cm (2 ft) above the drum set – the distance blends the individual sounds together.
- Above the drummer's head, angled towards the center of the drum set – the drummer listens from this position, and plays so the balance is good to them.

Two mics? Try a stereo pair of overheads in either of the above positions. Or try adding a kick drum mic to a single overhead – to add bottom, thump, and low frequency power to the sound. In a relatively live space you could also try two mics 1 m (3 ft) in front of the drum set, 1 m (3 ft) off the ground, spaced 45 cm (1.5 ft) from the center of the kick to form an equilateral triangle (ribbon mics work well for this).

Three mics? Stereo overheads plus a kick drum mic, to record a stereo picture of the drum set with increased kick drum presence.

Four mics? Adding a snare mic will give you balance and perspective control and allow individual processing of the kick and snare (the two loudest and most important elements of a rock drum set). The snare mic will also pick up plenty of hi-hat spill – the relative balance of which can be adjusted by changing the polar pattern, distance, and angle of the snare mic to the hi-hat.

Every drum is different, every drummer plays differently, and different musical and song styles demand different drum sounds. It's important to really listen before tracking, and experiment, in order to get the most desirable sound. There is no such thing as a universal "best mic position" for any type of drum!

AUDIO EXAMPLES
Can be found on the companion website

Drums – Minimal Miking

Example 12.1: Drum set, single condenser mic, in front of and above the drum set.
Example 12.2: Drum set, single condenser mic, above the drummer's head.
Example 12.3: Drum set, stereo pair, above and in front of the drum set.
Example 12.4: Drum set, stereo pair above the set, plus a kick drum mic.

12.3 Kick Drums

A dynamic microphone is typically used on the kick drum because a kick drum is:

- Very loud.
- Does not contain any essential super-high frequency content.
- Can beneficially be made a little bigger and beefier by the transient compression characteristics of a dynamic mic.

Many general purpose instrument mics sound great on kick drums, but there are also low frequency instrument mics specifically designed for kick drums. General purpose instrument mics may not sound as instantly big and powerful, but if the specific characteristics of a low frequency instrument mic don't suit the drum, a flat mic is a better choice – you don't want to use EQ to fight the sound of an already heavily contoured mic.

Rock kick drums usually have a hole in the front to allow a mic to be positioned inside the drum, where it picks up a snappy, punchy sound. Generally:

- The closer the mic is positioned to the center of the beater side drum head, the more snap and click the sound will have.
- The further back from this position the mic is, the beefier and boomier the sound will be.

The mic does not have to be placed directly on-axis with the center of the drumhead, but if it is off center it is usually angled so that it is aiming towards where the beater impacts the head. The mic can go anywhere from 6 or 10 cm (2 or 3 in) from the beater side head, to close to the resonant head (the head with the hole in it). Moving the mic a small distance has a huge effect on the recorded sound.

Avoid positioning microphone capsules right at the sound hole, and be careful positioning mics just outside it. A lot of air moves in and out of that small hole at high velocity, and can produce plosive-like booms and pops in some mics.

Kick drums can also be multi-miked:

- A mic added to the outside of, but close to the front head will pick up a less clicky, boomier sound that can be blended with the mic inside.
- A more distant kick drum mic, 30 cm to 2 m (1 ft to 6 ft) in front of the drum, can be isolated from excessive spill by constructing a tunnel from the kick drum to behind the mic using wide sturdy cylindrical objects, or heavy blankets, duvets, or comforters draped over chairs, mic stands, or other supports.
- Reverse wired loudspeaker "mics" (either home built or commercial products) can be positioned directly in front of the drumhead. These "microphones" pick up a *very* boomy sound that in isolation is not desirable, but can really add weight when blended with a conventional mic inside the kick drum.
- A boundary mic placed inside the drum can also produce a big beefy sound that can be blended with a conventional kick mic.
- A condenser mic, or even a ribbon mic can be used (carefully) outside the kick drum – and blended with a mic inside the drum. *Never position a ribbon mic on-axis or close to a kick drum though – the SPL and air currents can perforate the fragile ribbon. Angle it down at 45°, at least 45 cm (1.5 ft) away.*

AUDIO EXAMPLES
Can be found on the companion website

Kick Drum Miking

Example 12.5: Kick drum, miked with a flat kick drum mic.
Example 12.6: Kick drum, miked with a loudspeaker cone mic, alone.
Example 12.7: Kick drum, the flat mic, and loudspeaker cone mic combined.
Example 12.8: Kick drum, miked with a low frequency instrument mic.

Figure 12.1 **Top Left:** Looking through the hole into a kick drum with a mic typically placed for single mic recording, several inches from the beater-side head. **Top Center:** A mic about 15 cm (6 in) outside the center of a jazz kick drum. Avoid miking at or close to the outside of a hole as small as the one in this drum – a very strong air current is forced through a hole this small. **Top Right:** A combination of a loudspeaker cone mic, and a conventional mic inside the drum. **Bottom:** An additional mic is placed several feet from the kick drum, and comforters, duvets, blankets or whatever is on hand is used to isolate it from excessive spill in a "tunnel." Short mic stands are used here, in "T" formation, to support the blankets.

Low frequency wavelengths are so long that the distance between the inside and outside mics doesn't usually cause comb filtering problems in the beef and body of the sound. However, the click can suffer. To fix any phasing in the click sound you can either:

Change the distance between the mics slightly until the problem is minimized.

OR:

- Measure the distance between the mics.
- Calculate the delay time it creates: 2.9 ms per meter (0.9 ms per foot).
- Either: Add that delay electronically to the mic inside the drum.
- Or in a DAW:
    - Either: Nudge the outside mic audio track forwards so the peaks in the waveform line up to be sample-accurate with the inside mic waveform.
    - Or: Nudge the inside mic audio track back so the peaks in the waveform line up to be sample-accurate with the outside mic waveform.

- *Listen carefully to see which solution interacts best with the snare and overhead mics.* With outside mics more than 30 cm (1 ft) or so in front of the drum, it is advantageous to nudge the outside mic audio forwards, so you don't change the phase relationship between the inside kick and snare mics too much.

NO HOLE?

Jazz kick drums do not usually have a hole in the front. The role and sound of a jazz kick drum are different to a rock drum – they are much warmer, more boomy, less clicky, and more of an accent marker than relentless rhythm maker. A kick drum with no hole will never sound like one with a hole in it, regardless of microphone technique and any processing applied!

For the most defined sound on a kick with no hole, a flatter mic should be positioned on-axis, in the center of the drumhead. Positioned too close, the sound will be very muddy and boomy due to proximity effect – so avoid positioning the mic less than 10 cm (4 in) away. The mic distance can be increased to 30 cm (1 ft) or more for less proximity effect and a more natural sound – but increasing the distance will increase spill.

If a drum is particularly dull and boomy sounding, and has no hole for the mic, a second mic can be positioned close to and pointing towards the beater impact point *on the beater side* of the kick drum – where it will pick up more click and snap. This mic should be polarity reversed during tracking, and the combined sound of both mics carefully auditioned prior to pressing the record button. Keep this mic and stand out of the drummer's way, and position the mic to minimize spill from the bottom of the snare drum. Listen for, and fix any drum pedal squeaks, rattles, and clunks before hitting the record button.

Condenser mics can also be used on kick drums, particularly for lighter styles such as jazz, folk, and other acoustic music. Some manufacturers make condenser mics designed specifically for this purpose. The SPL inside or close to a kick drum can be in excess of 140 dB, exceeding the capabilities of many regular condenser mics, so before using a condenser mic, check the maximum SPL listed on its spec sheet to see what it can and can't handle, engage the pad, and listen carefully for any distortion.

PAD IT!

- The pad switch found on many mics can prevent the mic's electronics, and the preamp, from being overloaded and distorting at high SPLs.
- The pad button found on channel strips or a preamp prevents the preamp from being overloaded and distorting at high SPLs, but not the mic.
- Pads will not prevent the diaphragm itself from over-excursion and distortion – only a lower source SPL will.

A condenser kick drum mic will produce a more detailed sound, with a brighter "snap" than a dynamic mic – but there will probably be more spill in the sound. For jazz music the additional spill may not be a problem (it can add to the organic quality of the drum set sound), but it may not be desirable in rock or pop projects.

12.4 Snare Drums

Dynamic microphones are the go-to mics for snare drums because:

- A snare drum is very loud – especially at a distance of a few inches!
- The mic is directly in harm's way – near fast moving drumsticks which could easily damage or destroy a more fragile mic capsule.
- There is little extreme high frequency content in the sound produced by the top head of a snare drum.
- The transient compression characteristics of a dynamic mic can make the sound bigger.

Make sure you position the mic and its mounting hardware so that it is unlikely to be accidentally struck by the drummer – this generally means that the mic is boomed in from behind the drum, away from the drummer, so it is pointing towards the drummer from the far side of the drum. The mic should be positioned with the capsule several centimeters (an inch or two) above the drum, either over the rim or just inside, angled down towards the center of the drum. The angle of inclination can be varied:

- A steeper angle, with the mic "looking down" towards the center of the bottom head, will pick up more tone and body for a potentially beefier sound.
- A less severe angle, the mic aimed towards where the stick hits the head, will potentially produce more attack.

In jazz, the snare drum is not played as loudly as in rock styles, so a directional condenser mic (with good off-axis pick-up characteristics) is also an option. In addition to the increased detail, the increased spill puts the drum "in the set" more, decreasing its sense of isolation and individuality – usually desirable in traditional jazz recording.

When using a condenser mic it is even more important to have a great sounding drum and a mic which complements its sound. EQ will change not only the sound of the snare drum, but also the sound of all the instruments in the increased spill. For example, you might make the snare sound great, but that same EQ might negatively affect the hi-hat sound when the EQ'd spill is combined with the hi-hat mic itself.

You can also try adding a mic to the side of the drum, near or across from the tone hole, to capture a blend of the top and bottom sound.

DOUBLE MIKING

A single mic above a snare drum tends to produce a dull, boxy sound. The snares that produce the bright noise content of the drum are on the bottom, far from the top mic. Double miking the snare drum by adding a second mic, 10 to 15 cm (4 to 6 in) underneath the drum, pointing towards the center of the snares, will add brightness to the sound.

- A dynamic mic with extended high frequency response, or a condenser mic are both suitable "snare bottom" mics.
- The kick drum is close to the bottom snare mic. By cutting the low frequencies from the bottom snare mic when mixing, some of the kick spill will be reduced.
- Generally, the body and weight will come from the top mic, and a smaller amplitude of the bottom mic will add brightness and snap.

Because the bottom mic is pointing in the opposite direction to the top mic, each mic's perspective of the drumhead movement is opposite – as the heads move downwards they cause the top mic's diaphragm to move outwards, but the bottom mic's diaphragm to move inwards. When the mics are summed together in the mixer or DAW, sound common to both mics (the body and beef of the sound) will cancel out significantly. During recording, the bottom mic should be polarity reversed using the Ø button on the preamp or console.

It is vitally important to check how the mics sum together before pushing record – so that you know you have the snare sound the project needs. You do not want to find out there is a problem during mixing, after the drummer has gone home and all the other musicians have already laid down their tracks to that drum recording!

Double miking can also add body, beef, and sustain to a tom tom – although the difference is less dramatic than the effect a bottom mic has on the snare drum. Adding a beater-side mic to a kick drum may add some snap to rescue a dull boomy kick (but there will also be a lot of snare spill added). The "opposite side" mic on any drum must be polarity reversed.

AUDIO EXAMPLES
Can be found on the companion website

Snare Drum Miking

Example 12.9: Snare drum, top mic only.
Example 12.10: Snare drum, bottom mic only.
Example 12.11: Snare drum, top and bottom mics blended, but the bottom mic is (incorrectly) not polarity reversed.
Example 12.12: Snare drum, top and bottom mics blended, with the bottom (correctly) polarity reversed.

Figure 12.2 A double miked snare drum (dynamic mic on the top, small diaphragm condenser on the bottom) showing the different top mic angles described in the main text. (You would not normally use two identical top mics at different angles simultaneously.) The mic stands are boomed in from behind the snare drum, in-between the tom tom and hi-hat. (The hi-hat is not shown, for picture clarity.)

12.5 Hi-Hats

Small diaphragm condenser microphones are generally preferred for hi-hats because of their accuracy and high frequency response. Brightness and sizzle (high frequency detail) are essential components of a hi-hat sound, and a few inches away, the SPL of a hi-hat is well within a condenser mic's capabilities.

- The mic should be positioned above the top hi-hat cymbal in order to pick up the impact of the sticks on the instrument. The mic should be positioned at least 12 cm (5 in) above the cymbal, and can be up to about 20 cm (8 in) away. Make sure the area struck by the sticks is in the mic's pick-up. The mic can be pointed straight down; or angled in from slightly to the side of where the sticks hit the instrument; or positioned further out of harm's way, towards the outside edge of the instrument, but still over the instrument, angled in towards where the sticks hit it.
- A mic positioned too close to the edge of the hi-hat will produce a brash crunchy sound.
- A mic positioned too near the bell in the center of the cymbal will produce a clangy, ringy sound.

The hi-hat should *not* be miked with the mic to the side or underneath. From the side, pointing towards the edges of the hi-hat, the sound is brash and brittle, and there is a large rush of air as the cymbals close. From underneath, there is a lot of snare drum spill, and a lack of definition.

12.6 Tom Toms

Tom toms can be approached similarly to snare drums, with a microphone 5 to 8 cm (2 to 3 in) above the rim of the top drum head (or slightly over the drum head), angled down into

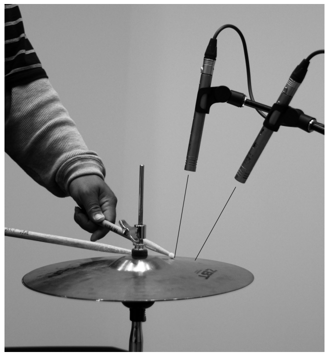

Figure 12.3 Some possible hi-hat mic positions.

Figure 12.4 A rack tom and floor tom miked from the top.

the drum. Generally, a steeper angle produces more tone and body, and a shallower angle increases attack. A condenser mic will pick up the drum's attack transients most accurately, but a good dynamic mic is also a good choice. Low frequency instrument mics can even be used on larger, lower pitched tom toms. Double miking can be applied to tom toms – the body and resonance of the sound can be increased by adding a bottom mic positioned 10 to 15 cm (4 to 6 in) below the bottom head, angled up towards the center of the bottom head.

Positioning a single mic between two tom toms is *not* recommended – both tom toms are off-axis to the mic, and subject to its muddy off-axis coloration. Panning options are also reduced when one mic is used to pick up multiple drums. Also *not* recommended is positioning a bidirectional mic, aimed across the top of the tom toms. This reverses the polarity of one tom tom and much of the drum set in that mic, causing phase problems when combined with the other drum mics.

12.7 Cymbals and Overheads

Condenser mics are preferred as overheads because cymbals contain a lot of bright splashy high frequencies. Small transducer ribbon mics are also a popular choice, because of their smooth high frequency characteristics. The “body” of a cymbal sound is surprisingly low though – down to around 250 Hz, so it’s not just about picking up high frequencies.

- Flat response small diaphragm mics are favored for their extended high frequency accuracy when a faithful recording is desired.
- Very small diaphragm mics are a great option too, producing an extra level of detail and transparency.
- Many large diaphragm mics also make good overhead mics – from flatter more accurate models, to characterful mics which give the overheads some sonic hype, or mellow the brightness a little.

The function of the cymbal or overhead mics depends on how the rest of the drum set is miked. If the tom toms (and any other instruments) are not miked, the overheads are also responsible for picking up those instruments:

- If the tom toms are frequently and dramatically used, there can be a distracting perspective shift from the closely miked kick, snare, and hi-hat, to tom toms which are only picked up by the overhead mics.
- Without tom tom mics it may not be possible to get the tom toms loud or present enough without the cymbals becoming too loud.
- If the overheads are responsible for picking up more than the cymbals, any EQ that benefits the cymbals may negatively affect the tom toms and other instruments.

Any of the stereo arrays discussed earlier in this book can be used as overheads, each with their unique imaging characteristics. Overheads can be positioned anywhere from 40 to 90 cm (1.5 to 3 ft) above the cymbals:

- The closer the mics are, the wider the stereo image – but the mics zoom in and overemphasize the cymbals to which they are closest.
- The further away the mics are, the more blended the picture of the whole set – but the narrower the stereo image, and the more room sound is picked up.

On wider drum sets, spaced pairs allow the microphones to be positioned so they evenly pick up all of the cymbals. AB spaced pairs are commonly set up with omnidirectional mics spaced 40 to 60 cm (15 to 24 in) apart. This doesn’t usually put the mics in an ideal position though – they are too close together, not centered over the left and right banks of cymbals on either side of the drum set. A wider spacing is typical, to center each mic over each bank of cymbals. This wider spacing may cause imaging problems in the form of a hole in the center, and “ghost” images to appear on opposite sides – for example, a crash cymbal struck on the right is heard loudly on the right, but with a distinct, separated, blurry version appearing to the left of the stereo image.

Figure 12.5 A spaced array with a mic positioned over each bank of cymbals, at the first and third "quarter" divisions of the width of the set.

Directional mics can be used instead of omnidirectional mics in spaced overhead arrays – this is particularly desirable in rooms that are not acoustically ideal, have low ceilings, or if there are sources of spill in the room. Directional mics will increase the apparent width of the image, which can be good… or bad, if the left and right become too separated. Widely spaced overheads can be angled into the drum set to put more sound sources on-axis and reduce off-axis coloration. Very large drum sets may also benefit from having more than two overheads – an added center overhead mic, or individual mics positioned closer to specific cymbals.

Be sure to check your overhead image, and push the mono button to check mono compatibility before pushing record!

MEASURE IT!

The snare drum is usually one of the most used, and most important instruments the overheads will pick up – whether the snare is individually miked or not. It is important that the snare drum sound in the overhead mics will blend well with the individual snare mics. We tend to pan dedicated snare drum mics center, or *very* close to center. Rarely is a snare drum actually set up exactly centered in the entire drum set – it is usually at least a little to one side, and not equidistant to both spaced overhead mics. This makes it appear panned off to one side of the overhead mics' image, and makes the stereo image quite asymmetrical. Also, the snare drum sounds less focused and precise when its overhead image (to the side) is combined with the close mic(s) panned center. Snare drum phase and comb filtering problems can also occur when the overhead mics are summed to mono.

One way to anchor and focus the overhead's snare drum sound is to make the distance from the center of the top snare drum head to each spaced overhead mic equidistant. Then the snare drum sound arrives at both overheads at the same time, creating a centered phantom image. The sound is in phase in both overhead mics (increasing mono compatibility), and the overhead's snare drum sound blends better with the close mics.

Figure 12.6 Both overhead mics are positioned equidistantly from the center of the snare drum's top head. If you don't have a measuring tape, you can use a length of mic cable to measure the same distance.

AUDIO EXAMPLES
Can be found on the companion website

Drums – Time Aligning Overhead Mics

Example 12.13: A spaced pair of overhead mics, positioned symmetrically over the entire drum set. The snare drum appears too far to the right in the recording.
Example 12.14: The snare drum image in the overheads is centered and solidified by raising the overhead mic closest to the snare drum, so that both overheads are the same distance from the snare drum.

Coincident, near-coincident, and MS arrays positioned above the center of the drum set can produce stereo images with less of the undesirable ghost imaging spaced arrays can be prone to. The width of the image depends on:

- The array used.
- The angle between the mics.
- The polar pattern of the mics.
- The height of the array.

In many cases these arrays will produce more cohesive and accurate images than a spaced array – but do not use them too close as they may focus on the cymbals nearest to them too much.

If all of the drums are individually miked, the overheads become dedicated cymbal mics. If this is the case:

- Directional mics can be positioned to focus on each cluster of cymbals to the left and right (and maybe in the center).

Figure 12.7 Near-coincident overheads.

- Cymbals can even be miked individually, at least 15 cm (6 in) above the edge of each cymbal, aimed at a point about one-third of the radius from the back edge of the cymbal – where the sticks are not likely to hit the mic. Cymbals move a lot when hit hard, so make sure they won't hit the mics.
- The ride cymbal is usually quieter than crash cymbals, so it can often benefit from its own dedicated mic. Individual cymbal mics can be at a shallow angle positioned above and at the rear edge of the cymbal, or angled more steeply from directly over the cymbal – you have to experiment and see how the cymbals sound, whether the attack transient "tink tink" of the sticks hitting the cymbal is present enough, balanced with there not being too much clangy "bell" sound.
- Individual cymbal mics are not stereo arrays optimized for natural imaging accuracy – instead they should pick up the best individual cymbal sounds, which are then panned to create the stereo image artificially.

Positioning microphones underneath the cymbals is a technique used in live concert sound either for visual aesthetics, or to reduce spill and isolate each cymbal's sound. It removes percussive attack, and picks up off-axis tom tom and snare drum spill, so is avoided for studio recording.

AUDIO EXAMPLES
Can be found on the companion website

Drums – Full Set Miking

Example 12.15: Drum set, five mics: kick (dynamic), snare (dynamic), hi-hat (condenser), two overheads (condenser). The tom toms sound distant and ambient compared to the kick, snare, and hi-hat.
Example 12.16: Drum set, all instruments miked, and the snare drum double miked. Dynamic mics are added on the tom toms.

GLYN JOHNS TECHNIQUE

Glyn Johns is the recording engineer responsible for the huge drum sounds on many records by The Eagles, Led Zeppelin, The Rolling Stones, and The Who, to name but a few. The technique he developed uses only four mics! Because natural balances and the room are the majority of the sound, this technique does require great sounding drums, a good drummer, and a great sounding room.

- A large diaphragm condenser overhead is positioned 90 to 120 cm (3 to 4 ft) directly above the snare. Listen for the whole drum set to be well balanced and adjust the mic position until you achieve that. If the cymbals are too brash, raise the mic. If the rack tom toms are too quiet, position or angle the mic more towards them.
- Measure the distance from the center of the top snare head to the overhead mic.
- Position an identical large diaphragm condenser mic 15 to 30 cm (6 to 12 in) above the floor tom, so that it is the same distance from the center of the snare drum head as the overhead mic you previously set up. Yes, this one is probably lower than the cymbals, and definitely out to the side of the floor tom!
- These two mics should be panned left and right.
- Add close kick and snare mics to fill in the bottom and weight of the kick drum, and body of the snare. The kick mic can be inside or outside the drum. These mics are not the foundation of the sound – the close mics should supplement the sound of the overheads.

Figure 12.8 Glyn Johns technique.

AUDIO EXAMPLES
Can be found on the companion website

Drums – Glyn Johns Technique

Example 12.17: A drum set miked with the Glyn Johns technique, in a large live room. No additional room mics are used.

12.8 Room Mics

Even though the close mics pick up mainly direct drum sound, the sound of the room the drums are recorded in *will be* present in those mics – it is impossible to remove the sound of the room from a recording. *Room mics* are additional dedicated mics to picking up the reflected sound in the room, so that it can be beneficially blended with the close mics, instead of relying on artificial reverb. If the drums are being recorded in a good, bright and live, medium to large room, room mics can:

- Add a layer of excitement, power, and glue to the drum sound.
- Act as glue that bonds the close mics together.
- Turn small, flat, compact sounds into bigger, more exciting, powerful sounds.

Because details and reach are important, condenser or ribbon mics with good off-axis characteristics should be used as room mics. Room mics can be:

- Any kind of stereo array. However, in smaller spaces, spaced techniques are common because of the exaggerated sense of width and space they produce. In larger spaces, other mic arrays can be effectively used.
- Positioned 2 to 5 m (6 to 15 ft) in front of the drum set, at more-or-less any height, the array pointing towards the center of the drum set or towards the edges of the room for a more diffuse sound.
- Positioned more than 15 feet away in larger rooms with reverberant acoustics. Directional mics facing the outer edges of the room will focus on the reverberant sound.
- Often you do not need additional cymbal content in the room mics – but you do want more beef from the drums. To achieve this, try positioning a low room mic (or pair) 10 to 60 cm (4 to 24 in) above the floor, several feet in front of the drums.

Try not to put room mics too close to a boundary – you want them to concentrate on the general room sound, and not the reflections from that single surface.

PANNING AND SQUASHING ROOM MICS

Aggressively compressing room mic tracks will make their sound even bigger and more powerful. This is particularly effective for drum sounds, because it raises the amplitude of the ambient reflections in-between the drum hits – producing a stereotypically large rock drum sound.

A mono room mic can also be effective, particularly for solidifying the center panned elements of a drum set – you might not want the snare's room sound to be spread wide, but to keep it compact and focused with the narrow snare image. Try a condenser or ribbon mic at least 2 m (6 ft) in front of the drum set, 60 to 90 cm (2 to 3 ft) above the floor, pointing at the drum set. Compress it aggressively, by at least 12 dB, with a fast attack and fast release. This will reduce mic distance related phase problems, and bring out a solid and powerful sustaining room tone between the drum hits.

Room mics pick up some similar content to the close mics, but the sound is delayed in the room mics due to the increased distance from the drum set – so it is important to check for phase problems between the close and room mics. In a larger live room, the reflected sound in the room mics should be of much greater amplitude than any dry sound, so phase problems should be minimal. But, if you do notice a thinness or phasey sound as you turn up the room mics, reposition them to minimize this. Some engineers will time align the room mics by advancing them so that they are in phase with the snare mic. This is easy to do visually in a DAW, by moving the room mic tracks so an obvious kick or snare transient lines up in both the room mic tracks and its respective close miked track. To bring out the sound of the room even more, some moderate to massive compression can be applied to the room mics during mixing.

Figure 12.9 A fully miked drum set, including room mics.

WHERE IN THE ROOM?

Where in a room do you set up a drum set (or any other sound source)? Typically, a fairly central location works best for drums – no one wall is much closer or further away than the others, creating more evenly distributed early reflections, and the room and the resulting reverb are more symmetrical in the recorded image. In less-than-acoustically-perfect rooms, resonances can be greatest right in the center, so do listen for problems in the mics.

As well as trying different mics and mic techniques, do experiment with positioning the same instrument in different places in the room, because:

- The reflected sound the mic(s) pick up will change depending upon position.
- The spill picked up will change depending upon room position.
- The apparent low frequency output of the instrument will be increased by positioning it close to walls.
- Some rooms do exhibit more resonance problems in the center, so be prepared to move the drums if you hear room generated ringing frequencies (that are not the drums themselves ringing), particularly in the overhead mics.

AUDIO EXAMPLES
Can be found on the companion website

Drums – Room Mics

Example 12.18: Drum set, room mics are added to the individually miked drum set from previous examples.

Example 12.19: The close and room mics mixed together. The room mics are heavily compressed in this example. The room sound is wide and expansive.

Example 12.20: The close mics are blended with a heavily compressed, single, less distant room mic, 3 m (9 ft) from the drums, 75 cm (2.5 ft) high, pointing at the drums. The room sound is tight and focused within the drum set image.

12.9 Rock Drums vs Acoustic Jazz Drums

In most pop and rock production styles, each drum is treated as a single unit, and those instruments are miked as individually as possible. The engineer then gives those discrete sounds their own precise place in the mix by making different level, EQ, and panning adjustments, and applying different dynamic processes and effects to each. For example, to make a drum set bigger and more powerful than it is naturally:

- A short reflective room reverb can be applied globally, to most of the mics. But why not do this naturally with room mics?
- A short snappy snare, or small sounding tom toms can be turned into more powerful sounds by adding a longer, more obvious, stylized reverb to them. Do not apply the same processing to the kick, hi-hat, or cymbals – kick drums sound like they're in a tunnel, and high frequency sounds take on constant hissy characteristics when long reverbs are added to them.
- Ultra-heavily compressed or limited low-height room mics can add power and sustain to the kick, snare, and toms.
- Compression can be added to the kick drum to bring out more snap, punch, and attack, or conversely to emphasize the tail and tone.
- Compression can be added to the snare (and toms) to reduce the attack, and bring out and sustain the body of the sound, or conversely to emphasize the attack and crack of the sound.

In lighter acoustic styles such as jazz and folk, the drum set is treated as a more organic, unified whole. Each sound is less discrete, both sonically and functionally. There are different ways to achieve this goal:

- Use fewer mics. This means that the overheads are relied on to pick up a more complete picture of the drum set, and not just the cymbals. It is important to spend time positioning them so that the cymbal vs rest of the drum set balance is optimized – with fewer mics, balances are less adjustable during mixing.
- Use condenser mics instead of dynamic mics on the snare, tom toms, and even the kick drum. This will produce a brighter, more detailed recording – with more spill in each mic! *Embrace the spill!* Good sounding spill can contribute to a more blended sound. Mics with good off-axis pick-up characteristics are essential to capture good spill. With lighter styles of music, the SPLs produced by the drum set should be within the capabilities of most modern condenser mics, and less aggressive drummers are (hopefully) less likely to inflict mic damage from an erroneously placed fast moving drum stick!
- Use omnidirectional mics to give the sound a less muddy, more accurate, open, and transparent character. Less EQ will be required because there is no proximity effect clouding up the sound. In addition, each drum sound is made less isolated because of the better sounding off-axis response and increased pick-up of the sounds around the desired source. The picture painted by panning and balancing individual omnidirectional mics across the soundstage will be less precisely focused than one recorded using directional mics – each instrument will have more girth to it, and the picture of the entire drum set will be smoother and less detached.

CONSIDER THIS…

A jazz drum set sounds very different from a rock drum set – the instruments and source sounds are very different. *You just cannot make a jazz drum set sound like a rock drum set, or make a rock drum set sound like a jazz set!* You must start with the right source instruments, in the right room, and mic them appropriately.

It is also essential to consider how the drums will be mixed for different musical styles, so that the correct sounds are captured during recording:

- For tight modern rock and pop styles the close mics tend to be the primary sources, are introduced into the drum mix first, and are loudest. Then the overheads (which are primarily cymbal mics) are blended into that image, before the room mics are added at a relatively low level. If the room mic sound is stylized with heavy compression or limiting to add power and duration to the drum sounds, then the room mics should be introduced while working on those individual drum mics, because they are an integral part of those sounds.
- For big roomy rock drum sounds, the room mics can be introduced earlier in the process, and used louder than in the modern rock style above. The drums might be "fitted into the room sound" more.
- For jazz and other acoustic styles, the overhead mics (which should be picking up a balanced image of the *entire* drum set, and sound wonderful by themselves) should be the primary sources, and then the close mics blended into that image to add definition and focus (and not for power and volume – which would produce more of a rock drum sound).

AUDIO EXAMPLES
Can be found on the companion website

Omnidirectional Condenser Drum Mics

Example 12.21: Drum set, the same drum set as the examples earlier in this chapter, but miked with omnidirectional kick, snare, hi-hat, tom tom, and overhead mics. Compare this with Example 12.16, which uses directional mics.

MIC POLAR PATTERNS AND NULL POINTS

A multi-miked drum set uses a lot of mics in a relatively small space. There are many sounds coming from all around each mic – the desired sound, and spill from other parts of the drum set. For maximum rejection of spill, select polar patterns carefully based on the physical drum set-up:

- Cardioid mics have nulls point directly behind them – making them ideal if there is another sound source directly behind the mic.
- If there is no sound source directly behind a mic, but there is one "not quite directly behind" the mic, a hyper-cardioid (with its null points at about 120°–140°) may offer better rejection of that source of spill.
- If there is no sound source directly behind a mic, but there is to the sides, a bidirectional mic's extreme rejection at 90° can help isolate the desired source.

12.10 Drum EQ Frequencies

It is impossible to tell you what EQ to use! The examples below are fairly common, but the mic, instrument, performance, and mix context are all unknown variables that will change any EQ necessary. Where add or cut suggestions are given, don't be surprised if you end up having to do the exact opposite!

Kick Drum

- Add *thump* – 60 to 80 Hz.
- Reduce low-mid *boxiness* – 200 to 600 Hz.
- Add *snap* and *punch* – 2.5 to 6 kHz.

Make sure the weight and body of the kick drum and the bass guitar do not mask each other because of too much similar low frequency content. Focus the low frequency energy of each on slightly different frequency ranges by using subtractive EQ on different frequencies for each instrument during mixing.

Snare Drum

- *Weight* – 100 to 250 Hz.
- Low-mid *boxiness* – 400 to 600 Hz.
- *Punch* – 1 to 2 kHz.
- Add *crispness* – 4 to 8 kHz.

A high pass filter can be used to roll off low frequency spill from the kick and low toms.

Hi-Hat

- Cut low frequencies below about 300 Hz.
- Add *fizz* and *sizzle* above 6 kHz.

A high pass filter can be used to aggressively cut low frequency spill – but make sure the hi-hat doesn't sound too thin.

Tom Toms

- Rack tom *body* – around 240 Hz.
- Floor tom *body* – 80 to 120 Hz.
- *Attack* – around 5 kHz.

The tom tom mics pick up a lot of spill, and the tom toms are usually infrequently played, making them a prime candidate for gating if you're mixing rock or pop.

Cymbals and Overheads

- Reduce the *low gong* sound below 500 Hz using a shelf or high pass filter – but make sure it doesn't get too shrill – there is desirable body down there!
- Reduce *clank* – 800 to 1 kHz.
- Add *fizz* and *brightness* above 7 kHz.

13 Drum Tuning

In This Chapter:

13.1 Why Learn to Tune Drums?

Why learn how to tune drums? Because the recording process is “garbage in – garbage out”, and not all drummers know how to tune their own instruments! EQ, compression, and effects can only improve mediocre drum sounds so much, and will never turn poorly tuned drums into great sounding instruments. Truly horrible drums are impossible to rescue, and sample replacement is not suitable or appropriate for many musical styles.

There is no one way to tune drums – ask a dozen drummers how they tune their drums and you will probably get a dozen different answers. Their approach and process will be different, as will the pitches and tones they aim for. It can take drummers many years to master drum tuning, and in the process they develop their own technique and sound. This short chapter aims to give you an understanding of the concepts and process, so that if you are confronted with a disaster of a drum set you can at least improve it and make a recording of acceptable quality.

AUDIO EXAMPLES
Can be found on the companion website

Tuning Drums Makes a Difference!

The examples throughout this chapter show what a difference a quick tuning can make. The drum set used is deliberately "less-than-perfect" – similar to what some amateur and weekend-warrior bands bring into home and project studios. It will never be made "perfect," but it can certainly be made tighter and "better."

Example 13.1: A multi-miked drum set before tuning.
Example 13.2: A multi-miked drum set after tuning.

13.2 Fundamental vs Lug Frequencies

The *fundamental frequency* of a drum is the lowest pitch you should hear when both heads are tuned, free to vibrate, and the drum is struck in the middle of the top head. Fundamental frequencies are approximately 0.55 x the lug frequency, but this depends on the depth of the drum, and other variables.

The *lug frequencies* are a higher harmonic that is heard when the drum head is tapped close to the rim, near each tension rod. Lug frequencies are usually approximately 1.75 x the fundamental frequency, but again this depends on the specific drum equipment.

Touching the center of one of the heads gently with a finger, while tapping near a lug, will mute the fundamental and attenuate lug frequencies from elsewhere on the drum, allowing that specific lug frequency to be heard more clearly.

AUDIO EXAMPLES
Can be found on the companion website

Fundamental and Lug Frequencies

Example 13.3: Striking a tom tom in the center of the top head produces its *fundamental frequency*. 114 Hz in this example.
Example 13.4: Striking the edge of the drum head near a tension rod produces its higher harmonic *lug frequency*. 193 Hz in this example.

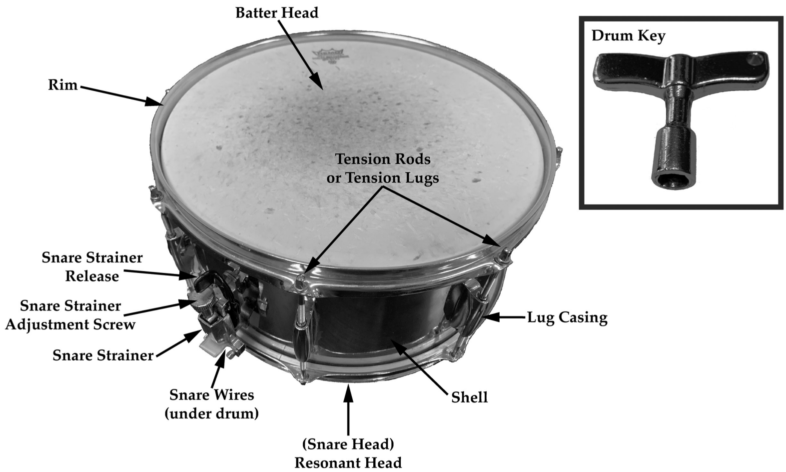

Figure 13.1 Snare drum parts, and a drum key. The same parts exist on tom toms and kick drums, with the exception of the snare parts.

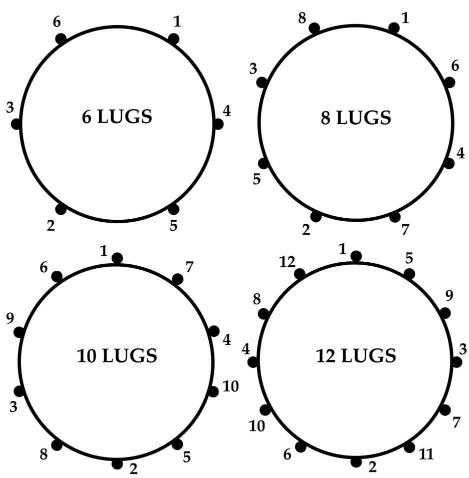

Figure 13.2 Tension rods should be tightened by equal, small amounts in numerical order, and the process repeated until the head is at an appropriate pitch.

13.3 Drums and Tuning Concepts

The parts of a drum referred to throughout this chapter are labelled in **Figure 13.1**.

To start tuning a drum, loosen both the batter and resonant heads so the tension rods are able to be freely twisted between your fingers and thumbs. Next, twist the tension rods so the drum head is tightened barely finger tight. Finally, use a drum key to tighten the rods in small increments in the order shown in **Figure 13.2**. Repeat this entire sequence, tightening a quarter or half turn at a time, so the head is gradually tensioned and raised in pitch, and the tension is kept even over the entire drum head. Even tensioning produces a better sound, hopefully free of rings, pings, and unpleasant overtones.

Some drummers prefer to listen to the fundamental pitch of the drum head while doing this initial tensioning. Others prefer to identify their goal by listening to the lug pitches.

Once the drum head has been tuned to approximately the desired fundamental or lug pitch, tap the drum near each lug, moving around in a circular motion, adjusting each tension rod to make all the lug pitches as identical as possible. If the resulting fundamental pitch is not as desired, re-tension the lugs using the patterns in the figure above, and then re-tune them to each other in the circular pattern.

Regardless of which method is used, a round or two of readjustments are usually necessary:

- If the lug frequencies are tuned first, and then the fundamental identified, the lug frequencies may well need further adjustment to produce the desired fundamental, followed by another round of fine tuning to make all the lug frequencies identical.
- If the fundamental frequency is tuned first, fine tuning the lug frequencies afterwards will also change the fundamental, so further adjustment may also be necessary.

The fundamental frequency gives each drum its pitch, and the infinite combinations of batter and resonant head tunings change the pitch motion, attack, and tone of the drum.

BATTER HEADS AND RESONANT HEADS

The *batter head* is the head that gets struck when played – the top of a snare drum or tom tom, and the kick drum head the beater strikes. The *resonant head* is the one on the bottom of a snare drum or tom tom, and the side of the kick drum that faces the audience.

Get a Key!

A standard drum key would be a great thing to have with you on every drum session, in case the drums need tuning. They are universal and fit most drum set instruments. Timpani, Latin and world percussion instruments will need different, unique keys.

TUNING RANGES

Each drum has a preferred tuning range that is dictated by the size of the drum, and the thickness and materials used for the head and shell. Trying to tune a drum outside of its effective range produces poor results – dead drum sounds, excess ringing and overtones, no stick rebound etc. Novices tend to overtighten the heads – very small pressure or tension changes create dramatic sonic differences. Some example fundamental tuning ranges are:

Kick drum:

- 40–80 Hz

Snare drum:

- 14 in (36 cm): 160–250 Hz.

Tom Toms:

- 8 in (20 cm): 165–210+ Hz
- 10 in (25 cm): 125–170 Hz
- 12 in (30 cm): 100–140 Hz
- 13 in (22 cm): 85–120 Hz
- 14 in (36 cm): 75–105 Hz
- 16 in (41 cm): 65–95 Hz
- 18 in (46 cm): 55–75 Hz

13.4 Kick Drum Tuning

For most rock and pop drum kits, a big, fat, beefy sound is desirable. This is obtained by tuning both the resonant head and batter head towards their lowest pitches:

- Start with both heads loose, the rods able to turn freely between your fingers and thumbs.

Then it doesn't matter which head you work on first:

- Finger-tighten all the tension rods to center the drum head – they should be barely tight, just so the slightest resistance to turning is felt. Tighten them two at a time, in opposite pairs by twisting the (usually) threaded part of the rods between your fingers and thumbs. Do this in the lug-order given earlier in this chapter.
- Press the head in the center to stretch it (if it is new) and re-center it, and repeat the step above.
- While tapping the drum in the center of the head, use a drum key to tighten the tension rods (in the same pattern as previously described) a quarter or half turn at a time until all the wrinkles disappear, and the drum starts to come up to a noticeable pitch.
- Tapping the drum head near the lugs, go around the drum and adjust the tension rods to make the pitch at each lug the same. The drum will start to resonate and "sing" as the pitches and tension get more even.

Next, work on the other head.

The tuning of each head can either be the same, or up to a fifth different, with either head higher or lower than the other. Jazz kick drums are generally tuned to higher pitches, with each head at a similar pitch.

AUDIO EXAMPLES
Can be found on the companion website

Example 13.5: A kick drum before tuning.
Example 13.6: A kick drum after tuning.

13.5 Snare Drum Tuning

If the snare wires are attached, loosen them and lift them away from the drum head. Use a drum key to loosen all the tension rods so that both heads have no tension and can move around a little. The resonant head is often tuned first, but the order really doesn't matter:

- Finger-tighten the tension rods to center the drum head – barely tight, with just the slightest resistance to turning. Tighten them two at a time, in opposite pairs by twisting the (usually) threaded part of the rods between your fingers and thumbs. Proceed around the drum in the lug-order given earlier in this chapter.
- Press the head in the center to stretch it (if it is new) and re-center it, and repeat the step above.
- Tapping the center of the drum head lightly with a drum stick should produce a low pitched, flat, "thunk" sound.
- Use the drum key to tighten each lug by the same amount, a quarter or half turn at a time, in the same lug pattern.
- Tap the drum head lightly near each lug (moving in a circular pattern around the drum), and adjust each lug's pitch up or down so that they are all identical. This will require going around the drum several times, because adjusting the tension at one lug has an effect on the others. This could require several quarter-turns on some lugs, and very slight adjustments on others.

Apply the same method to the other head.

For a higher, less "phat" sound, tighten the batter head more than the resonant head. A generic approach is to tune the top and bottom heads to approximately the same pitch. The musical and drumming style will dictate the tension and tuning – quite loose for big, "phat" sounds, or tighter for snappier, higher pitched sounds.

The tension applied to the snare wires, and the character and duration of the snare noise, can be changed using the adjustment screw on the side of the drum. If the snare is too loose it can buzz for an excessively long time, sounding unfocused and sloppy. If it is too tight, in addition to potentially over-stretching and damaging it, the snare noise will be too small and short.

AUDIO EXAMPLES
Can be found on the companion website

Example 13.7: A snare drum before tuning.
Example 13.8: A snare drum after tuning.

13.6 Tom Tom Tuning

Loosen all the tension rods on both heads so that the head has no tension and can move around a little.

The natural resonant frequency or "shell pitch" of the drum should be identified. This is often given in product literature, or can be heard by tapping the outside of the drum shell – the "knocking" sound will have a pitch component to it. This pitch can be used as an initial tuning target.

On the first head (it doesn't matter which one):

- Finger-tighten all the tension rods to center the drum head and remove ripples from the surface – barely tight, with just the slightest resistance to turning. They should be tightened two at a time, in opposite pairs by twisting the (usually) threaded part of the rods between your fingers and thumbs, in the lug-order given earlier in this chapter.
- Press the head in the center to stretch it (if it is new) and re-center it, and repeat the step above.
- Tapping the center of the drum head lightly with a drum stick should produce a low pitched, flat, "thunk" sound.
- Use a drum key to tighten each tension rod a quarter-turn in the same lug pattern, until the head begins to resonate at a pitch close to the shell pitch identified earlier.
- Tapping the drum head lightly near each lug (moving in a circular pattern around the drum), adjust each lug's pitch up or down so that they are all identical. This will require going around the drum several times, because adjusting the tension at one lug has an effect on the others. This could require several quarter-turns on some lugs, and very slight adjustments on others.

Apply the same process to the other head.

There are different ways to tune the top and bottom heads relative to each other:

- For a pure tone and long sustain, tune the top and bottom heads to the same pitch.
- For pitch-drop motion, tune the bottom head lower than the top head.

MUSICAL INTERVALS

Fundamental pitches are very important to tom toms – musical relationships can be created between the different sized toms. Pitch spacing between tom toms is usually some consistent musical interval – minor thirds, major thirds, perfect fourths, or perfect fifths between each drum. A common approach is to start with the lowest tom and work up from there, making sure the desired pitch of each higher drum is within its effective pitch range.

AUDIO EXAMPLES
Can be found on the companion website

Example 13.9: A rack tom drum before tuning.
Example 13.10: A rack tom after tuning.
Example 13.11: A floor tom drum before tuning.
Example 13.12: A floor tom after tuning.

13.7 Fixing Rings and Resonances

Rings and Pings

Rings and pings are overtones caused by different parts of the drum head being at different tensions – the head is out of tune with itself. To remove them, loosen all the lugs and start again, being sure to apply an equal amount of turns to each lug. Depending on the drum, this may not completely remove the rings or pings. Commercial gel products can be applied near the edge of the drum head to reduce these overtones, or in an emergency, gaff tape something like a credit card near the edge of the drum head, as shown in **Figure 13.3**. Any

Figure 13.3 A "credit card" taped near the edge of the drum head is an effective way to tame stubborn rings and overtones that tuning (and even commercial gel products) may not remove.

damage to a head (dents, holes, tears, defects because of over-tightening etc.) can make tuning difficult and getting a great drum sound impossible.

Sympathetic Resonances

If two drums are tuned to a similar (or sometimes harmonically related) pitch, hitting one drum can cause a head on the other drum to produce a droning or ringing pitch. This sympathetic resonance will be magnified by a close microphone. Snare buzzing can occur when a tom tom is played, or a tom tom can ring when the snare drum is played. The kick drum can cause a floor tom or even a large rack tom to resonate.

The solution is to re-tune one of the affected drums so their fundamental frequencies are less similar or less related. If the tom toms are tuned to a musical interval concept, then re-tune the snare drum or kick drum. If the toms are not tuned to any particular musical interval concept, then it is possible to re-tune the tom tom.

13.8 Tuning Devices and Apps

There are many hardware drum tuning devices, and apps for mobile phones available. These can help you learn the tuning process, and increase the accuracy of the process. Some devices use microphones to analyze the sound, and others measure drum head tension without the need to make sound. Some phone apps include tutorial content which can be a great resource.

Apps, such as the one shown in **Figure 13.4**, can have knowledge of typical fundamental and lug frequencies for each type and size of drum, as well as how those frequencies and

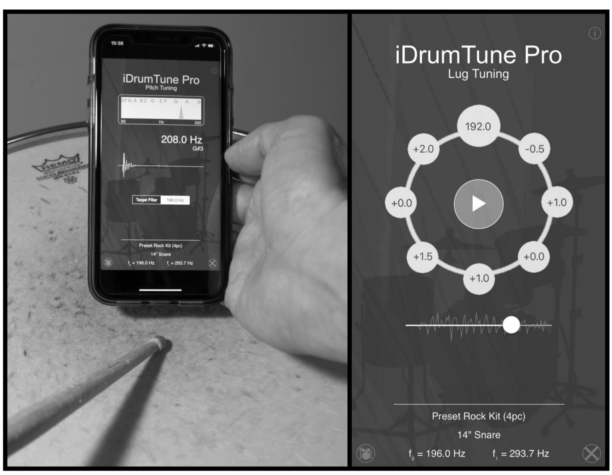

Figure 13.4 A drum tuning app. **Left:** The fundamental pitch tuning page shows the actual pitch of the drum, and a suggested target frequency. **Right:** The lug tuning page shows the relationship (in Hz) of each lug compared to a first (reference) lug.

the relative tuning of the batter and resonant heads typically relate to each other for different drum sounds and tuning effects.

YOU MUST HAVE A GOAL

Any app can help speed up the processes of making all the lug frequencies the same, and achieving a target fundamental frequency – but to purposefully create a personal drum sound, you must have a tuning concept in mind. This means understanding the resonant and tonal possibilities of each drum, the possible pitch relationships between all the drums in the set, and the batter/resonant head frequency combinations that create different effects.

Tune for the Song

Different songs, musical styles, and key signatures usually benefit from different drum sounds – both in terms of the timbre created by top/bottom head tunings, and the pitch relationships between each drum of the kit. Good drummers will tune the drums “to the song.”

14 Guitars, Basses, and Keyboards

In This Chapter:

14.1 The Role of the Rhythm Section

What type of music are you recording? To a certain extent this will suggest the relative importance of the rhythm section instruments.

The drums and bass work together to lay down a foundation to which the other instruments and vocals respond. In most commercial styles the drums and bass are equally important.

Guitars and keyboards are musically less important in many rock, pop, and jazz styles unless the instrument is taking a solo, in which case it replaces the lead vocal in terms of focus and importance. They are, to a certain extent, "fillers," often multi-layered, providing harmonic information and adding textural interest to the groove and foundation provided by the drums and bass. Most importantly, these instruments should never get in the way of the intelligibility of the vocals – which they will certainly do if they are used too loudly in a mix, or they have thin timbres that over-emphasize the 2 to 4 kHz range. A common cause of vocal levels getting out of control and needing to be turned up and up in order to be heard and *understood*, is that the harmonic rhythm section instruments are too loud or have too much competing frequency content.

14.2 Electric Guitar

An electric guitar is nothing without the amplifier and loudspeaker cabinet it is plugged into – its sound is dependent upon them. Change either one, and the sound will be dramatically different. Unlike good studio monitors, guitar cabinet loudspeaker cones and enclosures have all kinds of frequency and amplitude quirks that make them characteristic and desirable!

When miking a guitar cabinet:

- Dynamic microphones are the usual go-to choice. There is little desirable frequency content above about 6 kHz in most electric guitar cabinet speakers, so the extended frequency response of a condenser mic is not essential.
- Condenser mics can be used for a brighter sound – but they are sometimes a little too aggressive and scratchy.
- Ribbon mics, with their smoother high frequency characteristics, can sound great! Remember that the front and back of many ribbon mics have slightly different characteristics – so do try them both out and decide which gives you the most pleasing sound.
- Radically different sounding microphones, for example, a dynamic mic and a condenser mic, will each accentuate different frequency components – so it can sometimes be beneficial to use multiple mic technologies, each chosen and positioned to favor specific timbres which can then be blended together.

A mic placed a few centimeters (an inch or two) from the loudspeaker cone will have more proximity effect than one placed 15 cm (6 in) from the grille. Most dynamic mics start sounding quite thin when they are more than 20 cm (9 in) from the sound source, so adjust the mic distance until you find the best sound. Because of the loudspeaker cone's unpredictable response and distortions, moving the mic slightly to the left, right, up, or down, can change the sound quite significantly, so take the time to experiment and find the sweet-spot:

- For the brightest sound, position a mic on-axis in the center of the loudspeaker cone.
- For a warmer sound, position the mic off-center, towards the edge of the cone.

If a guitar cabinet has more than one loudspeaker cone, get your ears close, and listen to each in turn – *at a safe listening volume!* Even better – try the mic on all of them, and see

Figure 14.1 A guitar cabinet loudspeaker miked on-axis for the brightest sound, and off-axis for a warmer sound.

how each sounds from the mic's perspective. Decide which one sounds most suitable for the project. Listen carefully for any buzzes or crackles indicative of a bad cone. Unless loudspeaker cones sound radically different, there is limited value to miking multiple cones. You can try mics with very different characteristics to blend together in mono, but there are rarely enough differences to produce a wide stereo effect if they are panned left and right.

You want to record a sound that will work *in the mix*, and not necessarily the bright piercing sound a gigging guitarist might dial in so they can be heard in the far reaches of a bar or club. If the guitar sound coming from the monitor speakers isn't what you need, have the guitarist come in to the control room to listen to what the mic is picking up. They are likely to know how to desirably change the sound once they hear the miked sound.

WHERE IS THE LOUDSPEAKER CONE?

Sometimes you can't actually see loudspeaker cones – they are hidden behind grilles. Shine some light on the situation – use a flashlight or flashlight app on your phone so you can see through the grille and find the cone's exact position.

For loud heavily distorted sounds, the interaction between the loudspeaker cabinet and the room is a big factor in the recorded sound. There is so much sonic energy bouncing around, and *lots* of this reflected content gets into the mics. Boomy, muddy guitar sounds are unfortunately typical when recording in a room that is generically treated with broadband acoustical absorption, and not enough bass trapping – the high and upper-mid frequency reflections are attenuated by the treatment, but the lower frequencies still roll around in the room.

The position of the loudspeaker cabinet in the room will affect the sound:

- Is the cabinet in the corner, or close to a wall? Or is it on the floor? In which case low frequencies will be boosted, and the sound made boomier. At the very least, try raising the cabinet off the floor for a clearer, brighter sound.

- Is the cabinet near reflective, or absorptive surfaces? Reflective surfaces will make the sound brighter and more present, but there is a chance the sound picked up by the mics will be slightly comb filtered due to the reflections. Absorptive surfaces will deaden and darken the sound.
- Try the cabinet and mic set-up in different places, and facing different directions in the room. The reflected content will change in each position, and will interact differently with the direct sound coming from the cabinet.

DOUBLE TRACK IT!

If you only record one rhythm guitar track it often ends up A) panned centrally – which is boring, and where it competes for physical space with more important sounds, or B) off to one side – which can cause asymmetry in the mix. Although a little cliché, double tracking is a quick and effective way to create a wide expansive rhythm guitar image that exploits the less utilized extremes of the stereo image. It frees up the center, increasing the clarity of more important center panned sounds.

- Record two separate takes of the rhythm guitar part, with or without *slight* performance and tonal variations between them.
- Pan one left, and the other right.

Even with the same mic on the same speaker and the guitarist trying to play the same material, there are enough subtle pitch and timing differences in the performance to create an expansive, wide stereo image when the tracks are opposingly panned. The performance differences are usually enough that comb filtering when summed to mono is not a problem – but be sure to check mono compatibility before the performer leaves the studio!

AUDIO EXAMPLES
Can be found on the companion website

Double Tracking Rhythm Guitar

Example 14.1: An electric guitar, one take, two speaker cones miked individually and panned hard left and hard right. This produces a very narrow, slightly smeared stereo image. It is not effective or an exploitative use of the soundstage.

Example 14.2: An electric guitar, two takes, double-tracked and panned wide left and right. This produces a dramatic and wide image.

A close mic and distance mic on a guitar cabinet may sound impressively wide when panned hard left and right – but when summed to mono, you may hear phase problems and comb filtering. You cannot ignore mono compatibility, and need ensure the sound sums to mono to play back well on mono systems.

Comb filtering should be minimized in both single mic plus reflections, and close mic plus distance mic situations. Changing the distance between the two mics will change the frequencies at which comb filtering occurs:

- If you have an assistant, you can have them move the distance mic while you listen to it. Blend the close mic in and out with the distance mic in different positions. Listen for timbral shifts each time the distance is changed, and pick a spot where the sound is fullest when they are combined together.
- If you are working alone, put a pair of isolating headphones on in the recording room, and listen to the blend of both mics as you move the more distant mic around.

COMB FILTERING POLARITY TRICK

A method to reduce comb filtering between mics at different distances from the same sound source is to polarity reverse one of the mics, and listen to them both, panned identically and at equal amplitudes (trim them so the metered levels are equal, and put both faders at unity). As you move the distance mic, listen for where the combined signal is the lowest amplitude. They will never cancel completely – listen for the least amount of this difference signal. This position is where the mics are picking up the most identical, or least different frequency content. Once this mic position has been found, undo the polarity reversal and see if you get a fuller, less comb filtered sound.

A mic can be positioned behind the loudspeaker cone if a loudspeaker cabinet is an "open" design and you can see the rear of the loudspeaker cone. The signal picked up by this mic will be polarity reversed, so if combined with a front mic, there will definitely be cancellation. Before recording, the polarity reverse button on the rear mic's preamp, mixer, or DAW channel should be engaged, and the summed signal checked to ensure that there is no thinning, hollowing, or phasing.

DON'T DRAPE YOUR MIC!

Hanging an end address mic down the front of the loudspeaker grille is not recommended. They end up pointing at the floor, with the loudspeaker cone off-axis.

A side-address mic in this position will be on-axis, and although not subject to off-axis coloration, the mic is *very* close to the loudspeaker cone – so proximity effect will be excessive. The recording will also have very little natural depth. As long as spill is not a problem, a more natural sound can usually be achieved by increasing the distance between the mic and loudspeaker cone a little. That said, sometimes mics are positioned very close to loudspeaker cabinets (using stands, and not draped) for deliberately close perspectives, or so proximity effect will increase the weight of a sound that is too light or thin. Try a close omnidirectional mic to avoid low frequency build-up if a very close sound is desired.

Big powerful amps and loudspeaker stacks are not the best choice for recording purposes. A medium-sized combo that exhibits colorful distortions at lower volumes is a better choice. In many cases, the bigger and louder the amp/loudspeaker is, the noisier it is – and noticeable hiss, hum, and buzz can build up quite quickly when multiple guitar tracks are recorded. A very small practice amp/loudspeaker combo will always sound very small, so is not a good choice either.

Some guitar amps are notorious for only sounding "really good" when turned up to "11"! Particularly tube amps. This high SPL is another reason that dynamic mics are commonly used on guitar cabinets. *Power-soaks* are devices that are put between the amplifier and the loudspeaker cone. They allow the amplifier to operate at high levels and impart its character on the sound, but cause the loudspeaker cone to output lower SPLs because some of the electrical power is "soaked up." They make recording safer for the mic and people in the same room, and are essential if you're recording in a small room that is easily overloaded with sound. The only drawback is that causing the loudspeaker cone to work less hard may reduce its own characteristics.

AUDIO EXAMPLES
Can be found on the companion website

Miking Electric Guitars

Example 14.3: Electric guitar cabinet, dynamic mic, center of the cone, just a few inches away.
Example 14.4: Electric guitar cabinet, dynamic mic, near the edge of the cone, just a few inches away.
Example 14.5: Electric guitar cabinet, condenser mic, close, center of the cone.
Example 14.6: Electric guitar cabinet, condenser mic, two feet from the loudspeaker.

14.3 Creative Comb Filtering

Comb filtering *can* also be a creative tool!

Deliberately used, comb filtering can create wonderfully strange and unnatural effects.

- By blending two microphones, each a different distance from the cabinet, together as an identically panned mono point source image, the overall EQ or frequency balance of a sound can be changed dramatically.

- The distance between the mics dictates the specific frequency bands affected, and the amplitude of each mic changes the amount of the effect.
- Engaging the polarity reverse button on one of the channels will produce the difference between the mics – this is an even more intense effect.

These types of effects are probably *not* suitable for a standard rhythm guitar track, but they can be used on creative solo flourishes or as sporadic effects – ear candy to turn a demo tape into an interesting recorded production.

Creative comb filtering usually benefits from using mics that are as different as possible – different technologies, different frequency responses, and different off-axis pick-up characteristics. This makes the interactions between them more random, and you're more likely to get them to work together beneficially.

Always make sure that creative phase and comb filtering artifacts are:

- Controlled and deliberate.
- Appropriate.
- Mono compatible.

2 + 1

It's common to use a dynamic mic and a condenser mic as close mics on a guitar or bass loudspeaker – each producing different but equally useful characteristics. Because they are both positioned the same distance from the loudspeaker, their phase relationships are identical, and there is no comb filtering between them. Adding a third, drastically different mic, at a distance of at least 30 cm (1 ft) creates a trio of faders, and each can be adjusted independently to produce different characters and EQ effects.

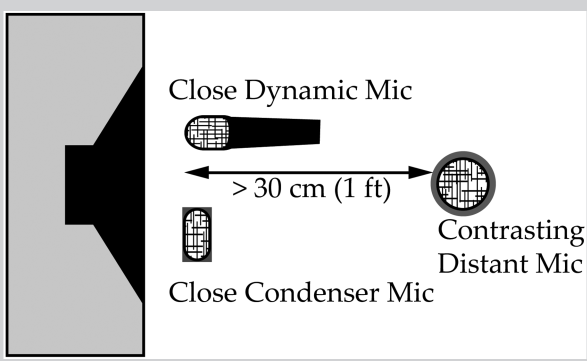

Figure 14.2 A top view of a trio of mics on a guitar cabinet.

14.4 Direct Boxes

DI, or *Direct Injection* boxes take an unbalanced signal (usually line or instrument level, but sometimes even loudspeaker level) and convert it to a balanced mic level signal. *Passive* DI boxes do not require any power to operate. *Active* DI boxes require power – phantom power or sometimes batteries. Good active DI boxes generally have better frequency response, a more detailed and transparent sound, higher output levels (requiring less preamp gain), and are cleaner and less noisy than passive DI boxes. But a passive DI with a high quality transformer can certainly sound better than a poor quality active DI.

Any instrument or sound source with a line or instrument level output can be DI'd – making it easy to plug the quarter-inch line level outputs on a synthesizer into the XLR mic inputs on a preamp, studio wall plate, or snake system.

Some guitar and bass amps have "direct outs" or "line outs" built in. These should be checked for noise and quality before use – a good studio DI might sound better. Do ensure that you are not plugging a post-amplifier loudspeaker level output into the DI box, unless the DI box is specifically designed to handle a loudspeaker level input, and all the necessary buttons to do so are pushed.

DI boxes are usually characterless. This works well for things like synthesizers, which have all their character built in. But they do not work well for guitars and basses, which rely on preamp and amplifier electronics, and speaker cabinets and speaker cones to produce most of their character. A miked speaker cabinet produces the most characterful electric guitar or electric bass sound. But if miking isn't possible, or all you have to work with is a bland DI track, then guitar and bass preamps, amp simulators, and reamping are possible solutions.

LIFTS AND ISOLATORS

When connecting instruments and recording devices together using DI's, sometimes low pitched hums or buzzes can be heard – particularly when the devices on each end are plugged into different electrical circuits. Plugging the bass/guitar rig into the same circuit as the recording equipment is the easiest solution. If that's not possible, or doesn't fix the problem, engaging the DI's *ground lift* (*earth lift*) to disconnect the audio ground/earth connection between the devices may remove or reduce the noise (assuming there are no problems inside the musical/audio equipment).

For extreme hum/buzz problems, *transformer isolators* can be used.

Never remove the earth or ground pin on a power plug, or use a 3-prong to 2-prong power adapter to remove hums and buzzes. In case of equipment fault, these leave potentially fatal power voltages nowhere safe to go (to trip breakers), so the high voltage travels the path of least resistance and ends up at the mic, audio equipment, or instrument, and then the performer or engineer – *resulting in serious electrical shock or even death*.

14.5 Reamping

A guitar sound that is not "dirty" enough, or a clean DI'd bass sound, can be *reamped* in order to add amplifier and loudspeaker character after recording. A previously recorded (clean) sound is sent (via a mixer aux, or DAW send to an interface output) to a reamp box, and then to a guitar or bass amp, with or without effect pedals, as shown in **Figure 14.3**. The speaker cabinet is miked and recorded while the clean/DI'd sound plays through it. This captures *real* amp, cabinet, and loudspeaker cone distortions which can be blended with or replace the original clean track. It is always a good idea to have the sound of the cabinet recorded, even if you don't end up using it.

The output of a DAW, mixer, or multitrack recorder is line level, and therefore the wrong level and impedance to connect to a guitar or bass amplifier or pedal directly, so a *reamping box* is essential if you want the correct tone, and don't want to risk burning up the instrument amplifier! Be sure to adjust the gain control on the reamping box (if it has one) to avoid overloading the amplifier inputs and causing undesirable distortion. Start with a clean sound on the amplifier to make sure there is no input overload distortion, before proceeding to dial in the character you want the amp and speaker to provide.

14.6 Amp and Pedal Simulation

There are many hardware devices and plug-ins available that simulate amp, loudspeaker, and pedal characteristics – so a clean guitar or bass track (or any other track for that matter) can be given character or dirtied up without an amp, cabinet, or pedals, while tracking, or in a DAW after it is recorded. Of course the result is not quite the same as using real amps, cabinets, loudspeaker cones, analog pedals, and microphones – and some simulators do

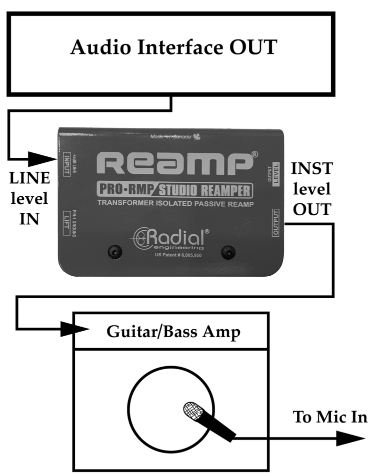

Figure 14.3 A reamp box plugged in between a mixer aux send or DAW send output, and a guitar or bass amplifier input.

sound cheap and "digital." Good simulation *can* save the day if all that is otherwise available is a bad amp, cabinet, pedal, and/or a bad room.

Hardware simulators are generally used while tracking, so the sound is "dialed in" and committed to during recording. Plug-ins allow the sound to be changed and refined during mixing. If you use simulation plug-ins, dial in something at least very close to the sound you think you'll need for the mix, before pushing the record button. Deferred decision making is not in the best interests of a good mix. The electric bass is a fundamentally important part of the "bottom" of the mix. How can you really know what sounds you need from the instruments around and above it, if you defer deciding on its sound until mixing? Same with electric guitar sounds and the other sounds that go around it.

Reamping provides many varied creative possibilities instead of simply recalling a preset in a simulator. "Everybody else" with the same simulator has that same preset, so it doesn't help you obtain your own unique sound.

14.7 Electric Bass

One way to achieve a good beefy recording of an electric bass is to mic the speaker cabinet – to capture the characteristics of the amplifier, the cabinet design, and the loudspeaker cone.

- Regular dynamic mics or low frequency instrument dynamic mics are the go-to mics for bass cabinets.
- Condenser mics and ribbon mics can also be used, but they may pick up too much snap and high frequency content.
- In a good room, which is free of low frequency problems, and in which there are no sources of spill, an omnidirectional mic will produce a more neutral, less muddy sound, than a directional mic.

The position of the loudspeaker cabinet in the room will have a huge impact on the recorded sound:

- If the cabinet is on the floor, or close to any walls, its perceived low frequency output will be boosted, and the sound may get too muddy and muffled. Conversely, this may help a small speaker cabinet sound bigger – but listen for low-mid frequency muddiness, and EQ it out on the amplifier if necessary.
- The cabinet and mic position in the room will change the recorded sound – so do experiment with positioning the cabinet in the center, a little off center, or more to the end or side of the room, and facing it in different directions. In a room with low frequency standing wave problems, the mic could be in a cancellation node of the fundamental frequency (or a significant harmonic) of the key of the song, or it could be at

a boosted antinode, where that frequency booms too loudly. If some notes seem louder or quieter than others (and it's not the player or instrument causing that) try moving the mic and cabinet a few feet.

A good sounding, reasonably large bass amp is necessary for a good miked bass sound. A small practice amp will always sound "small." You don't need a very loud amp/cabinet combination for a good recording, but you do need something that can produce weighty low frequency content. Twelve or 15 inch loudspeakers are ideal – 18 inch loudspeakers can lack focus and definition.

- A mic can be placed anywhere from a few centimeters (an inch or two) in front of the loudspeaker cone, making sure it is not touching the grill, to 60 cm (2 ft) or more away.
- The closer the mic, the greater a directional mic's proximity effect, and the boomier and muddier the sound will be – but spill from adjacent sound sources will be reduced.
- Moving the mic 15 to 30 cm (6 to 12 in) away will produce a clearer sound with reduced low frequency build-up.
- Moving the mic a meter (3 ft) away will produce a sound more similar to the bass sound we hear when listening from a distance away – but it will also pick up more spill and room reflections, and probably lack definition in most rooms.
- If the mic is positioned close to the loudspeaker, the brightest, crispest sound is picked up by a mic pointed directly at the center of the cone.
- A warmer, more rounded sound is produced when the mic is positioned towards the edge of the cone.

DESIGN THE BASS SOUND FOR THE MIC!

A close mic hears the bass amp very differently to how we hear it from a distance away. The mic then adds its own character. Adjust the sound of the bass amp so it sounds great through the mic and monitor loudspeakers in the control room, before pressing the record button. Do not use the bass player's default settings and plan on "fixing it in the mix."

- If the bass sound is too boomy, turn down the lows on the instrument or amp, move the mic further away, or move the cabinet off the floor or away from the walls.
- If it is too small and boxy, turn down the mids on the instrument or amp, or even try putting it closer to a wall, or in the corner of a room to boost its low frequency output.
- If it is too twangy, turn down the highs on the instrument or amp.

DI'S & BASS PREAMPS

Another approach to record a bass sound is to use a DI from the bass output or immediately following any effects pedals. A DI'd sound is crisper and snappier than a mic on a cabinet – but most generic DI's by themselves sound small and bland, lacking real bottom and weight.

Bass preamp devices *can* produce great results – particularly, characterful tube preamps! Using a good bass preamp instead of a cabinet and mic also removes a significant source of spill from the studio. Important to making a direct sound work is to experiment with the 120 to 350 Hz frequency range, as this is where a lot of the character of the sound is defined.

A DI (or bass preamp) signal can be blended with a cabinet mic – and both should be panned identically, usually to the center:

- The mic often provides the real lows, weight, and girth of the sound.
- The DI commonly adds punch, clarity, and focus.

The direct and mic signals should be recorded simultaneously – *not* as two separate takes, which would have unavoidable performance differences. The set up for this is shown in **Figure 14.4**.

Check there is no phase cancellation between the two signals. Compare one channel solo, to both channels combined. Make sure the sound is beefier in the low frequencies when both channels are combined. If not, then try polarity reversing one of them and see if the LF gets bigger, or in your DAW, zoom in and nudge one of them to manually line up the

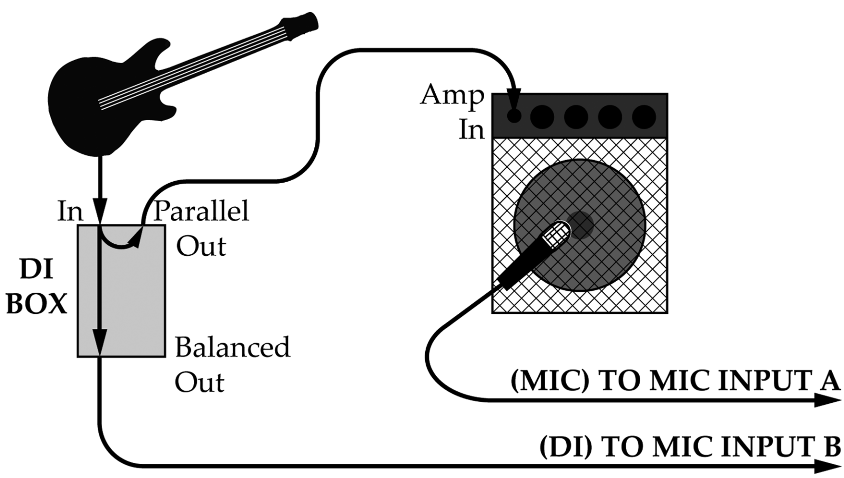

Figure 14.4 Recording a bass mic and DI simultaneously.

waveforms. If any mid or high frequencies get attenuated or phasey when both signals are combined, that should be corrected by changing the mic distance before recording.

AUDIO EXAMPLES
Can be found on the companion website

Miking Electric Bass

Example 14.7: Bass cabinet, dynamic mic, center of the cone, about 6 inches away.
Example 14.8: Bass cabinet, DI only.
Example 14.9: Bass cabinet, the mic, and DI mixed together.

14.8 More on Guitar and Bass Cabs

Multiple Drivers?

Many cabinets have more than one loudspeaker. Some bass cabinets have multiple cones, sometimes of different sizes, and maybe even a tweeter. The "whole" sound of a multi-driver cabinet is a blend of all of the drivers – and the phase anomalies between them. A single mic, close to one loudspeaker cone, does not pick up these characteristics – the sound may be "in-your-face" and have some good grunt, but it can also be a little flat and one-dimensional.

A mic 20 to 60 cm (8 in to 2 ft) in front of a single speaker cabinet will pick up the sound of the entire speaker cone. A mic 1 to 2 m (3 to 6 ft) away is a true distance mic (and is only a sensible option if the instrument is recorded in isolation in a dead acoustic). Distance miking is particularly useful on cabinets with multiple loudspeakers, as the increased distance allows the mic to pick up the blended sound of all the loudspeakers – but the greater distance also increases the reflected content picked up:

- The sound of a distant mic will be less boomy and muddy, and more open and transparent than a close mic. But the sound may get too "roomy."
- Floor reflections picked up by a distant mic may cause phasing and comb filtering. The distance between the cabinet and mic, and height of the mic and/or cabinet, should be adjusted to minimize this. The closer the mic and/or cabinet are to the floor, the greater the amount of comb filtering, the higher the affected frequencies are (maybe pushing them out of problematic ranges), but the more bass boost there is.

Depending upon the specific room acoustics, a distance mic can sound great by itself. A distance mic can also be particularly effective when blended with a close mic or DI:

- The close mic or DI provides the weight and grunt of the sound.
- The distance mic adds power and penetration.

Regardless of whether distance mics are used in isolation or combined with close mics, the effects of phasing and comb filtering need to be listened for and the set-up and sound corrected before pressing the record button.

If a bass cabinet has a tweeter in addition to woofers, use multiple close mics – one on the woofer cone and one on the tweeter. For most bass sounds that are not slapped and popped, the tweeters are not necessary – if a cabinet has multiple 12 inch cones plus a tweeter, combining a mic on the best sounding 12 inch with a DI is usually sufficient.

Boundary Loading

A *boundary* is a wall, floor, or ceiling. The low frequency output of a sound source *or microphone* will be boosted by approximately 3 dB by each boundary that it is positioned close to. A bass or guitar speaker cabinet put on the floor in the corner of a room will have its low frequency output boosted by up to 9 dB! If a mic is positioned close to that source, it is also effectively boundary loaded, so there will be up to another 9 dB boost, *plus* any proximity effect – making the sound dull, dark, congested, and confused.

- Avoid setting equipment up in the corners of rooms – unless the sound needs a significant bass boost.
- To remove "mud," try raising speaker cabinets off the floor with either commercial cabinet stands, sturdy chairs, or small tables.
- Due to the longer wavelengths of lower frequencies, a bass cabinet needs to be positioned further from boundaries than a guitar cabinet in order to reduce the effects of boundary loading.

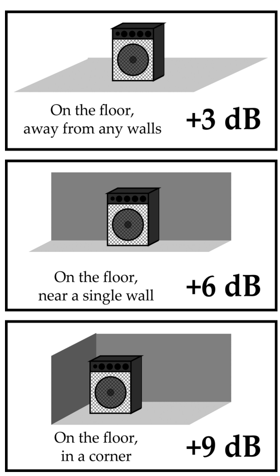

Figure 14.5 Speaker cabinet positions and the effects of boundary loading.

Boundary loading *could help* a small, thin sounding guitar or bass cabinet by giving it some added weight – but ultimately, a small speaker cabinet is always going to sound like a small speaker cabinet!

14.9 Acoustic (Upright) Bass

The acoustic bass is a unique miking challenge because:

- The instrument is relatively quiet, which makes spill problematic unless the instrument is recorded in isolation.
- Sound dispersion is very uneven over the instrument – different parts of it radiate radically different frequency content in different directions.
- The instrument's desirable sound radiates over a very narrow focused angle directly in front of the instrument, but the sound is not a natural unified "whole" close to the instrument.

A single close mic is unlikely to capture a good bass sound. A combination of two (or more) of the following techniques will produce the best recording of the instrument:

1 Listen for the sweet-spot. Every instrument has a point in front of it where its constituent sonic components combine to form great sound. With your ears 30 to 60 cm (1 to 2 ft) from the instrument, move around while listening for a sound that is big, warm, and full, with definition and some string attack, but not too bright, muffled, or cloudy. When you find the sweet-spot, put a mic in that position. Experimentation is key, as no two basses are identical, and different player's coax different sounds from the instrument.
    - A good starting point is anywhere from 30 cm to a meter (1 to 3 ft) away, with the mic angled down towards just above the bridge or top of the f-holes.
    - The sounds to the left, center, or right of the instrument, especially up closer, are each different – so be sure to try these different positions to see which sounds best.
    - Sometimes a more distant mic like this can sound a little scratchy, and lack bottom end – so blending it with a closer mic is desirable.

2 A significant amount of sound comes from the f-holes, which can be used to provide "bottom end" when mixed with an additional mic that provides definition and detail.
    - A mic placed 9 to 15 cm (3 to 6 in) from an f-hole will pick this up.
    - Both f-holes sound different – so be sure to try them both.
    - F-holes do tend to lack higher frequency details and sound muddy and undefined – by themselves they are unlikely to produce the best sound.

3 DI the pick-up. For jazz and folk styles, contact pick-ups on the instrument's bridge are commonly used. If they are good pick-ups, correctly installed, they *can* sound quite good, and provide an in-your-face intensity that a mic doesn't have. Bad pick-ups sound

truly awful and are unusable! Good pick-ups can be blended with an additional mic – a more distant mic, or an f-hole mic, or both. Like an electric bass DI, the pick-ups themselves do not portray the true weight and acoustic depth of the instrument.

4 Bridge miking. Some players have a small (omnidirectional) lavalier mic installed on the bridge instead of pick-ups. These can provide a beneficial layer of sound, or be unusable. Between the strings and the amplifying body of the instrument, this very close mic picks up an incomplete frequency picture of the instrument – so it is best used as one component of a multi-mic approach. Engineers have been known to take a studio condenser or dynamic mic, put it in a large foam windscreen, and wedge it into the hole in the bridge, pointing upwards or into the instrument – a conventional mic in this position produces a very boomy and inaccurate sound that needs combining with another mic which will fill in the missing detail and high frequencies. There are some commercial mics, transducers in foam, designed for this position on string instruments.

MULTI-MIC PANNING

Regardless of how many mics, DI boxes, preamps, and/or pick-ups are used simultaneously and blended to form a bass sound, they should always be panned identically. The bass is a narrow point source that needs to be horizontally compact, cohesive, and powerful in the mix. Panning elements of the sound differently will result in an unfocused, weak, smeared sound.

What mic should you use? Because basses are "bass instruments," dynamic mics are a common choice for close miking – "flat" general-purpose mics, or "low frequency instrument" mics. But, the acoustic bass is *not* loud, and *does* contain high frequency and low level details (string noise, finger and pluck noise, and the details of the bowed sound, etc.) so

Figure 14.6 Acoustic bass mic positions **1** & **2** described in the main text.

condenser and ribbon mics are excellent mic choices if a more detailed, transparent sound is desired. A good, quiet, high gain preamp is also essential.

Each mic used should make the sonic element it is providing sound great. For example:

- The body, definition, and stringy details should be accurately represented by a mic 30 to 60 cm (1 to 2 ft) away. That distant mic may sound a little lacking in low frequency weight, but that's okay – that's not the job of this more distant mic.
- A mic much closer to the f-hole or bridge is going to sound muffled, and lack the high frequency or "whole instrument" details – but what's important in that mic is that the low frequency weight of the sound is good.

Those two mics should then be complementary to each other when combined.

If a mic isn't getting the sound you want, try something quite different – crazy even! Surprising results can come from unlikely suspects!

LOW FREQUENCY ROOM PROBLEMS

Low frequencies are very difficult to control, and acoustically improper rooms can have very uneven low frequency characteristics – causing the bass to sound extra boomy on some notes and thin on others, or some notes will appear louder than others. Make some homemade bass traps to position around the room and instrument to reduce these problems. Change the instrument and mic's position in the room to see if you can find an area which provides a more even response.

AUDIO EXAMPLES
Can be found on the companion website

Acoustic Bass

Example 14.10: Acoustic bass, a condenser mic placed about a foot away, just above the bridge of the instrument.
Example 14.11: Acoustic bass, a dynamic mic positioned a few inches away from an f-hole.
Example 14.12: Acoustic bass, the two mics above are blended together for a more complete picture of the instrument's sound.

14.10 Acoustic Guitar

The acoustic guitar is a naturally wide source that can fill the stereo soundstage. But should you record it in stereo or mono?

- If there are multiple acoustic guitar players in a band, or multiple tracks of a single player are going to be recorded and panned to produce a wide image in a busy mix, mono guitar tracks are most useful. Multiple stereo acoustic guitar tracks tend to sound imprecise when each is panned to a different position in the stereo image. Panned on top of each other, they sound messy and confused.
- If just one single acoustic guitar track is being recorded, stereo miking is appropriate.

Condenser and ribbon microphones are ideal for recording acoustic guitar because of the high frequency content, low-level details, and transient information essential to its sound:

- Small diaphragm condenser mics are the go-to option for the most accurate and potentially brightest sound.
- Characterful large diaphragm condenser mics give the instrument a little extra power, and hype.
- Cardioid mics are most typically used, but in a great sounding room, omnidirectional mics will produce a more open and transparent sound, or bidirectional ribbon mics can be used for their smooth character and warmer top end.

Remember that the closer the mic is, the more it will zoom in and focus on only a small part of the instrument's whole sound – so with close mics, you will probably need more than one, each positioned so that the particular component of the guitar's sound they are responsible for sounds great. As mic distance increases, the sound becomes more natural and balanced, which is usually desirable – but undesirable room acoustics or spill can make miking closer than ideal a necessity.

Some good single mic starting points are shown in **Figure 14.7**. These techniques can be used in isolation, or mixed together in mono or stereo:

1. For a natural balanced sound, position a mic 15 to 60 cm (6 to 24 in) away, in front of the guitar, pointing towards where the fingerboard and body of the guitar overlap.
2. A mic positioned close to the sound hole will pick up a lot of boomy bass which needs rolling off or EQ'ing. This extra boom is in addition to any proximity effect. Even an omnidirectional mic positioned close to the sound hole will not pick up much definition or string detail – but another mic can be positioned somewhere else to add definition.
3. A mic positioned further up the neck of the guitar will produce a thinner, weaker sound, with increased string, finger and fret noise, and less body and fullness. By itself this may not be appealing, but it can add zing and details to a body or sound hole mic.
4. A mic positioned beyond the bridge, down towards the larger end of the body will pick up a warm, full sound – less boomy than the sound hole, and with even less string, pick, and finger detail. This can be a good sound to blend with a mic accentuating the fingerboard components.
5. In front of the guitar isn't the only place to consider putting a microphone. A second mic positioned *behind* the larger end of the guitar body will pick up a different perspective.

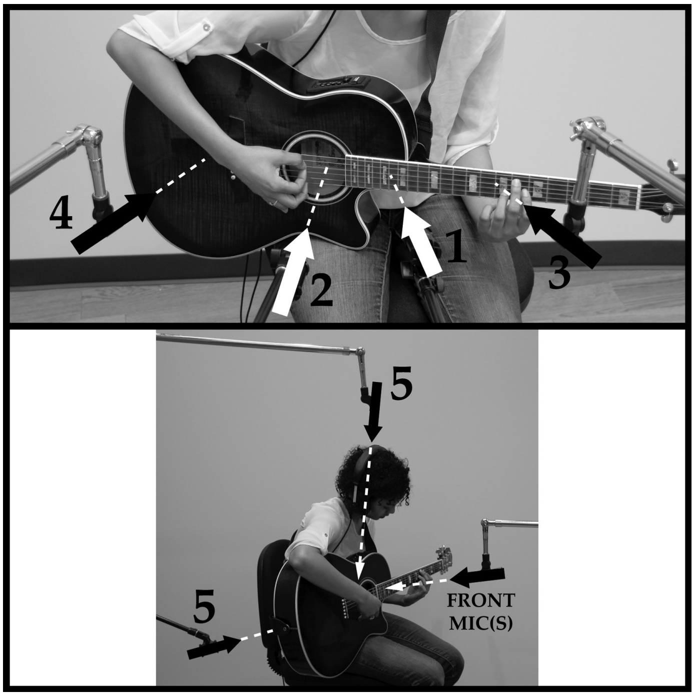

Figure 14.7 The acoustic guitar mic techniques described in the main text.

Another position that can produce great results when combined with closer mics, is above the player's shoulder, pointing down towards the instrument – this is near where the performer listens from, and evaluates their own sound.

Recording an acoustic guitar in stereo gives a single rhythm guitar track (or a solo track) width, depth, and space. Two spaced mics, one towards the fingerboard end of the guitar body, and the other just beyond the bridge, panned left and right will produce a large, wide, balanced stereo image, but any of the stereo arrays discussed earlier in this book can be used.

ACOUSTIC GUITAR PICK-UPS

Acoustic guitar pick-ups tend to produce a flat, veiled, artificial "stringy" sound, which has no depth or transparency. They are rarely used by themselves. Pick-ups can, however, be blended at a low level with an additional mic or two, to add another character component to a mono or stereo guitar sound. Comb filtering can occur between the pick-ups and mics, so do listen carefully to the combined sound, and adjust mic distances to fix any problems.

A fairly reflective room will have lots of natural reflections that are picked up by close mics. As long as the room is not too small, the sound of those reflections can really suit some musical and production styles. A dry, less reflective room, with its more in-your-face sound will suit other styles. Regardless of the general room acoustic, floor reflections are incredibly beneficial to an acoustic guitar's recorded sound – they add a layer of sharpness, impact, and excitement, particularly in small, less reflective rooms. So do try removing any carpet or rugs if there is a hard reflective floor underneath, or put some plywood over installed carpet.

AUDIO EXAMPLES
Can be found on the companion website

Acoustic Guitar Miking

Example 14.13: Acoustic guitar, a single mic positioned in the sweet-spot between the sound hole and neck of the instrument for a natural, balanced sound.
Example 14.14: Acoustic guitar, a single mic positioned closer to the sound hole.
Example 14.15: Acoustic guitar, a single mic positioned over the neck of the instrument.
Example 14.16: Acoustic guitar, two mics in similar positions to those in the previous two examples are combined and panned to produce a wide stereo guitar sound.

14.11 Grand Piano

ONLY A GOOD SOUNDING PIANO WILL SOUND GOOD!

Before even thinking about recording a real piano (grand or upright), *make sure it is well maintained and in tune!* Nothing can remove the out of tune honkey-tonk bar piano sound from a recording!

A piano is a physically large sound source that naturally takes up space. It lends itself to multi-mic techniques, to provide both multiple characteristics to blend together, and panning options that create width and fill the soundstage. A mono piano track is usually very difficult to incorporate into a stereo mix effectively – it is either panned center, where it competes with many other fundamental sound sources, or panned off to one side, where it creates asymmetry unless balanced with another sound on the opposite side of the soundstage. Either way, an instrument that is naturally very large is made to sound small and narrow in mono.

Details, dynamics, and accuracy are paramount to the sound of a piano, and being a percussion instrument there are significant transients created by the hammers hitting the strings – so condenser or ribbon mics are recommended. Cardioid, wide-cardioid, and omnidirectional mics are all options. Omnidirectional mics will produce the most open and transparent sound, but if the lid is closed or partially open, the room acoustics are not ideal, or there are other sound sources in the room, cardioid mics are preferred – so that reflections from the lid, the undesirable room sound, or spill are minimized.

THE LID. OPEN OR CLOSED?

- For rock and pop music, it is usually preferable to record the grand piano with the lid open, or completely removed.
- It is best to isolate the piano in its own room, or overdub it separately, so that the piano mics do not pick up loud spill.
- Putting the lid on its short-stick will only reduce the spill a little, and the sound picked up by the mics inside the instrument will not be as good as with the lid open.
- Keeping the lid closed decreases spill, but produces honky resonances that will be picked up by the mics inside the instrument.
- If it is not possible to isolate or overdub the piano, use mic techniques inside the piano with the lid on full stick (or even short stick), and drape comforters, duvets, or heavy blankets over the lid openings, rather than record with the lid closed. The sound will be less open, and more "dead" than without the isolating materials, but there will be less spill and the piano will be more mixable.
- Short-sticks and closed lids are sometimes necessary in live sound reinforcement situations, but they are undesirable in the studio.

There are as many different piano mic techniques as there are recording engineers! The following suggestions, shown in **Figure 14.8**, should serve as starting points.

1. A stereo array positioned 15 to 30 cm (6 to 12 in) above the hammers will produce the brightest, punchiest sound. A centered, near-coincident array produces a good stereo image, with individual notes localizing precisely throughout. The extreme high and low notes may sound a little weak due to the distance between those strings/hammers and the mic array.
2. A spaced pair over the hammers will provide more even coverage and a wider image, with a clearer transition of low notes to high notes from one side to the other (if the mics are not too close). Divide the piano keyboard into quarters, and position a mic on the division between the first and second, and the third and fourth quarters. If the mics favor a particular range of notes too much, they are probably too close.
3. A widely spaced pair of *outriggers* can be added to a coincident or near-coincident pair centered over the hammers. This will give the extreme low and high notes the same perspective and proximity as the middle notes closer to the center array, and widen the stereo image.

Figure 14.8 The grand piano mic techniques described in the main text.

4 A warmer, more resonant, but still punchy jazz sound with a wide stereo image can be achieved by adding a center panned mic over the low strings, towards the opposite end of the piano from the keyboard.

5 A crossed pair of "very-near-coincident" directional mics with just 10 to 15 cm (4 to 6 in) between the capsules, positioned mid-way down and about 30 cm (1 ft) above the strings, will pick up a full close sound with a good balance of punch and resonance. The stereo image will be wide, but individual notes will not clearly localize to discrete and precise positions.

6 If a slightly less close perspective is desired, try a pair of spaced mics positioned just inside the piano rim, 15 to 45 cm (6 to 18 in) above the strings – one towards the high strings, and one towards the low strings. This array picks up some punch in the high and mid notes, and the warmth of the mid and low notes. It produces an expansive, but imprecise stereo image – it is wide, but there is less accuracy in the localization of individual notes.

7 A very natural sound can be achieved by positioning a spaced pair of mics 30 to 60 cm (1 to 2 ft) outside the curved lid opening (with the lid attached, and on full stick), above the rim, angled down into the instrument. This image is still broad, but there is

no precise note localization (which there isn't when we listen to the piano from a few feet away anyway).

8 For the most completely natural or "classical" type of sound, any stereo array can be positioned 4 feet to many feet away from the curved lid opening, angled down into the instrument – with the lid attached and on full stick. The mics should be high enough to "look down" into the instrument, meaning that the further from the piano they are, the higher they are positioned.

9 Mics can also be positioned near the pianist – above, or to either side of their head. The performer plays so the piano sounds good to their ears – so try some mics near their ears! Some mics spaced to the side of the pianist can even be used to pick up key noise, if that natural component of the piano sound is desirable. Closer mics are often blended with these more distant mics to add some detail and definition to the recorded sound.

PRACTICAL EXERCISE

Before setting up mics, explore the piano with your ears. With the lid open, get your ears inside the piano and listen. Move your head around, trying to anticipate the sound that will be picked up by mics in different positions and at different distances. Try the following:

- ► Listen from several feet in front of the curved opening. The sound should be a natural piano heard acoustically in the room. The different sounds coming from all over the instrument have been given distance over which to blend and combine into the "whole" piano sound.
- ► Listen directly above the hammers. The sound should be typical of a modern pop sound – very bright, punchy, and percussive, with the ability to cut through a busy mix.
- ► Listen from above the low strings at the opposite end of the piano. The sound should be warmer and more resonant.
- ► Listen from directly above one of the sound holes in the iron plate of the piano. A lot of sound comes from the sound holes, but it is not usually flattering, and not usually a mic position.
- ► Listen for a more balanced sound – a blend between the punch, body, and resonance, but still with a close perspective. Try above the edge of the curved rim of the frame.
- ► Listen from under the piano – particularly if there is a hard reflective floor. There is a lot of aggressive sound down here, which is not usable by itself, but can be an excellent "grunt and power" component added to other mics.

AUDIO EXAMPLES
Can be found on the companion website

Miking the Grand Piano

Example 14.17: Grand piano, a coincident stereo pair over the hammers.
Example 14.18: Grand piano, a near-coincident stereo pair over the hammers.
Example 14.19: Grand piano, a spaced pair over the hammers.
Example 14.20: Grand piano, a mic on the low strings is added to a spaced pair over the hammers.
Example 14.21: Grand piano, a crossed pair of directional mics positioned centrally, mid-way down the strings.
Example 14.22: Grand piano, a spaced pair over the rim of the piano.
Example 14.23: Grand piano, the natural piano sound captured by a near-coincident pair six feet in front of the open piano lid.

Different musical styles call for different recorded piano sounds:

- Pop, rock, electronica, etc. – bright and punchy.
- Classical – natural, more distant, a true acoustic piano sound.
- Jazz – a hybrid of those two extremes. Warmer and rounder than a pop piano, but still with a fairly close perspective.

Different songs on the same album, by the same artist, can be given different moods and textures by exploiting different piano sounds:

- Upbeat or beat-driven tracks might benefit from a brighter piano sound.
- Ballads generally suit a mellower sound.

Even think about giving different sections of *the same song* different piano sounds!

- During the verse, use mics that provide a darker, mellower sound.
- To give the chorus a lift, and an extra dimension, add or switch to mics that produce a brighter sound.

Subtle balance adjustments are all that is needed – not night and day differences! Making sure you check mono compatibility, don't limit yourself to two mics! It's easy to end up using seven or more!

14.12 Upright Piano

The upright piano is *very* different to the grand piano. An upright will *never* sound like a grand – no matter how good it is, or what mic technique is used. Uprights always sound

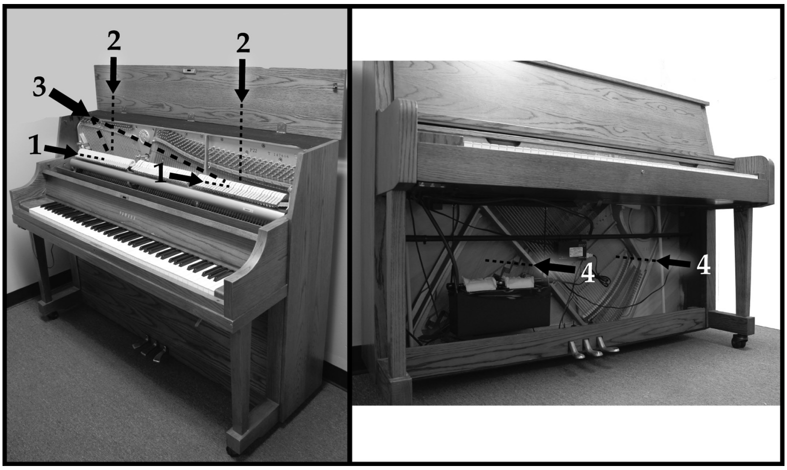

Figure 14.9 The upright piano mic techniques described in the main text. The black box in this piano is a humidity regulation system, which should be turned off while recording.

smaller and have less finesse, power, clarity, and resonance. The instrument is more difficult to mic because the strings and hammers are not as accessible as on a grand piano. Some possible mic techniques, shown in **Figure 14.9**, include:

1 With the top door of the piano removed, position a spaced pair of mics 15 to 30 cm (6 to 12 in) in front of the low and high strings respectively. This will produce the brightest, punchiest sound.
2 For a slightly less big and punchy, but still bright sound, remove the top lid of the piano (preferably also remove the front top door) and place a spaced pair 10 to 30 cm (4 to 12 in) over the opening. This technique will have a less clear stereo image than the mics in front of the hammers.
3 A more natural but distant sound can be achieved by removing the top lid and front door, and placing a coincident or near-coincident pair above the player's head, angled down towards the hammers inside the piano.
4 A warmer sound can be captured with the bottom door of the piano removed, and the lower portions of the strings miked each side of the performer.

Great results are rarely achieved by miking the backside of the piano – which is the amplifying soundboard. The sound is usually muddy and honky, and lacks attack and detail.

14.13 Electric Keyboards and Synthesizers

There are generally three categories of keyboard and synthesizer sounds:

- Simulations of acoustic instruments – pianos, harpsichords, church organs, orchestral instruments etc.

- Simulations of electronic instruments – guitars, electric basses, Hammond organs etc.
- Non-natural sounds such as synth pads and special effects.

Most acoustic instrument simulations, as well as generic synth pads and special effects are best recorded direct rather than through a miked keyboard amp. They will be cleaner, fuller frequency range, stereo, and provide the most flexibility when mixing.

Simulations of electronic instruments can often sound a bit artificial and sterile if they are recorded direct. Actual electric guitars and basses need the character of their amp and speaker cabinet to produce their full sonic character, and the same can often be said of synthesizer guitar and bass sounds. Running a bass synth sound through a bass amp or bass preamp can bring that sound to life and make it more convincing. Running a guitar sound through pedals, and/or a guitar amp and cabinet can make it more lifelike too. Equally important to realism is that the musician plays the sound like the instrument it is supposed to be – a synth guitar sound has to be played like a guitarist would actually play a guitar in order to make it convincing.

Hammond organs rely on their Leslie speaker cabinets for their sound character. So running generic synthesizer Hammond organ patch, or even a dedicated clone instrument through a miked Leslie speaker will greatly enhance the authenticity of the recorded sound – even more than using a Leslie simulator, either external or built in to the instrument.

AMP IT?

Keyboard amplifiers aren't really necessary for recording, they are more generic amplification for the gigging musician. They are not designed to impose any kind of essential character like a guitar or bass amp adds to those instruments.

But, amplifiers and speaker cabinets can be used creatively on synthesizer and keyboard sounds! Sending a keyboard piano sound into a miked guitar amp and cabinet, for example, produces unique tonal colorations and distortions that are difficult or impossible to recreate accurately with EQ and amp simulation. (A reamp box should be used between the keyboard or synthesizer's line output and a guitar or bass amp's instrument level input.)

Pedal It?

Guitar pedals can be used interactively and creatively on keyboard and synth sounds as well – with a reamp box before the first pedal input. Make a Rhodes piano sound funkier with a wah-wah pedal! Crunch an electric piano sound with some analog distortion! Make that synth guitar scream with some additional overdrive and compression! Get creative, and try all kinds of combinations, and see how much more interesting or "different" sounds can become.

Some classic vintage synthesizers and modern "retro" models aren't creatively stereo – they just use delay or chorus effects to widen the stereo image. If this is the case, turn off the built in effects, because studio effects units or plug-ins may be a better solution. If a newer synthesizer or keyboard is really stereo, *always record it in stereo* – you will have more creative and interesting mixing options, and the sounds can exploit less congested non-central spaces in the stereo image.

SOFTWARE SYNTHS?

Many keyboard players use software synths and samplers in conjunction with a MIDI controller keyboard. Others might even provide loops and pre-recorded tracks, or beats from their own computer systems. They might already have instruments and effects chained together in their software.

It is always best to have those tracks transferred into your DAW as separate tracks, and to avoid having to mix from their stems, or worse, an overall stereo mix containing every sound from their system on one stereo track. This is easy if they are using the same DAW as you, and their MIDI or virtual instrument tracks and effects can be bounced as individual audio tracks – the session can then be opened in your DAW.

Always bounce or transfer each track in stereo unless you are certain the source sound is a mono-synth with no stereo effects added to it. Bounce at the sample rate and bit depth that the studio DAW session will be using.

If different DAWs are in use, you will need to:

- EITHER: Bounce individual tracks in the keyboard players system and move the bounced files to your DAW session.
- OR: Record them in real time, from individual outputs to individual inputs, into your DAW making sure the two systems are synced as they record, or manually line them up afterwards.

To help manually line tracks up that were recorded during separate passes, it can be useful to put a single percussion hit (like a snare drum or rim shot) at exactly the same place in all tracks, before or after the actual musical content. The initial transient of this event can then be lined up to the same time location on all the tracks.

If you have to do a real-time transfer of the keyboard player's audio tracks, try to use a good quality external audio interface, and not the low quality built-in audio outputs on their computer. If the interface has multiple analog outputs, and the studio interface has multiple analog inputs, you can to transfer multiple stereo pairs simultaneously using different output and input pairs. To avoid unnecessary D to A and A to D conversion between the two interfaces, digital connections such as ADAT lightpipe (multi-channel) or SPDIF (stereo only) can be used if the two systems are running at the same sample rate.

14.14 Leslie Speakers and the Hammond Organ

Leslie speakers have a rotating horn that reproduces high frequencies, and either a stationary sideways (or downwards) firing woofer, or another rotating horn for lower frequencies. It is necessary to mic both of these. The player controls the horn rotation – they can be stationary, or rotate at various speeds, producing effects ranging from subtle tremolo through chorus, to thick grunge.

Because speakers can be very loud, and the horn's rotation causes strong air currents, dynamic mics are a "safe" close miking option. Condenser mics will produce a more open and detailed sound, but they may pick up too much mechanical noise, and are susceptible to distortion and popping because of the air currents created by the horn's movement – so windscreens or pop filters must be used. To minimize noise potential, position dynamic mics at least 10 cm (3 in) away from the horns, and condenser mics at least 6 inches away. It's a good idea to feel for where the air currents are at a minimum when the horns are rotating, and position the mics in those spots.

Some Leslie horn enclosures are open on both sides, others only on one side. The following mic techniques are shown in **Figure 14.10**.

1 For the clearest and widest swirl, try miking the rotating high frequency horns in stereo from the open side of the cabinet, with a spaced pair of mics (although any stereo array can be used).
2 For a fuller, phatter, juicier sound, but a less precise, yet more immersive width, try miking the wider baffled side in stereo.
3 For a less obvious swirl, but differently dramatic and expansive stereo image, put mics on opposite short sides of the cabinet.

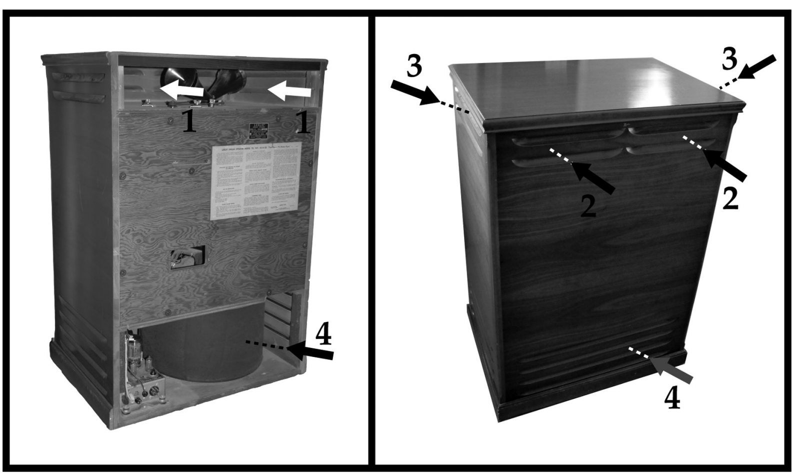

Figure 14.10 The Leslie speaker mic techniques described in the main text.

4 A side firing rotating low frequency driver can be miked with the baffle left in place, or removed. There is less of a need to mic a rotating low frequency driver in stereo because swirling low frequencies can be too distracting.

If the Leslie cabinet has a downwards firing low frequency loudspeaker, the mic should be positioned with its capsule underneath the cabinet, angled up towards the loudspeaker.

14.15 Accordions

Accordions come in different shapes and sizes. Some have a piano style keyboard and a buttonboard, and some have only buttonboards. In either case, sound comes from both ends of the instrument. A detailed recording of the accordion is usually desirable, so condenser or ribbon mics are preferable for studio recording as long as spill is not significant:

1 A single microphone can be placed about 60 cm (2 ft) in front of the accordion, at a height in the upper half of the instrument, aimed down towards the center of the instrument. Positioning the mic directly in front, or slightly over the keyboard, or slightly over the bellows produces subtly different sounds.
2 For a stereo recording, use a spaced pair of mics, each placed about 15 cm (6 in) in front of, and 15 to 25 cm (6 to 10 in) wider than the reed vents on each end of the instrument with the bellows extended, aimed towards the vents. Panning the mics hard left and right is usually inappropriate, and an exaggeration of the instrument's width, so pan them appropriately to form the instrument's place in the context of the mix you're working on.
3 If there are other sound sources in the room, a dynamic mic can be placed 10 to 15 cm (4 to 6 in) in front of the front reeds, above the keyboard, approximately one-third of the way down from the "low keys" end of the instrument.
4 Try a vertical stereo array (panned L/R) across the reed vents on the keyboard side plus a center panned mic on the bellows.

14.16 EQ Frequencies

It is impossible to tell you what EQ will be necessary! The examples below will help you start to identify frequency ranges and their characteristics. Where add or cut suggestions are given, they are fairly commonly required with close directional mics, but don't be surprised if you do end up having to do the exact opposite!

Figure 14.11 The accordion mic techniques described in the main text.

Electric Bass

- *Weight* and *phatness* – 60 to 100 Hz.
- *Body* – around 400 Hz.
- *Definition* – 600 Hz to 1 kHz.
- *Snap* and *pop* – 2.5 to 5 kHz.

Acoustic Bass

Similar to the electric bass, with the *slap* of the strings above 4 kHz.

Electric Guitar

- Reduce *boominess* – below 250 Hz.
- *Body* or *boxiness* – 400 to 800 Hz.
- *Bite* – 2 to 6 kHz.

Try separating two similar guitar tracks by accentuating different *bite* frequencies in each using *subtractive* EQ – attenuate different bite frequency ranges in each, and most importantly, make sure the guitars don't mask the vocals.

Acoustic Guitar

- *Bottom* and *weight* – around 120 Hz.
- *Body* and *thickness* – around 240 Hz.
- *Harshness* – 1 to 3 kHz.
- *Brightness* – 5 to 8 kHz.
- *Shimmer* – above 8 kHz.

Acoustic Piano

- *Weight* – 80 to 120 Hz.
- *Body* – around 200 Hz.
- Reduce *muddiness* – 300 to 800 Hz.
- Reduce *honkey-tonk* – 1 to 2 kHz.
- *Presence* – 2.5 to 6 kHz.
- *Crispness* – above 5 kHz.

Hammond Organ

- *Weight* – below 120 Hz.
- *Body* – around 250 Hz.
- *Presence* – around 2.5 kHz.

Accordion

- *Body* – around 200 Hz.
- *Reediness* – 700 Hz to 1.4 kHz.
- *Bite* – around 2 kHz.
- *Shimmer* – 4 to 9 kHz.

15 Strings, Winds, Brass, and Percussion

In This Chapter:

15.1 Orchestral String Instruments

Violin and Viola

Reflections are essential to the sound of a string instrument. It is difficult to make an orchestral string instrument sound natural if it is recorded in a room that is too small or too dead. In a room that is too small, in order to eliminate the small room's undesirable sound, the mic will end up too close – and there is no way to convincingly make it sound less close miked during mixing. A room that is too dead is an alien acoustic to most musicians, so their performance may suffer. Artificial reverbs that are more able to take a very dry sound and convincingly make it sound like it was recorded in a larger space are certainly not stock DAW plug-ins.

To pick up the most detailed, natural, and uncolored sound of violins and violas, flat condenser mics are the best choice – the presence peaks of highly characterful microphones can over-emphasize unflattering frequency ranges. The closer a mic is placed, the scratchier and more abrasive the sound is – but in order to minimize spill from adjacent instruments, sometimes closer-than-ideal miking plus corrective EQ is necessary. Some suggested mic positions include:

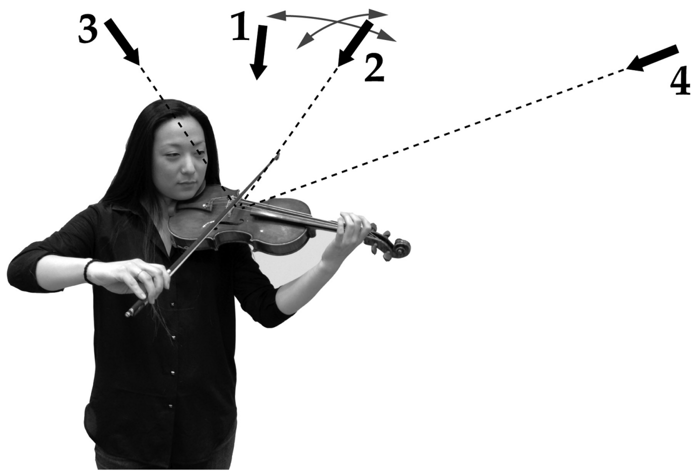

Figure 15.1 The violin and viola mic techniques described in the main text.

1 A mic positioned 45 to 60 cm (18 to 24 in) above a solo instrument (or an instrument in a section) angled towards the bridge, strings, and f-hole area, will pick up a natural but fairly close perspective.
2 Frequencies below about 500 Hz propagate fairly omnidirectionally around violins and violas, but frequencies above this are focused in a narrow 30° cone, upwards from the top soundboard. A mic directly in this beam can often sound too brittle, so to ensure you're getting the best frequency balance it is important to experiment with mic positions above the instrument.
3 An overhead mic can be placed above the performer's shoulder and head, pointing down towards the center of the instrument. This produces a slightly fuller, less edgy sound.
4 In a good sounding room, a warm and natural sound can be achieved by positioning a mic about 60 cm (2 ft) above and about 1 to 2 m (3 to 6 ft) in front of the instrument. The sound will be bright but less resonant, and there will be increased room sound compared to the closer techniques.

Blues, jazz, and rock players may have pick-ups on their instruments. These do not have the same sense of detail, transparency, or space that a well miked instrument has – but they do have a unique sound that may be stylistically appropriate. Pick-ups can be DI'd and blended with a mic – the technique is similar to the bass mic plus DI technique discussed in a previous chapter. However, due to the smaller wavelengths of the higher frequencies of the violin and viola, phase cancellation and comb filtering between the DI and mic is more significant and needs to be minimized with careful mic positioning.

WHICH POLAR PATTERNS TO USE

Wide-cardioid and omnidirectional mics will pick up the most natural instrument sound. Their wider pick-up and reduced proximity effect allow closer mic positions to capture a more natural sound than a cardioid or hyper-cardioid mic. They also beneficially pick up good room sound. If the room is too small or too reflective, cardioid mics may be necessary to reduce this ambience – but more directional mics need to be positioned further away in order to pick up a natural sound, and in doing so will end up picking up more colored off-axis room sound. Closer-than-ideal directional mic placement is necessary if there is too much spill or bad room sound – but the instrument's sound will be compromised and require corrective EQ.

Cello

Much of a cello's characteristic sound comes from its interaction with the wood floor or the wooden sound box that cellos (and basses) are sometimes played on. The wood floor or sound box amplifies the sound waves transmitted through the metal spike at the base of the instrument, to produce a resonant, full bodied sound. Positioning a mic closer than about 60 cm (2 ft) away doesn't pick up this passively radiated and reflected content. For a natural sound, a mic positioned up to 2 m (6 ft) away would be entirely appropriate if spill from adjacent sound sources is not problematic. The beauty of a cello's sound is in the subtleties and details – so flat response condenser microphones are a good choice.

1. Positioning a microphone too close, directly in front of a cello often produces a boxy, nasal, or scratchy sound.
2. A close mic positioned on the high string side of the instrument, slightly off-center to about 40° to the side, picks up a smoother, more balanced sound.
3. A mic positioned up to 2 m (6 ft) away, about 60 cm to 1 m (2 to 3 ft) high, angled below the bridge can pick up a combination of instrument and reflected sound. Different timbres can be picked up – body and floor reflections can be increased by aiming the mic further down the instrument, and the string and definition can be increased by aiming the mic further up the instrument. If the sound directly in front of the instrument is too thin, try moving the mic to the side in 15 cm (6 in) increments – and compare both sides to see which side sounds best.
4. A mic fairly high above the instrument, a little in front of it, pointing down will capture a sound more similar to the sound the player hears.
5. A cello is a relatively wide baritone source, and can also be spot or solo miked in stereo – using an array 60 to 90 cm (2 to 3 ft) in front of, and about 75 to 90 cm (2 ½ to 3 ft) high, aimed down into the instrument. Coincident, near-coincident, spaced AB, or MS techniques can be used. A spaced pair gives you creative control over the image width

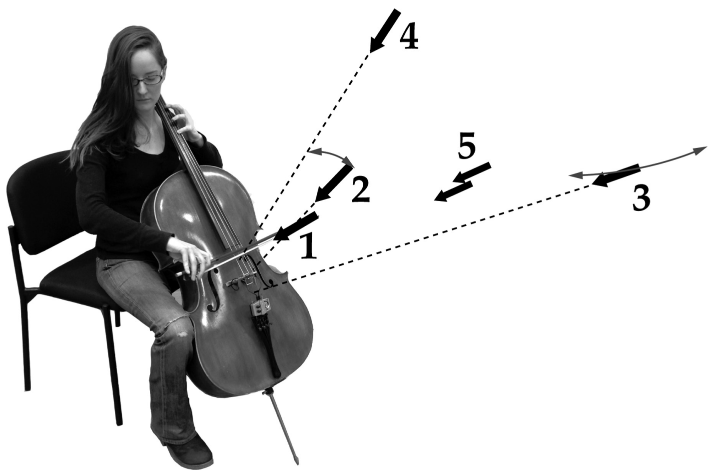

Figure 15.2 The cello mic techniques described in the main text.

via mic spacing – which is important, because you do not want the instrument to take up the entire width of the soundstage. Rather than spacing omnidirectional AB mics widely and panning them narrowly (which is prone to phase problems), space them about 15 cm (6 in) apart, and pan them almost or hard left and right for an appropriate width.

For isolation and spill control purposes, it is sometimes unfortunately necessary to mic closer than 60 cm (2 ft) away. If this is necessary, be sure to keep the mic out of the performer and bow's way – and experiment with mic position to obtain a sound that will require the least amount of corrective EQ.

String Bass

Close miking techniques for string bass are discussed in more detail earlier in this book.

For an acoustic, orchestral bass sound which has a fairly close perspective:

- Find a natural sounding sweet-spot about a meter (3 ft) away.
- Add some bottom and roundness with a closer mic on the best sounding f-hole.

For a less close perspective:

- A mic can be positioned up to 1.5 to 2 meters (4 to 6 ft) away, 1 to 1.5 m (3 to 4 ft) above the floor, angled towards the body below the bridge of the instrument – so that it picks up a blend of the direct instrument sound and floor reflections. Experiment with moving the mic to the sides about a foot at a time to find the sweet-spot that produces a warm and defined sound.

- An f-hole mic 30 to 60 cm (1 to 2 ft) away can add some low frequency grunt that might be missing from a more distant mic.

THE GOOD AND BAD OF CLOSE MICS

Good: For non-classical projects, the intimacy, "bow on string," plucking details, and reduced spill of close miking can be desirable – but you should experiment, moving the microphone a few inches at a time, to find the sweet-spot where the sound requires the least corrective EQ.

Bad: Close mics will also pick up more undesirable performer noises – fingers on the fingerboard, breathing, clothes rubbing, etc.

15.2 Horn Section Instruments

Trumpet and Trombone

Trumpets and trombones are loud instruments – a trumpet can produce over 130 dB SPL close to the instrument's bell! Trumpets and trombones radiate a lot of their high frequency content very directionally, so positioning the microphone slightly off-axis will keep it out of the loudest SPLs, and mellow an otherwise overly bright or harsh sound. Position a trombone mic to the player's left to keep it out of the way of the instrument's slide. Use a small diaphragm condenser mic with high SPL capabilities for the most accurate recording, a large diaphragm condenser mic for more warmth and character, or a ribbon mic for a smoother and less harsh sound. Use the pad on the mic to prevent mic distortion when close miking.

1 For a bright contemporary sound, position a mic about 30 cm (1 ft) away, up to 35° off-axis, angled into the bell.

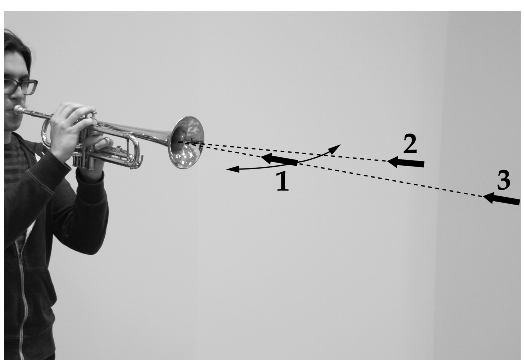

Figure 15.3 The trumpet mic techniques described in the main text.

2 For a more balanced, but still in-your-face sound, position a mic slightly off-axis, at a distance of about 1 m (3 ft).
3 In a good sounding room, without adjacent sources of spill, increase the mic distance to 2 to 3 m (6 to 9 ft) for a more classical sound, with de-emphasized brightness, and added warmth and body.

CLOSE MICS ON BRASS

Dynamic and condenser mics (with high SPL capabilities) can be positioned on-axis, less than 30 cm (1 ft) from trumpet and trombone bells. The sound will be:

- Brighter than we usually hear acoustically.
- Contain little buzzy details and breath and spit noises that we don't hear from a distance.
- Have increased amplitude spikes compared to more distant mics.

Like a close miked lead vocal, a track recorded this way will need more compression, EQ, and automation during mixing – but spill or room sound will be minimized, and the sound will be more "in-your-face."

Accuracy or Warmth?

Trumpets and trombones have huge dynamic ranges and contain large waveform peaks relative to their average level. Some tube microphones and preamps, or for pop and rock, high quality flat response dynamic mics, will "eat" some of these peaks, smooth the waveform, reduce crispy brightness, and produce a warmer, fatter, but less detailed sound than condenser mics.

Saxophone

Saxophones come in different shapes and sizes, the most common being the soprano, alto, tenor, and baritone. Even though the instruments have a shiny bell similar to brass instruments, sound does in fact come from the holes *all over* the instrument – depending on which holes are open or closed as the instrument is played. Large diaphragm condenser mics are a great choice for detail, power, and character. Ribbon mics can also sound great. Remember that the closer the mic, the more it will zoom in on just part of the instrument's sound, and accentuate breathing, performer noise, key clicks, and other instrument noises.

1 One place a mic should *not* be pointed is directly into the bell – the sound is thin, honky, and inappropriately loud when the lowest few notes are played.

2 For a close miked perspective of a curved instrument, a mic can be placed from 15 to 30 cm (6 to 12 in) in front of, and slightly to 30 cm (up to 8 in) above the bell, pointing back *past* the top of the bell towards the keys on the instrument – imagine the mic's pick-up pattern sweeping across a good portion of the horn, and not just into the bell. The higher the mic, the more it should be angled down towards just above the top/back of the bell. Higher mic positions can favor the "left hand notes" of the instrument, while lower mic positions can get more "honky" on the low notes.

3 For a more natural perspective suited to acoustic jazz and classical recording, take the longest dimension of the instrument, and position the mic about that far away (possibly just a little closer) and anywhere from a little above the bell to not much higher than the left hand thumb key on the back of the instrument. The higher the mic, the more it should be angled down towards just above the top/back of the bell.

4 Straight instruments require two mics for close miking. One should be positioned about 30 cm (1 ft) in front the instrument, angled towards a point 1/4 to 1/3 up from the bell of the instrument – this will pick up warmth and body, but lack brightness and the lowest notes. The second mic should be placed about 15 to 30 cm (6 to 12 in) directly below the bell, pointed up towards it. This mic will pick up a very thin, nasal, buzzy sound that is nasty by itself, but blends well with the other mic, and gives the lowest notes balance.

5 If only one mic is available for a straight saxophone, it should be angled towards a point 1/4 to 1/3 up from the bell, at a greater distance, 18 to 24 inches away (45 to 60 cm) to allow some floor reflections from the bell to bounce into the mic to brighten the sound and fill in the lowest notes.

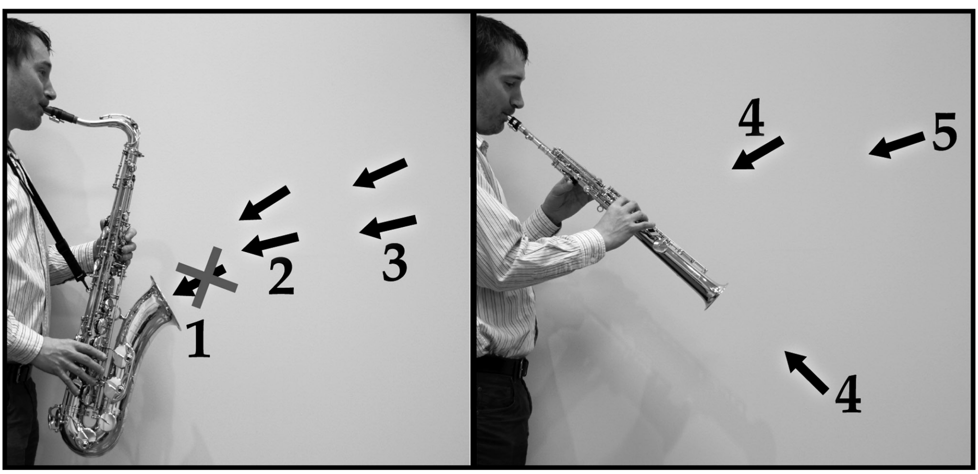

Figure 15.4 The saxophone mic techniques described in the main text.

LIVE SOUND VS RECORDING

In live concert sound, mics are frequently placed closer than described in this book. This is to:

- ► Maximize the amount of amplification possible before feedback – by getting as much sound into the mic as possible.
- ► Minimize the amount of spill from other sound sources on the stage.

When miking so close, the sound *is* compromised – it is not the best or most natural sound. Extensive corrective EQ is required. This is a necessary trade-off in order to get the show loud enough, and have individual control of each instrument in the mix. The recording studio is a much more controlled environment in which the goal is to capture *the best sound*, not just *the most* sound.

Flute

The sound of a flute comes from all over the instrument, depending on the note that is being played. A good quality condenser mic will best pick up the subtle details of a breathy flute sound.

1 For a natural, classical sound, a mic 1 to 2 m (3 to 6 ft) in front of, and slightly above the performer's head, angled down towards the center of the instrument is a good starting point.
2 The flute can also be miked from slightly behind the performer, at head height, to pick up something similar to the sound the performer hears.
3 For a closer jazz or pop perspective, a mic can be placed about 30 cm (1 ft) in front of, and slightly above the instrument, mid-way between the mouthpiece and left hand keys – avoiding the wind currents and breath noise produced when the player blows across the mouthpiece. Omnidirectional mics are less troubled by wind currents and pops, so are a good choice for this position.
4 For a stereotypical rock or jazz flute sound, a mic (dynamic or condenser) *with a windscreen* can be positioned within a few inches of the mouthpiece, just below the strong wind current coming from the player's mouth.

Stereo miking is also an option. A coincident or near-coincident array can be positioned similarly to the third option above, but halfway down the instrument.

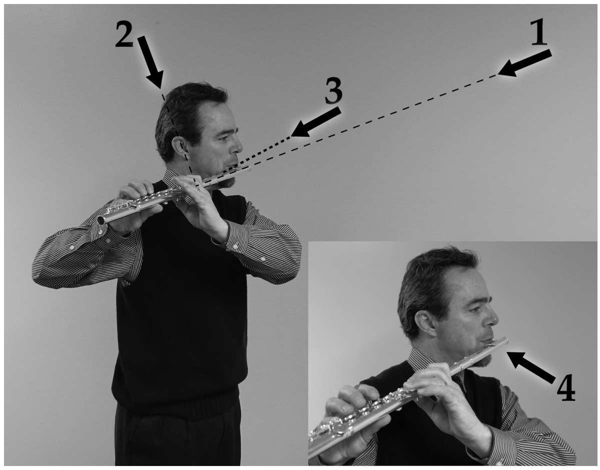

Figure 15.5 The flute mic techniques described in the main text.

15.3 Other Wind and String Instruments

Banjo

1 A close miked perspective of a banjo can be achieved by positioning a mic 10 to 30 cm (4 to 12 in) in front of where the neck and body of the instrument meet – move the mic around and listen for the sweet-spot where there is a good blend of body, resonance, stringiness, and fingerboard sound.
2 Positioning a mic further up the neck produces a thinner, more metallic, buzzy sound.
3 A mic behind the bridge will produce a thicker, fuller, more resonant sound, with a dark "pluck."
4 A coincident or near-coincident stereo array can be positioned in the sweet-spot.
5 A pair of spaced omnidirectional mics about 30 cm (1 ft) from the instrument are also a good stereo option – when transparency is desired, and the room sound is good.
6 Raising the mic above (and in front of) the instrument, and angling it down towards the body will increase the floor (and possibly room) reflections picked up – producing a more vibrant, exciting sound.

Mandolin and Ukulele

Mandolins and ukuleles are small and quiet instruments, so miking options are a little more limited than for guitars and banjos. To accentuate the body of their light sound, angle the mic towards the sound hole regardless of where it is positioned. A good starting point is 10 to 15 cm (4 to 6 in) in front where the neck and body meet.

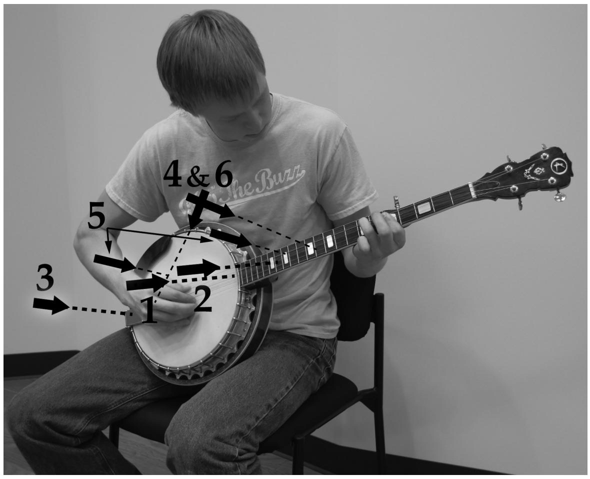

Figure 15.6 The banjo mic techniques described in the main text.

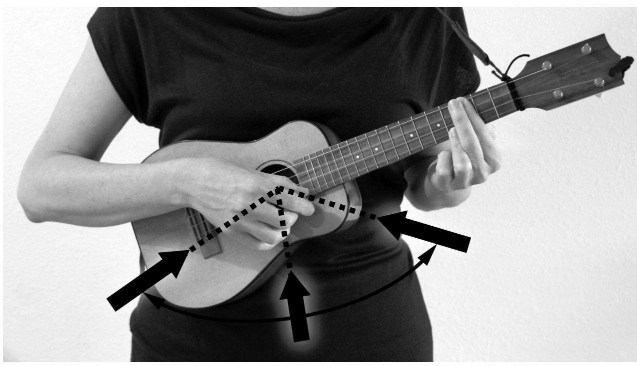

Figure 15.7 Angle a ukulele (or mandolin) mic towards the sound hole regardless of where it is positioned.

Dobro (Resonator Guitar)

The dobro is frequently played like a slide guitar, flat on its back on the player's knees. Try positioning a mic about 30 cm (1 ft) above and to the left or right of the metal resonator, angled into the resonator. There will be more body and "pluck" from the bridge side, and more bright stringiness from the neck side. If the dobro is isolated, and the room has good characteristics, pulling the mic a little further away will give the sound a more transparent and natural proximity and tonality.

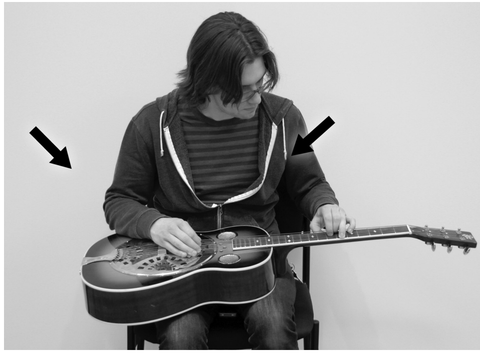

Figure 15.8 Angle a dobro mic towards the resonator. Different sounds can be achieved from each side of the resonator. Keep the mic on the neck side out of the player's breath currents by positioning it slightly in front of the instrument.

Bassoon

As with most other woodwind instruments, much of a bassoon's sound comes from each note's lowest open hole, and not the bell. Bassoons are large, and different sounds come from all over the instrument. Miking close presents challenges.

1 A more distant, natural sounding approach is to use a condenser mic with a wide pick-up pattern, positioned 1 to 2 m (3 to 6 ft) away, at about the player's eye level, angled slightly down so the pick-up pattern sweeps across the entire instrument.
2 For a closer perspective, a mic can be positioned about 30 cm (1 ft) away, slightly to the outside of the instrument, a quarter to a third from the bottom of the instrument, and blended with a second mic above and slightly in front of the bell, angled down so that it picks up the bell and lower notes.

Omnidirectional mics will pick up the beef and resonance of the lowest frequencies created by bassoons and contra-bassoons best – as long as the room reflections sound good, and spill from adjacent sound sources is not problematic.

Clarinet

The approach to miking a clarinet is similar to miking a straight soprano sax, ideally with two microphones. The first mic should be positioned about one-third of the instrument's length from the bell, about 30 cm (1 ft) away, angled towards the finger holes on the bottom half of the instrument. If only one mic is available, this is a good position to start with, but the mic distance should be increased to allow more floor and room reflections to be picked up. The second mic should be positioned about a foot below the bell of the instrument, and blended with the other mic to add brightness.

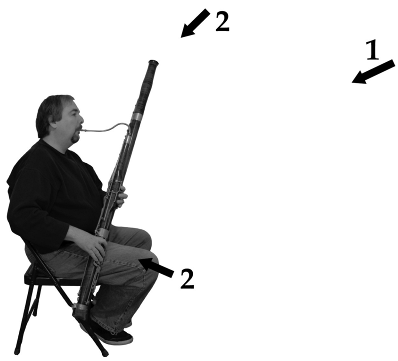

Figure 15.9 The bassoon mic techniques described in the main text.

Oboe and English Horn (Cor Anglais)

For a natural perspective, take the longest dimension of the instrument and position the mic that far away, pointed one-third of the length from the bell. A closer perspective can be achieved with less mic distance – but no less than 15 to 20 cm (6 to 8 in) – pointed at the same finger holes. If the sound is a little dull and unexciting, add a second mic under the bell, similar to that described for the clarinet and soprano sax.

French Horn

The French horn is unique in that its bell faces backwards. In orchestral settings the reflective wall behind the performers, or reflective panels on stage, reflect the instrument's sound back towards the audience. Unlike other brass instruments, the player's hand is typically placed inside the bell to either gently mute the sound, or give it a piercing nasal quality. The high frequency radiation pattern from the bell is naturally quite narrow, but it is diffused more widely by the hand in the bell.

A mic positioned 30 to 60 cm (1 to 2 ft) behind the player, slightly off-axis to the bell will accurately capture the sound. If there is a reflective wall behind the instrument, a bidirectional or omnidirectional mic will pick up a blend of direct and reflected content (which could cause phase problems).

Harmonica

1 The classic blues harmonica mic is the Shure "Green Bullet," cradled by the player's hands. A decent dynamic vocal mic (because of its built in windscreen) can be used to get a similar hand cradled sound.
2 A stand mounted mic can be positioned immediately in front of the player's hands, which are then free to be opened and closed to create timbral effects.

Figure 15.10 Miking a French horn.

Figure 15.11 The harmonica miking techniques described in the main text.

3 For a lighter, more transparent sound, with a less close perspective, a condenser mic with a pop filter or windscreen can be positioned about 10 cm (4 in) in front of the harmonica.

Harp

Harps need to be recorded in good sounding large rooms, and the mics need to pick up some room reflections. In small rooms harps become too bassy. They are an orchestral instrument designed for orchestral performance spaces!

1 For a natural perspective, a mic (or stereo array) can be placed about 2 m (6 ft) away, directly in front to 45° to the side opposite the player, 90 to 120 cm (3 to 4 ft) high, angled into the soundboard of the instrument.

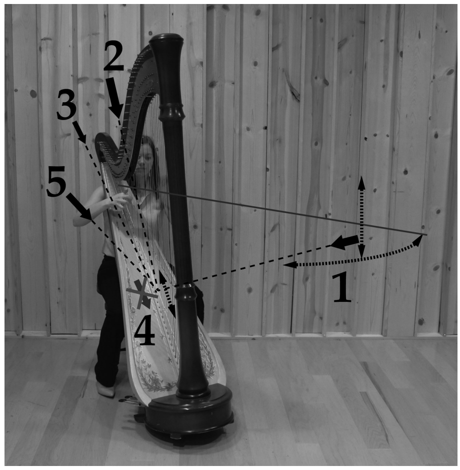

Figure 15.12 The harp miking techniques described in the main text.

Closer mics can be added – or used alone if more isolation is necessary:

2 Brightness and presence will be present in a mic positioned at least 30 cm (1 ft) to the side of the top of the pillar (opposite the player), angled down towards the soundboard.
3 A fuller sound can be achieved by positioning a mic level with the top of, but at the rear of the instrument, on the opposite side of the harp to the player's head, angled down and forwards into the strings and soundboard.
4 Avoid placing a mic too close, directly in front of the soundboard – it will be too muddy and bassy.
5 Instead, for a close and beefier sound, position a mic at least 45 cm (1.5 ft) to the side and 45 cm (1.5 ft) above the soundboard, pointing towards the base of the strings in the center of the soundboard, on the opposite side of the instrument to the player.

Tuba and Sousaphone

In terms of range and musical function, the tuba and sousaphone are similar, but a tuba's bell points upward, and a sousaphone's bell points forwards.

1 For a natural tuba sound, a mic can be positioned about 60 cm (2 ft) above the bell. On a sousaphone the mic should be positioned in front of the bell.
2 A slightly off-axis position produces a less buzzy, warmer sound.

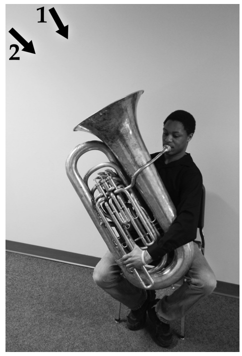

Figure 15.13 Miking a tuba.

These are not close mic techniques, so condenser or ribbon mics are preferable to dynamic mics. An omnidirectional mic will pick up the very low fundamental frequencies these instruments can produce better than a directional mic.

15.4 Percussion Instruments

Mallet Percussion Instruments

The marimba and vibraphone are both large horizontally wide instruments. Their sound can effectively use the width of the stereo soundstage. A common approach is a pair of condenser mics spaced about 60 to 75 cm (2 to 2.5 ft) apart (or dividing the instrument into thirds), about 60 cm (2 ft) above the instrument's bars. Other stereo arrays can be used – but a spaced array generally provides more even coverage of these wide instruments.

A xylophone is a little smaller, so coincident and near-coincident pairs are also practical options. If a spaced pair is used, the mics should be spaced a little over 30 cm (1 ft) wide, and 20 to 30 cm (8 to 12 in) above the instrument.

Hand Percussion Instruments

Shakers, claves, tambourines, triangles, and other small instruments can generally be miked with a single small diaphragm condenser mic (for the best detail and high frequency definition) positioned 15 to 45 cm (6 to 18 in) away, maybe slightly above the instrument, angled into it. Sometimes, close miking these instruments makes them hard to fit "back" in a mix. Increasing the mic distance to several feet (or more) captures them at a perspective they might naturally want to be at in the mix, rather than having to push a close miked sound back electronically and artificially using EQ and effects.

Figure 15.14 Stereo microphone techniques directly above a vibraphone's bars. You would only use one of these techniques at a time. **1:** A spaced pair. **2:** A near-coincident pair (or a coincident pair). **3:** A slightly wider spaced pair, angled in towards the center of the instrument will produce a wider image – but listen carefully for a hole in the middle of the image.

Larger membranophone instruments, including conga and bongo drums can be approached similarly to snare drums and tom toms for a close-up perspective. Use a close stereo array, or a mic on each drum to give yourself cool panning options. For a less close perspective, that naturally mixes itself further back in the stereo soundstage (where these drums can be very effective), use a coincident or near-coincident stereo array 30 to 100 cm (1 to 3 ft) away and slightly above the instrument.

15.5 EQ Frequencies

It is impossible to tell you what EQ will be necessary, so the center frequency examples below are just suggestions to help you identify the frequency ranges of different characteristics. The specific mic, mic technique, room, instrument, and the way it is played are all unknown variables that will alter what EQ might be necessary.

- 50 Hz: *Woof*. Adds fullness to low frequency instruments. Reduces excess boom on low frequency instruments.
- 100 Hz: The *boom-box fake-bass* range. *Boomy*. Adds fullness and warmth to some lower instruments. Attenuate to reduce boom and muddiness in others.
- 200 Hz: Low-mid *mud*. Adds fullness to some instruments. Attenuate to reduce muddiness (therefore increase definition) in others.
- 400 Hz: *Hollow* sound. Reduces *boxiness*.
- 800 Hz: Reduces *honk*.
- 1 kHz: *Megaphone*, *bullhorn*, or *AM radio* character. *Nasal* thinness.
- 2 kHz: Just plain *nasty*. The *ear fatigue* range. Increases *diction* and *intelligibility* of a sound when boosted. Decreases *clarity* and *bite* when attenuated. Increases or decreases *attack*, *snap*, and *punch*. The lowest of the *presence* frequencies.

- 3 to 4 kHz: *Trashy, ugly, not-quite-high-frequencies*. Increase or decrease *definition, cut, snap,* and *presence*.
- 5 kHz: *Attack, punch, harshness, presence*.
- 7 to 8 kHz: *Brightness, sibilance, sizzle, shimmer. Scratchy* or *harsh* in some sounds.
- 10 kHz and above: *Air, breath, float, lightness*.

16 Setting Up the Studio

In This Chapter:

16.1 The Three or Seven P's

> **P**re-**P**roduction **P**lanning. Or more assertively – **P**roper **P**rior **P**lanning **P**revents **P**iss **P**oor **P**erformance!

Before a recording session, you need to research the band or musicians you are recording:

- What is their musical style?
- How many musicians are there?
- What instruments do they each play?
- What is their "sound"?
- How do their previous recordings sound? How do they want this one to sound?

Another very important question to consider is:

- Does the band have studio quality instruments?

The recording process will put their performance and equipment under a microscope and magnify any problems with an instrument or amp cabinet's sound. Some studios carry drum and amp inventory, so if there are any band equipment problems they can arrange to use those (if they haven't sourced alternatives prior to the session). Some studios have rental arrangements with local companies. If the client, or you as the engineer or producer, decides it is necessary to rent gear, make sure it is adequately provided for in the project budget.

Questions you need to ask include:

- How will the session be set up, and how will it be run?
- Does the available equipment inventory (mics, stands, preamps, inputs, etc.) support the ideal approach?
- Do the physical facilities and limitations of the studio support or impact the set-up and approach?
- Will any compromises be needed? If so, ensure they will not negatively impact the recording process or end result.

Having answered those questions, you can decide how many mics, which mics, and what specific mic techniques you will use. You should create an "ideal" *input list* and *room plot*. You can't make any final decisions until you hear the instruments in the studio, but having a strategy enables you to efficiently start work, and appear relaxed and professional. This sets the tone for the session and puts your clients at ease. The input list and room set-up will probably change, and that is perfectly normal.

Be sure to keep good notes, and document any changes. Weeks or months from the session, you want to be able to remember what you did, what worked well, and what could have been better so you don't do it again! Extensive notes, including exact positioning of mic stands and the height and angle of the mics, are essential in case it is necessary to re-record anything at a later date. Take pictures of mic and instrument set-ups with a phone camera and store the files with the rest of the session documents. This is a good job for an assistant, intern, or a knowledgeable and reliable "extra" at the session!

HOW GOOD ARE YOU AS AN ENGINEER?

The saying goes:

> *"You're only as good as your last project…"*

But that's not true. Every project you do is advertising for future business, and part of your portfolio. You never know which of your (older) projects clients or prospective clients will hear, or when and where they will hear them! The saying should be:

> *"You're only as good as every project of yours that somebody might hear."*

INPUT LIST

PROJECT: ____________
DATE: ____________
ENGINEER: ____________
PRODUCER: ____________

INSTRUMENT	MIC	PREAMP	INPUT No.	PROCESSING	NOTES
Kick	D6	API 1	1		Short Round Base
Snare Top	SM 57	NEVE 1	2	2:1 Comp	Medium Boom
Snare Bottom	Beta 57	NEVE 2	3	2:1 Comp	Medium Boom
Hi-Hat	M44	NEVE 3	4	2:1 Comp	Tall Boom
Rack Tom 1	M88	NEVE 4	5	2:1 Comp	Tall Boom
Rack Tom 2	M88	NEVE 5	6	2:1 Comp	Tall Boom
Floor Tom	MD421	NEVE 6	7	2:1 Comp	Medium Boom
Overhead L	414 XLS	NEVE 7	8	2:1 Comp	Large Atlas
Overhead R	414 XLR	NEVE 8	9	2:1 Comp	Large Atlas
Bass Mic	MD421	API 2	10		Short Round Base. Gobo Around Amp.
Bass DI	Hot Box	API 3	11		
Guitar	R 121	SSL 1	12		Medium Boom. Gobo around Amp.
Synth L	Hot Box	SSL 2	13		
Synth R	Hot Box	SSL 3	14		
Sax	U87	GRACE 1	15		Large Atlas
Scratch Vocal	SM58	API 4	16		Tall Boom
Vocal	U87	API 1	1		Large Atlas
Vocal Doubles	U87	API 1	1		Large Atlas

recording and sound reinforcement

Figure 16.1 An input list for a recording session.

16.2 Bands – Small Room Recording

In a recording studio with a smaller recording space, there is probably not enough room to set the entire band up at once. Even in a medium-sized room the set-up may be cramped, uncomfortable, and not conducive to a great performance. If it's too cramped, mics will end up too close together, resulting in too much spill and less than ideal tracks. Better results in a small room studio are usually produced by overdubbing each instrument (or voice) separately.

The benefits of this layer-by-layer approach include:

- Better isolation between the instruments produces a cleaner mix. Balance, EQ, and stereo image accuracy can all be more precisely controlled, and predictable results can be more easily obtained.
- If recorded in a relatively dead room, artificial reverbs and effects can be controlled during mixing – the recording room itself does not impose significant acoustics of its own on all the tracks.
- Not all of the musicians need to be in the studio simultaneously. This may be helpful when scheduling busy musicians.
- As the engineer, you only have to set up a small number of mics and headphone mixes at any one time. For a novice engineer this can make a session more manageable and approachable.

CHECK EVERYTHING!

It is imperative that each microphone is picking up good sound – and there are no problems such as distortion, crackly cables, or too much spill, etc. The more mics there are in use at any one time, the easier it is for an issue to go unnoticed until after the musicians have left the studio – by which time it is a major problem. Use the solo buttons on the console or DAW to check each mic sounds good, and watch all preamp *and* recording levels continuously.

The drawbacks of this small room approach include:

- The musicians have to play in smaller units or individually – not together. This can be difficult if they're not used to doing so, or do not know the songs well enough. Some bands need to play together in order to groove properly – so the musicality of the recording could be impaired.
- If the only room available is relatively dead, some instruments will lack the life and punch provided by a more acoustically live environment.
- Distance and room mic creativity is limited in smaller rooms.
- The performers rely entirely on the headphone mix in order to perform, so it is important that you can provide the mix they need, and make adjustments quickly and correctly.

The usual approach when using this recording method is to start with the rhythm section:

- Create a click track if necessary, and make sure the drummer plays to it. Most DAWs can do this as a real-time plug-in, or alternatively you can record the output of a metronome or other sound source the band provides to a spare track. A loop from a drum machine or a drum loop sample can also be used, but make sure the loop isn't changing the feel of the recording – if the band is not comfortable with it, get rid of it! Also, make sure you mute any click track as soon as the drums are recorded and it's no longer necessary for musical time keeping. It's easy to get used to a click being there, becoming part of the song – and then you miss it and the song sounds empty when it's turned off!
- Record the drums first. The drummer will need a *guide*, or *scratch* track or two to play along with. These scratch tracks could be a hand held vocal mic, and a guitar or bass DI, performed in the control room. Scratch tracks are just there so the musician being recorded knows where in the song they are, and can deliver their part confidently and correctly. The scratch tracks will be replaced with better quality recordings later on.
- Record the bass while monitoring the previously recorded drums and scratch tracks. It's usually better to mute any scratch version of the same track being recorded – it will make it more difficult for the musician to hear their actual performance, and mask musical and technical problems in the track.

- Record the guitars, keyboards, and other rhythm section instruments – in whatever order makes most sense for the musicians.
- Record the vocals.
- Record any horn and solo parts. (These last two steps may be reversed.)

DRUMS FIRST…

A click track provides macro-time information – bars and beats. The drummer supplies micro-time information – beat subdivision nuances that create *groove*. The other musicians lock to both of these when they play, so it is important to record drums first. It is *very* difficult for any other instrument to supply groove first, and then for a drummer to gel with them afterwards. Even with a click track it rarely sounds convincing, and the energy of the song, and tightness of the band are usually compromised.

Another approach to small room recording is to DI everything it is possible to DI – using amp simulators if necessary, and to have the piano player play an electronic keyboard or synthesizer. This allows more of the band to be in the recording room with the drum set (if the room is big enough) or in the control room, playing live together, and gelling and grooving better. Although these DI tracks may be compromised in sound, they can be used for the final mix, or some (or all) of them can be treated as scratch tracks and replaced with properly miked tracks in a similar process to the one described above. The difference is that having recorded the whole band in the first place, any overdubbed musician has the "band as a whole" scratch tracks to play along with, so will probably be more comfortable and perform better.

TALK MICS

In both small and large room sessions, give each performer a "talk" mic if they don't already have a vocal mic. Drums, speaker cabinets, and horns are loud sound sources, requiring their mic's gains to be relatively low – so they will not pick up quiet talking well. Any kind of vocal dynamic mic can function as a talk mic – plugged into a spare mixer channel, or interface input routed to either a DAW aux track or audio track set to "input monitor." The mic will not be recorded, but it will give the performer(s) the ability to talk to the control room, and to each other's headphones.

16.3 Bands – Large Room Recording

Bands usually play better when they perform together and can hear all the parts they are used to hearing – the musicians can gel and groove, instantaneously feeding off each other's energy and musicality. A band-at-once session can be daunting. It's a lot of simultaneous inputs and headphone mixes! Make sure you take the time to check the sound and input levels of each mic. In addition to checking channels individually, also confirm that multi-miked instruments sound good when all their mics are combined.

The benefits of recording most of the band at the same time include:

- The band locks together more tightly, and the feel of the performance is usually much more organic and natural.
- It takes less time to record, because more instruments are recorded at the same time. This can be advantageous in terms of scheduling and budget.
- The spill (if properly controlled to sound as good as possible, and/or minimized) can glue the sounds together, producing a more organic sound, less separated than if each instrument was recorded in isolation.

The drawbacks of this approach include:

- More mics and inputs are necessary.
- Spill and isolation are harder to manage.
- It takes more time to set up the recording equipment and sound check the entire band.
- More experimentation is necessary to get mic positioning optimized for sound sources and spill.
- There are a lot more channels happening simultaneously, so a lot more things to continuously check for problems.
- Distance and room mics are impractical because they become "everything" mics.
- The producer has to pay attention to more musical parts and events simultaneously, ensuring the performances are correct and as good as possible. With so many things happening at once it's easier for mistakes to slip through the cracks, only to be noticed later.
- Headphone mixes can be more of a challenge because each performer may need different things in their headphones – "a one mix fits all" approach usually never works. Multiple headphone mixes can be set up on the mixer or in the DAW using auxes if multiple independent headphone outputs are available, or the performers can dial in their own mixes if a dedicated headphone monitoring system is being used.

Setting Up

A bad room, or poor positioning of a sound source in a room, can impair the sound of the recording. Instruments can spring to life in a good position in a great sounding room. One approach to setting up in a large room is to:

- Find the optimum place for the drums.
- Set up the other instruments around the drums.

Drums generally sound best located somewhat centrally in a room, but trial and error is really the only way to optimize their position. One technique to streamline the process is to set up only the snare drum or kick drum. Position it in different places in the room, and listen to it from the performer's perspective. Where in the room does it sound best/biggest/liveliest, etc.? Once this has been decided, put a mic on it and verify the sound is what you want. Then set up and mic the rest of the drums. Be sure to listen for room resonances from potential overhead and room mic positions while doing this, and move the set-up to avoid them.

It's often impractical to put the drums exactly in the center of the room, and it's necessary to move them back slightly towards one wall in order to fit the other musicians into the room comfortably. If this is necessary, try to position the drums closer to a longer wall and keep them centered between the farthest spaced walls.

Have mics for the other musicians on stands, and cables ready to go – but don't try to anticipate exactly where the other musicians will want to set up unless this has been discussed during pre-production meetings or conversations. Their place in the room(s) will be a combination of making them comfortable (most important), and what you, the engineer

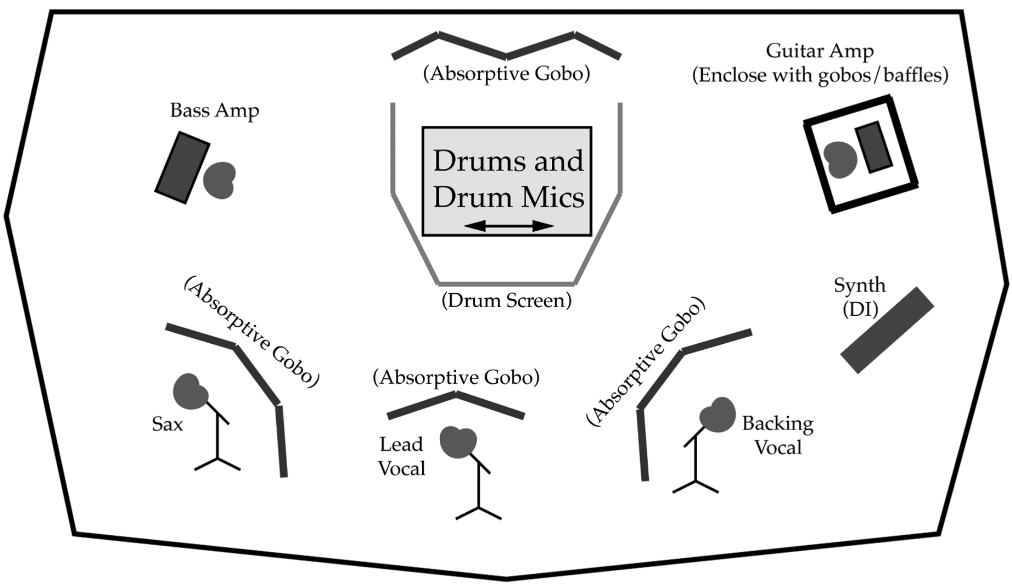

Figure 16.2 A band set-up in a large room. The musicians all face the center so they can see each other. Mic nulls are shown aimed at the drums. A transparent drum screen/absorptive gobo combination can be used around the drums to reduce drum spill into other mics in the room. Absorptive gobos can be positioned around the quieter sources to reduce spill into their mics. They can also be used behind (and even around) the singer when overdubbing the actual vocals to reduce the large room sound in the vocal mic if necessary.

need (secondary). Try to position sound sources and their respective mics to minimize spill – and choose mics with polar patterns that best do this. Think about positioning sound sources and mics by aiming the mic's null points towards loud sources of spill.

For most rock and pop projects the vocals are usually the focal point, and *must* be 100 percent correct – so it's best to overdub them after the rest of the music has been recorded. Using a good vocal mic it is difficult to get an isolated vocal sound when there are other loud sources also playing in the same room at the same time. Using a less sensitive mic, and having the singer "eat" the mic is as option, but not a preferred one – because the technical quality of the recording will be compromised. Overdubbing vocals also means that the rest of the band do not have to be present after their parts have been recorded. It *is* usually necessary to record a scratch vocal while the instrumental parts are being tracked though. This can be done from the control room (using a dynamic mic to avoid feedback and spill from the monitors), or in an isolation booth if one is available, or even in the main room with the rest of the band if the singer is not so loud that their voice will spill into other mics in the room (which will create a ghostly doubling effect when the "keeper" vocals are tracked later).

ABOUT CABLES

- Cheap cables and cheap connectors are a false economy. They will quickly become unreliable. There is signal degradation over long cable runs (including interference and high frequency loss), and the lower the quality of the cable, the worse it sounds.
- You don't need the ultra-over-priced-super-luxury brands, just good quality cables with rugged connectors on them. A good retailer/dealer will be able to advise you.
- Use the shortest cables possible. The shorter the cable, the less signal loss there is down the cable. There's also less of a coil lying somewhere for somebody to stand or trip on.
- Use good quality sub-snakes to minimize mic cable spaghetti. Ten long mic cables running from a wallplate to a drum set are an inconvenience to set up, and a trip hazard once set up. A single sub-snake running to the drum set, with ten much shorter mic cables connected to it not only looks better, but is also safer.
- Don't use cheap low quality snakes! There's no point in having good mic cables, and then plugging them into a cheap snake. If only cheap, discount or off-brand snakes are available, running multiple mic cables *would be* preferable.

16.4 Iso-Rooms and Multi-Room Studios

Quieter sound sources such as vocals, acoustic guitar, acoustic piano, and other acoustic instruments (horns, brass, strings, etc.) are difficult to record well if there are also much louder sound sources in the same room. These relatively quiet sources should be recorded in separate isolation rooms to decrease the bleed of the other sound sources into their mics:

- If a band has only one or two quieter sources, then it is best to isolate the quieter sources in separate rooms, or a single separate room if it is large enough.
- If a band is an acoustic ensemble, with only drums as a louder sound source, it may be easier to isolate the drums in their own room, leaving the majority of the performers together in the large room.

Smaller *iso-rooms* or *iso-booths*, built or treated to be acoustically dead and dry, work well for vocals because they produce a very intimate, close up sound, and allow different reverbs to be added during mixing. Small and medium sized rooms which are more live (but controlled and not too "roomy" and unflattering) are better for acoustic instruments. If there is a lack of bass trapping, basses and electric guitars often sound uneven or boomy in smaller rooms.

A HOME MULTI-ROOM SET-UP

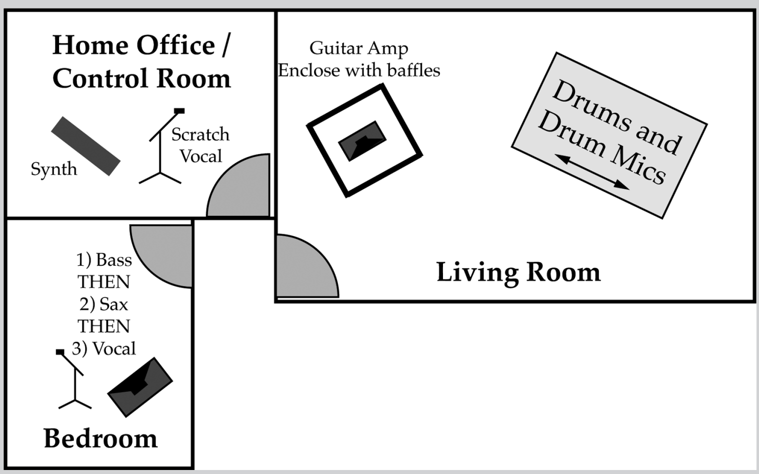

Figure 16.3 A possible set-up for multi-room home recording.

In a multi-room or home studio, it may be possible to record a band all at the same time, and maintain isolation between the different instruments by putting them in different rooms. Some of the acoustical home remedies discussed in an earlier chapter can temporarily transform multiple rooms into workable recording spaces. You don't need too many different rooms:

- Electric keyboards and synthesizers can be DI'd from the control room.
- A scratch vocal can be performed in the control room, and then it and any solos overdubbed from another room later.

- Drums can be recorded in a larger, more live room.
- Guitars can be recorded in a live or dead room, depending on the desired sound.
- Horns can be recorded in the less live rooms.
- Vocals and basses should be recorded in the deadest rooms.

It is also possible to DI and use amp simulators on guitars and bass, and have them play from the control room – but remember, the sound will not be the same as miking real cabinets.

16.5 Gobos and Sound Barriers

Gobos are large acoustically absorbent barriers. They can be purchased commercially or easily home built from acoustical foam or mineral fiber, wood, and casters. They come in a variety of sizes, from 1 m to over 2 m high (4 to 8 ft), 90 to 120 cm wide (3 to 4 ft), and some even have windows in them to maintain sight lines between performers. Smaller stand mounted screens and acoustical deadeners are also commercially available.

- Gobos placed in front of, or around a particular instrument, can tighten the recorded sound, making it less influenced by room reflections.
- Gobos placed in-between instruments can reduce the spill of the sound sources into each other's mics.
- Gobos placed near walls can reduce reflections from sources close to that wall – deadening the recorded sound, and potentially reducing phasing and comb filtering problems at the mic.

If there are multiple instruments in a room and you need to reduce spill, remember that gobos will change the sound of the instrument they are positioned around. For example, if drums and piano are in the same room and there is too much drum spill in the piano mics, one solution would be to build a cage of gobos around the drums to reduce the drum sound getting to the piano mics. But if the drums sound amazing, positioning gobos too close to them might remove too much of the room sound that makes them amazing – so it would be better to position the gobos closer to the piano to isolate it from the drum spill.

Applying diffusion products to one side of a gobo turns it into a mobile diffusor that can be used to break up reflections rather than absorb them – keeping the sound more live. In a large room, portable diffusors can be positioned around a sound source to reduce the natural room reflections, but create and add smaller space reflections from the closer boundaries they form. A few gobos with absorption on one side and diffusion on the other can be very useful!

Figure 16.4 Gobos set up around drums isolate the close drum mics from the room, reducing the room sound in those mics. This can be good or bad depending upon the drum sound desired – if the drums sounded amazing without the gobos, why would you compromise that sound?

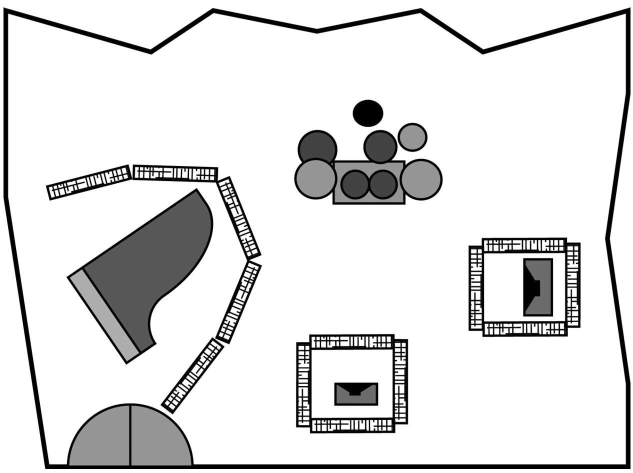

Figure 16.5 Gobos positioned around the piano, guitar, and bass cabinets isolate those instruments' mics from spill without changing the drum sound. The room sound is a less essential component of those instruments' sounds. A gobo "lid" or "roof" can also be placed over gobo "walls" to maximize isolation.

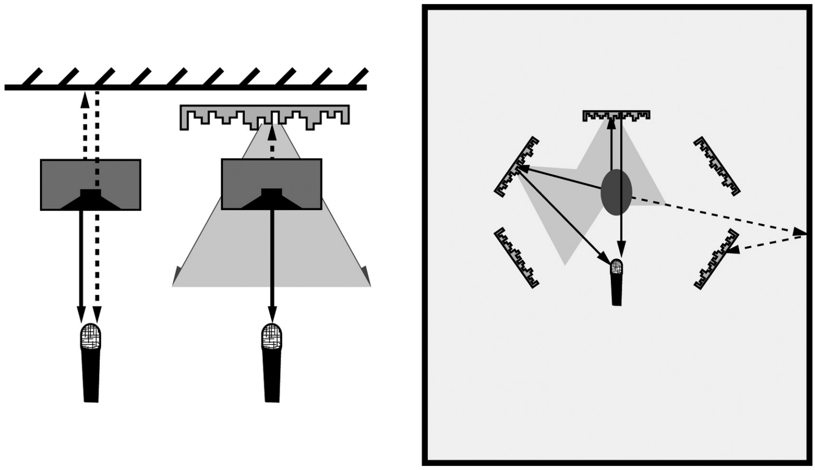

Figure 16.6 **Left:** A portable diffusor behind a speaker cabinet scatters reflections that could otherwise cause comb filtering at the microphone. **Right:** Diffusion gobos around an instrumentalist reduce natural large room reflections but keep the sound live by adding a smaller, more intimate environment to the sound picked up by the mic.

16.6 Drum Rooms and Drum Screens

Drum Rooms

A bad sounding or dead room will suck the life out of a recorded drum sound! The best drum rooms are medium to large reflective rooms. They produce a lot of early reflections and a short, bright reverb tail. The room sound is clearly present in the close mics and gives the drums excitement and power, as well as acting like glue to blend the individual drum and cymbal sounds together.

The smaller the drum room, the more reflective it needs to be in order to give the sound similar life and punch to a larger room, but the sound will always be smaller and more "roomy." Diffusion applied to the walls and ceilings of a slightly smaller than ideal room will break up and scatter reflections randomly, creating the illusion of a slightly larger space. Common bedroom sized room (or smaller), or rooms with less than 12 ft ceilings are going to generate "small room" early reflections that need absorbing rather than diffusing – but even then the drums will sound like they are in a small space.

Ultimately, big, powerful, punchy drum sounds will always be more easily recorded in larger drum rooms with high ceilings.

- In a studio with one large room and a smaller room or two, it is usually best to record the drums in the larger room, possibly with some other instruments also in there if drum spill can be managed so it is tolerable in other instrument mics.
- If a drum isolation room is necessary, then make sure it is also a more live room, with a higher ceiling, and not a small dead vocal room!

- If a super-clean and isolated drum sound is desired, then recording the drums in their own separate room, or recording them first and then overdubbing the other parts are ways to achieve that sound.

Drum Screens

Simple acrylic drum screens are made from large transparent panels that wrap around the front and sides of the drums. They reduce the level of direct higher frequency drum sound getting into other mics in the room. They are not comprehensive enough for recording though.

Suitable drum screen systems have absorptive material covering the bottom (and even the top) of the panels, leaving sight lines open, and an absorptive roof. They put the drums in a much smaller and drier acoustic than the larger room they are in. The thin absorption panels are not effective low frequency absorbers, but they do dramatically improve the performance and sound of the screen system – significantly reducing the level of drum sound escaping the cage. The most aggressive and loudest frequencies in the drum sound are mids and highs – these screen systems do effectively reduce spill in that frequency range. But because they reflect some higher frequency energy back into the drum mics, and of course isolate the drums from room reflections, they do contribute to a different recorded sound.

Behind the drummer, a wall of absorptive gobos can be set up, to reduce the level of drum sound reflected back into the drum mics and any other mics in the room. If the wall behind the drum cage is already absorptive, then these extra gobos are not necessary.

With a large drum set and lots of microphones, it can get cramped inside the drum screen. Make sure the system you're going to use, or purchase, is going to be large enough!

GAPS AND HOLES…

To clean up the recording significantly, it is not necessary to completely remove drum spill from the other mics in the room. Drum screen systems have many large gaps and holes in them, and sound does escape and get into other mics. Lower frequency energy passes straight through the screens and panels. But this is not a problem!

If the screen system reduces the level of the drum spill by only 6 dB it will have a noticeable impact on the band-at-once recording – the drum spill will be *not-quite-half-as-loud* in those other mics. If the screen reduces the drum spill by 9 dB, it will be *half as loud* as it would otherwise have been. A 6 or 9 dB reduction in all the other mics in the room adds up to a significant difference in the character of the recording.

Figure 16.7 A drum screen set up around a drum set.

16.7 String Sections

Live string sections (in a rock or pop setting) are best recorded separately from the rest of the band. They are much quieter than rhythm section instruments and they themselves take up a lot of space in the room. If a separate room is not available, overdub them later.

1 Close mics will produce the most individualized "Hollywood" sound, with a lot of definition and grunt.
2 A more blended, less individual, yet still close and contemporary "section" sound can be obtained by positioning a mic above each desk (pair) of violin or viola players, either from behind or in front, at a height of 60 to 100 cm (2 to 3 ft) above the instruments. A mic just under the music stand of each desk of cello or bass players will give them a similar perspective.
3 Area mics can be positioned higher, in front of blocks of string players, for more of a blended section sound.
4 A more distant front stereo array can also be used to capture the natural "in the room" ensemble sound.

Close (or desk), area, and distant techniques can be used simultaneously on the same instruments, and blended together to produce various perspectives – not only between different songs or recording projects, but even within songs. The closer mics can be electronically time aligned to the distant mics for a more seamless blend, or not, leaving the closer mic's character more obvious and forwards.

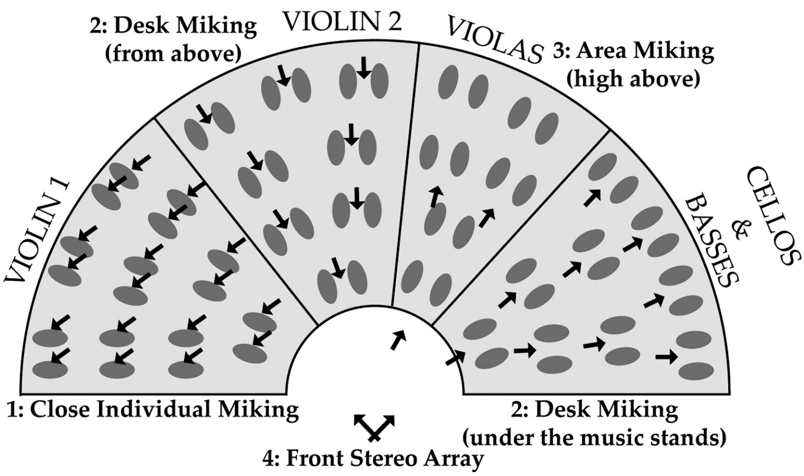

Figure 16.8 The close, desk, area, and stereo array mic techniques described in the main text are illustrated on different instruments of a string section. (This is just an illustration – the techniques used would usually be more uniform and not so different between the sections.)

16.8 Horn Sections

The saxes, trumpets, and trombones found in commercial horn sections are best miked individually so that their relative balances, processing, and the stereo image can be created and optimized for the project during mixing. Horn sections, like multi-tracked vocal images, can be panned quite narrowly for a more powerful and assertive sound, or widely for a more interesting and expansive, but less aggressive sound.

17 Miking Large Ensembles

In This Chapter:

17.1 Orchestras and Large Instrumental Ensembles

It's the lucky few who record professional large ensembles on a regular basis, and while large ensemble recording may not be a goal for everyone, you never know when you might get a call because somebody wants a concert recorded. This chapter should give you enough information to be able to say "yes" to that call, and approach a large ensemble concert recording successfully.

LISTEN, LISTEN, LISTEN!

You cannot record, mix, or balance any recording project if you are not familiar with the type of ensemble or musical genre. For large ensemble projects it is essential that you are familiar with how an orchestra, choir, or jazz big band sounds – both as a live acoustic experience, and in a recorded format. So – experience live performances and analyze respected commercial recordings!

Condenser mics are the obvious choice for this type of recording. Quality and reach are essential – because the sound sources are distant from the microphones. Small diaphragm condenser mics usually produce the most transparent, uncolored sound with the best transient and high frequency details. Large diaphragm condenser mics can be used when a bigger, more characterful sound is desired. Flat mics, or those with smooth, wide, gentle presence peaks are generally preferred over those with more aggressive presence peaks (which can emphasize unflattering string instrument characteristics).

Some engineers prefer flat, neutral preamps that emphasize accuracy, while others prefer preamps with more strident characteristics.

In a live recording, if there is an audience present, or the event is being filmed, microphone choices and techniques may be more restricted – the audience and cameras need their view to be as unobstructed as possible, so small diaphragm condenser mics are usually used.

ENSEMBLE ARRANGEMENT AND THE STEREO IMAGE

Large ensembles can be set up in different ways. A couple of orchestral string section and choir arrangements are shown below.

Different physical arrangements produce different recorded stereo images. The second of both of the diagrams in **Figure 17.1** spread pitch content over the entire stereo image more evenly. Choirs can also set up with the different voice types even more randomly scattered.

In a (non-concert) recording session, performer positioning should be optimized specifically for the recording. Physical arrangement should be changed from piece to piece – to produce the most appropriate, balanced, and interesting image for each work.

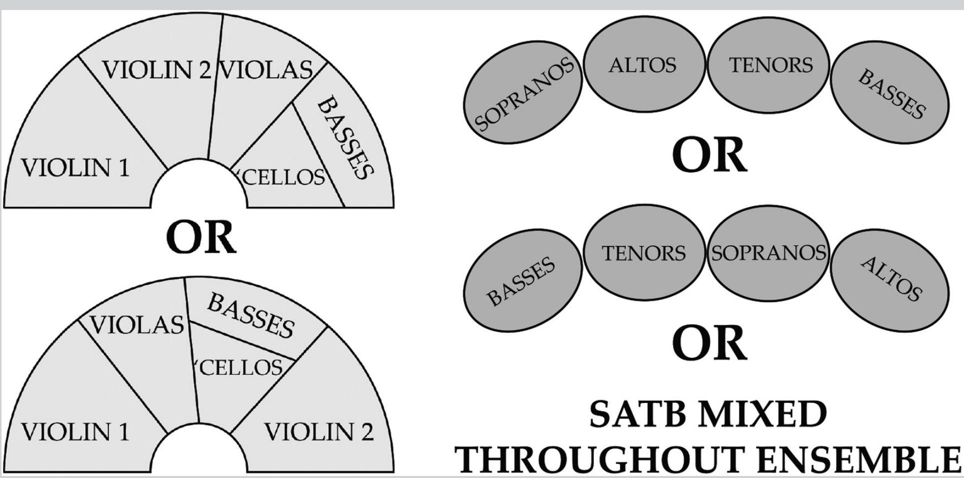

Figure 17.1 Different ensemble layouts spread pitch content and activity over the stereo image differently.

17.2 Main Arrays

In a good acoustic, a great recording of a small to medium-sized ensemble is possible with just a single stereo microphone array. The skill and artistry of this minimalist technique is in the placement of the array, and matching the choice of array to the ensemble and the venue's characteristics. Any stereo array can be used in front of a large ensemble.

Main array options include:

1. Coincident pairs, which produce very mono-compatible, but the narrowest, and least expansive and enveloping stereo images. *Do try them though – they can sound great on the right ensemble in the right venue!*
2. Near-coincident pairs, which produce wider, and less muddy images than coincident pairs – but they are a little less mono-compatible.
3. Spaced omni pairs produce a bigger, beefier low end – which results in big resonant string bass, timpani, and bass drum sounds. The stereo image can be incredibly expansive – but at the expense of imaging accuracy and precision.
4. Three spaced omnis – a center mic, plus two additional matching mics set up symmetrically about half to two-thirds of the distance from the center mic to the edges of the ensemble. Although not the most accurate or precise image, this technique does produce a very big, enveloping image. The center mic channel does not have to be used at the same level as the left and right mics – it should be used at a level that produces an evenly spread stereo image that does not have a hole in the middle, or become too mono, narrow, and center dominated.
5. Middle-Side technique. As well as offering the utmost in image finesse and precision, MS technique can be very useful in less than ideal monitoring situations – because the non-matrixed M and S signals can be recorded raw, and then matrixed to stereo (and the image width precisely controlled) in a studio environment after recording.
6. A Decca Tree, which produces a very immersive sound, but is a big awkward array to set up if space or visual aesthetics are a concern – which they might be in a live concert setting.

The main front mics and outriggers should be positioned at the same height, 3.5 to 5 m (10 to 15 ft) above the stage or studio floor, so that they pick up the front and rear of the orchestra as evenly as possible, and do not focus on the front rows of performers. To pick up a well blended sound and avoid a "scratchy" string sound, they should be at least about 2 m (6 ft) in front of the first row of musicians. They can be up to about 5 m (15 ft) away, but the sound will be a lot more diffuse at that distance. They should be moved forwards or backwards in order to find the sweet-spot where the wet/dry balance, stereo image width, orchestral balance, perspective, and string timbre are appropriate and focused. Decca Trees however, are generally positioned much closer to the ensemble.

Increasing the separation between the mics in near-coincident and spaced pairs increases the width of the image – but only up to a point, after which the stereo image falls apart, and develops a "hole" in the middle. Adding a center mic to a spaced pair to form three spaced omnis can compensate for this.

Increasing the angle between coincident and near-coincident mics will increase the perceived image width – however centrally located sound sources are then more off-axis, and subject to the mic's off-axis coloration. Hyper-cardioid mics are an alternative to cardioid mics, and they will exaggerate image width more than cardioid mics – but centrally located sound sources are subject to increased off-axis coloration.

Moving the main array closer to the ensemble will increase the width of the stereo image – but positioning the mics too close will produce a scratchy, unblended, individualized string sound – particularly with "less than professional" ensembles! It will also increase the direct sound, and reduce the reverb or room acoustic picked up. Moving the mic array further from the ensemble produces a more unified, smoother, blended string sound, increases in the level of reverb and room acoustic, but decreases the image width.

MAIN ARRAY POSITIONING – ART OR SCIENCE?

Main array placement is an art rather than a science – there are so many variables that all interact to change the perspective of the sounds recorded, the image width, and wet/dry balance, including:

- The type of array used.
- Microphone choice.
- The angles of incidence between the mics in the array.
- The distances between the mics in the array.
- The distance and elevation of the mic array.
- The room acoustics.
- The sound of the ensemble, and how it interacts with the room.

It is impossible to accurately predict which techniques will be most suitable for every situation. Experimentation during rehearsals is essential!

17.3 Outriggers

If only a single stereo array is used, instruments located towards the edges of a large ensemble can sound distant and unfocused compared to those located directly in front of the array. A more even pick-up might be achieved by moving the mics further away from the ensemble – but the stereo image will be narrowed, and the perspective will be compromised. A pair of *outriggers*, positioned in line with, and at the same height as the main stereo array, symmetrically spaced about two-thirds of the distance from the main pair to the edges of the ensemble will solidify the imaging of the sound sources towards the extremes of the stereo image. See **Figure 17.2**.

Omnidirectional mics are usually the preferred choice for outriggers because off-axis sound sources sound better, and the more diffuse image they create is easier to blend with the main array. Additionally, they offer better low frequency performance. However, if audience or HVAC noise is present, cardioid or wide cardioid mics can be used to reduce extraneous noises – although you might have to do an EQ boost of the extreme low frequencies to compensate for their natural low frequency roll-off.

The main array (usually panned hard left and hard right, but not always) should be used for the bulk of the pick-up, and to create a sense of the overall image width. The outriggers should be blended into the main array image to solidify the edges, and subtly increase its width:

- If the outriggers are not turned up enough, the edges of the image will be weak and fuzzy. Sound sources positioned towards the edges of the ensemble will not be as present or defined as they should.
- If the outriggers are turned up too much, the image will become extreme L and R heavy, and develop an unfocused center.

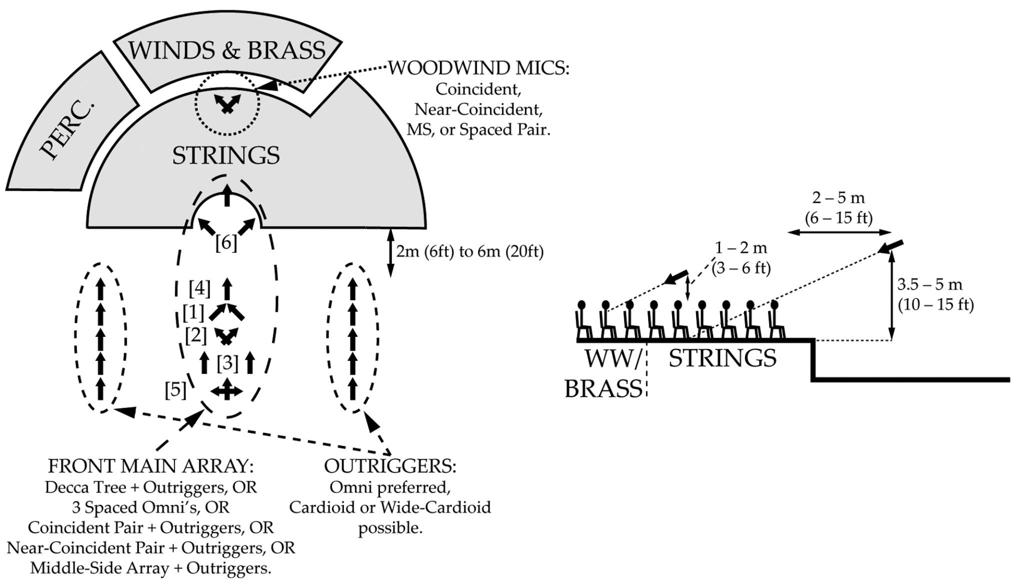

Figure 17.2 Top and side views of the orchestra or large ensemble microphone techniques discussed in the main text.

SETTING OUTRIGGER BALANCE

You should be able to clearly hear activity evenly spread throughout the entire stereo soundstage when the main array and outriggers are mixed together. To achieve this:

- ► With the main array faders at unity (and the mics usually panned hard L and R), turn the outriggers up until you start to hear them changing the stereo image. Make a note of that level.
- ► Continue to turn the outriggers up until the extremes of the stereo image become too loud, and the center seems weak. Make a note of this level.
- ► A good starting balance for the outriggers will be somewhere between these two levels.
- ► Mute and unmute the outriggers at different levels between the two positions previously noted. Increased presence and solidity of the extremes, and a little widening of the stereo image is what you're listening for. If there is a huge amplitude change, or the image gets dramatically wider when the outriggers are unmuted, they are probably too loud.

17.4 Woodwind Mics

The further away an instrument is from the front mics, the more distant its recorded perspective becomes. Woodwind instruments are relatively quiet, and don't project as well as the brass or percussion, so they are prone to lack definition and sound distant in the front mic's image.

A coincident or near-coincident array, or two or more (usually directional) spaced microphones can add definition to the woodwinds. They should be positioned several feet in front of the first row of woodwinds, and several feet above the player's heads, angled down into the center of the woodwind section, as shown in **Figure 17.2**. Position them so the array's effective pick-up covers as much of the section as possible – from both the left/right and front/back perspectives.

The woodwind mics may or may not need to be panned all the way left and right. Naturally, the woodwind section is not as wide as the whole orchestra, so it doesn't need to spread over the entire image width.

To balance and blend the woodwind mics, repeatedly mute and unmute them, listening for amplitude and perspective changes. There should be a little woodwind amplitude increase, and the woodwinds should jump forwards slightly – sounding less distant, and less ambient and reverberant. They are too loud if they jump out or sound so dry that they don't blend with the rest of the orchestra. Use a balancing process similar to that described for the outriggers.

17.5 Natural Reverb and Room Acoustics

Large classical ensembles sound best in reverberant concert halls, and those same halls are the preferred environments for recording them – without using artificial reverbs. A large ensemble recorded in a dry room without reverb sounds strange, unnatural, and even unpleasant – because the instruments sound too individual, and are not glued together by the room acoustic.

The *critical distance* of a hall is the location at which the level of dry direct sound is equal to the level of reflected reverberant content. The critical distance is not simply half-way back in the hall – it cannot be measured with a tape measure. It can be measured with acoustical measurement tools, or you can get a rough idea by walking from the front to the back of the hall (facing the back) listening for when the sound from the stage gets quite diffuse and unfocused. Reverb or room mics should be placed beyond the critical distance so that they pick up primarily reflections, and not a time-delayed version of the sound picked up by the front mics.

Reverb mics should be high in the air, pointed towards the rear of the hall, and probably upwards, so they focus on where the best sounding reverb reflections are coming from. They should not be too far back or too close to any walls or surfaces, because then they will focus on the reflections from only those surfaces and not the reverb characteristic as a whole. Reverb gives a concert hall, and therefore the recording, character – so take the time to experiment and find the best position for the reverb mics.

Any stereo array can be used for reverb mics, but those that produce wider images are recommended, because reverb promotes a sense of space, width, and envelopment.

FREE FIELD AND DIFFUSE FIELD MICROPHONES

Free field microphones are meant for use relatively close to a sound source, where they pick up predominantly direct sound. Their frequency response is usually flat, or designed with presence peaks for character and projection. Most mics are designed for free field use.

Diffuse field mics are meant to be used at greater distances, where the mic picks up predominantly reflected content – room and reverb mics, for example. Diffuse field mics have a slight high frequency boost to compensate for the natural attenuation of high frequencies over distance. This brightens up the sound so that the reverb is not dull or muddy. It also has the effect of bringing the sound forwards in the image, giving it a less “distant” perspective.

17.6 Audience Mics

Without an audience and the noise the audience generates, a concert hall or large church is ideal for large ensemble recording (if there is also no HVAC, traffic, or other undesirable noise). With front, woodwind, and reverb mics you should be able to achieve good results. However if the recording is of a live concert, it is also necessary to capture good sounding applause and audience responses. The audience is behind the front array and outrigger mics, so directional front mics don't do a good job of it. Rear room/reverb mics will pick up only a small part of the audience, with a very wet sound – so it is essential to set up dedicated audience mics.

AUDIENCE NOISE

Audiences are an essential part of a live concert experience. However, they generate a lot of undesirable noise during performances, so only use the audience mics when there is applause, and between pieces – bring them into the mix smoothly before applause starts, and then fade them down slowly and subtly before the next piece starts.

Extraneous Noises

Make sure the venue you're planning to record in is quiet. Although there is excellent software that can reduce the level of undesired continuous noise (like HVAC noise) and the occasional non-continuous noise (a cough, or vehicle driving past), it is expensive, can be time consuming to use well (and transparently), and it can impact the timbre of the recording when it is made to process heavily. Software noise removal should be considered a last resort. If you think you'll need to remove continuous noise, you will need to record some of that noise alone, so the software can analyze it and know what to remove. Unfortunately, many churches are close to roads, so if possible try to do a recording session at a time when traffic noise can at least be minimized.

Rear facing microphones, in front of the critical distance, make good audience mics. Options (some shown in **Figure 17.3**) include:

- Directional mics (even shotgun mics) on the front of the stage, left and right, pointed into the center of the audience. Make sure they are high above the audience, angled down into the center of the audience so they don't focus on only the front rows.

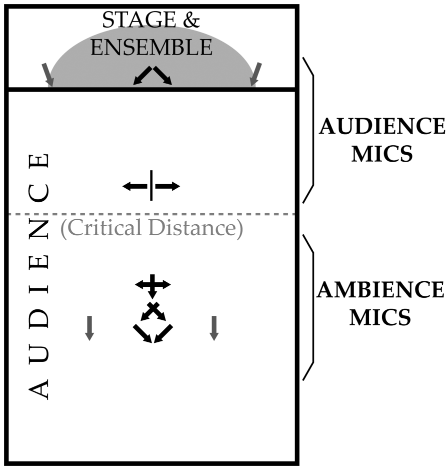

Figure 17.3 Ambience/reverb mics, and audience mics.

- Coincident, near-coincident, or MS arrays positioned centrally, in front of (and above) the audience.
- A baffled AB array positioned *in* the audience will provide a less distant perspective.
- In larger halls, or if a closer audience perspective is desired, zone techniques can be employed – multiple audience mics or arrays positioned throughout the hall, with outriggers to the sides in a wide hall. Unlike reverb mics, these audience mics should be pointed down into the audience.

17.7 Spot Mics

Spot mics are close perspective microphones used in addition to the main front array, outriggers, and woodwind mics. Spot mics serve several different functions:

- They allow the balance of an instrument or voice to be increased so that important or subtle solos can be heard more clearly.
- They can change the perspective of an instrument or voice – bringing it closer and more forwards in the stereo image.
- They can increase the clarity, focus, and definition of a sound source that is far from the main mic arrays. Blended in at a low level, a spot mic can add definition and details without making the instrument or voice louder.

Commonly used spot mics include:

- Soloists (vocal, or instrumental) in the ensemble, and in front of the ensemble. When deciding what soloists to spot mic, get the conductor or director's input.

Read the musical score while mixing and editing. Fade soloist mics in and out subtly to avoid noticeable shifts in the stereo image, and only use them when necessary – that way the natural balance, perspective, and imaging of the ensemble are not permanently changed. If an occasionally used solo mic does noticeably change the image of the main arrays, then to avoid dramatic image shifts as the soloist comes in and out, it is best to leave them in all the time, but faded down by up to 6 dB when the soloist is not featured. Fade them up and down slowly and subtly!

- Double basses. Spot mics on one or more of the bass section add weight and low frequency definition to the more distant sound picked up by the main array.
- Timpani. Most of the attack and definition of the timpani radiates upwards, and never makes it to the main mic arrays. A single mic, or a stereo pair, several feet above the timpani will restore these characteristics.
- French horns. With their rear facing bells, French horns can sometimes sound more distant than the rest of the brass section (which is often the goal). A spot mic or two behind the horn section will add projection and definition to their recorded sound.
- Harp and piano. If they are not soloists in front of the orchestra, harps and pianos can sound quite distant in the main mic arrays – they are relatively quiet instruments which do not project definition, attack, and details well. Even if a piano soloist is in front of the orchestra, the main mic array (which is focused on the orchestra) may not do a good job of picking up the piano sound.

Hyper-cardioid and cardioid mics focus tightly on the sound in front of them, isolating it from its surroundings. Sounds that are too isolated from their context are difficult to blend into the stereo image. Wide-cardioid and omnidirectional mics are preferred as spot mics, because they provide a less isolated sound that blends into the image more naturally.

But a single spot mic of any polar pattern is mono – compressing sound that comes from a wide physical space into a narrow point source panned somewhere into in the stereo image created by the main mic arrays. Mono spot mics do not accurately represent the context of a sound source's location, nor do they make it possible to most accurately position the spot miked instrument with the sounds around it, in the stereo image. On more important sound sources, or physically larger instruments, stereo spot miking captures the sound with its own sense of width and space, and allows it to be panned and blended more naturally into the stereo image.

Stereo spot miking requires more set-up time, equipment, and preamp or mixer channels, so is not always an option – but it is essential for large sound sources such as pianos and sections of instruments or voices. Stereo mics, which have two capsules (either a Left and Right, or Middle and Side) in the same body, are a great choice for faster set-up and less visual distraction.

EQ SPOT MICS FOR THE MIX

EQ spot mics so they sound great *in the mix*. Do not spend too much time EQ-ing them so they sound great soloed out – you're not going to use them soloed out! It is often necessary to EQ spot mics quite severely, so that they provide only the sonic characters that the main array is missing. For example, on timpani spot mics, the low and low-mid frequencies can usually be aggressively attenuated because there is plenty of that picked up by the main array. The spot mic is EQ'd so that it provides the upper-mid and high frequency attack and definition that is missing from the main mics.

In traditional classical recording, spot mics are not usually used to radically change balances. In modern contemporary productions however, the balance changing possibilities of spot mics are exploited – there's just no way a singer can compete with a fully scored orchestra playing loudly, without being amplified via a solo spot mic!

17.8 To Time Align or Not to Time Align?

There can be relatively large distances between spot mics and the front mics – 6 to 10 m (20 to 30 ft) or more. This means that the main mics pick up spot miked sounds 20 or 30 ms after the spot mics. When the mics are combined, this can cause the spot miked sounds to jump forwards too much, and sometimes cause phasing and comb filtering. Some engineers compensate for this by time aligning, or delaying the spot mics.

This is easy to do using DAW delay plug-ins, and on most digital mixers:

- Measure the distance (mic capsule to mic capsule) between the spot mic and main mic array.
- Delay the spot mic by entering that measurement in meters or feet, or calculating the delay time from that distance measurement if necessary.

Sound travels 35 cm (1.13 ft) per ms at room temperature, meaning that it takes 2.9 ms to travel 1 m (or 0.88 ms to travel 1 ft). So if a spot mic is 10 m (30 ft) from the main array, the required delay on the spot mic is 26.5 milliseconds.

$$\text{Required delay time (ms)} = 2.9 \times \text{distance } \textit{in meters}$$

OR

$$\text{Required delay time (ms)} = 0.88 \times \text{distance } \textit{in feet}$$

Because of the relatively low levels of identical frequency content between front, woodwind, and spot mics placed throughout an orchestra, phasing and comb filtering artifacts are not

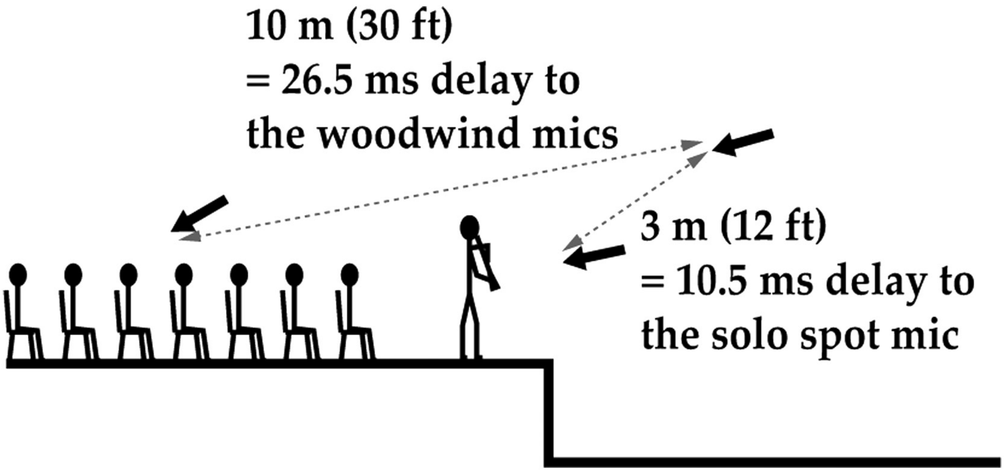

Figure 17.4 Delay times calculated for a solo spot mic in front of the orchestra, and the woodwind mics.

usually a huge problem. In fact some engineers refuse to time align mics – instead allowing the law of first wavefront to bring closer miked sounds forwards in the mix. This means that less amplitude is needed to make the effect of the spot mics obvious and appropriate in the mix. Other engineers feel that time aligning the woodwind, section, and spot mics to the main front array gives the stereo image more clarity, clearing up some blurriness and smearing.

BALANCING ORCHESTRA MICS

- The bulk of the stereo image is provided by the main stereo array, which is panned and balanced to produce a smooth and even "center to mid-wide" portion of the stereo image. The main array faders should be at unity.
- Add the panned outriggers to the main array to increase the image width, and solidify the extremes of the image.
- Bring up the appropriately panned woodwind mics to focus the woodwinds, enhance their stereo imaging, and make them less distant.
- Add the reverb/room mics to make the sound more natural and glue the orchestra together.
- EQ the continuously used spot mics (double bass, timpani, French horn, harp, and piano) so they provide what's missing in the main arrays, and creep them up in the mix until they fill in what's missing. Make sure they are panned to the instrument's location in the main array.
- Finally, EQ and add the non-continuously used soloist mics to the mix, so they provide a slight amplitude boost where necessary, and focus the perspective of their sound sources. They should be panned to the instrument's location in the main array.

17.9 Artificial Reverb

If the natural reverb in a hall does not have the right sound or character (it may be too long, too short, too dull or too bright), artificial reverb can and should be used – possibly combined with the natural reverb. *Reverb for classical recording needs to be top quality though!* Cheaper reverb units or plug-ins usually have obvious and grainy characteristics which work well in pop music production, but not for classical projects. Smooth natural sounding reverb that blends seamlessly with the acoustic sound and accurately recreates a real concert environment is preferred.

ARTIFICIAL REVERB TIPS

It is common to dial in more reverb on the closer and drier spot mics than on the main mics, to help them blend in.

Most classic outboard hardware and plug-ins are *algorithmic reverbs*. This type of reverb uses mathematical equations (algorithms) to generate the artificial reverb sound. As a generalization, these reverbs are more "noticeable" than "natural." Algorithmic reverb is great in non-classical productions, where it is used for stylistic and obvious effect.

Convolution reverb plug-ins simulate real spaces by applying processing based on the analysis of impulses recorded in real acoustical spaces. These *impulse response* sources can be sine wave sweeps played via a loudspeaker in the space (preferred nowadays), or transient events such as popping a balloon, or firing a starter pistol. Convolution reverbs are ideal for classical music because they are so natural sounding – more "natural" than "noticeable." Convolution reverbs are less editable than algorithmic reverbs.

17.10 The Hollywood Sound

Orchestral film scores and "pops" recordings have a very different sound from traditional classical concert hall recordings – the orchestra as a whole has a closer perspective, and each instrument is more discrete. The following techniques can be used to achieve this sound:

- A Decca Tree, positioned above the conductor as the main array. It is closer to the orchestra than other main arrays that are usually located in front of the orchestra and behind the conductor.
- Outriggers in front of the orchestra.

- Additional section/zone mics. These could be:
  - A couple of mics on each section of the orchestra, a few feet above, and slightly in front of the performers they're aimed at.
  - A mic on each desk of string players, plus individual mics on the woodwinds and brass, and multiple mics across the percussion section.
- Alternatively, each instrument can be individually miked.

For live sound reinforcement, miniature clip-on instrument mounted mics or pick-ups/transducers (on string instruments) can be combined with traditional free field mics, but clip-on mics and pick-ups are not commonly used in recording studios because they don't sound as good as traditional mics.

In order to increase isolation, and decrease the spill of louder instruments into the string and woodwind section mics, gobos, acoustic baffles, or acrylic screens can be positioned in front of the brass and percussion sections, and the physical arrangement of the ensemble is often changed so that the brass and percussion are located further away from the strings and woodwinds.

Scoring stages are large studio recording rooms used for orchestras. They are much smaller than concert halls and do not have the same reverb characteristics. They can be live, semi-dead, or anywhere in-between – but they are large rooms and not cavernous concert halls, so artificial reverb is essential.

While six main mics plus a few spot mics can produce great traditional, natural sounding orchestral recordings, many more mics are necessary to get an intimate "Hollywood" or "pops" sound.

17.11 Large Choirs

A large choir has significant physical width, but is only usually a few people deep. One approach to miking this type of choir is to use a main front stereo array, plus outriggers – similar to the front mics on an orchestra.

A good starting point for the main stereo array is 3.5 to 5 m high (11 to 15 ft), and 3 to 7 m (10 to 20 ft) in front of the choir, angled down towards the middle row of singers, if not just in front of that point to pick up more floor reflections and potentially increase presence and liveness.

- If the mics are too close they will pick up individual voices rather than a blended ensemble – but the sound will be drier and less reverberant, and the stereo image will be wider.
- If the mics are further away, the sound will be more organic and blended – but the stereo image will be narrower, and the mics will pick up more reverb.

The addition of outriggers will:

- Increase the width of the stereo image.
- Give the singers on the edges a more similar balance and distance perspective to those closer to the main array.
- Solidify the imaging of the singers on the extreme edges of the ensemble.

The outriggers should be positioned at the same height and distance as the main array. To balance the outriggers, follow similar steps to those discussed for orchestra recording earlier in this chapter. Reverb or room mics can also be added if the recording is taking place in a concert hall or large auditorium.

A greater number of closer microphones can be used when a closer perspective is desired, or if a venue has inferior acoustics (necessitating the use of artificial reverb). Sections of singers can be miked – either groups corresponding to their voice part, or more arbitrary clusters of singers. Use the 3:1 rule described in an earlier chapter in order to minimize the effects of phasing and comb filtering between multiple closer mics.

Less is often more. For large vocal ensemble recording, too many mics generally sound worse than too few mics:

- The more mics you use, the closer they have to go to adhere to the 3:1 rule.
- Too many mics, positioned too close, pick up voices too individually and not as a blended ensemble.
- Sound sources are picked up by multiple microphones, with different time arrivals at each – compounding phasing, comb filtering, and mono compatibility problems, negatively impacting clarity, timbre, and the stereo image.

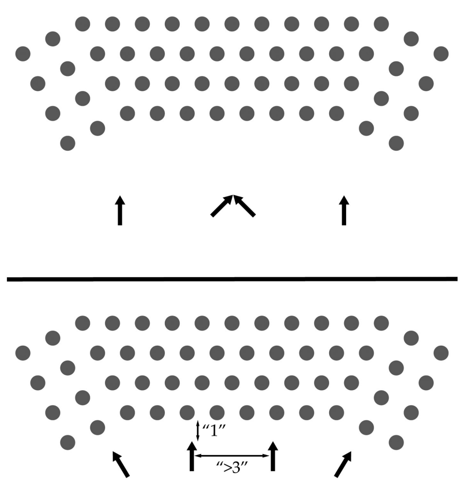

Figure 17.5 **Left:** Minimalist miking on a choir with just a front stereo array and outriggers. Any stereo array can be used. **Right:** A multi-miked "zone" approach, adhering to the 3:1 rule.

Spot mics can be added for soloists within or in front of the choir. Miking a soloist with a stereo mic will make it easier to blend into the image even though the singer is a point source.

17.12 Jazz Big Bands – Concert Seating

In a live concert, big bands usually set up as shown in **Figure 17.6**, with the horn sections in rows behind each other (saxophones, trombones, and trumpets), and the rhythm section to the side (piano, guitar, bass, and drums).

Drum Set

For a modern contemporary sound, rhythm section miking can be approached using the techniques described earlier in this book – but remember that a jazz drum set does not sound like a rock drum set. Jazz drums are smaller, and the kick often doesn't have a hole in the front – it produces a woofier, boomier sound that is an occasional accent, and has less attack and click than a rock kick drum. In rock and pop music the kick and snare provide time, and propel the rhythm forwards – in jazz, the hi-hat and ride cymbal are the timekeepers and driving forces, so make sure they can be clearly heard. Understand the jazz drum set and its sound before you try to capture and mix it faithfully.

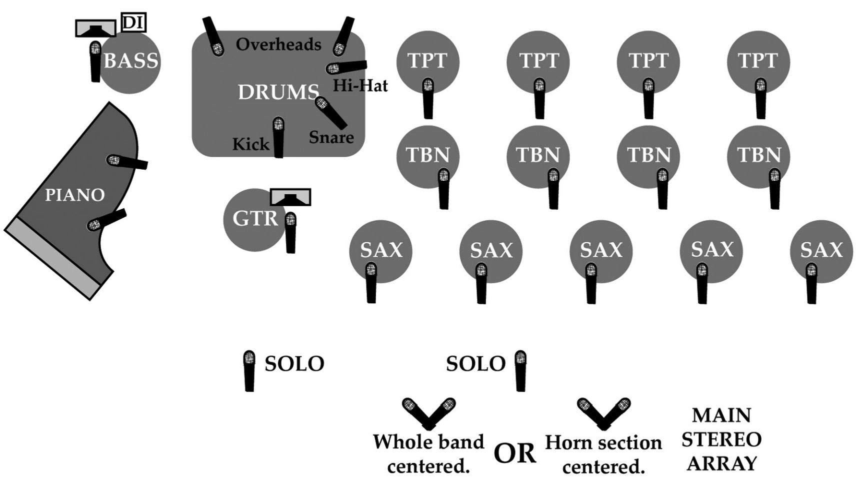

Figure 17.6 A miking strategy for a big band concert.

- Five mics are a desirable minimum – kick, snare, hi-hat, and stereo overheads.
- A kick mic and two overhead mics would be an absolute minimum.
- Tom mics and a bottom snare mic can be also be added.

Over-processing the drums sounds is inappropriate because the jazz drum sound is more of a natural organic whole than a gated, compressed, heavily processed rock sound. Inevitable spill should be embraced to create blend and not gated out to create separation.

Acoustic and Electric Bass

- Acoustic bass: If the bass is isolated, and the band are wearing monitor headphones, a bass amp is usually unnecessary. Blend a close bass mic, or mics, with a DI of the instrument's pick-up. If the bass is not isolated, or an amp is necessary, blend a mic on the bass cabinet with a DI of the instrument's pick-up.
- Electric bass: Blend a loudspeaker cabinet mic with a DI of the instrument's output.

Jazz Guitar

Jazz guitar sounds are fat, warm, and clean. Unless soloing, they are often barely audible behind the band. While a dynamic mic is a respectable choice, a condenser or ribbon mic will coax a little more bite and attack out of a sound that doesn't have much bite and attack – giving the guitar more clarity and exposure.

- Mic the loudspeaker cabinet – the amp and cone are essential parts of the sound.
- A DI can be blended with a loudspeaker cabinet mic, but *do not* use only a DI on the guitar – the sound will be flat and characterless.

Piano

A jazz piano sound is somewhere in-between a punchy pop sound and a more distant classical sound:

- Close mics and corrective EQ are inevitable in a live concert or single room recording session. If the lid is open, the piano mics will pick up too much spill (particularly from the drums). Close, or half-stick the lid, and have somebody play scales up and down the piano. Listen for notes or ranges that stick out or sound different. Tweak the mic positions until all the notes sound similar and even, and phasing (when summed to mono) is minimized between the mics.

- Try leaving the piano on full stick, and drape a heavy comforter or duvet over the openings. This will allow the mics to be positioned about a foot above the strings for more even coverage, and a bigger, warmer and more natural sound, requiring less corrective EQ.

Rhythm section mics can be panned to their actual positions towards the left of the image, or a more symmetrical image can be forced by panning the bass and guitar centrally, and the drums and piano across the stereo image.

The Horn Section

An accurate picture of a naturally well balanced horn section can be achieved using a single stereo array, 3 to 5 m (10 to 15 ft) above floor level, 2.5 to 4 m (8 to 12 ft) in front of the front row of saxophones. The distance, height, and angle of this array really affect the balance of the instruments picked up – particularly whether the saxes are too "close," too loud or too quiet, and whether the more distant trumpets are loud enough – so some testing and adjusting is necessary. Neither individual equalization nor compression of each instrument is possible using this technique, and it will also pick up a significant amount of rhythm section spill (particularly the drums). This section pick-up can be augmented with spot mics for soloists.

An alternate method is to individually mic each instrument. This is beneficial when:

- The horn section will require rebalancing during mixing.
- A closer perspective is desired.
- Individual EQ and dynamic processing of each instrument is desired.
- A more modern, commercial sound is desired.

Individual mics may also be blended with a main stereo array, and the individual or solo spot mics can also be time aligned to the main array.

HORN SECTION MICS – CONDENSER OR DYNAMIC?

Condenser mics are recommended for close horn section miking, not only because of their superior sound, but because their increased reach means that the performers do not have to "get on" the mics as closely as they would with less sensitive dynamic mics. They do not have to adjust them so much if they change from seated ensemble playing to standing for solos.

A drawback of this concert seating arrangement is that in the horn section, only the trumpets have any real isolation in their mics – the trombone mics also pick up the trumpets, and the sax mics also pick up the trombone and trumpets. The clarity, timbre, and balance of the horn section can be compromised. One solution is to use a main stereo array as the horn section pick-up, and augment that with individual mics for solos and quiet passages – but you need to have a top notch, well balanced horn section to do this effectively. Alternatively, the horn section mics can be time aligned to the main array, or to the sax mics if there is no main array.

17.13 Jazz Big Bands – Studio Isolation Seating

In a recording studio, the horns can be set up so that there is more isolation between their mics, allowing:

- More creative and clearer panning of each instrument.
- Individual equalization of each instrument.
- Compression to be applied to each instrument individually.

This can be achieved by setting the horns and director up in a square, as shown in **Figure 17.7**.

The piano and drums are best isolated – either each in their own room, or using gobos, cages, and sound barriers. At a minimum, try to isolate the drums in a different room or a drum cage. If necessary, the piano, bass, and guitar can all be put in one room, and barriers used to isolate them from each other. Amplifier cabinets can be completely "boxed" in a cage of sound barriers.

Each instrument in the horn section should be individually miked. Cardioid condenser mics for the two sections that are to the left and right, opposite each other – this puts each

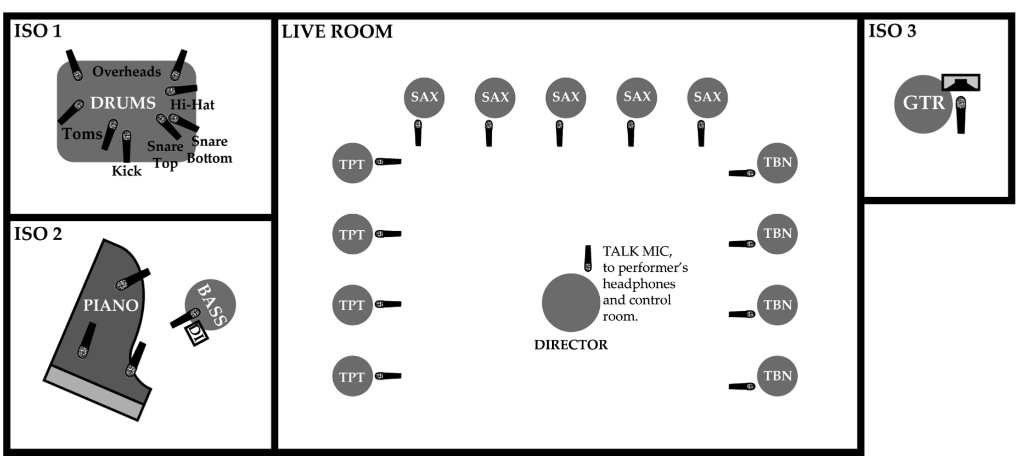

Figure 17.7 A miking strategy for a big band studio recording session.

section in the opposite mic's null points. For the section in the center, hyper-cardioid mics will minimize spill from the immediate sides – but with close positioning 30 to 45 cm (12 to 18 in) from each horn, spill should not be much of a problem.

Different horn section images can be created from recordings made in this way:

- The horn sections can be overlapped on top of each other, each across the entire width of the stereo soundstage – emulating a live set-up with the players in rows behind each other.
- Each section can be panned similarly to how they are seated.

Solos can be tracked live, as the band performs, or overdubbed later. If a solo is recorded live, and the soloist isn't happy with it, the whole band needs to re-record that solo section due to the spill in the room, and the solo getting into every mic in the room. Overdubbing solos after recording the band backgrounds allows the solo to be worked on and perfected without holding up the whole band, but the rhythm section and soloists do not get to interact and play off each other as they would if they were playing together.

Jazz and jazz recording styles are about cohesive units of sound. Spill should be embraced as a useful sound-shaping and blending tool, and not treated as an enemy to be eliminated as it is in a lot of rock and pop production styles.

18 Putting It All Together

In This Chapter:

18.1 Recording for the Mix

The purpose of this book is to give you the knowledge to record tracks that sound good, and sound the way you anticipate using them in the mix. Tracks "mix themselves" much more if you approach recording this way – mixing becomes easier and quicker, and doesn't rely on after-the-fact processing to "force sounds into the mix." For example:

- If you know you want a lead vocal to be front and center of the soundstage, record it so it sounds like that.
- If you know you want a sound to be further back in the mix and exploit the depth of the soundstage, mic it so it sounds like that and don't rely on EQ and effects processing to put it there.
- If you know you want a stereo sound source to be wide and expansive, or narrower and more focused, use stereo mic techniques that create the desired image characteristics.

YOU DON'T HAVE TO USE EVERY TRACK YOU RECORD

Analog tape machines have a limited number of tracks. DAWs don't! It doesn't take very long to set up a couple of additional mics to record some different perspectives of a sound source. A more distant perspective sometimes needs a little closer definition added to it for clarity and focus – so record that as well, and mix it in at a low level. (Checking for and correcting phase problems before you record of course!) Or you can decide you don't need the extra track, and don't use it. That's better than not recording the tracks and wishing you did have them!

With typical rock band instrumentation (vocals, guitar, bass, and drums) the only source that naturally has any width is the drums – via panned overheads, hi-hat, and tom toms. That could leave many of the other sounds parked on top of each other in the center, fighting for clarity and frequency space. Mixing becomes challenging, and the result boring to listen to. So record what you need to make the mix interesting:

- Double track (or more) rhythm guitars. The symmetrical panning options created relocate a lot of low-mid and high-mid frequency content to the sides of the image where it doesn't compete with other sounds. It will be much easier to keep the center panned bass, kick, snare, and vocals clear and intelligible.
- Record multiple vocal parts. You can then create different images and intensities depending on whether they are multiple unisons or harmonies, and whether they are all panned identically, slightly, or widely.
- Always record synthesizers and keyboards (except mono vintage synths) in stereo. They can be used to fill the stereo image. If there are multiple keyboard parts, panning some to one side of the image, and some to the other may produce a clearer mix than panning them all hard left and right – depending upon the sound and musical role of each part. Aim for symmetry – balancing an instrument or activity on one side of the image with one of a similar intensity of activity on the other side. This applies to stereo guitars too.

If you are stuck with limited instrumentation and minimal miking, use effects to widen the image. Delay effects, small room reverbs, chorus or pitch shift effects added to a single mono guitar track give you elements that can be panned widely to relocate the guitar from the center of the image to the wider extremes.

Do not simply set up mics in generic stock positions because somebody told you that was the only correct technique. And think beyond single mic, one take techniques.

Record for the mix!

18.2 Ear Candy

A recording should be more than a "demo," and more than gigging musicians presenting a song as they might in a live gig. There are no visuals to accompany an audio recording – so the music and production aesthetics alone have to sustain the listener's attention. Successful producers and engineers create sonic art that:

- Is compelling to listen to.
- Draws the listener in and retains their interest.
- Encourages the listener to play the track again because it is musically interesting.
- Encourages the listener to play it again because it is sonically interesting, and they wonder if there's any other cool *ear candy* they have yet to discover.

Ear candy includes things like dramatic and effective stereo image use, animated and active stereo images, creative effects, extra-musical events that happen "between" the main musical lines, layering of sounds, creative textures, and any unexpected "cool" moments!

Make sure you *record what you need for the mix!*

18.3 Pre-Production

Pre-production involves the producer working with the band before they come into the studio, so they can efficiently record a great piece of art. It includes:

- Refining musical ideas, song structure, and arrangement.
- Working out backing vocals and harmonies.
- Selecting the instruments and sounds to be used.
- Deciding whether to rent other instruments and hire additional performers.
- Developing and giving the artist's performance the necessary direction.
- Discussing the approach to and order of recording and overdubbing, use of click tracks etc.

It's much cheaper to enter the studio with these details worked out, and with a clear vision of the songs, than to spend hours or days paying for studio time while doing it. Depending on budgets, pre-production time can range from none to several weeks. Pre-production can also be as little as the producer/engineer attending a live show, or maybe a rehearsal, then suggesting some things for the band to work on themselves, or having to refine a song during the first few run-throughs in the studio.

18.4 The Pre-Mix

In addition to making sure there are no musical or technical problems with the tracks you are recording, you need to make sure all the sounds will work together in the final mix. The rough mix you build while tracking, the *pre-mix*, enables you to do this:

- Major EQ issues, that have not been addressed through mic choice or placement during soundchecks or rehearsal run-throughs prior to recording, should be fixed.
- Gentle compression should be added to tracks that have a wide dynamic range (vocals and bass would be prime candidates) – but do not record this compression.
- Tracks that will be panned away from center (drum overheads, stereo pairs, keyboards, and synths, etc.) should be panned approximately.
- Each track's amplitude should be appropriately balanced on the return faders. Good, musically appropriate pre-mix levels allow you to judge each sound in context, and have a good starting point for the performer's monitor mixes.
- Some generic, or stylistically suitable reverb should be added to the vocals, solos, horns, and any other instruments that are obvious candidates. Any sounds that will rely on reverb or other effects as part of their character should have at least an approximate version of that effect applied from the outset, so you can evaluate those sounds in context.

18.5 The Headphone Mix

With most hardware and software recording systems, the easiest mix to send to the performer's headphones is the one that is being heard in the control room. But this is rarely the best plan:

- If you compromise the "on the fly" pre-mix in the control room, you may not hear technical or musical problems that need correcting while recording.
- Each performer will also probably want specific things in their personal mix, so a "one mix fits all" approach won't keep everybody happy.

A better approach is to create a separate mix for each performer, and dial in only the things they *need* to hear. This can be done using pairs of aux sends on a hardware console (Aux 1 is left, Aux 2 is right, etc.), or in a DAW use stereo busses routed to dedicated pairs of interface outputs connected to the headphone system.

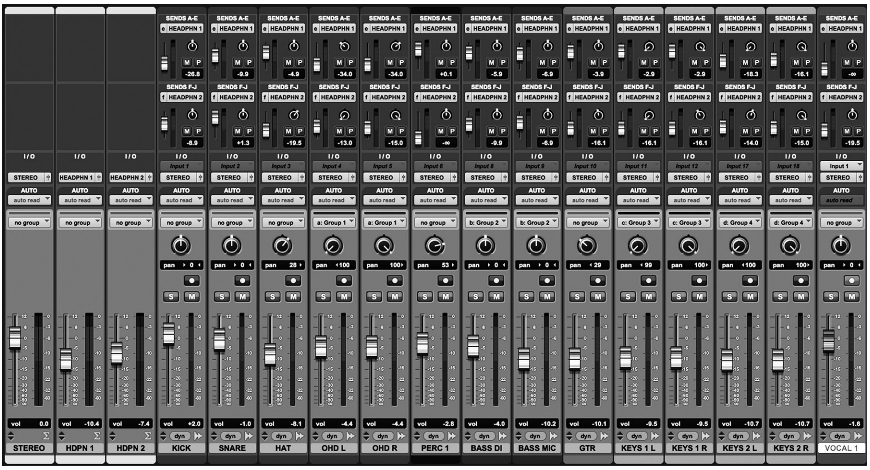

Figure 18.1 Two headphone monitor mixes in a DAW. Two stereo sends (the top two small faders of each channel) are assigned across all tracks, with different mixes (amplitude and panning) dialed in for each. On the left of the mix window are the master faders for the main stereo mix, and the two headphone mixes.

BUILDING A PRE-SET HEADPHONE MIX FOR A PERFORMER

- Put the performer themselves in their headphones so they are comfortable.
- Add the bass, and then the harmonic instruments (guitars/keyboards) to their mix so they can match pitch but still hear themselves adequately.
- Finally, add the drums so they can hear time, rhythm, and groove. Make sure the drums don’t drown out the pitch information previously added.

If the performer is having trouble with rhythm, try increasing the ratio of rhythmic sounds in their mix. If they are having trouble with pitch, try increasing the level of pitched content, and possibly themselves. If you turn something up, try to turn some other elements down – so that the headphone volume doesn’t creep up and get out of control.

Do listen to the headphone mix yourself as you create it – either in the room with the performer, or on identical headphones in the control room. This way you can hear problems, create solutions, and ensure the mix is at a safe listening level.

Personal headphone monitor systems allow each musician to customize the mix they hear themselves, using a personal mixer box. Using these systems *you* don’t have to get everybody’s mix right, or create a single “one mix for all.” But, they can take longer to set up than simple stereo-only systems, and might sometimes give inexperienced musicians too many choices – resulting in endless tweaking and little satisfaction! Sometimes it’s better to create their mix for them, making adjustments based on what they tell you they need – or what you infer they need, if they are unable to clearly explain it.

SETTING UP PERSONAL HEADPHONE MONITOR MIXERS

- Use stereo post-fader sends to build stems for the headphone mix, and duplicate mix panning to those sends. This way, if you're hearing a good pre-mix in the control room, setting all the post-fader sends to unity transfers that same mix to the headphones.
- For a typical pop/rock project, the stereo stems sent to the headphone system might be: Drums, Bass, Guitars, Keyboards, Backing Vocals, Lead Vocals, and Talkback. Effects can either be on an additional stem, or embedded with each existing stem.
- Label the channels on each monitor mixer.
- On the monitor mixers, set each stem's level control to an identical mid position. Set each stem's panning to be about 50 percent – some stereo image, but not too wide or distracting. This will give the musicians a good starting mix they can tweak to their satisfaction.

Too much going on in a mix can be distracting to performers, so it's a good idea to strip down a busy mix to the essential elements needed to promote a great performance. Many layers may add textural interest to the final mix, but they will probably not benefit a headphone monitor mix.

There are mixed opinions on whether a mono or stereo headphone mix is better. The benefits of a mono mix are:

- It's quicker to create.
- The stereo image does not distract the performer – wide and active mixes can be off-putting to some performers.
- If the performer removes one headphone muff (or uses a single muff headphone), everything is still in the muff they are listening to.

The benefits of a stereo headphone mix are:

- Clarity is improved. It is easier for the performer to hear all elements of the mix because they each have their own physical space and are not all panned on top of each other.
- The performer will have more of a sense of being "in the song" and more intimately connected to it.
- The performer will have more of a sense of their place "in the mix," and with some stereo reverb added, the sound will be more natural and familiar.

A compromise is to create a headphone mix with *some* panning (not extreme hard left and right) and keeping focal points and solo elements centered. As with any good mix, make sure the headphone mix is symmetrically balanced.

REVERB IN THE HEADPHONES?

Some artists like a dry mix, others like some reverb. A little generic reverb will produce a more organically blended mix that the artist may prefer to sing or play to. Reverb on their own sound (particularly vocals) while recording, also gives them important feedback about their pitch. Most singers and musicians prefer performing in reverberant spaces, so reverb usually helps – but you don't want so much that it prevents the artist from hearing the details of their performance. The addition of the extra layer of acoustical feedback some reverb tail provides may mean that less of the performer is necessary in the headphone mix.

To perform well, performers need to hear rhythm, pitch, and form:

- Rhythm is provided by the drums. In pop and rock, the kick and snare are the fundamental timekeepers. If the hi-hat and cymbals are too loud, the micro-groove information they supply can become over-powering, causing performers to try too hard to be on top of them, rather than relax *with* them. In jazz, the hi-hat and ride cymbal are most important, and take the place of the kick and snare in rock music.
- Pitch information is provided by a combination of the bass and harmonic instruments (guitars and keyboards). The bass provides the musical fundamental of the chord – the note from which the singer calculates the melody's pitch. The harmonic instruments provide chord type and closer pitch-matching information.
- "Where in the song" should be obvious if the performers can hear the instruments providing rhythm and pitch, and the arrangement of the song is effective!

Most importantly, the performers need to hear themselves! Some performers like lots of themselves in the headphone mix, while others will remove or partially remove a headphone muff so they hear lots of their natural sound in the room. For some performers, too much of themselves in the headphone mix is unfamiliar and will prevent them from giving their best performance.

Gently compress singers, wind, brass, and orchestral string soloists in the headphone mix, as you are doing the pre-mix – but do not record the compression. This allows the performers to "go for it" without suddenly hearing too much of themselves in the headphones.

Ultimately, the better the headphone mix is, the better the performance will be. A good headphone mix is especially critical for overdubs – final vocals, horn solos, string parts, etc. When the performer is drawn in by what they are hearing, their performance is always better.

OUT OF TUNE? TURN THEM UP. OR DOWN.

Turning some performers up in the headphone mix helps them stay in tune. But for others, turning them up makes things worse! If a performer doesn't notice they are out of tune, hearing themselves loudly in the headphone mix may suggest that they are "correct" because they are dominant. Turning the performer down can make the other instruments dominant and reinforce the desired correct pitch.

Is the Singer Constantly Flat?

- Reduce their level in the headphones – if they are okay with the quieter volume. This encourages them to sing louder – and in the process, open up and potentially sing a little sharper.
- Turn the entire headphone mix down. Quiet sounds are perceived as being slightly sharper than loud sounds, so this can make the singer think that they are even flatter and raise their pitch to compensate.
- A little vocal compression while recording (but not recorded) will allow the singer to "go for it" a little more, and possibly hit the notes more accurately without hearing too much of themselves during louder passages.

Is the Singer Constantly Sharp?

- Raise their level in the headphones. This can encourage them to back off a little, and possibly sing a little flatter.
- Turn up the entire headphone mix – if they're okay with that, and it's still at a safe listening level of course. Loud sounds are perceived as slightly flatter than quiet sounds, so this could make the singer think they are sharp and lower their pitch to compensate.

Single muff or double muff headphones? Many performers like to remove one muff from their ear so they can hear some of their natural sound in the room. Whatever they need to do to get a great performance is okay – with a few words of caution:

- An unsealed ear muff leaks sound which the microphone can pick up – so listen carefully for tinny headphone spill (and even feedback if the mic is cranked up in the headphone mix). Using single muff headphones will reduce this problem, and they are more comfortable than "half wearing" a muff.
- By removing an ear muff, the performer doesn't only lose 50 percent of their mix – they actually lose 60 percent or more because of the way our hearing works. Approximately

40 percent of the auditory information is supplied by each ear, and the remaining 20 percent by the brain's processing of this information. Removing one muff probably means that the mix needs to be turned up more in the other ear – louder than it would need to be if both muffs were worn. This puts the performer at an increased risk of hearing damage in that ear.

18.6 Click Tracks

The drummer is usually the rhythmic timekeeper of the band – everybody else synchronizes to the drummer. Less than ideal drummers may push and pull the time quite significantly. Timing imperfections in the drum part will be revealed and magnified by the recording process and exist forever! Unfortunately for a sloppy drummer, the best solution might be to get another drummer! Or can a click track help?

Part of the pre-production discussions with the musicians should be regarding the use of click tracks, and whether they are comfortable performing with them. Musicians who are not used to performing with click tracks usually have difficulty doing so. The recording studio is not the time to try it out for the first time. Only the drummer may need to perform to a click – as long as they provide time for the rest of the band *throughout the entire song*. If the drummer is having difficulty playing to a click track you will probably get more musical results by turning it off – even though the tempo may fluctuate.

Click tracks can help maintain rhythmic accuracy in songs with tempo variations (which need to be programmed in to the DAW session), or where the drums lay out for sections of the song – again, assuming that the musicians are familiar with playing to a click track. Many DAWs default to providing beeps as the click track. Some musicians prefer more musical sounds that blend better, while others prefer the artificial obviousness of the beeps. You need to be able to provide click track sounds that the drummer or other musicians are comfortable with.

One problem of using a click track is that the tempo becomes static. Great songs often speed up slightly for the chorus, then slow back down for the next verse. This gives the chorus more excitement and energy. These subtle changes can be programmed into your click track!

If you're using a DAW, it's best to have the DAW supply the click track, rather than an external metronome. The bars/beats grid of the DAW session will line up musically with the bars and beats of the performance, making editing much easier, as shown in **Figures 18.2** and **18.3**.

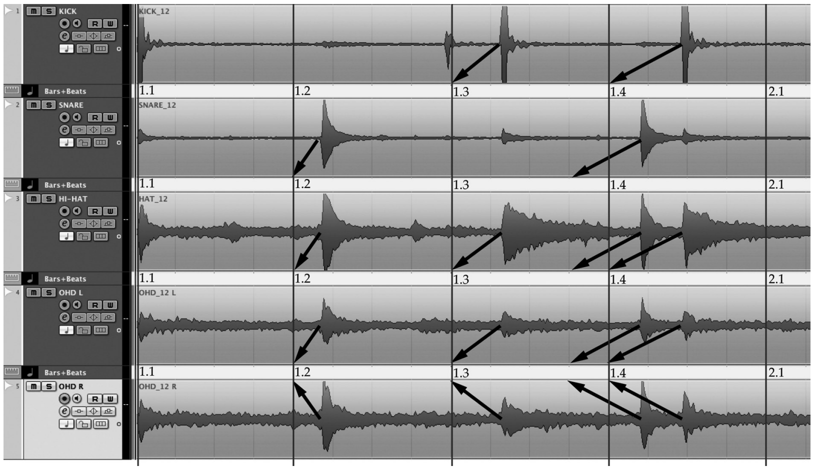

Figure 18.2 A recording made to an external click track (not shown) does not sync with the DAW tempo track, therefore the recorded drum tracks line up with themselves, but not the "bars and beats" grid – making editing more difficult. The arrows show where the events should be to line up with the bars and beats grid.

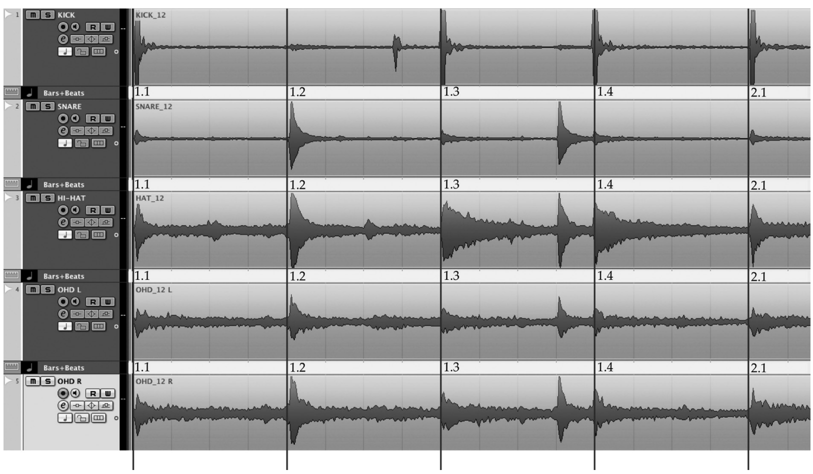

Figure 18.3 Drums recorded to the DAW's internal click (not shown) result in a performance that lines up with the DAW's grid – making editing much easier.

18.7 Knowing the Song

If the musicians cannot play their part of the song individually, and don't know their part upside down and inside out, they have no business being in the studio – although, billing by the hour, studio "practice" sessions can be very lucrative for the studio owner! If the musicians have not had much studio experience, discuss the recording process with all of the band members before the session – and make sure they know whether the process will be one instrument at a time, or the whole band at once.

Punching in and overdubbing can correct mistakes or build up a track section by section – but the more punches and edits there are, the greater the chances of volume, timbre, energy, and intensity inconsistencies between takes, making mixing more challenging.

ARE YOU THE ENGINEER OR A PRODUCER/ENGINEER?

If you are the engineer your job is to record (and possibly mix), but not to express musical, performance, or other artistic opinions. That said, on hearing something that they know is not in the best interests of the project, most engineers will find a polite, tactful way to suggest a more suitable alternative. Methods of doing this include using subtle suggestions so the performers think they came to the conclusion or solution themselves, or taking a receptive band member aside for a discussion, and having them bring the issue up with the other band members.

In home, project, and many lower budget situations, there is no dedicated producer. Discuss your role with the band before the session to see if they will be receptive to your input as a producer/engineer, and then coach the performers, giving them the guidance and suggestions necessary so they deliver a magical take!

If you are an assistant engineer it is best to have *no opinion* and keep your mouth shut! You should do what the engineer asks you to do, and proactively do *everything* to set up the session and ensure it goes smoothly. Assistants *should* answer client questions concerning the studio and how to best use it, but *avoid* saying anything that might be contrary to anything any other member of the production team has said or might say. "I'm sorry, I don't know. Why don't you ask the producer or engineer?" will politely move the question to a higher ranking member of the production team.

18.8 Don't Give Everything Away at Once

The arrangement and recording and mixing processes are uniquely intertwined – *together* they produce a successful record or a hit song. A band hoping to have a hit record need a producer familiar with crafting arrangements that promote listener retention, and an engineer who knows how to deliver the sonic goods appropriate to the musical style, target demographic, and market.

A successful song leads the listener to its climax. If all the instruments thrash away similarly throughout an entire song it will be boring, with little sense of form, progression, dynamics, or building excitement. The studio is an expensive place for a band to work out an arrangement – so ideally a "recording arrangement" of the song should be thoroughly worked out and practiced before coming into the studio.

If a professional producer is not part of the project, it's important that the band:

- Objectively evaluate their song forms, structures, instrumental and vocal arrangements.
- Be receptive to outside advice from listeners and critics who are not obsessed fans.
- Be receptive to the input of the engineer – who hopefully has vast experience of making successful recordings, and understands the musical genre and the target audience.

You may be able to tactfully suggest that a guitar or keyboard drop out for a verse, then come back in for the chorus, for example. But if the recorded arrangement of a song is not perfect, you do have some important tools at your disposal to help create tension, release, and climax:

- The mute button!
- The copy/paste/delete capabilities of DAWs!

Experiment with muting textural sounds, additional riffs, and delaying the entry of percussion sounds until later in the song. Each chorus should have something new and extra added to make it bigger than the previous one. Before the final and biggest chorus, a breakdown can be effective – drop down to minimal instrumentation, then build back up to the biggest and final refrain.

A little bit of unpaid overtime is a fact of life for recording engineers. After the band have left you could quickly create an alternate version with your mutes and re-arrangements, and see what the band think of it at their next session. If they like it, they'll value your input more and should hire you again. If they don't like it, hopefully there's no harm done and they'll be impressed you cared enough to go to that effort.

Percussion parts like congas, bongos, triangles, shakers, and claves, give extra lift and intensity to bridges, choruses, and hooks. Multiple guitar takes can be multi-tracked, using different sound and mic combinations to provide additional textures that can thicken and intensify the guitar sound during bigger sections of the song.

A recording is a totally different aesthetic to a live gig. It is missing all the visual cues that allow us to perceive sounds more clearly. So, if a musical part is not really contributing anything useful to the song, try removing it – clearing up musical, image, and frequency space can reveal other more important elements of the song.

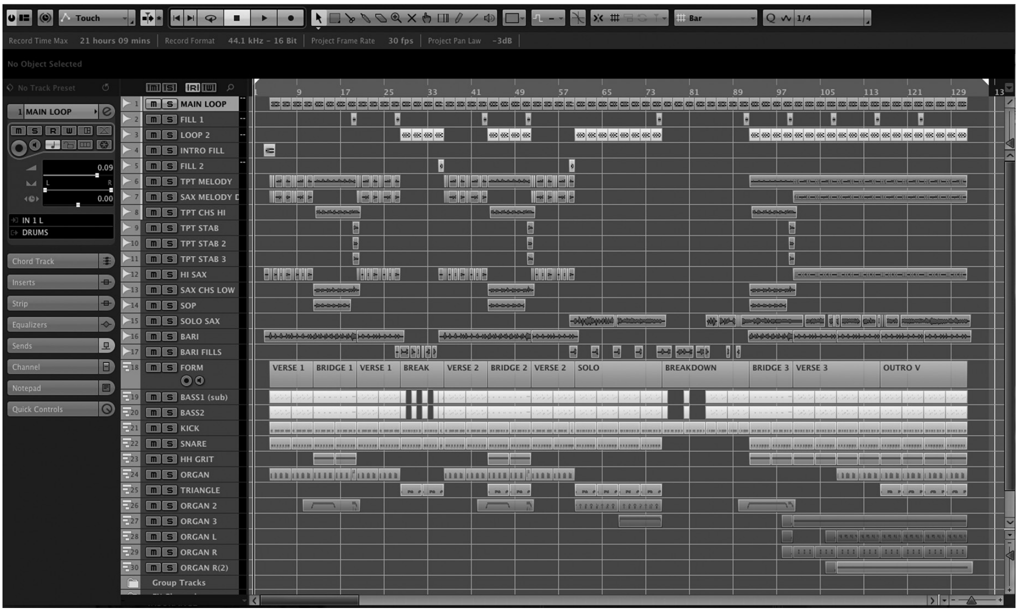

Figure 18.4 The edit window of a DAW shows an arrangement map of a song. Note how the bridge (chorus) sections are more heavily orchestrated than the verses, the sparse breakdown, and the build to the end of the song.

18.9 Correct Problems Early (But Know When to Let Them Go)

If you hear something is technically or musically not quite right – fix it. Immediately. Don't settle for second best. *"We'll fix it in the mix" is not a valid approach*, and technical problems are rarely fixable. A small glitch ignored until, or discovered during mixing can be a big problem – the quality of the final product may be compromised. The additional set-up, studio, engineer, and performer time necessary to bring everybody back to correct the problem is expensive. Correcting the problem immediately takes just a few minutes, and is time well spent.

Great musicality is what sells records to the majority of consumers. Only audiophiles are persuaded to buy something because it's a technically stunning recording. Getting a magical take can be more important than trying to fix a minor technical glitch in that take if the musicality is unrepeatable. Most listeners aren't going to notice a very minor technical issue – their focus is on being able to understand and sing along to the words, and not on the few milliseconds of slight distortion at the beginning of a syllable. You can do what you can to mitigate the problem – use a restoration plug-in if available, automate it to a lower level, make it less obvious through momentary use of EQ, cover it with reverb, re-draw the waveform in an editor or whatever – but know that *you* are always likely to notice the problem, because unlike a consumer, you know it's there and you're listening for it.

FIRST TAKES

You should want to create a perfect recording, and strive for that goal, but in the real and billable world, it's important to step back and realize that sometimes something that could be a little technically better is musically unique and wonderful – and worth keeping, because only we audio professionals find the problem distracting or annoying. *First takes often have freshness and magic to them, and even if you record three takes after that, you often find most of what you use comes from the first take!*

18.10 Fixing Dull or Small Sounding Tracks

Myth: Compression Cures All

PROBLEM: A common mistake made by novice engineers is using too much compression – thinking it's a cure for small, dull, or lifeless tracks, and that some magical hardware unit or plug-in will bring that track to life.

RESULT: What usually happens is that too much fast attack (or "auto") compression is applied. The sound gets smaller and less punchy, but because the output gain is turned up, it is louder. *But it's not better. It's just different*. It's easy to convince yourself something sounds "better" because you spent a lot of time working on it, or it is a little louder.

SOLUTION: Compression, correctly applied and controlled by a skilled engineer with a thorough understanding of the concepts and controls, can indeed make things bigger and punchier. It doesn't necessarily need a lot of compression. Good, artistic applications of compression include:

- The "sound" of a classic compressor.
- Dynamic control.
- Deliberate creative effect.

Trying to rescue a track that sounds poor because of bad instruments, a bad performer, or bad recording technique should be a last resort. It's much better to fix an imperfect sound at the source, and make sure it sounds great during the recording stage. As the saying goes, "garbage in, garbage out…"

AUDIO EXAMPLES
Can be found on the companion website

Compression

Example 18.1: A vocal, with no compression.
Example 18.2: The same vocal compressed fairly heavily to even out the short-term irregularities in the performance.
Example 18.3: The same vocal, processed gently with a classic tube limiter known for its phatness and warmth.
Example 18.4: A snare drum, with no compression.
Example 18.5: The same snare drum, squashed with too much fast attack compression. The life is sucked out of the sound.
Example 18.6: The same snare drum, with about 6 dB of slow attack/medium release compression to bring out the initial attack of the sound.
Example 18.7: The same snare drum, with about 6 dB of fast attack/fast release compression to exaggerate the body and sustain of the sound.

Microphone Distance

PROBLEM: Directional microphones are the most commonly used type in small rooms, or larger rooms if there is significant spill from other sound sources. Remember that directional mics have proximity effect, and the see-saw effect of EQ means that too many low frequencies have a similar effect to turning down higher frequencies – making the sound dull, dark, and muddy.

SOLUTION: Reduce proximity effect by moving mics further away. This will instantly increase transparency and definition, and produce a brighter, more vibrant track. The trade-off is that there will be more spill from adjacent sound sources, and more room sound – but the better direct sound may outweigh those drawbacks. Isolation can be increased by using sound barriers and gobos.

AUDIO EXAMPLES
Can be found on the companion website

Microphone Distance

Example 18.8: Electric guitar cabinet. Miked too close it sounds dull and lifeless.
Example 18.9: Moving the mic back about 30 cm (1 ft) produces a more vibrant sound.

Microphone Focus

The natural, acoustic sound of an instrument is the sum of all the sounds radiated from its entirety. The closer a mic is to a sound source, the more it zooms in on the sound coming from where it is pointed, and the less it picks up all the constituent components of the sound. While corrective EQ may be able to somewhat correct frequency imbalances, it cannot put back components that are either not there, or severely misrepresented. You must understand the instrument or sound source you are recording, and how it radiates its different sonic components. With this knowledge mics can be better positioned.

Moving the mic further away de-focuses it, allowing it to capture more of the sound of the whole instrument. As discussed in previous chapters, some instruments radiate certain frequency ranges very narrowly – so correct placement is critical, and small changes have a radical effect on the sound the mic picks up.

DON'T FORGET TO LISTEN…

The most important tools you have are your ears – not your mics. Before setting up a mic, listen closely to the instrument or sound source from different positions. Move your head (and ears) around the instrument as it is played, and listen to how the sound changes. From where does it sound most accurate, or best for the project? Put a mic there. Because a mic is more subjective in what it picks up than your ears, it will not sound the same, but it should be a good starting point.

Then experiment with the microphone's placement, moving it around close to that starting position – and anywhere else you feel like trying! If you are working with an assistant, have them carefully move a mic around the instrument while you listen in the control room. If you are working alone, send the mic's sound back into a pair of isolating headphones in the recording room – so you can get an idea of the sound the mic is picking up as you move it around the instrument yourself.

AUDIO EXAMPLES
Can be found on the companion website

Microphone Focus

Example 18.10: An acoustic guitar. Miked too close to the sound hole, the mic focuses on only part of the sound – the boom of the sound hole.

Example 18.11: Moving the mic back about 45 cm (1.5 ft) produces a more balanced sound incorporating the body and brightness coming from different parts of the instrument.

Using the Room and Early Reflections

Another cause of dead, dull, and unexciting tracks (particularly drum and guitar tracks) can be a lack of room reflections – the sound is simply too dry. We are used to hearing certain sounds in natural acoustics of some kind, and they become unnatural, bland, and lifeless when they are too dry.

A room treated with too much, or thin absorption is a common cause of dull sound. Good live rooms are designed to be reflective but controlled (with no standing waves), and have a balanced frequency response – they make things sound huge and full of life!

A potential benefit of increasing the distance between the sound source and microphone is the increased pick-up of important and beneficial room reflections. Positioning mics so that floor reflections are picked up will make some sounds instantly bigger, punchier, and more exciting. If your recording room has reflective floors (hardwood, tile, or concrete for example), remove any rugs that may be present. If your room has a carpeted floor, try putting sheets of plywood over it – the difference is night and day, particularly on acoustic guitars and drums.

THE ROOM MATTERS – EVEN WHEN CLOSE MIKING!

While it’s true that the closer a mic is to a sound source, the less reflected sound it picks up, the presence of those reflections has a drastic effect on the recorded sound, *even when close miking*.

Room reflections can be exaggerated by setting up room mics far from the sound source – listening carefully for, and correcting any phasing and comb filtering by adjusting the distances between the close mics and room mics, or by using electronic alignment techniques in a DAW.

As a last resort, reflections can be added artificially using an effects unit or plug-in. A reverb program that is predominantly early reflections, or a very short room, can work wonders on small, dull, compact, lifeless tracks. But of course, if you have a room with good sounding reflection characteristics it’s best to capture that natural acoustic.

AUDIO EXAMPLES
Can be found on the companion website

Floor Reflections

Example 18.12: The same acoustic guitar as in the previous example, recorded with a reflective floor (no carpet or rug). The reflected content makes the sound brighter and more powerful.

18.11 "Polishing a Turd"

Technical or musical problems can only be fixed and polished so much. Most sound sources can be buffed over and made to sound a little better by adding some compression, EQ, and effects. But, bad melodies or poor song structures don't result in successful records. Pitch and time correction software has its limitations and is time consuming to use properly. And some stereo image or gross distortion problems are impossible to fix. If the drums or guitar amp were just naturally horrible sounding, or the singer's performance was poor, no matter how much polish and buffing you try, those problems will still exist. It can be endlessly frustrating trying to make them less obvious – and in the end they will probably still stink…. Sometimes the only real fix is to re-record a better performance, or better sounding instruments or vocals.

18.12 Exploration and Experimentation

This book does not pretend to give you all the answers! That is impossible! There is no single way to achieve a great recording. But what this book should give you is the ability to:

- Think about the characteristics of a sound source and anticipate the recorded perspective the microphones will translate that sound into.
- Be able to use recording and mic techniques to capture accurate or stylized sound and timbres, different distance perspectives, and different stereo images.

You must have the mix in mind before you start setting up mics – otherwise how do you know what sound you need to capture, and how to set up the mics to capture it? Thinking about the mix as you record, and recording for the mix will result in:

- Technically better sound that needs less processing during mixing.
- A more streamlined mix process.
- Knowing immediately during tracking whether a sound is going to work or not.

Artistic and technical skills you should be comfortable with before pressing the record button include:

- Understanding the source sound, its musical role, and its place in the mix.
- Being able to choose an appropriate room, or adjust the room characteristics, to capture the best sound.
- Understanding mics, mic specifications, mic characteristics, and mic and recording techniques so you can choose suitable equipment and techniques for the job – guided by the mix you are recording for.
- Listening critically to the source sound, and using appropriate techniques to capture it faithfully or for its role and place in the mix.

Combining the miking suggestions in the latter part of this book with the theoretical knowledge presented earlier on will enable you to capture the sound you're really after by adapting the techniques described, and importantly, being able to anticipate the effect of those changes. Just like a musician builds their musical vocabulary by emulating others, practicing, and eventually synthesizing many influences into their own unique style, you should explore and experiment to build up your own collection of preferred recording techniques and solutions.

Basic rules and techniques should be mastered in theory and experientially before completely "doing your own thing" – because they'll help you recognize desirable sonic traits and avoid bad ones. But experimentation is key to developing your own style and being able to pair appropriate recording techniques with the aesthetic goals of different musical and production styles.

PEOPLE SKILLS

As important as these technical and artistic skills are, they're nothing without good people and inter-personal skills. You need to be able to put clients at ease, quickly establish good working relationships, be able to work as part of a team, and be able to coax memorable performances from the musicians. Without clients and great performances, the best mics and techniques in the world are worthless!

You must always aim for the highest audio quality possible – unless making a deliberate artistic statement. If we as professionals don't aim for the best, we set precedents that lead to a decline in accepted audio quality, and don't provide the next generation of future audio professionals with desirable material to analyze and learn from. True, many common consumer dissemination formats (lossy compressed audio files, played through bad sounding earbuds or tiny mobile phone speakers) do not reveal all the sonic details and definition we work hard to capture in a recording. But as streaming speeds and storage device capacities increase, and listening habits change (as they always do), eventually better sounding dissemination formats and listening systems will become mainstream. And they *will* reveal deficiencies in your work, giving your product a short lifespan and impairing your reputation as an audio professional.

19 Audio for Video

In This Chapter:

19.1 Why Audio for Video?

Most of this book is related to music recording, so what is a chapter on audio for video doing here? Versatility is key to finding and establishing an audio career path, and video production companies need audio people to capture good sound. This chapter introduces some of the concepts, technologies, and techniques of recording sound for picture, so if you get a call asking "can you do this?," you can. Some of the technologies and techniques are also relevant to theater and musical theater.

19.2 Types of Audio

Three types of audio make up video soundtracks, whether it is a film, TV show, documentary, or a corporate video production. They are *dialog, music, and effects,* abbreviated to *DME*.

Dialog can be recorded on set in a studio, or on location. Boompole, pistol-grip, or lavalier mics can be used. For documentary or news productions a handheld mic may be held by the talent. Depending on the type of production, mics can be visible or hidden. Dialog that tells the story – the narrative – used to be called *A-roll*, but that term is not used much anymore.

Music, as its name implies, is the musical part of the soundtrack – there to enhance the storytelling, and give the listener emotional cues.

Effects are additional sounds used to reinforce events happening on-screen or off-screen. They should make the storytelling and narrative more compelling and convincing.

- *Ambiences* are longer duration sounds that reinforce the setting of the on-screen action.
- *Spot effects* are shorter sounds that match up with a specific on-screen actions such as a door slam or gunshot.
- *Foley effects* are sounds "performed" by foley artists. The sounds synchronize with on-screen action, are created using various objects in real time. They are recorded in a foley studio. Footsteps, bones breaking, the sound clothes make as a person moves, and drinks being poured are some examples of possible foley effects.

Footage that enhances the storytelling but is not directly related to the dialog is known as *B-roll*. An interview might be part of a documentary, but watching the interviewee on-screen for two minutes may not be visually compelling – so the picture may cut away to something else, over which the interviewee's dialog continues. That B-roll video footage needs ambient B-roll sound that is going to enhance its visual content and make it believable, but not get in the way of the dialog.

Different sounds and sonic perspectives are needed for different length camera shots.

- *Extreme wide* or *extreme long shots* are very wide angle views containing entire buildings or landscapes.
- *Long shots* frame people from head to toe, and include their immediate surroundings.
- *Medium shots* feature only the upper half of the talent's body.
- *Close-ups* feature just the talent's head (or less).

If the video is an extreme wide shot of a crowd of people outside a building, we expect dialog to sound more distant than if it was recorded with a close mic, and also to hear more ambient/environmental sound. Conversely, we expect an extreme close up of an actor's face and mouth to be accompanied by a close perspective voice sound, and possibly less ambience.

Pre-Production happens before filming, and includes things like: scouting locations to be used for shooting; going to meetings to understand each scene, video shot, and actor movements; discussing wardrobe and clothing; and even starting sound design. This is the time to discuss technical issues, and equipment and technique choices – to ensure the audio and video gear will work together, and the most appropriate sound gets captured on set.

Production Audio is the recording of audio as filming or shooting is happening.

Post-Production refers to the creation and processing of audio that synchronizes with pre-existing video content – it includes dialog replacement, music editing, sound design, foley, voice-overs, and mixing.

CAPTURING APPROPRIATE SOUND

It is important to understand the storytelling and narrative, and to record sounds which best reinforce that, while making sure the dialog is clear and can be easily understood.

Dialog is often recorded with multiple microphone systems – a boompole mic and a lavalier mic for example. This provides options during post-production, and also redundancy in case one system fails.

The audio being recorded should always be monitored (on headphones) while shooting is taking place. It is important to make sure the sound is appropriate for the scene, and to immediately identify and fix any problems – it's cheaper to replace a bad mic cable before shooting, than to have to pay for an actor and studio time to replace lines after filming.

19.3 Recording Systems

The mics built in to video cameras and DSLRs are not suitable for professional use! In-camera mics can be upgraded by attaching better quality mono or stereo mics of various polar patterns to the camera. For example:

- A shotgun mic can be used for close-up shots so the sound "zooms in" to match the picture.
- A cardioid condenser mic can be used to pick up a broader field which might better suit a wide shot.

Audio from a camera-mounted or built-in mic can serve as a reference track to which better quality audio can be synchronized during post-production – so it *is* worth recording, even if it is not heard in the production.

An external on-camera mic will sound distant and roomy when the camera is more than a short distance from the speaker, making on-camera mics unusable for actual production quality audio most of the time.

External mics (and a mixer if necessary) connected to the audio input of a camera, can capture production quality dialog that remains synchronized with the picture – because the video and audio are recorded on the same device, and ingested into the editing software together.

Figure 19.1 A four channel preamp/recorder, in its over-the-shoulder bag.

Professional *production* or *location sound engineers* use dedicated hardware recorders to capture the audio from boompole mics, lav mics, and even multi-channel ambient sound mics simultaneously. Location mixers and recorders have between two and twelve analog mic inputs, and there are devices with many more digital inputs. Smaller devices can be battery powered and conveniently worn and operated in an over-the-shoulder bag, as shown in **Figure 19.1**. (If using a battery powered device, have *lots* of spare battery packs on hand!) On a small production, the same person can operate the recorder and boompole, and monitor the audio being recorded in real time via headphones. Larger devices are rack mounted in an *audio cart* – a portable rack-on-wheels which contains preamps, recorders, wireless receivers, the headphone system, as well as miscellaneous accessories.

Laptops and interfaces are rarely used on location by professional production sound engineers. They are not built to withstand the environment and abuse, or portable enough, and computers crash occasionally. Computer based systems are only suitable for the more controlled environment of an in-studio shoot.

19.4 Synchronization

If a separate audio recorder is used, the video and audio content of the production exist as multiple separate files which need to be ingested into editing software. They must be accurately lined up and synchronized so the audio plays back at the correct time relative to the picture.

Manual Synchronization

This is the most time consuming method, but with appropriate preparation can be fairly fast – and perfectly adequate for most amateur, semi-professional, and lower-budget professional productions.

- A *slate* (*clapperboard*, *sync slate*, or *sound marker*) is used at the beginning of each take. It makes a "clap" or "click" sound as it is closed. The cameras record a video frame in which the slate closes completely, and the audio recordings have a large transient spike in them at that moment in time. The slate should be in-shot on *all* cameras, and its sound picked up by *all* microphones.
- If a slate is not available, a person can substitute with a large, visually clear, and loud clap – that is easily visible on all cameras.
- After the director or lead videographer cues video and audio to start recording, the audio recordist should give a verbal response to confirm that audio is recording – for example, "sound, rolling!"
- After the "mark" cue from the director, the in-shot person operates the slate or provides the substitute clap.
- In post-production the separate video and audio files are ingested into the editor. The initial transient of the clap waveform can be quickly identified and lined up with the video frame in which the slate closes completely, or the hands touch together.

Over very long continuous shots (30 minutes or an hour or more) the audio and video may drift apart slightly because the clocks controlling video frame rate in the camera, and audio sample rate in the audio recorder may not be completely identical. It is sometimes necessary to separate the audio into smaller sections, and re-sync it to the picture by nudging subsequent sections backwards or forwards using some other audio and visual event that was recorded.

Timecode

Timecode is a continuous signal generated by a timecode generator clock in either the audio recorder, video camera or video recorder, or an external dedicated timecode generator. Timecode enables audio and video assets recorded in separate devices to be easily, quickly, and even automatically synchronized in the editing/mixing software. Timecode can also be used to start and stop multiple devices during recording. Some slates also have a display that syncs to and shows timecode values, which can facilitate synchronization of audio and video assets in post-production. The sound department often supply the timecode on set – timecode settings and responsibilities should be discussed with the video crew during pre-production meetings.

TIMECODE FORMATS

Non-drop frame timecode is displayed as a clock that counts in hours, minutes, seconds, and frames, with a colon ":" between each field:

HH:MM:SS:FF

A *drop frame* video format is indicated by a semicolon ";" between the seconds and frames fields:

HH:MM:SS;FF

Timecode frame rates must be set identically on all connected audio and video devices – but the timecode frame rate can be different to the video frame rate.

Common frame rates (in FPS, frames per second) are: 23.976, 24, 25, 29.97, 29.97 DF (drop frame), 30, and 30 DF.

Devices are connected together electronically using timecode (TC) in and out connections (usually BNC connectors and coaxial cable) on the audio and video equipment. Wireless systems are also available.

An audio recorder can lock to its own internal timecode clock, or synchronize to an external timecode source, either continually, or only at specific moments in time, by selecting different timecode modes in each device:

- *Internal (INT) Free Run*: The audio recorder generates its own timecode continually throughout the day, regardless of whether the device is actively recording or stopped. The start time of each file can be embedded into the audio file as a timecode value, as well as being transmitted to other devices via the TC out connection.
- *Internal (INT) Record Run*: The audio recorder uses its own clock to generate timecode, but only outputs timecode when the device is actually recording and the clock is moving forwards. This mode can be used to remotely start and stop other attached devices.
- *External (EXT)*: The audio recorder locks to timecode coming from another device, and regenerates its own internal timecode continuously based on that incoming timecode. This syncs the timecodes of both devices together, and can be used to bypass an inferior timecode clock on a cheaper audio device.
- *External (EXT) Auto Record*: If the audio recorder receives forwards moving timecode it will automatically start recording. It will stop recording when timecode from the external device stops moving forwards.

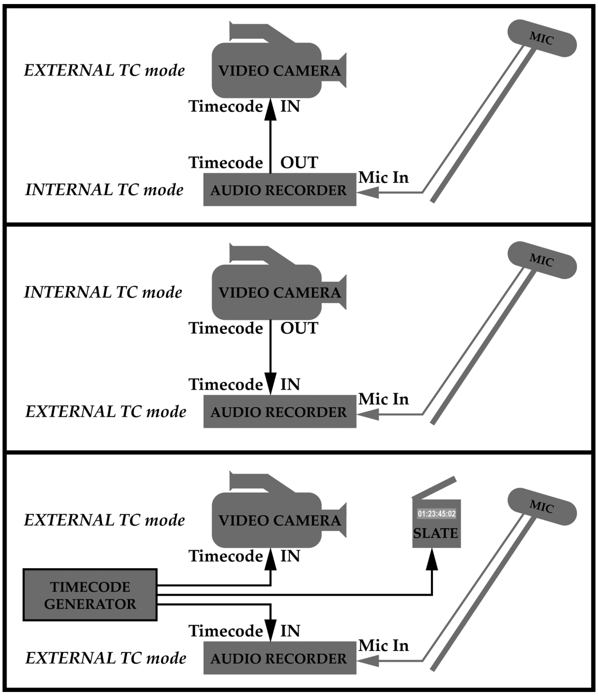

Figure 19.2 **Top:** The video camera syncs to the timecode generated by the audio recorder. **Middle:** The audio recorder syncs to timecode generated by the video camera. **Bottom:** The video camera, audio recorder, and slate display, sync to timecode generated by an external timecode generator.

JAMMING

Jamming is a process used in *free run* modes. The receiving device listens for, and synchronizes to a short incoming burst of timecode signal, and then freewheels and regenerates its internal timecode after that initial "reset" reference. Because individual clocks can drift over long periods of time, free run equipment should be re-jammed every couple of hours.

TEST, VERIFY, RECORD

► *Test your equipment*. Before a shoot, hook up the audio and video systems and sync them together, comparing timecode displays on all devices – so you know the planned synchronization works. Do this on a day before everybody is on set (and payroll!) so there is time to diagnose and fix any problems.

- *Verify timecode displays*. It is impossible to accurately determine frame-accurate sync on multiple devices by simply looking at the timecode displays – the numbers change too quickly! With the devices rolling, use the camera on your phone to take a single picture containing both displays next to each other, and check both displays read the same timecode in that picture.
- *Record .bwf Broadcast Wave Files*. Regular .wav and .aif files do not support burned in timecode values. BWF files store the file's start timecode as metadata, so post-production software can read it.

Synchronization Software

Professional video editing and post-production software can read timecode metadata in both the video and audio files, and automatically synchronize them.

If timecode was not used, or was set up incorrectly, there are standalone software utilities that can ingest video-with-audio files along with separate audio files, and automatically synchronize the audio files to the video file's audio, based on waveform analysis. These features are also available in some post-production DAWs and professional video editing software suites.

19.5 Shotgun Microphones

Shotgun mics are commonly used for picking up dialog, particularly for film and non-documentary TV shows. They are also used by sound recordists and sound designers to record sound effects and other source sounds. Shotgun mics are made to be more directional than super-cardioid and hyper-cardioid mics by placing an *interference tube* in front of a directional capsule. The interference tube results in a *lobar* polar pattern shown in the **left of Figure 19.3**.

Sound directly on-axis to the mic enters the open end of the interference tube. Off-axis sound enters via the carefully positioned cut-out slots down the length of the tube. The sound passing through each slot travels a different distance to the capsule, and arrives at the capsule with a different phase relationship. This causes cancellation of off-axis sound at the diaphragm – predominantly in the higher frequencies. Note the slight side and rear pick-up, which is a byproduct of using an interference tube to make the super or hyper-cardioid capsule more directional.

The specific directionality of shotgun mics varies greatly. Longer interference tubes produce increased directionality (narrower pick-up patterns), and more directionality at lower frequencies than shorter mics. Standard length (fairly short) shotgun mics are more directional than hyper-cardioids only above about 2 kHz. Longer mics can be more directional than hyper-cardioid mics down to about 500 Hz. At higher frequencies shotgun mic pick-up patterns can become quite uneven and lobar to the sides, as shown in the right of **Figure 19.3**.

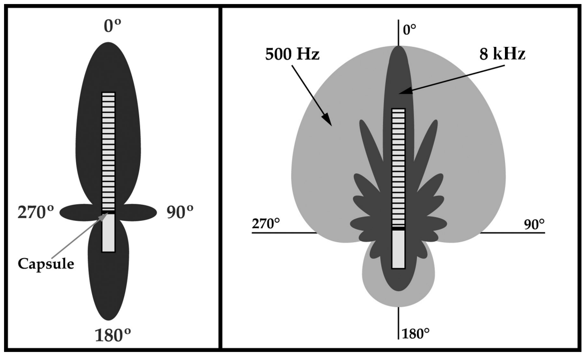

Figure 19.3 **Left:** A theoretical simplification of a shotgun mic's lobar pick-up pattern. **Right:** In reality, shotgun mics without built in DSP become more hyper-cardioid at lower frequencies, and the interference tubes cause increasingly lobar polar response at higher frequencies.

More expensive dual diaphragm designs incorporating digital signal processing can provide the directionality of longer mics in shorter, more compact formats, and also much more even directionality down to below 250 Hz.

Although shotgun mics can be mounted to any stand, and even on cameras, for sound-for-video, they are commonly used on the end of boompoles or hand-held using pistol-grips.

19.6 Blimps

Blimps are combined windshield and shock mount systems that reduce both wind and handling noise, an example is shown in **Figure 19.4**. The microphone is suspension mounted inside the blimp – the suspension mount reduces the level of vibrations caused by the

Figure 19.4 A blimp and pistol grip system. The microphone attaches to a shock mount inside the blimp, which is attached to the pistol grip. This pistol grip can be attached to a boompole, or hand-held. The blimp can be removed from the shock mount for indoors, no-wind situations. The dead cat windscreen can be stretched over the blimp in windy conditions.

handler, the boompole, and mic cable. Most blimps have a two-part windshield system – there is usually a foam base-layer built in, which can be augmented by stretching an artificial fur *dead cat* type of wind muff over the blimp. These muffs are made of artificial fur, and designed to minimize wind noise but remain fairly acoustically transparent. Fleece layers can also be used between the blimp and wind muff in high wind conditions.

For indoor use where there are no fast air currents, and the mic itself is not moved quickly during recording, windshields may not be necessary – the outer cylindrical part of the blimp can be removed, exposing the mic and shock mount. The mic should always remain in the shock mount to reduce handling noise.

19.7 Boompoles and Pistol Grips

Boompoles

Boompoles allow microphones to be positioned above (and slightly in front of) actor's heads, where a very natural, clear sound can be picked up. The boompole operator can continuously move the pole to keep the mic in the best position. It is also possible to move the mic rapidly between two actors. Boompoles range in length from about 2 meters (6 ft) to over 3 meters (11 ft). Less expensive models are commonly made from aluminum, and more expensive lighter-weight models from carbon fiber.

When setting up a boompole, the sections should not be completely extended – there should be about 6 cm (2 in) of tube overlap, to give the pole strength and stability. To do this, extend each section fully, then back the thinner section back into the larger section a little.

Just because a boompole has a rubberized or foam grip, do not assume that is the best place to hold it! If the boompole is extended to its full length, it can be held in the middle instead of at the end, possibly making it more balanced and less tiring to hold.

To position a mic overhead, the boompole should be supported by the bottom of the "V" between the thumb and fingers of the hand farthest up the pole, and lightly gripped. The other hand, closest to the grip on the bottom end of the pole, is used for steering – again a light grip and light touch is all that's required, so it's not too fatiguing. The boompole is usually held somewhat "sideways" relative to the operator's body.

To avoid fatigue, it is important that the pole feels balanced. If more than just a little downwards pull is required from the steering hand to keep the mic up, move the supporting hand a little farther up the pole. The arms can be more "upright" with the hands about two segments apart, or the hands can be farther apart. Work out what is comfortable for you given where you can stand for each camera shot, and what requires the least effort to keep the mic in position.

To steer the mic:

- Move the rear arm up and down to tilt the mic up and down.

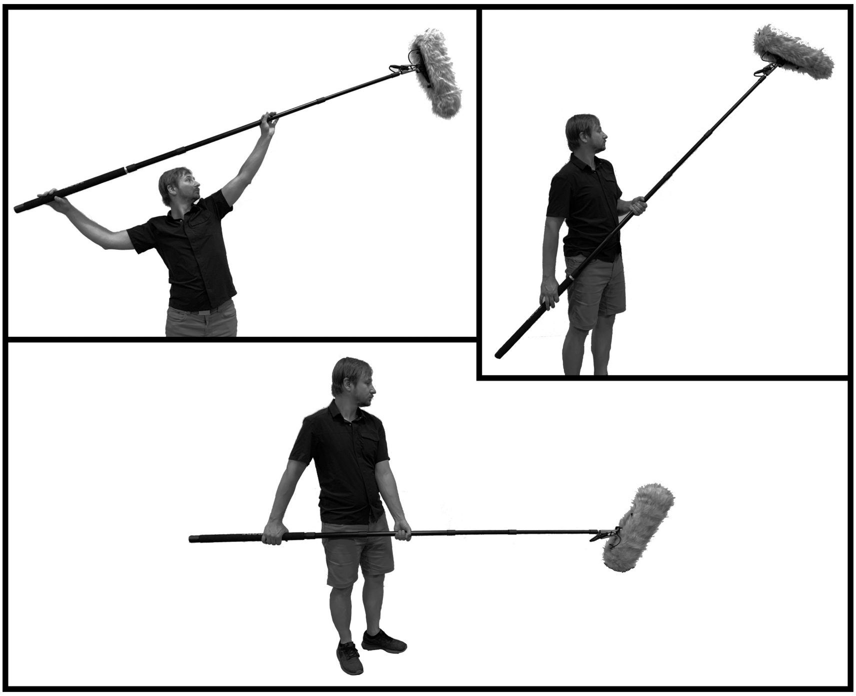

Figure 19.5 **Top:** Two methods of holding a boompole to position the mic out of shot, over an actor's head. **Bottom:** Holding the boompole to position the mic below the camera frame.

- Move the rear arm forwards and backwards relative to your body to swing the mic left and right.
- If there are two actors in a scene, the mic can be positioned above and between them. You can rotate the boompole with the steering hand to aim the mic at each specific actor as they deliver their lines.
- Both arms can be moved in and out along the trajectory of the pole to move the mic closer or further away.

The mic should be above and in front of the actor, aimed at their nose, at an angle of about 50º to 80º. If the mic is aimed too low, the sound can lose definition. If it is aimed too high, the sound can lose body.

If there are multiple actors, ideally they speak at a similar volume and the mic is positioned the same distance from each, and the timbre, proximity effect, and direct/ambient noise picked up remains consistent between them. This might mean moving the boompole from one actor to another, as well as rotating the mic to aim it correctly at each actor's nose. If two actors speak at very different levels, then the mic should be moved closer to the quieter person when they speak, so their lines are picked up more evenly, and the level of background noise remains constant between them. Turning the quiet person up when mixing will undesirably bring up spill and background noise when they talk.

HANDLING NOISE

To avoid handling noise and vibrations getting into the mic and recording, always use a shock mount system with boompoles and pistol grips.

Pistol Grips

Pistol grips are used to hand hold a shotgun or other condenser mic. See **Figure 19.4**. They are not used as much as boompoles for dialog recording, because they limit the mic to upwards firing positions from below the camera frame, and can only be used for close-up shots where it is possible to keep the operator and equipment out of frame. Pistol grips can be useful for sound effect recording where it is not practical to use a mic stand.

19.8 Lavalier Microphones

Lavalier mics (also known as *lav*, *lapel* or *body mics*) are miniature and subminiature mics that are clipped to clothing and ties, hidden in clothing or hair, or stuck to an actor's body. Lav mics are used in documentary productions, as well as in film and TV shoots where boompoles are impractical – for example, extreme long shots, when it would not be possible to keep a boompole and operator out of the frame. In a studio or documentary shoot, wired lav mics can be used, but in a scene involving actor movement, wireless lav mics are necessary.

Potential advantages of lav mics include:

- They maintain their position relative to the actor's mouth – so the sound is more consistent as the actor moves.
- They are positioned closer to the speaker than a boompole mic – so there is less background noise, and they capture a close proximity sound which is often desirable in documentaries and news gathering.

Potential disadvantages of lav mics include:

- Chest mounted mics can sound a bit "full" and lack the clarity and intelligibility of a boompole mic (although EQ and careful positioning can mitigate most of this).
- Clothing rubbing against the mic can produce rustling sounds (particularly in active movement scenes).
- Delicate clothing might be damaged by the mic clips/mounts.
- The talent can forget they are wearing the mic and knock or rub it accidentally.

Most lav mics are omnidirectional, although directional models are also available and useful in noisy environments. Most mounting techniques do not locate the mic in an optimal on-axis position with the mouth, so omni mics are usually easier to use. Also, omnidirectional mics are not as susceptible to plosives or wind current noise as directional mics. Even when using an omnidirectional mic it is important to point it towards the speaker's mouth – remember that all mics become more directional at higher frequencies, and this can certainly impact sibilance and intelligibility. Many lav mics are designed with a slight high frequency boost to compensate for off-axis positioning.

Positioning Lavalier Mics

There are many ways to mount and position lav mics, some are shown in **Figure 19.6**. For documentary, interviews and news gathering, it is generally permissible that the mic is visible – and mounting is easy, using "off-the-shelf" clips. For film and TV shows, mics must be hidden so they do not break the illusion and storytelling of the production. This often requires more customization and ingenuity. Most clothing, as long as it's not too thick and heavy, tends to act like a windscreen and not change the sound too much. Mics are frequently carefully hidden behind clothing, although it is always best to try to have some of the capsule exposed (but invisible to the camera).

Some visible mic positions and mounting methods include:

- Clipping between the buttons of a dress shirt.
- Clipping to the edge of a tie.
- Using a *vampire clip* to secure the mic if there is no clothing edge for a conventional clip.

If a conventional clip is used, a loop of cable is made around the clip so it is easier to direct the cable under the clothing to hide it.

To hide the mic, gaff tape or commercial lav mic adhesive products like "stickies" can be used. Some common positions include:

- Under a shirt or blouse collar, preferably near the opening so the capsule is not completely covered.
- Between dress shirt buttonholes, between the button and buttonhole layers. It may be possible to position the capsule so it is in, or slightly through a buttonhole, and not completely covered.
- Under a jacket lapel. For better sound it might be possible to position the capsule so it is in or pokes through the buttonhole.
- In the center of a bra, under the layer of clothing covering it.
- Above and forwards of an ear, hiding the mic under long hair, fixing it in place with medical tape or a hair clip.
- At the front hairline, hiding the mic in long hair, using a hair clip to keep it in place.
- Hiding the mic under the brim of a hat or cap.

WHAT'S THE BEST SOUNDING LAV MIC POSITION?

Choosing the best lav mic position is a combination of what the talent's clothing, hair, and movements allow, and what the camera will see in the scene – rarely is it based on where the best sound will be picked up! Headset positions are included here, even though headsets are rarely used for film and TV – but they are used for theater and sport event broadcasts.

1) A headset position, with the mic directly in front of the mouth produces the most accurate, intelligible sound. However, directional lav mics are prone to plosive booms, and proximity effect.
2) Headsets which locate the mic at the mid cheek line, to the side of the mouth, perform quite well in the essential upper-midrange intelligibility frequencies – however directional sibilant frequencies are attenuated because the mic is significantly off-axis to the mouth. A high frequency EQ boost, or a mic with a built in HF boost is desirable in this position.
3) Forehead and center hairline positions can be surprisingly good and clear sounding. Of the off-mouth-axis positions they offer the flattest and most natural pick-up, often requiring the least compensatory EQ.
4) Positions near the top of the ear or over the ear near the upper cheek bone put the mic very off-axis to the mouth, and have significant high frequency reduction above 1 kHz. Compensatory EQ, or a mic with a built in HF EQ boost is definitely necessary in this position.
5) Chest positions sound full and congested and lack intelligibility due to a large reduction of the 1 kHz to 5 kHz range, and the mic's location right on top of lower frequency chest resonances. For chest-mounted positions, the mic is usually placed between the nipple-line and chin.

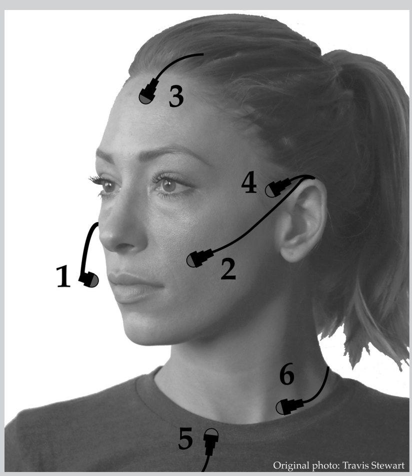

Figure 19.6 Lav mic positions discussed in the text box.

6) Neck positioning, below the chin, on or in a collar, puts the mic significantly off-axis to the mouth, and close to the chest resonances, causing greatly reduced pick-up above 1 kHz, dramatically reducing intelligibility.

There are various products available to separate fabric layers and discourage rubbing and rustling when mounting a lav mic under or in clothing, including:

- *Stickies* – double-sided adhesive patches used to attach the mic to clothing.
- *Undercovers* and *overcovers* – furry layers stuck between the mic and clothing layers.

If these are not available, a triangle of gaff tape on each side of the mic can substitute, as shown in **Figure 19.7**.

Medical tapes and dressings such as *Transpore* and *Tegaderm* can also be used to stick lav mics directly to skin. Use one piece to form a base layer on the skin, put the mic in position over it. Secure it with another piece of tape stuck over the first one, as shown in **Figure 19.8**. If a lav mic needs to be positioned on a particularly hairy or sweaty chest, *chest straps* are available. An improvised strap can be created by using a loop of tape around the entire torso, with the mic fixed to that with another layer of tape.

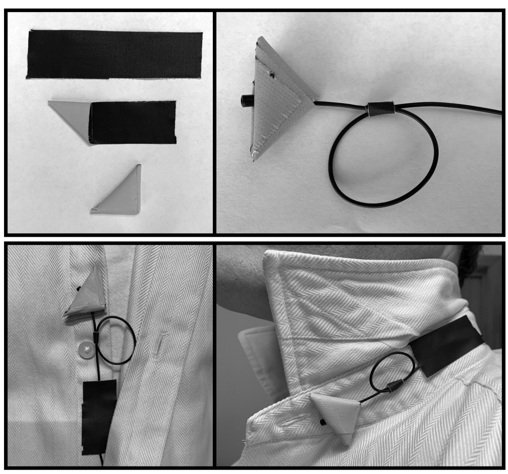

Figure 19.7 Gaff tape triangles on each side of a lav mic capsule help prevent fabric rubbing and rustling noises. **Top Left:** Take a 9 cm (3 in) strip of 3 cm (1 in) tape, and fold it corner over corner repeatedly to form the triangle shape, making sure the adhesive is on the outside. **Top Right:** Sandwich the mic between two of these triangles. **Bottom:** These or commercial products can be used to mount lav mics between layers of clothing, including between shirt buttons, and under shirt collars – the mics will be invisible when the buttons are done up and the collar is folded down.

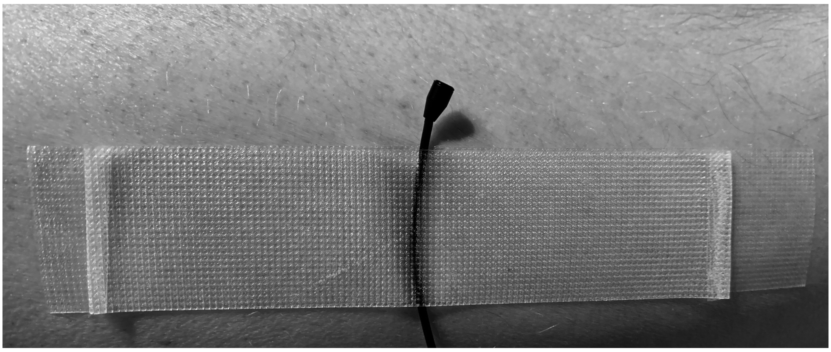

Figure 19.8 Mounting a lav mic directly to skin. The edges of the second layer are folded over for visibility in this picture. A loop of cable and second strain-relief strip of tape are below this, out of the picture.

LOOPS AND SLACK

Always form a small loop of cable near the mic to reduce cable noises. Additionally, be sure to leave a little slack cable between the mic and transmitter pack, so the mic is not moved or pulled off as the actor moves their body.

It is not always possible to use a lav mic on an actor – clothing (or lack of it) and actor action being just two reasons. Long, wide video shots make it impossible to use boompole mics close enough. So get creative! Lav mics, or other small microphones can be hidden in set pieces or prop items located close to the actors. For the best sound, the mic should be kept on-axis and relatively close – for example hidden in a potted plant on a table the actor is sitting at. Ensure necessary sound is picked up effectively by concealing whatever mic you can, wherever it can be hidden!

AUDIO EXAMPLES
Can be found on the companion website

Production Dialog Recording

These examples are deliberately recorded in an environment with significant background noise.

Example 19.1: Dialog recorded with an on-camera mic.
Example 19.2: Dialog recorded with a cardioid condenser mic on a boompole.
Example 19.3: Dialog recorded with a video shotgun mic on a boompole.
Example 19.4: Dialog recorded with a lavalier mic attached to the speaker's shirt.

Wireless Transmitter Placement

If wireless lav mics are used, their transmitter packs should be carefully placed:

- A convenient place to put a transmitter pack is in an inside jacket pocket. The mic cable can be hidden by running it under the shirt from the mic, all the way down to the waist, then outside the shirt, up and under the jacket to the transmitter.
- Transmitter packs can be clipped to waistbands and chestbands hidden under clothing.
- If it's not possible to attach the transmitter to clothing, bandages or straps can be used to wrap the transmitter around the talent's leg.
- Antennas must be free and straight – flexible antennas don't work well when bundled up or trapped behind the transmitter pack.
- Sweat and body contact can drastically reduce the effectiveness of radio frequency (RF) transmission, so the transmitter and antenna should be located where the talent will not be getting wet. Put dry layers of clothing between the actor's skin and the antenna.
- Transmitter packs can be protected from sweat using wireless transmitter sheaths, or unlubricated condoms. They should be put on so the antenna hangs out of the open end, which will become the bottom of the pack and hang down, strapped to the actor.
- Transmitters and receivers should be in "line of sight" of each other. RF signals do not travel through bodies or walls well.
- There *is* such a thing as "too much" RF signal, and it can cause random noises and dropouts. If the RF level meter on the receiver is at maximum, either reduce the transmitter RF power (if it is switchable), or move the receiver further from the transmitter.

Always run the lav mic cable first, and then attach it to the transmitter – it's just too awkward to organize if the two parts are connected together. If a film or TV production has a wardrobe or costume department and a particular character, actor or costume makes it difficult to keep the cables hidden, talk to the costumers – they may be able to make hidden holes in the clothing for cables to run through. A *Lav-Bullet* or fishing weights can be temporarily attached to mic cables to encourage them to run down or between clothing layers – it's easier to run cables from top to bottom using gravity as assistance!

WORKING WITH THE TALENT

Positioning a lav mic on someone's body, bare skin, or under clothing, makes miking talent a potentially intimate process. It's important to be professional and respect boundaries, personality types, and any guidelines you get from the talent.

- Be polite and calm, but confident.
- Talk to the talent, tell them what you are about to do, and why you are doing it.

- Ask the talent what they're comfortable with you doing when you need to run cables under their clothes, potentially touch them, or clip or stick mics to their bodies or to clothing they are already wearing.
- The talent may be happy to help put the mic on and run the cable themselves.
- The costuming or wardrobe department may be able to help when dressing the talent.
- Be sure to check everything feels good and will not impede the talent's performance. Have them move around a little to check the mic, transmitter pack, and cables don't move or come loose.
- If using a wireless mic, show the talent the mute button on their transmitter pack so they can have privacy during a personal conversation, phone call, or toilet break!

19.9 Voice-Overs and Dialog Replacement

Voice-Overs

Voice-overs consist of commentary or explanatory dialog from an off-screen character. This is usually recorded and added to a production after filming. Voice-overs are also used in video games.

- *Asynchronous* or *timed* voice-overs have to be an exact length (they have specific start and stop times), however the exact timing of the words and sentences between the start and stop times is not critical.
- *Synchronous* or *Synchronized* voice-overs have to be a specific duration, *and* elements within the script have to line up and synchronize with specific on-screen actions or hit points.

Voice-overs are read from a script, and the recording equipment setup is similar to other vocal recording – in a booth in a studio, frequently with a large diaphragm condenser mic. In rooms with less desirable acoustics, or low-level background noise, a good quality dynamic mic, used close, can capture a good voice sound with less extraneous noise. Always use a music stand or other device to support any scripts. The voice talent can sit or stand – they should be comfortable. If the voice talent is looking down at a script or monitor screen, the mic should be positioned lower, so they are still talking "into" or just across the top or side of the mic when they are looking at the script or screen.

For asynchronous material it can be useful for the talent to see a clock or timing device while recording. For synchronous material it is essential that they can see the on-screen action they are trying to sync lines or words with. It's best to get it as correct as possible during recording, rather than rely on editing and re-timing in post-production.

Automated Dialog Replacement

Automated Dialog Replacement or *ADR* is the process of replacing actor dialog recorded on set, with dialog re-recorded in a studio. The original or a different actor may do the ADR. A sound editor or ADR supervisor will check all of a production's dialog and decide what needs to be re-recorded. There are many reasons ADR might be necessary, including:

- The dialog recorded on set may have technical problems or undesirable spill, not noticed or impossible to get rid of during filming.
- There may be dialog changes decided upon after filming.
- There may be unsuitable but unavoidable extraneous sounds in the original recording that cannot be in the production.
- Diction, accent, timing, or other delivery characteristics of the lines might need improving.
- The actor might have made mistakes not noticed during filming.
- Expletives may need to be replaced for "family friendly" broadcasts.
- The production may be getting *dubbed* into another language for foreign distribution.

ADR involves looping a short section of video while the actor watches it and delivers lines in sync with the on-screen action. If they are trying to recreate the original performance, it helps if they can hear the original recording while rehearsing. Cue beeps or on-screen timing visuals are used prior to the line, to signal the actor to be ready and start delivery. DAW's playback video, making them suitable for ADR. DAWs made specifically for post-production have specialized ADR features.

When setting up for ADR sessions, it is important to:

- Use a mic that will match the timbre of the actor's on set dialog.
- Position the mic so there is a similar proximity to the sound, both in terms of proximity effect, timbre, and also room acoustic (or plan to use artificial reverb to match room acoustic).
- If the actor is looking at a screen or script, position the mic so it is on-axis and between their mouth and the screen or script as they are looking at it, but without obscuring their view.
- Put scripts on a music stand or other support to avoid picking up rustling and other noises when handheld.
- Use a pop filter.

- Have notes and clearly annotated scripts, so you know exactly what needs replacing and what doesn't.
- Know the material and the goals of the ADR session… Why do the lines need replacing? What is desired in the replacement recordings? If a director or ADR/editing supervisor is not present, this also means producing the session and coaching the talent to deliver the required performance.

ADR Editing and Alignment

Sometimes an actor can almost, "but not quite" synchronize their delivery with the mouth movements of a facial close-up. If automatic alignment software or plug-ins are not available, tools found in every DAW can be used to manually edit and tweak the timing of specific lines:

- Separate the line into smaller fragments: those that are OK, and those that drift.
- Nudge the fragment that drifted – forwards or backwards so its beginning or conclusion starts or ends in sync with the picture.
- Time compress or expand the fragment to put the other end of the region in the right place, and more closely match the mouth movements on screen.
- Apply short crossfades where the fragments join, to avoid clicks and glitches.

Alignment software or plug-ins can automate this process. Some post-production DAWs have these features built in. The process requires the original dialog audio track (the *guide* or *reference*) to be clear enough for software analysis, so that the ADR track (the *dub* or *target*) can be processed and aligned to the original.

19.10 Isolation and Restoration Software

If the delivery and performance of the original on set dialog is okay, there are several software packages and plug-ins that can very easily and quickly remove light to moderate problems such as:

- Excessive background noise (continuous noises, or not enough dialog isolation with variable extraneous noises).
- Clicks and pops.
- Wind noise.
- Plosive booms.

- Distortion.
- Too much reverb.
- Excessive breath sound.
- Clothes rustling on a lav mic.

It is even possible to adjust the tone and intonation contours of lines to make them more confident, or more questioning. The most flexible full versions of this type of software are rarely cheap, but they can cost less than needing studio time and personnel for multiple ADR sessions!

AUDIO EXAMPLES
Can be found on the companion website

Production Dialog Isolation

Example 19.5: The previous shotgun mic dialog processed (heavily) with de-noise software.

19.11 Sound Effect Recording

Sound effects, or SFX, are needed for many reasons, including:

- Some sounds don't exist in the real world!
- It's not possible to mic some necessary sounds on set.
- There's too much extraneous noise on set.
- The sound of an object on set is not suitable or dramatic enough.
- It's not cost effective to mic and record on set sounds with a full crew present.

Sound effects can be:

- Real sounds recorded somewhere.
- Created by layering together many individual sounds to form a more complex texture, or sequence of sound events.
- Performed by foley artists and recorded in a foley studio.
- Sourced from a commercial sound effects library, or the sound designer's own collection.

SONIC AUTHENTICITY?

Sound effects can come from anywhere and be recorded anywhere! For most film and TV productions, suggestion, emotion, impact, and storytelling are more important than sonic accuracy.

- SFX can be recordings of sound sources directly related to the on screen action – for example, if a scene is shot in a train station, the ambient sounds are also recorded in a train station.
- They can also be unrelated sounds that simply suggest the right meaning and reinforce the on screen activity – a breaking bone in a horror movie is not actually a breaking bone, but might be a layering of sounds made by snapping vegetables.

Sound designers build up libraries of source sounds they can edit and process into spot effects and ambiences for different productions they work on. A cheap portable recorder with built in mics (or even a mobile phone with a decent mic attached) might be good enough to record louder close up sounds. Better mics and preamps are definitely required for quiet sounds and distant ambient sounds. Shotgun mics on stands, or suspension mounted in pistol grips, are convenient for location recording and sound gathering. In the studio, condenser mics produce accurate detailed recordings.

A *foley studio* has a large assortment of every day and unusual objects that can produce the sounds needed to make video scenes convincing. For example:

- If a glass is being filled on screen, it needs sound. The live sound may not have been picked up well by any mics on set, so a *foley artist* "performs" and records the glass filling in a foley studio.
- A monster smashing somebody's head and removing their brains on screen can be created by close miking a melon being hit with a hammer, and then squishing the insides of the melon with a hand.

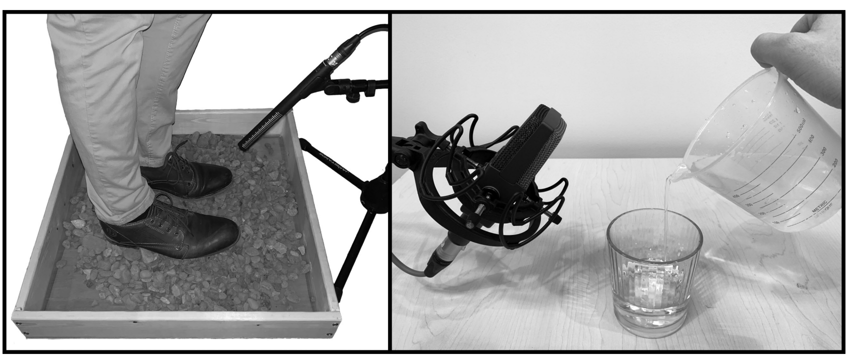

Figure 19.9 **Left:** A shotgun mic set up to record footsteps in a makeshift foley pit. **Right:** A condenser mic set up to record a close perspective of liquid being poured into a glass.

Foley studios have *foley pits* containing different materials used to make ground-related and other assorted sounds, each pit containing different materials:

- To make a walking scene more realistic, a foley artist could walk (or move shoes) in one of the gravel, dirt, or leaf pits in time with on screen foot movement.
- The foley artist could rub clothes together close to a mic, in time with an on screen actor's movements, to create and add the sound of moving clothes, and add intimacy to the scene.

B-roll sound is often recorded in multi-channel surround or immersive formats. This requires a multichannel recording device. Convenient and compact single-body 5.1, 7.1, and Ambisonic microphones are available.

VIDEO PEOPLE NEED AUDIO PEOPLE!

Try watching a dramatic or action-packed TV show or movie *with the sound turned off*. So much is missing! That's the importance of good location and post-production audio – it gives meaning to the visuals, and tells the viewer what their emotional reaction should be.

20 Tips From the Professionals…

In This Chapter:

20.1 To Conclude…

Specific models of microphone or other technologies have not been mentioned much up to this point – so that you can apply the information presented throughout this book, listen closely to the results of your own experimentation, and come to your own conclusions on how to capture the best sound for the project you're working on. Audio professionals do love to talk about their favorite gear though! So let's close with some invaluable real world tips and tales from some professionals representing a variety of sectors of the industry. As you read these interviews you'll notice common practices, as well as widely varying opinions! One of the many great things about recording is that there's no single correct piece of equipment for a particular job, and no one correct way of doing something. As the interviews show, the essential skills any successful recording engineer should have are pretty universally acknowledged – but they are not all technical skills….

20.2 Lenise Bent

PRODUCER/ENGINEER, POST-PRODUCTION EDITOR/MIXER/FOLEY, INSTRUCTOR, CONSULTANT, ARCHIVIST – LOS ANGELES, CALIFORNIA, USA

- Engineer: Blondie – *Autoamerican*. (First woman engineer to receive an RIAA Platinum record)
- Worked on: Steely Dan – *Aja*; Supertramp – *Breakfast in America*
- Producer/Engineer: Primal Kings – *Primal Kings* (all analog project, 2020)
- Foreign Dubbing Supervisor: Dreamworks.

What type of recording work do you generally do?
"I do all kinds of recording work, from music recording and mixing, to foley, podcasts, and presentations. I like to work with real musicians and bands performing live. I've been doing a few projects to tape as well lately. Analog!"

Where do you generally work?
"I mostly work in a 'proper' recording studio for tracking and overdubs – which can also be remotely recorded, or in my own home studio/recording environment. With the portability of technology and all of the software for creative online collaboration, I can potentially record anywhere. Recently, I recorded foley footsteps in puddles and real grass outside – that was

Figure 20.1 Lenise Bent

fun! For sound effects, if I hear a noise or an ambience I want to capture I can just use my iPhone. I also have a Zoom and a Tascam digital recorder. They both can record 5.1."

How long does a typical project last?
"There is no typical project, it depends on what is being recorded. I've worked on projects taking a few hours, or a few days, or several months."

Anything mic technique-related that unexpectedly pleasantly surprised you?
"On one tracking date I'd chosen two or three mics to shoot out and audition as the inside kick drum mic – an AKG D112, a Blue, and a Coles. A Shure SM7 was sitting about 3 ft away, on a tiny stand. In the control room, the faders for the D112 and SM7 just happened to be up on the console when I walked in – and the sound was so rich and fat! I hadn't intended to use them both, and the SM7 hadn't been deliberately positioned, but the sound was so good I said, 'Let's do that!'"

"On another occasion, I'd just finished assisting on a session at The Village Studios in West Los Angeles, and had to tear that down and set up for a morning tracking date that included a Hammond B3 organ and a Leslie. I'd put some AKG 414's on the Leslie, side and bottom. While I was checking the mics in the control room, a colleague was noodling on the piano, which was 20 ft away in the opposite corner to the Leslie mics. The piano sound from these distant mics on the Leslie was so weird and freaky! I loved it! I've constantly looked for an opportunity to use this technique, but haven't had one yet…!"

How about any mic techniques that didn't work as well as you hoped?
"So many singers have heard about the Neumann U47 tube mic, and they want to use it. So I will do a blind test with the artist, using a U47 and some other mics I've determined might sound good on their voice. The favorite so far has never been the U47! It's usually the U67."

"I used a U47 with Debbie Harry, for an old time, 20s sound. I'll often mix up vocal mics throughout an album – Neumann U67, U87, or even an AKG C414 for vocals on a brighter song."

What is the key to recording great sound?
"Capturing the performance the best way you possibly can, meaning using the right mic for that instrument or voice, put in the right place. And I prefer to work quickly while the energy is up."

What's the most important part of the recording signal chain?
"It's the source – without the source, nothing else matters. If you don't have a good singer or guitar player, or if you're recording a great piano player playing a lousy piano, you're not capturing anything good. Also, if it's not a good song, even great musicians and singers can't make it a good song."

What pre-production do you do before a session?

"I determine what songs, and how many we're going to record. We rehearse song structures, tempos, keys, and adjust lyrics to get the basic structure intact before we go into the studio. Once I know instrumentation, I figure out what mics, what studio to record in, and where in the studio to set things up. What will the basic tracking be, and what will the overdubs be? I like to get the basics down so we don't have to think about that in the studio – and then the magic can happen."

What are your favorite mics and preamps for some different applications?

"My favorite mic, if I could choose only one, would be a Neumann U67. It's quite universal, great on vocals, guitar amps, mandolins, stand-up bass – just about anything. I realize they are a bit precious pricewise, and after a few hours of vocalizing they can start to break up as the capsule gets moist, but just about everybody sounds good on one. There are several emulations out now, the Sterling ST69 being one, and it is surprisingly similar and affordable. I also like AKG 414s, Shure SM7s, and Neumann KM84s."

"It depends on what I'm recording and what's in the mic locker where I'm recording. For preamps I'm a big fan of Neve 1073s. The studio is selected based on the quality of the room, the gear, and the budget."

"I prefer not to use any EQ when I record, and I keep things as simple as possible so I can mix as I'm recording. I work very quickly and want to keep the session moving smoothly to keep creative energy and spirit intact. People forget that the technical equipment is there to serve the art, not the other way around. It's not about the gear – it's about the music, the performance, and capturing the emotion the best way possible."

What's your approach to recording in a less than ideal room?

"It's never completely impossible, though it may be easier to record in a better room. It's important to really know your mics, and where to put them for the best sound. Use your ears! Use comforters or pillows to reduce standing waves and reflections in a bedroom, and turn off the AC."

What's your attitude to using less than top-of-the-line equipment?

"You just have to think 'how good can we make it?' Determine what's less than ideal, and how you can upgrade the situation to be professional. I'm not going to pretend I can polish every situation, but it's got to be pretty bad for me not to be able to improve on it at all, because my style of recording is pretty simplistic, and there aren't a lot of elements involved."

"I have standards. Not only for me, but for the people who taught me, like Roger Nichols."

What do you want people to think when they hear your work?

"That it sounds great! The quality of the recording, mixing, and mastering are all equally important to me. I'm already thinking about the mastering when I'm cutting basic tracks."

20.3 David V. R. Bowles

SWINESHEAD PRODUCTIONS, LLC – SAN FRANCISCO, CALIFORNIA, USA

- Recording producer and engineer: Philharmonia Baroque Orchestra and Chorale
- Guest Lecturer: New York University – Steinhardt, Music Technology
- Grammy and Juno award nominations, many "recording of the year" accolades

What type of recording work do you generally do?
"Classical music recording, either live or in the studio."

Where do you prefer to work?
"The best recording rooms are meant for performances. Studios can sound too sterile, and are generally too small to allow for a diffuse reverb tail. A concert hall is better for the performers. They can all hear themselves and are comfortable with that acoustic, making for a better recording. These days a lot of orchestral recordings are compiled from different live performances, sometimes in different locations. That creates unique challenges!"

How long does a typical session last?
"A three-hour session is good for a large ensemble or orchestra. You'll get about fifteen minutes of 'finished product' from each session. For chamber music, two four-hour sessions, with breaks work well. In both cases you can get a complete project done in two days. Another possibility for live recordings is a post-performance patch session, usually lasting

Figure 20.2 David V. R. Bowles

two hours. This allows coverage for inevitable audience noises and passages which weren't optimal during the concerts."

Anything related to mic techniques that unexpectedly pleasantly surprised you?
"Orchestras have a lot going on in the middle of the ensemble. I've found that spaced omni stereo pairs tend to have a hole, or wandering of the image, in the middle. Adding a center cardioid or figure-8 mic [to the main stereo array] helps to solidify the middle of the image, allowing me to space the left and right further apart, and get better separation between the left and right of the orchestra. That was a happy accident I've used ever since."

"I've always used some kind of coincident pair in orchestral recordings, to pick up winds or chorus. I've found this results in a very stable center of the image, as well as being a very accurate way to spot mic that central section."

"The 'conductor's position' [i.e., as they would hear the ensemble] is not an optimal listening position. However, a mic placed about 1.5 meters above their head is a good recording position."

Any mic techniques that didn't work as well as you hoped?
"I've found the ORTF technique to be unsatisfactory. It works fine for listening on headphones, but on loudspeakers there's not enough width, and it tends to be too dry. If I moved the mics back, there could be too much ambience, especially in a bad sounding room. With a spaced pair of omni microphones, you get accurate low frequency response and, I think, a more natural high frequency sound. Omnis do pick up a lot of ambience so generally need to be placed closer, but not so close that one hears too many individuals. It goes back to needing good room acoustics to start with, and having enough of a blend coming through without having to move the mics back too far."

"I have found it difficult to get a pair of small diaphragm cardioids to have the same smoothness as a pair of omnis. Cardioids can work well as spot mics, but they can also bring out too much sibilance or 'fizz.' In general, avoid an omni or cardioid that has a high frequency bump around 2 to 3.5 kHz."

What's the key to recording great sound?
"Listening to the performers, and asking myself, 'how do I go about reproducing that?' It is dishonest to *want* a kind of sound without listening to the performers first, then striving to bring out the individuality of each performance and acoustic. Though I start with generic techniques for picking up vocal or piano sounds, these are starting points rather than fixed in stone. For this reason it's difficult to formulate what will work for different pieces, even performed by the same group in the same acoustic – this is why I change my technique slightly from recording to recording."

What's the most important part of the recording signal chain?
"The microphones, preamps, and A to D converters are all equally important. Look for good mics. There are fewer good mics than preamps, that is for certain. What colors do

you want to get, and what do you want to avoid? Similarly, with A to D conversion, look for stable clocking, no jitter and very low self-noise specs. Test charts and measurements are important, but it's also important to try out equipment and ask colleagues what they think."

What pre-production do you do before a session?

"The questions I ask include what kind of music is being recorded, and what space is appropriate? I get copies of the printed music, and become familiar with it. If the ensemble has even a bad recording of a live performance, that will help me hear what they do with that particular music."

"Come into the studio with a mic list, set-up and seating plan, and be prepared to change out mics if necessary. Performers need to be comfortable. I'll have a seating plan, but sometimes they'll react negatively to how I want them to sit – so I have to adjust accordingly."

"Discuss the recording schedule with the artists, be flexible with how much time is allotted to each piece, but keep in mind how much music needs to be covered in each session. Always ask discreetly which music poses the most difficulty, in order to allow enough time to cover that section adequately. For initial balance, choose a section which they are comfortable with, but has loud and soft sections."

What are your favorite mics for some different applications?

"Josephson C617SET for spaced omnis. For spot mics, the Sennheiser MKH series, or Neumann TLM 170. For percussion, the cardioid or hyper-cardioid models from the Schoeps MK series."

What's your approach to recording in a less than ideal room?

"I would probably mic a little closer than usual, and monitor with added [artificial] reverb so you get an idea of how you can transform that acoustic space into something more pleasing."

What's your attitude to using less than top-of-the-line equipment?

"I would make it clear to the client that if the equipment is less than optimal, the result will be less than optimal, and that more post-production work will be necessary to improve the sound."

What do you want people to think when they hear your work?

"The ideal listener reaction is to how good the performers are, how unique their interpretation is, and then how good the sound is. In that order."

"In one project I did, different sections of the same pieces were recorded in different locations, and then edited together. It was a challenge to match the sounds from each location. No reviewers picked up that this had been done, which to me is the ultimate compliment."

20.4 Joel Hamilton

PRODUCER, ENGINEER, CO-OWNER: STUDIO G BROOKLYN – NEW YORK, USA

- Producer/Mix Engineer: Highly Suspect – The Boy Who Died Wolf, MCID
- Recording/Mix Engineer: Aaron Neville – Apache
- Engineer: Bonobo – Migration
- Six time Grammy and one time Latin Grammy nominated

What type of recording work do you generally do?
"If I'm in a room when any mics are being set up at all, generally I'm the producer, and I also work as the engineer. Then I work as a mix engineer, but it's rare I work just as an engineer, meaning that if there are microphones happening, and the multitrack is running, chances are I'm going to want to be able to give advice and help guide people."

Where do you generally work?
"I work in studios all over the world, but I prefer to work in Studio G Brooklyn, which is a multi-room facility including one of the largest orchestral rooms in New York City. Because the place grew organically from the needs that I had in sessions, it's got everything to serve me – I'm really spoiled!"

Figure 20.3 Joel Hamilton

How long does a typical project last?
"It varies so greatly… from one week for an EP tracked and mixed, up through six months for the latest major label project I worked on."

Anything related to mic techniques that unexpectedly pleasantly surprised you?
"There was a microphone during the recording of the Sparklehorse record that the primary songwriter had found at a dump! It was a little Silvertone microphone with a quarter-inch jack. We gaffed it to the body of a Fender P bass. When I went into the control room to listen to what I was getting using this little crappy thing, used almost like a contact microphone, I thought I was listening to something else because I was getting massive low end, but in this cool Hoffner-ey way. It sounded like a hollow body bass but with giant low end!"

How about any mic techniques that didn't work as well as you hoped?
"Every time I've ever put up a microphone in a room where the people were uncomfortable. Every time I've ever put up a microphone with no intention, and just kind of threw it up, and it didn't suit the context in which I was looking for a particular result."

What is the key to recording great sound?
"Having a vision for the word 'great,' and being able to define that for any given project. Then engineering in pursuit of that greatness. Without that vision I don't even understand how people would choose a microphone."

What's the most important part of the recording signal chain?
"Me…! [laughs] The engineer and the producer, without a doubt, because nothing else in that room is going to have a vision for whether it is doing a great job or not. Subjectively… [the performance] is going to be processed by me, and I'll either change the microphone choice, move the microphone, or change what I'm doing in other stages."

What pre-production do you do before a session?
"Pre-production is essentially discussing how we're going to flatter the microphones in the room by adjusting small things about the performance or the arrangement… about what the humans are doing in front of the microphones."

"If I have a band that can be in town for a few days before recording, I'll do rehearsals of the recording process. I might put up a single overhead and a kick mic, and I'll start miking things up. I work to tape, with Pro Tools on the back of the tape machine, so I'm not running the tape machine to start with. When the band has got the 'shape' of it, I'll ask them to start going for 'execution.' Then we start spinning the tape."

"If the band is not from New York, or can't afford to be here for a week before we start recording, that same concept will come into play, but I'll just do it as the first two times through the song, then I'll have them come into the control room where it's much easier to discuss when I have control over the volume."

What are your favorite mics for some different applications?

"I wind up finding that things with incredibly fast transient response on something like a kick drum have become my friend more and more over the years. Something like an Earthworks TC30 on the beater side gives me some low frequency response that's there and then not – really fast."

"I use a Neumann SM69 [stereo microphone] quite a bit, just for ease of setting up… I'll have that as an ambient microphone or piano microphone, or something where I just need a great XY or Blumlein pair. It just works! An AEA R88, dual ribbon stereo, I use for the room on drums all the time."

"I really like small diaphragm condensers on toms. They could be Octava MK-012s or Josephson E22s."

"I use a Revox M3500 on the snare top. Something ridiculous on snare bottom – a condenser, like an AKG 451 with three –10 dB pads screwed into it!"

"I have a Soundelux U95 that I've recorded some of the most compelling vocals of my career on – from Sparklehorse to Norah Jones to Mos Def."

What's your approach to recording in a less than ideal room?

"I find that most of what I record is in a less than ideal room, whatever studio it's in! Often rooms don't work for what I'm doing. I end up with gobos, temporary treatments, blankets, foam, moving blankets, a couch, pillows… it looks like a hobo moved into the drum kit whenever I record, because I'm going to want ambience in the sectional mics [overheads and room mics] and isolation in the spot mics. I've never just put up microphones and not had to make it like 'Foam-Henge' around the drums to have it do what I need it to do to serve the production vision."

What's your attitude to using less than top-of-the-line equipment?

"I think there's an idiomatic response developed from a thing that didn't break when AC/DC showed up, or the thing that wasn't broken when David Bowie arrived, or the thing that was working and turned on when somebody did something legendary in front of it. A Neve 1073 preamp is a perfect example of something that carries provenance more than actually being 'better' for the record. Using things that are inferior has its moment – exploiting what's great about them. The frequency response is probably OK, the transient response is probably ok, but it will probably break before a 1073 under the rigors of what a professional studio puts it through."

What do you want people to think when they hear your work?

"I want people to feel the compositional intention was respected and considered when choosing the mics and other equipment that I used to exalt that gesture… I want people to feel like when a musician is extracting all they can from their instrument, that I extracted all I could from the musician.… And that requires microphone choices and equipment choices that flatter the moment and are used in the service of the vision of the project."

20.5 Kerry-Anne Kubisa

FREELANCE LOCATION SOUND RECORDIST – UK

- Former Certified Apple Logic Pro and Avid ProTools Trainer
- Recording Engineer for the Leeds International Piano Competition, 2009 and 2012
- Founder and past Chair, Audio Engineering Society – North of England section

What type of recording work do you generally do?
"Either live concerts, on location, usually in concert halls or churches, or location recordings with studio based production and without an audience. The genre of music can be varied, from opera to early music ensembles, from orchestras to string quartets, jazz to electro-acoustic and contemporary ensembles."

How long does a typical session or project last?
"Sessions can vary from four to five hours for a simple one-hour live concert, including set-up, soundcheck, the concert itself and pack-up, to sessions which can last 12–15 hours, or even days."

Anything related to mic techniques that unexpectedly pleasantly surprised you?
"When I was a student, recording a band in a pub and wanting to capture the liveness and feel of the environment, I put up a pair of Sennheiser ME-80 shotgun mics. I positioned them to the left and right of the room, either side but behind the house PA, and aimed them at the audience. There wasn't anything more technical to it than wanting a mic that

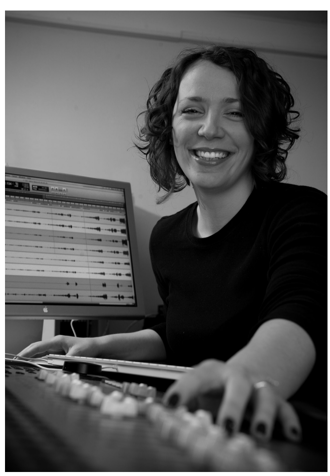

Figure 20.4 Kerry-Anne Kubisa

would pick up sounds further away. They are not the best sounding shotgun mics out there, but when I got the recording back in the studio, panned them hard left and hard right, and brought them up in the mix, wow, they gave me exactly what I was looking for. Obviously I hadn't invented that type of mic technique, it's quite common, but at the time, I felt like I'd found the holy grail of recording!"

"I was working on an electro-acoustic/contemporary ensemble concert, which combined a front of house PA and some instruments that played acoustically without any amplification. Doing a recording wasn't a requirement of the concert but I did feel this was an opportunity to do some mic technique-related experimentation! I wanted the production to be very natural, to be recorded from the perspective of the audience, and not to interfere with the stage set-up in anyway. So I put a pair of miniature DPA 4060s, with their windshields on, in my ears and voilà, I had a dummy head! I sat very, very still in the sweet-spot of the hall for an hour – and the results were astonishing! There was no transference noise from them being placed in my ears, the stereo imaging was spot on, and the balance of the instruments was perfect."

How about any mic techniques that didn't work as well as you hoped?

"In the beginning of my classical recording career I was asked to record historically performed twentieth-century violin technique performances for an academic project to be submitted to the British Library. I proposed that we record from the perspective of the listener, with no close miking. I setup a Soundfield ST250 in B-format, a pair of DPA 4009 as a spaced pair, a pair of DPA 4011 in ORTF, and a binaural KU100 Neumann Dummy Head. Once we got back into the studio to listen to the initial recordings it became clear that there was just too much going on, and phasing between the pairs was a big problem. In the end we didn't use all the stereo pairs for each ensemble, just picked the best ones for each and that resolved the phasing. We also experimented with close up miking, and in the end went for a typically modern approach – a combination of stereo recording techniques *and* close up miking."

What's the key to recording great sound?

"Listening! You can have the best of everything but without the ability to listen to the music or sound of what you're recording you will never be able to get a good sound. You must also listen to what the artist wants. There's no point going all guns blazing with what you think is the best way of doing things if the artist goes away unhappy."

What's the most important aspect of the recording signal chain?

"Reliability! Well-maintained, solid equipment is key to a successful recording. Without faith that your gear will work time and time again, seeds of doubt and uncertainty set in. The last thing you want to be worrying about is, 'mmmm that input was a bit temperamental last time I used it, please don't die on me now!'"

What pre-production do you do before a session?

"Before I commit to any recording I sit down and discuss what the artists' expectations are, to make sure that my skills and approach to recording that genre of music, and what the artist wants from the recording are the same. You don't want to find out down the line that you can't meet the artists' expectations."

"I ask for the score and listen to recordings of the pieces the artist wants to record, so I know the repertoire thoroughly before I start."

"A visit to the venue, if I haven't worked there before, is also essential. You need to look at access, power outlets, cable runs, health and safety issues, mic placement, etc."

What are your favorite mics and preamps for some different applications?

"My favourite preamp is the Earthworks 1024, and the mics I would usually reach for are Neumann and DPA, though AKG C414s are put on anything that's brass."

"I would typically use DPA 4009 or DPA 4006 for spaced pairs, DPA 4011 for ORTF, XY, etc."

"I love KM184s on strings and grand piano. They sound amazing and are my first choice for anything string related."

What's your approach to recording in a less than ideal room?

"I am often put in the position of recording drums *live* in a classical concert hall. The obvious way to avoid the drums becoming a muddy, noisy, over-reverberant mess is close miking. Maybe not so obvious is for the drummer to adapt what they play – instead of playing like 'Animal' from Sesame Street, they can use rim shots, brushes, tinkle the cymbals, or simplify the rhythms."

What's your attitude to using less than top-of-the-line equipment?

"Tricky. I think I would have to weigh up the negative and positive impact the project would have on my reputation, and either run away, or jump in feet first, depending on the odds. If you've got good ears, I think you could make the right sound, in the right room, sound great with a bag of SM58s and SM57s!"

What do you want people to think when they hear your work?

"I would like to say that I hope that they would hear the passion and commitment that went into the recording! Really though, what they should hear is the music. The listener should feel immersed in it, like they are sitting in the venue it was recorded in, surrounded by the audience and ambiance of the space, and forget about everything else!"

20.6 Wes "Wesonator" Maebe

ENGINEER/PRODUCER: SONIC CUISINE, RAK STUDIOS – LONDON, UK

- Mix Engineer: UB40 – A Real Labour of Love; New Model Army – Something Like Jesus, 30th Anthology
- Recording/Mix/Mastering Engineer, Producer: 10 Gauge – Rain's Coming
- Mix/Mastering Engineer: Erin Dickens (founder, Manhattan Transfer) – Vignettes

What type of recording work do you generally do?
"Anything from baroque to death metal, blues, jazz, pop, film scores, old hard rock, and stonking techno!"

Where do you generally work?
"A lot of my studio-based tracking happens at RAK Studios and British Grove, and then it's back to my Sonic Cuisine facility for overdubs and mixing. Classical sessions usually take place in concert halls and even stately homes. I've found there's always a road near churches, but stately homes are usually in the middle of nowhere, so they're quieter!"

How long does your typical project last?
"The budget dictates the length. Most are a couple of days of tracking in a studio, and then back to Sonic Cuisine. The longest single job I worked on was about two months of recording and mixing at RAK."

Figure 20.5 Wes "Wesonator" Maebe

Anything related to mic techniques that unexpectedly pleasantly surprised you?
"When I first used RAK's Studio 2.… It has a very high ceiling, with the control room upstairs from the live room. The height was amazing for drums! I tried the Glyn Johns technique, with one overhead, one mic across from the floor tom, and a kick mic, plus some other spots. I put two Coles ribbon mics all the way up in the ceiling. I didn't need any artificial reverb. At the time, I hadn't done that before."

"Also on drums, I've put an SM57 or SM58 behind the drummer's elbow, in addition to a closer snare mic. I've had drummers comment 'Are you out of your mind? That's going to phase.…' I've used that mic as the main snare mic, and got the best sounds out of it!"

"I like to mic guitar cabs from the back, flipping the phase of the mic. It's been good on a Vox AC30 amp, but worked best on a Fender Concert."

How about any mic techniques that didn't work as well as you hoped?
"I started out as a classical balance engineer. Nothing was close miked in those sessions. My first time in the studio I close miked everything expecting full control of everything. But nothing breathed. There was no space. It was claustrophobic. The solution was to de-rig and start again. I went back to a classical approach."

"I once had a teacher tell me you need to walk the room and listen – that's really it. When I'm recording, I'm running around so much I lose weight! A lot of engineers sit and reach for EQ – I run out and move mics."

What's the key to recording great sound?
"Good players. The musician is more important than the mic. You can put a shit mic in front of a decent player and it will still sound cool. The best mics don't help bad musicians. The musicians must know how to deal with their instruments.… An amazing drummer with shit drums, or a good guitarist with a bad sound, it doesn't matter where you put the mic, it translates as bad sound."

"An engineer who knows how to use a mic is more important than the mic."

What's the most important part of the recording signal chain?
"That's a tough one, but I think the mic and preamp. I track virtually clean and don't use compression or EQ when recording. That's the same for analog and digital."

What pre-production do you do before a session?
"If it's a band, I go to rehearsals if I can. Hopefully we can tidy up some loose ends like song structures, and figure out a way to play songs so they pick-up into the chorus. If the

band has mediocre equipment, it gives me an idea of what we'll need in the studio, so I'll come up with a 'hire list,' or ask the studio if we can use their gear. I also put together a 'wish list' of mics, preamps, compressors, DIs, and EQs, and funnily enough, mic stands! You'd be surprised how quickly you run out of little kick drum stands. For the Glyn Johns technique I use two heavy Brauner VM1s, so I also need enough heavy stands for heavy mics."

What are your favorite mics for some different applications?
"Brauner VM1s on drums. My favorite [stereo] technique for anything else is MS with the VM1s. I like the AKG 451 on hi-hat and acoustic guitar. I find the Neumann U47 valve mic very versatile on many vocalists. The BeesNeez Jade has never let me down on vocals. Royer R121 and Neve RNR1 on guitar cabs. Sennheiser MD421 on toms, Shure SM57 on snare. I also like a Neumann KM184 just above the rim of the snare, horizontally, blended with a 57. That's fun because it freaks drummers out…."

What's your approach to recording in a less than ideal room?
"It's possible. Here in my little Sonic Cuisine studio I record vocals using the SE reflection screen, which kind of works. I also use the pencil trick. [Rubber bands, or tape hold a pencil down the center of the mic, over the diaphragm, instead of a pop filter.] I hang up throws to get rid of reflections from windows."

"I had to record an eight piece jazz band in a very small room. I used only eight mics, because I only had eight inputs. It was a terrible room. I miked closer, and used polar patterns more, using figure-8 and hyper-cardioids more than I usually do, and moved the musicians around to exploit those polar patterns more."

What's your attitude to using less than top-of-the-line equipment?
"Fine! Let's do it! Blaming the equipment is the easy way out. You can make decent recordings with half decent equipment. If I know in advance, I'll take a few things with me, like my LunchBox [a powered chassis, originally developed by API, which can house a customized selection of boutique preamp, EQ, and compressor modules], and a few mics."

What do you want people to think when they hear your work?
"They should hear open, organic sounds. I've been described as producing an edgy British sound, combined with smooth American production values. I like straddling those two markets."

20.7 Matt Ross-Spang

PRODUCER/ENGINEER/MIXER: SOUTHERN GROOVES – MEMPHIS, TENNESSEE, USA

- Co-producer/Engineer/Mixer: Margo Price – Midwest Farmer's Daughter, All American Made
- Engineer: John Prine – The Tree of Forgiveness
- Co-producer/Engineer/Mixer: Calexico/Iron and Wine – Years to Burn
- Engineer: Jason Isbell – Something More Than Free, The Nashville Sound

What type of recording work do you generally do?
"I record bands of all genres, but mostly Americana, rock'n'roll, and rockabilly! Nowadays genres are so blurred anyway. Mostly it's the whole band live, with minimal overdubs."

Where do you generally work?
"Ninety percent at Sun Studio. It's a very historic location, unchanged from the 50s. There's no isolation, so it's all about mic technique. I use the same type of gear they had: 1936 RCA tube console with five mic inputs going live to a mono Ampex tape machine or lathe. I also have a little Pro Tools set up, a 16 input Studer console, and a one-inch 8 track Studer A800, if the band wants that sound, or they're scared of live to mono. I rarely use more than eight mics at a time because of the nature of the room. The vocal mic is also the drum room mic – everything multipurposes!"

Figure 20.6 Matt Ross-Spang

Anything related to mic techniques that unexpectedly pleasantly surprised you?
"Tambourine is one of the hardest things to record and fit in a mix, especially when you put the tambourine right on the mic. I left an RCA 77 ribbon mic up in the room as a vocal mic, and the guy was playing tambourine 20 feet back from that mic while I was patching in the control room. It sounded amazing! Now I always do tambourine that way – it mixes perfectly."

"The quieter you get someone to play drums, the farther away you can mic it and the bigger you can make it sound! That seems backwards, but if they play too loud, and you mic it too close, you get too much attack and not enough overtones. The distance gives you depth which you can't fake with a plug-in."

How about any mic techniques that didn't work as well as you hoped?
"When I was first engineering I thought in terms of name brands, monetary value, and industry staples. I would have bigger clients in and I'd try to use those mics because I thought I should – but I didn't like them, and they soon gathered dust in my closet. I love mics that have a unique character."

"Don't spend $5,000 and think you've got to use that mic on lots of things. A lot of times something cheaper will sound better. Do whatever you think sounds good."

What's the key to recording great sound?
"Because the room at Sun is so small, I have the whole band play together when sound-checking. I don't spend a lot of time EQ-ing or compressing. Instrument and mic positioning is a game of inches, and you have to treat the room as an instrument. Get it right up front, and you know you're golden later on."

"I spend an hour, not a day, getting sounds dialed in, and I spend that time running around moving mics until I'm happy!"

What's the most important part of the recording signal chain?
"Whatever you're putting the mic on. And it's really the person playing. I have an old 60s blue drum set at Sun. It can sound amazing some nights, and others it can really bug me. It depends who's playing it – the tone is in the fingers of the player. And if you know your microphones, you know what to put up to accentuate the instrument's tone."

"A great microphone or preamp isn't going to save a bad acoustic guitar. But a great acoustic guitar played right will sound great miked with just about anything, if you mic it right."

What pre-production do you do before a session?
"Little to none. I think the art of pre-production is lost on most bands these days. Forty-five minutes after the band comes in we're usually cutting."

"Sometimes bands send me demos, but I don't necessarily listen to all of them. I like my first impressions to be in the studio since Sun is so different than most studios – that's why they come. You have to adjust to the room, and while I'm miking we're talking about arrangements."

What are your favorite mics and preamps for some different applications?
"I don't have many mics, but my go-to mics are: Western Electric 639B [a ribbon and dynamic microphone from the 40s], Electro-Voice 666, RCA 77D, RCA 44, Altec M11, SM7, Ampex 1101, good ol' SM57, and my modded American R331 is the most beautiful acoustic guitar mic I've heard in my life."

"I love the RCA and Studer preamps on the consoles, but also have some tube Ampex 351s and funky old portable tube tape machines I use for certain sounds."

What's your approach to recording in a less than ideal room?
"A 'less than ideal room' has different meanings to different people. When I listen to major hits these days I don't hear character or the uniqueness of the room – it's all isolated tracks. Studios in the 1950s and 60s, like Sun and Royal Studios have a tone I really like – you can't fake that with digital plug-ins. You could tell where a record was cut by the echo chamber! I miss that!"

"If the artist is inspired, then the room is now ideal, and there's a way to mic it and make it work."

"People are recording at home all the time, in less than ideal rooms, but they are getting great performances and often times those rooms have character. Sure, you may have to work harder sound-wise, but that's a good thing!"

What's your attitude to using less than top-of-the-line equipment?
"The equipment should be invisible, and if you don't have the nicest microphone or whatever, that's not an issue. The goal is not to have nice toys to record with, the goal is to get a great performance. You just have to work really hard, and move the microphones until you get a good sound that the artist is happy with. Most artists don't get excited about the gear, they are excited about their songs."

What do you want people to think when they hear your work?
"Well first, I hope they think it sounds good! I'm such a fan of recordings by Sam Phillips, Rick Hall, Chips Moman, and Willie Mitchell, because five seconds into the track you know who engineered it and where it was cut, and then the rest of the song you think 'how the hell did they do that?' Not to mention they managed to do that while always serving the song! I hope my sounds inspire others the way those guys inspired me!"

20.8 Mark Rubel

DIRECTOR OF EDUCATION AND INSTRUCTOR: THE BLACKBIRD ACADEMY – NASHVILLE, TENNESSEE, USA

Owner: Pogo Studio. Engineer, Producer, Author, Musician, Audio Expert Witness

- Author: "The Great American Recording Studios of the 1960s and '70s"
- Made hundreds and hundreds of records – continually learning from the musical, interpersonal, and cultural process that involves
- Being part of two fantastic musical communities: Champaign, Illinois, and Nashville, Tennessee – and the audio community at large

What type of recording work do you generally do?
"I record all different styles, from death metal to Korean music to bluegrass, and everything in-between. I have also done a good deal of classical recording in concert halls – everything from soloists to symphony orchestras. I mostly do music recording, but also voice-over work, video game soundtracks, and forensic audio – you name it!"

Where do you generally work?
"Mostly I record at my own Pogo Studio. I like to use it because I think of it as a musical instrument – I know its capabilities and I'm able to quickly get whatever sound I'm looking for."

Figure 20.7 Mark Rubel

How long does your typical session last?
"Infinitely variable! It might be an hour, or more lengthy, and like everyone else in this business I've also had a periodic 32 hour session."

Anything related to mic techniques that unexpectedly pleasantly surprised you?
"Hanging a small diaphragm condenser mic in a five gallon water jug, in front of a bass drum. I use that on almost every session. It gives you as much sub as you get from using a speaker in reverse, but it has more resonance so the low frequency tone lasts longer – it makes the bass drum sound something like a Roland 808!"

"Occasionally for fun I will use a cell phone as a delay unit. I'll have a cell phone calling another cell phone, one sitting near the drum kit, the other isolated so I can mic its little speaker. You get a delay from the process of transmission, and because it goes through so many filters that are meant for voice you often get odd, syncopated, alien vocal sounds."

How about any mic techniques that didn't work as well as you hoped?
"I once did a workshop called 'Wacky Tracking' – the idea was to try and come up with as many silly, ridiculous, but potentially useful techniques as we could! It was actually difficult to come up with something new – I found just about everything I thought of had been done by Joe Meek in the 60s."

"In trying to come up with the most outlandish thing I could think of, I decided to put a microphone in a helium balloon, thinking it would be a pitch transposer. It turns out it doesn't work. I thought it was going to be fantastic, but it just sounded like a microphone in a balloon!"

What's the key to recording great sound?
"Listening. Listen to the instrument one's recording in order to know what the source sound is, even if you're going to change it. Listen to it from all directions and understand where the sound you like or don't like is. Listen to what the musicians have to say about their music."

"Listen to your internal voice – the process of recording is realization, trying to get the sound you have in your mind into the real world. Listen to your gut and move instinctively – is a sound pleasing you, or affecting you emotionally? If not, how can you get it there?"

What's the most important part of the recording signal chain?
"The instrument is the origin of the signal, and it's most important. If it sounds wonderful and you want to capture it 'as is,' that determines the rest of the signal chain. If the source

isn't what you hear in your head, what can you do to the source or put in the signal chain to get it that way?"

What pre-production do you do before a session?

"Generally, get together with the client, listen to music and demos, and talk about the process. They might play me other music they like, so we're thinking in the same direction. We'll talk about arrangements; instrumentation; who plays what; when, where, and how we're going to record; what will be overdubbed; and are we using their instruments or the studio's instruments?"

"I prepare for the session by having everything set up, in place, tested and ready to go – so that they can walk in and make music as quickly as possible."

What are your favorite mics and preamps for some different applications?

"I like the concept of using a consistent set of preamps – it's more cohesive, and gives the record a tone. I love the sound of my vintage API."

"Electric Guitar: Beyer M88, fairly close, where the dome meets the cone."

"Vocals: It varies, but I'm never unhappy with a Neumann U67."

"Drum overheads: A small capsule stereo tube mic like the Neumann SM2, or a pair of U67s, or a pair of ribbon mics, depending on the sound I'm looking for."

"An RCA 77-DX is a perfect match for saxophone."

What's your approach to recording in a less than ideal room?

"Generally, to move the mic closer to the instrument so we're not hearing the room as much. Don't set up in the middle of the room where most resonance problems are, and don't set up in a corner unless the problem is not enough bass."

"I would use more directional microphones, but also be careful to not mic too closely."

What's your attitude to using less than top-of-the-line equipment?

"I make it work. Something that is less than hi-fi can still be interesting. It depends upon the style of music and the sound we're going for. I would experiment with placement and settings to minimize what might not be great about it and maximize what's interesting about it."

What do you want people to think when they hear your work?

"I want them to think that the music speaks to them. That the recording is the frame around the picture, and the music is the picture. I would hope that the recording would serve the music to the point that it doesn't intrude."

20.9 Catherine Vericolli

ENGINEER, STUDIO OWNER: FIVETHIRTEEN RECORDING – PHOENIX, ARIZONA, USA

- Successful studio owner since 2005
- Board member: SoundGirls.org
- Audio Archivist: Useful Industries, Nashville, Tennessee, USA
- Conference presenter and article author

What type of recording work do you generally do?
"When I do recording work now, after 15 years of having a studio, I tend to do things that I'm really interested in, which these days is 'the simpler the better'."

"I love recording instrumental music. Obviously I've recorded a ton of vocals in my lifetime. I think I got burned out on vocal editing and tuning. As I got older I really started to move towards instrumental music."

Where do you generally work?
"I can work anywhere, but I prefer to work in my studio. The biggest reason why is if something is weird or wrong as far as troubleshooting is concerned, I'm the only one to blame – I put the studios together myself."

Figure 20.8 Catherine Vericolli

How long does your typical project last?

"It depends on how prepared the band is. It also depends on the musicianship of the band, the instrumentation, and how picky they are. The last record I did I tracked in a day and half, and mixed in another day. But that was a two-piece band recorded live."

"Some records I spent a year making. In the early days of the studio, we spent two years making some."

Anything related to mic techniques that unexpectedly pleasantly surprised you?

"I remember early on we were tracking a really talented musician, and we wanted some really gritty back-up vocals. We ran a microphone through a cool effects pedal into this old Magnatone amp, with a ton of reverb and a ton of delay on it – it was all just out of the amp, miked up. We decided to put some drums on a couple of tracks. We moved things out of the way, and that microphone was still up. It was the coolest accidental tone ever."

"Pleasantly surprising mic techniques are always accidents. Leave everything up and running – you never know what you're going to get."

How about any mic techniques that didn't work as well as you hoped?

"Most of them! We have our standards – stereo techniques we've learned. There's a tendency to really deviate from those standards because we want to be creative and push boundaries, but more often than not, letting that science fall by the wayside, you're not going to get great tones. I'm going to go out on a limb and say that most of the time, deviating from standard miking techniques I've had bad experiences."

What is the key to recording great sound?

"To do everything you can to promote the quality of the source. When tracking vocals you have to do everything possible to make the vocalist feel comfortable. What mic you're using isn't really as important as creating an environment so they can deliver the best performance they can, or didn't know they could."

"Is the snare drum tuned to absolutely the best ability of the drummer or engineer? Are the musicians comfortable, and are you communicating properly? It's more about the people, including yourself as an engineer, than any recording equipment or piece of gear that you're touching."

What's the most important part of the recording signal chain?

"The source. Second to the source, the placement of the microphone – listening, finding the right spot. I still go in the room and put my ear up to things to find a place that sounds best to me, and that's where I'm starting, where a microphone's going. I try to translate what that musician is going for or is wanting. It's a combination of the source, the room, and where you put the microphone."

What pre-production do you do before a session?

"If the band or artist can afford it, I do as much pre-production as possible. Working in a market like Phoenix, budgets can be pretty small and pre-production can be a new concept for bands if they're inexperienced. A lot of the time, the only pre-production we're getting is going and seeing a show. We've gotten really good at making decisions based on that. If we have a band with the time and budget, we'll spend a couple of weeks talking about arrangements. It's contingent on the band. If I had my way we would have at least two days, if not a week."

What are some of your favorite mics for some different applications?

"I'm a huge fan of ribbon microphones. If I could have a studio full of Coles 4038s, I would – they are my favorite drum overhead and percussion microphone. Percussion is my favorite thing, and can really save a song, or push a chorus – so a 4038 for tambourine."

"How can I get the best quality recording with the least amount of mics possible? I love being able to put up one fader and go 'oh yeah'! If there's 22 microphones on the drum kit, I can't focus on the artist."

What's your approach to recording in a less than ideal room?

"I think my live rooms I've been recording in forever are less than ideal. A good approach is to stick to fundaments – place the mics where you know they should go, or where you've had experience with success, as a starting point. From there either take away or add, move and reassess. Use whatever you can get away with using – if you're in a living room maybe the couch can serve as a baffle, and there are lots of things you can use to get some diffusion or absorption."

What's your attitude to using less than top-of-the-line equipment?

"My attitude has really changed as I've gotten older and I've worked more. I used to be really not interested in working with equipment that was less than ideal, or that wasn't expensive."

"My attitude now is 'use whatever it is that works for you.' If you're lucky enough to have access to beautiful analog consoles and expensive mics, awesome! But if you're making recordings you're happy with and that are successful, in Ableton on a laptop with USB microphones, then I'm all about it!"

What do you want people to think when they hear your work?

"I don't want people to think anything. I want people to experience it in whatever way they experience it. If they don't like it, that's perfectly fine, that's their experience. It's just as powerful as somebody who thinks it really great. I'd love to put out my work with a caveat about the recording situation, or the budget or timescale – but you can't do that. People will have their experience, and I hope it's a good one. But if it isn't, it's just as valid."

20.10 Paul "Willie Green" Womack

PRODUCER/ENGINEER, OWNER: THE GREENHOUSE RECORDING CO. – BROOKLYN, NEW YORK, USA

- Producer/Engineer: The Roots, Billy Woods, Wiz Khalifa, Open Mike Eagle, and more…
- Audio Engineering Society member

What type of recording work do you generally do?
"Most often I'm recording vocals, usually rap and R&B. But I've got a nice large booth where I can record drums, small horn and string groups, and more."

Where do you generally work?
"I primarily work out of my studio, I like to have my own space. The more intimately I know the studio, the tech component disappears and I can just focus on the music. Real estate is at a premium in New York City of course, so my space is limited. If I need to do a full band or string ensemble, I'll have to go to a larger tracking studio. There's a good number of Neve rooms around Brooklyn with nice mic lockers, so there are options."

Figure 20.9 Paul "Willie Green" Womack

How long does a typical session last?
"I don't usually like to go longer than five to six hours for a vocal session. If things are going well I'll go over, but after that I tend to get diminishing results. In a rap session, that's often multiple songs – provided they were written beforehand. For an R&B or pop session, that time is spent on one song. First, I'm focusing on leads and stacks, and backgrounds and harmonies on a good day. It's all dictated by whatever speed the artist is comfortable with. The key is to be capable of being as fast as possible, but able to slow down to match the artist's vibe."

Anything mic technique related that unexpectedly pleasantly surprised you?
"I've always found it fascinating looking at pictures of sessions from the 1950s and 60s and you'd see the great balladeers like Nat King Cole or Sinatra singing two to three feet back from the microphone. It intrigues me because while you can add room sound of course with reverb, what you can't add with a plug-in is that space between the vocalist and the mic. You've got to have the right room for this, and the right singer of course, but with the right circumstances and vocal power, that distant vocal can be really magical."

How about any mic techniques that didn't work as well as you hoped?
"Honestly I never have much luck with ribbon mics on vocals. With the exception of one singer who it sounded magical on, ribbons fail for me every time. I've tried all different kinds – vintage, modern – it just never seems to work out."

What is the key to recording great sound?
"Making the artist comfortable is first and foremost. Standing in front of a mic in a studio is a very vulnerable place, and if the artist isn't safe and comfortable they'll never give a good performance. The lighting should be right, the headphones should sound great, everything should be ready when they walk in the door."

What's the most important part of the recording signal chain?
"Mic selection is really important, picking the mic with the right tonality means less EQ work later. Of course each component in your signal path affects the result, but everything is based off of the mic choice."

What pre-production do you do before a session?
"I start with my template, which has a number of Lead, BG and Hook tracks set up and routed to busses, as well as some go-to reverb and delays so I can get a vibe on a vocal quickly. Most tracking sessions I do are generally cut to a 2-track of a beat (with multitracks to come later for mixing) so I load the track into my DAW and get the tempo locked, and drop markers to denote the song form. This process is important because it makes navigating the session, editing and flying vocals much easier on a grid, and I'm also getting to know the song so I can more easily communicate with the artist."

What are some of your favorite mics and preamps for some different applications?
"My go to for rap vocals has been the Shure SM7B with a Cloudlifter into a DIYRE CP5, usually with their Cinemag Transformer option. The Cloudlifter gives the mic a nice openness, as well as reducing any noise floor issues. The preamp is super clean, but gives me the option to push the transformer as hard as I want or not at all."

What's your approach to recording in a less than ideal room?
"It's a matter of being honest about the space and what you want to accomplish, then trying to minimize the effects as much as possible. You may love distant miking techniques, but that might have to wait for another record. For vocals I find the isolation shields work well, although the various designs do have differing sonic effects, so that is something to consider."

What's your attitude to using less than top-of-the-line equipment?
"We've come a long way with affordable quality equipment, and it's a great way to get started and make some good sounding records. The key is to do your research and find the best in a crowded sea of products. Expensive gear costs more for good reason however, and as you gain experience, your now nuanced ears will appreciate that quality more. Spending a bit more to buy something quality once, rather than a number of cheap stop-gaps usually makes more sense in the long run."

What do you want people to think when they hear your work?
"The best compliment I ever got was that my music felt physical, in a way that evoked a similar reaction. Anyone can put up a mic and record a sound, but I want to capture the intangibles of a vocal, that's where the emotion is. I want people to connect with the artist's emotion when they listen – that emotion is what turns 'sound' into 'art'."

Index