

Digital Sound Effects

Introduction

Sound effects for musical audio have been around since the advent of electric instruments, where effects were implemented with electronics that directly modified the analog audio signal. Nowadays, **digital signal processing** allows us to implement these effects and many others more flexibly and cheaply by writing code to manipulate digital representations of analog signals. In this activity, we will be exploring three effects, how they're implemented, and what each does to an audio signal in time and frequency domains. The three effects are **echo**, **distortion**, and **filters**.


This activity sheet contains three sections, each with instructions for setting up the experiment. We ask that you make changes to each app only as instructed, and wait for the instructor's directions before and after each section. Feel free to ask any of the instructors for clarification, and have fun!

Echo

Most of us are familiar with the sound phenomenon that we call echo. In the musical world, different types of echo effects also go by the names *reverb* and *delay*. *Chorus* and *flangers* are also types of delay effects. In this section, we will investigate the source of echo in the physical world and examine two computer simulations of echo effects. Using the iPad DAW (digital audio workstation) **Garageband** and modular synthesizer app **Dingsaller**, we will:

- Introduce a model of a single echo in space/time
- Experiment with modeling a range of physical environments by varying echo parameters
- Modify the models toward a greater complexity and realism by adding feedback

Activity 0:



Let's take a walk, listen, and think about the sounds we hear

Activity 1.1: Single Echo Visualization

1. Open the GarageBand app and open the project file "Echo - Activity 1"
2. By now you should be somewhat familiar with the playback controls. Press play and listen to the clap sound at the beginning of the track a few times.

Note: At any time, you can press the question mark icon in the upper-right corner for explanations of the various features, or just ask an instructor!

3. With the track selected, tap again to bring up the menu shown and copy the track.



Figure 1: Clip Menu

Tap anywhere in the workspace below the track to bring up another menu to paste the track. The beginning of the track will be aligned with the current position of the playhead.

4. You can shift the new track copy left and right by selecting and dragging the track. Zoom in on the timeline with a reverse two-finger pinch so you can move the track by smaller time steps.

Listen to the echo resulting from large and small time shifts. Does this sound like any of the echoes we heard in Activity 0? Add a third copy of the track. Does it sound like a real echo yet?

5. Swipe the tab on the column on the left side of the screen to open the track controls.
6. Use the volume controls on the three tracks to change the relative volumes of the time-shifted tracks; experiment with different values and more tracks. How many copies should you add to get a natural sounding echo like the one we heard in the long hallway? How should the amplitudes of the copies be set?



Figure 2: Track Controls

Activity 1.2: Echo Effect in a Signal Flow Model

Dingsaller is a graphical modular synth environment of the iPad. Using Dingsaller, we can simulate a sound source running through some effects, much like you would plug a guitar into a stompbox, then into an amplifier. This type of simulation is called a **signal-flow model**. We can apply effects to the outputs of virtual instruments and control effect parameters in real-time.

At the bottom of the workspace is the output object. We can connect other objects like instruments, effects, or other sound sources when we want the sound to go to the iPad speakers or headphones.




Figure 3: Instruments Menu

1. Let's begin by placing a virtual instrument in the workspace. Open the **instruments menu** by swiping the tab in the upper-left corner of the workspace, and tapping on the violin icon.
2. Add the **SimpleSampler** to the workspace. Double-tap on the SimpleSampler block to bring up its parameter menu.
3. Tap **Load File** in the upper-left corner of the sampler's parameter menu. From the list of samples that appears, choose the **Click** sample.

4. Now we need something to trigger the sample. Again, swipe the same tab you used to bring up the instruments menu. Now, tap on the left-most icon to bring up the **Notes** menu. Tap **Keyboard** to add it to the workspace.
5. Note we still haven't connected the SimpleSampler to the output. Also notice the sampler object has green and orange terminals. In Dingsaller, the orange terminals carry **audio signals**, while the green terminals carry **control signals**. We want to use the keyboard to control the sampler, and send the sampler's output to the iPad speakers. Can you figure out how to connect the blocks?
6. Once you've got the sampler connected, double-tap the keyboard block to bring up the on-screen keyboard. Play a few notes on the keyboard and note the dry, anechoic sound of the single click. Now, we want to connect some effects! Swipe to enter the menu in the upper-left of the workspace again, and now select the right-most icon to bring up the **Effects** menu. Add a **MultiTap Delay** to the workspace.
7. Connect the MultiTap Delay between the sampler and output blocks.
Note: Grab a terminal and swipe up to break existing connections.
8. Your connections should look something like Figure 4

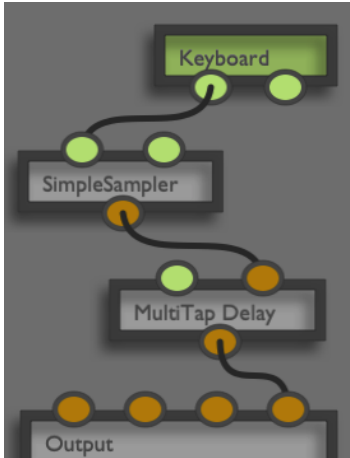


Figure 4: Completed Connections

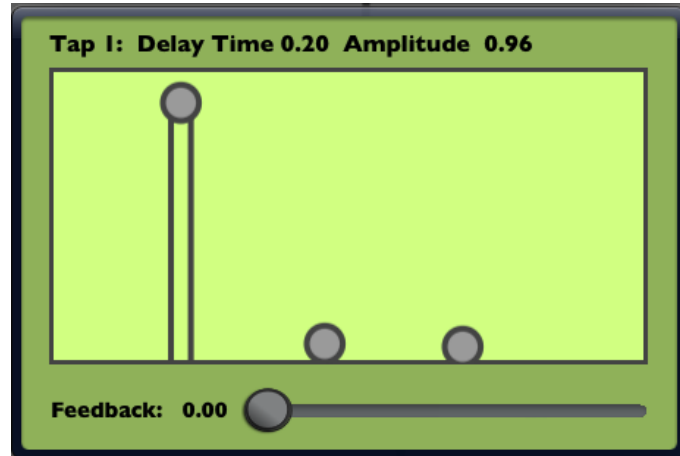


Figure 5: Delay Parameter Menu

9. Double-tap on the **MultiTap Delay** block to open its parameter menu. The three gray dots (called *taps* in Dingsaller) inside the light green square represent three time-shifted copies of the input signal. Drag the taps left and right to change the delay between the original sound and the echo. Drag the taps up and down to change the intensity or amplitude of the delayed copy.
10. The **feedback** slider controls the amount of delayed signal that gets sent back to the Multi-Tap Delay's input. If you've experienced a microphone in too-close proximity to a PA system, then you've experienced feedback. Feedback helps us create more realistic sounding echoes with only a few delay taps, since it effectively adds more time-shifted copies to the signal.

Note: Be careful with the feedback slider! Certain combinations of tap amplitudes and feedback slider values can quickly cause runaway feedback that can damage your hearing.

11. Try setting the feedback slider with *only a single-tap echo* so that it sounds like our multiple echoes in GarageBand. What is a good value for the feedback slider?
12. Which sounds more natural: delay times set to integer multiples (e.g. 0.2 seconds and 0.4 seconds) or unrelated numbers (e.g. 0.31 and 0.5)? Which sounds better to you?
13. What is the minimum delay time for a *single echo* that you actually perceive two distinct sounds instead of one?

Harmonic Distortion

Distortion, as the term is generally used in signal processing, means an alteration of a signal, with the implication that the alteration is undesirable. When musicians refer to distortion, they are most often talking about a specific kind of distortion called *harmonic distortion*, which since the 1950s has gone from being an unwanted design-flaw in electric guitar amplifiers to a highly desirable effect.

Harmonic distortion is a process that introduces frequencies to a signal that are **harmonically related**, meaning the added frequencies are **integer multiples** of the original frequency or frequencies.

We'll begin by opening the **SoundFX** app.

Activity 2: Hard Clipping

1. Add a single fundamental frequency (the first harmonic slider), and boost the amplitude a little with the master volume slider to the left of the time-domain plot. Keep the amplitude under 0.75.
2. Now pinch the time domain plot as if you were zooming out of a photo. This brings down two red lines. These lines show a signal amplitude at which anything above that amplitude is clipped off the top and bottom of the waveform.




Figure 6: Hard Clipping Thresholds

3. What happens to the spectrum of the signal when the tops and bottoms of the waveform are clipped off? Which harmonics are generated? What does the time-domain waveform look like?
4. Try distorting a more complex combination of sine waves. Look at the harmonic amplitudes of the signal before distortion (the gray dots on the spectrum) vs. the spectrum plotted in purple while you introduce distortion. Does hard clipping always generate odd harmonics for every frequency? Is there some other simple relationship between input harmonic amplitudes and the distorted harmonic amplitudes like in the case of a single sine wave?

Note: Make sure when adding harmonics, that you use the master volume slider to keep the overall signal level within the bounds of the plot or it will introduce distortion beyond your ability to raise the hard-clipping threshold.

Filtering

We've seen that echo and harmonic distortion are both created by manipulating a signal in the **time domain**, now we'll examine an effect that manipulates a signal in the **frequency domain**.

Filters allow us to amplify certain frequencies in a signal without affecting (or possibly attenuating) other frequencies. The "tone" knob on a guitar is one example of a filter. Any type of audio **equalization** (like the EQ on your home stereo) is also performing filtering. Effects units like **envelope filters**, **wah-wah pedals**, and **graphic EQs** are obvious applications of filters, but most effects have some type of filtering involved.

Activity 3: Drawing Filter Shapes

1. Let's begin by synthesizing a sawtooth wave (do you remember how?). Once you've gotten something resembling a sawtooth wave, press the **Draw** button above the spectrum plot.
2. Now, in draw mode, you can draw a filter shape with your finger over the spectrum plot. At each point on the x-axis (frequency), the value on the y-axis (amplitude) of this filter will be multiplied with the signal's amplitude *at that particular frequency*. This is how filters are used to selectively amplify or attenuate certain ranges of frequencies without affecting others.
3. Draw a filter that you can use to isolate **only the fundamental frequency** of the sawtooth wave. When you're done drawing, press **Draw** again to return to the normal spectrum plot. The filter curve will be displayed in blue. Does your filtered waveform look like a single sine wave?
4. Now, using two fingers, we can slide this filter curve back and forth in frequency. Slowly slide it up in frequency, encompassing more harmonics as you go. Observe how the wave gradually transitions from a sine back to its original sawtooth shape when all harmonics are included. What happens to the difference between the source harmonic amplitudes (gray dots) and the filtered harmonic amplitudes (purple plot).
5. Now draw a filter that passes only a small range of middle frequencies and suppresses the rest as shown in Figure 7. Slide it up and down in frequency. Does this sound like anything familiar?

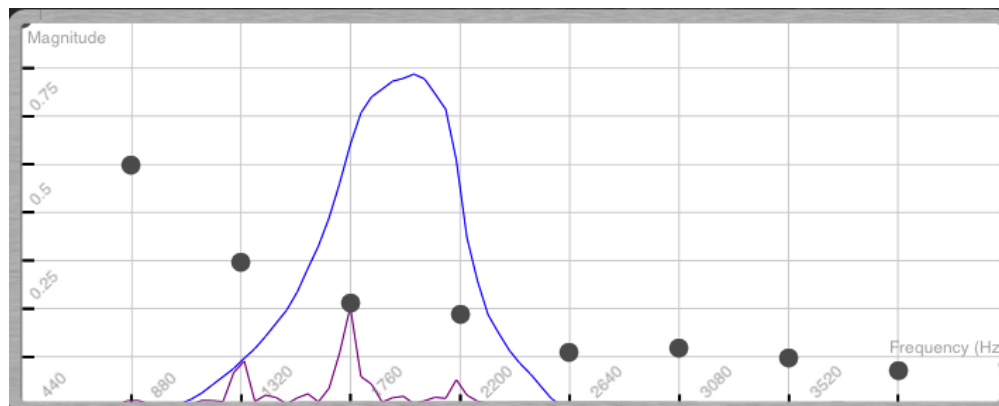


Figure 7: Mystery Filter