

GEORGIA INSTITUTE OF TECHNOLOGY  
SCHOOL of ELECTRICAL and COMPUTER ENGINEERING

**EE 2200 Winter 1999**  
**Lab #5: Potpourri: A/D and D/A**

Date: week of 8 Feb 1999

---

This is *the official* Lab#5 description.

**Lab Quiz #2 will be given during the week of 15-Feb-99.**

The Warm-up section of each lab must be completed in Lab and the steps marked *Instructor Verification* must also be signed off **during the lab time**. One of the laboratory instructors must verify the appropriate steps by initialing on the **Instructor Verification** line. When you have completed a step that requires verification, simply raise your hand and demonstrate the step to the instructor.

The lab report for this week will an **Informal Lab Report**. Staple the **Instructor Verification** sheet to the end of your lab report as evidence that the appropriate steps were witnessed by the instructor.

The report will **due during the week of 15-Feb at the start of your lab**.

---

## 1 Debugging

I heard the following (or something similar) from a veteran programmer who had developed a real-time OS for data acquisition system that handled signals from a variety of sensors for oil-field exploration:

*“Any fool can write a computer program, but it takes brain power to debug one.”*

Testing and debugging code is more than half the job, as you know if you been staying up late on the first few 2200 labs. Almost any modern programming environment provides a *symbolic debugger* so that break-points can be set and variables examined in the middle of program execution. Of course, many programmers (and 2200 students) insist on using the old-fashioned method of inserting print statements in the middle of their code (or the MATLAB equivalent, leave off a few semi-colons). This is akin to riding a tricycle to commute around Atlanta.

In order to learn how to use the MATLAB tools for debugging, try `help debug`. Here is part of what you'll see:

```
dbstop      - Set breakpoint.
dbclear     - Remove breakpoint.
dbcont      - Resume execution.
dbstack     - List who called whom.
dbstatus    - List all breakpoints.
dbstep      - Execute one or more lines.
dbtype      - List M-file with line numbers.
dbquit      - Quit debug mode.
```

When a breakpoint is hit, MATLAB goes into debug mode. On the PC and Macintosh the debugger window becomes active and on UNIX and VMS the prompt changes to a `K>`. Any MATLAB command is allowed at the prompt. To resume M-file function execution, use `DBCNT` or `DBSTEP`. To exit from the debugger use `DBQUIT`.

For experimentation, set a breakpoint in one of the functions used in your music synthesis project, and show the TA how you are able to examine one of the variables inside that function. Probably you should pick the function that does the sinusoidal generation and exhibit the variable that holds the frequency information. Note: next week we'll do a short verification on debugging.

## 1.1 Emergency 911

One of the most useful modes of the debugger, is setting it up so that whenever an error happens the program jumps into “debug mode.” This is done by typing:

```
dbstop if error
```

With this mode active, you can snoop around after an error has occurred and you will still be in the function where it happened so you can examine local variables that probably causes the error. It's sort of like an automatic call to 911 when you've gotten into an accident. Try `help dbstop` for more information.

## 2 Overview

There are three objectives/activities in this lab:

1. One objective of this lab is to introduce more complicated signals that are related to the basic sinusoid. These signals which implement frequency modulation (FM) and amplitude modulation (AM) are widely used in communication systems such as radio and television, but they also can be used to create interesting sounds that mimic musical instruments. There are a number of demonstrations on the CD-ROM that provide examples of these signals for many different conditions.
2. Use linear-FM chirps to exhibit aliasing
3. Use digital images to illustrate the D-to-A reconstruction process. In this lab, we will show a commonly method of reconstruction that gives “poor” results—a later lab will revisit this issue and do a better job.



### 2.1 Frequency Modulated Signals

We will also look at signals in which the frequency varies as a function of time. In the constant-frequency sinusoid  $A \cos(2\pi f_0 t + \phi)$  the argument of the cosine is also the exponent of the complex exponential, so the phase of this signal is the exponent  $(2\pi f_0 t + \phi)$ . This phase function changes *linearly* versus time, and its time derivative is  $2\pi f_0$  which equals the constant frequency of the cosine.

A generalization is available if we adopt the following notation for the class of signals with time-varying phase:

$$x(t) = A \cos(\psi(t)) = \Re\{Ae^{j\psi(t)}\} \quad (1)$$



The time derivative of the phase from (1) gives a frequency

$$\omega_i(t) = \frac{d}{dt}\psi(t) \quad (\text{rad/sec})$$

but we prefer units of hertz, so we divide by  $2\pi$  to define the *instantaneous frequency*:

$$f_i(t) = \frac{1}{2\pi} \frac{d}{dt}\psi(t) \quad (\text{Hz}) \quad (2)$$

## 2.2 Chirp, or Linearly Swept Frequency

A *chirp* signal is a sinusoid whose frequency changes linearly from some low value to a high one. The formula for such a signal can be defined by creating a complex exponential signal with quadratic phase by defining  $\psi(t)$  in (1) as

$$\psi(t) = 2\pi\mu t^2 + 2\pi f_0 t + \phi$$

The derivative of  $\psi(t)$  yields an instantaneous frequency (2) that changes *linearly* versus time.

$$f_i(t) = 2\mu t + f_0$$

The slope of  $f_i(t)$  is equal to  $2\mu$  and its intercept is equal to  $f_0$ . If the signal starts at  $t = 0$ , then  $f_0$  is also the starting frequency. The frequency variation produced by the time-varying phase is called *frequency modulation*, and this class of signals is called FM signals. Finally, since the linear variation of the frequency can produce an audible sound similar to a siren or a chirp, the linear-FM signals are also called “chirps.”

We have already seen that FM defines the signal  $x(t)$  to have a time-varying phase

$$x(t) = A \cos(\psi(t))$$

and that the instantaneous frequency changes according to the derivative of  $\psi(t)$ . If  $\psi(t)$  is linear,  $x(t)$  is a constant-frequency sinusoid; whereas, if  $\psi(t)$  is quadratic,  $x(t)$  is a chirp signal whose frequency changes linearly in time.

## 2.3 Advanced Topic: Spectrograms

It is often useful to think of signals in terms of their spectra. A signal’s spectrum is a representation of the frequencies present in the signal. For a constant frequency sinusoid the spectrum consists of two spikes, one at  $2\pi f_0$ , the other at  $-2\pi f_0$ . For more complicated signals, the spectra may be very interesting and, in the case of FM, the spectrum is considered to be time-varying. One way to represent the time-varying spectrum of a signal is the *spectrogram* (see Chapter 3 in the text). A spectrogram is found by estimating the frequency content in short sections of the signal. The magnitude of the spectrum over individual sections is plotted as intensity or color on a two-dimensional plot versus frequency and time.

There are a few important things to know about spectrograms:

1. In MATLAB the functions `specgram` and `plotspec` will compute the spectrogram, as already explained in the previous lab. Type `help specgram` `help plotspec` to learn more about these functions and their arguments *which are similar but not exactly the same*.
2. Spectrograms are numerically calculated and only provide a *numerical estimate* of the time-varying frequency content of a signal. There are theoretical limits on how well they can actually represent the frequency content of a signal. Lab 11 (in the book) treats this problem in the course of programming a method that uses the spectrogram to extract the frequencies of piano notes.

## 2.4 Digital Images

In this lab we introduce digital images as a signal type for studying the effect of sampling and reconstruction. An image can be represented as a function  $x(t_1, t_2)$  of two continuous variables representing the horizontal ( $t_2$ ) and vertical ( $t_1$ ) coordinates of a point in space.<sup>1</sup> For monochrome images, the  $x(t_1, t_2)$  would be a scalar function of the two spatial variables, but for color images the function would be a vector function of the two

<sup>1</sup>The variables  $t_1$  and  $t_2$  are confusing since they *do not denote time*.



CD-ROM

Spectrograms  
&  
Sounds:  
Wide-  
band  
FM



CD-ROM

Sounds  
& Spec-  
trograms

variables.<sup>2</sup> Moving images (TV) would add a time variable to the two spatial variables. Monochrome images are displayed using black and white and shades of gray, so they are called *gray-scale* images. In this lab we will consider only sampled gray-scale still images.

A sampled gray-scale still image would be represented as a two-dimensional array of numbers of the form

$$x[m, n] = x(mT_1, nT_2) \quad 1 \leq m \leq M, \text{ and } 1 \leq n \leq N$$

Typical values of  $M$  and  $N$  are on the order of 256 or 512; e.g., a  $512 \times 512$  image which has nearly the same resolution as a standard TV image. In MATLAB we can represent an image as a matrix consisting of  $M$  rows and  $N$  columns. The matrix entry at  $(m, n)$  is the sample value  $x[m, n]$ —called a *pixel* (short for picture element).

An important property of light images such as photographs and TV pictures is that their values are always non-negative and finite in magnitude; i.e.,

$$0 \leq x[m, n] \leq X_{\max}$$

This is because light images are formed by measuring the intensity of reflected or emitted light which must always be a positive finite quantity. When stored in a computer or displayed on a monitor, the values of  $x[m, n]$  have to be scaled relative to the maximum value  $X_{\max}$ . Usually an eight-bit integer representation is used. With 8-bit integers, the maximum value can be  $X_{\max} = 2^8 - 1 = 255$ , and there are  $2^8 = 256$  gray levels for the display.

## 2.5 Displaying Images

As you will discover, the correct display of an image on a gray-scale monitor can be tricky, especially after some processing has been performed on the image. We have provided the function `show_img.m` in the *DSP First Toolbox* to handle most of these problems, but it will be helpful if the following points are noted:



1. All image values must be non-negative for the purposes of display. Filtering may introduce negative values, especially if differencing is used (e.g., a high-pass filter).
2. The default format for most gray-scale displays is eight bits, so the pixel values  $x[m, n]$  in the image must be converted to integers in the range  $0 \leq x[m, n] \leq 255 = 2^8 - 1$ .
3. The actual display on the monitor is created with the `show_img` function.<sup>3</sup> The `show_img` function will handle the color map and the “true” size of the image. The appearance of the image can be altered by running the pixel values through a “color map.” In our case, all three primary colors (red, green and blue, or RGB) are used equally, so we get a “gray map.” In MATLAB the `gray` color map is created via

```
colormap(gray(256))
```

which gives a  $256 \times 3$  matrix where all 3 columns are equal. The MATLAB function `colormap(gray(256))` creates a linear mapping, so that each input pixel amplitude is rendered with a screen intensity proportional to its value (assuming the monitor is calibrated). For our experiments, non-linear color mappings would introduce an extra level of complication, so we won’t use them.

<sup>2</sup>For example, an RGB color system needs three values for red, green and blue at each spatial location.

<sup>3</sup>If the MATLAB function `imagesc.m` is used to display the image, two features will be missing: (1) the color map may be incorrect because it will not default to gray, and (2) the size of the image will not be a true pixel-for-pixel rendition of the image on the computer screen.

- When the image values lie outside the range [0,255], or when the image is scaled so that it only occupies a small portion of the range [0,255], the display may have poor quality. In this lab, we will use `show_img.m` to *automatically rescale the image*: This requires a linear mapping of the pixel values:<sup>4</sup>

$$x_s[m, n] = \alpha x[m, n] + \beta$$

The scaling constants  $\alpha$  and  $\beta$  can be derived from the min and max values of the image, so that all pixel values are recomputed via:

$$x_s[m, n] = \left\lfloor 255.999 \left( \frac{x[m, n] - x_{\min}}{x_{\max} - x_{\min}} \right) \right\rfloor$$

where  $\lfloor x \rfloor$  is the floor function, i.e., the greatest integer less than or equal to  $x$ .

### 3 Warm-up

The instructor verification sheet may be found at the end of this lab.

#### 3.1 Warm-up: Display of Images

You can load the images needed for this lab from \*.mat files. Any file with the extension \*.mat is in MATLAB format and can be loaded via the `load` command. To find some of these files, look for \*.mat in the *DSP First* toolbox or in the MATLAB directory called `toolbox/matlab/demos`. Some of the image files are named `lenna.mat`, `echart.mat` and `zone.mat`, but there are others within MATLAB's demos. The default size is  $256 \times 256$ , but alternate versions are available as  $512 \times 512$  images under names such as `lenna512.mat` and `zone512.mat`. After loading, use the command `whos` to determine the name of the variable that holds the image and its size.

Although MATLAB has several functions for displaying images on the CRT of the computer, we have written a special function `show_img()` for this lab. It is the visual equivalent of `soundsc()`, which we used when listening to speech and tones; i.e., `show_img()` is the "D-to-C" converter for images. This function handles the scaling of the image values and allows you to open up multiple image display windows. Here is the help on `show_img`:

```
function [ph] = show_img(img, figno, scaled, map)
%SHOW_IMG    display an image with possible scaling
% usage:  ph = show_img(img, figno, scaled, map)
%   img = input image
%   figno = figure number to use for the plot
%           if 0, re-use the same figure
%           if omitted a new figure will be opened
% optional args:
%   scaled = 1 (TRUE) to do auto-scale (DEFAULT)
%           not equal to 1 (FALSE) to inhibit scaling
%   map = user-specified color map
%   ph = figure handle returned to caller
%-----
```

Notice that unless the input parameter `figno` is specified, a new figure window will be opened. This feature must be turned off if you want to display several images together in `subplot`. Using `subplot` is the preferred method of *printing* multiple images (e.g., for comparisons in lab reports) because one `subplot` can easily hold four images. In addition, having several images on one printed page makes it much easier to discuss comparisons in a lab report.

<sup>4</sup>The MATLAB function `show_img` has an option to perform this scaling while making the image display.



### 3.2 Display Test

In order to probe your understanding of image display, generate a simple test image in which all of the rows are identical.

- (a) Now load and display the “lenna512” image from `lenna512.mat`. The command `load lenna512` will put the sampled image into the array `xx`. Use `whos` to check the size of `xx`.
- (b) Use the colon operator to extract the 200<sup>th</sup> row of the “lenna512” image, and make a plot of the 200<sup>th</sup> row as a 1-D discrete-time signal. Observe that the range of values is between 0 and 255—which ones represent white and which ones black?

**Instructor Verification** (separate page)

## 4 Lab Exercises: Sampling, Aliasing and Reconstruction

### 4.1 FM Signals: Chirps

The following MATLAB code will synthesize a chirp:

```
fsamp = 11025;
dt = 1/fsamp;
dur = 1.0;
tt = 0 : dt : dur;
psi = 2*pi*(1200*tt.*tt);
xx = real( exp(j*psi) );
soundsc( xx, fsamp );
```

- (a) Determine the mathematical formula for the chirp signal created by this function.
- (b) Make a sketch by hand of the instantaneous frequency versus time.
- (c) Determine the range of frequencies (in hertz) that will be synthesized by this MATLAB script. What are the minimum and maximum frequencies that will be heard?
- (d) Display the spectrogram of your chirp using the MATLAB function: `specgram(xx, [], fsamp)`.
- (e) Listen to the signal to verify that it has the expected frequency content (use `soundsc()`).

### 4.2 Chirps and Aliasing

- (a) Use the previous MATLAB code as a model to synthesize the following “chirp” signal:
  - (i) A total time duration of 3 secs.
  - (ii) The *desired* instantaneous frequency starts at 0 Hz and ends at 18,000 Hz.
  - (iii) Use a sampling rate and D/A conversion rate of  $f_s = 8000$  Hz.
  - (iv) Include the code with your lab report.
- (b) Listen to the signal. What comments can you make regarding the sound of the chirp (e.g., is it linear)? Does it chirp down, or chirp up, or both?
- (c) Create a spectrogram of your chirp signal. Use the sampling theorem (from Chapter 4 in the text) to help explain what you hear and see.

- (d) In addition, make some theoretical calculations by hand:
- (i) Determine the range of frequencies (in hertz) that will be synthesized by your MATLAB script.
  - (ii) Make a sketch by hand of the instantaneous frequency versus time.
  - (iii) Explain how *aliasing* affects the instantaneous frequency that is actually heard.
  - (iv) Listen to the signal again to verify that it has the expected frequency content.

### 4.3 Sampling of Images

Images that are stored in digital form on a computer have to be sampled images because they are stored in an  $M \times N$  array (i.e., a matrix). The sampling rate in the two spatial dimensions was chosen at the time the image was digitized (in units of samples per inch if the original was a photograph). For example, the image might have been “sampled” by a scanner where the resolution was chosen to be 300 dpi (dots per inch).<sup>5</sup> If we want a different sampling rate, we can simulate a *lower* sampling rate by simply throwing away samples in a periodic way. For example, if every other sample is removed, the sampling rate will be halved (in our example, the 300 dpi image would become a 150 dpi image). Sometimes this is called *sub-sampling* or *down-sampling*.

*Down-sampling:* If the vector `x1` represents a signal such as a row of an image, we can reduce the sampling rate by a factor of 4 by simply taking every 4<sup>th</sup> sample. In MATLAB this is easy with the colon operator, i.e., `xdown = x1(1:4:length(x1))`. The vector `xdown` will be one fourth the length of `x1`.

- (a) *Down-sampling* throws away samples, so it will shrink the size of the image. This is what is done by the following scheme:

$$\mathbf{x}_p = \mathbf{xx}(1:p:M, 1:p:N);$$

One potential problem with down-sampling is that aliasing might occur. This can be illustrated in a dramatic fashion with `zone512` image, and to a lesser degree with the `lenna512` image.

Load the `zone512` image which has the image stored in a variable called `zone`. Now down-sample the `zone512` image by a factor of 2. Notice the aliasing in the down-sampled image, which is surprising since no new values are being created by the down-sampling process. Describe how the aliasing appears visually.<sup>6</sup>

- (b) Down-sample the `lenna512` image by a factor of 2, and also by factors of 4 and 8. Notice that this image seems to be relatively unaffected by the down-sampling by 2 process, but eventually the larger down-sampling factors make a difference. What can you say about the frequency content of the `lenna` image as opposed to the `zone_plate` image?

### 4.4 Reconstruction of Images

When an image has been sampled, we can fill in the missing samples by doing interpolation. For images, this would be analogous to the examples shown in Chapter 4 for sine-wave interpolation which is part of the reconstruction process in a D-to-A converter. We could use a “square pulse” or a “triangular pulse” or other pulse shapes.

<sup>5</sup>The Sampling Theorem applies to digital images, so there is a *Nyquist Rate* that depends on the maximum *spatial* frequency in the image.

<sup>6</sup>One difficulty with showing aliasing is that we must display the pixels of the image exactly. This almost never happens because most monitors and printers will perform some sort of interpolation to adjust the size of the image to match the resolution of the device. In MATLAB we can override these size changes by using the function `truesize` which is part of the Image Processing Toolbox. In the *DSP First Toolbox*, an equivalent function called `trusize.m` is provided.

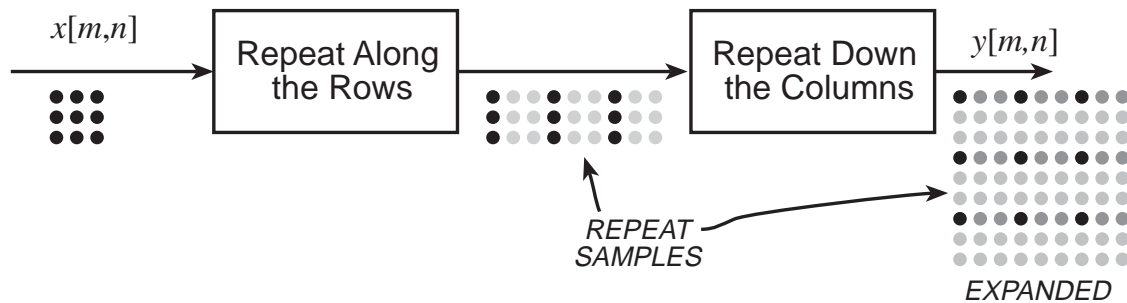


Figure 1: 2-D Interpolation broken down into row and column operations: the gray dots indicate repeated data values created by a zero-order hold; or, in the case of linear interpolation, they are the interpolated values.

For these reconstruction experiments, use the `lenna512` image, down-sampled by 4 (from the previous part). The objective will be to reconstruct the `lenna512` image which is 512 from the  $128 \times 128$  down-sampled image.

For your lab report, it would be wise to put the original image and the filtered reconstructions on the screen in different figure windows for easy comparison. However, when making hardcopy for your lab report, use `subplot` to print multiple images on the same page

- (a) The simplest interpolation would be reconstruction with a square pulse which produces a “zero-order hold.” Here is a method that works for a one-dimensional signal (i.e., one row or one column of the image), assuming that we start with a row vector `xx`.

```
L = length(xx);
nn = floor( 1:(1/4):L+0.75 );
yy = xx(nn);
```

Plot the vector `yy` to verify that it is a zero-order hold version derived from `xx`. Explain what values are contained in the indexing vector `nn`. Your lab report should include the explanation for this part, but no plots are required.

- (b) Now return to the down-sampled `lenna512` image, and process all the rows of `xx` to fill in the missing points. Use the idea from part (a). Call the result `yhold`. Display `yhold` as an image, and compare it to the original image `xx`.
- (c) Now process all the columns of `yhold` to fill in the missing points in each column and compare the result to the original image `xx`. Include your code for parts (b) and (c) in the lab report.
- (d) *Linear interpolation:* Carry out a linear interpolation operations on both the rows and columns using MATLAB’s `interp1` function with the `*linear` option. Call the interpolated output `ylin`. Compare `ylin` to the original image `xx` and to the square pulse interpolated image from part (c). Comment on the visual appearance of the two “reconstructed” images. Include your code for this part in the lab report.

When unsure about a command, use `help`.

- (e) Compare the linear interpolation result to the zero-order hold result.
- (f) *Comment:* You might use zooming to show the “zero-order hold” effect in a small patch the output image. See `help zoom`. However, zooming also does its own interpolation—probably it uses a zero-order hold.



## 4.5 More about Images in MATLAB (Optional)

For more information on the image processing functions in MATLAB, try help:

```
help images
```

but keep in mind that the Image Processing Toolbox, which is available in the CoC labs, may not be on your computer.

### 4.5.1 Zooming in Software

If you have used an image editing program such as Adobe’s “Photoshop,” you might have observed how well or how poorly image zooming (i.e., interpolation) is done. For example, if you try to blow up a JPEG file that you’ve downloaded from the web, the result is usually disappointing. Since MATLAB has the capability to read lots of different formats, you can apply the image zooming via interpolation to any photograph that you can acquire. The MATLAB function for reading JPEG images is `imread( )` which would be invoked as follows:

```
xx = imread('foo.jpg','jpeg');
```

Since `imread( )` is part of the image processing toolbox, this test can be done in the CoC computer labs, but may not be possible on your home computer.

### 4.5.2 Warnings

Images obtained from JPEG files might come in many different formats. Two precautions are necessary:

1. If MATLAB loads the image and stores it as 8-bit integers, then MATLAB will use an internal data type called `uint8`. The function `show_img( )` cannot handle this format, but there is a conversion function called `double( )` that will convert the 8-bit integers to double-precision floating-point for use with filtering and processing programs.

```
yy = double(xx);
```

2. If the image is a color photograph, then it is actually composed of three “image planes” and MATLAB will store it as a 3-D array. For example, the result of `whos` for a  $545 \times 668$  color image would give:

Name	Size	Bytes	Class
xx	545x668x3	1092180	uint8 array

In this case, you should use MATLAB’s image display functions such as `imshow( )` to see the color image. Or you can convert the color image to gray-scale with the function `rgb2gray( )`. For more information on the image processing functions in MATLAB, try help:

```
help images
```

**Lab #5**

**EE-2200**

**Winter-99**

**INSTRUCTOR VERIFICATION PAGE**

Staple this page to the end of your Lab Report.

Name: \_\_\_\_\_

Date of Lab: \_\_\_\_\_

Part 3.2 Load and display a digital image; extract one row of the image:

Verified: \_\_\_\_\_

Date/Time: \_\_\_\_\_