

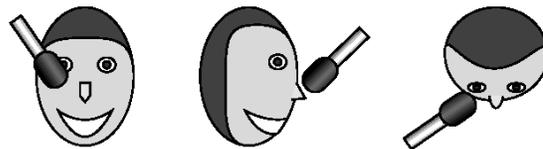
from Chapter 8: Voice-Over and ADR

✂✂✂ You'll need a flexible way to position the mic. The most common is a floorstand with a 2-1/2' boom; these usually sell for about \$50 at broadcast suppliers and music stores. You don't need a shockmount or heavy windscreen; the mounting clamp and foam windscreen supplied with the mic should be sufficient. While some engineers prefer pop screens—4" diameter frames filled with nylon mesh, that clamp in front of the mic—I've never found them useful for v-o work. They don't stop pops any better than using a good mic position, and are one more large item in the talent's face to break their concentration. (They can be helpful for singers who prefer to belt a performance directly into the mic.)

Ask the talent to start reading, so you can see how they hold their head when working with a script. Then try one of the positions below. Listen through speakers and adjust as necessary. Closer will be more intimate but can pick up more vocal problems. As you move off-axis from the mouth, sibilance and pops decrease but so can warmth. The best compromise is often to be off-axis but close, with the working end of the mic about 9" from the talent's mouth. Don't get so close that quality changes as announcers turn their head slightly to read the bottom of a page.

For a small diameter cardioid or hypercardioid, start with the mic angled down toward the mouth from just below eye level. The mic comes from the side, so doesn't get in the way of reading the script. This position seems very good for getting an intimate sound without pops or sibilance. Figure 6 shows this from three angles.

This position also works with large diaphragm cardioids. Or as an alternative, try at mouth level and slightly toward the side. The



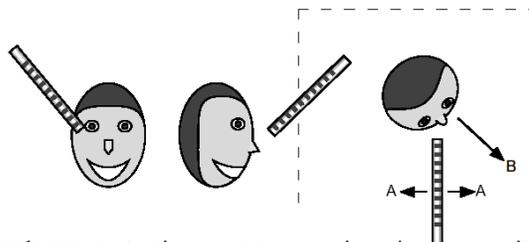
08fig06 A good starting position for a small mic.

front of the mic, of course, should be angled toward the mouth. Figure 7 shows this placement for a large cardioid that accepts sound from the side. If the mic is a type that accepts sound from the top, rotate it 90°.



08fig07 An alternate position for a large side-address mic. The arrow indicates the front of the mic, pointing to the mouth.

You can use a short shotgun in the same position as a small cardioid, though a couple of inches farther away (figure 8). Angle things so the nearest reflective surfaces are about 90° from the mic's axis (direction A), though as far from the mic as practical. For best results, the talent should be facing the largest dimension of the room (direction B). The dashed lines indicate walls, at least 3' away, and shows how the best position might be diagonal.



08fig08 A similar position works when recording v-o with a short shotgun, and lets you avoid acoustic problems in the room.

These drawings are intended as starting positions. Take time to learn how your mic sounds best with the announcers you use. As you're experimenting, there are a couple of positions to avoid:

- Don't put a large diaphragm mic directly in front of the talent's mouth (figure 9A). Some announcers assume this is the only way to work—it's popular in radio—but it's sure to pick up pops and clicks. ❌ ❌ ❌

❌ ❌ ❌ While ADR is an accepted practice, it's frowned upon by feature directors and sound supervisors because the acting is seldom as good as during the original performance. Producers don't like ADR either, because it's expensive. Many actors hate the process. They have to read the scene line by line without other performers to work against. It's difficult to recall the original emotions or motivations for a reading, and to keep an appropriate projection level. It can also be intimidating to have to perform while listening to cues in a headphone and watching yourself on a screen.

Gotcha: ADR Avoidance

If looping is difficult in Hollywood, where they've got experienced technicians and the right equipment, it's even harder when you have to do it ad-hoc. The best strategy is to do as little as possible. Some of this has to take place long before post.

- Strive for the best production sound. Renting better equipment or hiring a good boom operator is usually cheaper than re-recording later. Ten minutes of extra setup time at the shoot can save a day of ADR. If you must replace dialog, the session will go smoother if you have decent guide tracks. If the original was recorded indoors with a camera mic, ADR can be long and painful because the guide track will be blurred by echoes.
- If you think ADR will be necessary because of a noisy setting, block the scene so lip movement isn't too obvious. Have extras mime

their conversations. If you can, rewrite the scene with a lot of back-and-forth dialog: short lines are easier to replace.

- While you're shooting, pay attention to intermittent noises like car horns or airplane flyovers. It's cheaper to retake a scene immediately than try to record it later. Even if the second performance isn't as good, you may be able to lift just enough dialog to save the first one. Or forget about picture. Have the actors record problem lines wild, right after the bought take. They'll still be in the right mood, and the acoustics will match.

ADR Equipment

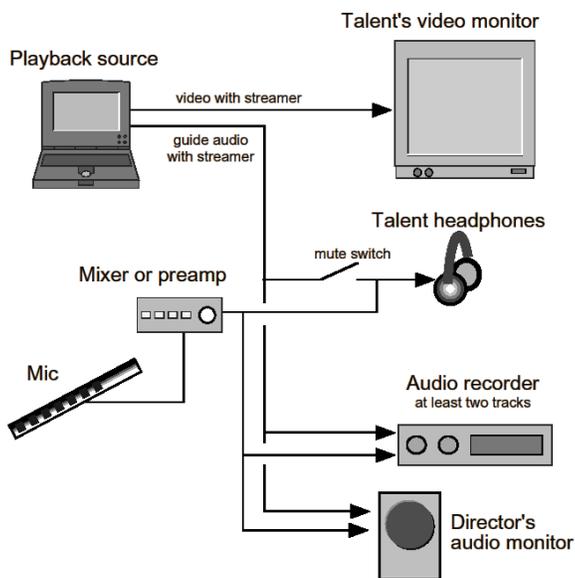
You'll need an a/v playback device: an NLE is ideal, since it can cue quickly. If you can't isolate its fan noise, use a camera or a VCR instead and prepare a tape with each line repeated multiple times. You'll also need a separate audio recorder with at least two tracks, so you can record the original audio on one channel as a sync reference. Timecode helps but isn't necessary, so a consumer MiniDisc may be adequate for this. You'll need the same kind of mics that were used at the shoot. You'll need headphones for the actor; ideally isolating ones so that cue tracks aren't picked up by the mic.

Hook everything up as in figure 12 (next page). Studios that specialize in ADR have more complicated setups and specialized software to drive it, but this will do the job for casual sessions.

You'll also need a quiet room, with little or no reverberation. Even though the production dialog has natural room reverb, every room sounds different and you don't want replaced lines to stand out. If the new audio is echo-free, you'll be able to simulate reverb that matches the shoot when you mix. If the new audio has its own echo, you probably won't.

The room should be larger than a typical announce booth, unless the production track used lavs and was low-key. New dialog has to be recorded at the same mic distance and with the same energy as the original, and most

booths aren't big enough to support this. So plan on a lot of sound blankets or other absorbers. If you have to simulate exteriors, the room must be absolutely dead. Outdoor echoes are nothing



like the ones you get indoors.

08fig12 A setup for casual ADR. Streamers are described in the next section.

The mic should be mounted on a boom, at about the same distance and position as it would be used at the shoot. Some ADR recordists prefer two boom mics, one slightly closer and one slightly farther. They record each to a separate track, and choose the most appropriate one in post. If the shot originally used a lav, of course, use a lav for the ADR.

ADR Technique

The first step is to do your homework. Watch the edited project, with a script handy, and mark which lines need replacing. Jot the line's timecode on the script, along with why you're replacing it and anything else that might be useful. Knowing whether a shot was interior or exterior and how it was miked will save time at the session.

Then check for alternate takes. You may be able to lift a few words from a cleaner track, and slip them under the existing video.

If you must loop, turn each line into a separate a/v clip. If characters overlap, make two clips. Record each separately and mix them later. Besides, it's better to book each character for a separate session: ADR skills vary, and there's no sense making one actor wait while another blows takes.

Video-dominant ADR

There are two principal ways to record ADR. The traditional method is to make picture dominant. Each clip should start a few words before the desired line, and end a few words after.

You'll need to add visual streamers to each clip, stopping on the frame before the target line. Originally, this was done by scribing a diagonal line across a few dozen frames of the film itself; when projected, it would appear as a white line moving across the screen from left to right. These days, streamers are usually generated in video by the ADR equipment, which may be part of a specialized audio workstation. If you're trying ADR on a desktop NLE system, you'll need to add the streamers as a super. Also add an audio streamer to the track: three short beeps about 20 frames apart, forming a rhythm that ends as the new line starts.

The actor watches a playback of the clip on a large monitor, with production audio in headphones, to get comfortable both with the delivery and with the way the streamer leads into the line. Then the guide track is muted. The actor records a new version while watching the screen, but without hearing the original. (The director may want to listen to both tracks, with production audio at a slightly lower volume, as a guide for sync.) Multiple takes will probably be needed, and it's important that they form a predictable rhythm. Start-and-stop ADR, where it takes a random time to cue each pass, puts extra stress on the actor. Eventually, the new performance will start to match the guide track. You don't have to achieve perfect lipsync: small errors can be fixed by editing.

The advantage of this method is that actors can feel freer to experiment with delivery, since the original version isn't being played in their headphones during the actual take.

Audio-dominant ADR

Traditional ADR can be intimidating to an actor. A modern alternative is to make sound dominant. Set the playback device to play the clip and streamer in a continuous loop. Let the actor hear the rhythmic beeps and line over and

over, and then record them speaking along with the original until the new and old voices match. The advantage to this method is that it can be faster with performers who aren't used to looping, often produces a tighter match, and since it isn't necessary to concentrate on picture — you don't even need the video monitors — the actor can be looking at a script instead.



from Chapter 16: **Noise Reduction**

Other noise reduction facts

There are plenty of myths and misconceptions about the process of noise reduction. If you understand what's really going on, you'll usually get better results.

Noise reduction without a noise reducer

Noise reduction software often attempts to take care of things automatically. This can do more damage, in many cases, than using other techniques manually. If you've been following the examples on this book's CD, you've already heard some fairly effective noise reduction that relies on other kinds of processors.

- Get rid of whistles using the equalization techniques in chapter 12, and demonstrated on track 51.
- Tracks 54 and 55 show how downward expansion can improve modestly noisy interview tracks (chapter 13).
- Track 62 virtually eliminates dimmer buzz using a comb filter (chapter 14).

And of course, the ultimate noise reduction for extremely bad dialog recordings is ADR (chapter 8).

Go back and re-read the appropriate sections. These tools are often the first defense, and should be considered before whipping out general-purpose noise reduction software.

Editing can also be used for noise reduction: replace the noisy part with something else. That's one of the principal uses for room tone,

though noise-reduction editing can be as subtle as replacing a few waves with clones of adjacent ones. It can even involve editing individual samples, drawing over clicks or other transient sounds with the pencil tool in an audio editing program. A few high-priced DSP-based noise reducers can do these things automatically, though they need precise tuning to sound good.

Nulling noise?

Folks who know a little acoustic theory sometimes ask, "why can't I create or capture a sample of the noise, invert the polarity, and use that to cancel the noise out?" The idea behind this is basically correct: the comb filter, for example, works by delaying symmetrical noise exactly one-half cycle. This lines up the negative side of the wave against the positive, nulling out the noise. But that's a special case, relying on a characteristic of the noise itself and using a precise delay instead of a capture or copy and paste.

For the sample-and-invert technique to work, the noise would have to be of absolutely consistent waveform and volume, and the sample would have to be pasted back exactly in sync with the continuing noise. We don't have any way to do that with today's technology. If there is the tiniest error in timing or if the noise has changed since the sample, this operation would increase the noise instead of removing it.

There are some noise reduction processes that rely on a sample of the noise, but they're not using an inverted sample. Instead, they take

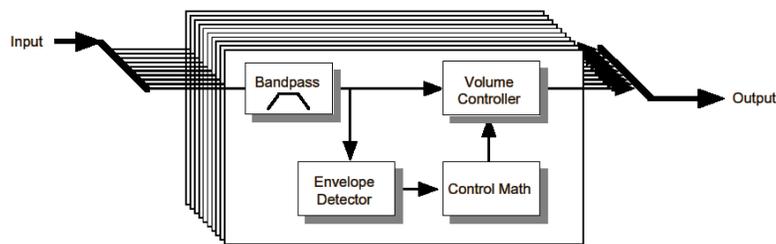
a spectral fingerprint of it and use that to control other techniques. Nulling isn't involved.

the frequency of the band as well, to avoid distortion.

~~~~~  
 ✂ ✂ ✂ **Multi-channel masking**

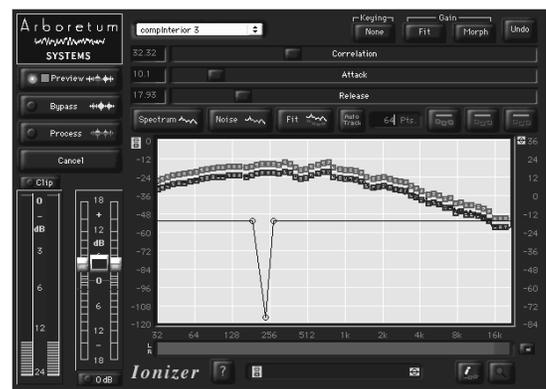
Imagine a bank of hundreds of downward expanders, each preceded by a narrow bandpass filter. Each one would look at a small slice of dialog frequencies, sometimes smaller than a single formant. When there's voice energy in a particular narrow band, the expander would pass it through. Any noise in that band would be masked by the voice. When things are silent in that band, its overall level would be reduced, lowering the noise, even though the voice may be speaking (and masking noise) in other bands. Figure 9 shows a simplified version; actual multi-band noise reducers often have as many as 512 bands.

This could present the user with a nightmare of controls. Or a manufacturer could preset most of them for "average" noise—whatever the programmer thinks that is—leaving the system less effective on other tracks. Arboretum's lonizer plug-in for Mac and Windows (figure 10) seems to make the best compromise; Sonic Foundry's Windows-only Noise Reduction software is a close second. In both, attack and release time, and the amount of linkage between bands (correlation or overlap) is adjusted globally, and each band's threshold and floor can be automatically fitted to a sample of the noise or adjusted manually. lonizer also lets you



16fig09 A multi-band noise reducer, simplified.

While this approach can yield very effective, almost amazing noise reduction, it can also cause problems. Pure tones and narrow band signals can cause a flanging effect in the noise, as single expanders open while their neighbors remain closed. Good noise reducers include a way to partially link adjacent bands, so each expander responds mostly to its own energy but also considers the state of the expanders around it. But too much linkage will destroy the noise-reduction effect. Time constants are also important. If the release time is too fast you'll hear a chirping artifact as bands respond to transient sounds. If it's too long, noise can remain after the masking sound is gone. Both attack and release have to consider



adjust the knee for each band. Instructions for using these programs are in the cookbook section of this chapter.

16fig10 Arboretum's lonizer plug-in is one of the few high-powered noise reducers that combines a full set of user controls with mostly automatic setup .

**Get it: A great freebie!**

Audacity ([audacity.sourceforge.net](http://audacity.sourceforge.net)) is an open-source, freeware audio editor for Windows, Macintosh, and Linux. Among its features is a fairly good noise reduction processor.

The program has some limitations, including the inability to preview effects. On the other hand, being open-sourced, new features are sure to appear. ✂ ✂ ✂

from Chapter 19: **After the Mix**

## ✂ ✂ ✂ Streaming

Web formats such as RealMedia, Windows Media, and QuickTime are designed to stream, so they start playing before all the data is downloaded. Streams can be played easily in most browsers or free add-ons, and the fast response is gratifying to users. But the biggest advantage may be that most users don't know how to save streams as files; this appeals to the security-conscious.

There are plenty of self-contained streaming formats. Some, such as Macromedia Shockwave, rely on garden-variety mp3 with a slightly different header. But you can stream mp3 files directly with a little extra web programming. Or you can convert an mp3 file to streaming QuickTime easily, with no quality loss.

### Streaming by mime type

Most web browsers accept a mime type of audio/mpegurl. This is a text file with the extension .m3u, containing a single line pointing to the URL of an mp3 file. Link to the m3u with a standard <A HREF> tag. When the browser loads the m3u file, it passes the URL to an application that can handle mp3 streams, typically RealPlayer. The application loads a buffer and the file starts playing.

A typical m3u might look like this:

```
Filename: intro.m3u
Content:
http://64.33.79.241/aes/audio/intro.mp3
```

You can also specify the mp3's location with a domain name, but IP addresses generally load faster. Some older servers may not support this mime type and you'll have to link directly to the mp3 file. Or go to plan B:

### Streaming via QuickTime

QuickTime doesn't save mp3 audio—Apple doesn't have the Fraunhofer encoding license—but it can play the files. You can trick the

player into streaming raw mp3 by putting it in a QuickTime shell. You'll need Apple's QuickTime Pro, a \$30 upgrade to their free Mac and Windows player.

- 1) If you're making an audio/video movie, save the track as 16 bit linear audio. If it's to be an audio-only file, find the source file and skip to step 3.
- 2) Export the a/v movie's track as an AIFF or WAVE, 16 bit linear file. Leave the original movie file open.
- 3) Open the audio file and convert it to mp3, using a good encoder.
- 4) Choose QuickTime's File: New Player command. Then open the mp3 file in that player.
- 5) Choose the Edit: Select All, and then Copy. This puts the mp3 audio data on the clipboard. If you're making an a/v movie, skip to step 7.
- 6) If you're making an audio-only movie, select File: New Player again. This creates an empty movie.
- 7) Click on the original a/v movie or the new player you created in step 6. Choose Edit: Select All, then Edit: Add. The mp3 data is now in the movie. If you're doing audio-only, skip to step 9.
- 8) Choose Edit: Delete Tracks and, in the window that appears, delete Audio Track 1. This gets rid of the original 16-bit linear data.
- 9) Chose File: Export. In the window that appears, select Export: to Quick Time Movie. Click Options, and in that window turn on Prepare for Internet Streaming, Quick Start. Click Okay, name your movie, and click Save.

The QuickTime with its mp3 track can now be put on a web page with a standard <A HREF> command, or you can <EMBED> the movie. ✂ ✂ ✂

Want to learn more?  
<http://www.dplay.com/book/app>