

Music 421 Final Project : This is not a comb filter

Jorge Herrera

03/16/2011

1 Introduction

This is not a comb filter is my final project for MUSIC 421 (Audio Applications of the Fast Fourier Transform). It is an image driven real-time audio filtering application for iOS, inspired in Magritte's painting *This is not a pipe*¹.

The user is presented with an image (by default the one showed in Figure 1). When two fingers are placed onto the image, the pixel values along the line defined by the two fingers are used as the sampled magnitude response of a time-varying filter that processes the microphone input in real-time. As the user moves the fingers on the screen, the filter's magnitude response changes.

The user change the "source" image either by browsing the image library on the phone, by browsing the internet, or by taking a picture using the integrated camera.

2 Design

Figure 2 shows a block diagram of the most relevant parts in the system, identifying the different processing steps and the interaction between visual and audio systems.

2.1 Image processing

Since this is an image processing is not the central of the present writeup, I will briefly point out one key aspect.

The main problem to solve in this step is to accurately sample the image to generate enough information to derive the filter's spectral envelope while updating the line channel information

¹see http://en.wikipedia.org/wiki/The_Treachery_of_Images



Figure 1: Application's user interface

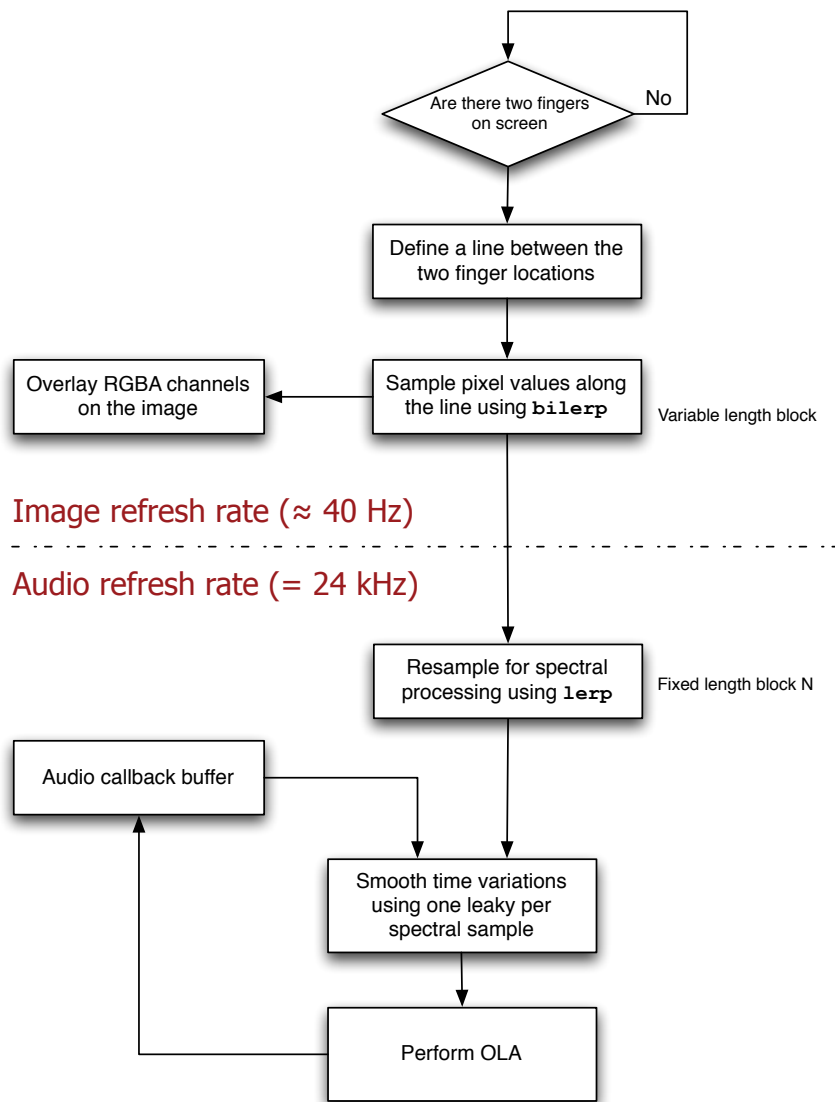


Figure 2: System overview

visually on screen in real-time, providing a coherent audiovisual feedback. Note that visual and audio systems operate at different sampling/refreshing rates. Therefore, they impose different resolution constraints.

I decided to keep a constant “visual resolution” in this image processing step. In other words, to sample the image at intervals of pixel on the screen. This is easy when the line is aligned exactly vertically or horizontally, but it requires interpolation in any other orientation. Bilinear interpolation (`bilerp`) was used to derive these values. While sufficient for visual accuracy, this sampling is not necessarily enough for audio processing purposes. I’ll address this problem later.

It is also important to note that four visual channels (RGBA) are interpolated simultaneously. In the current implementation only one of the application, only one of the channels –the red one– is being used to drive the filter.

2.2 Audio processing

To implement a real-time filtering in the spectral domain a STFT is required. There are at least two options to implement the STFT: 1) Overlap And Add (OLA); 2) Filter Bank Summation (FBS). Although originally I was planning to implement the FBS solution, I decided to switch to an OLA design because of the time-varying nature of the filter.

As mentioned before, the Spectral Envelope derived in the image processing step does not necessarily have enough frequency resolution for the audio processing. In particular, for a N long FFT we need $N/2$ spectral envelope values. That’s why it is necessary to implement a second interpolation of the “visually interpolated” line. A simple linear interpolation (`lerp`) was selected for this purpose.

Other design decisions/restrictions made to keep things simple in a first iteration of the application are:

1. A 50% overlap in this first implementation
2. Parametrized FFT length (currently 4 times the window length)
3. Fixed linear frequency mapping
4. Zero-phase filter assumption
5. Monophonic processing

Most of these restrictions are very easy to modify/relax, to make the application more interesting and increase the audio quality.

3 Implementation

3.1 UI and Image processing

For the image processing, the display of both the image and the 4 channel lines was done in C++/OpenGL. The rest of the visual system was done using the iOS SDK/Objective-C/Interface Builder, with the help of an open source code found at vocaro.com².

Figure 3 presents two examples of the visual feedback provided to the user. It is possible to see 4 colored lines representing the corresponding RGBA channel (alpha channel is represented in white). The ref “profile” is being used to derive the spectral envelope of the filter.

²see <http://vocaro.com/trevor/blog/2009/10/12/resize-a-uiimage-the-right-way/>



Figure 3: Two screenshots of the application, using different images to drive the filter

3.2 Audio processing

The Mobile Music toolkit³ takes care of the low level audio routing, simplifying the interaction between graphics, audio and touchscreen interactions.

The core of the audio processing is the OLA implementation, which was encapsulated in its own class (see Appendix 5.1).

Since the audio callback provides contiguous non-overlapping blocks, the OLA class needs to split the M -samples long input block into two, in order to perform the 50% overlap. The first half is combined with the the second half of the previous audio input block and the processed. The returned values are added according to the OLA procedure and the first half (R samples) of the output audio bock are computed. The second R samples of the output audio block are computed by spectrally processing the whole current input block and performing the add step. Finally, the second half of the input block is saved for the next iteration.

The processing step just mentioned is as follows: First, the block is windowed, zero-padded and its forward FFT is calculated; then, the spectral envelope is smoothed over time using one leaky integrator per spectral bin; the “time-smoothed envelope” multiplies, sample by sample, the FFT of the input signal; finally, the inverse FFT is performed and the output is returned, to bused in the addition step.

The integration between image processing and audio processing is simple. Every time the line changes—every time time the user moves one or two fingers—a new spectral envelope is computed. The visual interpolation of the red channel is resampled get $N/2$ magnitude values. This defines a “target envelope” that drives the leaky integrators mentioned in the previous paragraph.

³see <http://momu.stanford.edu/>

The main parameters used in the application are:

- Sampling frequency (F_s): 24.000 Hz
- Window length (M): 512
- Overlapp length (R): 256
- FFT length (N): 4096 ($4*M$)

Although optimally is required that $F_s \geq 44000$, the selected value of F_s provides a good trade-off between audio quality and real-time performance.

4 Conclusions and Future Work

This first iteration of the application successfully completed the audio/visual integration, providing a real-time, image driven filtering of the microphone input in an iPhone.

Although all the parts work together, many improvements could be made to improve the quality and usability of the application. Some ideas are:

- Implement minimum phase filter (instead of the zero-phase one)
- Combine different channels in interesting ways (e.g. Red channel mapped Left channel and Blue channel mapped to Right channel)
- Change frequency to a perceptual scale (e.g. BARK scale)
- Parametrize overlap % and windows
- Image carrousel to select images

5 Appendix

5.1 Relevant source code

```
1  /*
2   * OLA.h
3   * CombSynth
4   *
5   * Created by Jorge Herrera on 3/15/11.
6   * Copyright 2011 Stanford. All rights reserved.
7   *
8   */
9
10 /*
11
12   Based on SASP, p. 231
13
14 */
15
16 #ifndef __OLA_H__
17 #define __OLA_H__
18
19 #import "mo_fft.h"
20
21 #define RHO 0.5 // leaky integrator constant
22
23
24 // implements 50% OLA using a Hamming window
25 class OLA {
26
27 // public methods
28 public:
29   /*
30   Constructor:
31   Performe some sanity checking, initializes some instance variables and
32   allocates memory for the different arrays
33   */
34   OLA( unsigned int M, unsigned int N );
35
36   /*
37   Destructor:
38   Free up memory used by the instance arrays
39   */
40   ~OLA();
41
42   /*
43   processBlock:
44   Does all the processing related to an stereo input buffer. For now it
45   converts the input into a mono signal (takes only the left channel)
46
47   Arguments:
48   block:      input block, interleaving left and right samples (the same
49               format received from the audio_callback)
50   numFrames:  number of samples in a channel
```

```

51     output:   array of floats to hold the output block associated with
52               the input block
53     */
54     void processBlock( float * block, unsigned int num_frames, float * output );
55
56     /*
57     When the user changes the finger placing in the image, the target spectral
58     envelope changes accordingly
59     */
60     void updateSpectralEnvelope( float * new_envelope, unsigned int num_points );
61
62
63     // private methods
64     private:
65     /*
66     Helper method that takes a monophonic input and returns the corresponding
67     output, after performing the spectral processing. It takes care of the
68     actual OLA step and related book keeping.
69
70     Arguments:
71     input:   N samples long monophonic signal. The first M samples are
72             considered. The rest is a place holder for zero padding
73             first and later it will contain required time-domain samples
74             to be used in the OLA step.
75     output:  M/2 samples long array to return the processed block
76             (M/2 is because of the fix 50% overlap)
77     */
78     void processOverlappedFrame( float * input, float * output );
79
80     /*
81     Helper method to perform the spectral modification on a block
82
83     Takes a N samples long time domain input (the first M samples being actual
84     samples and the rest is a place holder for zero padding).
85
86     Performs the following operations:
87
88     1. Apply the window to the first M samples
89     2. Zeropad the rest of the input signal
90     3. Computes a forward FFT (in place)
91     4. Applies spectral modifications
92     5. Computes the inverse FFT (in place)
93
94     All N returned samples must be used in the OLA step
95
96     */
97     void spectralProcess( float * input );
98
99
100
101     // private attributes
102     private:
103     unsigned int M;           // Window size (in samples)

```

```

104 unsigned int N; // FFT size
105 unsigned int Nover2; // Half FFT size
106 unsigned int R; // Number of overlapping samples (given that I'm
107 // using a fix 50% overlap, R = M/2)
108
109 float * previous_half; // Previous block required to perform "block
110 // decomposition" required for the overlapp
111 float * window; // Actual window used per processed block
112 float * ola_buffer; // Circular buffer to perform the OLA
113 float * current_block; // Used to "construct" the overlapped time domain
114 // block before going into the freq. domain
115
116 unsigned int ola_buffer_head; // Needed to implement the circula buffer
117
118 float * spectral_envelope; // Spectral envelope (derived from the image)
119 // to be imposed onto the audio signal
120 float * last_envelope; // Last envelope (used to smooth out the changes
121 // in the spectral envelope, using a
122 // Leaky-integrator)
123
124 float a1, b0; // leaky integrator constants used to smooth out spectral
125 // envelope changes
126
127 };
128
129
130 #endif // end __OLA_H__

```

```

1 /*
2  * OLA.mm
3  * CombSynth
4  *
5  * Created by Jorge Herrera on 3/15/11.
6  * Copyright 2011 Stanford. All rights reserved.
7  *
8  */
9
10
11 #ifndef __OLA_MM__
12 #define __OLA_MM__
13
14 #import "OLA.h"
15 #import "utils.h"
16
17 OLA::OLA( unsigned int M, unsigned int N ) : M(M), N(N), R(M/2) {
18
19     assert( M && !(M & (M - 1)) ); // fail if M is not a power of 2
20
21     // if N is not a power of two, make it a power of two
22     if (!( N && !(N & (N - 1)) )) {
23         printf("N = %d is not a power of 2\n", N);
24         float l2 = log( N ) / log( 2 ); // compute the log2(N)
25         N = 1 << (int)ceil( l2 );
26         printf("making N = %d\n", N);

```



```

27     this->N = N;
28 }
29
30 Nover2 = N/2;
31
32 int i;
33
34 // compute the hanning window to use
35 window = (float *)malloc( M * sizeof(float) );
36 MoFFT::hamming( window, M );
37 for ( i = 0; i < M; i++) window[i] /= 1.08; // compensate for the 1.08 COLA
38 // constant of the hamming window
39
40 // allocate memory for the different arrays used in the OLA process
41 previous_half = (float *)malloc( R * sizeof(float) );
42 ola_buffer = (float *)malloc( N * sizeof(float) );
43 current_block = (float *)malloc( N * sizeof(float) );
44 ola_buffer_head = 0;
45
46
47 // inititalize necessary arrays
48 for ( i = 0; i < R; i++) previous_half[i] = 0.f;
49 for ( i = 0; i < N; i++) ola_buffer[i] = 0.f;
50
51 // allocate memory for the spectral envelope and related array and
52 // initialize them
53 spectral_envelope = (float *)malloc( Nover2 * sizeof(float) );
54 last_envelope = (float *)malloc( Nover2 * sizeof(float) );
55 for ( i = 0; i < Nover2; i++) {
56     spectral_envelope[i] = last_envelope[i] = 0.f;
57 }
58
59 // initializes spectral envelope leaky integrator constants
60 b0 = RHO;
61 a1 = 1 - RHO;
62
63 }
64
65
66 OLA::~OLA(){
67     SAFE_DELETE( window );
68     SAFE_DELETE( previous_half );
69     SAFE_DELETE( ola_buffer );
70     SAFE_DELETE( current_block );
71     SAFE_DELETE( spectral_envelope );
72     SAFE_DELETE( last_envelope );
73 }
74
75
76
77
78 void OLA::processBlock( float * block, unsigned int num_frames, float * output ) {
79

```

```

80     assert( M == num_frames );
81
82     int i;
83
84     // using the previous half plus the first half of the new input
85     for ( i = 0; i < R; i++) current_block[i] = previous_half[i];
86     // the input array has 2 channels, but we are handling monophonic OLA
87     for ( i = R; i < M; i++) current_block[i] = block[2*(i-R)];
88     this->processOverlappedFrame( current_block, &output[0] );
89
90
91     // using exclusively the new block data
92     // the input array has 2 channels, but we are handling monophonic OLA
93     for ( i = 0; i < M; i++) current_block[i] = block[2*i];
94     this->processOverlappedFrame( current_block, &output[R] );
95
96
97     // store the previous half, to use it with the next input block
98     for ( i = R; i < M; i++) previous_half[i-R] = block[2*i];
99
100 }
101
102
103 void OLA::processOverlappedFrame( float * input, float * output ) {
104
105     // go to the freq. domain, perform spectral modifications and come back to
106     // time domain
107     this->spectralProcess( input );
108
109     int i;
110
111     // update the ola_buffer
112     for ( i = 0; i < N; i++)
113         ola_buffer[(i+ola_buffer_head) % N] += input[i];
114
115     for ( i = 0; i < R; i++) {
116         // we have all the overlaps needed for the "first" R samples in the
117         // ola_buffer, so we output them
118         output[i] = ola_buffer[(i+ola_buffer_head) % N];
119         // zero-out the "first" R samples in the ola_buffer, to be ready for the
120         // next iteration
121         ola_buffer[(i+ola_buffer_head) % N] = 0.f;
122     }
123     // update the ola_buffer read head
124     ola_buffer_head = (ola_buffer_head + R) % N;
125 }
126
127
128 void OLA::spectralProcess( float * input ) {
129     // window the input to the first M samples
130     MoFFT::apply_window( input, window, M );
131
132     int i;

```

```

133
134 // zero-padding
135 for (i = M; i < N; i++) input[i] = 0.f;
136
137 // forward FFT
138 MoFFT::rfft( input, N, true );
139
140
141 // Smooth out the spectral envelope using a Leaky Integrator
142 for (i = 0; i < Nover2; i++)
143     last_envelope[i] = b0*spectral_envelope[i] + a1*last_envelope[i];
144
145
146 // Apply spectral processing
147 for (i = 0; i < Nover2; i++) {
148     // TODO: spectral processing in dB space
149     // spectral processin in dB space
150     // quick & dirty hack: squaring the filter values (they are between 0-1)
151     // to make the differences more dramatic.
152     input[2*i] *= last_envelope[i]*last_envelope[i];
153     input[2*i+1] *= last_envelope[i]*last_envelope[i];
154 }
155
156 // inverse FFT
157 MoFFT::rfft( input, N, false );
158
159 }
160
161
162 void OLA::updateSpectralEnvelope( float * new_envelope, unsigned int num_points ) {
163
164     if (num_points == 0) return;
165
166     // This implementation accepts input envelopes of a different size (not
167     // necessarily Nover2). If the size doesn't match, inptpolation is
168     // performed. For now, a simple linear interpolation is being used. If I
169     // have time I'll try to implement a more elegant interpolation
170     int i;
171     if ( num_points == Nover2 ) {
172         for (i = 0; i < Nover2; i++)
173             spectral_envelope[i] = new_envelope[i];
174     } else {
175         float step = ( (float)num_points - 1 ) / ( (float)Nover2 - 1 );
176         double int_part;
177         double frac_part;
178         double idx = 0.f;
179         i = 0;
180         while (idx <= num_points) {
181             frac_part = modf( idx , &int_part );
182             if ( frac_part == 0 ) {
183                 spectral_envelope[i] = new_envelope[(int)int_part];
184             } else {
185                 spectral_envelope[i] = lerp( new_envelope[(int)int_part],

```

```
186         new_envelope[(int)(int_part + 1)],
187         frac_part );
188     }
189
190     if ( spectral_envelope[i] > 1.f) spectral_envelope[i] = 1.f;
191     else if ( spectral_envelope[i] < 0.f) spectral_envelope[i] = 0.f;
192
193     idx += step;
194     i++;
195 }
196 }
197
198 }
199
200
201
202 #endif // end __OLA_MM__
```