DUE DATE: APRIL 3, 2009

On the course website are links to some sound files from several different instruments, such as a trumpet and a flute.

1. Convert the .wav file to a human readable ascii file. There are several different ways to do this. For example, you can use the sox command in linux and it is as simple as

   ```
   tiki:amat415$ sox bassoon.wav > bassoon.dat
   ```

   What you should find in the resulting .dat file are three columns, the first being the time stamp and the second two being a decimal between -1 and $+1$, which is the conversion of a signed 16 bit binary number. Count the number of lines in the resulting .dat file, and list the first ten of them.

2. Make a new file, which contains the middle column of the .dat file. Call it, say, bassoon.level. It is easy to do this on linux,

   ```
   tiki:amat415$ awk '{$2}' basson.dat > bassoon.level
   ```

   after first removing the header information, which you will see in the previous part (this is the first couple of lines in the file, which tell the sampling rate etc.) It is also easy to do in Octave, matlab or even emacs.

3. Use an FFT to locate the fundamental frequency of the sample. Using 440 Hz as a reference pitch for A, decide which note the instrument is trying to play and decide whether the sample is sharp or flat, and by how many cents. Here assume that we are using the well tempered scale which goes by frequency multiples of $2^{1/12}$. You may also find it useful to listen to the sound file first, and just take out a known number of lines of the converted data file, such as 44,100 of them, or a number of lines which is a large power of two. For example, the following commands in Octave will get you on your way:

   ```
   octave > x = load(bassoon.level)
   ```

   ```
   octave > plot(abs(fft(x)))
   ```

   However you decide to do the calculation, you must describe what you did to arrive at your conclusion.