Delay

Topics Addressed

- Delay
- Frequency Response
- Linearly Interpolated Delay
- Recirculating Delay
- Comb Filtering White Noise

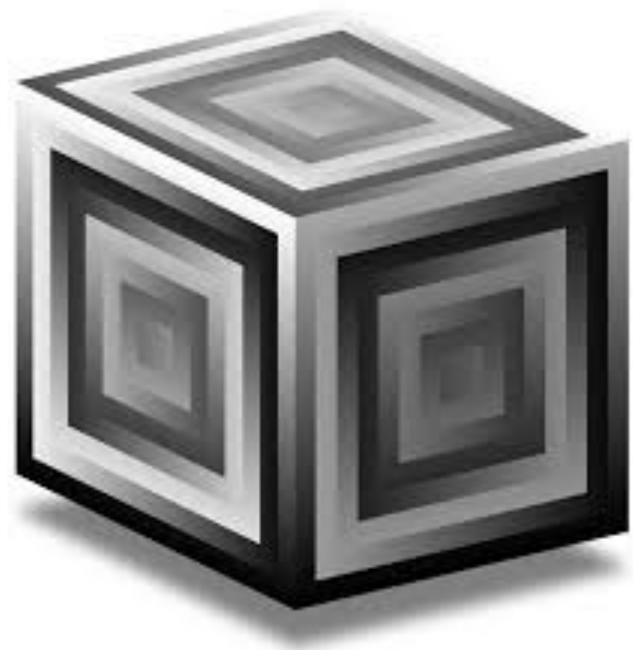

Delay

- Delay is a basic concept in electronic music that means to playback a signal after some unit of time.
- Delay is ubiquitous throughout electronic and popular music.
  - U2 – Where The Streets Have No Names
  - Lady Gaga – Just Dance
  - Etc.
- Delay is a foundational building block of two other important digital music concepts
  - Filtering
  - Reverb

Modeling Delay

This delay is called a feedforward comb filter

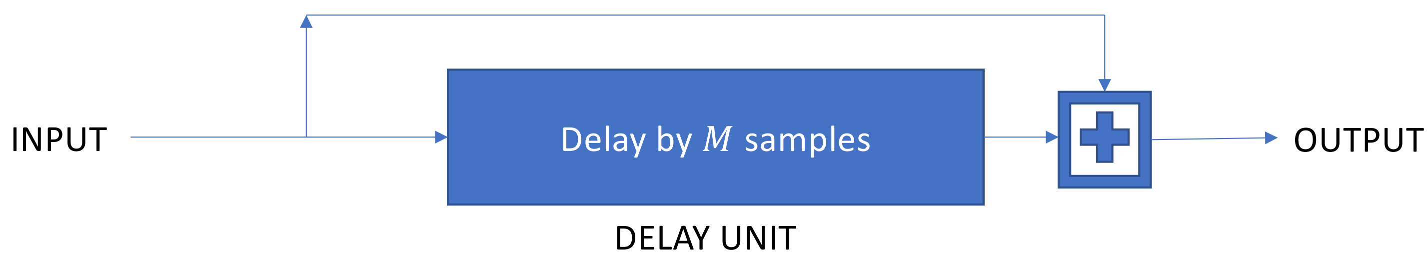

In the basic delay setup, an input signal is added to a copy of the same signal that has been passed through a delay unit which delays the incoming signal by some unit of time. For now, we will deal with *audible* delays, meaning ones where we can perceive the two signals as separate but simultaneous streams of sound. When the delay between signals gets incredibly small other interesting effects can occur.

Computing Delays

- In the digital world, a signal is simply an array of samples.
  - Given a signal $x[n]$ where n is the nth sample in the signal x, then the delayed signal of $x[n]$ by M samples is $x[n-M]$.
  - How can we convert M to a measure of time? Simply divide M by the sample rate to convert to seconds
- To compute the resulting signal based on the model of the previous slide, we simply need to compute $x[n] + x[n-M]$
- Oftentimes, the delayed signal is scaled to a smaller amplitude to create an echo effect. Mathematically, this is $x[n] + g_f x[n-M]$ where g_f is the feedforward coefficient.

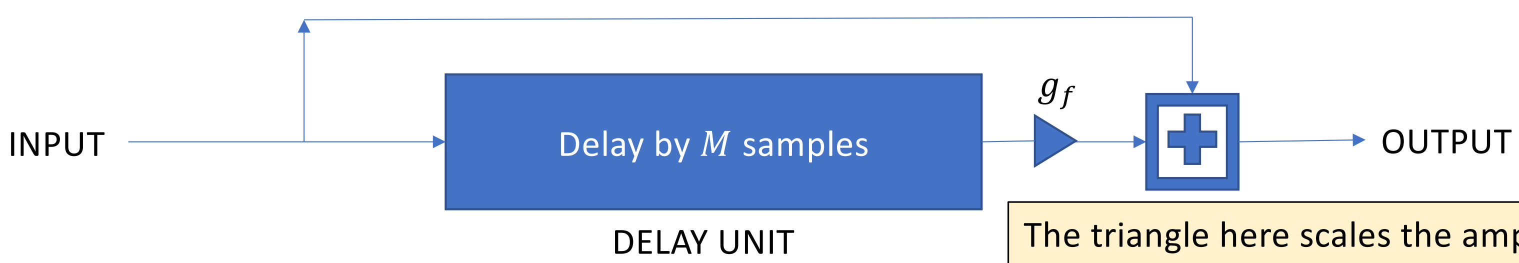

The triangle here scales the amplitude of the delay. It is equivalent to multiplication.

Simple Example

Assume that $M = 2$

$x[n] =$

Sample #	$x[0]$	$x[1]$	$x[2]$	$x[3]$	$x[4]$	$x[5]$
Amplitude	0.5	0.4	0.3	0.2	0.1	0.0

We will simply state the amplitude of negative indices as 0.

$x[n - M] =$

Sample #	$x[0 - 2]$	$x[1 - 2]$	$x[2 - 2]$	$x[3 - 2]$	$x[4 - 2]$	$x[5 - 2]$
Amplitude	0	0	0.5	0.4	0.3	0.2

$x[n] + x[n - M] =$

Sample #	$x[0] + x[-2]$	$x[1] + x[-1]$	$x[2] + x[0]$	$x[3] + x[1]$	$x[4] + x[2]$	$x[5] + x[3]$
Amplitude	0.5	0.4	0.8	0.6	0.4	0.2

Delay in SC

```
(
SynthDef(\simpleDelay, {
	// Outputs a stereo signal
	arg out = 0, delayTime = 0.5, g_f = 0.4, in;
	var sig, sigDelayed;
	sig = In.ar(in, 2);
	sigDelayed = DelayN.ar(sig, 1, delayTime, g_f);
	Out.ar(out, sig + sigDelayed);
}).add;
)
```

- DelayN is the basic delay in SuperCollider. It simply delays the start of signal.
- To implement the model on slide 4, we must add the signal from DelayN to the original signal.
- DelayN is implemented using a buffer to store the samples from the signal until they are ready to be outputted after the delayed time. The second argument, the maximum delay time, is necessary to determine the size of that buffer. The third argument should **never** exceed the former.

Samples from a triangle wave

Amplitude Response

Sample *n*	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
Amplitude	1	0.5	0	-0.5	-1	-0.5	0	0.5	1	0.5	0	-0.5

- Suppose we pass the samples above through a feedforward comb filter with delay of 8 samples. What would happen to our signal?

- Suppose we pass the samples above through a feedforward comb filter with delay of 4 samples. What would happen to our signal?

Conclusion: we can see in this small example that we were able to double or zero out all the frequencies of a triangle wave. Depending upon the delay time, certain frequencies will be boosted or attenuated. We call this the amplitude response.

Amplitude Response

- Formally, the amplitude response (sometimes called the magnitude frequency response) is the factor of amplitude gain as a function of frequency.
- We will eschew the math to deduce the frequency response of the feedforward comb filter and present it in graphical form.

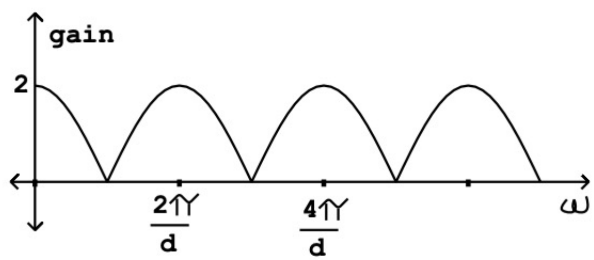

w is referred to as angular frequency and is simply a scaled version of frequency. $w = 2\pi f$

The delay time determines which frequencies will be amplified in a complex input signal

Picture from Miller Puckette's book

Filter

- Any electronic or digital system that boosts/attenuates frequencies in a given signal is called a filter.
- Delay networks always provide some level of filtering.
- This delay system is sometimes called a feedforward comb filter.
  - Important aside: when most people use the term comb filtering they are usually referring to a different delay system called a feedback comb filter. For clarity, specify either feedforward or feedback comb filtering.
- This filter is called an FIR filter (finite impulse response) because there are finite number (i.e., 1) of echoes from the filter.

Delay Time Quantization

- Most of the delay UGens in SuperCollider are controlled through time in seconds. Given some delay time t and sample rate f_s, the delay in samples can be computed as tf_s. Note though that tf_s is not guaranteed to be a natural number.
- For the UGen `DelayN`, when tf_s is not a natural number, the delay in samples is quantized to the nearest whole number. In this instance, quantization simply means rounding.
- Other UGens like `DelayL` and `DelayC` use interpolation to estimate the proper delay. The most common is linear interpolation and works analogously to linear interpolation in wavetable synthesis.

Linear Interpolated Delay

- Suppose we have a signal $x[n]$ that is delayed by some $tf_s = M + \Delta m$ where t is the desired delay in seconds, f_s is the sample rate, M is $\lfloor tf_s \rfloor$ and Δm will be between 0 and 1 which is the fractional number of samples.
- We can't determine $x[n - M - \Delta m]$ directly because we only have integer numbered samples and no fractional samples.
- We will estimate $x[n - M - \Delta m]$ with linear interpolation. We will call this $\hat{x}[n - M - \Delta m]$.
  - As a side note the carrot hat in machine learning and statistics signifies predictions/estimates
- Linear interpolation has the following definition:
  - $\hat{x}[n - M - \Delta m] = x[n - M] + \Delta m * (x[n - M - 1] - x[n - M])$
  - When the index into x is negative, we can simply say the amplitude is 0.

Linear Interpolated Delay

Sample *n*	**0**	**1**	**2**
Amplitude	1	0.5	1

Let's estimate $x[n - tf_s]$ where $n = 2$ and $tf_s = 1.3$ samples

- $x[2 - 1.3] = x[0.7]$
- $x[0.7] \approx \hat{x}[0.7]$
- We know $M = 1, \Delta m = 0.3, n = 2$
- $\hat{x}[0.7] = x[1] + 0.3 * (x[0] - x[1]) = 0.5 + 0.3 * (1 - 0.5) = 0.65$

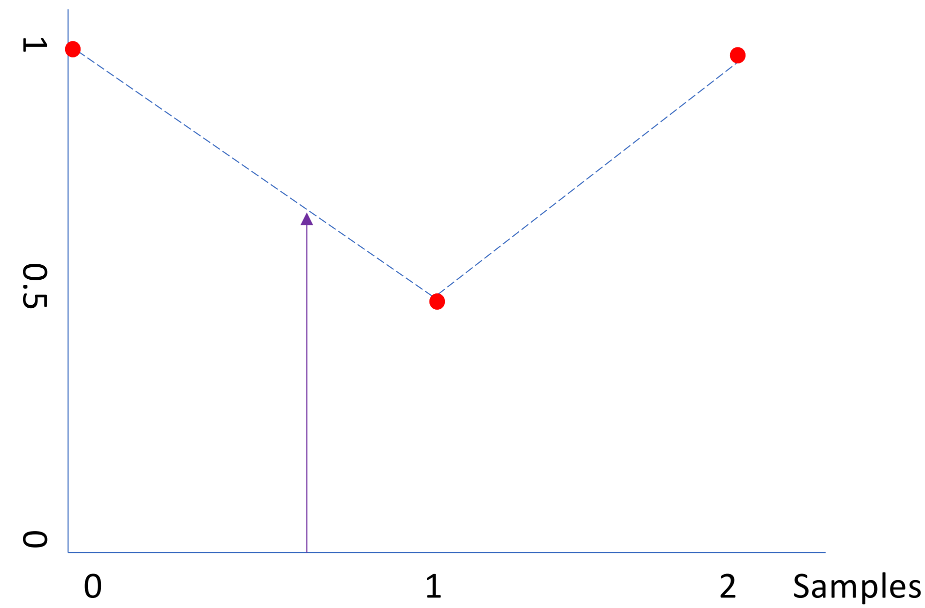

$$\hat{x}[n - M - \Delta m] = x[n - M] + \Delta m * (x[n - M - 1] - x[n - M])$$

Linear Interpolated Diagram

$$\hat{x}[n - M - \Delta m] = x[n - M] + \Delta m * (x[n - M - 1] - x[n - M])$$

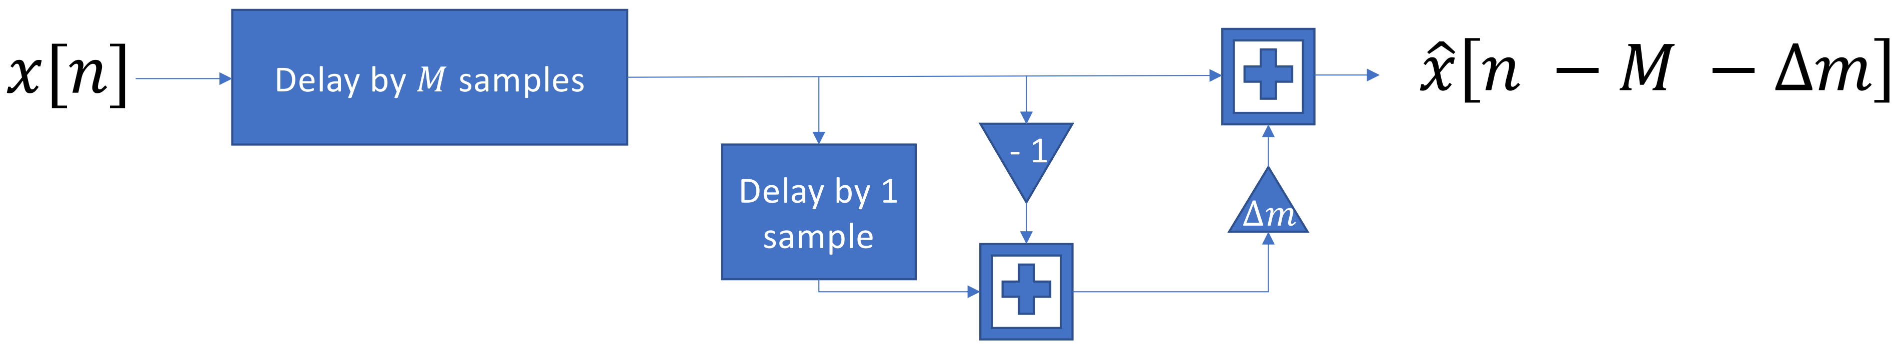

A block diagram of the linearly interpolated delay. The triangles are equivalent to multiplication.

Linearly Interpolated Delay

- Linear Interpolation tends to work best when the sample rate is high.
  - Why? Less error when there are more samples per period of a given frequency.
- SuperCollider offers cubic interpolation.
  - Provides more accurate samples
  - More expensive computationally

Exercise

- Given the signal $x[n] = [0, 1, 0, -1, 0, 1, 0, -1]$, what are the first eight samples of $x[n \ - 1.5]$?

- What is the amount of delay in seconds of 1.5 samples if the sample rate is 44,100Hz?

Recirculating Delay

This delay is called a feedback comb filter

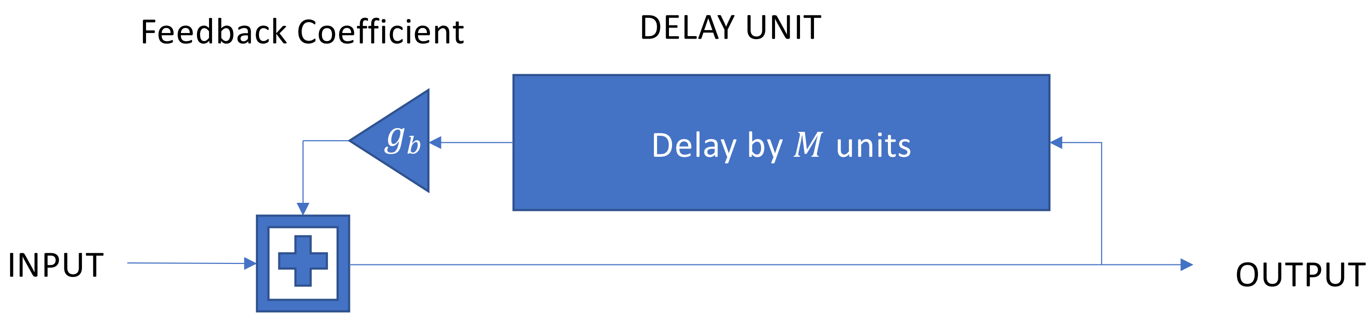

This is also referred to as a recirculating comb filter

$$y[n] = x[n] + g_b y[n - M]$$

Impulse Response

- One way to test this new design is to use an input signal of a single sample at amplitude 1. This is what is termed an impulse and can be used to see how the delay network responds.
- Impulses are used in the testing of delays, filters, and reverb

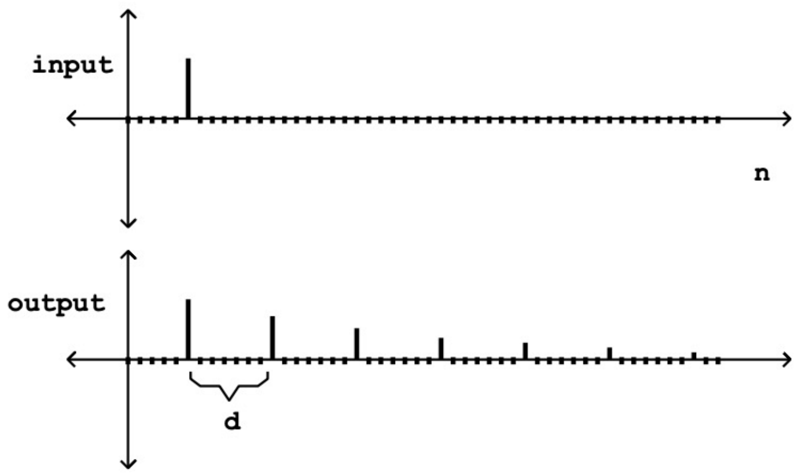

Picture from Miller Puckette's book

Recirculating Delay

Let's assume that the delay M is set to two samples and the feedback coefficient is 0.9.

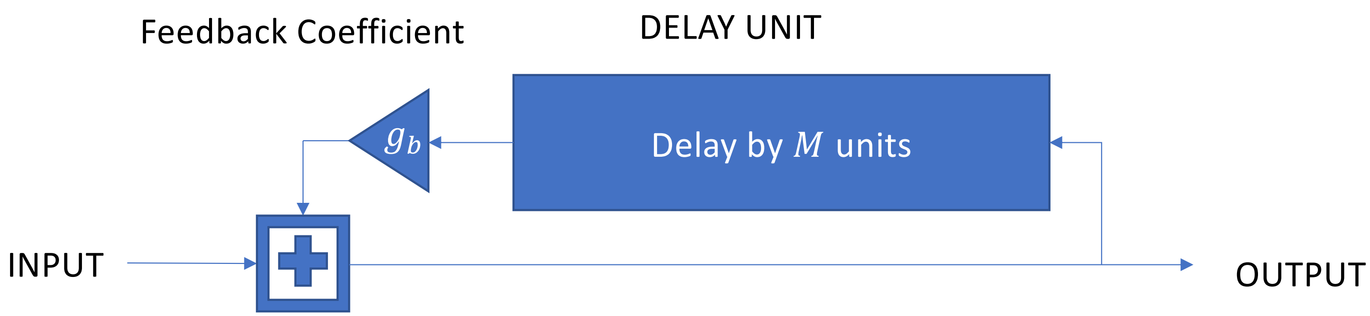

Sample *n*	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Amplitude	1	0	0.9	0	0.81	0	0.729	0	0.6561

From impulse

$$y[n] = x[n] + g_b y[n - M]$$

Feedback Coefficient

- The feedback coefficient allows the feedback comb filter to control the strength of the echoes. We will express it with the variable g_b.
- When the feedback coefficient is less than one, the network is *stable*
  - The strength of the echoes falls down toward zero
- When the feedback coefficient is one or greater, the network is *unstable*.
  - The signal will continue to grow in amplitude
  - Clipping and distortion will occur once the amplitude exceeds 1.0/-1.0
  - You should not do this.
- The feedback comb filter is an example of an IIR filter (Infinite Impulse Response) because the feedback process will create an endless number of decaying impulses.
  - All IIR filters use feedback. As a result, these are sometimes referred to as recursive filters.

Comb Filtering Amplitude Response

- Like the feedforward comb filter that was first introduced, certain frequencies are amplified/attenuated when a signal is passed through a feedback comb filter.
  - Because of the changes to gain at certain frequencies, the feedback comb filter is a filter.
  - It has an amplitude response that expresses the ratio of that amplification as a function of frequency.

Feedback Comb Filter Frequency Response

Peak Gain Ratio = $1/(1 - g_b)$

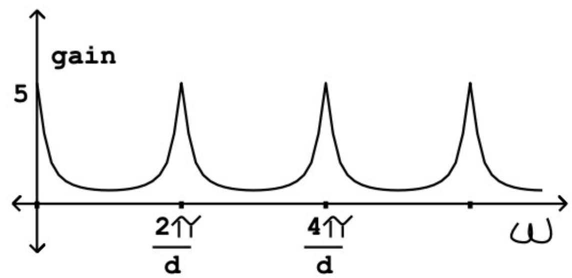

w is referred to as angular frequency and is simply a scaled version of frequency.
$w = 2\pi f$

The delay time determines which frequencies will be amplified in a complex input signal

The sharp peaks of the feedback comb filter explain the origin of the name. In general, when musicians/scientists refer to comb filtering, they are referring to feedback comb filtering or some combination of feedback/feedforward comb filtering.

Picture from Miller Puckette's book

Comb Filtering in SC

```
(
SynthDef(\fbComb, {
	arg out = 0, delayTime = 0.2, decayTime = 1, in;
	var sig, sigCombed;
	sig = In.ar(in, 2);
	sigCombed = CombN.ar(sig, 2, delayTime, decayTime);
	Out.ar(out, sigCombed);
}).add;
)
```

- The combed signal here also uses a buffer to store the delayed signal until it is to be outputted. The second argument is the maximum delay time which determines the length of the buffer.
- The feedback coefficient is not provided directly but is instead inferred by the decay time. See the documentation for exactly how that is calculated.

Comb Filtering Noise

- An interesting application of comb filtering is to pass noise through a comb filter.
  - Noise is defined as equal energy throughout the harmonic spectrum and can be implemented by using random samples between -1 and 1.
  - This is what is termed white noise. There are other kinds of noise that changes the energy throughout the frequency spectrum.
- Because comb filters so strongly resonate certain frequencies as determined by the feedback coefficient, we can actually extract pitches from noise.
- See the lecture code for today for examples.

Conclusions

- Delays and filtering are intimately tied together. Creating one almost always invariably creates the other.
  - Delay is the process of delaying a signal. It's definition speaks for itself
  - Filtering is the process of changing the gain/phase of a signal as a function of frequency.
    - More on phase soon
- Nearly all mastered/mixed audio you hear in songs has undergone a complex web of filtering.