Music 192B Studio Exercise 1

This exercise will allow the use of all available mixing processes to be used in a mix. We will continue using Logic X Pro as the main software platform however you are also encouraged to try Pro Tools 12 if you wish. The goal is to become familiar with the use of the tools available to fine-tune a mix, which include equalization, dynamic range processing, effects and automation. Modern DAW systems allow great precision in adjusting the mix but there is a learning curve that can extend for a lifetime as we try different ways of achieving what might be an elusive goal: making a compelling presentation of our music that stands up no matter how a listener reproduces our work. There are as many ways to approach this task as there are mix engineers. The Internet is full of how-to videos that detail the various approaches. There's still no alternative to trying options yourself in order to understand what works best for you – that's what we will try to do.

Both Logic X Pro and Pro Tools 12 are host-based programs, meaning that all processing runs on the host computer – a Mac Pro. We also have Universal Audio plug-ins available that run on UA hardware connected via Thunderbolt but this is transparent to the user. I recommend using the built-in plug-ins at first, as they are quite good on their own before exploring the more specialized ones from Universal Audio and McDSP. For this exercise, we have several multi-track songs to mix, some live recordings and some studio recordings. The tracks should be imported into your session from the Audio One disk folder "!192b Exercise session soundfiles". To start with, there are four to choose from: John Mayall "Hideaway", Mike London "Tell Me, Tell Me", Papa Grows Funk "Ratta Tang Tang" and SoulXcetera "Ain't Too Proud To Beg". You only need to mix one for the exercise but are free to do more.

Using these tracks, we will learn how to organize a mixing session and how to use the functions of compression/limiting and automation in addition to equalization and panning. Import the desired files into a new session: you can assign them to new tracks or place them in the Media folder and drag them to the Logic Arrange window individually. Directly sending them to new tracks automatically creates new mix tracks with the proper file names so this might be preferable. You can then drag them into the order you prefer. I usually place drum files at the top, bass and rhythm instrument files next and vocals at the bottom but the reverse also works – see what makes it more intuitive for you as you will be doing a lot of moving around in the session to do automation and editing if necessary. (We don't need to do any editing here.)

One way of starting to assess the mix is to put up all faders up equally (at a level that doesn't overload the system) and listen to the result. You can start to hear what comes through and what doesn't. (In some cases, compression may have been applied in the recording process but we will assume the tracks we are working with are uncompressed.) Listen to how the tracks interact: do some sounds "come and go" as other sounds drown them out? This would be a good reason to use dynamic range management. Once we establish a need for compression and limiting, there are as many ways to accomplish the goal as there are compressors and/or dynamic range plug-in options. It is therefore necessary to understand the basic types of dynamic range manipulation that can be employed and each can produce rather different results.

The compression plug-ins we find are mostly designed to emulate one or another of the hardware compressors that have been used since their invention in the last century. Several of the traditional hardware approaches have produced an identifiable and distinctive sound. Although compressors essentially do the same thing, their exact behaviors differ depending on several factors, notably the type of gain-modulating circuitry and signal amplitude measurement circuitry. In the analog world in which the majority of hardware compressors were developed there were limited ways of creating an electronically modulated gain stage. Electro-optical devices were one option, a technique exploited by the LA-2A compressor/limiter. An FET could also be used as a gain element as in the 1176 compressor. The variable-mu vacuum tube could be used as a variable gain element as it was in the Fairchild 660/670 compressor. Voltage controlled amplification (VCA) could be accomplished using integrated circuits that made complicated circuits available as IC chips like those used in the dbx series of compressors. The specific characteristics of the gain element contribute significantly to the sound of the compressor.

The built-in Logic X compressors model some of the options listed above:

Platinum Digital	Modern digital compressor with synthetic distortion if desired
Studio VCA	Focusrite Red 3 clean VCA with added distortion if desired
Studio FET	Universal Audio 1176 Rev E (“Blackface”)
Classic VCA	dbx 160
Vintage VCA	SSL G bus compressor
Vintage FET	Universal Audio 1176 Rev H (“Silverface”)
Vintage Opto	Universal Audio LA-2A

(The one style no longer represented in the stock Logic compressors is the variable-mu type, an unfortunate development.) Just using the built-in compressor plug-ins, there are plenty of possible options that will require significant exploration before we can be confident about the best option for a given job. To start, we will use just the built-in compressors.

In addition to the gain element, another circuit can affect the sound of the compressor – the signal measurement circuit. The issue of how loud a signal will sound is complicated and not easily determined, but we need a control signal to send to the gain element that accurately tracks the amplitude. Several circuit types can be used to find the instantaneous amplitude of the signal: average, peak and RMS. Average measurement is the simplest, it can be generated using an op-amp full-wave rectifier followed by a low-pass filter. This will not reflect instantaneous changes in the amplitude so peaks in the signal can be missed. A peak-tracking circuit works faster in order to continuously find the instantaneous amplitude but it is more complicated than the simple rectifier. The most complicated approach is to measure the actual root-mean-squared amplitude, which best reflects our perception of loudness. It also fails to track peak amplitudes but better reflects how we would perceive the signal's loudness with our ears. Each of these detector circuits imparts a different sound in addition to the contribution of the gain element itself. These circuits are part of the digital simulations for the particular compressors modeled in the plug-ins.

While the difference between a compressor and a limiter is one of degree, they can sound quite different. Limiters have a much higher compression ratio than compressors and allow a higher threshold of action while keeping the dynamic range the same. A compressor threshold set lower means that more of the range of amplitudes is altered than with a limiter set to hold the maximum output to the same level. This results in a more audible change in the relationship between the soft and louder components than we get with a limiter. The limiter affects only the highest peaks, drum hits for example, while making everything below the peaks linearly louder.

In addition to level-related settings, most compressors have attack and release time adjustments. The attack time determines how long it takes the gain reduction element to start reducing the gain and the release time determines how long it takes the gain to return to its original level following the signal drop below threshold. While it may not be obvious, this can dramatically affect how the compressor sounds. Much of the character of a particular compressor comes from how fast it reduces the gain after it detects a signal crossing the threshold level. Sounds in general contain a lot of information in the onset part of the signal, as opposed to the steady state part. Since so much of the sound character is contained in the onset, we often want the compressor to let that part through before clamping down on the longer part of the sound. An attack time of several tens of milliseconds is often necessary to preserve the identity of the compressed sound. In the case of limiters, however, fast attacks are required to prevent early transients from overloading the system. Release times are also adjustable and depend on the sound's time course to prevent audible pumping as the gain is changed.

The complementary process to compression is expansion: creating an increase in dynamic range. The ultimate example is a noise gate that cuts off the signal entirely when its level drops below the gate's threshold level. The purpose is often to remove low-level noise but expansion can be used to tailor the envelope of a signal as well. For example, often a microphone inside a kick drum produces a very filtered version of the ambient sound as it resonates with the room sound. An expander that allows the drum hits to exceed the threshold while the quieter leakage is removed can help this. This sometimes works very well. While a noise gate is either on or off, an expander can have variable gain below threshold so low-amplitude sounds are not completely eliminated but are reduced in amplitude with an input/output ratio below threshold. This is the inverse of the compressor, where gain is reduced when the signal exceeds the threshold.

Compression and limiting are powerful tools but they have their limits. Very large swings in amplitude can overload the plug-in and cause audible distortion and overly intrusive compression on very loud parts. Often, we want to automate the signal level before it reaches the compressor (or other channel plug-in). Careful automation of vocals, for instance, help the compressor maintain a more constant sound as the signal level contains no large excursions. In Logic (and Pro Tools), the audio tracks' plug-ins are pre-fader. This means the signal goes to the compressor before it hits the fader, hence any volume changes introduced by fader automation occur after the compressor (or any plug-in). Often we want to use fader automation before the compressor – how do we do it? The answer is to leave the audio track without plug-ins and introduce whatever level automation we desire there. The audio track output is then routed not to the stereo output but to a bus that is routed to the input of an aux track with the desired post-fader plug-in inserted. If you find a compressor incapable of taming a particularly wide dynamic range signal, level automation can help. In Logic, aux tracks are created automatically when a new bus is selected. In Pro Tools, they are created from the new track menu.

Automation is very powerful but can be quite time-consuming. It allows us to micro-manage the sounds by changing amplitude, gating, and remembering almost any other mixing change we wish to make. The trick is to use the minimum amount that gets the job done. Sometimes we cannot find a threshold for a gate that doesn't cut off some of the wanted output signal – in that case we can try mute automation to manually set the track on and off. This certainly takes more time than adjusting a gate threshold but it gives us a precision we lack with most expanders. Automation can be accomplished "on-the-fly" by recording fader moves or other mouse-modified parameters while the song plays. It can also be done using a graphical representation that can be altered with the mouse. Since amplitude automation is most widely used, we can look at how it is accomplished and use similar methods to automate any parameter in a mix or plug-in.

On-the-fly automation is accomplished by playing the mix with one (or more) of the channel automation buttons set to touch or write (see below).

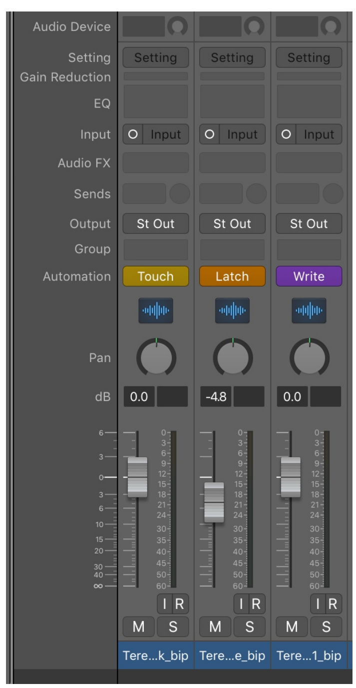

Touch starts writing new automation whenever the fader or other control is moved. When stopped, the automation returns to whatever was previously written, if anything. In Latch mode, the current value is written when recording is stopped. Write overwrites whatever was there previously whether or not a value is changed and can be confusing. Once automation is written, the button can be set to “Read”. Most often, Volume and Mute automation are used but any parameter of any plug-in can also be automated, as can sends to busses. While it is possible to write automation in real-time, it can also be done graphically while the audio data is displayed, making small changes that are linked to the audio waveform.

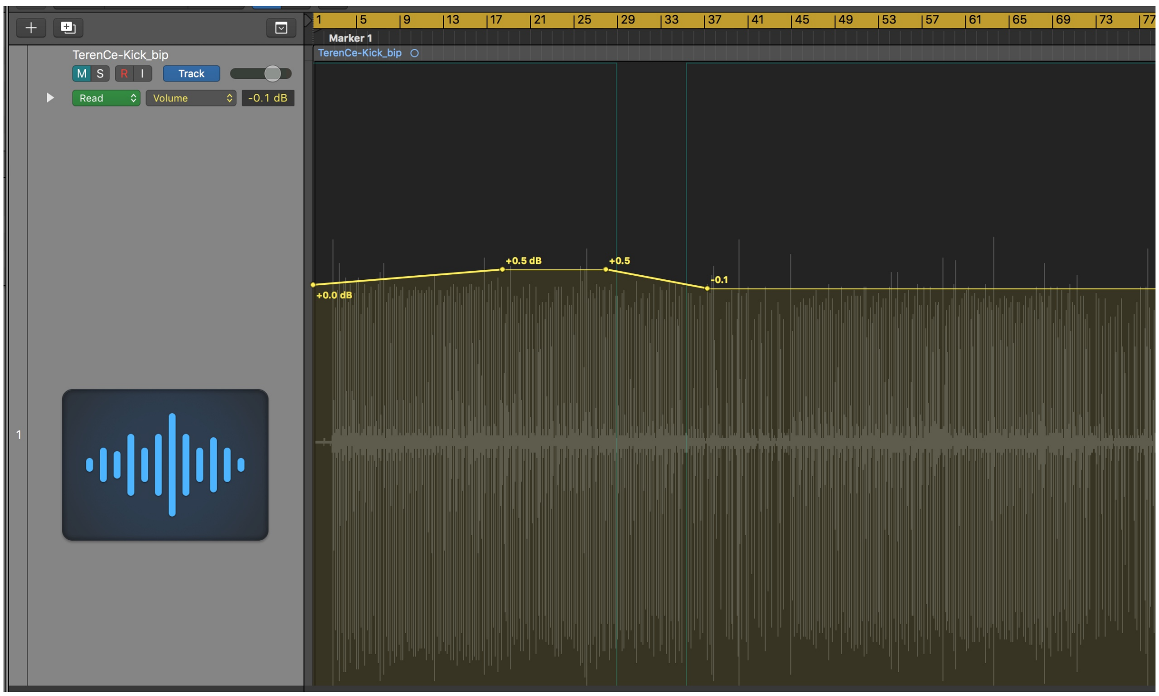

We see Volume automation displayed graphically above. Also displayed (light blue lines) is Mute automation. The target for automation is selected with the automation menu button – in this case it says “Volume”. Automation values can be set by clicking and dragging the break points and lines. Since the waveform is visible, necessary changes to volume or other parameters can be made graphically in synchrony with the sound, something not possible in the analog automation world.

In addition to processing the original audio tracks, we can add effects like reverb, chorus and flange either individually to tracks or globally by using aux tracks and busses to select the desired inputs. There is a tendency to add these effects to separate tracks and for some applications this is fine, but particularly for reverb we often want a single reverb for vocals, another for drums, etc. By using a global effect, we save processor power as we only need to run one instantiation of the effect to serve multiple tracks. We also get a “glue” that can make vocals sit together, for example.

These are some suggestions for getting started – the more you mix the better you will get. We will present our mixes in class so we can learn from each other.