RUTGERS UNIVERSITY
The State University of New Jersey
School of Engineering
Department of Electrical and Computer Engineering

332:348 — Digital Signal Processing Laboratory

DSP Lab Manual

Sophocles J. Orfanidis

Spring 2014
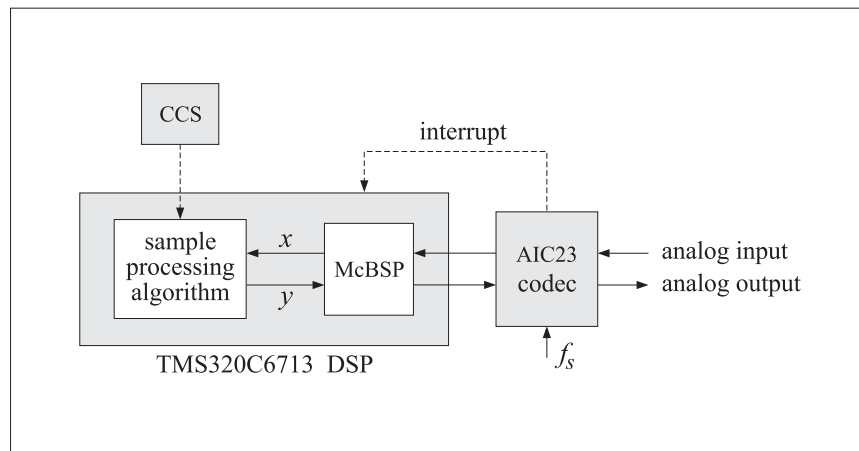
Lab Schedule – Spring 2014

Week	Group	Labs
2/03 2/10	A B	Lab1 – CCS introduction, aliasing, quantization, data transfers, guitar distortion
2/17 2/24	A B	Lab2 – CCS, sinusoids, wavetables, AM/FM, ring modulators, tremolo
3/03 3/10	A B	Lab3 – Delays, circular buffers, FIR filters, voice scrambler
3/17		
3/24	A & B	Lab4 – Sampling, aliasing, FIR filtering (software lab)
4/07 4/14	A B	Lab5 – Digital audio effects, reverb, multi-delay, guitar strings, flangers, vibrato
4/21	A & B	Lab6 – IIR filtering experiments (software lab)

Notes

1. Labs meet in room ELE-209.

2. The lab sessions have a duration of two periods. Attendance in all labs is required (it is not possible to get an "A" in the lab course if one of these sessions is missed.) Due to the limited number of workstations, missed hardware labs cannot be made up.

3. Each lab section has been split into two groups, A & B, that meet on alternate weeks as shown on the above schedule. The groups are as follows, divided according to student last names (please note that these may change until registration is closed):

Section	Group A	Group B
Section–1, Th 10:20 AM – 1:20 PM	Agarwal – Loh	MacPhee – Zieder
Section–2, W 3:20 PM – 6:20 PM	Ali – Patel	Peirce – Young
Section–3, F 8:40 AM – 11:40 AM	Beard - Wei	
TA	S. Sun / T. Wu	T. Wu / S. Sun

Contents

Lab 0 – Introduction

The DSP lab consists of four of hardware experiments illustrating the programming of real-time processing algorithms on the Texas Instruments TMS320C6713 floating-point DSP. Programming of the DSP chip is done in C (and some assembly) using the Code Composer Studio (CCS) integrated development environment. In addition, the lab includes two MATLAB-based software experiments on digital filtering.

Familiarity with C programming is necessary in order to successfully complete this lab course. All of the C filtering functions in the textbook [1] translate with minor changes to the CCS environment. MATLAB is also necessary and will be used in the software experiments and to generate input signals to the DSP and to design the filters used in the various hardware examples.

The hardware experiments are real-time sample-by-sample processing examples and include aliasing and quantization effects; the circular buffer implementation of delays, FIR, and IIR filters; voice scramblers; wavetable generators; and several digital audio effects, such as comb filters, plain, allpass, and lowpass reverberators, Schroeder's reverberator, and several multi-tap, multi-delay, and stereo-delay type effects, tremolo, vibrato, flangers, wah-wah filters and phasers, as well as the Karplus-Strong string algorithm; various guitar distortion effects, such as fuzz and overdrive.

The lab assignments contain a short introduction to the required theory. More details, as well as several concrete C and MATLAB implementations, may be found in the book [1], which may be freely downloaded from the web page:

```
http://www.ece.rutgers.edu/~orfanidi/intro2sp/
```

0.1. Lab Guidelines

Attendance is *required* in all lab sessions (see the lab schedule at the beginning of this manual.) It is not possible to receive a grade of "A" if one of these sessions is missed. Due to the limited number of workstations and tight space, missed hardware labs cannot be made up. In addition, a 1–2 page lab report on each hardware lab must be submitted at the next lab session. A full multi-page reports is required for the software experiments (Lab-4 and Lab-6).

Students work in pairs on each workstation. Each lab section section has been split into two groups, A & B, that meet on alternate weeks (see lab schedule on the lab web page). Please make sure that you attend the right group (if in doubt please contact your TA).

0.2. Running C Programs

Most of the C programs will be written and run under the CCS IDE. However, practicing with and learning C can be done on any departmental computer in ELE-103. Computer accounts on `ece.rutgers.edu` may be obtained by contacting the system administrator of the ECE department, Mr. John Scafidi.

C programs may be compiled using the standard Unix C compiler `cc` or the GNU C compiler `gcc`. Both have the same syntax. It is recommended that C programs be structured in a modular fashion, linking the separate modules together at compilation time. Various versions of GCC, including a Windows version, and an online introduction may be found in the web sites:

```
http://gcc.gnu.org/
http://www.delorie.com/djgpp/
http://www.network-theory.co.uk/docs/gccintro/
```

Some reference books on C are given in Ref. [3]. As an example of using `gcc`, consider the following main program `sines.c`, which generates two noisy sinusoids and saves them (in ASCII format) into the data files `y1.dat` and `y2.dat`:

```
/* sines.c - noisy sinusoids */

#include <stdio.h>
#include <math.h>
```

```
#define L    100
#define f1  0.05
#define f2  0.03
#define A1   5
#define A2   A1

double gran();                          /* gaussian random number generator */

void main()
{
    int n;
    long iseed=2001;                    /* gran requires iseed to be long int */
    double y1, y2, mean = 0.0, sigma = 1.0, pi = 4 * atan(1.0);
    FILE *fp1, *fp2;

    fp1 = fopen("y1.dat", "w");         /* open file y1.dat for write */
    fp2 = fopen("y2.dat", "w");         /* open file y2.dat for write */

    for (n=0; n<L; n++) {               /* iseed is passed by address */
        y1 = A1 * cos(2 * pi * f1 * n) + gran(mean, sigma, &iseed);
        y2 = A2 * cos(2 * pi * f2 * n) + gran(mean, sigma, &iseed);
        fprintf(fp1, "%12.6f\n", y1);
        fprintf(fp2, "%12.6f\n", y2);
        }

    fclose(fp1);
    fclose(fp2);
}
```

The noise is generated by calling the gaussian random number generator routine gauss, which is defined in the separate module gran.c:

```
/* gran.c - gaussian random number generator */

double ran();                            /* uniform generator */

double gran(mean, sigma, iseed)          /* x = gran(mean,sigma,&iseed) */
double mean, sigma;                      /* mean, variance = sigma^2 */
long *iseed;                             /* iseed passed by reference */
{
    double u = 0;
    int i;

    for (i = 0; i < 12; i++)             /* add 12 uniform random numbers */
        u += ran(iseed);

    return sigma * (u - 6) + mean;       /* adjust mean and variance */
}
```

In turn, gran calls a uniform random number generator routine, which is defined in the file ran.c:

```
/* ran.c - uniform random number generator in [0, 1) */

#define  a    16807                               /* a = 7^5 */
#define  m    2147483647                          /* m = 2^31 - 1 */
#define  q    127773                              /* q = m / a = quotient */
#define  r    2836                                /* r = m % a = remainder */
```

```
double ran(iseed)                                     /* usage: u = ran(&iseed); */
long *iseed;                                           /* iseed passed by address */
{
   *iseed = a * (*iseed % q) - r * (*iseed / q);          /* update seed */

   if (*iseed < 0)                                    /* wrap to positive values */
         *iseed += m;

   return (double) *iseed / (double) m;
}
```

The three programs can be compiled and linked into an executable file by the following command-line
call of gcc:

```
gcc sines.c gran.c ran.c -o sines -lm            (unix version of gcc)
gcc sines.c gran.c ran.c -o sines.exe -lm        (MS-DOS version of gcc)
```

The command-line option -lm links the math library and must always be last. The option -o creates the
executable file sines (or, sines.exe for MS-DOS.) If this option is omitted, the executable filename is
a.out (or, a.exe) by default. Another useful option is the warning message option -Wall:

```
gcc -Wall sines.c gran.c ran.c -o sines -lm
```

   If the command line is too long and tedious to type repeatedly, one can use a so-called response file,
which may contain all or some of the command-line arguments. For example, suppose the file argfile
contains the lines:

```
-Wall
sines.c
gran.c
ran.c
-o sines
-lm
```

Then, the following command will have the same effect as before, where the name of the response file
must be preceded by the at-sign character @:

```
gcc @argfile
```

To compile only, without linking and creating an executable, we can use the command-line option -c:

```
gcc -c sines.c gran.c ran.c
```

This creates the object-code modules *.o, which can be subsequently linked into an executable as follows:

```
gcc -o sines sines.o gran.o ran.o -lm
```

0.3. Using MATLAB

The plotting of data created by C or MATLAB programs can be done using MATLAB's extensive plotting
facilities. Here, we present some examples showing how to load and plot data from data files, how to
adjust axis ranges and tick marks, how to add labels, titles, legends, and change the default fonts, how
to add several curves on the same graph, and how to create subplots.

   Suppose, for example, that you wish to plot the noisy sinusoidal data in the files y1.dat and y2.dat
created by running the C program sines. The following MATLAB code fragment will load and plot the
data files:

```
load y1.dat;                            % load data into vector y1
load y2.dat;                            % load data into vector y2

plot(y1);                               % plot y1
hold on;                                % add next plot
plot(y2, 'r--');                        % plot y2 in red dashed style

axis([0, 100, -10, 10]);                % redefine axes limits
set(gca, 'ytick', -10:5:10);            % redefine yticks
legend('y1.dat', 'y2.dat');             % add legends
xlabel('time samples');                 % add labels and title
ylabel('amplitude');
title('Noisy Sinusoids');
```
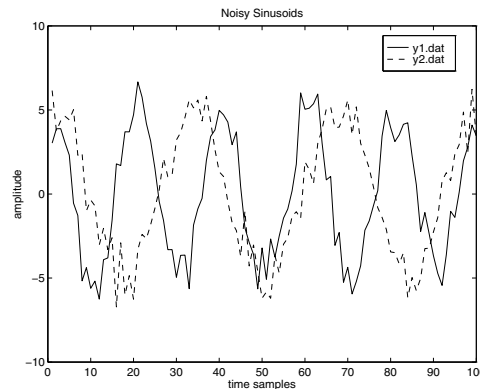
The resulting plot is shown below. Note that the command `load y1.dat` strips off the extension part of the filename and assigns the data to a vector named `y1`.

The command `hold on` leaves the first plot on and adds the second plot. The axis command increases the y-range in order to make space for the legends. The legends, labels, and title are in the default font and default size (e.g., Helvetica, size 10 for the Windows version.)

A more flexible and formatted way of reading and writing data from/to data files is by means of the commands `fscanf` and `fprintf`, in conjunction with `fopen` and `fclose`. They have similar usage as in C. See Ref. [2] for more details.

The next example is similar to what is needed in Lab-1. The example code below generates two signals $x(t)$ and $y(t)$ and plots them versus t. It also generates the time-samples $y(t_n)$ at the time instants $t_n = nT$. All three signals $x(t)$, $y(t)$, $y(t_n)$ span the same total time interval $[0, t_{\max}]$, but they are represented by arrays of different dimension ($x(t)$ and $y(t)$ have length 101, whereas $y(t_n)$ has length 11). All three can be placed on the same graph as follows:

```
tmax = 1;                               % max time interval
Nmax = 100;                             % number of time instants
Dt = tmax/Nmax;                         % continuous-time increment
T = 0.1;                                % sampling time interval

t  = 0:Dt:tmax;                               % continuous t
x = sin(4*pi*t) + sin(16*pi*t) + 0.5 * sin(24*pi*t);   % signal x(t)
y = 0.5 * sin(4*pi*t);                        % signal y(t)

tn = 0:T:tmax;                          % sampled version of t
yn = 0.5 * sin(4*pi*tn);                % sampled version of y(t)

plot(t, x, t, y, '--', tn, yn, 'o');    % plot x(t), y(t), y(tn)

axis([0, 1, -2, 2])                     % redefine axis limits
```

```
set(gca, 'xtick', 0:0.1:1);              % redefine x-tick locations
set(gca, 'ytick', -2:1:2);               % redefine y-tick locations
set(gca, 'fontname', 'times');           % Times font
set(gca, 'fontsize', 16);                % 16-point font size
grid;                                    % default grid

xlabel('t (sec)');
ylabel('amplitude');
title('x(t), y(t), y(tn)');

axes(legend('original', 'aliased', 'sampled'));   % legend over grid
```

The following figure shows the results of the above commands. Note that the *x*-axis tick marks have been redefined to coincide with the sampled time instants $t_n = nT$.
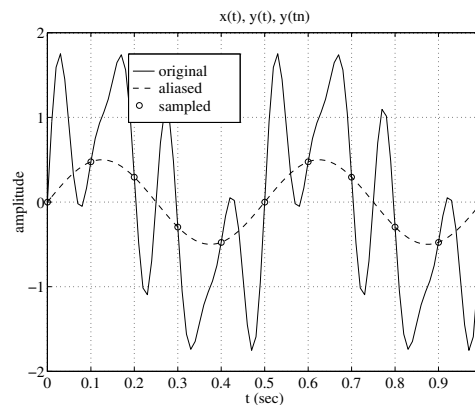
The 'o' command plots the sampled signal $y(t_n)$ as circles. Without the 'o', the plot command would interpolate linearly between the 11 points of $y(t_n)$.

The font has been changed to Times-Roman, size 16, in order to make it more visible when the graph is scaled down for inclusion in this manual. The command `axes` creates a new set of axes containing the legends and superimposes them over the original grid (otherwise, the grid would be visible through the legends box.)

The next program segment shows the use of the command `subplot`, which is useful for arranging several graphs on one page. It also illustrates the `stem` command, which is useful for plotting sampled signals.

```
subplot(2, 2, 1);                        % upper left subplot

plot(t, x, t, y, '--', tn, yn, 'o');     % plot x(t), y(t), y(tn)

xlabel('t (sec)');
ylabel('amplitude');
title('x(t), y(t), y(tn)');

subplot(2, 2, 2);                        % upper right subplot

plot(t, y);                              % plot y(t)
hold on;                                 % add next plot
stem(tn, yn);                            % stem plot of y(tn)

axis([0, 1, -0.75, 0.75]);               % redefine axis limits

xlabel('t (sec)');
ylabel('y(t), y(tn)');
title('stem plot');
```
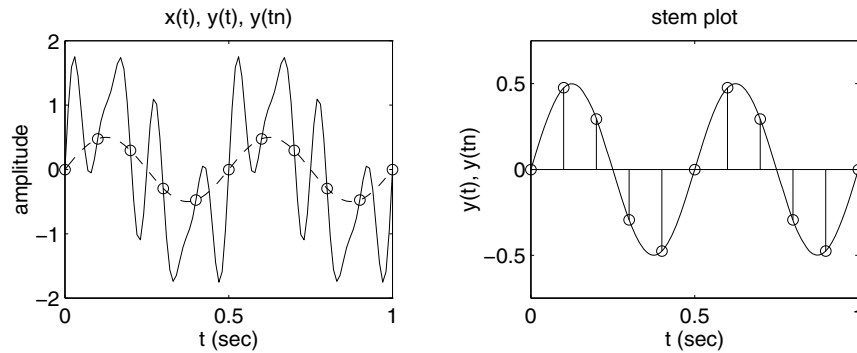
The resulting graph is shown below. Note that a 2×2 subplot pattern was used instead of a 1×2, in order to get a more natural aspect ratio.

Finally, we mention some MATLAB resources. Many of the MATLAB functions needed in the experiments are included in Appendix D of the text [1]. Many MATLAB on-line tutorials can be found at the following web sites:

```
http://www.mathworks.com/academia/student_center/tutorials/index.html
http://www.eece.maine.edu/mm/matweb.html
```

0.4. References

[1] S. J. Orfanidis, *Introduction to Signal Processing*, online book, 2010, available from:
     `http://www.ece.rutgers.edu/~orfanidi/intro2sp/`

[2] MATLAB Documentation: `http://www.mathworks.com/help/techdoc/`

[3] B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, 2nd ed., Prentice Hall, Englewood Cliffs, NJ, 1988.

     S. P. Harbison and G. L. Steele, *C: A Reference Manual*, Prentice Hall, Englewood Cliffs, NJ, 1984.

     A. Kelly and I. Pohl, *A Book on C*, 2nd ed., Benjamin/Cummings, Redwood City, CA, 1990.

Lab 1 – TMS320C6713 DSK and Code Composer Studio

1.1. Introduction

The hardware experiments in the DSP lab are carried out on the Texas Instruments TMS320C6713 DSP Starter Kit (DSK), based on the TMS320C6713 floating point DSP running at 225 MHz. The basic clock cycle instruction time is $1/(225\text{ MHz}) = 4.44$ nanoseconds. During each clock cycle, up to eight instructions can be carried out in parallel, achieving up to $8\times225 = 1800$ million instructions per second (MIPS).

The C6713 processor has 256KB of internal memory, and can potentially address 4GB of external memory. The DSK board includes a 16MB SDRAM memory and a 512KB Flash ROM. It has an on-board 16-bit audio stereo codec (the Texas Instruments AIC23B) that serves both as an A/D and a D/A converter. There are four 3.5 mm audio jacks for microphone and stereo line input, and speaker and head-phone outputs. The AIC23 codec can be programmed to sample audio inputs at the following sampling rates:

$$f_s = 8,\ 16,\ 24,\ 32,\ 44.1,\ 48,\ 96\ \text{kHz}$$

The ADC part of the codec is implemented as a multi-bit third-order noise-shaping delta-sigma converter (see Ch. 2 & 12 of [1] for the theory of such converters) that allows a variety of oversampling ratios that can realize the above choices of f_s. The corresponding oversampling decimation filters act as anti-aliasing prefilters that limit the spectrum of the input analog signals effectively to the Nyquist interval $[-f_s/2, f_s/2]$. The DAC part is similarly implemented as a multi-bit second-order noise-shaping delta-sigma converter whose oversampling interpolation filters act as almost ideal reconstruction filters with the Nyquist interval as their passband.

The DSK also has four user-programmable DIP switches and four LEDs that can be used to control and monitor programs running on the DSP.

All features of the DSK are managed by the CCS, which is a complete integrated development environment (IDE) that includes an optimizing C/C++ compiler, assembler, linker, debugger, and program loader. The CCS communicates with the DSK via a USB connection to a PC. In addition to facilitating all programming aspects of the C6713 DSP, the CCS can also read signals stored on the DSP's memory, or the SDRAM, and plot them in the time or frequency domains.

The following block diagram depicts the overall operations involved in all of the hardware experiments in the DSP lab. Processing is interrupt-driven at the sampling rate f_s, as explained below.
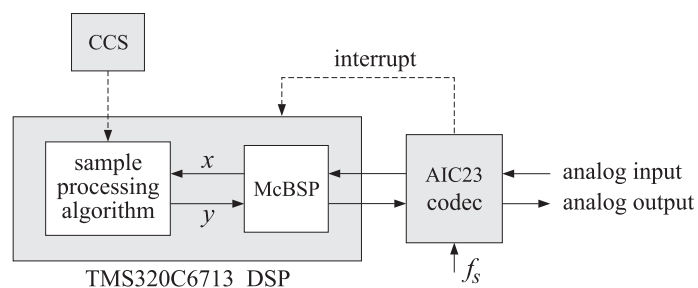
The AIC23 codec is configured (through CCS) to operate at one of the above sampling rates f_s. Each collected sample is converted to a 16-bit two's complement integer (a **short** data type in C). The codec actually samples the audio input in stereo, that is, it collects two samples for the left and right channels.

At each sampling instant, the codec combines the two 16-bit left/right samples into a single 32-bit unsigned integer word (an **unsigned int**, or **Uint32** data type in C), and ships it over to a 32-bit receive-register of the multichannel buffered serial port (McBSP) of the C6713 processor, and then issues an interrupt to the processor.

Upon receiving the interrupt, the processor executes an interrupt service routine (ISR) that implements a desired sample processing algorithm programmed with the CCS (e.g., filtering, audio effects, etc.). During the ISR, the following actions take place: the 32-bit input sample (denoted by x in the diagram) is read from the McBSP, and sent into the sample processing algorithm that computes the corresponding

32-bit output word (denoted by y), which is then written back into a 32-bit transmit-register of the McBSP, from where it is transferred to the codec and reconstructed into analog format, and finally the ISR returns from interrupt, and the processor begins waiting for the next interrupt, which will come at the next sampling instant.

Clearly, all processing operations during the execution of the ISR must be completed in the time interval between samples, that is, $T = 1/f_s$. For example, if $f_s = 44.1$ kHz, then, $T = 1/f_s = 22.68$ μsec. With an instruction cycle time of $T_c = 4.44$ nsec, this allows $T/T_c = 5108$ cycles to be executed during each sampling instant, or, up to $8\times5108 = 40864$ instructions, or half of that per channel.

Resources

Most of the hardware experiments in the DSP lab are based on C code from the text [1] adapted to the CCS development environment. Additional experiments are based on the Chassaing-Reay text [2].

The web page of the lab, `http://www.ece.rutgers.edu/~orfanidi/ece348/`, contains additional resources such as tutorials and user guides. Some books on C and links to the GNU GCC C compiler are given in Ref. [5].

As a prelab, before you attend Lab-1, please go through the powerpoint presentations of Brown's workshop tutorial in Ref. [3], Part-1, and Dahnoun's chapters 1 & 3 listed in Ref. [4]. These will give you a pretty good idea of the TMS320C6000 architecture and features.

The help file, `C:\CCStudio_v3.1\docs\hlp\c6713dsk.hlp`, found in the CCS installation directory of each PC, contains very useful information on the C6713 processor and DSK kit. The following pictures are from that help file:
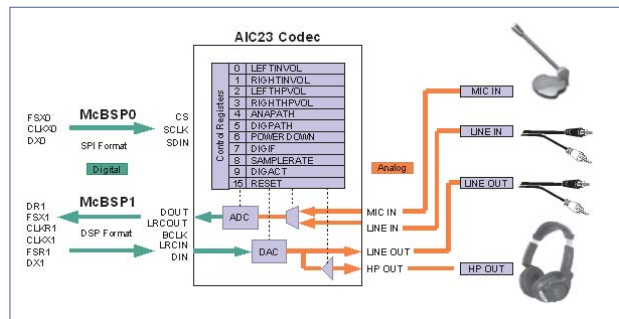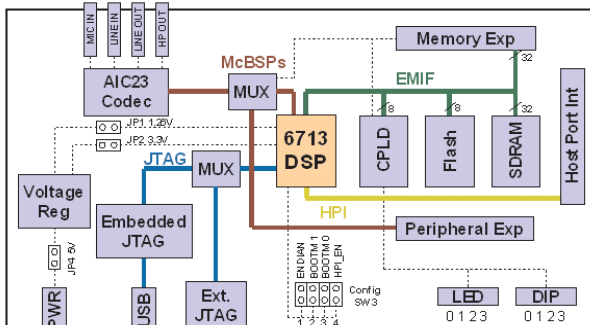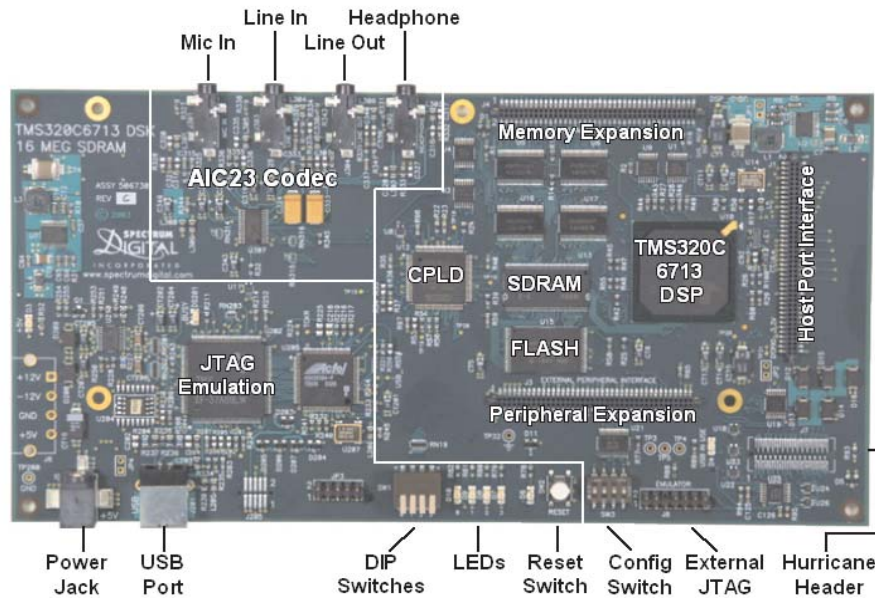
1.2. Lab Tasks

In this lab, you will learn how to use some basic features of the Code Composer Studio (CCS), such as creating projects, compiling and linking them to the run-time libraries, loading them for execution on the DSP chip, using GEL files for changing program parameters during run-time.

You will hear what aliasing effects sound like (i.e., distortions arising from using the wrong sampling rate). You will hear what quantization effects sound like (i.e., when you use too few bits for your audio samples). You will find out how the stereo A/D converter packs the two 16-bit samples from the left and right audio channels into a 32-bit word and sends it over to the processor, and how it gets unpacked into the two individual 16-bit left/right words by the processor. You will also study panning between speakers, and several nonlinear input/output functions such as fuzz (hard clipping) and tube amplifier (soft clipping) for guitar distortion.

1.3. Template Program

You will begin with a basic talkthrough program, listed below, that simply reads input samples from the codec and immediately writes them back out. This will serve as a template on which to build more complicated sample processing algorithms by modifying the interrupt service routine `isr()`.

```c
// template.c - to be used as starting point for interrupt-based programs
// ------------------------------------------------------------------------------

#include "dsplab.h"          // DSK initialization declarations and function prototypes

short xL, xR, yL, yR;        // left and right input and output samples from/to codec
float g=1;                   // gain to demonstrate watch windows and GEL files

// here, add more global variable declarations, #define's, #include's, etc.
// ------------------------------------------------------------------------------

void main()                  // main program executed first
{
  initialize();              // initialize DSK board and codec, define interrupts

  sampling_rate(8);          // possible sampling rates: 8, 16, 24, 32, 44, 48, 96 kHz
  audio_source(LINE);        // LINE or MIC for line or microphone input

  while(1);                  // keep waiting for interrupt, then jump to isr()
}

// ------------------------------------------------------------------------------

interrupt void isr()         // sample processing algorithm - interrupt service routine
{
   read_inputs(&xL, &xR);    // read left and right input samples from codec

   yL = g * xL;              // replace these with your sample processing algorithm
   yR = g * xR;

   write_outputs(yL,yR);     // write left and right output samples to codec

   return;
}

// ------------------------------------------------------------------------------
// here, add more functions to be called within isr() or main()
```

The template has three sections. In the top section, global variables are declared and defined, such as the left/right input/output audio samples x_L, x_R, y_L, y_R, whose scope is the entire file and are known to all functions in the file. Additional `#define` and `#include` statements, such as `#include <math.h>`, and additional global variable declarations may be added in this section.

The second section consists of the function `main()`, which is executed first, and performs the initialization of the DSK board, sets the sampling rate, selects the audio input, and then goes into an infinite loop waiting for an interrupt. Upon receiving the interrupt, it jumps to the function `isr()`. Additional local variables and other preliminary operations, such as the zeroing of delay-line buffers, may be added in this section before the `wait(1)` statement.

The third section consists of the interrupt service routine `isr()`, which implements the desired sample processing algorithm. Note that the keyword **interrupt** has been added to the C language implementation of the CCS. In the template file, the ISR function reads the left/right input samples, process them by multiplying them by a gain, sends them to the output, and returns back to `main()`.

The reading and writing of the input and output samples are done with the help of the functions `read_inputs()` and `write_outputs()`, which are declared in the header file `dsplab.h` and defined in `dsplab.c`. These two files must always be included in your programs and reside in the common directory `C:\dsplab\common\`.

Besides the above three basic sections, other sections may be added that define additional functions to be called within `isr()` or `main()`.

Working with CCS

For each application to be run on the C6713 processor, one must create a "project" in the Code Composer Studio, which puts together all the information about the required C source files, header files, and C libraries, including all the compiler and linker build options. During the lab session, you will be working in the temporary folders:

```
C:\labuser
C:\labuser\dsplab
C:\labuser\dsplab\template
C:\labuser\dsplab\examples
```

Before you start work, please double-click on the desktop icon called `cleanup`. This refreshes the above `labuser` directory and removes all files created by previous students. Consequently, before you leave the lab, you must save your work files in a thumb-drive or email them to yourselves.

To save you time, the project file, `template.pjt`, for the above template has already been created, and may be simply edited for all other projects. To proceed, copy the following three files from the template directory `C:\labuser\dsplab\template`:

```
template.c
template.pjt
template.gel
```

into your temporary working directory, e.g., `C:\labuser`, and double-click the project file, `template.pjt`, which will open in an ordinary text editor. The first few lines of that file are shown below:

```
[Project Settings]
ProjectDir="C:\dsplab\template\"
ProjectType=Executable
CPUFamily=TMS320C67XX
Tool="Compiler"
Tool="CustomBuilder"
Tool="DspBiosBuilder"
Tool="Linker"
Config="Debug"
Config="Release"

[Source Files]
Source="C:\CCStudio_v3.1\C6000\cgtools\lib\rts6700.lib"
Source="C:\CCStudio_v3.1\C6000\csl\lib\csl6713.lib"
Source="C:\CCStudio_v3.1\C6000\dsk6713\lib\dsk6713bsl.lib"
Source="C:\dsplab\common\dsplab.c"
Source="C:\dsplab\common\vectors.asm"
Source="template.c"
```

Only the second and bottom lines in the above listing need to be edited. First, edit the project directory entry to your working directory, e.g.,

```
ProjectDir="C:\labuser"
```

Alternatively, you may delete that line—it will be recreated by CCS when you load the project. Then, edit the source-file line `Source="template.c"` to your new project's name, e.g.,

```
Source="new_project.c"
```

Finally, rename the three files with your new names, e.g.,

```
new_project.c
new_project.pjt
new_project.gel
```

Next, turn on the DSK kit and after the initialization beep, open the CCS application by double-clicking on the CCS desktop icon. Immediately after it opens, use the keyboard combination "ALT+C" (or the menu item *Debug -> Connect*) to connect it to the processor. Then, with the menu item *Project -> Open* or the key combination "ALT+P O", open the newly created project file by navigating to the project's directory, e.g., `C:\labuser`. Once the project loads, you may edit the C source file to implement your algorithm. Additional C source files can be added to your project by the keyboard combination "ALT+P A" or the menu choices *Project -> Add Files to Project*.

Set up CCS to automatically load the program after building it, with the menu commands: *Option -> Customize -> Program/Project Load -> Load Program After Build*. The following key combinations or menu items allow you to compile and load your program, run or halt the program:

```
compile & load:  F7,        Project -> Build
run program:     F5,        Debug -> Run
halt program:    Shift+F5,  Debug -> Halt
```

It is possible that the first time you try to build your program you will get a warning message:

```
warning: creating .stack section with default size of 400 (hex) words
```

In such case, simply rebuild the project, or, in the menu item *Project -> Build Options -> Linker*, enter a value such as `0x500` in the stack entry.

When you are done, please remember to save and close your project with the keyboard combinations "ALT+P S" and "ALT+P C", and save your programs in your account on ECE.

Lab Procedure

a. Copy the template files into your temporary working directory, edit the project's directory as described above, and build the project in CCS. Connect your MP3 player to the line input of the DSK board and play your favorite song, or, you may play one of the wave files in the directory: `c:\dsplab\wav`.

b. Review the template project's build options using the menu commands: *Project -> Build Options*. In particular, review the Basic, Advanced, and Preprocessor options for the Compiler, and note that the optimization level was set to `none`. In future experiments, this may be changed to `-o2` or `-o3`.

For the Linker options, review the Basic and Advanced settings. In particular, note that the default output name `a.out` can be changed to anything else. Note also the library include paths and that the standard included libraries are:

```
rts6700.lib     (run-time library),     C:\CCStudio_v3.1\C6000\cgtools\lib\rts6700.lib
csl6713.lib     (chip support library), C:\CCStudio_v3.1\C6000\csl\lib\csl6713.lib
dsk6713bsl.lib  (board support library), C:\CCStudio_v3.1\C6000\dsk6713\lib\dsk6713bsl.lib
```

The run-time library must always be included. The board support library (BSL) contains functions for managing the DSK board peripherals, such as the codec. The chip support library (CSL) has functions for managing the DSP chip's features, such as reading and writing data to the chip's McBSP. The user manuals for these may be found on the TI web site listed on the lab's web page.

c. The gain parameter g can be controlled in real-time in two ways: using a watch window, or using a GEL file. Open a watch window using the menu item: *View -> Watch Window*, then choose *View -> Quick Watch* and enter the variable g and add it to the opened watch window using the item *Add to Watch*. Run the program and click on the g variable in the watch window and enter a new value, such as $g = 0.5$ or $g = 2$, and you will hear the difference in the volume of the output.

d. Close the watch window and open the GEL file, `template.gel`, with the menu *File -> Load GEL*. In the *GEL* menu of CCS a new item called "gain" has appeared. Choose it to open the gain slider. Run the program and move the slider to different positions. Actually, the slider does not represent the gain g itself, but rather the integer increment steps. The gain g changes by $1/10$ at each step. Open the GEL file to see how it is structured. You may use that as a template for other cases.

e. Modify the template program so that the output pans between the left and right speakers every 2 seconds, i.e., the left speaker plays for 2 sec, and then switches to the right speaker for another 2 sec, and so on. There are many ways of doing this, for example, you may replace your ISR function by

```
#define D 16000      // represents 2 sec at fs = 8 kHz
short d=0;            // move these before main()

interrupt void isr()
{
   read_inputs(&xL, &xR);

   yL = (d < D)  * xL;
   yR = (d >= D) * xR;

   if (++d >= 2*D) d=0;

   write_outputs(yL,yR);

   return;
}
```

Rebuild your program with these changes and play a song. In your lab write-up explain why and how this code works.

1.4. Aliasing

This part demonstrates aliasing effects. The smallest sampling rate that can be defined is 8 kHz with a Nyquist interval of $[-4, 4]$ kHz. Thus, if a sinusoidal signal is generated (e.g. with MATLAB) with frequency outside this interval, e.g., $f = 5$ kHz, and played into the line-input of the DSK, one might expect that it would be aliased with $f_a = f - f_s = 5 - 8 = -3$ kHz. However, this will not work because the antialiasing oversampling decimation filters of the codec filter out any such out-of-band components before they are sent to the processor.

An alternative is to decimate the signal by a factor of 2, i.e., dropping every other sample. If the codec sampling rate is set to 8 kHz and every other sample is dropped, the effective sampling rate will be 4 kHz, with a Nyquist interval of $[-2, 2]$ kHz. A sinusoid whose frequency is outside the decimated Nyquist interval $[-2, 2]$ kHz, but inside the true Nyquist interval $[-4, 4]$ kHz, will not be cut off by the antialiasing filter and will be aliased. For example, if $f = 3$ kHz, the decimated sinusoid will be aliased with $f_a = 3 - 4 = -1$ kHz.

Lab Procedure

Copy the template programs to your working directory. Set the sampling rate to 8 kHz and select line-input. Modify the template program to output every other sample, with zero values in-between. This can be accomplished in different ways, but a simple one is to define a "sampling pulse" periodic signal whose values alternate between 1 and 0, i.e., the sequence $[1, 0, 1, 0, 1, 0, \dots]$ and multiply the input samples by that sequence. The following simple code segment implements this idea:

```
yL = pulse * xL;
yR = pulse * xR;

pulse = (pulse==0);
```

where `pulse` must be globally initialized to 1 before `main()` and `isr()`. Why does this work? Next, rebuild the new program with CCS.

Open MATLAB and generate three sinusoids of frequencies $f_1 = 1$ kHz, $f_2 = 3$ kHz, and $f_3 = 1$ kHz, each of duration of 1 second, and concatenate them to form a 3-second signal. Then play this out of the PCs sound card using the `sound()` function. For example, the following MATLAB code will do this:

```
fs = 8000; f1 = 1000; f2 = 3000; f3 = 1000;
L = 8000; n = (0:L-1);
A = 1/5;                           % adjust playback volume

x1 = A * cos(2*pi*n*f1/fs);
x2 = A * cos(2*pi*n*f2/fs);
x3 = A * cos(2*pi*n*f3/fs);

sound([x1,x2,x3], fs);
```

a. Connect the sound card's audio output to the line-input of the DSK and rebuild/run the CCS down-sampling program after commenting out the line:

```
pulse = (pulse==0);
```

This disables the downsampling operation. Send the above concatenated sinusoids to the DSK input and you should hear three distinct 1-sec segments, with the middle one having a higher frequency.

b. Next, uncomment the above line so that downsampling takes place and rebuild/run the program. Send the concatenated sinusoids to the DSK and you should hear all three segments as though they have the same frequency (because the middle 3 kHz one is aliased with other ones at 1 kHz). You may also play your favorite song to hear the aliasing distortions, e.g., out of tune vocals.

c. Set the codec sampling rate to 44 kHz and repeat the previous two steps. What do you expect to hear in this case?

d. To confirm the antialiasing prefiltering action of the codec, replace the first two lines of the above MATLAB code by the following two:

```
fs = 16000; f1 = 1000; f2 = 5000; f3 = 1000;
L = 16000; n = (0:L-1);
```

Now, the middle sinusoid has frequency of 5 kHz and it should be cutoff by the antialiasing prefilter. Set the sampling rate to 8 kHz, turn off the downsampling operation, rebuild and run your program, and send this signal through the DSK, and describe what you hear.

1.5. Quantization

The DSK's codec is a 16-bit ADC/DAC with each sample represented by a two's complement integer. Given the 16-bit representation of a sample, $[b_1 b_2 \cdots b_{16}]$, the corresponding 16-bit integer is given by

$$x = (-b_1 2^{-1} + b_2 2^{-2} + b_3 2^{-3} + \cdots + b_{16} 2^{-16}) 2^{16} \tag{1.1}$$

The MSB bit b_1 is the sign bit. The range of representable integers is: $-32768 \le x \le 32767$. As discussed in Ch. 2 of Ref. [1], for high-fidelity audio at least 16 bits are required to match the dynamic range of human hearing; for speech, 8 bits are sufficient. If the audio or speech samples are quantized to less than 8 bits, quantization noise will become audible.

The 16-bit samples can be requantized to fewer bits by a right/left bit-shifting operation. For example, right shifting by 3 bits will knock out the last 3 bits, then left shifting by 3 bits will result in a 16-bit number whose last three bits are zero, that is, a 13-bit integer. These operations are illustrated below:

$$[b_1, b_2, \ldots, b_{13}, b_{14}, b_{15}, b_{16}] \quad \Rightarrow \quad [0, 0, 0, b_1, b_2, \ldots, b_{13}] \quad \Rightarrow \quad [b_1, b_2, \ldots, b_{13}, 0, 0, 0]$$

Lab Procedure

a. Modify the basic template program so that the output samples are requantized to B bits, where $1 \le B \le 16$. This requires right/left shifting by $L = 16 - B$ bits, and can be implemented very simply in C as follows:

```
yL = (xL >> L) << L;
yR = (xR >> L) << L;
```

Start with $B = 16$, set the sampling rate to 8 kHz, and rebuild/run the program. Send a wave file as input and listen to the output.

b. Repeat with the following values: $B = 8, 6, 4, 2, 1$, and listen to the gradual increase in the quantization noise.

1.6. Data Transfers from/to Codec

We mentioned in the introduction that the codec samples the input in stereo, combines the two 16-bit left/right samples x_L, x_R into a single 32-bit unsigned integer word, and ships it over to a 32-bit receive-register of the McBSP of the C6713 processor. This is illustrated below.
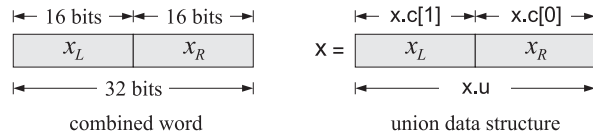
The packing and unpacking of the two 16-bit words into a 32-bit word is accomplished with the help of a union data structure (see Refs. [2,3]) defined as follows:

```
union {                 // union structure to facilitate 32-bit data transfers
    Uint32 u;           // both channels packed as codec.u = 32-bits
    short c[2];         // left-channel = codec.c[1], right-channel = codec.c[0]
} codec;
```

The two members of the data structure share a common 32-bit memory storage. The member `codec.u` contains the 32-bit word whose upper 16 bits represent the left sample, and its lower 16 bits, the right sample. The two-dimensional short array member `codec.c` holds the 16-bit right-channel sample in its first component, and the left-channel sample in its second, that is, we have:

```
    xL = codec.c[1];
    xR = codec.c[0];
```

The functions `read_inputs()` and `write_outputs()`, which are defined in the common file `dsplab.c`, use this structure in making calls to low-level McBSP read/write functions of the chip support library. They are defined as follows:

```
// ------------------------------------------------------------------------------

void read_inputs(short *xL, short *xR)              // read left/right channels
{
   codec.u = MCBSP_read(DSK6713_AIC23_DATAHANDLE);   // read 32-bit word

   *xL = codec.c[1];                               // unpack the two 16-bit parts
   *xR = codec.c[0];
}

// ------------------------------------------------------------------------------

// ------------------------------------------------------------------------------

void write_outputs(short yL, short yR)              // write left/right channels
{
   codec.c[1] = yL;                                // pack the two 16-bit parts
   codec.c[0] = yR;                                // into 32-bit word

   MCBSP_write(DSK6713_AIC23_DATAHANDLE,codec.u);   // output left/right samples
}

// ------------------------------------------------------------------------------
```

Lab Procedure

The purpose of this lab is to clarify the nature of the union data structure. Copy the template files into your working directory, rename them `unions.*`, and edit the project file by keeping in the source-files section only the run-time library and the main function below.

```
// unions.c - test union structure

#include <stdio.h>
#include <stdlib.h>
#include <math.h>

void main(void)
{
   unsigned int v;
   short xL,xR;

   union {
      unsigned int u;
      short c[2];
   } x;

   xL = 0x1234;
   xR = 0x5678;
   v  = 0x12345678;

   printf("\n%x  %x  %d  %d\n", xL,xR, xL,xR);

   x.c[1] = xL;
   x.c[0] = xR;
   printf("\n%x  %x  %x  %d  %d\n", x.u, x.c[1], x.c[0], x.c[1], x.c[0]);

   x.u = v;
   printf("%x  %x  %x  %d  %d\n", x.u, x.c[1], x.c[0], x.c[1], x.c[0]);
```

```
x.u = (((int) xL)<<16 | ((int) xR) & 0x0000ffff);
printf("%x  %x  %x  %d  %d\n", x.u, x.c[1], x.c[0], x.c[1], x.c[0]);
```

The program defines first a union structure variable x of the codec type. Given two 16-bit left/right numbers xL,xR (specified as 4-digit hex numbers), it defines a 32-bit unsigned integer v which is the concatenation of the two. The first printf statement prints the two numbers xL,xR in hex and decimal format. Note that the hex printing conversion operator %x treats the numbers as unsigned (some caution is required when printing negative numbers), whereas the decimal operator %d treats them as signed integers.

Next, the numbers xL,xR are assigned to the array members of the union x, such that x.c[1] = xL and x.c[0] = xR, and the second printf statement prints the contents of the union x, verifying that the 32-bit member x.u contains the concatenation of the two numbers with xL occupying the upper 16 bits, and xR, the lower 16 bits. Explain what the other two printf statements do.

Build and run the project (you may have to remove the file vectors.asm from the project's list of files). The output will appear in the stdout window at the bottom of the CCS. Alternatively, you may run this outside CCS using GCC. To do so, open a DOS window in your working directory and type the DOS command djgpp. This establishes the necessary environment variables to run GCC, then, run the following GCC command to generate the executable file unions.exe:

```
gcc unions.c -o unions.exe -lm
```

Repeat the run with the following choice of input samples:

```
xL = 0x1234;
xR = 0xabcd;
v = 0x1234abcd;
```

Explain the outputs of the print statements in this case by noting the following properties, which you should prove in your report:

$$(\text{0xffff0000})_{\text{unsigned}} = 2^{32} - 2^{16}$$

$$(\text{0xffffabcd})_{\text{unsigned}} = 2^{32} + (\text{0xabcd})_{\text{signed}}$$

$$(\text{0xffffabcd})_{\text{signed}} = (\text{0xabcd})_{\text{signed}}$$

1.7. Guitar Distortion Effects

In all of the experiments of Lab-2, the input/output maps are memoryless. We will study implementation of delays in a later lab. A memoryless mapping can be linear but time-varying, as was for example the case of panning between the speakers or the AM/FM wavetable experiments discussed discussed in another lab. The mapping can also be nonlinear.

Many guitar distortion effects combine delay effects with such nonlinear maps. In this part of Lab-1, we will study only some nonlinear maps in which each input sample x is mapped to an output sample y by a nonlinear function $y = f(x)$. Typical examples are hard clipping (called fuzz) and soft clipping that tries to emulated the nonlinearities of tube amplifiers. A typical nonlinear function is $y = \tanh(x)$. It has a sigmoidal shape that you can see by the quick MATLAB plot:

```
fplot('tanh(x)', [-4,4]); grid;
```

As suggested in Ref. [6], by keeping only the first two terms in its Taylor series expansion, that is, $\tanh(x) \approx x - x^3/3$, we may define a more easily realizable nonlinear function with built-in soft clipping:

$$y = f(x) = \begin{cases} +2/3, & x \geq 1 \\ x - x^3/3, & -1 \leq x \leq 1 \\ -2/3, & x \leq -1 \end{cases} \tag{1.2}$$

This can be plotted easily with

```
fplot('(abs(x)<1).*(x-1/3*x.^3) + sign(x).*(abs(x)>=1)*2/3', [-4,4]); grid;
```

The threshold value of $2/3$ is chosen so that the function $f(x)$ is continuous at $x = \pm 1$. To add some flexibility and to allow a variable threshold, we consider the following modification:

$$y = f(x) = \begin{cases} +\alpha c, & x \geq c \\ x - \beta c (x/c)^3, & -c \leq x \leq c \ , \\ -\alpha c, & x \leq -c \end{cases} \qquad \beta = 1 - \alpha \qquad (1.3)$$

where we assume that $c > 0$ and $0 < \alpha < 1$. The choice $\beta = 1 - \alpha$ is again dictated by the continuity requirement at $x = \pm c$. Note that setting $\alpha = 1$ gives the hard-thresholding, fuzz, effect:

$$y = f(x) = \begin{cases} +c, & x \geq c \\ x, & -c \leq x \leq c \\ -c, & x \leq -c \end{cases} \qquad (1.4)$$

Lab Procedure

First, run the above two `fplot` commands in MATLAB to see what these functions look like. The following program is a modification of the basic `template.c` program that implements Eq. (1.3):

```c
// soft.c - guitar distortion by soft thresholding
// -------------------------------------------------------------------------------

#include "dsplab.h"          // init parameters and function prototypes
#include <math.h>

// -------------------------------------------------------------------------------

#define a 0.67               // approximates the value 2/3
#define b (1-a)

short xL, xR, yL, yR;        // codec input and output samples

int x, y, on=1, c=2048;      // on/off variable and initial threshold c

int f(int);                  // function declaration

// -------------------------------------------------------------------------------

void main()                  // main program executed first
{
  initialize();              // initialize DSK board and codec, define interrupts

  sampling_rate(16);         // possible sampling rates: 8, 16, 24, 32, 44, 48, 96 kHz
  audio_source(LINE);        // LINE or MIC for line or microphone input

  while(1);                  // keep waiting for interrupt, then jump to isr()
}

// -------------------------------------------------------------------------------

interrupt void isr()         // sample processing algorithm - interrupt service routine
{
    read_inputs(&xL, &xR);   // read left and right input samples from codec

    if (on) {
       yL = (short) f((int) xL);  yL = yL << 1;     // amplify by factor of 2
       yR = (short) f((int) xR);  yR = yR << 1;
       }
```

```
        else
          {yL = xL; yR = xR;}

        write_outputs(yL,yR);      // write left and right output samples to codec

        return;
    }

    // --------------------------------------------------------------------------------

    int f(int x)
    {
        float y, xc = x/c;            // this y is local to f()

        y = x * (1 - b * xc * xc);

        if (x>c)  y =  a*c;           // force the threshold values
        if (x<-c) y = -a*c;

        return ((int) y);
    }

    // ----------------------------------------------------------------
```

a. Create a project for this program. In addition, create a GEL file that has two sliders, one for the on variable that turns the effect on or off in real time, and another slider for the threshold parameter c. Let c vary over the range $[0, 2^{14}]$ in increments of 512.

   Build and run the program, load the gel file, and display the two sliders. Then, play your favorite guitar piece and vary the slider parameters to hear the changes in the effect. (The wave file turn-turn3.wav in the directory c:\dsplab\wav is a good choice.)

b. Repeat the previous part by turning off the nonlinearity (i.e., setting $\alpha = 1$), which reduces to a fuzz effect with hard thresholding.

1.8. References

[1] S. J. Orfanidis, *Introduction to Signal Processing*, online book, 2010, available from:
    http://www.ece.rutgers.edu/~orfanidi/intro2sp/

[2] R. Chassaing and D. Reay, *Digital Signal Processing and Applications with the TMS320C6713 and TMS320C6416 DSK*, 2nd ed., Wiley, Hoboken, NJ, 2008.

[3] D. R. Brown III, 2009 Workshop on Digital Signal Processing and Applications with the TMS320C6713 DSK, Parts 1 & 2, available online from:
    http://spinlab.wpi.edu/courses/dspworkshop/dspworkshop_part1_2009.pdf
    http://spinlab.wpi.edu/courses/dspworkshop/dspworkshop_part2_2009.pdf

[4] N. Dahnoun, "DSP Implementation Using the TMS320C6711 Processors," contained in the Texas Instruments "C6000 Teaching Materials" CD ROM, 2002-04, and available online from TI:
    http://www.ti.com/ww/cn/uprogram/share/ppt/c6000/Chapter1.ppt
    http://www.ti.com/ww/cn/uprogram/share/ppt/c6000/Chapter2.ppt
    http://www.ti.com/ww/cn/uprogram/share/ppt/c6000/Chapter3.ppt

[5] B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, 2nd ed., Prentice Hall, Englewood Cliffs, NJ, 1988.

    S. P. Harbison and G. L. Steele, *C: A Reference Manual*, Prentice Hall, Englewood Cliffs, NJ, 1984.

    A. Kelly and I. Pohl, *A Book on C*, 2nd ed., Benjamin/Cummings, Redwood City, CA, 1990.

    GNU gcc, http://gcc.gnu.org/
    DJGPP - Windows version of GCC, http://www.delorie.com/djgpp/
    GCC Introduction, http://www.network-theory.co.uk/docs/gccintro/

[6] C.R. Sullivan. "Extending the Karplus-Strong Algorithm to Synthesize Electric Guitar Timbres with Distortion and Feedback," *Computer Music J.*, **14**, 26, (1990).

Lab 2 – Wavetable Generators, AM/FM Modulation

2.1. Lab Tasks

The concept of a wavetable is introduced and applied first to the generation of sinusoidal signals of different frequencies and then to square waves. AM and FM examples are constructed by combining two wavetables. Ring modulation and tremolo audio effects are studied as special cases of AM modulation.

2.2. Wavetable Generators

Wavetable generators are discussed in detail in Sect. 8.1.3 of the text [1]. A wavetable is defined by a circular buffer **w** whose dimension D is chosen such that the smallest frequency to be generated is:

$$f_{\min} = \frac{f_s}{D} \quad \Rightarrow \quad D = \frac{f_s}{f_{\min}}$$

For example, if $f_s = 8$ kHz and the smallest desired frequency is $f_{\min} = 10$ Hz, then one must choose $D = 8000/10 = 800$. The D-dimensional buffer holds one period at the frequency $f_{\min}$ of the desired waveform to be generated. The shape of the stored waveform is arbitrary, and can be a sinusoid, a square wave, sawtooth, etc. For example, if it is sinusoidal, then the buffer contents will be:

$$w[n] = \sin\left(\frac{2\pi f_{\min}}{f_s}\, n\right) = \sin\left(\frac{2\pi n}{D}\right), \quad n = 0, 1, \dots, D-1$$
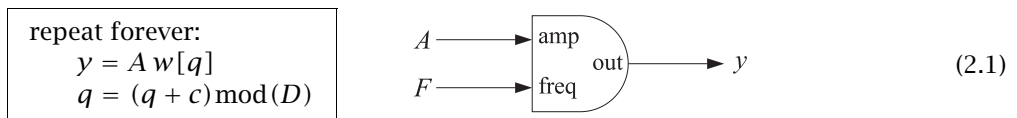
Similarly, a square wave whose first half is $+1$ and its second half, -1, will be defined as:

$$w[n] = \begin{cases} +1, & \text{if} \quad 0 \le n < D/2 \\ -1, & \text{if} \quad D/2 \le n < D \end{cases}$$

To generate higher frequencies (with the Nyquist frequency $f_s/2$ being the highest), the wavetable is cycled in steps of c samples, where c is related to the desired frequency by:

$$f = c f_{\min} = c\,\frac{f_s}{D} \quad \Rightarrow \quad c = D\,\frac{f}{f_s} \equiv D F, \quad F = \frac{f}{f_s}$$

where $F = f/f_s$ is the frequency in units of [cycles/sample]. The generated signal of frequency f and amplitude A is obtained by the loop:

$$\boxed{\begin{array}{l} \text{repeat forever:} \\ \quad y = A\, w[q] \\ \quad q = (q + c)\,\text{mod}\,(D) \end{array}} \qquad \begin{array}{c} A \longrightarrow \text{amp} \\ \quad\quad \text{out} \longrightarrow y \\ F \longrightarrow \text{freq} \end{array} \tag{2.1}$$

The shift c need not be an integer. In such case, the quantity $q + c$ must be truncated to the integer just below it. The text [1] discusses alternative methods, for example, rounding to the nearest integer, or, linearly interpolating. For the purposes of this lab, the truncation method will suffice.

The following function, `wavgen()`, based on Ref. [1], implements this algorithm. The mod-operation is carried out with the help of the function `qwrap()`:

```
// -----------------------------------------------
// wavgen.c - wavetable generator
// Usage: y = wavgen(D,w,A,F,&q);
// -----------------------------------------------

int qwrap(int, int);

float wavgen(int D, float *w, float A, float F, int *q)
{
    float y, c=D*F;
```

```
     y = A * w[*q];

     *q = qwrap(D-1, (int) (*q+c));

     return y;
  }

  // ------------------------------------------------
```

We note that the circular index q is declared as a pointer to **int**, and therefore, must be passed by address in the calling program. Before using the function, the buffer **w** must be loaded with one period of length D of the desired waveform. This function differs from the one in Ref. [1] in that it loads the buffer in forward order and cycles the index q forward.

2.3. Sinusoidal Wavetable

The following program, `sinex.c`, generates a 1 kHz sinusoid from a wavetable of length $D = 4000$. At a sampling rate of 8 kHz, the smallest frequency that can be generated is $f_{min} = f_s/D = 8000/4000 = 2$ Hz. In order to generate $f = 1$ kHz, the step size will be $c = D \cdot f/f_s = 4000 \cdot 1/8 = 500$ samples.

In this example, we will not use the function `wavgen` but rather apply the generation algorithm of Eq. (2.1) explicitly. In addition, we will save the output samples in a buffer array of length $N = 128$ and inspect the generated waveform both in the time and frequency domains using CCS's graphing capabilities.

```c
// sinex.c - sine wavetable example
//
// 332:348 DSP Lab - Spring 2012 - S. J. Orfanidis
// -------------------------------------------------------------------------------

#include "dsplab.h"          // DSK initialization declarations and function prototypes
#include <math.h>
//#define PI 3.141592653589793

short xL, xR, yL, yR;        // left and right input and output samples from/to codec

#define D 4000               // fmin = fs/D = 8000/4000 = 2 Hz
#define N 128                // buffer length

short fs=8;                  // fs = 8 kHz
float c, A=5000, f=1;        // f = 1 kHz
float w[D];                  // wavetable buffer
float buffer[N];             // buffer for plotting with CCS
int q=0, k=0;

// -------------------------------------------------------------------------------

void main()                  // main program executed first
{
  int n;
  float PI = 4*atan(1);

  for (n=0; n<D; n++) w[n] = sin(2*PI*n/D);      // load wavetable with one period

  c = D*f/fs;                // step into wavetable buffer

  initialize();              // initialize DSK board and codec, define interrupts

  sampling_rate(fs);         // possible sampling rates: 8, 16, 24, 32, 44, 48, 96 kHz
 // audio_source(LINE);      // LINE or MIC for line or microphone input

  while(1);      // wait for interrupts
}
```

```
// -------------------------------------------------------------------------------

interrupt void isr()
{
   yL = (short) (A * w[q]);                   // generate sinusoidal output
   q = (int) (q+c); if (q >= D) q = 0;        // cycle over wavetable in steps c

   buffer[k] = (float) yL;                    // save into buffer for plotting
   if (++k >= N) k=0;                          // cycle over buffer

   write_outputs(yL,yL);                      // audio output
}

// -------------------------------------------------------------------------------
```

The wavetable is loaded with the sinusoid in `main()`. At each sampling instant, the program does nothing with the codec inputs, rather, it generates a sample of the sinusoid and sends it to the codec, and saves the sample into a buffer (only the last N generated samples will be in present in that buffer).

Lab Procedure

a. Create a project for this program and run it. The amplitude was chosen to be $A = 5000$ in order to make the wavetable output audible. Hold the processor after a couple of seconds (SHIFT-F5).

b. Using the keyboard shortcut, "ALT-V RT", or the menu commands *View -> Graph -> Time/Frequency*, open a graph-properties window as that shown below:
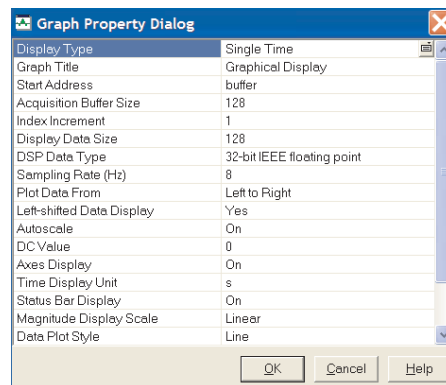
Select the starting address to be, `buffer`, set the sampling rate to 8 and look at the time waveform. Count the number of cycles displayed. Can you predict that number from the fact that N samples are contained in that buffer? Next right-click on the graph and select "Properties", and choose "FFT Magnitude" as the plot-type. Verify that the peak is at $f = 1$ kHz.

c. Reset the frequency to 500 Hz. Repeat parts (a,b).

d. Create a GEL file with a slider for the value of the frequency over the interval $0 \leq f \leq 1$ kHz in steps of 100 Hz. Open the slider and run the program while changing the frequency with the slider.

e. Set the frequency to 30 Hz and run the program. Keep decreasing the frequency by 5 Hz at a time and determine the lowest frequency that you can hear (but, to be fair don't increase the speaker volume; that would compensate the attenuation introduced by your ears.)

f. Replace the following two lines in the `isr()` function:

```
yL = (short) (A * w[q]);
q = (int) (q+c); if (q >= D) q = 0;
```

by a single call to the function `wavgen`, and repeat parts (a,b).

g. Replace the sinusoidal table of part (f) with a square wavetable that has period 4000 and is equal to +1 for the first half of the period and −1 for the second half. Run the program with frequency $f = 1$ kHz and $f = 200$ Hz.

h. Next, select the sampling rate to be $f_s = 96$ kHz and for the sinusoid case, start with the frequency $f = 8$ kHz and keep increasing it by 2 kHz at a time till about 20 kHz to determine the highest frequency that you can hear—each time repeating parts (a,b).

2.4. AM Modulation

Here, we use two wavetables to illustrate AM modulation. The picture below shows how one wavetable is used to generate a modulating amplitude signal, which is fed into the amplitude input of a second wavetable.
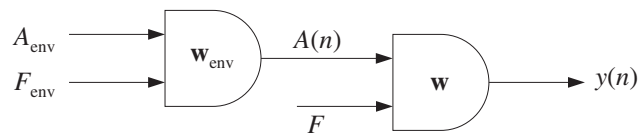
The AM-modulated signal is of the form:

$$x(t) = A(t)\sin(2\pi f t), \qquad \text{where} \quad A(t) = A_{\text{env}}\sin(2\pi f_{\text{env}}t)$$

The following program, `amex.c`, shows how to implement this with the function `wavgen()`. The envelope frequency is chosen to be 2 Hz and the signal frequency 200 Hz. A common sinusoidal wavetable sinusoidal buffer is used to generate both the signal and its sinusoidal envelope.

```
// amex.c - AM example
// --------------------------------------------------------------------------------

#include "dsplab.h"          // DSK initialization declarations and function prototypes
#include <math.h>
//#define PI 3.141592653589793

short xL, xR, yL, yR;         // left and right input and output samples from/to codec

#define D 8000               // fmin = fs/D = 8000/8000 = 1 Hz
float w[D];                   // wavetable buffer

short fs=8;
float A, f=0.2;
float Ae=10000, fe=0.002;
int q, qe;

float wavgen(int, float *, float, float, int *);

// --------------------------------------------------------------------------------

void main()
{
  int i;
  float PI = 4*atan(1);

  q=qe=0;

  for (i=0; i<D; i++) w[i] = sin(2*PI*i/D);        // fill sinusoidal wavetable

  initialize();
  sampling_rate(fs);
  audio_source(LINE);

  while(1);
```

```
   }

   // --------------------------------------------------------------------------

   interrupt void isr()
   {
     float y;

      // read_inputs(&xL, &xR);                // inputs not used

      A = wavgen(D, w, Ae, fe/fs, &qe);
      y = wavgen(D, w, A, f/fs, &q);

      yL = yR = (short) y;

      write_outputs(yL,yR);

     return;
   }

   // --------------------------------------------------------------------------
```

Although the buffer is the same for the two wavetables, two different circular indices, q, q_e are used for the generation of the envelope amplitude signal and the carrier signal.

Lab Procedure

a. Run and listen to this program with the initial signal frequency of $f = 200$ Hz and envelope frequency of $f_{env} = 2$ Hz. Repeat for $f = 2000$ Hz. Repeat the previous two cases with $f_{env} = 20$ Hz.

b. Repeat and explain what you hear for the cases:

$$f = 200 \text{ Hz}, \quad f_{env} = 100 \text{ Hz}$$
$$f = 200 \text{ Hz}, \quad f_{env} = 190 \text{ Hz}$$
$$f = 200 \text{ Hz}, \quad f_{env} = 200 \text{ Hz}$$

2.5. FM Modulation

The third program, fmex.c, illustrates FM modulation in which the frequency of a sinusoid is time-varying. The generated signal is of the form:

$$x(t) = \sin\left[2\pi f(t)t\right]$$

The frequency $f(t)$ is itself varying sinusoidally with frequency f_m:

$$f(t) = f_0 + A_m \sin(2\pi f_m t)$$

Its variation is over the interval $f_0 - A_m \le f(t) \le f_0 + A_m$. In this experiment, we choose the modulation depth $A_m = 0.3f_0$, so that $0.7f_0 \le f(t) \le 1.3f_0$. The center frequency is chosen as $f_0 = 500$ Hz and the modulation frequency as $f_m = 1$ Hz. Again two wavetables are used as shown below, with the first one generating $f(t)$, which then drives the frequency input of the second generator.
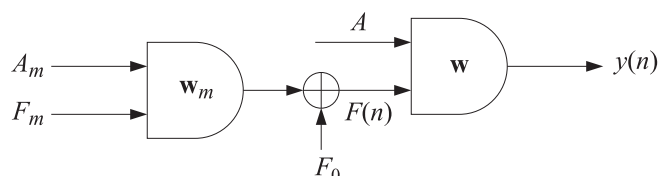
```
// fmex.c - FM example
// ------------------------------------------------------------------------------

#include "dsplab.h"           // DSK initialization declarations and function prototypes
#include <math.h>
//#define PI 3.141592653589793

short xL, xR, yL, yR;         // left and right input and output samples from/to codec

#define D 8000                // fmin = fs/D = 8000/8000 = 1 Hz
float w[D];                   // wavetable buffer

short fs=8;
float A=5000, f=0.5;
float Am=0.3, fm=0.001;
int q, qm;

float wavgen(int, float *, float, float, int *);

// -------------------------------------------------------------------------------

void main()
{
  int i;
  float PI = 4*atan(1);

  q = qm = 0;

  for (i=0; i<D; i++) w[i] = sin(2*PI*i/D);          // load sinusoidal wavetable
  //for (i=0; i<D; i++) w[i] = (i<D/2)? 1 : -1;      // square wavetable

  initialize();
  sampling_rate(fs);
  audio_source(LINE);

  while(1);
}

// -------------------------------------------------------------------------------

interrupt void isr()
{
  float y, F;

  // read_inputs(&xL, &xR);                           // inputs not used

  F = (1 + wavgen(D, w, Am, fm/fs, &qm)) * f/fs;       // modulated frequency

  y = wavgen(D, w, A, F, &q);                          // FM signal

  yL = yR = (short) y;

  write_outputs(yL,yR);

  return;
}

// -------------------------------------------------------------------------------
```

Lab Procedure

a. Compile, run, and hear the program with the following three choices of the modulation depth: $A_m = 0.3f_0$, $A_m = 0.8f_0$, $A_m = f_0$, $A_m = 0.1f_0$. Repeat these cases when the center frequency is changed to $f_0 = 1000$ Hz.

b. Replace the sinusoidal wavetable with a square one and repeat the case $f_0 = 500$ Hz, $A_m = 0.3f_0$. You

will hear a square wave whose frequency switches between a high and a low value in each second.

c. Keep the square wavetable that generates the alternating frequency, but generate the signal by a sinu-soidal wavetable. To do this, generate a second sinusoidal wavetable and define a circular buffer for it in `main()`. Then generate your FM-modulated sinusoid using this table. The generated signal will be of the form:

$$x(t) = \sin\left[2\pi f(t)t\right], \qquad f(t) = 1 \text{ Hz square wave}$$

2.6. Ring Modulators and Tremolo

Interesting audio effects can be obtained by feeding the audio input to the amplitude of a wavetable generator and combining the resulting output with the input, as shown below:
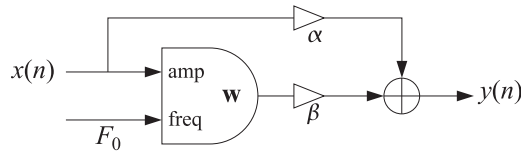
For example, for a sinusoidal generator of frequency $F_0 = f_0/f_s$, we have:

$$y(n) = \alpha x(n) + \beta x(n)\cos(2\pi F_0 n) = x(n)\left[\alpha + \beta\cos(2\pi F_0 n)\right] \tag{2.2}$$

The *ring modulator* effect is obtained by setting $\alpha = 0$ and $\beta = 1$, so that

$$y(n) = x(n)\cos(2\pi F_0 n) \tag{2.3}$$

whereas, the *tremolo* effect corresponds to $\alpha = 1$ and $\beta \neq 0$

$$y(n) = x(n) + \beta x(n)\cos(2\pi F_0 n) = x(n)\left[1 + \beta\cos(2\pi F_0 n)\right] \tag{2.4}$$

The following ISR function implements either effect:

```
// ------------------------------------------------------------------------------------

interrupt void isr()
{
  float x, y;

   read_inputs(&xL, &xR);

   x = (float) xL;

   y = alpha * x + beta * wavgen(D, w, x, f/fs, &q);

   yL = yR = (short) y;

   write_outputs(yL,yR);

   return;
}

// ------------------------------------------------------------------------------------
```

Lab Procedure

a. Modify the `amex.c` project to implement the ring modulator/tremolo effect. Set the carrier frequency to $f_0 = 400$ Hz and $\alpha = \beta = 1$. Compile, run, and play a wavefile with voice in it (e.g., `dsummer.`)

b. Experiment with higher and lower values of f_0.

c. Repeat part (a) when $\alpha = 0$ and $\beta = 1$ to hear the ring-modulator effect.

2.7. Scrambler as Ring Modulator

In the frequency domain, Eq. (2.3) is equivalent to frequency translation:

$$Y(f) = \frac{1}{2}\Big[X(f - f_0) + X(f + f_0)\Big] \tag{2.5}$$

As f_0 is chosen closer and closer to the Nyquist frequency $f_s/2$, the shifted replicas begin to resemble the inverted spectrum of $X(f)$. In particular, if $f_0 = f_s/2$, then,
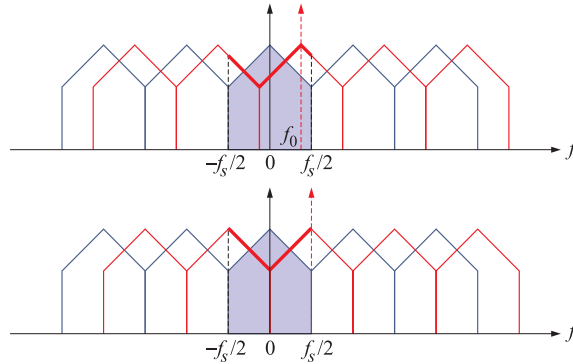
$$Y(f) = \frac{1}{2}\Big[X(f - f_s/2) + X(f + f_s/2)\Big]$$

Using the periodicity property $X(f \pm f_s) = X(f)$, we then obtain the equivalent expressions:

$$Y(f) = \frac{1}{2}\Big[X(f - f_s/2) + X(f + f_s/2 - f_s)\Big] = X(f - f_s/2), \quad 0 \le f \le \frac{f_s}{2}$$

$$Y(f) = \frac{1}{2}\Big[X(f - f_s/2 + f_s) + X(f + f_s/2)\Big] = X(f + f_s/2), \quad -\frac{f_s}{2} \le f \le 0$$

which imply that the positive (negative) frequency part of $Y(f)$ is equal to the negative (positive) frequency part of $X(f)$, in other words, $Y(f)$ is the inverted version of $X(f)$. This is depicted below.

Because in this case $F_0 = f_0/f_s = (f_s/2)/f_s = 1/2$, the carrier waveform is simply the alternating sequence of ± 1:

$$\cos(2\pi F_0 n) = \cos(\pi n) = (-1)^n$$

and the modulator output becomes

$$y(n) = (-1)^n x(n) \tag{2.6}$$

Lab Procedure

Modify the `template.c` program to implement the frequency-inversion or scrambling operation of Eq. (2.6). This can be done easily by introducing a global index:

```
int q = 1;
```

and keep changing its sign at each interrupt call, i.e., after reading the left/right codec inputs, define the corresponding codec outputs by:

```
yL = q * xL;
yR = q * xR;

q = -q;
```

Compile and run this program. Send the wave file `JB.wav` into it. First comment out the line `q = -q`, and hear the file as pass through. Then, enable the line, recompile, and hear the scrambled version of the file.

The scrambled version was recorded with MATLAB and saved into another wave file, `JBm.wav`. If you play that through the scrambler program, it will get unscrambled. In Labs 3 & 4, we will implement the frequency inversion in alternative ways.

2.8. References

[1] S. J. Orfanidis, *Introduction to Signal Processing*, online book, 2010, available from:
     `http://www.ece.rutgers.edu/~orfanidi/intro2sp/`

[2] R. Chassaing and D. Reay, *Digital Signal Processing and Applications with the TMS320C6713 and TMS320C6416 DSK*, 2nd ed., Wiley, Hoboken, NJ, 2008.

[3] F. R. Moore, *Elements of Computer Music*, Prentice Hall, Englewood Cliffs, NJ, 1990.

[4] C. Dodge and T. A. Jerse, *Computer Music*, Schirmer/Macmillan, New York, 1985.

[5] J. M. Chowning, "The Synthesis of Complex Audio Spectra by Means of Frequency Modulation," *J. Audio Eng. Soc.*, **21**, 526 (1973).

[6] M. Kahrs and K. Brandenburg, eds., *Applications of Digital Signal Processing to Audio and Acoustics*, Kluwer, Boston, 1998.

[7] Udo Zölzer, ed., *DAFX – Digital Audio Effects*, Wiley, Chichester, England, 2003. See also the DAFX Conference web page: `http://www.dafx.de/`.

Lab 3 – Delays and FIR Filtering

3.1. Introduction

In this lab you will study sample by sample processing methods for FIR filters and implement them on the TMS320C6713 processor. Once you know how to implement a multiple delay on a sample by sample basis, it becomes straightforward to implement FIR and IIR filters. Multiple delays are also the key component in various digital audio effects, such as reverb.

Delays can be implemented using linear or circular buffers, the latter being more efficient, especially for audio effects. The theory behind this lab is developed in Ch. 4 of the text [1] for FIR filters, and used in Ch. 8 for audio effects.

3.2. Delays Using Linear and Circular Buffers

A D-fold delay, also referred to as a delay line, has transfer function $H(z) = z^{-D}$ and corresponds to a time delay in seconds:

$$T_D = DT = \frac{D}{f_s} \quad \Rightarrow \quad D = f_s T_D \tag{3.1}$$

where T is the time interval between samples, related to the sampling rate by $f_s = 1/T$. A block diagram realization of the multiple delay is shown below:
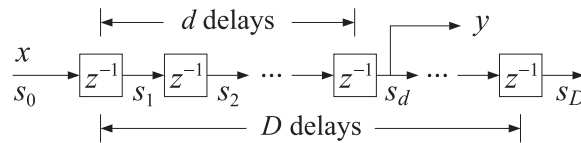

Fig. 3.1 Tapped delay line.

There are D registers whose contents are the "internal" states of the delay line. The dth state s_d, i.e., the content of the dth register, represents the d-fold delayed version of the input, that is, at time n we have: $s_d(n) = x(n-d)$, for $d = 1, \ldots, D$; the case $d = 0$ corresponds to the input $s_0(n) = x(n)$.

At each time instant, all D contents are available for processing and can be "tapped" out for further use (e.g., to implement FIR filters). For example, in the above diagram, the dth tap is being tapped, and the corresponding transfer function from the input x to the output $y = s_d$ is the partial delay z^{-d}.

The D contents/states $s_d, d = 1, 2, \ldots, D$, and the input $s_0 = x$ must be stored in memory in a $(D+1)$-dimensional array or buffer. But the manner in which they are stored and retrieved depends on whether a linear or a circular buffer is used. The two cases are depicted below.
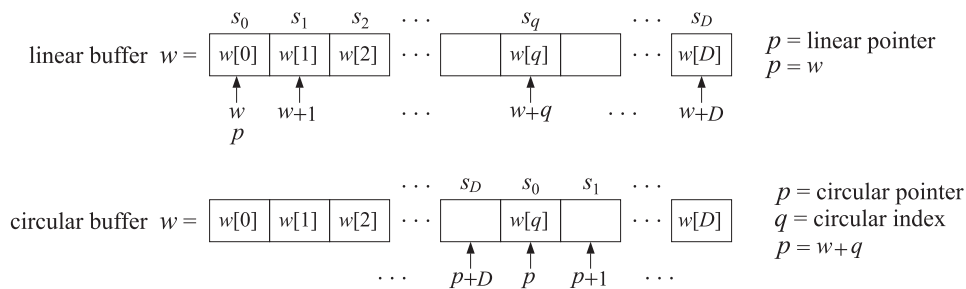

Fig. 3.2 Linear and circular buffers.

In both cases, the buffer can be created in C by the declaration:

```
float w[D+1];
```

Its contents are retrieved as $w[i]$, $i = 0, 1, \ldots, D$. Thinking of w as a pointer, the contents can also be retrieved by $*(w + i) = w[i]$, where $*$ denotes the de-referencing operator in C.

In the linear buffer case, the states are stored in the buffer sequentially, or linearly, that is, the ith state is:

$$s_i = w[i] = *(w + i), \quad i = 0, 1, \ldots, D$$

At each time instant, after the contents s_i are used, the delay-line is updated in preparation for the next time instant by shifting its contents to the right from one register to the next, as suggested by the block diagram in Fig. 3.1. This follows from the definition $s_i(n) = x(n - i)$, which implies for the next time instant $s_i(n+1) = x(n+1-i) = s_{i-1}(n)$. Thus, the current s_{i-1} becomes the next s_i. Since $s_i = w[i]$, this leads to the following updating algorithm for the buffer contents:

$$\text{for } i = D \text{ down to } i = 1, \text{ do:}$$
$$w[i] = w[i - 1]$$

where the shifting is done from the right to the left to prevent the over-writing of the correct contents. It is implemented by the C function `delay()` of the text [1]:

```
// delay.c - linear buffer updating
// ------------------------------

void delay(int D, float *w)
{
    int i;

    for (i=D; i>=1; i--)
        w[i] = w[i-1];
}

// ------------------------------
```

For large values of D, this becomes an inefficient operation because it involves the shifting of large amounts of data from one memory location to the next. An alternative approach is to keep the data unshifted but to shift the beginning address of the buffer to the left by one slot.

This leads to the concept of a circular buffer in which a movable pointer p is introduced that always points somewhere within the buffer array, and its current position allows one to retrieve the states by $s_i = *(p + i)$, $i = 0, 1, \ldots, D$. If the pointer $p + i$ exceeds the bounds of the array to the right, it gets wrapped around to the beginning of the buffer.

To update the delay line to the next time instant, the pointer is left-shifted, i.e., by the substitution $p = p - 1$, or, $--p$, and is wrapped to the right end of the buffer if it exceeds the array bounds to the left. Fig. 3.3 depicts the contents and pointer positions at two successive time instants for the linear and circular buffer cases for $D = 3$. In both cases, the states are retrieved by $s_i = *(p + i)$, $i = 0, 1, 2, 3$, but in the linear case, the pointer remains frozen at the beginning of the buffer, i.e., $p = w$, and the buffer contents shift forwards, whereas in the circular case, p shifts backwards, but the contents remain in place.
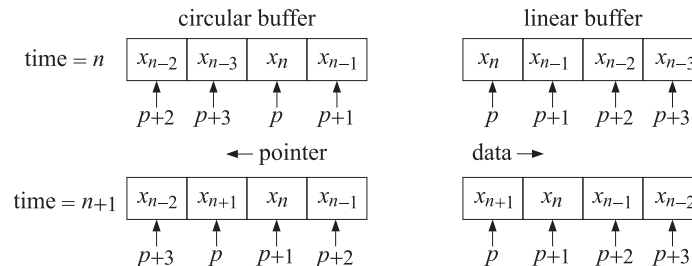

Fig. 3.3 Buffer contents at successive time instants for $D = 3$.

In the text [1], the functions `tap()` and `cdelay()` are used for extracting the states s_i and for the circular back-shifting of the pointer. Although these two functions could be used in the CCS environment, we prefer instead to use a single function called `pwrap()` that calculates the new pointer after performing the required wrapping. The function is declared in the common header file `dsplab.h` and defined in the file `dsplab.c` in the directory `C:\dsplab\common`. Its listing is as follows:

```
// pwrap.c - pointer wrapping relative to circular buffer
// Usage: p_new = pwrap(D,w,p)
// ----------------------------------------------------

float *pwrap(int D, float *w, float *p)
{
    if (p > w+D)
        p -= D+1;

    if (p < w)
        p += D+1;

    return p;
}

// ----------------------------------------------------
```

The ith state s_i and the updating of the delay-line can be obtained by the function calls:

$$s_i = *\text{pwrap}(D, w, p + i), \quad i = 1, 2, \dots, D$$

$$p_{\text{next}} = \text{pwrap}(D, w, --p)$$

We will use this function in the implementation of FIR filters and in various audio effects. It will allow us to easily translate a sample processing algorithm expressed in pseudo-code into the actual C code. As an example, let us consider the circular buffer implementation of the partial delay z^{-d}. The block diagram of Fig. 3.1 and the pseudo-code computational algorithm are as follows:
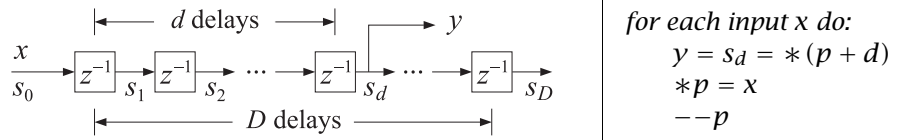
We may translate this into C by the following operations using `pwrap`:

```
y = *pwrap(D,w,p+d);          // delay output
*p = x;                       // delay-line input
p = pwrap(D,w,--p);           // backshift circular buffer pointer
```

In the last line, we must pre-decrement the pointer inside `pwrap`, that is, `--p`, instead of post-decrementing it, `p--`. Why? By comparison, the linear buffer implementation, using a $(D+1)$-dimensional buffer, is as follows:

```
y = w[d];                     // delay output
w[0] = x;                     // delay-line input
for (i=D; i>=0; i--)          // update linear buffer
    w[i] = w[i-1];
```

An alternative approach to circular buffers is working with circular indices instead of pointers. The pointer p always points at some element of the buffer array w, that is, there is a unique integer q such that $p = w + q$, with corresponding content $*p = w[q]$. This is depicted in Fig. 3.2. The index q is always bound by the limits $0 \leq q \leq D$ and wrapped modulo-$(D+1)$ if it exceeds these limits.

The textbook functions `tap2()` and `cdelay2()`, and their corresponding MATLAB versions given in the Appendix of [1], implement this approach. Again, however, we prefer to use the following function, `qwrap()`, also included in the common file `dsplab.c`, that calculates the required wrapped value of the circular index:

```
// qwrap.c - circular index wrapping
// Usage: q_new = qwrap(D,q);
// -----------------------------------

int qwrap(int D, int q)
{
      if (q > D)
            q -= D + 1;

      if (q < 0)
            q += D + 1;

      return q;
}

// -----------------------------------
```

In terms of this function, the above *d*-fold delay example is implemented as follows:

```
qd = qwrap(D,q+d);              // (q+d) mod (D+1)
y = w[qd];                     // delayed output
w[q] = x;                      // delay-line input
q = qwrap(D,--q);              // backshift pointer index
```

We note that in general, the *i*th state is:

$$s_i = *(p + i) = *(w + q + i) = w[q + i]$$

where $q + i$ must be wrapped as necessary. Thus, the precise way to extract the *i*th state is:

$$q_i = \mathrm{qwrap}(D, q + i), \quad s_i = w[q_i], \quad i = 1, 2, \ldots, D$$

Lab Procedure

A complete C program that implements the above *d*-fold delay example on the TMS320C6713 processor is given below:

```
// delay1.c - multiple delay example using circular buffer pointers (pwrap version)
// --------------------------------------------------------------------------------

#include "dsplab.h"          // init parameters and function prototypes

short xL, xR, yL, yR;        // input and output samples from/to codec

#define D 8000               // max delay in samples (TD = D/fs = 8000/8000 = 1 sec)
short fs = 8;                // sampling rate in kHz
float w[D+1], *p, x, y;      // circular delay-line buffer, circular pointer, input, output
int d = 4000;                // must be d <= D

// --------------------------------------------------------------------------------

void main()                          // main program executed first
{
   int n;

   for (n=0; n<=D; n++) w[n] = 0;    // initialize circular buffer to zero
   p = w;                            // initialize pointer

   initialize();                     // initialize DSK board and codec, define interrupts

   sampling_rate(fs);                // possible sampling rates: 8, 16, 24, 32, 44, 48, 96 kHz
   audio_source(MIC);                // use LINE or MIC for line or microphone input

   while(1);                         // keep waiting for interrupt, then jump to isr()
}
```

```
// ----------------------------------------------------------------------------

interrupt void isr()         // sample processing algorithm - interrupt service routine
{
    read_inputs(&xL, &xR);   // read left and right input samples from codec

    x = (float) xL;          // work with left input only

    y = *pwrap(D,w,p+d);     // delayed output - pwrap defined in dsplab.c
    *p = x;                  // delay-line input
    p = pwrap(D,w,--p);      // backshift pointer

    yL = yR = (short) y;

    write_outputs(yL,yR);    // write left and right output samples to codec

    return;
}

// ----------------------------------------------------------------------------
```

Note the following features. The sampling rate is set to 8 kHz, therefore, the maximum delay $D = 8000$ corresponds to a delay of 1 sec, and the partial delay $d = 4000$, to $1/2$ sec. The circular buffer array w has dimension $D + 1 = 8001$ and its scope is global within this file. It is initialized to zero within `main()` and the pointer p is initialized to point to the beginning of w, that is, $p = w$.

The left/right input samples, which are of the **short int** type, are cast to **float**, while the **float** output is cast to **short int** before it is sent out to the codec.

a. Create and build a project for this program. Then, run it. Give the system an impulse by lightly tapping the table with the mike, and listen to the impulse response. Then, speak into the mike.

Bring the mike near the speaker and then give the system an impulse. You should hear repeated echoes. If you bring the mike too close to the speakers the output goes unstable. Draw a block diagram realization that would explain the effect you are hearing. Experimentally determine the distance at which the echoes remain marginally stable, that is, neither die out nor diverge. (Technically speaking, the poles of your closed-loop system lie on the unit circle.)

b. Change the sampling rate to 16 kHz, recompile and reload keeping the value of d the same, that is, $d = 4000$. Listen to the impulse response. What is the duration of the delay in seconds now?

c. Reset the sampling rate back to 8 kHz, and this time change d to its maximum value $d = D = 8000$. Recompile, reload, and listen to the impulse response. Experiment with lower and lower values of d and listen to your delayed voice until you can no longer distinguish a separate echo. How many milliseconds of delay does this correspond to?

d. Set $d = 0$, recompile and reload. This should correspond to no delay at all. But what do you hear? Can you explain why? Can you fix it by changing the program? Will your modified program still work with $d \neq 0$? Is there any good reason for structuring the program the way it was originally?

e. In this part you will profile the computational cost of the sample processing algorithm. Open the source file `delay1.c` in a CCS window. Locate the `read_inputs` line in the `isr()`, then right-click on that line and choose *Toggle Software Breakpoint*; a red dot will appear in the margin. Do the same for the `write_outputs` line.

From the top menu of the CCS window, choose *Profile -> Clock -> View*; a little yellow clock will appear on the right bottom status line of CCS. When you compile, load, and run your program, it will stop at the first breakpoint, with a yellow arrow pointing to it. Double-click on the profile clock to clear the number of cycles, then type F5 to continue running the program and it will stop at the second breakpoint. Read and record the number of cycles shown next to the profile clock.

f. Write a new program, called `delay2.c`, that makes use of the function `qwrap` instead of `pwrap`. Repeat parts (a) and (e).

g. Next, write a new program, called `delay3.c`, that uses linear buffers. Its `isr()` will be as follows:

```
interrupt void isr()
{
   int i;

   read_inputs(&xL, &xR);

   x = (float) xL;

   w[0] = x;               // delay-line input
   y = w[d];               // delay output
   for (i=D; i>=0; i--)    // update linear buffer
      w[i] = w[i-1];

   yL = yR = (short) y;

   write_outputs(yL,yR);

   return;
}
```

Build the project. You will find that it may not run (because the data shifts require too many cycles that over-run the sampling rate). Change the program parameters D, d to the following values $D = 2000$ and $d = 1000$. Rebuild and run the program. Repeat part (e) and record the number of cycles. Change the parameters D, d of the program `delay1.c` to the same values, and repeat part (e) for that. Comment on the required number of samples using the linear vs. the circular buffer implementation.

3.3. FIR Comb Filters Using Circular Buffers

More interesting audio effects can be derived by combining several multiple delays. An example is the FIR comb filter defined by Eq. (8.2.8) of the text [1]:

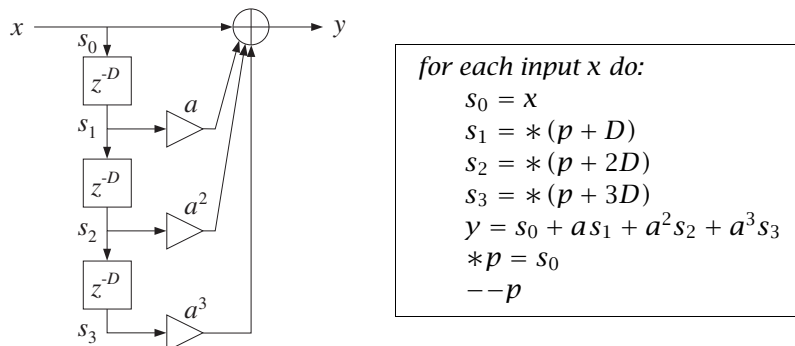$$y_n = x_n + a\, x_{n-D} + a^2 x_{n-2D} + a^3 x_{n-3D}$$

Its transfer function is given by Eq. (8.2.9):

$$H(z) = 1 + a z^{-D} + a^2 z^{-2D} + a^3 z^{-3D}$$

Its impulse response has a very sparse structure:

$$\mathbf{h} = [1, \underbrace{0, 0, \dots, 0}_{D-1 \text{ zeros}}, a, \underbrace{0, 0, \dots, 0}_{D-1 \text{ zeros}}, a^2, \underbrace{0, 0, \dots, 0}_{D-1 \text{ zeros}}, a^3]$$

The comb-like structure of its frequency response and its zero-pattern on the z-plane are depicted in Fig. 8.2.5 of [1]. Instead of implementing it as a general FIR filter, a more efficient approach is to program the block diagram directly by using a single delay line of order $3D$ and tapping it out at taps 0, D, $2D$, and $3D$. The block diagram realization and corresponding sample processing algorithm are:


```
for each input x do:
    s₀ = x
    s₁ = *(p + D)
    s₂ = *(p + 2D)
    s₃ = *(p + 3D)
    y = s₀ + a s₁ + a² s₂ + a³ s₃
    *p = s₀
    −−p
```

The translation of the sample processing algorithm into C is straightforward and can be incorporated into the following `isr()` function to be included in your main program:

```
interrupt void isr()
{
    float s0, s1, s2, s3, y;            // states & output

    read_inputs(&xL, &xR);              // read inputs from codec

    s0 = (float) xL;                    // work with left input only
    s1 = *pwrap(3*D,w,p+D);             // extract states relative to p
    s2 = *pwrap(3*D,w,p+2*D);           // note, buffer length is 3D+1
    s3 = *pwrap(3*D,w,p+3*D);
    y = s0 + a*s1 + a*a*s2 + a*a*a*s3;  // output sample
    *p = s0;                            // delay-line input
    p = pwrap(3*D,w,--p);               // backshift pointer

    yL = yR = (short) y;

    write_outputs(yL,yR);               // write outputs to codec

    return;
}
```

Lab Procedure

Set the sampling rate to 8 kHz and the audio source to microphone. Choose the delay to be $D = 4000$, corresponding to $T_D = 0.5$ sec, so that the total duration of the filter is $3T_D = 1.5$ sec, and set $a = 0.5$.

a. Write a C program called `comb.c` that incorporates the above interrupt service routine. You will need to globally declare/define the parameters D, a, p, as well as the circular buffer w to be a $3D+1$ dimensional float array. Make sure you initialize the buffer to zero inside `main()`, as was done in the previous example, and also initialize $p = w$.

   Build and run this project. Listen to the impulse response of the filter by tapping the table with the mike. Speak into the mike. Bring the mike close to the speakers and get a closed-loop feedback.

b. Keeping the delay D the same, choose $a = 0.2$ and run the program again. What effect do you hear? Repeat for $a = 0.1$. Repeat with $a = 1$.

c. Set the audio input to LINE and play your favorite wave file or MP3 into the input. Experiment with reducing the value of D in order to match your song's tempo to the repeated echoes. Some wave files are in the directory `c:\dsplab\wav` (e.g., try `dsummer`, `take5`.)

d. The FIR comb can also be implemented *recursively* using the geometric series formula to rewrite its transfer function in the recursive form as shown in Eq. (8.2.9) of the text:

$$H(z) = 1 + az^{-D} + a^2z^{-2D} + a^3z^{-3D} = \frac{1 - a^4z^{-4D}}{1 - az^{-D}}$$

This requires a $(4D+1)$-dimensional delay-line buffer w. The canonical realization and the corresponding sample processing algorithm are shown below:

for each input x do:
$$s_1 = *(p + D)$$
$$s_4 = *(p + 4D)$$
$$s_0 = x + a s_1$$
$$y = s_0 - a^4 s_4$$
$$*p = s_0$$
$$--p$$
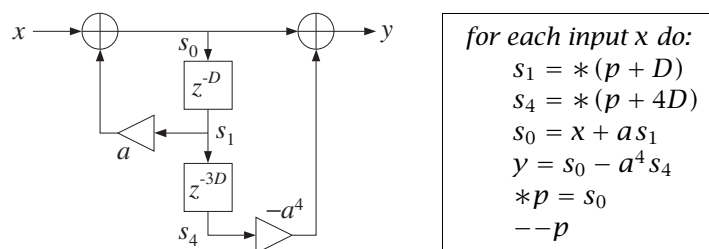
Write a new program, `comb2.c`, that implements this algorithm. Remember to define the buffer to be a $(4D+1)$-dimensional float array. Using the values $D = 1600$ (corresponding to a 0.2 sec delay) and $a = 0.5$, recompile and run both the `comb.c` and `comb2.c` programs and listen to their outputs.

In general, such recursive implementations of FIR filters are more prone to the accumulation of round-off errors than the non-recursive versions. You may want to run these programs with $a = 1$ to observe this sensitivity.

3.4. FIR Filters with Linear and Circular Buffers

The sample-by-sample processing implementation of FIR filters is discussed in Sect. 4.2 of the text [1]. For an order-M filter, the input/output convolutional equation can be written as the dot product of the filter-coefficient vector $\mathbf{h} = [h_0, h_1, \ldots, h_M]^T$ with the state vector $\mathbf{s}(n) = [x_n, x_{n-1}, \ldots, x_{n-M}]^T$:

$$y_n = \sum_{m=0}^{M} h(m)x(n-m) = [h_0, h_1, \ldots, h_M] \begin{bmatrix} x(n) \\ x(n-1) \\ \vdots \\ x(n-M) \end{bmatrix} = \mathbf{h}^T \mathbf{s}(n), \quad \mathbf{s}(n) = \begin{bmatrix} x(n) \\ x(n-1) \\ \vdots \\ x(n-M) \end{bmatrix}$$

A block diagram realization for the case $M = 3$ is depicted below.
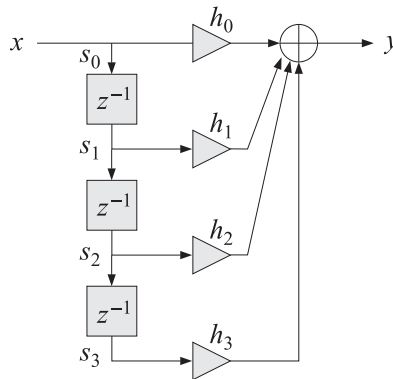
We note that the ith component of the state vector is $s_i(n) = x(n-i)$, $i = 0, 1, \ldots, M$, and therefore, the states are the tap outputs of a multiple delay-line with M delays. Thus, the definition of the delay line and its time updating remains the same as in the previous sections. To realize the FIR filter, we must use the tapped outputs s_i from the delay line to calculate the dot product, and then update the delay line to the next time instant.

In this lab, we consider five implementations of FIR filters and study their relative efficiency in terms of machine cycles at different levels of compiler optimization:

```
y = fir(M, h, w, x);          - linear buffer implementation
y = firc(M, h, w, &p, x);     - circular buffer with pointers
y = firc2(M, h, w, &q, x);    - circular index with updating in loop
y = firq(M, h, w, &q, x);     - circular index with updating outside loop
y = fira(w, h, Lh, Nb, q);    - circular buffer in linear assembly
```

These functions are defined below. The function, `fir`, implements the linear buffer case:

```
// fir.c - FIR filter in direct form with linear buffer
// Usage: y = fir(M, h, w, x);
// --------------------------------------------------------------

float fir(int M, float *h, float *w, float x)
{
   int i;
```

```
      float y;                        // y=output sample

      w[0] = x;                       // read current input sample x

      for (y=0, i=0; i<=M; i++)       // process current output sample
          y += h[i] * w[i];           // dot-product operation

      for (i=M; i>=1; i--)            // update states for next call
        w[i] = w[i-1];                // done in reverse order

      return y;
    }

    // ----------------------------------------------------------------
```

The function `firc` implements the circular buffer version using the pointer-wrapping function `pwrap`:

```
    // firc.c - FIR filter implemented with circular pointer
    // Usage: y = firc(M, h, w, &p, x);
    // ---------------------------------------------------------------------

    float *pwrap(int, float *, float *);            // defined in dsplab.c

    float firc(int M, float *h, float *w, float **p, float x)
    {
       int i;
       float y;

       **p = x;                       // read input sample x

       for (y=0, i=0; i<=M; i++) {
          y += (*h++) * (**p);        // i-th state s[i] = *pwrap(M,w,*p+i)
          *p = pwrap(M,w,++*p);       // pointer to state s[i+1] = *pwrap(M,*p+i+1)
          }

       *p = pwrap(M,w,--*p);          // update circular delay line

       return y;
    }

    // ---------------------------------------------------------------------
```

The function `firc2` is a circular buffer version using the pointer-index-wrapping function `qwrap`:

```
    // firc2.c - FIR filter implemented with circular index
    // Usage: y = firc2(M, h, w, &q, x);
    // ---------------------------------------------------------------------

    int qwrap(int, int);                            // defined in dsplab.c

    float firc2(int M, float *h, float *w, int *q, float x)
    {
       int i;
       float y;

       w[*q] = x;                     // read input sample x

       for (y=0, i=0; i<=M; i++) {
          y += *h++ * w[*q];          // i-th state s[i] = w[*q]
          *q = qwrap(M,++*q);         // pointer to state s[i+1] = w[*q+1]
          }

       *q = qwrap(M,--*q);            // update circular delay line

       return y;
    }

    // ---------------------------------------------------------------------
```

In both `firc` and `firc2`, the circular pointer or index are being wrapped during each pass through the for-loop that computes the output sample y. This is inefficient but necessary because C does not support circular arrays.

All modern DSP chips, including the C6713, support circular addressing in hardware, which does the required wrapping automatically without any extra instructions. The following function, `firq`, emulates the hardware version more closely by avoiding the repeated calls to `qwrap` inside the for-loop—it performs only one wrapping when it reaches the end of the buffer and wraps the index back to $q = 0$; furthermore, it wraps once more after the for-loops in order to backshift the pointer index:

```
// firq.c - FIR filter implemented with circular index
// Usage: y = fircq(M, h, w, &q, x);
// -----------------------------------------------------------------------

float firq(int M, float *h, float *w, int *q, float x)
{
   int i, Q;
   float y;

   Q = M - (*q);                        // number of states to end of buffer

   w[*q] = x;                           // read input sample x

   for (y=0, i=0; i<=Q; i++)            // loop from q to end of buffer
      y += h[i] * w[(*q)++];

   (*q) = 0;                            // wrap to beginning of buffer

   for (i=Q+1; i<=M; i++)               // loop to q-1
      y += h[i] * w[(*q)++];

   (*q)--;  if (*q == -1) *q = M;       // backshift index

   return y;
}

// -----------------------------------------------------------------------
```
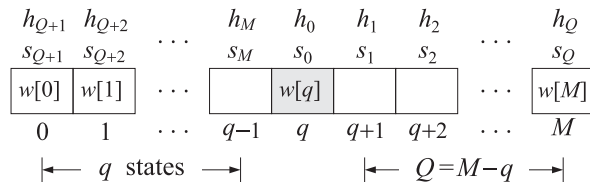
The for-loop is split into two parts, the first part starts at position q and loops until the end of the buffer, then it wraps to the beginning of the buffer; the second part loops till $q - 1$. The required states s_i of the FIR filter and their association with the filter coefficients h_i are depicted below.

Finally, we consider the *linear assembly* function, `fira.sa`, listed below, that exploits the hardware implementation of circular buffers on the C6713 processor. It is based on the function `convol1.sa` of Ref. [3], adapted here to our convention of counting the states and filter coefficients in forward order and updating the circular index by backshifting. Linear assembly functions have an extension `.sa` and may be included in a project just like C functions. The linear assembly optimizer determines which particular hardware registers to assign to the various local variables in the function.

```
; fira.sa - linear assembly version of FIR filter with circular buffer
;
; extern float fira(float *, float *, int, int, int);
;
; float w[Lw];
; #pragma DATA_ALIGN(w, Lb)
;
```

```
; usage: w[q] = x;                      read input sample
;         y = fira(w,h,Lh,Nb,q);        compute output sample
;         q--; if (q==-1) q = Lw-1;     update circular index by backshifting
;
; M                      = filter order
; Lh  = M+1              = filter length
; Nb >= 1 + ceil(log2(Lh))  = circular buffer bytes-length exponent
; Lb  = 2^(Nb+1)         = circular buffer length in bytes
; Lw  = Lb/4 = 2^(Nb-1)  = circular buffer in 32-bit words
; -------------------------------------------------------------------

        .global _fira
_fira   .cproc  w, h, Lh, Nb, q       ; function arguments
        .reg  Y, P, si, hi            ; local variables

     ADDAW   w, q, w      ; point to w[q] = x = current input
                          ; set up the circular buffer
     SHL Nb, 16, Nb       ; shift Nb to BK0 field
     set Nb, 8,8, Nb      ; set circular mode, BK0, B4
     MVC Nb, AMR          ; load mode into AMR

     ZERO  Y              ; output

loop: .trip 8, 500        ; assume between 8 and 500 taps

     LDW *w++, si         ; load i-th state, si = x(n-i)
     LDW *h++, hi         ; load i-th filter coeff, h(i)
     MPYSP  si,hi,P       ; multiply single precision, P = hi*si
     ADDSP P,Y,Y          ; Y = Y + P = accumulate output

  [Lh] SUB Lh, 1, Lh      ; decrement, Lh = Lh-1
  [Lh] B  loop            ; loop until Lh=0

     .return Y            ; put sum in A4 - C convention

     .endproc

; -------------------------------------------------------------------
```

Lab Procedure

A lowpass FIR filter of order M and cutoff frequency f_0 can be designed using the Hamming window approach by the following equations (see Ch.11 of [1]):

$$w(n) = 0.54 - 0.46 \cos\left(\frac{2\pi n}{M}\right), \quad h(n) = w(n) \frac{\sin(\omega_0(n - M/2))}{\pi(n - M/2)}, \quad 0 \le n \le M$$

where $\omega_0 = 2\pi f_0/f_s$, and $w(n)$ is the Hamming window.

a. Design such a filter with MATLAB using the following values: $f_s = 8$ kHz, $f_0 = 2$ kHz, and filter order $M = 100$. Then, using the built-in MATLAB function freqz, or the textbook function dtft, calculate and plot in dB the magnitude response of the filter over the frequency interval $0 \le f \le 4$ kHz. The designed filter response is shown in Fig. 3.4 in absolute units and in dB.

b. The designed 101-long impulse response coefficient vector **h** can be exported into a data file, h.dat, in a form that is readable by a C program by the following MATLAB command:

```
C_header('h.dat', 'h', 'M', h);
```

where C_header is a MATLAB function in the directory c:\dsplab\common. A few lines of the resulting data file are shown below:
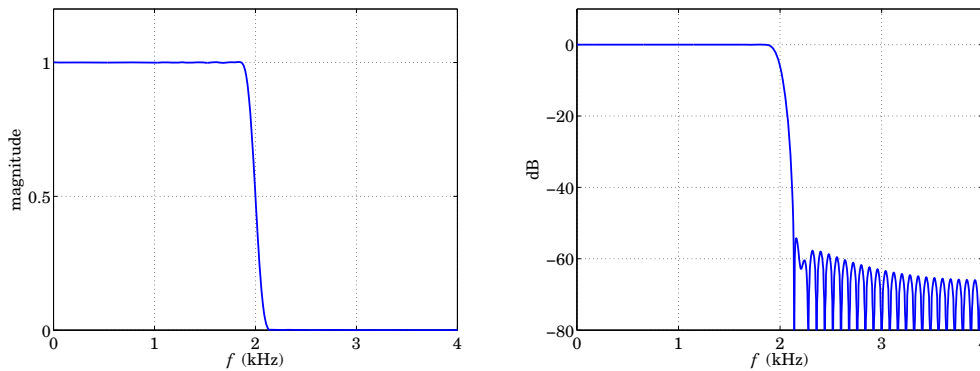
Fig. 3.4 Magnitude response of lowpass filter.

```
// h.dat  -  FIR impulse response coefficients
// exported from MATLAB using C_header.m
// ----------------------------------------

#define M 100          // filter order

float h[M+1] = {
   -0.000000000000000,
    0.000525586170315,
   -0.000000000000000,
   -0.000596985888698,
    0.000000000000000,
    0.000725138232653,
       --- etc. ---
   -0.000596985888698,
   -0.000000000000000,
    0.000525586170315,
   -0.000000000000000
    };

// ----------------------------------------
```

The following complete C program, `firex.c`, implements this example on the C6713 processor. The
program reads the impulse response vector from the data file `h.dat`, and defines a 101-dimensional
delay-line buffer array `w`. The FIR filtering operation is based on any of the choices, `fir`, `firc`,
`firc2`, `firq`, depending on which lines are uncommented.

```
// firex.c - FIR filtering example
// -----------------------------------------------------------------------------

#include "dsplab.h"          // DSK initialization declarations and function prototypes

//float fir(int, float *, float *, float);
//float firc(int, float *, float *, float **, float);
//float firc2(int, float *, float *, int *, float);
//float firq(int, float *, float *, int *, float);

short xL, xR, yL, yR;        // left and right input and output samples from/to codec

#include "h.dat"     // contains M+1 = 101 filter coefficients

float w[M+1];        // filter delay lines
int on = 1;          // turn filter on
//float *p;
//int q;

// -----------------------------------------------------------------------------
```

```
void main()
{
  int i;

  for (i=0; i<=M; i++) w[i] = 0;     // initialize delay-line buffer
  //p = w;                           // initialize circular pointer
  //q = 0;

  initialize();                 // initialize DSK board and codec, define interrupts

  sampling_rate(8);             // possible sampling rates: 8, 16, 24, 32, 44, 48, 96 kHz
  audio_source(LINE);           // LINE or MIC for line or microphone input

  while(1);                     // keep waiting for interrupt, then jump to isr()
}

// -----------------------------------------------------------------------------

interrupt void isr()
{
   float x, y;                          // filter input & output

   read_inputs(&xL, &xR);               // read audio samples from codec

   if (on) {
      x = (float)(xL);                  // work with left input only

      //y = fir(M,h,w,x);
      //y = firc(M,h,w,&p,x);
      //y = firc2(M, h, w, &q, x);
      //y = firq(M, h, w, &q, x);

      yL = (short)(y);
      }
   else                                 // pass through if filter is off
      yL = xL;

   write_outputs(yL,yL);                // write audio samples to codec

   return;
}

// -----------------------------------------------------------------------------
```

Create and build a project for this program. You will need to add one of the functions `fir, firc, firc2, firq` to the project. Using the following MATLAB code (same as in the aliasing example of Lab-1), generate a signal consisting of a 1-kHz segment, followed by a 3-kHz segment, followed by another 1-kHz segment, where all segments have duration of 1 sec:

```
fs = 8000; f1 = 1000; f2 = 3000; f3 = 1000;
L = 8000; n = (0:L-1);
A = 1/5;                          % adjust playback volume

x1 = A * cos(2*pi*n*f1/fs);
x2 = A * cos(2*pi*n*f2/fs);
x3 = A * cos(2*pi*n*f3/fs);

sound([x1,x2,x3], fs);
```

First, set the parameter `on=0` so that the filtering operation is bypassed. Send the above signal into the line input of the DSK and listen to the output. Then, set `on=1` to turn the filter on using the linear buffer version, `fir`, recompile and run the program, and send the same signal in. The middle 3-kHz segment should not be heard, since it lies in the filter's stopband.

c. Create breakpoints at the `read_inputs` and `write_outputs` lines of the `isr()` function, and start the profile clock. Run the program and record the number of cycles between reading the input samples and writing the computed outputs.

d. Uncomment the appropriate lines in the above program to implement the circular buffer versions using the functions `firc, firc2, firq`. You will need to add these to your project. Recompile and run your program with the same input.

   Then, repeat part (c) and record the number of computation cycles.

e. The compiler optimization thus far was set to "none". Using the keyboard combination "ALT-P P", or the CCS menu commands *Project -> Build Options*, change the optimization level to `-o0, -o1, -o2, -o3`, and for each level and each of the four filter implementations `fir, firc, firc2, firq`, repeat part (c) and record the number of cycles in a table form:

	none	-o0	-o1	-o2	-o3
fir					
firc					
firc2					
firq					
fira					

f. Add to the above table the results from the linear assembly version implemented by the following complete C program, `firexa.c`, and evaluate your results in terms of efficiency of implementation and optimization level.

```
// firexa.c - FIR filtering example using circular buffer with linear assembly
// --------------------------------------------------------------------------------

#include "dsplab.h"          // DSK initialization declarations and function prototypes

extern float fira(float *, float *, int, int, int);

short xL, xR, yL, yR;        // left and right input and output samples from/to codec

#include "h.dat"     // contains M+1 = 101 filter coefficients

#define Nb 8         // circular-buffer length (bytes) exponent, Nb = 1 + ceil(log2(M+1)) = 8
#define Lb 512       // circular-buffer length (bytes) = 2^(Nb+1)
#define Lw 128       // circular-buffer length (words) = 2^(Nb-1) = Lb/4
#define Lh 101       // filter length = M+1

float w[Lw];                 // circular buffer

#pragma DATA_ALIGN(w, Lb)    // align buffer at byte-boundary

int q;                       // circular-buffer index
int on = 1;                  // filter is ON or OFF

// --------------------------------------------------------------------------------

void main()
{
  int i;

  for (i=0; i<Lw; i++) w[i] = 0;    // initialize circular buffer to zero
  q = 0;                            // initialize index into buffer

  initialize();                // initialize DSK board and codec, define interrupts

  sampling_rate(8);            // possible sampling rates: 8, 16, 24, 32, 44, 48, 96 kHz
  audio_source(LINE);          // LINE or MIC for line or microphone input
```

```
   while(1);                    // keep waiting for interrupt, then jump to isr()
 }

 // ------------------------------------------------------------------------------

 interrupt void isr()
 {

   float x, y;                  // filter input & output

   read_inputs(&xL, &xR);

   if (on) {
     x = (float)(xL);           // work with left input only

     w[q] = x;                  // put x into w[q],

     y = fira(w, h, Lh, Nb, q); // fira does not update q

     q--;  if (q == -1) q = Lw-1;  // backshift to update q for next time instant

     yL = (short)(y);
     }
   else
     yL = xL;

   write_outputs(yL,yL);

   return;
 }

 // ------------------------------------------------------------------------------
```

3.5. Voice Scrambler

A simple voice scrambler works by spectrum inversion. It is not the most secure way of encrypting speech, but we consider it in this lab as an application of low pass filtering and AM modulation. The main operations are depicted below.
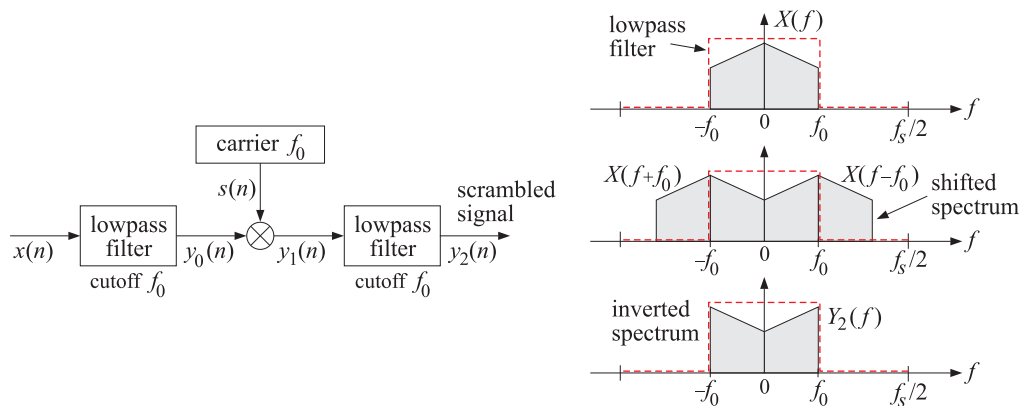
First, the sampled speech signal $x(n)$ is filtered by a lowpass filter $h(n)$ whose cutoff frequency f_0 is high enough not to cause distortions of the speech signal (the figure depicts an ideal filter). The sampling rate f_s is chosen such that $4f_0 < f_s$. The filtering operation can be represented by the convolutional equation:

$$y_0(n) = \sum_m h(m)x(n - m) \tag{3.2}$$

Next, the filter output $y_0(n)$ modulates a cosinusoidal carrier signal whose frequency coincides with

the filter's cutoff frequency f_0, resulting in the signal:

$$y_1(n) = s(n) y_0(n), \quad \text{where} \quad s(n) = 2\cos(\omega_0 n), \quad \omega_0 = \frac{2\pi f_0}{f_s} \tag{3.3}$$

The multiplication by the carrier signal causes the spectrum of the signal to be shifted and centered at $\pm f_0$, as shown above. Finally, the modulated signal $y_1(n)$ is passed through the same filter again which removes the spectral components with $|f| > f_0$, resulting in a signal $y_2(n)$ with inverted spectrum. The last filtering operation is:

$$y_2(n) = \sum_m h(m) y_1(n - m) \tag{3.4}$$

To unscramble the signal, one may apply the scrambling steps (3.2)–(3.4) to the scrambled signal itself. This works because the inverted spectrum will be inverted again, recovering in the original spectrum.

In this lab, you will study a real-time implementation of the above procedures. The lowpass filter will be designed with the parameters $f_s = 16$ kHz, $f_0 = 3.3$ kHz, and filter order $M = 100$ using the Hamming design method:

$$h(n) = w(n) \frac{\sin(\omega_0 (n - M/2))}{\pi(n - M/2)}, \quad 0 \le n \le M \tag{3.5}$$

where $\omega_0 = 2\pi f_0/f_s$, and $w(n)$ is the Hamming window:

$$w(n) = 0.54 - 0.46\cos\left(\frac{2\pi n}{M}\right), \quad 0 \le n \le M \tag{3.6}$$

The following C program, scrambler.c, forms the basis of this lab. It is a variation of that discussed in the Chassaing-Reay text [2].

```
// scrambler.c - voice scrambler example
// --------------------------------------------------------------------------------------

#include "dsplab.h"                   // initialization declarations and function prototypes
#include <math.h>
#define PI 3.14159265358979

short xL, xR, yL, yR;                 // left and right input and output samples from/to codec

#define M 100                         // filter order
#define L 160                         // carrier period, note L*f0/fs = 160*3.3/16 = 33 cycles

float h[M+1], w1[M+1], w2[M+1];       // filter coefficients and delay-line buffers
int n=0;                              // time index for carrier, repeats with period L
int on=1;                             // turn scrambler on (off with on=0)

float w0, f0 = 3.3;                   // f0 = 3.3 kHz
short fs = 16;                        // fs = 16 kHz

// --------------------------------------------------------------------------------------

void main()
{
  int i;
  float wind;

  w0 = 2*PI*f0/fs;                            // carrier frequency in rads/sample

  for (i=0; i<=M; i++)  {                     // initialize buffers & design filter
     w1[i] = w2[i] = 0;
     wind = 0.54 - 0.46 * cos(2*PI*i/M);      // Hamming window
     if (i==M/2)
        h[i] = w0/PI;
     else
        h[i] = wind * sin(w0*(i-M/2)) / (PI*(i-M/2));
     }
```

```
   initialize();               // initialize DSK board and codec, define interrupts

   sampling_rate(fs);          // possible sampling rates: 8, 16, 24, 32, 44, 48, 96 kHz
   audio_source(LINE);         // LINE or MIC for line or microphone input

   while(1);                   // keep waiting for interrupt, then jump to isr()
}

// ------------------------------------------------------------------------------------

interrupt void isr()           // sample processing algorithm - interrupt service routine
{
   int i;
   float y;

   read_inputs(&xL, &xR);      // read left and right input samples from codec

   if (on) {
      y = (float)(xL);                    // work with left input only

      w1[0] = y;                          // first filter
      for (y=0, i=0; i<=M; i++)
         y += h[i] * w1[i];
      delay(M,w1);

      y *= 2*cos(w0*n);                   // multiply y by carrier
      if (++n >= L) n = 0;


      w2[0] = y;                          // second filter
      for (y=0, i=0; i<=M; i++)
         y += h[i] * w2[i];
      delay(M,w2);

      yL = (short)(y);
      }
   else
      yL = xL;                            // pass through if on=0

   write_outputs(yL,yL);

   return;
}

// ------------------------------------------------------------------------------------
```

Two separate buffers, w_1, w_2, are used for the two lowpass filters. The filter coefficients are computed on the fly within `main()` using Eqs. (3.5) and (3.6). A linear buffer implementation is used for both filters. The sinusoidal carrier signal is defined by:

$$s[n] = 2\cos(\omega_0 n), \quad \omega_0 = \frac{2\pi f_0}{f_s}$$

Since $f_s/f_0 = 16/3.3$ samples/cycle, it follows that the smallest number of samples containing an integral number of cycles will be:

$$L = \frac{16}{3.3} \cdot 33 = 160 \text{ samples}$$

that is, these 160 samples contain 33 cycles and will keep repeating. Therefore, the time index n of $s[n]$ is periodically cycled over the interval $0 \le n \le L - 1$.

Lab Procedure

a. Explain why the factor 2 is needed in the carrier definition $s(n) = 2\cos(\omega_0 n)$. Explain why f_0 must

be chosen such that $4f_0 < f_s$ in designing the lowpass filter.

b. Create and build a project for this program. The parameter `on=1` turns the scrambler on or off. Create a GEL file for this parameter and open it when you run the program.

c. Play the following two wave files through program:

```
JB.wav
JBs.wav
```

When you play the second, which is a scrambled version of the first, it will get unscrambled.

d. Open MATLAB and generate three sinusoids of frequencies 300 Hz, 3000 Hz, and 300 Hz, sampled at a rate of 16 kHz, each of duration of 1 second, and concatenate them to form a 3-second signal. Then play this out of the PCs sound card using the `sound()` function. For example, the following MATLAB code will do this:

```
fs = 16000; f1 = 300; f2 = 3000; f3 = 300;
L = 16000; n = (0:L-1);
A = 1/5;                              % adjust playback volume

x1 = A * cos(2*pi*n*f1/fs);
x2 = A * cos(2*pi*n*f2/fs);
x3 = A * cos(2*pi*n*f3/fs);

sound([x1,x2,x3], fs);
```

Play this signal through the DSK with the scrambler off. Then, play it with the scrambler on. What are the frequencies in Hz of the scrambled signal that you hear? Explain this in your report.

e. Instead of actually computing the cosine function at each call of `isr()`, a more efficient approach would be to pre-compute the L repeating samples of the carrier $s[n]$ and keep re-using them. This can be accomplished by replacing the two modulation instructions in `isr()` by:

```
y *= s[n];                  // multiply y by carrier
if (++n >= L) n = 0;
```

where $s[n]$ must be initialized within `main()` to the L values, $s[n] = 2\cos(\omega_0 n)$, $n = 0, 1, \ldots, L - 1$.

Re-write the above program to take advantage of this suggestion. Test your program.

In Lab-4, we will reconsider the scrambler and implement the required spectrum inversion using the FFT.

3.6. References

[1] S. J. Orfanidis, *Introduction to Signal Processing*, online book, 2010, available from: http://www.ece.rutgers.edu/~orfanidi/intro2sp/

[2] R. Chassaing and D. Reay, *Digital Signal Processing and Applications with the TMS320C6713 and TMS320C6416 DSK*, 2nd ed., Wiley, Hoboken, NJ, 2008.

[3] S. A. Tretter, *Communication System Design Using DSP Algorithms with Laboratory Experiments for the TMS320C6713 DSK*, Springer, New York, 2008, code available from: http://www.ece.umd.edu/~tretter

Lab 4 – Sampling, Aliasing, FIR Filtering

This is a software lab. In your report, please include all Matlab code, numerical results, plots, and your explanations of the theoretical questions. The due date is *one week* from assignment.

4.1. Sampling and Aliasing – Sinusoids

The aim of this lab is to demonstrate the effects of aliasing arising from improper sampling. A given analog signal $x(t)$ is sampled at a rate f_s, the resulting samples $x(nT)$ are then reconstructed by an *ideal* reconstructor into the analog signal $x_a(t)$. Improper choice of f_s will result in a different signal, $x_a(t) \neq x(t)$, even though the two agree at their sample values, that is, $x_a(nT) = x(nT)$. The procedure is illustrated in the following figure:
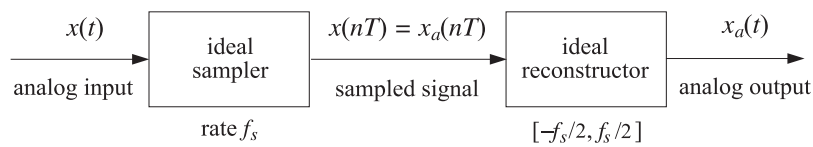
Lab Procedure

a. Consider an analog signal $x(t)$ consisting of three sinusoids of frequencies of 1 kHz, 4 kHz, and 6 kHz:

$$x(t) = \sin(2\pi t) + 2\sin(8\pi t) + 3\sin(12\pi t)$$

where t is in milliseconds. Show that if this signal is sampled at a rate of $f_s = 5$ kHz, it will be aliased with the following signal, in the sense that their sample values will be the same:

$$x_a(t) = 2\sin(2\pi t)$$

On the same graph, plot the two signals $x(t)$ and $x_a(t)$ versus t in the range $0 \le t \le 2$ msec. To this plot, add the time samples $x(t_n)$ and verify that $x(t)$ and $x_a(t)$ intersect precisely at these samples.

b. Repeat part (a) with $f_s = 10$ kHz. In this case, determine the signal $x_a(t)$ with which $x(t)$ is aliased. Plot both $x(t)$ and $x_a(t)$ on the same graph over the same range $0 \le t \le 2$ msec. Verify again that the two signals agree at the sampling instants, $x_a(nT) = x(nT)$. See example graphs at the end.

4.2. Sampling and Aliasing – Square Wave

Consider a periodic pulse wave $x(t)$ with period $T_0 = 1$ sec, as shown below. Let $p(t)$ denote one basic period of $x(t)$ defined over the time interval $0 \le t \le 1$:

$$p(t) = \begin{cases} 1, & \text{if } 0.125 < t < 0.375 \\ -1, & \text{if } 0.625 < t < 0.875 \\ 0.5, & \text{if } t = 0.125 \quad \text{or} \quad t = 0.375 \\ -0.5, & \text{if } t = 0.625 \quad \text{or} \quad t = 0.875 \\ 0, & \text{otherwise} \end{cases} \tag{4.1}$$

This periodic signal admits a Fourier series expansion containing only sine terms with odd harmonics of the basic period $f_0 = 1/T_0 = 1$ Hz, that is, the frequencies $f_m = mf_0, m = 1, 3, 5, \ldots$ Hz:

$$x(t) = \sum_{m=1,3,5,\ldots} b_m \sin(2\pi mt) = b_1 \sin(2\pi t) + b_3 \sin(6\pi t) + b_5 \sin(10\pi t) + \cdots \tag{4.2}$$

The Fourier series coefficients are given as follows, for $m = 1, 3, 5, 7, \ldots$

$$b_m = \frac{\cos(\pi m/4) - \cos(3\pi m/4) - \cos(5\pi m/4) + \cos(7\pi m/4)}{\pi m}$$

The reason why the signal $x(t)$ was defined to have the values ± 0.5 at the discontinuity points is a consequence of a theorem that states that any finite sum of Fourier series terms will always pass through the mid-points of discontinuities.
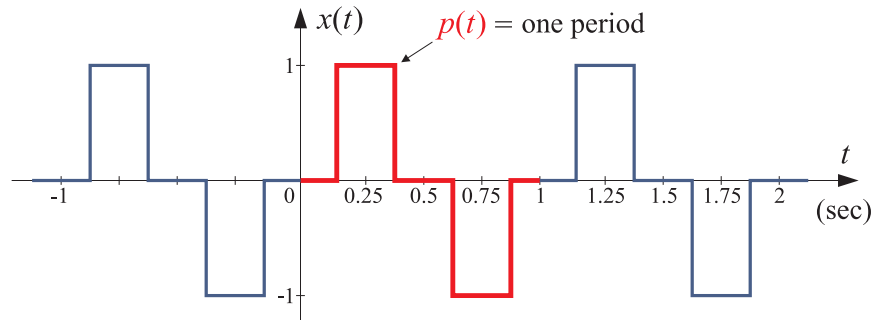
Lab Procedure

a. Define the function of Eq. (4.1) in MATLAB using a one-line anonymous function definition of the form:

```
p = @(t) ...          % one period of the square wave
```

using vectorized relational operations, such as, (0.125<t & t<0.375).

b. To understand the nature of the approximation of the square wave by the Fourier series sum, truncate the sum to a finite number of terms, that is, with M odd,

$$x_M(t) = \sum_{m=1,3,5,\dots}^{M} b_m \sin(2\pi m t) = b_1 \sin(2\pi t) + b_3 \sin(6\pi t) + \cdots + b_M \sin(2\pi M t) \tag{4.3}$$

Evaluate and plot $x(t)$ and $x_M(t)$ over one period $0 \le t \le 1$, for $M = 21$ and $M = 41$.

c. The pulse waveform $x(t)$ is now sampled at the rate of $f_s = 8$ Hz and the resulting samples $x(nT)$ are reconstructed by an *ideal* reconstructor resulting into the aliased analog signal $x_a(t)$.

The spectrum of the sampled signal consists of the periodic replication of the harmonics of $x(t)$ at multiples of f_s. Because $f_s/f_0 = 8$ is an even integer, all the odd harmonics that lie outside the Nyquist interval, $[-4, 4]$ Hz, will be wrapped onto the odd harmonics that lie inside this interval, that is, onto $\pm 1, \pm 3$ Hz. This can be verified by listing a few of the odd harmonics of $x(t)$ and the corresponding wrapped ones modulo f_s that lie within the Nyquist interval:

```
1   3   5   7   9   11   13   15   17   19   21   23   25   27   29 ...
1   3  -3  -1   1    3   -3   -1    1    3   -3   -1    1    3   -3 ...
```

where the bottom row is obtained by subtracting enough multiples of $f_s = 8$ from each harmonic until it is brought to lie within the interval $[-4, 4]$ Hz. This means then that the aliased signal will consist only of sinusoids of frequencies $f_1 = 1$ and $f_3 = 3$ Hz,

$$x_a(t) = A \sin(2\pi t) + B \sin(6\pi t) \tag{4.4}$$

Determine the coefficients A, B by setting up two equations in the two unknowns A, B by enforcing the matching equations $x_a(nT) = x(nT)$ at the two sampling instants $n = 1, 2$.

On the same graph, plot one period of the pulse wave $x(t)$ together with $x_a(t)$. Verify that they agree at the eight sampling time instants that lie within this period. Because of the sharp transitions of the square wave, you must use a very dense time vector, for example,

```
t = linspace(0,1,4097);
```

Also, if you wish, you may do part (c) and part (d), as special cases of part (e).

d. Assume, next, that the pulse waveform $x(t)$ is sampled at the rate of $f_s = 16$ Hz. By considering how the out-of band harmonics wrap into the Nyquist interval $[-8, 8]$ Hz, show that now the aliased signal $x_a(t)$ will have the form:

$$x_a(t) = a_1 \sin(2\pi t) + a_2 \sin(6\pi t) + a_3 \sin(10\pi t) + a_4 \sin(14\pi t)$$

where the coefficients a_i are obtained by the condition that the signals $x(t)$ and $x_a(t)$ agree at the first four sampling instants $t_n = nT = n/16$ Hz, for $n = 1, 2, 3, 4$. These four conditions can be arranged into a 4×4 matrix equation of the form:

$$\begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{bmatrix} = \begin{bmatrix} * \\ * \\ * \\ * \end{bmatrix}$$

Determine the numerical values of the starred entries. Then, using MATLAB, solve this matrix equation for the coefficients a_i. Once a_i are known, the signal $x_a(t)$ is completely defined.

On the same graph, plot one period of the pulse waveform $x(t)$ together with $x_a(t)$. Verify that they agree at the 16 sampling time instants that lie within this period.

e. The methods of parts (c,d) can be generalized to any sampling rate f_s such that $L = f_s/f_0$ is an even integer (so that all the out-of-band odd harmonics will wrap onto the odd harmonics within the Nyquist interval). First show that the number of odd harmonics within the positive side of the Nyquist interval is:

$$K = \text{floor}\left(\frac{L + 2}{4}\right)$$

This means that the aliased signal will be the sum of K terms:

$$x_a(t) = \sum_{k=1}^{K} a_k \sin(2\pi(2k - 1)t) \tag{4.5}$$

By matching $x_a(t)$ to $x(t)$, or $p(t)$, at the first K sampling instants $n = 1, 2, \ldots, K$, set up a linear system of K equations in the K unknowns a_k, i.e., with $t_n = nT$,

$$\sum_{k=1}^{K} a_k \sin(2\pi(2k - 1)t_n) = p(t_n), \quad n = 1, 2, \ldots, K \tag{4.6}$$

and solve it with Matlab. Once you have the coefficients a_k, evaluate and plot $x(t)$ and $x_a(t)$, and add the sampled points on the graph. Repeat this for the following eight, progressively larger, sampling rates:

$$f_s = [4, 8, 16, 24, 32, 40, 48, 64]$$

It should be evident that even though the square wave is not a bandlimited signal, it can still be sampled adequately if the sampling rate is chosen to be large enough.

4.3. FIR Filtering

The objective of this lab is to implement your own version of the built-in function **filter** adapted to FIR filters. The IIR case will be considered in a future lab. The documentation for **filter** states that it is implemented using the transposed block diagram realization. For example, for an order-3 FIR filter that realization and the system of difference equations implementing it are:

$$y(n) = h_0 x(n) + v_1(n)$$
$$v_1(n+1) = h_1 x(n) + v_2(n)$$
$$v_2(n+1) = h_2 x(n) + v_3(n)$$
$$v_3(n+1) = h_3 x(n)$$
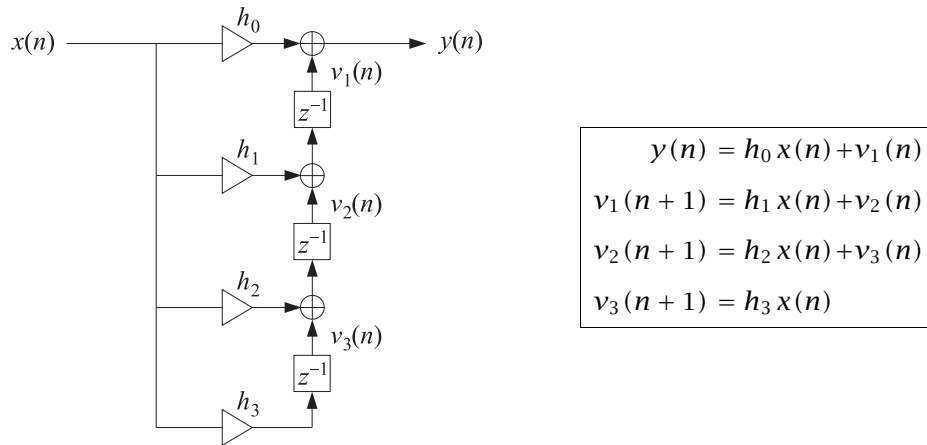
For an Mth order filter, the computational algorithm is,

$$y(n) = h_0 x(n) + v_1(n)$$
$$v_1(n+1) = h_1 x(n) + v_2(n)$$
$$v_2(n+1) = h_2 x(n) + v_3(n)$$
$$\vdots$$
$$v_{M-1}(n+1) = h_{M-1} x(n) + v_M(n)$$
$$v_M(n+1) = h_M x(n)$$

$$\mathbf{v}(n) = \begin{bmatrix} v_1(n) \\ v_2(n) \\ \vdots \\ v_M(n) \end{bmatrix} = \text{state vector}$$

The state vector $\mathbf{v}(n)$ represents the current contents of the M delays. The algorithm uses the current state $\mathbf{v}(n)$ to compute the current output $y(n)$ from the current input $x(n)$, and then, it updates the state to the next time instant, $\mathbf{v}(n+1)$. Usually, the state vector is initialized to zero, but it can be initialized to an arbitrary vector, say, $\mathbf{v}_{\text{init}}$.

Lab Procedure

a. Write a MATLAB function, say, **firtr.m**, that implements the above algorithm and has the possible syntaxes:

```
y = firtr(h,x);
[y,vout] = firtr(h,x,vin);

%   h   = (M+1)-dimensional filter vector (row or column)
%   x   = length-N vector of input samples (row or column)
%   y   = length-N vector of output samples (row or column)
% vin   = M-dimensional vector of initial states - zero vector, by default
% vout  = M-dimensional final state vector, i.e., final contents of delays
```

Test your function with the following case:

```
x = [1, 1, 2, 1, 2, 2, 1, 1];
h = [1, 2, -1, 1];
y = [1, 3, 3, 5, 3, 7, 4, 3]    % expected result
```

b. The function **firtr** can be run on a sample by sample basis to generate the successive internal state vectors, for example, using a loop such as,

```
v = zeros(1,M);                         % initial state vector
for n=1:length(x)
   [y(n),vout] = firtr(h,x(n),v);       % recycled state vector v
   v = vout;                            % next state
end
```

Add appropriate **fprintf** commands before, within, and after this loop to generate the following table of values for the above example,

```
n    x    y    v1    v2    v3
---------------------------
0    1    1    0     0     0
1    1    3    2    -1     1
2    2    3    1     0     1
3    1    5    4    -1     2
4    2    3    1     1     1
5    2    7    5    -1     2
6    1    4    3     0     2
7    1    3    2     1     1
8    -    -    3     0     1
```

where v_1, v_2, v_3 are the internal states at each time instant.

c. Write a function, **myconv.m**, that uses the above function **firtr** to implement the convolution of two vectors h, x. It should be functionally equivalent to the built-in function **conv** and have usage,

```
% y = myconv(h,x);
%
% h  = (M+1)-dimensional filter vector (row or column)
% x  = length-N vector of input samples (row or column)
% y  = length-(N+M) vector of output samples (row or column)
```

Test it on the following case:

```
x = [1, 1, 2, 1, 2, 2, 1, 1];
h = [1, 2, -1, 1];
y = [1, 3, 3, 5, 3, 7, 4, 3, 3, 0, 1]    % expected result
```

4.4. *Filtering of Noisy Signals*

A length-N signal $x(n)$ is the sum of a desired signal $s(n)$ and interference $v(n)$:

$$x(n) = s(n) + v(n), \quad 0 \le n \le N - 1$$

where

$$s(n) = \sin(\omega_0 n), \quad v(n) = \sin(\omega_1 n) + \sin(\omega_2 n), \quad 0 \le n \le N - 1$$

with

$$\omega_1 = 0.1\pi, \quad \omega_0 = 0.2\pi, \quad \omega_2 = 0.3\pi \quad [\text{radians/sample}]$$

In order to remove $v(n)$, the signal $x(n)$ is filtered through a bandpass FIR filter that is designed to pass the frequency ω_0 and reject the interfering frequencies ω_1, ω_2. An example of such a filter of order

$M = 150$ can be designed with the Fourier series method using a Hamming window, and has impulse response:

$$h(n) = w(n) \left[\frac{\sin(\omega_b(n - M/2)) - \sin(\omega_a(n - M/2))}{\pi(n - M/2)} \right], \quad 0 \leq n \leq M$$

where $\omega_a = 0.15\pi$, $\omega_b = 0.25\pi$, and $w(n)$ is the Hamming window:

$$w(n) = 0.54 - 0.46 \cos\left(\frac{2\pi n}{M}\right), \quad 0 \leq n \leq M$$

It has an effective passband $[\omega_a, \omega_b] = [0.15\pi, 0.25\pi]$. To avoid a computational issue at $n = M/2$, you may use MATLAB's built-in function **sinc**, which is defined as follows:

$$\operatorname{sinc}(x) = \frac{\sin(\pi x)}{\pi x}$$

Lab Procedure

a. Let $N = 200$. On the same graph plot $x(n)$ and $s(n)$ versus n over the interval $0 \leq n \leq N - 1$.

b. Filter $x(n)$ through the filter $h(n)$ using your function **firtr**, and plot the filtered output $y(n)$, together with $s(n)$, for $0 \leq n \leq N - 1$. Apart from an overall delay introduced by the filter, $y(n)$ should resemble $s(n)$ after the M initial transients.

c. To see what happened to the interference, filter the signal $v(n)$ separately through the filter and plot the output, on the same graph with $v(n)$ itself.

d. Using the built-in MATLAB function **freqz** calculate and plot the magnitude response of the filter over the frequency interval $0 \leq \omega \leq 0.4\pi$:
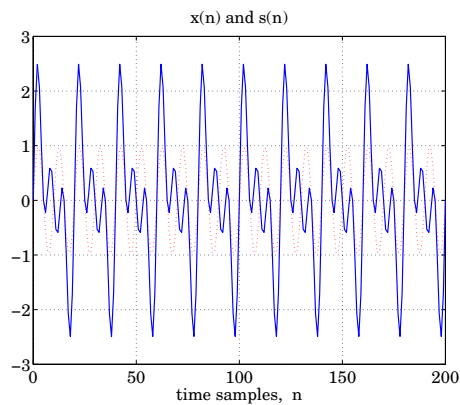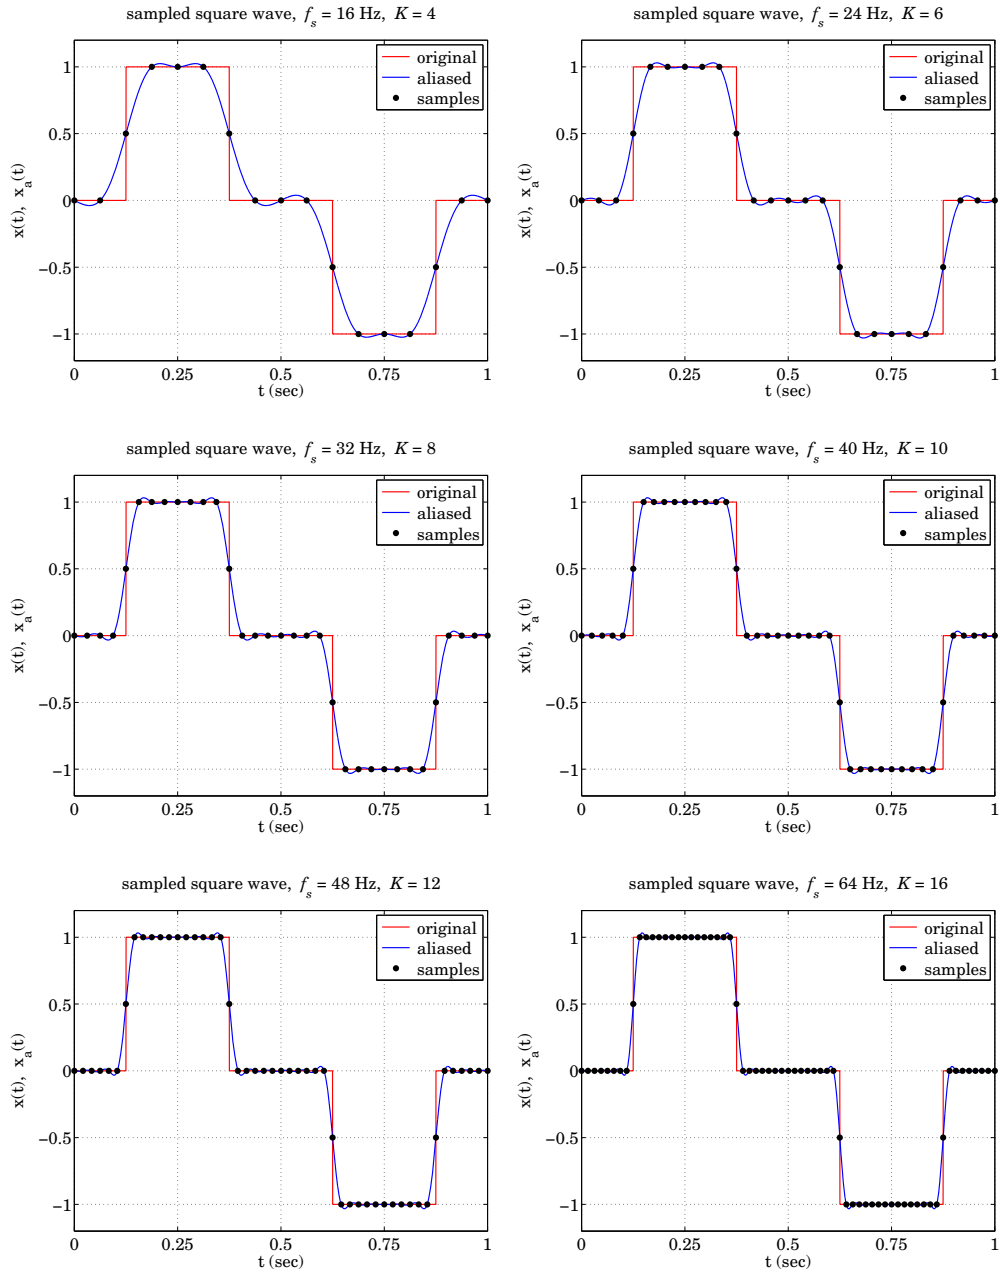
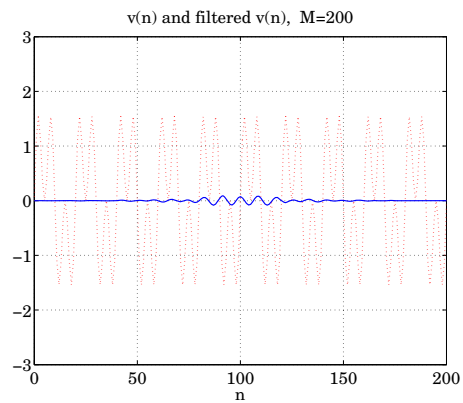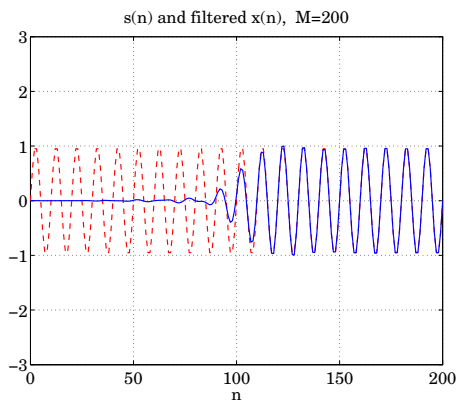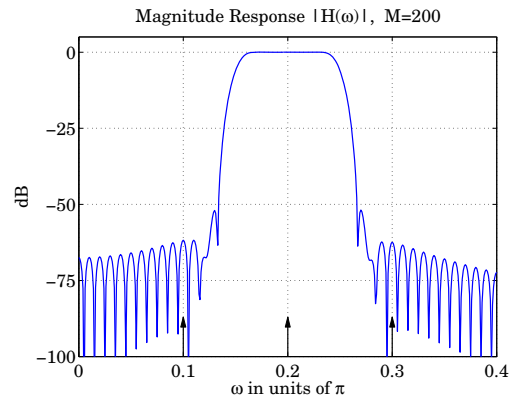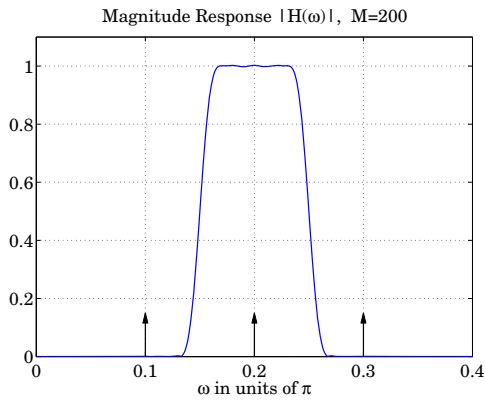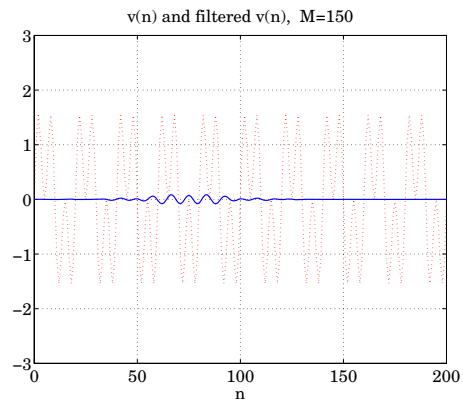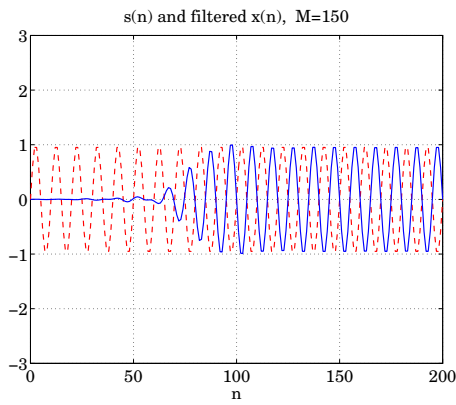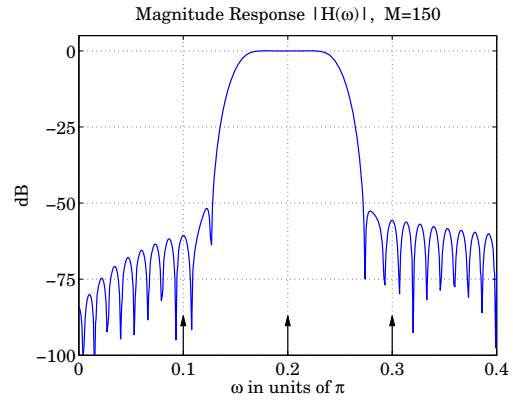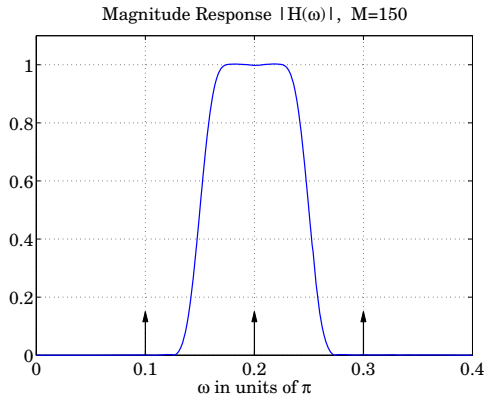$$|H(\omega)| = \left| \sum_{n=0}^{M} h(n) e^{-j\omega n} \right|$$

Indicate on that graph the frequencies $\omega_1, \omega_0, \omega_2$. Repeat the plot of $|H(\omega)|$ in dB units.

e. Redesign the filter with $M = 200$ and repeat parts (a)–(d). Discuss the effect of choosing a longer filter length.

Example Graphs

sampled square wave, $f_s = 16$ Hz, $K = 4$

sampled square wave, $f_s = 24$ Hz, $K = 6$

sampled square wave, $f_s = 32$ Hz, $K = 8$

sampled square wave, $f_s = 40$ Hz, $K = 10$

sampled square wave, $f_s = 48$ Hz, $K = 12$

sampled square wave, $f_s = 64$ Hz, $K = 16$

$x(n)$ and $s(n)$

Magnitude Response |H(ω)|, M=150

Magnitude Response |H(ω)|, M=150

s(n) and filtered x(n), M=150

v(n) and filtered v(n), M=150

Magnitude Response |H(ω)|, M=200

Magnitude Response |H(ω)|, M=200

s(n) and filtered x(n), M=200
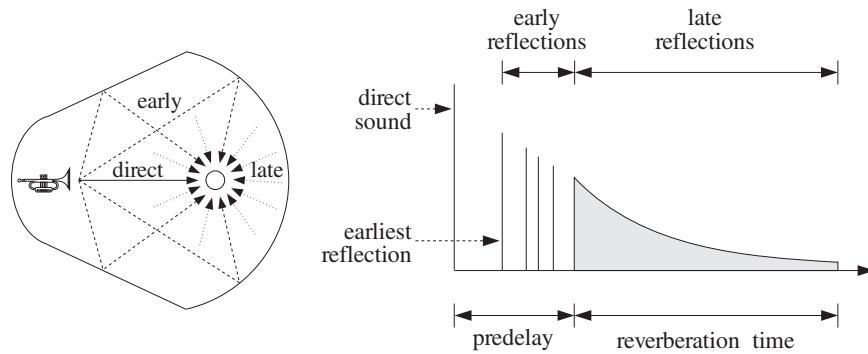
v(n) and filtered v(n), M=200

Lab 5 – Digital Audio Effects

5.1. Plain Reverb

The reverberation of a listening space is typically characterized by three distinct time periods: the direct sound, the early reflections, and the late reflections, as illustrated below:


The early reflections correspond to the first few reflections off the walls of the room. As the waves continue to bounce off the walls, their density increases and they disperse, arriving at the listener from all directions. This is the late reflection part.

The reverberation time constant is the time it takes for the room's impulse response to decay by 60 dB. Typical concert halls have time constants of about 1.8–2 seconds.

In this and several other labs, we discuss how to emulate such reverberation characteristics using DSP filtering algorithms. A *plain reverberator* can be used as an elementary building block for more complicated reverberation algorithms. It is given by Eq. (8.2.12) of the text [1] and shown in Fig. 8.2.6. Its input/output equation and transfer function are:

$$y(n) = ay(n - D) + x(n), \qquad H(z) = \frac{1}{1 - az^{-D}}$$

The comb-like structure of its frequency response and its pole-pattern on the z-plane are depicted in Fig. 8.2.7 of Ref. [1] and shown below.
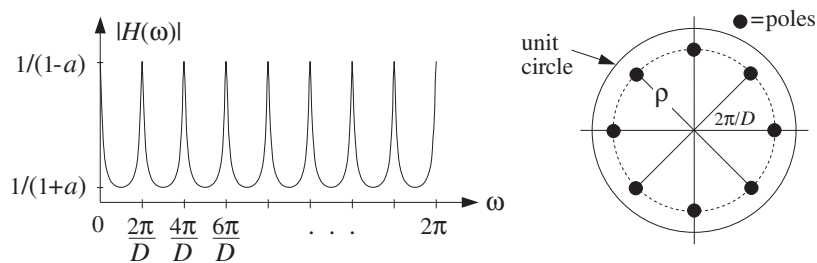
Its sample processing algorithm using a circular delay-line buffer is given by Eq. (8.2.14) of [1]:

for each input sample x do:
$$s_D = *(p + D)$$
$$y = x + a\,s_D$$
$$*p = y$$
$$--p$$
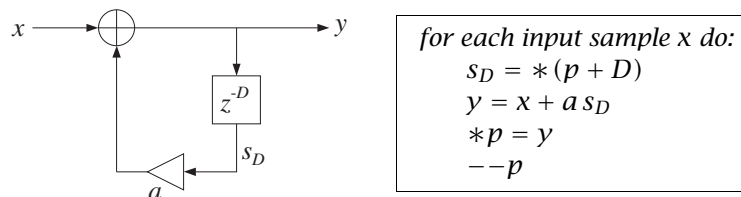
It can be immediately translated to C code with the help of the function `pwrap()` and embedded in the interrupt service routine `isr()`:

```
interrupt void isr()
{
    float sD, x, y;              // D-th state, input & output

    read_inputs(&xL, &xR);       // read inputs from codec

    x = (float) xL;              // process left channel only

    sD = *pwrap(D,w,p+D);        // extract D-th state relative to p
    y = x + a*sD;                // compute output sample
    *p = y;                      // delay-line input
    p = pwrap(D,w,--p);          // backshift pointer

    yL = yR = (short) y;

    write_outputs(yL,yR);        // write outputs to codec

    return;
}
```

Lab Procedure

a. Modify the template program into a C program. `plain1.c`, that implements the above ISR. Set the sampling rate to 8 kHz and the audio input to MIC. With the values of the parameters $D = 2500$ and $a = 0.5$, compile and run your program on the DSK.

   Listen to the impulse response of the system by lightly tapping the microphone on the table. Speak into the mike.

   Set the audio input to LINE, recompile and run. Play one of the wave files in the directory `c:\dsplab\wav` (e.g., `dsummer`, `noflange` from [3]).

b. Recompile and run the program with the new feedback coefficient $a = 0.25$. Listen to the impulse response. Repeat for $a = 0.75$. Discuss the effect of increasing or decreasing a.

c. According to Eq. (8.2.16), the effective reverberation time constant is given by

$$\tau_{\text{eff}} = \frac{\ln \epsilon}{\ln a} T_D, \qquad T_D = DT = D/f_s$$

   For each of the above values of a, calculate τ_{eff} in seconds, assuming $\epsilon = 0.001$ (which corresponds to the so-called 60-dB time constant.) Is what you hear consistent with this expression?

d. According to this formula, τ_{eff} remains invariant under the replacements:

$$D \to 2D, \qquad a \to a^2$$

   Test if this is true by running your program and hearing the output with $D = 5000$ and $a = 0.5^2 = 0.25$ and comparing it with the case $D = 2500$ and $a = 0.5$. Repeat the comparison also with $D = 1250$ and $a = \sqrt{0.5} = 0.7071$.

e. When the filter parameter a is positive and near unity, the comb peak gains $1/(1 - a)$ become large, and may cause overflows. In such cases, the input must be appropriately scaled down before it is passed to the filter.

   To hear such overflow effects, choose the feedback coefficients to be very near unity, for example, $a = 0.99$, with a corresponding gain of $(1 - a)^{-1} = 100$. You may also need to multiply the input x by an additional gain factor such as 2 or 4.

f. Modify the above ISR so that it processes the input samples in stereo (you will need to define two separate buffers for the left and right channels.) Experiment with choosing slightly different values of the left and right delay parameters D, or different values of the feedback parameter a. Keep the left/right speakers as far separated as possible.

5.2. Allpass Reverb

Like the plain reverberator, an allpass reverberator can be used as an elementary building block for building more complicated reverberation algorithms. It is given by Eq. (8.2.25) of the text [1] and shown in Fig. 8.2.17. Its I/O equation and transfer function are:

$$y(n) = ay(n-D) - ax(n) + x(n-D), \qquad H(z) = \frac{-a + z^{-D}}{1 - az^{-D}}$$

As discussed in [1], its impulse response is similar to that of the plain reverberator, but its magnitude response remains unity (hence the name "allpass"), that is,

$$\left| H(e^{j\omega}) \right| = 1, \quad \text{for all } \omega$$

Its block diagram representation using the so-called canonical realization and the corresponding sample processing algorithm using a circular delay-line buffer is given by Eq. (8.2.14) of [1]:
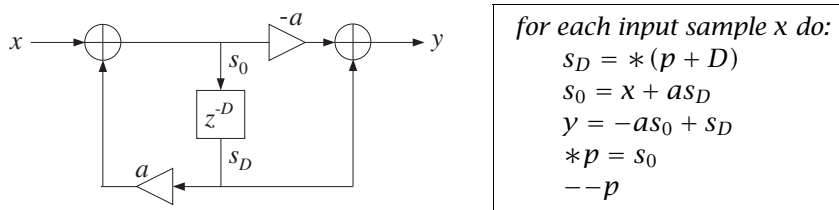
The algorithm can be translated immediately to C with the help of `pwrap()`. In this lab, we are going to put these steps into a separate C function, `allpass()`, which is to be called by `isr()`, and linked to the overall project. The function is defined as follows:

```c
// --------------------------------------------------------------------------
// allpass.c - allpass reverb with circular delay line - canonical realization
// --------------------------------------------------------------------------

float *pwrap(int, float *, float *);

float allpass(int D, float *w, float **p, float a, float x)
{
    float y, s0, sD;

    sD = *pwrap(D,w,*p+D);

    s0 = x + a * sD;

    y  = -a * s0 + sD;

    **p = s0;

    *p = pwrap(D,w,--*p);

    return y;
}
// --------------------------------------------------------------------------
```

The `allpass` function is essentially the same as that in the text [1], but slightly modified to use floats and the function `pwrap()`. In the above definition, the parameter p was declared as pointer to pointer to float because in the calling ISR function p must be defined as a pointer to float and must be passed passed by address because it keeps changing from call to call. The calling ISR function `isr()` is defined as follows:

```
interrupt void isr()
{
   float x, y;

   read_inputs(&xL, &xR);        // read inputs from codec

   x = (float) xL;               // process left channel only

   y = allpass(D,w,&p,a,x);      // to be linked with main()

   yL = yR = (short) y;

   write_outputs(yL,yR);         // write outputs to codec

   return;
}
```

Although the overall frequency response of the allpass reverberator is unity, the intermediate stage of computing the recursive part s_0 can overflow because this part is just like the plain reverb and its peak gain is $1/(1-a)$. Such overflow behavior is a potential problem of canonical realizations and we will investigate it further in a future lab.

The allpass reverberator can also be implemented in its transposed realization form, which is less prone to overflows. It is depicted below together with its sample processing algorithm:

for each input x do:
$$s_D = *(p + D)$$
$$y = s_D - ax$$
$$*p = s_0 = x + ay$$
$$--p$$
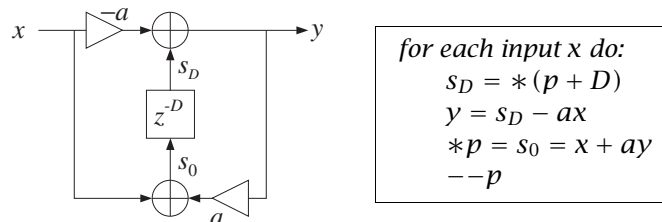
The following function `allpass_tr()` is the translation into C using `pwrap()`, where again p is defined as a pointer to pointer to float:

```
// ------------------------------------------------------------------------------
// allpass_tr.c - allpass reverb with circular delay line - transposed realization
// ------------------------------------------------------------------------------

float *pwrap(int, float *, float *);                         // defined in dsplab.c

float allpass_tr(int D, float *w, float **p, float a, float x)
{
   float y, sD;

   sD = *pwrap(D,w,*p+D);

   y  = sD - a*x;

   **p = x + a*y;

   *p = pwrap(D,w,--*p);

   return y;
}
// ------------------------------------------------------------------------------
```

Lab Procedure

a. Incorporate the above ISR into a main program, `allpass1.c`, and create a project. Remember to prototype the `allpass` function at the beginning of your program. Add the file that contains the `allpass` function to the project. Compile and run with the parameter choices: $D = 2500$, $a = 0.5$, with an 8 kHz sampling rate and LINE input.

b. Repeat part (a) using the transposed form implemented by the function `allpass_tr()`, and name your main program `allpass2.c`.

c. Choose a value of *a* and input gain that causes `allpass1.c` to overflow, then run `allpass2.c` with the same parameter values to see if your are still getting overflows.
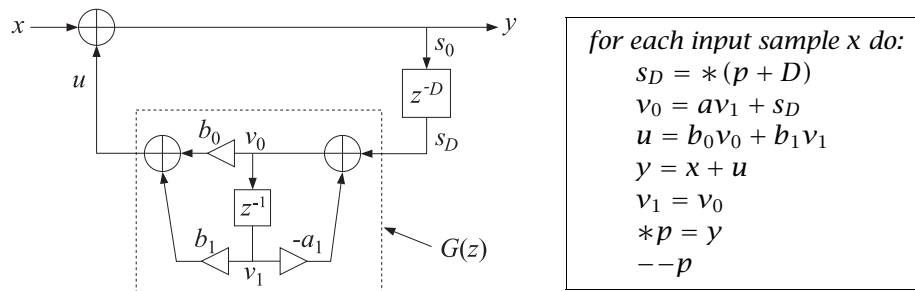
5.3. Lowpass Reverb

The lowpass reverberator of this experiment is shown in Figs. 8.2.20 and 8.2.21 of Ref. [1]. The feedback gain *a* of the plain reverb is replaced by a lowpass filter $G(z)$, so that one obtains the new transfer function by the replacement:

$$H(z) = \frac{1}{1 - az^{-D}} \quad \Rightarrow \quad H(z) = \frac{1}{1 - G(z)z^{-D}}$$

The filter $G(z)$ effectively acts as frequency-dependent feedback parameter whose value is smaller at higher frequencies (because it is a lowpass filter), thus attenuating high frequencies faster, and whose value is larger at lower frequencies, and hence attenuating those more slowly—which is a more realistic behavior of reverberating spaces. For this experiment, we will work with the simple choice:

$$G(z) = \frac{b_0 + b_1 z^{-1}}{1 + a_1 z^{-1}}$$

Setting $a = -a_1$, the corresponding sample processing algorithm is:


```
for each input sample x do:
    s_D = *(p + D)
    v_0 = av_1 + s_D
    u = b_0 v_0 + b_1 v_1
    y = x + u
    v_1 = v_0
    *p = y
    --p
```

The following is its C translation into the `isr()` function:

```
interrupt void isr()
{
   float x, y, sD, u;

   read_inputs(&xL, &xR);      // read inputs from codec

   x = (float) xL;             // process left channel only

   sD = *pwrap(D,w,p+D);

   v0 = a*v1 + sD;             // feedback filter G(z) = (b0 + b1*z^-1)/(1-a*z^-1)
   u = b0*v0 + b1*v1;          // feedback filter's output
   v1 = v0;                    // update feedback filter's delay

   y = x+u;                    // closed-loop output

   *p = y;

   p = pwrap(D,w,--p);

   yL = yR = (short) y;

   write_outputs(yL,yR);       // write outputs to codec

   return;
}
```

Lab Procedure

a. Create a project with this ISR. Choose an 8 kHz sampling rate and MIC input. Set the parameter values $D = 2500$, $a = 0.5$, $b_0 = 0.2$, $b_1 = 0.1$. Compile and run. Listen to its impulse response. Speak into the mike. Notice how successive echoes get more and more mellow as they circulate through the lowpass filter. Note that the DC gain of the loop filter $G(z)$, obtained by setting $z = 1$, and the AC gain at Nyquist, obtained by setting $z = -1$, are:

$$G(z)\big|_{z=1} = \frac{b_0 + b_1}{1 - a} = 0.6\,, \qquad G(z)\big|_{z=-1} = \frac{b_0 - b_1}{1 + a} = \frac{1}{15} = 0.0667$$

These are the effective feedback coefficients at low and high frequencies. Therefore, the lower frequencies persist longer than the higher ones.

Recompile and run with LINE input and play a wave file (e.g., `noflange`) through it.

b. Try the case $D = 20$, $a = 0$, $b_0 = b_1 = 0.495$. You will hear a guitar-like sound. Repeat for $D = 100$. What do you hear?

Repeat by setting the sampling rate to 44.1 kHz and $D = 100$.

This type of feedback filter is the basis of the so-called Karplus-Strong string algorithm for synthesizing plucked-string sounds, and we will study it further in another experiment.

5.4. Schroeder's Reverb Algorithm

A more realistic reverberation effect can be achieved using Schroeder's model of reverberation, which consists of several plain reverb units in parallel, followed by several allpass units in series. An example is depicted in Fig. 8.2.18 and on the cover of the text [1], and shown below.
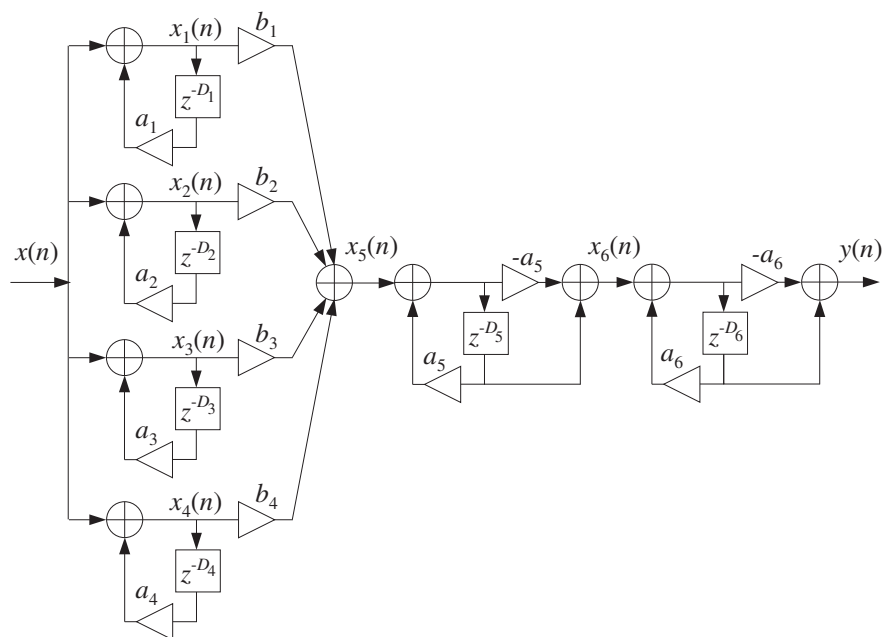
The different delays in the six units cause the density of the reverberating echoes to increase, generating an impulse response that exhibits the typical early and late reflection characteristics.

Its sample processing algorithm is given by Eq. (8.2.31) of [1]. It is stated in terms of the functions `plain()` and `allpass()` that implement the individual units:

$$
\boxed{
\begin{aligned}
&\textit{for each input sample } x \textit{ do:}\\
&\quad x_1 = \text{plain}(D_1, \mathbf{w}_1, \&p_1, a_1, x)\\
&\quad x_2 = \text{plain}(D_2, \mathbf{w}_2, \&p_2, a_2, x)\\
&\quad x_3 = \text{plain}(D_3, \mathbf{w}_3, \&p_3, a_3, x)\\
&\quad x_4 = \text{plain}(D_4, \mathbf{w}_4, \&p_4, a_4, x)\\
&\quad x_5 = b_1 x_1 + b_2 x_2 + b_3 x_3 + b_4 x_4\\
&\quad x_6 = \text{allpass}(D_5, \mathbf{w}_5, \&p_5, a_5, x_5)\\
&\quad y\ = \text{allpass}(D_6, \mathbf{w}_6, \&p_6, a_6, x_6)
\end{aligned}
}
\tag{5.1}
$$

There are six multiple delays each requiring its own circular buffer and pointer. The `allpass()` function was already defined in the allpass reverb lab section. The `plain` function is straightforward and implements the steps used in the plain reverb lab section:

```c
// -----------------------------------------------------
// plain.c - plain reverb with circular delay line
// -----------------------------------------------------

float *pwrap(int, float *, float *);

float plain(int D, float *w, float **p, float a, float x)
{
   float y, sD;

   sD = *pwrap(D,w,*p+D);

   y = x + a * sD;

   **p = y;

   *p = pwrap(D,w,--*p);

   return y;
}
// -----------------------------------------------------
```

The following (incomplete) C program implements the above sample processing algorithm in its `isr()` function and operates at a sampling rate of 44.1 kHz:

```c
// ----------------------------------------------------------------------
// schroeder.c - Schroeder's reverb algorithm using circular buffers
// ----------------------------------------------------------------------

#include "dsplab.h"          // init parameters and function prototypes

short xL, xR, yL, yR;        // input and output samples from/to codec

short fs = 44;               // sampling rate in kHz

#define D1 1759
#define D2 1949
#define D3 2113
#define D4 2293
#define D5  307
#define D6  313

#define a 0.88

float b1=1, b2=0.9, b3=0.8, b4=0.7;
float a1=a, a2=a, a3=a, a4=a, a5=a, a6=a;

float w1[D1+1], *p1;
float w2[D2+1], *p2;
float w3[D3+1], *p3;
```

```
      float w4[D4+1], *p4;
      float w5[D5+1], *p5;
      float w6[D6+1], *p6;

      float plain(int, float *, float **, float, float);        // must be added to project
      float allpass(int, float *, float **, float, float);

      // ------------------------------------------------------------------------------

      void main()
      {
         int n;
         for (n=0; n<=D1; n++) w1[n] = 0;              // initialize buffers to zero
         for (n=0; n<=D2; n++) w2[n] = 0;
         for (n=0; n<=D3; n++) w3[n] = 0;
         for (n=0; n<=D4; n++) w4[n] = 0;
         for (n=0; n<=D5; n++) w5[n] = 0;
         for (n=0; n<=D6; n++) w6[n] = 0;

         p1 = w1; p2 = w2; p3 = w3; p4 = w4; p5 = w5; p6 = w6;      // initialize pointers

         initialize();               // initialize DSK board and codec, define interrupts

         sampling_rate(fs);          // possible sampling rates: 8, 16, 24, 32, 44, 48, 96 kHz
         audio_source(MIC);          // LINE or MIC for line or microphone input

         while(1);                   // keep waiting for interrupt, then jump to isr()
      }

      // ------------------------------------------------------------------------------

      interrupt void isr()
      {
         read_inputs(&xL, &xR);          // read inputs from codec

         // ------------------------------------------------------------
         // here insert your algorithm implementing Eq.(5.1) given above
         // ------------------------------------------------------------

         write_outputs(yL,yR);           // write outputs to codec

         return;
      }

      // ------------------------------------------------------------------------------
```

Lab Procedure

a. Create a project for this program, compile and run it with audio input set to MIC. Listen to its impulse response and speak into the mike. To reduce potential overflow effects, you may want to reduce the input level by half, for example, by the statement:

```
      x = (float) (xL>>1);
```

b. What are the feedback delays of each unit in msec? Replace all the delays by double their values, compile, and run again. Compare the output with that of part (a). Repeat when you triple all the delays. (Note that you can just replace the constant definitions by #define D1 1759*2, etc.)

c. Repeat part (a) by experimenting with different values of the feedback parameter a.

5.5. *Stereo Reverb*

In some of the previous experiments, we considered processing in stereo, but the left and right channels were processed completely independently of each other. In this experiment, we allow the cross-coupling

of the two channels, so that the reverb characteristics of one channel influences those of the other.

An example of such system is given in Problems 8.22 and 8.23 and depicted in Fig. 8.4.1 of the text [1] and shown below.
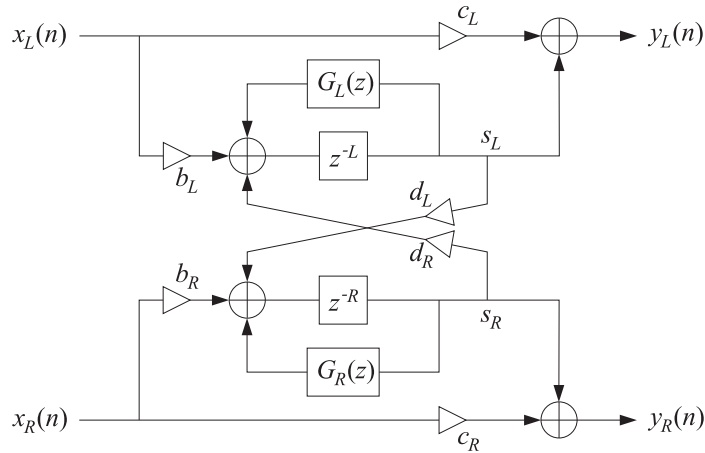
Here, we assume that the feedback filters are plain multiplier gains, so that

$$G_L(z) = a_L, \qquad G_R(z) = a_R$$

Each channel has its own delay-line buffer and circular pointer. The sample processing algorithm is modified now to take in a pair of stereo inputs and produce a pair of stereo outputs:

$$
\begin{aligned}
&\textit{for each input stereo pair } x_L, x_R \textit{ do:}\\
&\quad s_L = *(p_L + L)\\
&\quad s_R = *(p_R + R)\\
&\quad y_L = c_L x_L + s_L\\
&\quad y_R = c_R x_R + s_R\\
&\quad *p_L = s_{L0} = b_L x_L + a_L s_L + d_R s_R\\
&\quad *p_R = s_{R0} = b_R x_R + a_R s_R + d_L s_L\\
&\quad --p_L\\
&\quad --p_R
\end{aligned}
$$

where L and R denote the left and right delays. Cross-coupling between the channels arises because of the coefficients d_L and d_R. The following is its C translation into an `isr()` function:

```
interrupt void isr()            // sample processing algorithm - interrupt service routine
{
   float sL, sR;

   read_inputs(&xL, &xR);          // read inputs from codec

   sL = *pwrap(L,wL,pL+L);
   sR = *pwrap(R,wR,pR+R);
   yL = cL*xL + sL;
   yR = cR*xR + sR;
  *pL = bL*xL + aL*sL + dR*sR;
  *pR = bR*xR + aR*sR + dL*sL;
   pL = pwrap(L,wL,--pL);
   pR = pwrap(R,wR,--pR);

   write_outputs(yL,yR);           // write outputs to codec

   return;
}
```

Lab Procedure

a. Create a project whose main program includes the above ISR. Select an 8 kHz sampling rate and line input. Choose the following parameter values:

$$L = R = 3000, \quad a_L = a_R = 0, \quad b_L = b_R = 0.8, \quad c_L = c_R = 0.5, \quad d_L = d_R = 0.5$$

Compile and run this program. Even though the self-feedback multipliers were set to zero, $a_L = a_R = 0$, you will hear repeated echoes bouncing back and forth between the speakers because of the cross-coupling. Make sure the speakers are as far separated as possible, and play one of the wave files in `c:\dsplab\wav` (e.g., `take5`, `dsummer`).

b. Next try the case $d_L \neq 0$, $d_R = 0$. And then, $d_L = 0$, $d_R \neq 0$. These choices decouple the influence of one channel but not that of the other.

c. Next, introduce some self-feedback, such as $a_L = a_R = 0.2$. Repeat part (a). Vary all the parameters at will to see what you get.

5.6. *Reverberating Delay*

A prototypical delay effect found in most commercial audio effects processors was discussed in Problem 8.17 of the text [1]. Its transfer function is:

$$H(z) = c + b \frac{z^{-D}}{1 - az^{-D}}$$

Its block diagram realization and corresponding sample processing algorithm using a circular delay-line buffer are given below:
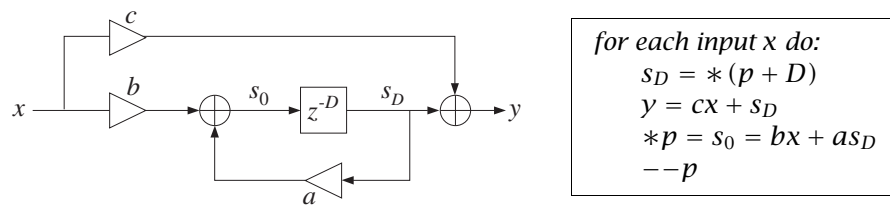
The following is its C translation into an `isr()` function:

```
interrupt void isr()              // sample processing algorithm - interrupt service routine
{
    float sD, x, y;               // D-th state, input & output

    read_inputs(&xL, &xR);        // read inputs from codec

    x = (float) xL;               // process left channel only

    sD = *pwrap(D,w,p+D);         // extract states relative to p
    y = c*x + sD;                 // output sample
    *p = b*x + a*sD;              // delay-line input
    p = pwrap(D,w,--p);           // backshift pointer

    yL = yR = (short) y;

    write_outputs(yL,yR);         // write outputs to codec

    return;
}
```

Lab Procedure

a. Create a project, compile and run it with 8 kHz sampling rate and MIC input. Choose the parameters:

$$D = 6000, \quad a = 0.5, \quad b = 1, \quad c = 0$$

Listen to its impulse response and speak into the mike. Here, the direct sound path has been removed, $c = 0$, in order to let the echoes be more clearly heard.

b. What values of b and c would you use (expressed in terms of a) in order to implement a plain reverberator of the form:

$$H(z) = \frac{1}{1 - az^{-D}}$$

For $a = 0.5$, calculate the proper values of b, c, and then compile and run the program. Compare its output with that of `plain1.c`.

c. Compile and run the case: $a = 1, b = c = 1$, and then the case: $a = -1, b = -1, c = 1$. What are the transfer functions in these cases?

5.7. Multi-Delay Effects

Here, we consider the multi-delay effects processor shown in Fig. 8.2.27 of the text [1]. We assume that the feedback filters are plain multipliers. Using two separate circular buffers for the two delays, the block diagram realization and sample processing algorithm are in this case:
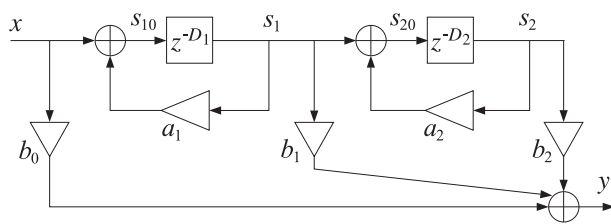
for each input x do:
$$s_1 = *(p_1 + D_1)$$
$$s_2 = *(p_2 + D_2)$$
$$y = b_0 x + b_1 s_1 + b_2 s_2$$
$$*p_2 = s_{20} = s_1 + a_2 s_2$$
$$--p_2$$
$$*p_1 = s_{10} = x + a_1 s_1$$
$$--p_1$$

Its C translation is straightforward:

```
interrupt void isr()          // sample processing algorithm - interrupt service routine
{
   float x, s1, s2, y;

   read_inputs(&xL, &xR);        // read inputs from codec

   x = (float) xL;               // process left channel only

   s1 = *pwrap(D1, w1, p1+D1);
   s2 = *pwrap(D2, w2, p2+D2);

   y = b0*x + b1*s1 + b2*s2;

   *p2 = s1 + a2*s2;
   p2 = pwrap(D2, w2, --p2);

   *p1 = x + a1*s1;
   p1 = pwrap(D1, w1, --p1);

   yL = yR = (short) y;

   write_outputs(yL,yR);         // write outputs to codec

   return;
}
```

Lab Procedure

a. Write a main program, `multidel.c`, that incorporates this ISR, compile and run it with an 8 kHz sampling rate and MIC input, and the following parameter choices:

$$D_1 = 5000, \quad D_2 = 2000, \quad a_1 = 0.5, \quad a_2 = 0.4, \quad b_0 = 1, \quad b_1 = 0.8, \quad b_2 = 0.6$$

Listen to its impulse response and speak into the mike. Then select LINE input and play a wave file (e.g., `dsummer`) through it.

b. Set $b_1 = 0$ and run again. Then, set $b_2 = 0$ and run. Can you explain what you hear?

5.8. *Multitap Delay Effects*

This experiment is based on the multi-tap delay line effects processor of Fig. 8.2.29 of the text [1]. Both this effect and the multi-delay effect of the previous section are commonly found in commercially available digital audio effects units.

The implementation uses a common circular delay-line buffer of order D_1+D_2, which is tapped out at taps D_1 and D_1+D_2. The sample processing algorithm is:
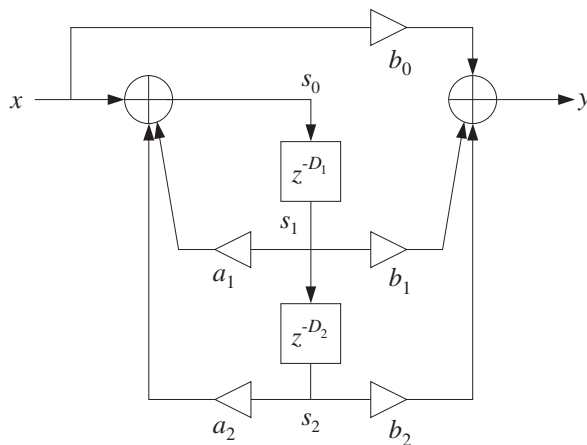
for each input sample x do:
$$s_1 = *(p + D_1)$$
$$s_2 = *(p + D_1 + D_2)$$
$$y = b_0 x + b_1 s_1 + b_2 s_2$$
$$s_0 = x + a_1 s_1 + a_2 s_2$$
$$*p = s_0$$
$$--p$$

The following ISR is its C translation:

```
interrupt void isr()            // sample processing algorithm - interrupt service routine
{
   float x, s0, s1, s2, y;

   read_inputs(&xL, &xR);       // read inputs from codec

   x = (float) xL;              // process left channel only

   s1 = *pwrap(D1+D2, w, p+D1);
   s2 = *pwrap(D1+D2, w, p+D1+D2);
   y = b0*x + b1*s1 + b2*s2;
   s0 = x + a1*s1 + a2*s2;
   *p = s0;
   p = pwrap(D1+D2, w, --p);

   yL = yR = (short) y;

   write_outputs(yL,yR);        // write outputs to codec

   return;
}
```

Lab Procedure

a. Write a main program, `multidel.c`, that incorporates this ISR, compile and run it with an 8 kHz sampling rate and MIC input, and the following parameter choices:

$$D_1 = 3000, \quad D_2 = 1500, \quad a_1 = 0.2, \quad a_2 = 0.5, \quad b_0 = 1, \quad b_1 = 0.8, \quad b_2 = 0.6$$

Listen to its impulse response and speak into the mike. Then select LINE input and play a wave file (e.g., `dsummer`) through it.

b. Repeat for the following values of the feedback parameters: $a_1 = a_2 = 0.5$, which makes the system marginally stable with a periodic steady output (any random noise would be grow unstable.)

Repeat also for the case $a_1 = a_2 = 0.75$, which corresponds to an unstable filter. Please reset the processor before the output grows too loud. However, do let it grow loud enough to hear the overflow effects arising from the growing feedback output s_0.

As discussed in Ref. [1], the condition of stability for this filter is $|a_1| + |a_2| < 1$. Interestingly, most commercially available digital audio effects units allow the setting of the parameters $D_1, D_2, a_1, a_2, b_0, b_1, b_2$ from their front panel, but do not check this stability condition.

5.9. *Karplus-Strong String Algorithm*

A model of a plucked string is obtained by running the lowpass reverb filter with zero input, but with initially filling the delay line with random numbers. These random numbers model the initial harshness of plucking the string. But, as the random numbers recirculate through the lowpass filter, their high frequencies are gradually removed, resulting in a sound that models the string vibration.

The model can be approximately "tuned" to a frequency f_1 by picking D such that $D = f_s/f_1$. The Karplus-Strong model [9] assumes a simple averaging FIR filter for the lowpass feedback filter as given by Eq. (8.2.40) of the text [1]. Here, we take the transfer function to be:

$$G(z) = b_0(1 + z^{-1})$$

with some $b_0 \lesssim 0.5$ to improve the stability of the closed-loop system. See Refs. [4–15] for more discussion on such models and computer music in general. The following program implements the algorithm. The code is identical to that of the lowpass reverb case.

The sampling rate is set to 44.1 kHz and the generated sound is the note A440, that is, having frequency 440 Hz. The correct amount of delay is then

$$D = \frac{f_s}{f_1} = \frac{44100}{440} \approx 100$$

The delay line must be filled with $D+1$ random numbers. They were generated as follows by MATLAB and exported to the file `rand.dat` using the function `C_header()`, e.g., by the code:

```
iseed = 1000; randn('state', iseed);
r = 10000 * randn(101,1);
C_header('rand.dat', 'r', 'D', r);
```

The full program is as follows:

```
// -------------------------------------------------------------------------------
// ks.c - Karplus-Strong string algorithm
// -------------------------------------------------------------------------------

#include "dsplab.h"            // init parameters and function prototypes

short xL, xR, yL, yR;          // input and output samples from/to codec

#define D 100
```

```
float w[D+1], *p;              // circular delay-line buffer, circular pointer

#include "rand.dat"            // D+1 random numbers

float a = 0;
float b0 = 0.499, b1 = 0.499;

float v0, v1;                  // states of feedback filter

short fs = 44;                 // sampling rate is 44.1 kHz

// ------------------------------------------------------------------------------

void main()                    // main program executed first
{
   int n;
   for (n=0; n<=D; n++)        // initialize circular buffer to zero
      w[n] = r[n];
   p = w;                      // initialize pointer
   v1 = 0;                     // initialize feedback filter

   initialize();              // initialize DSK board and codec, define interrupts

   sampling_rate(fs);         // possible sampling rates: 8, 16, 24, 32, 44, 48, 96 kHz
   audio_source(LINE);        // LINE or MIC for line or microphone input

   while(1);                  // keep waiting for interrupt, then jump to isr()
}

// ------------------------------------------------------------------------------

interrupt void isr()           // sample processing algorithm - interrupt service routine
{
   float y, sD, u;

   // read_inputs(&xL, &xR);        // inputs not used

   sD = *pwrap(D,w,p+D);

   v0 = a*v1 + sD;            // feedback filter G(z) = (b0 + b1*z^-1)/(1-a*z^-1)
   u = b0*v0 + b1*v1;        // feedback filter's output
   v1 = v0;                  // update feedback filter's delay

   y = u;                    // closed-loop output - with x=0

   *p = y;

   p = pwrap(D,w,--p);

   yL = yR = (short) y;

   write_outputs(yL,yR);        // write outputs to codec

   return;
}
// ------------------------------------------------------------------------------
```

Lab Procedure

a. Create a project, compile and run. The program disables the inputs and simply outputs the recirculating and gradually decaying random numbers.

b. Repeat for $D = 200$ by generating a new file rand.dat using the above MATLAB code. The note you hear should be an octave lower.

5.10. *Flangers and Vibrato*

As discussed in Ref. [1], a flanging effect is implemented as an FIR comb filter with a time-variable delay.
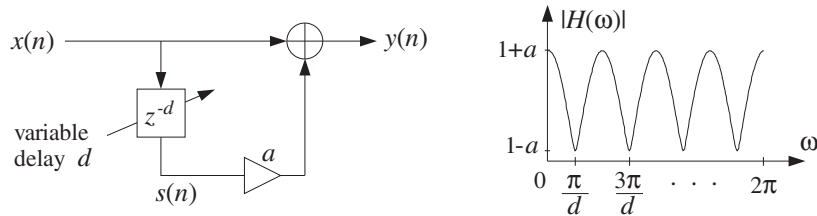
If the delay d varies sinusoidally between $0 \le d(n) \le D$, with some low frequency f_d, then

$$d(n) = \frac{D}{2}[1 - \cos(\omega_d n)], \quad \omega_d = \frac{2\pi f_d}{f_s} \quad \text{[rads/sample]}$$

and the flanger output is obtained by

$$y(n) = x(n) + ax(n - d(n))$$

If the delay d were fixed, the transfer function would be:

$$H(z) = 1 + az^{-d}$$

The peaks of the frequency response of the resulting time-varying comb filter, occurring at multiples of f_s/d, and its notches at odd multiples of $f_s/2d$, will sweep up and down the frequency axis resulting in the characteristic whooshing type sound called flanging. The parameter a controls the depth of the notches. In units of [radians/sample], the notches occur at odd multiples of π/d.

In the early days, the flanging effect was created by playing the music piece simultaneously through two tape players and alternately slowing down each tape by manually pressing the flange of the tape reel.

Because the variable delay d can take non-integer values within its range $0 \le d \le D$, the implementation requires the calculation of the output $x(n-d)$ of a delay line at such non-integer values. This can be accomplished easily by truncating to the nearest integer, or as discussed in [1], by rounding, or by linear interpolation. To sharpen the comb peaks one may use a plain-reverb filter with variable delay, that is,

$$y(n) = x(n) + ay(n-d), \qquad H(z) = \frac{1}{1 - az^{-d}}$$

Its sample processing algorithm using a circular buffer of maximum order D is:

$$\boxed{\begin{array}{l} \text{for each input } x \text{ do:} \\ \quad d = \text{floor}\left[(1 - \cos(\omega_d n))D/2\right] \\ \quad s_d = *(p + d) \\ \quad y = x + a\,s_d \\ \quad *p = y \\ \quad --p \end{array}}$$

Its translation to C is straightforward and can be incorporated into the ISR function:

```
interrupt void isr()            // sample processing algorithm - interrupt service routine
{
    float sd;

    read_inputs(&xL, &xR);         // read inputs from codec

    x = (float) xL;                // work with left input only

    d = (1 - cos(wd*n))*D/2;       // automatically cast to int, wd = 2*PI*fd/fs
```

```
    if (++n>=L) n=0;                // L = 16000 to allow fd = 0.5 Hz

    sd = *pwrap(D,w,p+d);           // extract d-th state relative to p
    y = x + a*sd;                   // output
    *p = y;                         // delay-line input
    p = pwrap(D,w,--p);             // backshift pointer

    yL = yR = (short) y;

    write_outputs(yL,yR);           // write outputs to codec

    return;
}
```

Lab Procedure

a. Create a project for this ISR. You will need to include `<math.h>` and define `PI`. Choose D to correspond to a 2 msec maximum delay and let $f_d = 1$ Hz and $a = 0.7$. Run the program and play a wave file through it (e.g., `noflange`, `dsummer`, `take5`). Repeat when $f_d = 0.5$ Hz.

b. Experiment with other values of D, f_d, and a.

c. Rewrite part (a) so that an FIR comb filter is used as shown at the beginning of this section. Play the same material through the IIR and FIR versions and discuss differences in their output sounds.

d. A *vibrato* effect can be obtained by using the filter $H(z) = z^{-d}$ with a variable delay. You can easily modify your FIR comb filter of part (c) so that the output is taken directly from the output of the delay. For this effect the typical delay variations are about 5 msec and their frequency about 5 Hz. Create a vibrato project with $D = 16$ (correspondoing to 2 msec at an 8 kHz rate) and $f_d = 5$ Hz, and play a wave file through it. Repeat by doubling D and/or f_d.

5.11. References

[1] S. J. Orfanidis, *Introduction to Signal Processing*, online book, 2010, available from:
    `http://www.ece.rutgers.edu/~orfanidi/intro2sp/`

[2] R. Chassaing and D. Reay, *Digital Signal Processing and Applications with the TMS320C6713 and TMS320C6416 DSK*, 2nd ed., Wiley, Hoboken, NJ, 2008.

[3] M. J. Caputi, "Developing Real-Time Digital Audio Effects for Electric Guitar in an Introductory Digital Signal Processing Class," *IEEE Trans. Education*, **41**, no.4, (1998), available online from:
    `http://www.ewh.ieee.org/soc/es/Nov1998/01/BEGIN.HTM`

[4] F. R. Moore, *Elements of Computer Music*, Prentice Hall, Englewood Cliffs, NJ, 1990.

[5] C. Roads and J. Strawn, eds., *Foundations of Computer Music*, MIT Press, Cambridge, MA, 1988.

[6] C. Roads, ed., *The Music Machine*, MIT Press, Cambridge, MA, 1989.

[7] C. Dodge and T. A. Jerse, *Computer Music*, Schirmer/Macmillan, New York, 1985.

[8] J. M. Chowning, "The Synthesis of Complex Audio Spectra by Means of Frequency Modulation," *J. Audio Eng. Soc.*, **21**, 526 (1973). Reprinted in Ref. [5].

[9] R. Karplus and A. Strong, "Digital Synthesis of Plucked String and Drum Timbres," *Computer Music J.*, **7**, 43 (1983). Reprinted in Ref. [6].

[10] D. A. Jaffe and J. O. Smith, "Extensions of the Karplus-Strong Plucked-String Algorithm," *Computer Music J.*, **7**, 56 (1983). Reprinted in Ref. [6].

[11] C. R. Sullivan, "Extending the Karplus-Strong Algorithm to Synthesize Electric Guitar Timbres with Distortion and Feedback," *Computer Music J.*, **14**, 26 (1990).

[12] J. O. Smith, "Physical Modeling Using Digital Waveguides," *Computer Music J.*, **16**, 74 (1992).

[13] J. A. Moorer, "Signal Processing Aspects of Computer Music: A Survey," *Proc. IEEE*, **65**, 1108 (1977). Reprinted in Ref. [5].

[13] M. Kahrs and K. Brandenburg, eds., *Applications of Digital Signal Processing to Audio and Acoustics*, Kluwer, Boston, 1998.

[15] Udo Zölzer, ed., *DAFX – Digital Audio Effects*, Wiley, Chichester, England, 2003. See also the DAFX Conference web page: `http://www.dafx.de/`.

Lab 6 – IIR Filtering Experiments

6.1. Transposed Realization

In Lab-4 you implemented your own version of the function **filter**, specialized to FIR filters, using the transposed realization. In this lab you will implement the IIR case. The transposed realization block diagram is shown in the MATLAB documentation for the function **filter**. An equivalent version is shown below for an order-3 filter.

$$H(z) = \frac{b_0 + b_1 z^{-1} + b_2 z^{-2} + b_3 z^{-3}}{1 + a_1 z^{-1} + a_2 z^{-2} + a_3 z^{-3}}$$
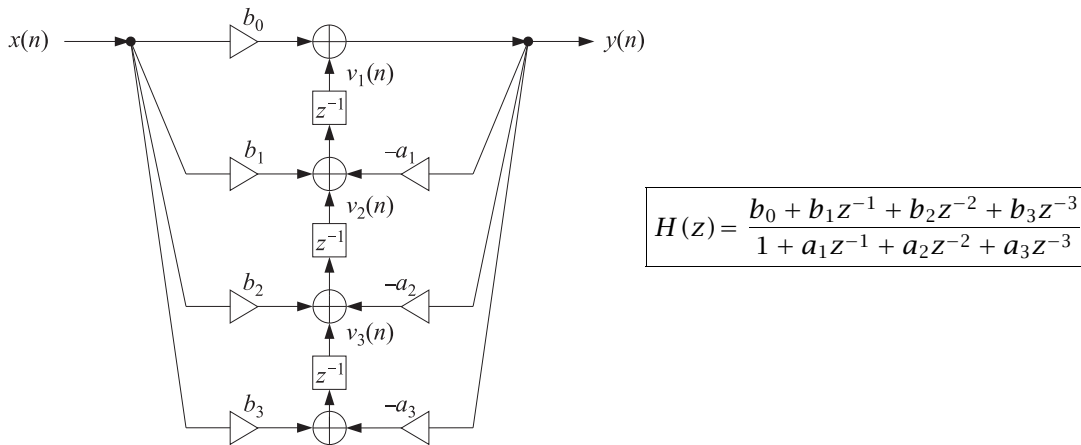
For an Mth order filter, the computational algorithm is,

$$y(n) = v_1(n) + b_0 x(n)$$
$$v_1(n+1) = v_2(n) + b_1 x(n) - a_1 y(n)$$
$$v_2(n+1) = v_3(n) + b_2 x(n) - a_2 y(n)$$
$$\vdots$$
$$v_{M-1}(n+1) = v_M(n) + b_{M-1} x(n) - a_{M-1} y(n)$$
$$v_M(n+1) = b_M x(n) - a_M y(n)$$

$$\mathbf{v}(n) = \begin{bmatrix} v_1(n) \\ v_2(n) \\ \vdots \\ v_M(n) \end{bmatrix} = \text{state vector}$$

The state vector $\mathbf{v}(n)$ represents the current contents of the M delays. The algorithm is an example of a state-space realization. It computes the current output $y(n)$ from the current input $x(n)$ and the current state $\mathbf{v}(n)$, and then, it updates the state to the next time instant, $\mathbf{v}(n+1)$. Usually, the state vector is initialized to zero, but it can be initialized to an arbitrary vector, say, $\mathbf{v}_{\text{in}}$.

Lab Procedure

a. Write a MATLAB function, say, **filtr.m**, that is functionally equivalent to the built-in function **filter** and implements the above algorithm and has the possible syntaxes:

```
y = filtr(b,a,x);
[y,vout] = filtr(b,a,x,vin);

%   b   = order-L numerator coefficient vector,   b = [b0, b1, b2, ..., bL]
%   a   = order-K denominator coefficient vector, a = [ 1, a1, a2, ..., aK]
%   x   = length-N vector of input samples (row or column)
%   y   = length-N vector of output samples (row or column according to x)
% vin   = M-dimensional vector of initial states - zero vector, by default
% vout  = M-dimensional final state vector, i.e., final contents of delays
```

Internally, yout function must extend **b, a** to the maximum required order $M = \max(L, K)$ by padding zeros. Test your function on the following case:

```
b = [1,   0.6, -0.4];              % numerator
a = [1, -1.3,   0.2,  0.4];        % denominator
x = [1, 1, 2, 1, 2, 2, 1, 1];
y = [1, 2.9000, 5.7700, 8.3210, 10.3033, 12.2221, 11.8997, 9.7038];    % expected result
```

b. The function **filtr** can be run on a sample by sample basis to generate the successive internal state vectors, for example, using a loop such as,

```
v = zeros(1,M);                        % initial state vector
for n=1:length(x)
   [y(n),vout] = filtr(b,a,x(n),v);    % recycled state vector v
   v = vout;                           % next state
end
```

Add appropriate **fprintf** commands before, within, and after this loop to generate the following table of values for the above example,

```
n    x      y         v1        v2        v3
--------------------------------------------------
0    1    1.0000    0.0000    0.0000    0.0000
1    1    2.9000    1.9000   -0.6000   -0.4000
2    2    5.7700    3.7700   -1.3800   -1.1600
3    1    8.3210    7.3210   -3.1140   -2.3080
4    2   10.3033    8.3033   -4.3722   -3.3284
5    2   12.2221   10.2221   -6.1891   -4.1213
6    1   11.8997   10.8997   -7.3657   -4.8888
7    1    9.7038    8.7038   -7.6688   -4.7599
8    -      -       5.5462   -7.1006   -3.8815
```

where v_1, v_2, v_3 are the internal states at each time instant.

c. The effective *time constant* of a stable and causal IIR filter, measured in samples, may be defined in terms of the maximum pole radius as follows,

$$n_{\text{eff}} = \frac{\ln \epsilon}{\ln R}, \quad R = \max_i |p_i| \tag{6.1}$$

where p_i are the poles of the filter, and ϵ is a small number, such as $\epsilon = 10^{-2}$ or $\epsilon = 10^{-3}$ for the so-called 40-dB or 60-dB time constants, respectively.

Calculate the 40-dB time constant of the example filter of part (a), and then using your function **filtr**, evaluate and plot its *impulse response* h_n over the interval $0 \le n \le n_{\text{eff}}$. Similarly, evaluate and plot the *unit-step response* over the same time interval. Note that the asymptotic value of the unit-step response (for a stable/causal filter) is given in terms of the filter coefficients by

$$u_\infty = \frac{\displaystyle\sum_{m=0}^{M} b_m}{\displaystyle\sum_{m=0}^{M} a_m}$$

Indicate this value on your graph of the unit-step response.

d. *Envelope detectors* are used extensively in communication systems and in audio effects such as compressors and expanders. A typical envelope detector is a lowpass filter acting on the absolute value of a signal.

In this part, you will determine the envelope of the impulse response h_n that was calulated in the part (c). Consider the following IIR and FIR lowpass filtering operations on the signal $|h_n|$:

$$\text{(IIR)} \qquad h_{\text{env}}(n) = a\,h_{\text{env}}(n-1) + (1-a)\,|h_n|$$

$$\text{(FIR)} \qquad h_{\text{env}}(n) = \frac{|h_n| + |h_{n-1}| + |h_{n-2}| + \cdots + |h_{n-N+1}|}{N} \tag{6.2}$$

where $0 < a < 1$ and N is an integer. It can be shown that the two filters are roughly equivalent if one makes the identification

$$N = \text{floor}\left(\frac{1+a}{1-a}\right)$$

For $a = 0.80$, calculate N, and carry out the filtering operations of Eq. (6.2). Plot $h(n)$ and $h_{\text{env}}(n)$ over the time interval $0 \le n \le n_{\text{neff}}$. Make different plots for the IIR and the FIR envelopes.

Repeat for $a = 0.94$. Comment on the tradeoffs between the two choices $a = 0.80$ and $a = 0.94$ in terms of their initial transients and their ability to extract the envelope.

6.2. Feedback System

A general feedback system is shown below, where the output of filter $H_1(z)$ is fed back into filter $H_2(z)$ and then back to the input, and where the delay z^{-1} can be positioned at the four locations A, B, C, or D.
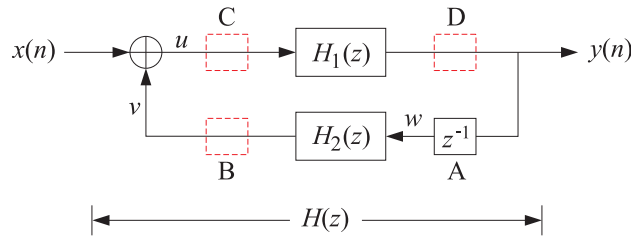
Lab Procedure

a. For each case A,B,C,D, determine the transfer function $H(z)$ of the overall closed-loop system from $x(n)$ to $y(n)$, expressed in terms of the transfer functions $H_1(z)$ and $H_2(z)$.

b. Choose the forward and feedback transfer functions as,

$$H_1(z) = \frac{1 + z^{-1}}{1 - 0.5z^{-1}}, \qquad H_2(z) = \frac{0.4 - 0.4z^{-1}}{1 - 0.4z^{-1}} \tag{6.3}$$

Show that for case A, the transfer function $H(z)$ is,

$$H(z) = \frac{1 + 0.6z^{-1} - 0.4z^{-2}}{1 - 1.3z^{-1} + 0.2z^{-2} + 0.4z^{-3}} \tag{6.4}$$

What are the transfer functions $H(z)$ for the cases B, C, D?

c. The objective of this part is to use the function **filtr** on a sample-by-sample basis, applied separately on the transfer functions $H_1(z)$ and $H_2(z)$, to implement the feedback system directly, without transforming it first into the equivalent transfer function $H(z)$.

As we saw earlier, the filtering function **filtr** can be invoked repetitively on a sample-by-sample basis using a for-loop that recycles the state vector:

```
s = zeros(1,M);                        % initial state vector
for n=1:length(x)
   [y(n),s] = filtr(b,a,x(n),s);       % recycled state vector s
end
```

To implement the transfer functions $H_1(z)$ and $H_2(z)$ of part (b), you need to define their filter coefficients and states, $\{\mathbf{b}_1, \mathbf{a}_1, \mathbf{s}_1\}$ and $\{\mathbf{b}_2, \mathbf{a}_2, \mathbf{s}_2\}$, and initialize the states to zero, as well as the content of the delay at locations A,B,C,D. For example, initially the content w of the delay at A shown in the above block diagram will be zero. As an example, the pseudo-code implementation of case A will be:

> initialize: $w = 0$, $s_1 = 0$, $s_2 = 0$
>
> for each input sample x do:
>
> $\quad [v, s_2]$ = output and state of H_2 with input sample w and state s_2
>
> $\quad u = x + v$
>
> $\quad [y, s_1]$ = output and state of H_1 with input sample u and state s_1
>
> $\quad w = y$

Implement this case in a for-loop in Matlab and evaluate the output $y(n)$ when the input is selected to be:

```
x = [1, 1, 2, 1, 2, 2, 1, 1];
```

Check your output by also computing it using the overall transfer function $H(z)$ of Eq. (6.4).

Similarly, formulate the computational steps using a for-loop for cases B, C, D, and calculate the corresponding output signals. Note that each case requires a different ordering of the computations.

Using appropriate **fprintf** commands, print all cases A, B, C. D in a table as shown below, where the second column y represents the output of case A computed with Eq. (6.4),

x	y	yA	yB	yC	yD
1.0000	1.0000	1.0000	1.0000	0.0000	0.0000
1.0000	2.9000	2.9000	2.9000	1.0000	1.0000
2.0000	5.7700	5.7700	5.7700	2.9000	2.9000
1.0000	8.3210	8.3210	8.3210	5.7700	5.7700
2.0000	10.3033	10.3033	10.3033	8.3210	8.3210
2.0000	12.2221	12.2221	12.2221	10.3033	10.3033
1.0000	11.8997	11.8997	11.8997	12.2221	12.2221
1.0000	9.7038	9.7038	9.7038	11.8997	11.8997

Example Graphs
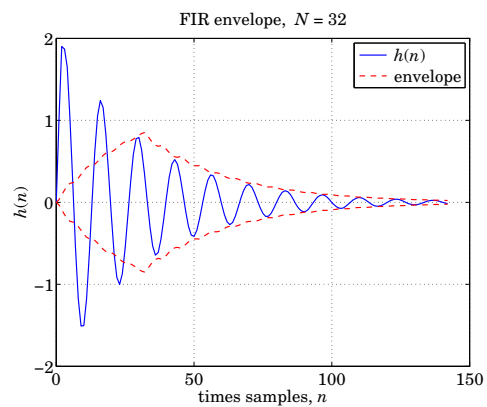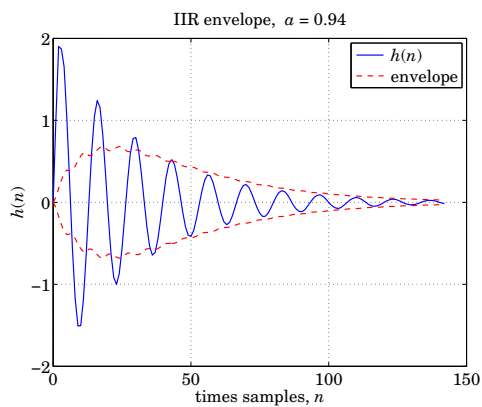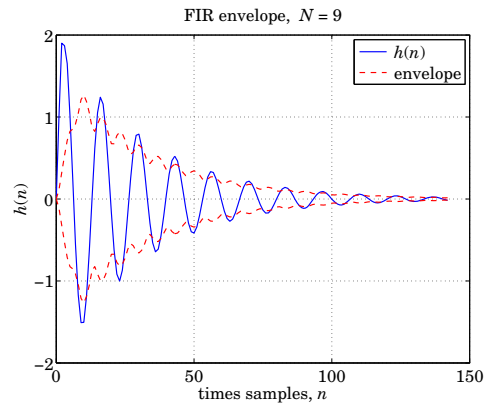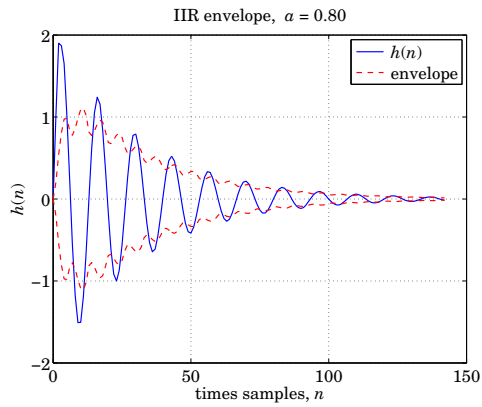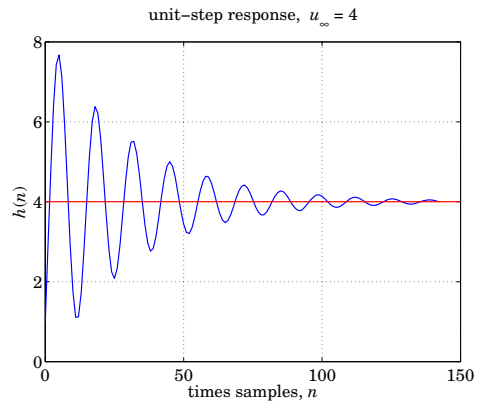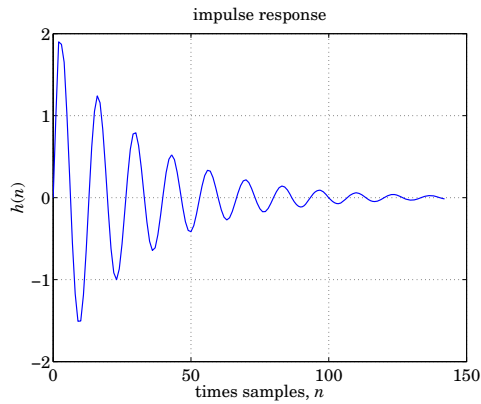
References

[1] A. V. Oppenheim, R. W. Schafer, *Discrete-Time Signal Processing*, 3/e, Prentice Hall, Englewood Cliffs, NJ, 2009.

[2] S. J. Orfanidis, *Introduction to Signal Processing*, online book, 2010, available from: `http://www.ece.rutgers.edu/~orfanidi/intro2sp/`.

[3] R. Chassaing and D. Reay, *Digital Signal Processing and Applications with the TMS320C6713 and TMS320C6416 DSK*, 2nd ed., Wiley, Hoboken, NJ, 2008.

[4] S. A. Tretter, *Communication System Design Using DSP Algorithms with Laboratory Experiments for the TMS320C6713 DSK*, Springer, New York, 2008, code available from: `http://www.ece.umd.edu/~tretter`

[5] MATLAB Documentation: `http://www.mathworks.com/help/techdoc/`