

Transferring Vinyl LPs to Digital: One Approach

by [Matt Neuburg](#)

For some years now, I've been using my computer to transfer analog recordings to digital. I started with all my cassette tapes, as they were physically deteriorating while they sat on the shelf. Having completed that task - I no longer have *any* cassette tapes - I've been transferring a number of vinyl LPs.

It's important to be clear on the reasons for transferring vinyl assets to digital, since they are *not* deteriorating. It isn't that I think I can improve the sound significantly. It's true that I can remove some of the annoying clicks and pops; but at the same time I'm one of those wacky audiophiles who thinks that vinyl sounds better than any digital derivation. (Compressed music, such as MP3 or AAC, is compressed lossily, and I can hear the loss; and even an uncompressed digital format like AIFF is sampled, and reconstructed during playback, with consequent loss, artifacts, and errors.)

No, I'm transferring vinyl for the clear and sane reason that I rarely *play* the vinyl, since setting up and attending to the record player is an elaborate business, and each playing wears down the record slightly. Digital music, on the other hand, is readily available and playable at any time, with no adverse effects to the medium. In short, digital is a lot more convenient!

While doing all this transferring, I've often thought of writing a long technical article, perhaps even something of book length, explaining how to do it. But writing such a piece seems like an insane amount of work. Besides, it would be ephemeral, since the technology changes constantly; and it would please no one, since no two people use exactly the same procedure. So I've decided instead to sketch (rather than even describing in detail) the procedure that I'm currently using.

The goal of this procedure, as it has developed over the years, involves keeping costs low and speeds high. I used to spend many days nursing a single recording, but now my priorities are more balanced, and I get results that are good enough in just the hour it takes to record two sides of vinyl and less than a second hour of processing. And although there exists extremely fine audio restoration software costing thousands of dollars, my method is sufficiently effective and suits my budget-conscious temperament.

So I present my procedure, for what it's worth, in case anyone is interested. As I've already said, this presentation deliberately avoids too much detail, since this isn't a

book, and makes no pretense to dictate to anyone else, let alone to enter the many realms of technical detail and near-religious controversy. I'm just going to tell you what I am currently doing, along with some of my thoughts about each step.

Externalities -- By "externalities" I mean everything used to play a record (or cassette tape or whatever it is) before digitization. I can't say anything very useful about this part of the process. Your externalities are, in effect, the very same equipment you would use to listen to the material in the first place. Clearly, the better it is, the better the sound will be. If you've been playing records on a super-expensive audiophile turntable with an expensive cartridge and a high-end pre-amp, so much the better; you can use it. If what you've been using is, shall we say, less high-end, you can use that too. My point is that the expense here is essentially zero in any case, since if you can listen to records now, you probably have the equipment you need to digitize them.

In case you don't have that equipment, or you're wondering whether what you already have is sufficient, I'll just quickly sketch the train of musical production. Assuming we're starting with a vinyl LP, there's a turntable, including a cartridge and needle. Then there's a pre-amplifier. This is an important component, for two reasons. First, the signal from a turntable is too weak to feed directly into a computer. Second, the direct signal from an LP is incorrectly equalized; because of the physical (analog) nature of vinyl LPs and turntables, records are made with high frequencies boosted (to help drown out the noise of physical imperfections in the medium) and low frequencies attenuated (to reduce the physical distance the needle must vibrate during playback). A dedicated pre-amp reverses this.

What I use is an old-fashioned standard ("integrated") amplifier that contains a "phono stage." In other words, my pre-amp is a special circuit inside my regular amplifier, that comes into play when I switch the amplifier to "phono." If and when it goes bad, I might have to purchase a dedicated pre-amp costing several hundreds of dollars, but so far that has not been necessary.

In the accelerated transfer procedure I'm describing here, I don't worry about the inexactness of the match between an individual record's built-in equalization and the reverse equalization of the phono stage. (It is likely to be inexact, because the so-called RIAA curve [1] was never a universal standard.) I just take what I get. I can always equalize further, if I like, while listening to the digital transfer.

Skipping merrily past the many questions, complications, and invitations to spend vast sums raised by the externalities - issues regarding the elimination of hum and other noise on the electrical line could, alone, occupy an entire article - the chief question you have to solve is how to get the sound out of your analog equipment and conduct it in the general direction of your computer. The idea is to get electrical signals representing the sound coming down some cables. My stereo amplifier has two holes in the back marked OUT (on other amps they might say TAPE OUT or similar). So I have cables with RCA jacks running out of those holes. If you're not

using a standard amplifier, a similar cable might be coming directly out of your pre-amp.

The question now is what to do with the other ends of those cables to get the signal into the computer.

Analog-Digital Conversion -- If you're going to spend a lot of money, this - the moment of analog-digital (A/D) conversion - is the point in the process at which to do it. You could, in theory, stick an RCA-to-1/8 adapter on the end of your cables and plug it directly into the computer's line-in / microphone port, if it has one. But in that case you're asking the computer to use its internal A/D conversion, which is the worst possible quality.

Instead, you want some sort of device that accepts analog signals and transforms them into digital and runs into your computer's USB port - or, in higher-end cases, into a FireWire port. It should be good quality and it should support, at a minimum, a 44.1 kHz sample rate and 24-bit bit rate. (You *really* want to record in 24-bit if possible, because it provides far better quality than 16-bit and lets you worry less about excessive headroom. More about that later.) Such a device is often termed a "sound card" even though it doesn't look like a card; you might also see terms like "audio interface" or DAC (for "digital-analog converter").

A visit to a site like [Sweetwater](#) [2] will give you a sense of the sort of equipment that's available. Bear in mind that you're going to get what you pay for. Your goal is to refrain from [cheaping out](#) [3] completely, while at the same time not spending more than is necessary for your purposes. If you are not also planning on starting up a home sound recording studio, mixing live vocals and instruments and that sort of thing, something like this [Tascam](#) [4] device would probably make a good starting place. Obviously you can spend considerably more if you want to. Whatever you're thinking of getting, make very certain there are drivers for your system! Also, if your device draws its power from the USB port, it might be a good idea to invest in an inexpensive powered USB hub, since some computers (especially portables) might not provide enough power directly.

Recording -- Once you've run your analog sound into an A/D conversion device and you've set that device as the sound input source on your Mac, you're ready to record. There are three issues here: what format to use; what software to use; and setting your recording levels. The three problems are related and not easy to disentangle, so I'll just cut to the chase and tell you that I use [Audio Hijack Pro](#) [5] (\$32; see "[Why Go Pro \(Audio Hijack Pro, That Is\)](#) [6]," 24 January 2005), for two main reasons: it is totally reliable, never giving me any latency or skipping problems; and, it permits me to use the free [Inspector](#) plug-in, which shows signal strength in real time and tracks headroom so you know if the gain is too loud. However, I'd be failing in my duty if I didn't point out that you could probably do just as well with the freeware [Audacity](#) [7], or even GarageBand, which may well have come with your computer (and there is a

splendid Take Control book [8] devoted to the topic of recording with GarageBand).

I recommend that you record as AIFF in 24-bit with at least a 44.1 kHz sampling rate; these settings must be made both for the A/D conversion device and for the recording software. The more information, the better; so if your device can sample at 96 kHz, you may be tempted to do so. But, in choosing a format, keep in view your final, archival format, since you don't want to resample unevenly later. (So, if you record at 96 kHz, you can later halve it to 48 kHz, but you should *not* resample to 44.1 kHz; if that was your final goal, you should have recorded at 88.2 kHz.)

Your A/D conversion device should have volume settings that allow you to set a good level. Your goal is to minimize headroom without clipping.

What I mean is this. Let's assume that your recording consists of both soft and loud sounds. Then the loudest sounds should be close to the 0.0 saturation point without actually reaching or going above it, since a wave peak that exceeds the saturation point is useless because the top is sliced off (clipping). But you don't want to be overly cautious in the other direction, either, since if you never get close to the saturation point - that is, you have too much headroom between the loudest sounds and the saturation point - you are wasting your capabilities, and not getting all the digital information you're entitled to, which can degrade the ultimate quality of the recording. (If your loudest sounds are recorded too quietly, imagine how poorly your quietest sounds are being captured!)

This is one reason why you really, really should record in 24-bit; those extra bits per sample mean, among other things, that having a little extra headroom is nowhere near as disastrous as it would be when you record in 16-bit, and thus you don't have to be so nervous about it. As a result, instead of recording many takes in an attempt to minimize headroom without clipping, a single recording is usually sufficient.

It's important to understand that not *all* clipping is bad. If the recording clicks because there is a loud crack or pop on the record, that's completely unimportant. The signal you are trying to capture here is the music. The click is not part of that signal; it's noise, and you're going to remove it later anyway. It's a pity it's there, but it's irrelevant to the business of setting levels. Moreover, if just one or two (or even, say, half a dozen) wave peaks of the music happen to exceed saturation point, that's not important either; you can mend those individual wave peaks later on, and indeed, the fact that this is happening is a good sign, because it means that the rest of your signal is minimizing headroom.

Check Peaks -- Once the recording is made, I open it with Amadeus Pro [9] (\$40). This is far and away my waveform editor of choice (see "Amadeus Pro: The Classic Continues [10]," 12 February 2007), and has many analysis and transformation functions that can be useful in this process.

Now we look for clipping. In Amadeus, press Command-Option-P to find a clip. Then

move the playhead to the right of the clip and do it again, running through the entire file. As I just said, if a clip is due to an extraneous click on the record (very obvious from its shape), it's fine, and if a clip is an isolated case of one wave peak being too tall, that's fine too, especially if there are only a few such cases in the file. What's bad is a stretch of a dozen saturated peaks in a row. This means the gain was just too high and the entire recording must be made all over again at slightly lower amplification.

Next we *fix* any clipping we just discovered. Start again at the beginning of the file, and this time, each time you find a clip, select the region around it and choose Effects > Interpolate. Amadeus's wonderful click repair function will remove the click or clipped wave.

At this point I often select the whole file and press Command-Option-P again, which, when there's a selection, means we locate the tallest peak in the selection. If this peak is a click, select it and remove it using Interpolate and start again, until you find the tallest actual wave peak. Now we can judge our headroom. If the tallest wave peak is insufficiently tall (too much headroom), and if this is supposed to be a loud sound, this might be a reason to make the recording again at slightly higher amplification. But if the headroom is pretty good (about -3 db, let's say), I don't get my knickers in a twist. We've recorded in 24-bit, so there's plenty of information present if we want to increase the amplitude artificially later on. On the other hand, if the headroom is significantly more than that, our softest signals may be so diluted that re-recording at higher amplification might be a good idea.

Remove Clicks and Pops -- If the recording is in pretty good shape, so that it has relatively few clicks and pops, I typically use a two-stage process to remove them: I remove very gross clicks semi-manually with Amadeus, and then I use ClickRepair to finish the job automatically.

So, first, I remove the grossest imperfections with Amadeus. Open the Repair Centre window, start at the beginning of the document, set the sensitivity quite low (less than 50 percent), and press N to find the first click. If it is obviously a real click and not a miscalculation - and, at this low sensitivity, it probably will be a real click - press R to repair it. You can repeatedly press N and R to run through the whole document, or if you're fairly confident, press Shift-R ("repair and find next") repeatedly. Unfortunately, Amadeus Pro has no command for continuing through the file in this way automatically (the earlier version of Amadeus did).

After I get to the end of the document in Amadeus, I start at the beginning and do the process again. The reason is that if two clicks are close to one another, one of them will have been missed during the first pass, due to the nature of Amadeus's click-detection algorithm.

We have now removed the grossest clicks and pops, thus paving the way for ClickRepair. So now I save the file and close it, and open it with ClickRepair [11] (\$40). This effective and inexpensive program will suppress the smaller clicks and

pops. I'm so confident in ClickRepair that I just set it to DeClick about 22 (no higher, or the resulting sound starts to degrade), no DeCrackle, completely Automatic, Pitch Protection on, Wavelet algorithm, and just process the entire file (sound output set to Off).

In the case of a recording that's in quite poor shape, I skip using Amadeus (because even at low sensitivity it finds just too many clicks) and use ClickRepair alone. In this case I might use a setting more like DeClick 22, DeCrackle 10, Pitch Protection off, Wavelet x2 algorithm. This might cause a little additional degradation of the sound, but the extent of the rescue from the scratchy and crackly overlay of the original is simply astounding.

Reduce Surface Noise and Hiss -- The use of ClickRepair in the previous step generated a new file, whose name ends by default with "cr" to distinguish it from the original. We are now going to process this new file with DeNoise [12] (\$40), another great program by the author of ClickRepair.

First, however, we must help DeNoise determine the level and nature of the underlying noise. The idea is to select a short stretch of pure surface noise and tell DeNoise to Sample Noise. If your LP recording starts with a blank area between the time the needle hits the record and the time the music starts, you might be able to find and use this area within DeNoise. But I find it easiest to open the file with Amadeus, select this area or some other area between tracks that's empty of any music, and save it off as a separate file. Then I open that file with DeNoise and use it to sample the noise. This causes DeNoise to set its parameters appropriately.

Now I open the complete ClickRepair-generated file with DeNoise and prepare to process it. I tend to take a fairly conservative approach to noise reduction, so I usually slide the Noise Floor slightly to the left of where DeNoise sets it, and I keep the "Limit reduction" slider somewhere around 8 db. It might be worth listening to the effect of different settings (and if you're going to do that, it might be worth investing in another piece of hardware, namely a good pair of reference headphones), but on the whole my goal here is to spend as little time as possible fussing, so I don't usually bother.

In praise of DeNoise, I want to say that it does nearly as good a job of reducing surface noise and hiss without degrading the music as high-end software costing literally thousands of dollars. The interface (and this is true for ClickRepair as well) takes some getting used to, and you will definitely have to read the manual, but once you've understood how to use the software, it processes a whole record side in just a few seconds.

It's important, however, to be technically and philosophically clear about what you're trying to accomplish during this step. Until you've practiced some noise reduction followed by listening under realistic conditions, there is a danger, especially if you're using sensitive reference headphones, that you'll be tempted to try to remove far too

much of the noise. Resist that temptation! The results will be sonically unnatural; and besides, the effort is completely unnecessary. Remember that you can't make a silk purse out of a sow's ear; even more important, you don't *need* to make a silk purse.

The goal is not to turn an LP into something magically resembling a CD. That goal is unreachable and unrealistic. You don't have the original tapes from which the LP was made, and you probably don't have the magic production ears of a professional. Your goal is simply to make listening to the digital form of the LP easy and pleasant. Your greatest ally here is not the software but the magic of the human ear (and brain). If the background noise is not obnoxious, your perceptions will effectively filter it out while you're listening. So you want to make the noise and hiss situation a *little* better, while not making the listening experience worse than listening to the actual LP.

By the way, DeNoise comes with a second application, DeNoiseLF (for "low frequency"), which can be used to reduce hum and subsonic rumble. If you're going to apply this, it's best to do so before using ClickRepair, as it improves ClickRepair's performance. (On the other hand, some unruly recordings may actually include a hum which isn't exposed until after you've used DeNoise, so you'll just have to remove it now.) Amadeus' wonderful Analyze > Spectrum command produces a frequency graph which can supplement your ears as a guide to whether your recording's silences include significant hum or rumble.

Save to Archival Form -- We are now ready to save the file into its archival form. The idea is to keep a pristine version of the file from which derivations, such as compressed formats, can be generated at any future time. That way, if I make a compressed version and later want to change the amount of compression, or if I want to experiment by adding a little equalization or other modifications, or if I just want a pristine listening experience, I always have the original uncompressed file as a starting point. We live in an age when a 2 TB hard disk costs less than \$160, so there is no good reason not to keep these archival files on hand.

My choice of archival format is FLAC [13]. That's because FLAC saves some space, but is non-lossy. (I do not favor Apple's "Apple Lossless" format, as it is proprietary and, in my experience, unreliable.) Of course I could save as AIFF, but FLAC can represent a space savings of nearly 50 percent, which is significant.

The only downside is that FLAC makes no provision for "markers," which are strings associated with specific points in time within the sound file. The reason we might need markers is in order to designate where individual tracks begin. There are two solutions to this problem: split the file into individual FLAC files, each file representing a track; or, keep the file as one large FLAC file, but accompany it with a CUE file, which is a text file in a special format describing the tracks. Amadeus makes either approach extremely easy.

So, I open the file generated by DeNoise (which, by default, has an "ns" suffix in its name) with Amadeus. I "top and tail" the file, deleting the initial and ending silences.

Now I place markers - one at the very start, and one at the start of each subsequent track. Finally, I either split the file at these markers while at the same time converting to FLAC (this is a single, easy move with Amadeus) or else I export the markers (which creates a CUE file) and then save the entire file as a single FLAC file.

Unfortunately, the CUE file generated by Amadeus is not legal, and many applications cannot read it properly. This is easily mended by hand. Open the CUE file with your favorite text editor and, after the first line, insert a second line specifying a PERFORMER parameter (the word PERFORMER followed by a space and some text in quotes). Also, in the third line, specifying the FILE parameter, change the value in quotes to match the exact name of the FLAC file.

Now place the CUE file and the FLAC file together in a single folder. That's it! The CUE file is a legal file and points correctly at the FLAC file. So, for example, you can listen to the archived tracks directly by opening the CUE file with Cog [14] or Play [15], and you can generate a compressed-format version of the tracks, suitable for use with iTunes or your iPod, by opening the CUE file with Max [16] or XLD [17] (all are freeware).

Conclusion -- This procedure has taken many words to describe, even in this sketchy form, but I assure you that it is calculated to be easy, reliable, and quick. There is no shortcut for initially recording the files, since this requires playing the LPs; but the time required to post-process the recordings is mostly the computer's time, not mine. Very little input is required from me, and I can easily accomplish other tasks while things progress automatically. A professional-quality transfer would require auditioning (listening to) the files during and after each step, which, along with careful manual intervention, can add *days* to the procedure. But although, in the past, I have done this, and might still do it in rare and exceptional cases, on the whole I no longer bother.

There are as many variations of my procedure as there are users. Nothing I've said here is written on tablets of jade. Legitimate divergences of taste or purpose are possible at every step. No matter what the variation, I've probably been there and done that. What I've described here is what I happen to be doing now. There is plenty of experimentation and sound reasoning behind my choices at every step, and the results suit me. They might also provide a starting place for someone puzzled or hesitating over how to accomplish a similar aim. If that's you, be reassured that it's possible to turn an hour of vinyl into very acceptable digital in just a couple of hours and without a capital outlay that will bust your budget.

[1]: http://en.wikipedia.org/wiki/RIAA_equalization

[2]: http://www.sweetwater.com/shop/computer-audio/audio_interfaces/

[3]: <http://www.sweetwater.com/store/detail/iMic2/>

[4]: <http://www.sweetwater.com/store/detail/US144mk2/>

[5]: <http://www.rogueamoeba.com/audiohijackpro/>

[6]: <http://db.tidbits.com/article/7961>

[7]: <http://audacity.sourceforge.net/download/mac>

- [8]: <http://www.takecontrolbooks.com/garageband-recording>
 - [9]: <http://www.hairersoft.com/AmadeusPro/AmadeusPro.html>
 - [10]: <http://db.tidbits.com/article/8857>
 - [11]: http://www.clickrepair.net/clicks_crackle.html
 - [12]: http://www.clickrepair.net/continuous_noise.html
 - [13]: <http://flac.sourceforge.net/>
 - [14]: <http://cogx.org/>
 - [15]: <http://sbooth.org/Play/>
 - [16]: <http://sbooth.org/Max/>
 - [17]: http://tmkk.hp.infoseek.co.jp/xld/index_e.html
-

Permanent article URL: <http://tidbits.com/article/11072>

This is a specially formatted version of this TidBITS article that removes colors, and optionally removes images; it's designed to minimize the use of ink and paper.

Unless otherwise noted, this article is copyright © 2010 TidBITS Publishing, Inc.. TidBITS is copyright © 2008 TidBITS Publishing Inc. Reuse governed by this Creative Commons License: <http://tidbits.com/terms/>.