

MUMT 618: Week #2

1 Time-Varying Delay Effects:

A variety of common audio effects are implemented using variable-length delay lines. In this section, we analyze several of these techniques. Note that a discussion of artificial reverberation will be “delayed” to a later section.

1.1 Flanging

- Flanging involves the summing together of a signal and a time-varying delayed version of itself.
- The input-output relationship for a flanger is given by:

$$y[n] = x[n] + gx[n - M[n]],$$

where $M[n]$ is the time-varying length of a delay line and g is the “depth” of the flanging effect. A flanger block diagram is shown in Fig. 1.




Figure 1: A digital flanger block diagram.

- Because the delay-line length, $M[n]$, must change continuously and smoothly through time, it is necessary to make use of an interpolating delay line.
- At any instant in time, the flanger is equivalent to a feedforward comb filter, which has a frequency response as shown in Fig. 2.
- For $g > 0$, there are M peaks in the frequency response, centered about the frequencies $\omega_k = 2\pi k/M, k = 0, 1, \dots, M - 1$. Between these peaks, there are M notches at intervals of f_s/M Hz.
- As M changes over time, the peaks and notches of the comb response are compressed and expanded. The spectrum of a sound passing through the flanger is thus accentuated and deaccentuated by frequency region in a time-varying manner.
- The delay-line length of a flanger is typically modulated by a low-frequency oscillator (LFO). Oscillator waveforms are typically sinusoidal, triangular, or exponential.
- For a sinusoidally varied delay,

$$M[n] = M_0 \cdot (1 + A \sin[2\pi f n T]),$$

where f is the flanger “rate” in Hz, A is the “excursion” (maximum delay swing), and M_0 is the average delay-line length that controls the average notch density.




Figure 2: Magnitude response of a feedforward comb filter with $M = 5$, $b_0 = 1$, and $g = b_M = 0.1, 0.5$, and 0.9 .

- For values of $-1 \leq g \leq 0$, the peaks and notches of the comb filter trade places. In practice, g is normally constrained to the interval $[0, 1]$ and the option of sign inversion is provided by a “phase inversion” switch.
- In the inverted mode, a notch is located at zero frequency. As a result, bass response will likely be weakened.
- Some flangers also implement feedback (in addition to the delayed feedforward path), which introduces spectral peaks and notches as previously described for feedback comb filters.
- The following MSP patch implements a flanger.

1.2 Phasing

- A *phaser*, or *phase shifter*, is similar to a flanger in that it sweeps notches through the spectrum of an input signal. But while a flanger provides only uniformly spaced notches, a phasor can modulate the frequencies of *non-uniformly* spaced notches.
- A phaser is implemented with allpass filters instead of delay lines, as shown in the block diagram of Fig. 3.




Figure 3: A digital phaser block diagram.

- Second-order allpass filters are particularly convenient to use because each can control a separate notch frequency and bandwidth. Second-order allpass filters have a transfer function given by

$$H(z) = \frac{a_2 + a_1 z^{-1} + z^{-2}}{1 + a_1 z^{-1} + a_2 z^{-2}}$$

where

$$\begin{aligned} a_1 &= -2R \cos(\theta) \\ a_2 &= R^2, \end{aligned}$$

R is the radius of each pole relative to the z -plane unit circle ($R=1$), and the pole angles are $\pm\theta$. The pole angle can be interpreted as $\theta = \omega_c T$ where ω_c is the frequency and T is the sampling period.

- The phaser will have a notch wherever the phase of the allpass chain is at π (180 degrees). This happens very close to the complex-conjugate pole pair angles.
- The instantaneous frequency response of a phaser created using 4 second-order allpass filters with notch frequencies set at 300, 800, 1000, and 4000 Hz and $R = 0.9, 0.98, 0.8,$ and 0.9 is shown in Fig. 4.




Figure 4: Instantaneous frequency response of a phaser created with 4 second-order allpass filters and notch frequencies set at 300, 800, 1000, and 4000 Hz.

- The depth of the notches can be varied together by changing the feedforward gain parameter g .
- To achieve the time-varying “phasing” effect, the notch frequencies are modulated with a periodic signal. Note that only a single filter coefficient need be changed in each allpass section to accomplish this.

1.3 Doppler Effect

- The Doppler effect occurs when a sound source and listener are moving relative to one another. When the source and listener move toward each other, the sound is perceived to increase in frequency. When the source and listener move away from each other, the sound is heard to decrease in frequency.
- The Doppler effect can be used to enhance the realism of simulated moving sound sources.
- The amount of frequency shift is given by:

$$f_l = f_s \frac{c + v_{ls}}{c - v_{sl}},$$

where f_s is the frequency of the source at rest, f_l is the frequency perceived by the listener, v_{ls} is the speed of the listener toward the source (zero if not moving), v_{sl} is the speed of the source toward the listener (zero is not moving), and c is the sound speed.

- We can simulate time-varying source and listener velocities using a time-varying digital delay line with separate read and write pointers. The write pointer corresponds to the source signal and the read pointer corresponds to the listener. If the source position is moving toward the listener, the write pointer increment should be changed from 1 to $1 + v_{sl}/c$. Likewise, if the listener is moving toward the source, the read pointer increment should be changed from 1 to $1 + v_{ls}/c$.
- Interpolated reads from a delay line (fractional delay lengths) were previously discussed. Interpolated writes are referred to as *de-interpolation*.
- Because Doppler shift is dependent only on the *relative* motion of a source and listener and because the de-interpolation process is generally more complicated to implement than interpolation, it is best to change only the read pointer increment when possible.
- A continuously varying delay can be implemented with a “growth parameter” $g = -v_{ls}/c$ such that the read pointer is incremented by $1 + g$ at each time step.
- Multiple read pointers can simulate multiple moving listeners.
- Multiple write pointers can simulate multiple moving sources.
- It is not possible to use a single delay line to simulate both multiple moving listeners and multiple moving sources. In this case, the number of necessary delay lines is the lesser of the number of listeners or sources.

1.4 Chorus Effect

- The simulation of many apparent similar sound sources occurring nearly in unison from a single, actual source is referred to as a chorus effect.
- The chorus effect is achieved by applying various independent modifications to multiple copies of a given input signal and summing them together to create the output signal.
- Signal modifications can include delay, frequency shift, and/or amplitude changes.
- Time-varying delay lines can simulate changes in delay and frequency shift.
- A single multi-tap interpolating delay line with modulated output tap locations can efficiently and effectively implement the chorus effect.
- The following MSP patch implements a chorus effect. Better results can be achieved with more “copies” of the input signal.

1.5 Pitch Shifting

- The Doppler effect causes a source signal to appear as though it has been pitch shifted. Pitch shifting of an input signal can thus be implemented with time-varying delay lines as described above.
- Because the read pointer of a pitch shifter is incremented at a constant non-integer rate, the read pointer will eventually “catch up to” or “fall back into” the write pointer location. To avoid discontinuity issues caused when the read and write pointers cross, a multiple read pointer cross-fade system can be used.

2 Synthesis ToolKit Introduction:

The *Synthesis ToolKit in C++ (STK)* is a set of open source audio signal processing and algorithmic synthesis classes written in C++. STK was designed to facilitate rapid development of music synthesis and audio processing software, with an emphasis on cross-platform functionality, realtime control, ease of use, and educational example code. STK currently runs with “realtime” support (audio and MIDI) on Linux, Macintosh OS X, and Windows computer platforms.

In this course, we will make use of STK to gain hands-on experience implementing acoustic and audio algorithms.

2.1 Getting Started

- STK provides a set of C++ classes that help you create audio signal processing programs.
- STK makes use of object-oriented programming methodology and hierarchy.
- STK does not specify particular rates (audio and/or control) or scheduling systems. It is up to the programmer to design a framework that best addresses a given processing context.
- STK provides support for either single-sample computation or vectorized computations.
- A short STK tutorial, as well as class documentation, is provided at <http://ccrma.stanford.edu/software/stk/>.

2.2 Installation and Compiling (Macintosh OS-X)

- For this course, we will compile programs that make use of STK classes by linking against a static library (`libstk.a`).
- The latest STK distribution can be downloaded from <http://ccrma.stanford.edu/software/stk/download.html>.
- In this course, it is assumed that the STK distribution is located in the `/Developer/stk-4.4.3/` directory and that a static library has been compiled in the `src` subdirectory. This is the case in the Music Technology Computer Lab (MTCL).
- Under these conditions, we can compile a simple C++ program file `myprogram.cpp` from any directory and link against the STK library as follows (further libraries may be necessary depending on the particular program):

```
g++ -o myprogram -I/Developer/stk-4.4.3/include/ -L/Developer/stk-4.4.3/src/ myprogram.cpp -lstk
```

2.3 Inheritance

Nearly all STK classes inherit from the `Stk` abstract base class, which provides common functionality related to error reporting, sample rate control, and byte swapping. Several other base classes exist that roughly group many of the classes according to function as follows:

- **Generator:** source signal unit generator classes [`Blit`, `BlitSaw`, `BlitSquare`, `Envelope`, `ADSR`, `Asymp`, `Noise`, `Modulate`, `SineWave`, `SingWave`, `Granulate`]
- **Filter:** digital filtering classes [`OneZero`, `OnePole`, `PoleZero`, `TwoZero`, `TwoPole`, `BiQuad`, `Fir`, `Iir`, `FormSwep`, `Delay`, `DelayL`, `DelayA`, `TapDelay`]
- **Function:** input to output function mappings [`BowTable`, `JetTable`, `ReedTable`]
- **Instrmnt:** sound synthesis algorithms, including physical, FM, modal, and particle models
- **Effect:** sound processing effect classes [`Echo`, `Chorus`, `PitShift`, `LentPitShift`, `PRCRev`, `JCRev`, `NRev`]
- **WvOut:** audio file and streaming output classes [`FileWvOut`, `RtWvOut`, `InetWvOut`]
- **WvIn:** audio file and streaming input classes [`FileWvIn`, `FileLoop`, `RtWvIn`, `InetWvIn`]

2.4 A Simple Example

- Audio and control signals throughout STK use a floating-point data type, `StkFloat`, the exact precision of which can be controlled via a typedef statement in the file `Stk.h` (the default is a *double*).
- Most STK classes and instruments expect input/output values in the range ± 1.0 .
- The following program will generate 20 random floating-point (`StkFloat`) values in the range -1.0 to +1.0:

```
#include "Noise.h"
using namespace stk;

int main()
{
    StkFloat output;
    Noise noise;

    for ( unsigned int i=0; i<20; i++ ) {
        output = noise.tick();
        std::cout << "i = " << i << " : output = " << output << std::endl;
    }

    return 0;
}
```

- The `StkFrames` class provides a general "mechanism" for handling and passing vectorized, multi-channel audio data.
- The `StkFrames` class also provides data interpolation functionality.
- The example from above could be performed as a "vectorized" computation in the following way:

```
#include "Noise.h"
using namespace stk;

int main()
{
    StkFrames output(20, 1); // initialize StkFrames to 20 1-channel frames
    Noise noise;

    noise.tick( output ); // perform vectorized computation
    for ( unsigned int i=0; i<output.frames(); i++ ) {
        std::cout << "i = " << i << " : output = " << output[i] << std::endl;
    }

    return 0;
}
```

2.5 "Ticking"

- The sample computations performed by each STK class are generally implemented in a function called `tick()`.
- Most STK unit generators take and/or produce monophonic data. Thus, the `tick()` function normally takes a single sample input value and/or outputs a single sample output value.

- Overloaded versions of the `tick()` function exist to allow vectorized computations.
- In a few cases, STK classes may input or output multi-channel data (ex., `WvIn`, `WvOut`, and their subclasses).
- In a multi-channel context, the `tick()` functions taking or returning `StkFloat` arguments output average values for all channels or input a single value to all channels.
- The `tick()` functions taking `StkFrames` arguments, on the other hand, expect multi-channel input/output data in the form of an `StkFrames` variable.

2.6 Audio File Input/Output

- The classes `FileRead` and `FileWrite` provide input/output support for RAW, WAV, AIFF, SND (AU), and MAT-file formatted soundfiles containing 8-, 16-, and 32-bit integer, as well as 32- and 64-bit floating-point, data.
- The `FileRead` and `FileWrite` classes offer an efficient means for querying audio file parameters because they do not read/write data during instantiation. Further, they never allocate internal memory. Rather, file data is read/written directly to/from an `StkFrames` object from/to the file.
- The `FileWvIn` and `FileWvOut` classes provide a “tick-level” interface to the `FileRead` and `FileWrite` classes.
- `FileWvIn` provides a variable read rate (including negative rates) using linear interpolation.
- The class `FileLoop` is a subclass of `FileWvIn` and provides looping functionality for any of the supported soundfile types mentioned above.

```
// sineosc.cpp STK tutorial program

#include "FileLoop.h"
#include "FileWvOut.h"
using namespace stk;

int main()
{
    // Set the global sample rate before creating class instances.
    Stk::setSampleRate( 44100.0 );

    int nFrames = 100000;
    FileLoop input;
    FileWvOut output;

    try {
        // Load the sine wave file
        input.openFile( "sinewave.raw", true );

        // Open a 16-bit, one-channel WAV formatted output file
        output.openFile( "hellosine.wav", 1, FileWrite::FILE_WAV, Stk::STK_SINT16 );
    }
    catch ( StkError & ) {
        exit(0);
    }

    input.setFrequency( 440.0 );
}
```

```

// Option 1: Use StkFrames
/*
StkFrames frames( nFrames, 1 );
try {
    output.tick( input.tick( frames ) );
}
catch ( StkError & ) {
    exit( 1 );
}
*/

// Option 2: Single-sample computations
for ( int i=0; i<nFrames; i++ ) {
    try {
        output.tick( input.tick() );
    }
    catch ( StkError & ) {
        exit( 1 );
    }
}

return 0;
}

```

- When opening a file, the `FileWvIn` class automatically adjusts the data “read” rate based on the current STK sample rate and the file rate. For example, if the sample rate is currently 44100 Hz and the file rate is 22050, the read rate will be set to 0.5. If you need to set an explicit read rate for a given `FileWvIn` instance, it is safest to change it *after* the file has been opened.
- Since STK release 4.3.1, sinusoidal oscillator functionality is also available using the `SineWave` class. This is generally “safer” to use because it does not depend on an external sine table file (`sinewave.raw`). `SineWave`’s internal table is statically allocated and shared amongst multiple instances of the class.

2.7 STK Filters

- STK provides a variety of classes for digital filtering, all inheriting from the `Filter` abstract base class.
- The class `Iir` provides generalized infinite impulse response filtering functionality similar to that of the filter function in Matlab.
- The class `Fir` provides generalized finite impulse response filtering functionality and is slightly more efficient than using the `Iir` class when only a feedforward structure is needed.
- All filter classes in STK inherit from the `Filter` class, including `Fir`, `Iir`, `OneZero`, `OnePole`, `PoleZero`, `TwoZero`, `TwoPole`, `BiQuad`, `Delay`, ...
- The `BiQuad` class provides functionality for resonance filtering, as used below.

```

#include "FileWvIn.h"
#include "FileWvOut.h"
#include "BiQuad.h"
using namespace stk;

int main( int argc, char *argv[] )
{
    FileWvIn  input;

```

```

FileWvOut output;

if ( argc != 2 ) {
    std::cout << "usage: " << argv[0] << " file" << std::endl;
    std::cout << "    where 'file' is an input soundfile to process" << std::endl;
    exit(0);
}

try {
    // Load the input file
    input.openFile( argv[1] );
}
catch ( StkError & ) {
    exit(0);
}

// Set the global STK sample rate to the input file sample rate.
Stk::setSampleRate( input.getFileRate() );

// Reset the input file reader increment to 1.0
input.setRate( 1.0 );

try {
    // Define and open a 16-bit, one-channel WAV formatted output file
    output.openFile( "filter2", 2, FileWrite::FILE_AIF, Stk::STK_SINT16 );
}
catch ( StkError & ) {
    input.closeFile();
    exit(0);
}

BiQuad filter;
filter.setResonance( 1000.0, 0.999, true ); // set resonance to 1000 Hz, pole radius = 0.999, an
StkFrames frame( 1, 2 ); // one frame of 2 channels

int i;
i = input.getSize(); // in sample frames
while ( i-- >= 0 ) {

    try {
        frame[0] = input.tick();
        frame[1] = filter.tick( frame[0] );
        output.tick( frame );
    }
    catch ( StkError & ) {
        goto cleanup;
    }
}

cleanup:
input.closeFile();
output.closeFile();

return 0;

```

```
}
```

2.8 Realtime Audio and MIDI Input/Output

- The STK-independent class `RtAudio` provides realtime audio input and output support for a variety of computer operating systems and APIs, including OS X (CoreAudio, Jack), Linux (OSS, ALSA, and Jack), and Windows (DirectSound and ASIO).
- The STK-independent class `RtMidi` provides realtime MIDI input and output support for a variety of computer operating systems and APIs, including OS X (CoreMIDI), Linux (ALSA sequencer), and Windows (WinMM).
- The class `RtAudio` provides both blocking and callback functionality for use in a realtime synthesis context. However, the blocking functionality does not perform well on systems for which the underlying audio API is based on callbacks (Macintosh OS-X, Linux Jack, and Windows ASIO). On these systems, it is best to use the callback functionality (as demonstrated below).

```
// rtex1.cpp
//
// Realtime audio output example using callback functionality.

#include "RtAudio.h"
#include <iostream>
#include <cmath>

const double PI      = 3.14159265359;
const double TWO_PI = 2 * PI;
const double frequency = 440.0;
const int sampleRate = 44100;
const double samplePeriod = 1.0 / sampleRate;
const double phaseIncrement = TWO_PI * frequency * samplePeriod;
double phase = 0;

int sin( void *outputBuffer, void *, unsigned int nBufferFrames,
         double, RtAudioStreamStatus, void * )
{
    // Cast the buffer to the correct data type.
    double *my_data = (double *) outputBuffer;

    // We know we only have 1 sample per frame here.
    for ( int i=0; i<nBufferFrames; i++ ) {
        my_data[i] = std::sin( phase );
        phase += phaseIncrement;
        if ( phase > TWO_PI ) phase -= TWO_PI;
    }

    return 0;
}

int main()
{
    unsigned int nBufferFrames = 256; // 256 sample frames
    unsigned int sampleRate = 44100;
    unsigned int nChannels = 1;
    RtAudio dac;
```

```

// Open the default realtime output device.
RtAudio::StreamParameters parameters;
parameters.deviceId = dac.getDefaultOutputDevice();
parameters.nChannels = nChannels;
try {
    dac.openStream( &parameters, NULL, RTAUDIO_FLOAT64, sampleRate, &nBufferFrames, &sin );
}
catch ( RtError &error ) {
    error.printMessage();
    exit( EXIT_FAILURE );
}

try {
    dac.startStream();
}
catch ( RtError &error ) {
    error.printMessage();
    exit( EXIT_FAILURE );
}

char input;
std::cin.get( input );

// Stop the stream.
try {
    dac.stopStream();
}
catch ( RtError &error ) {
    error.printMessage();
}

return 0;
}

```

- This example can be compiled on a Macintosh OS-X system with the following syntax:

```
g++ -I/Developer/stk-4.4.3/include/ -L/Developer/stk-4.4.3/src/ -D__MACOSX_CORE__ rtex1.cpp -lstk -
```

- The previous example can be easily modified to take realtime input from a soundcard, pass it through a comb filter, and send it back out to the soundcard as follows:

```

// rtex2.cpp
//
// Realtime audio input/output example with comb filter using callback functionality.
#include "RtAudio.h"
#include <iostream>
#include "Delay.h"
using namespace stk;

Delay delay;

int comb( void *outputBuffer, void *inputBuffer, unsigned int nBufferFrames,
          double, RtAudioStreamStatus, void * )

```

```

{
    // Cast the buffers to the correct data type.
    double *idata = (double *) inputBuffer;
    double *odata = (double *) outputBuffer;

    // We know we only have 1 sample per frame here.
    for ( int i=0; i<nBufferFrames; i++ ) {
        odata[i] = idata[i] + 0.98 * delay.lastOut(); // feedback comb
//    odata[i] = idata[i] + delay.tick( idata[i] ); // feedforward comb
        odata[i] *= 0.45;
        delay.tick( odata[i] ); // feedback comb
    }

    return 0;
}

int main()
{
    unsigned int nBufferFrames = 256; // 256 sample frames
    unsigned int sampleRate = 44100;
    unsigned int nChannels = 1;
    RtAudio adac;

    delay.setDelay( 1000 );

    // Open the default realtime output device.
    RtAudio::StreamParameters oParameters, iParameters;
    oParameters.deviceId = adac.getDefaultOutputDevice();
    iParameters.deviceId = adac.getDefaultInputDevice();
    oParameters.nChannels = nChannels;
    iParameters.nChannels = nChannels;
    try {
        adac.openStream( &oParameters, &iParameters, RTAUDIO_FLOAT64, sampleRate, &nBufferFrames, &comb
    }
    catch ( RtError &error ) {
        error.printMessage();
        exit( EXIT_FAILURE );
    }

    try {
        adac.startStream();
    }
    catch ( RtError &error ) {
        error.printMessage();
        exit( EXIT_FAILURE );
    }

    char input;
    std::cin.get( input );

    // Stop the stream.
    try {
        adac.stopStream();
    }
}

```

```
    catch ( RtError &error ) {  
        error.printMessage();  
    }  
  
    return 0;  
}
```