Columbia University **EE E4896 Music Signal Processing**
Department of Electrical Engineering **Spring 2005**

Homework 3
– Due April 25, 2005 –

Problem 1 – Panning

Paning is the process of distributing a given (mono) signal to a set of two or more channels for the purpose of creating a "panoramic" sound field effect. Assume we only have two speakers, L and R. Given a signal *x*, and a pan parameter *p* that ranges from 0 to 1 (0 corresponding to *x* being sent to the L channel, and 1 corresponding to the case where *x* is sent only to the R channel) we can define the panning process as:

$$L = f(p)x \quad \text{and} \quad R = g(p)x$$

There are (at least) two approaches to panning. The first and simplest one is "constant amplitude," in which the result of panning is such that the sum of the L and R signals yields back the original signal. In other words: $L + R = x \Rightarrow f(p) + g(p) = 1$. An obvious choice is to set $f(p) = p$ and $g(p) = 1 - f(p) = 1 - p$. A counter-argument to this is that our ear is sensitive to power levels, and hence what we should make sure is that the total power of the paned signal is the same as the power of the original signal. This is called "constant power" panning. Express mathematically the requirement of constant power and propose simple $f(\cdot)$ and $g(\cdot)$ functions to accomplish it. Returning to constant level panning, plot the total power gain/attenuation of the panned signals as a function of *p*. What is the value for $p = 1/2$ (center) and 0 or 1 (hard left, hard right)? Use MATLAB to implement a tempo-synced panner that accepts as an input a mono signal and a period value in samples and returns a stero signal in which the input is panned with a *p* that oscillates linearly between 0 and 1 with a period as specified in the input parameters. Apply your effect on a 1 kHz sine wave and listen to the result for various values of the panning period. (You can also apply your algorithm to a stereo file, including recorded music, by processing each of the two input channels separately. If you match the tempo of the music precisely, you can create a very interesting effect.)

Problem 2 – THD+N

Consider a – bad – 16-bit D/A converter with the following input/output relationship:
$y = sign(x)\sqrt{|x|}\sqrt{x_{\max}}$, where $x_{\max}$ is the largest possible value of the input. Plot this function. Using MATLAB, compute the THD for this system. Assuming that white noise is added to the analog output with equivalent power -70 dBFS, compute THD+N. (Note: THD+N measurements are made typically at 1 kHz and near 0 dBFS, and should only include the audible spectrum. For the noise signal, you can use your generator from the last homework.)

Problem 3 – Pro Tools

The following problem is Pro Tools related. Make sure you install the Pro Tools sample session "Be There" from the Pro Tools installation CD as we will be using audio files from it. Note that Pro Tools includes a well-organized PDF version of the Reference Manual (under "Help") which you will need to consult as you work through the problem.

In this exercise we will explore basic functions of Pro Tools. Create a new session (16 bits, 44.1 KHz, BWF), and use as the name your Columbia (AcIS) login name with the suffix "-hwk3" (e.g., ae12-hwk3) . In the Audio Regions area, import the (stereo) files "Virus Arpeg-00" and "monkeys loop(L)-TCEX-01-16" from the "Be There" sample session. Create two stereo audio tracks, naming them "monkey" and "virus". Insert the corresponding audio files into each track, so that they start at 0 sec. Duplicate the "monkey" region enough times to last as long as the "virus" track. Play back the session. You will notice that the two tracks are not perfectly aligned; the "virus" starts a bit later. You can easily see this from their waveforms: select a small portion at the beginning and zoom in. Insert a marker named "Step 1" near 0 sec saving the zoom settings along with it so that we know you did this.

We will now align the two tracks. There are two (at least!) ways to do this: either delete the initial silent portion of "virus" and align the remaining region to monkeys, or shift the "monkeys" loop so that it starts at the point where the "virus" audio kicks in. (Hint: to accurately position yourself, use the "tab-to-transient" feature.) Next, insert a 4-band EQ in the "virus" track and trim the low-mid at 700 Hz down 12 dB, and the high-mid at 1,2 KHz down 6 dB. Listen to the result; the "virus" sounds a bit cleaner. You may want to adjust the levels of the tracks to your liking.

Now bring in the file "LV 2-16" (from "Be There") into a new mono audio track called "vocals". Pan it 14% to the left. You can put it to start where the other tracks start, but we want the vocals to come in later in the piece. To make it easier to navigate ourselves in the session, it's useful to create a tempo structure for the song, i.e., add information about tempo in beats per minute, as well as the time signature. (If you don't know anything about time signatures, the "monkey" is in 4/4 – four quarters – and thus contains four beats in one bar.) Using the "Identify Beat" command ("Edit" menu) set the first "monkey" loop to correspond to the very first bar of the song (1 | 1 | 00 to 2 | 1 | 000). Note that Pro Tools will automatically compute the right tempo and assume it is the same for the entire piece. While true for "electronic" music, this is definetely not the case for live music, unless a click track has been used to guide the performers. In fact we can use a click track to confirm the accuracy of our tempo track: create a new mono aux track named "click" and insert a "Click" plug in in it. Make sure the "Click" option in the "MIDI" pull-down menu is selected, so that MIDI clicking information is generated and picked up by the Click plug-in. Play back the session hearing carefully the alignment of the clicking with the audio. Mute the "vocals" track so that it doesn't get in the way.

Now that we have a structure in terms of bars and beats, move the "vocals" track to start at bar 3. Our next task is to add reverb to the "vocals" track. Create a new stereo auxiliary track (name it "reverb") taking input from bus 15-16, and add a "D-Verb" plugin as insert A. Add a send A in "vocals" going to bus 15-16 at -3 dB. For the "D-Verb", use the default options. You can audition the audio with and without the reverb by muting the "Reverb" track (you can do this in real-time, while audio is playing). Experiment with the send level and set it so that the effect of the reverb is audible but not too much. Reorganize your tracks so that "vocals" is first, and the "reverb" is next-to-last, and "click" the last one. Now is a good time to re-examine your fader levels; adjust them for good balance and high volume but avoid clipping!

Now we will add automation to the “vocals” track. Using the “auto touch” fader automation mode, write automation so that the second verse (“so you’ll see in the bright side”) is at a slightly lower level. Next, use the graphical automation editing tool (switch your edit view from “waveform” to “volume”) so that the level during the word “dark” is reduced by about 5 dB. Since the duration of the word is small, you won’t be able to do so using the fader.

To finish off our piece, we’ll let the “virus” and “monkey” play for a bar or two after the vocals end, and then fade out the “virus” over a period of two bars, and then fade out the entire piece to silence in the last three bars. Trim the “virus” track so that it ends in bar 10. Then implement a fade-out from bar 8 to 10. Using the custom fade shapes, select the one that has the fastest initial attenuation (but not the step-down one).

Now we will implement a fade-out for the entire song. The best way to do this is to create a stereo master fader track (name it “master” and put it as the last track in your session). Since we are at it, add a “POW-r” dither plug-in with type 3 noise shaping at 16 bits (suitable for broadband material such as music). Using fader automation or by drawing volume automation implement a fade out as smooth as you can over the final three bars (starting at bar 11, one bar after the “virus” is gone).

Now we are ready to record our masterpiece to disk. Make a selection from the beginning of your piece until the end, and then bounce the whole thing to disk (with the name “final-mix”). Make sure the click track is muted, otherwise it will be recorded as well. Also, import the bounced track back into the session so that we can easily audition it.

To submit your assignment, zip the entire directory where the project is located (using the same name as the Pro Tools file) and upload it to the ftp server flavor.ee.columbia.edu/incoming, with the user name and password set for the “Downloads” area of the course web site. Email the instructor when the upload is complete.